Zurich Platform security

Last updated: 06/12/2026

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

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

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

servicenow

Table of Contents

Secure your instance

Platform security provides capabilities to secure the instance.

ServiceNow Vault

Use the ServiceNow Vault product set of data security tools that protect sensitive information from unauthorized access, corruption, or theft throughout its entire life cycle.

Platform Security

Platform security provides capabilities to secure the instance.

Platform Privacy

Platform Privacy enables you to mask the sensitive date on the instance.

Encryption

Protect your sensitive data and stay compliant with regulatory#equirements and standards.

Access Management

Access Management enables you to have access to ServiceNow instance securely.

Identity

Know more about the Identities in the instance.

Additional resources for Platform Security products and solutions

Explore additional Platform Security resources for learning.

ServiceNow Vault

ServiceNow Vault provides a single location to review and implement security tools for advanced security and privacy requirements.

Get started

Explore ServiceNow Vault

Learn about ServiceNow Vault and its features.

Configure ServiceNow Vault

Get started with installing and configuring ServiceNow Vault

ServiceNow Vault dashboard

Learn how to use the ServiceNow Vault dashboard to review your data security.

ServiceNow Vault tools

Learn about the tools used with ServiceNow Vault

Exploring ServiceNow Vault

Learn more about ServiceNow Vault and review the benefits it can provide for your data protection needs.

ServiceNow Vault overview

ServiceNow Vault is a set of data security tools that protect sensitive information from unauthorized access, corruption, or theft throughout its entire life cycle. You can use ServiceNow Vault to apply protections like encryption, secrets management, and data privacy for the redaction and auditing of sensitive information.

Note:

ServiceNow Vault is a paid plugin that ServiceNow personnel must activate on your production instance. ServiceNow Vault includes the components listed here.

To purchase a subscription, contact your ServiceNow account manager. When you purchase a subscription, certain plugins are activated automatically. If a paid plugin isn't activated automatically, you can manually activate it from the **All Applications** list in your instance.

ServiceNow Vault benefits

ServiceNow Vault benefits

Benefit	Feature	Users
Get a comprehensive ServiceNow Vault overview and review the metrics related to the discovery, classification, and protection of your data	ServiceNow Vault console dashboard	Platform Admins
Guided setups to easily start using ServiceNow Vault tools with applications	Guided Vault	Platform Admins

What to explore next

To learn more about configuring and using ServiceNow Vault, see:

- ServiceNow Vault console dashboard
- Configuring ServiceNow Vault
- Use guided setup for ServiceNow Vault

ServiceNow Vault tools

Encryption

Key Management and Field Encryption is a suite of highly configurable encryption modules

Code Signing

Help improve security by validating sensitive application configuration data and scripts before they are used.

Data Privacy

Use the Data Privacy plugin to remove personally identifiable information (PII) from user data when it is migrated from a production instance to a non-production instance.

Data Discovery

The Data Discovery plugin enables you to find personally identifiable information (PII) from user data. The data can then be classified for further security measures.

Logs

Improve security, performance, and user experience by importing ServiceNow log data into enterprise log analytics using the log export service.

Zero Trust Access

ServiceNow Session Access enables organizations to dynamically reduce user privilege in a web session

Configuring ServiceNow Vault

Learn how to install and configure ServiceNow Vault

Configuration overview

- Install Vault plugins

  Learn how to install ServiceNow Vault for your instance.

-

  Learn and setup the roles necessary to use ServiceNow Vault

- Install Vault plugins

  Install and verify the plug-ins needed to make the most of ServiceNow Vault.

Using guided setup to implement ServiceNow Vault

Guided setup provides a sequence of tasks that help you configure ServiceNow Vault on your ServiceNow instance. To open guided setup for ServiceNow Vault, navigate to **All** > **Vault** > **Vault console**.

For more information, seeGuided Vault.

Install Vault plugins

Learn how to review and install the necessary plugins for ServiceNow Vault.

Before you begin

Role required: sn_vault_console.vault_console_admin

About this task

ServiceNow Vault is best used with other ServiceNow security products. Install and configure the other plugins to make the most of your tools and metrics.

Procedure

1. Navigate to **All** > **Vault** > **Vault Console**.
2. Review the status of your tools.

Tool statuses

Status	Description
Limited access	Tool has a limited access license. Reported metrics may be affected or unavailable.
Included with platform	Tool is included by default with the platform
Needs license	Tool has no license. Reported metrics may be affected or unavailable.
Premium license	Tool has full access.

3. **Optional:** If a status is either limited or needs a license, review its respective installation page, and then install or activate the tool.

Tool installation

Tool	Installation page
Data Discovery	Activating Data Discovery
Data Classification	Included with the platform
Data anonymization	Activate data privacy **Note:** Data anonymization is installed as part of Data privacy
Cloud Encryption with Key Management	Requesting Cloud Encryption for instances with Now Support Service Catalog [KB1117369] article in the Now Support Knowledge Base
Field Encryption	Activate Field Encryption
○ Zero Trust Access (ZTA) ○ Continuous Authentication (CA)	○ Activate Zero Trust Access ○ Activate Continuous Authentication

Use guided setup for ServiceNow Vault

Use guided setup to begin using an application with ServiceNow Vault easily.

Before you begin

Role required: Elevate to sn_vault_console.vault_console_admin role.

Procedure

1. Navigate to **All** > **Vault** > **Vault console**.
2. In the **Guided vault** section, Select the **Get started** button for your application.
3. In **Select app data** use the table to select the application data to be used with ServiceNow Vault, select **Preview classes** when finished

Select app data table fields

Label	Description
Table	The table the data is located in.
Column	The column the data is located in
Existing Class	The current classification of the data
Recommended Class	The recommended classification for the data.
Application	The application scope of the data

4. In **Preview data classification** preview your data classification settings, when finished check agree and select the **Apply recommended classes** button.

Preview data classification table fields

Label	Description
Table	The table the data is located in.
Column	The column the data is located in
Final class	The class the data will be assigned
Application	The application scope of the data

5. In **Classification summary** review the results of the data classification, select **Next** when finished.

Trouble?
Use Data Classification to review any data that failed to classify.

6. In **Protect existing data** review the protection policies of the data.

Protect existing data table fields

Label	Description
Table	The table the data is located in.
Column	The column the data is located in
Anonymization	Reports if a data anonymization policy can be, or already is applied to the data.
Field encryption	Reports if a field encryption policy can be, or already is applied to the data.
Zero trust access	Reports if a zero trust access policy can be, or already is applied to the data.

7. **Optional:** You can select **Available** to begin applying that columns respective application data protection policy, review Vault tools and metrics for more information.
8. In **Protect real time data** review your real time data protection policies, select **Mark as complete** when finished.

Protect real time data fields

Label	Description
Name	The name of the ???
Type	The data channel type.
Real time anonymization	Reports if there is a real time anonymization policy available, or already applied to the data.

Note: Real time data will be protected across the entire instance and all applications,.

Result
The select application now has classified data and protection policies. As well, it will now report relevant metrics to the ServiceNow Vault console dashboard.

ServiceNow Vault console dashboard

Use the ServiceNow Vault console dashboard to track and manage your ServiceNow Vault security tools.

The dashboard provides an easy way to review the security, privacy and compliance of your sensitive data. The dashboard reports on ServiceNow Vault tools and their various metrics, as well as showing guided setups for applications compatible for ServiceNow Vault.

To access the ServiceNow Vault console dashboard, navigate to **All** > **Vault** > **Vault Console**

Vault console dashboard page

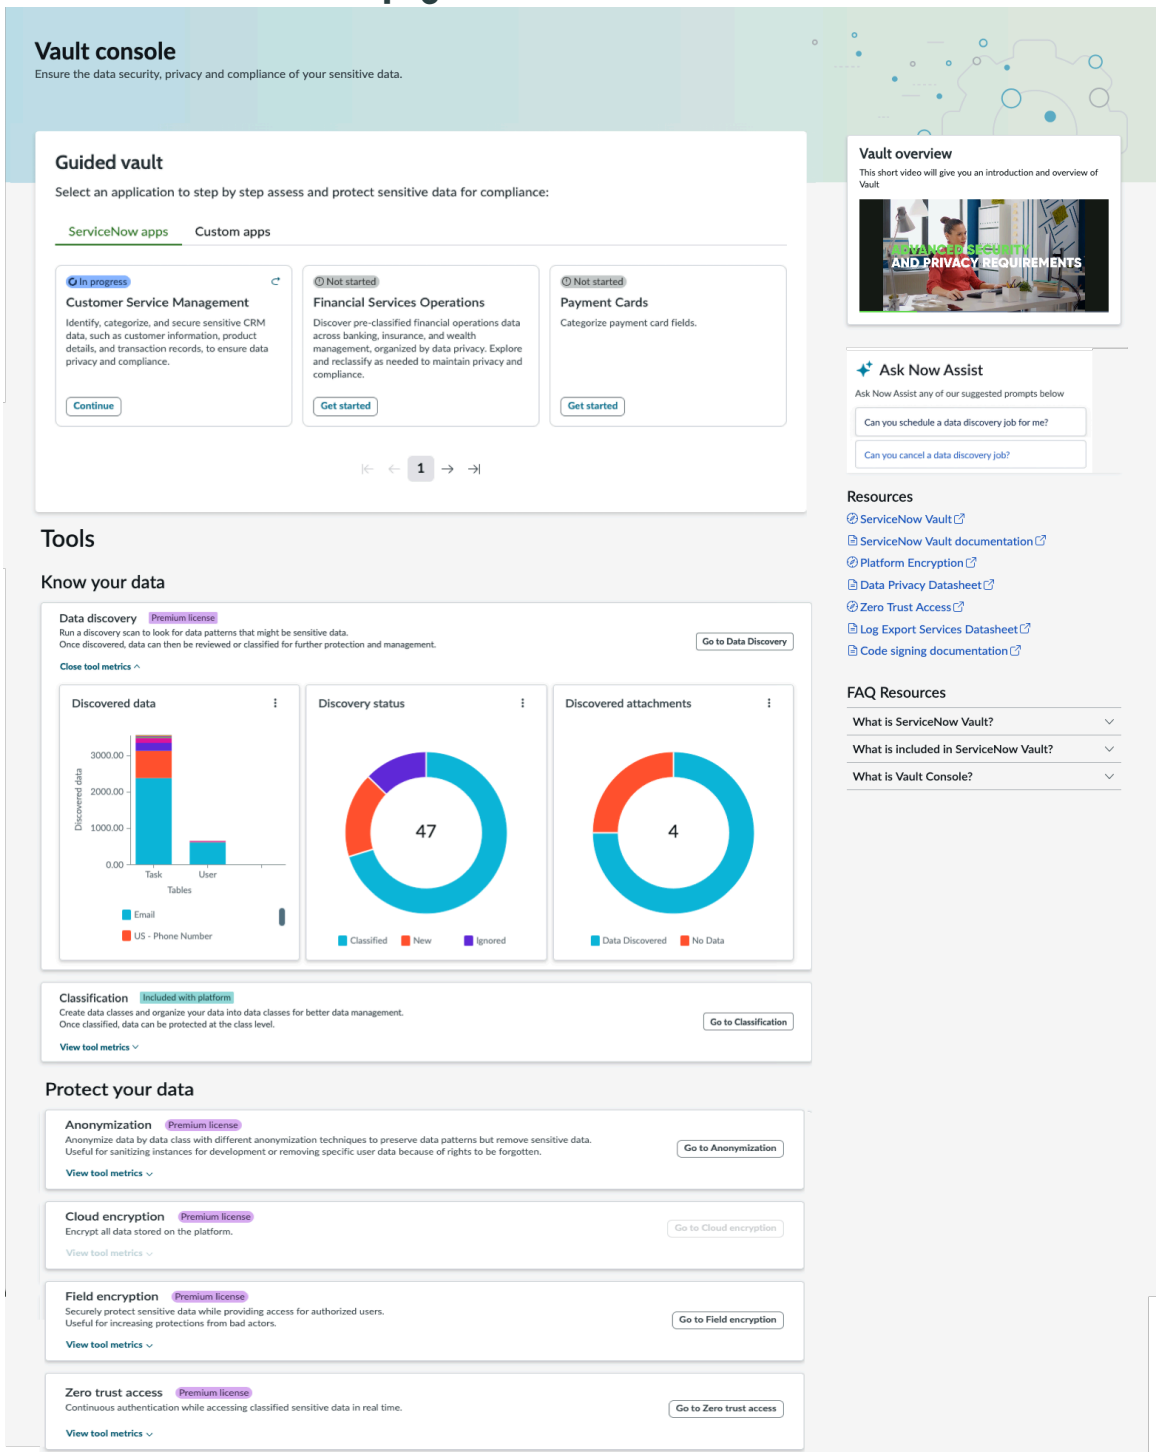

Section	Subsection	Description
Vault Overview, Resources, and Ask Now Assist	Vault Overview	Provides a video overview of ServiceNow Vault
	Resources	Offers additional resources for help using ServiceNow Vault
	FAQ Resources	Get answers to commonly asked questions about ServiceNow Vault

Section	Subsection	Description
	Ask Now Assist	Use generative AI to streamline your tasks in ServiceNow Vault. For example, you can ask Now Assist to schedule a data discovery job for you. Prerequisite: Enable skills in Now Assist Admin console. For more information, see Activate a Now Assist skill and Using Now Assist for Vault.
Guided Vault	Guided setups cards for applications	Select the **Get Started** button on an applications card to start using ServiceNow Vault with that application. Read more about guided setup here:Use guided setup for ServiceNow Vault.
Vault tools and metrics	Tool information	A brief description and licensing information for the tools used in ServiceNow Vault. Select the **Go to** button go to the tools homepage. Tools currently displayed on the dashboard are: • Data Discovery • Data Classification • Data anonymization • Cloud Encryption with Key Management • Field Encryption • Zero Trust Access (ZTA)
	Tool metrics	Select the dropdown to review various metrics and graphs about a tool. **Note:** Find out more information about the metrics by reviewing the respective tool's homepage.

Guided Vault

The guided vault setup offers a step by step process to assess and protect sensitive data for an application.

Guided vault setup offers an easy way to begin using ServiceNow Vault with your applications. The setup will use the selected applications data. Setup first helps the user review the application data for classification, and then helps implement protection policies on the classified data. You can review how to use the ServiceNow Vault guided setup at Use guided setup for ServiceNow Vault

There are 5 steps in a guided setup:

1. Select app data
2. Preview data classification
3. Classification summary
4. Protect existing data
5. Protect real time data

After a guided vault setup you can review application data metrics in the **Tools** section of the Vault console dashboard. For more information see Vault tools and metrics

Vault tools and metrics

Learn about the tools and metrics ServiceNow Vault uses to protect and discover sensitive data.

ServiceNow Vault integrates with several tools to provide a cohesive overview of your sensitive data security. You can hover over a widget to get further insight on the reported data. Select the **Go to** button on any tool to go to its respective page.

Know your data

ServiceNow Vault uses Data Discovery and Data Classification to help you understand and know your data.

Tools and metrics

Tool	Metric	Description
Discovery Use Data Discovery to run a discovery scan to look for data patterns that might be sensitive data. Once discovered, data can then be reviewed or classified for further protection and management.	Discovered data	Occurrences of sensitive data across tables in your instance, categorized by sensitive data pattern type.
	Discovery status	Current state of all discovered sensitive data patterns, including new findings pending review, classified, or marked as ignored.
	Discovered attachments	Total sensitive data occurrences in attachments across tables in your instance.
Classification Data Classification create data classes and helps organize your data into data classes for better management. Classified data can be protected at the class level.	Classifiable data	Proportions of classifiable data.
	Classified data	Proportions of classified data.

Protect your data

ServiceNow Vault uses data anonymization, cloud encryption, field encryption, and zero trust access to help secure and protect your data.

Tools and metrics

Tool	Metric	Description
Anonymization Anonymize data by data class with different anonymization techniques to preserve data patterns but remove sensitive data. Useful for sanitizing instances for development or removing specific user data because of rights to be forgotten.	Existing data	All classified data per workflow that is anonymized or not.
	Real time data	The amount of anonymized real time data.
	Anonymization run times	Run times in hours of data anonymized in real time by channel, such as Now Assist or Virtual Agent.
Cloud Encryption with Key Management Securely protect sensitive data in encrypted storage for your data using block encryption, along with enhanced key management.	Active cloud key	Total rotations of the active cloud key. **Note:** To view this data, you need the Key Management Framework admin role (sn_kmf.admin or sn_kmf.cryptographic_manager).
	Key rotation	Time elapsed between each rotation of active keys on your instance. Bar height measures how long a key was used before rotation. **Note:** To view this data, you need the Key Management Framework admin role (sn_kmf.admin or sn_kmf.cryptographic_manager).
Field Encryption Securely protect sensitive data while providing access for authorized users. Useful for increasing protections from bad actors.	Encrypted fields classification status	Classification status of all data protected with Field Encryption.
	Classes protected with Field Encryption	The proportion of classified data protected withField Encryption.
	Active encryption keys	Number of active Field Encryption keys in your instance. Ideally, the number of active keys matches the number of classifications.

Tools and metrics (continued)

Tool	Metric	Description
		Note: To view this data, you need the Key Management Framework admin role (sn_kmf.admin or sn_kmf.cryptographic_manager) and the security_admin role.
Zero Trust Access (ZTA) Continuous authentication while accessing classified sensitive data in real time.	Continuous Authentication classification status	Number of classifications that are protected due to the continuous authentication policies.
	Classes protected with Continuous authentication	Number of classes protected with continuous authentication, categorized by their class.

All ServiceNow Vault tools

Encryption

Key Management and Field Encryption is a suite of highly configurable encryption modules

Code Signing

Help improve security by validating sensitive application configuration data and scripts before they are used.

Data Privacy

Use the Data Privacy plugin to remove personally identifiable information (PII) from user data when it is migrated from a production instance to a non-production instance.

Data Discovery

The Data Discovery plugin enables you to find personally identifiable information (PII) from user data. The data can then be classified for further security measures.

Log Export Services

Improve security, performance, and user experience by importing ServiceNow log data into enterprise log analytics using the log export service.

Zero Trust Access

ServiceNow Session Access enables organizations to dynamically reduce user privilege in a web session

Now Assist for Vault

With Now Assist for Vault, you can generate custom data patterns, check role access for an encrypted column, and schedule data discovery jobs. Now Assist for Vault can make it easier for you to perform common tasks without going to multiple systems.

Get started

Explore	Configure	Use
		
Learn about the generative AI skills that are available in Now Assist for Vault	Configure the generative AI skills in Now Assist for Vault	Use the generative AI capabilities in Now Assist for Vault

Important:

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

Troubleshoot and get help

Some ServiceNow resources that can provide you with helpful information are:

ServiceNow Community

ServiceNow Community

Developer

developer.servicenow.com

Impact

http://impact.servicenow.com

ServiceNow University

ServiceNow University

Best Practices

Best Practices

Partner

https://www.servicenow.com/partners.html

ServiceNow

http://servicenow.com

ServiceNow Store

http://servicenow.com

Support

- https://support.servicenow.com/now
- Known Error Portal

AI limitations

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

Data processing

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

Data collection

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

Exploring Now Assist for Vault

With the Now Assist for Vault application, you can generate custom data patterns from text descriptions to streamline your workload, check role access for an encrypted column to monitor your instance's encryption access posture, and schedule Data Discovery jobs to detect sensitive data.

Ask Now Assist panel in ServiceNow Vault console

The Ask Now Assist panel in ServiceNow Vault console lists skills to accomplish common security tasks. Selecting a skill opens a conversational interface that lets you accomplish a task through prompts. For more information about the Now Assist conversational interface, see Now Assist panel.

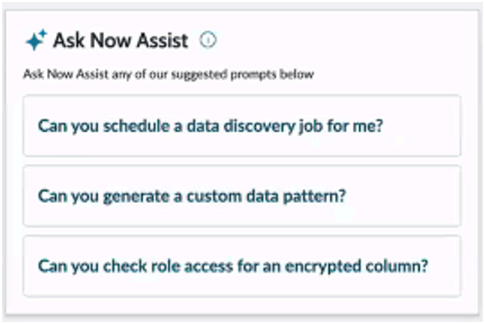

Configuring Now Assist for Vault

Install the ServiceNow® Now Assist for Vault application from the ServiceNow® Store to get Now Assist for Vault.

Before you begin

Review the Now Assist for Vault application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility. Now Assist for Vault installs the Now Assist for Vault application (sn_vault_gen_ai).

Role required: admin

Procedure

1. From the Now Assist for Vault application page on the ServiceNow Store, select **Buy**.
2. After approval has been granted, on your instance, navigate to **All > System Applications > All Available Applications > All**.
3. Using the search bar, search for the Now Assist for Vault application (sn_vault_gen_ai).
4. Select **Install**.
5. Verify that Now Assist for Vault is installed:
   a. Navigate to **All > Now Assist Admin > Now Assist Skills**.
   b. In the workflow list, select **Vault**.
   c. Verify that the Now Assist for Vault skills are active.

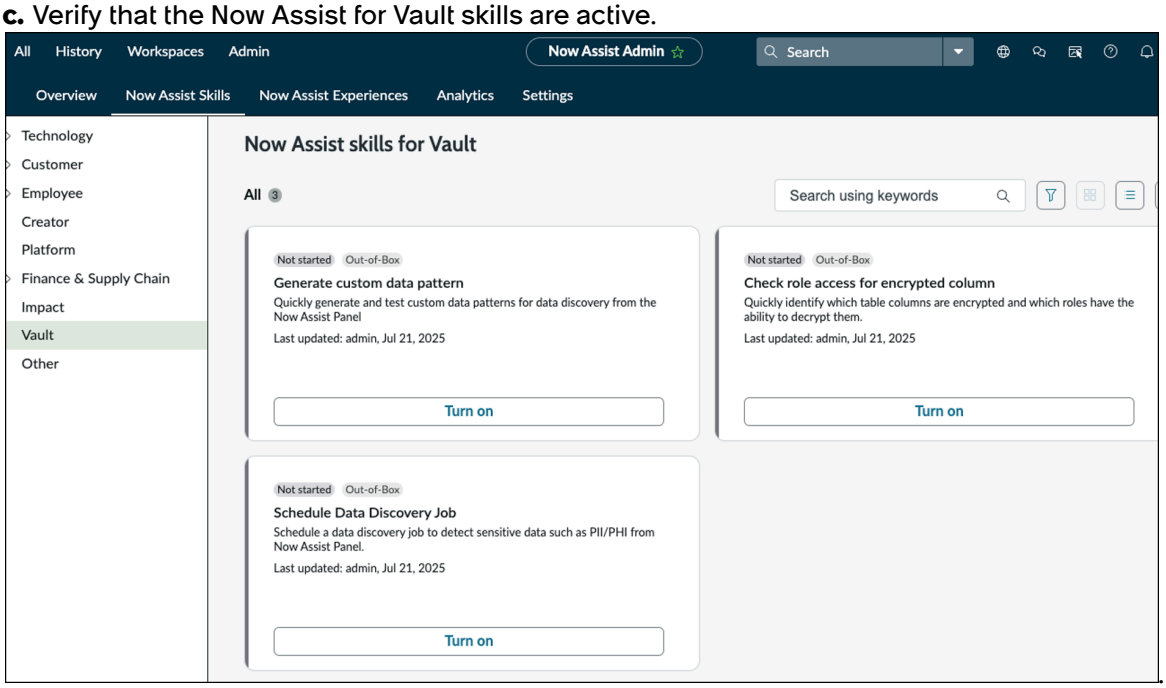

Using Now Assist for Vault

The Now Assist for Vault application includes the generative AI skills and features that enable you to streamline your administrative workload.

After you activate skills, they appear in the Ask Now Assist panel in ServiceNow Vault console.

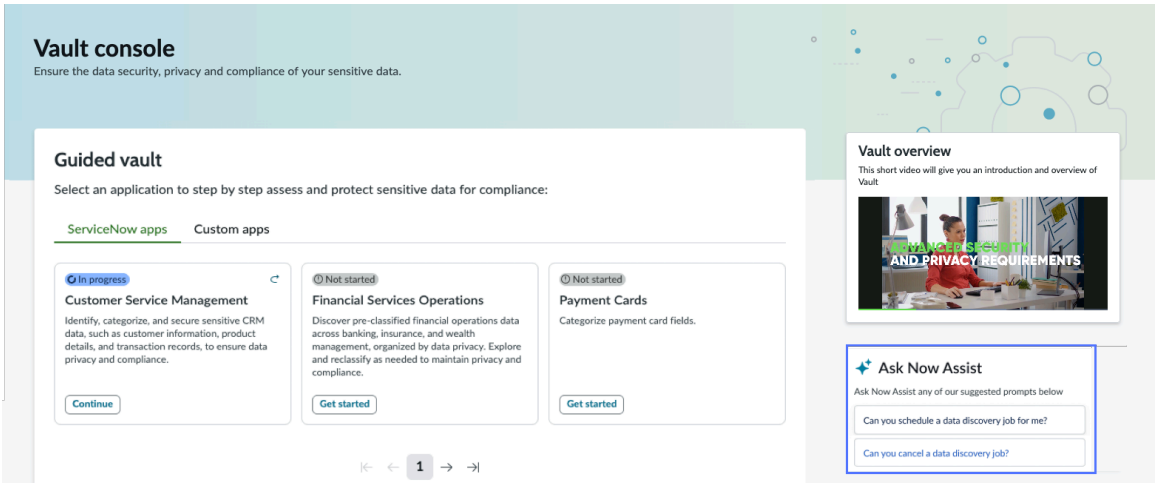

By default, all skills exist in the global domain. When you use Now Assist in a domain-separated environment, users are only able to access data in their domain. For example, if a user uses the summarization skill, Now Assist only uses material that exists in the user's domain when generating that summary.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

Generate a custom data pattern by using Now Assist for Vault

Use the generate custom data pattern skill to create a custom regular expression data pattern from your description and add it as an active data pattern to your instance.

Before you begin

- Install ServiceNow Vault. For more information, see Configuring ServiceNow Vault.
- Ensure that the generate custom data pattern skill is active. For more information, see Activate a Now Assist skill.

Role required: sn_vault_console.vault_console_admin

About this task

This skill improves your efficiency by reducing the time required to understand and write regular expressions.

Procedure

1. Navigate to **All** > **Vault** > **Vault console**.
2. In the Ask Now Assist panel, select **Generate custom data pattern** and specify the details.

Example

Example prompt: `I need help creating a regex to detect sensitive data for a Netherlands drivers license.`

Check role access for an encrypted column with Now Assist for Vault

Use the check role access for encrypted column skill to identify user roles that have access to encryption and decryption keys in your instance.

Before you begin

- Install ServiceNow Vault. For more information, see Configuring ServiceNow Vault.
- Ensure that the check role access for encrypted column skill is active. For more information, see Activate a Now Assist skill.

Role required: sn_vault_console.vault_console_admin

Procedure

1. Navigate to **All** > **Vault** > **Vault console**.
2. In the Ask Now Assist panel, select **Check role access for encrypted column** and specify the details.

Example
Example prompt: `Which roles have decryption key access to an encrypted column? Access includes read access.`

Schedule a Data Discovery job with Now Assist for Vault

Use the schedule data discovery job skill to schedule one-time or recurring Data Discovery jobs with Now Assist. Data Discovery jobs can detect sensitive data such as PII or PHI provided as input to the Now LLM.

Before you begin

- Install ServiceNow Vault. For more information, see Configuring ServiceNow Vault.
- Ensure that the schedule Data Discovery job skill is active. For more information, see Activate a Now Assist skill.

Role required: sn_vault_console.vault_console_admin

Procedure

1. Navigate to **All** > **Vault** > **Vault console**.
2. In the Ask Now Assist panel, select **Schedule data discovery job** and specify the job details.

Example
Example prompt: `Configure a data discovery job that scans the tables Task and Knowledge, looking for data patterns like Credit Card-Visa and Phone Number in columns such as Description and Text. Schedule a sample scan to occur on August 15th, 2025, with a time window from 10:00 AM to 5:00 PM. Name the job "Credit Card & Phone Number Scan" and describe it as "Scans for credit card and phone number data in Task and Knowledge tables."`

Platform Security

Platform security provides capabilities to secure the instance.

Security Center

Allow admins to continuously maintain the highest level of security posture and easily monitor for insecure events and behaviors.

Log Export Service

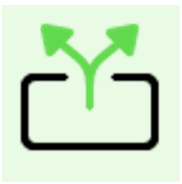

Log Export Service (LES) lets you seamlessly export your instance system and application logs into your enterprise security analytic tools. The service provides a highly scalable and near real-time integration with your analytic tools that is easy to setup and maintain.

Logs

Logs module provides a variety of logs that you can use to troubleshoot and debug transactions and events that take place within the instance.

Secrets Management

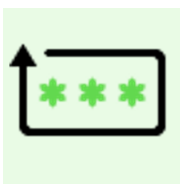

Use ServiceNow Secrets Management for granular management of access to your passwords to fit your business needs.

Code Signing

Code Signing creates digital signatures for the data which later are checked to confirm the authenticity and integrity of the data.

Antivirus Scanning Capabilities

Use Antivirus Scanning to help protect your instance against virus infections that can be introduced by file attachments to your system records, such as incidents, problems, and stories.

HTML Sanitizer

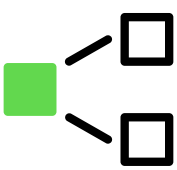

Auditing

High Security Settings

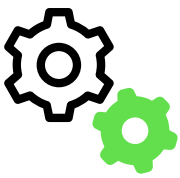

Remove unwanted code and protect against security concerns such as cross-site scripting attacks by sanitizing HTML markup in HTML fields and translated HTML fields.	Track record changes on auditing-enabled tables. By default, the system tracks changes to the incident, change, and problem tables, among others.	High Security Settings refer to several security options available in your instance.
Virtual Private Network (VPN)  Use a virtual private network (VPN) to integrate your instance with external data sources over the Internet.	**Hardening Settings**  Find detailed descriptions and compliance values for Security Center (SSC) hardening settings content	

Security Center

ServiceNow Security Center is an application that consists of a set of tools designed to help your organization maintain the security of your ServiceNow deployments. Using Security Center, you can improve security posture and strengthen compliance levels with a seamless user experience.

Security Center is a free application that administrators can download from the ServiceNow Store. It's installed by default starting with the Vancouver release. It's also made available in the ServiceNow Store every quarter, between family releases, to facilitate faster adoption of new functionality.

Important:

Instance Security Center (ISC) has reached the end of sales as of September 2024, and is no longer supported or available for activation.

Security Center (SSC) is the recommended solution going forward. For more information, see Instance Security Center to ServiceNow Security Center migration.

Get started

Security center landing

Use the Security Center landing page to find the information and tools you need to secure your instance.

Security configuration console

Use the Security configuration console to get a quick overview of the security posture of your instance on the overview page.

Security monitoring console

Supervise security notifications and metrics to stay informed about potential security risks on your instance.

Security posture console

Improve your security posture with comprehensive visibility and step-by-step instructions.

Security banner announcements

Supervise security notifications and metrics to stay informed about potential security risks on your instance.

Troubleshoot and get help

- Ask questions and explore other resources for Security Center in the Security Operations section of ServiceNow Community.
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Security Center landing page

Use the Security Center landing page to find the information and tools you need to secure your instance.

My Security Tasks

The my Security Tasks section displays a list of Security Tasks assigned to you. Open a task by selecting the number to review a tasks details, or select **View all** access the Security Tasks Manager, where all tasks are visible. For more information on Security Tasks, see Security Tasks.

Instance summary

The instance summary displays a high-level overview of the state of your instance's security.

Select the info (ⓘ) icon on any card for details on what each represents. You can also select any card to navigate to the relevant area of Security Center.

Identity and Access Management

Use the tools in the Identity and Access Management (IAM) section verify that your data is only accessible to the users and processes that need it. For details on these tools, see Identity and Access Management.

Security consoles

Select a card within the security consoles section to navigate to any of the three security consoles.

- Security configuration console
- Security monitoring console
- Security posture console

Tools

Use the tools section to navigate to the tools you need to manage your instance security. These tools are organized into three tabs displayed at the top of the section. Select a card to navigate to the selected tool.

Additional resources

Use the additional resources section on the edge of screen to navigate to documentation and videos relevant to learning Security Center as well as information on managing instance security.

Security center has been updated in Zurich to provide guidance to understand what the tool can do how to use it. This redesign is tailored to assist both new and infrequent users of Security Center. New modal windows appear within security center to demonstrate how these tools can help you achieve your security goals.

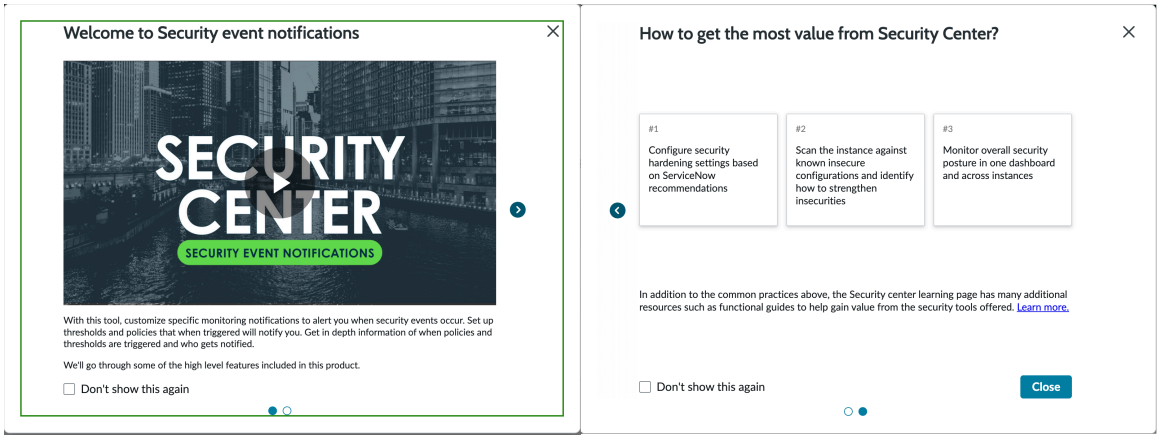

Identity and Access Management

Use the tools in the Identity and Access Management (IAM) section verify that your data is only accessible to the users and processes that need it.

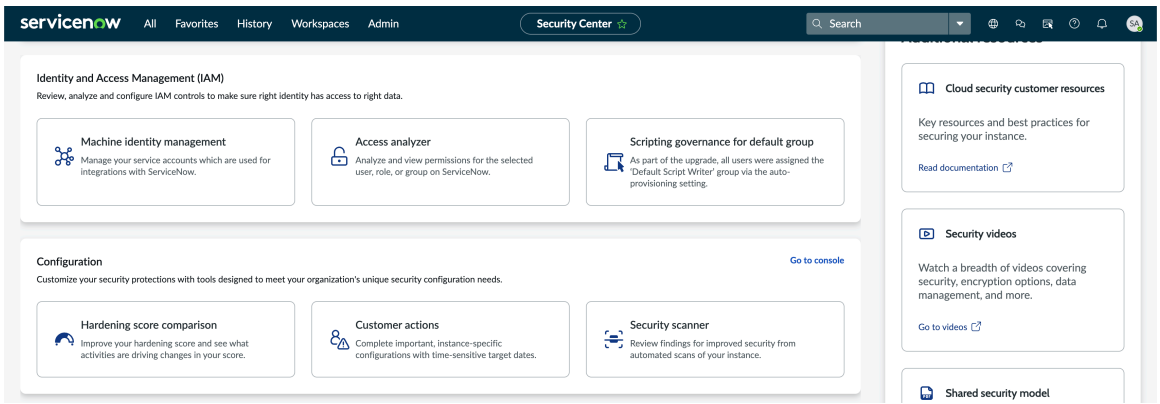

IAM consists of three tools you can use to manage access to the data on your instance.

Machine Identity Management

Machine identities are digital credentials such as certificates, keys, and tokens. Servers, applications, containers, and cloud services use these identities authenticate with each other. Use the **Machine Identity Console** to manage the machine identities used for integrations with ServiceNow.

Access analyzer

Use the **Analyze Access and Permissions** console to view permissions for a selected user, role, group, or compare access between two users.

Scripting Governance for default group

Use the Scripting Governance Tool to configure preferences for the Conditional Script Writer group. Users in this group are assigned the *`snc_required_script_writer_permission`* role, which allows users to access scripts and script-like fields across the platform.

With the settings on this console you can turn auto-assignment of this role on or off, as well manually assign users to the group. You can also see information on who is assigned, and scan your instance to find users who have scripted in a specific time frame.

Security configuration console

Use the security configuration page to get an overview of the security posture of your instance. View your hardening compliance score, discover graphical trends, analyze the top non-compliant hardening settings, and see the results of your security scans.

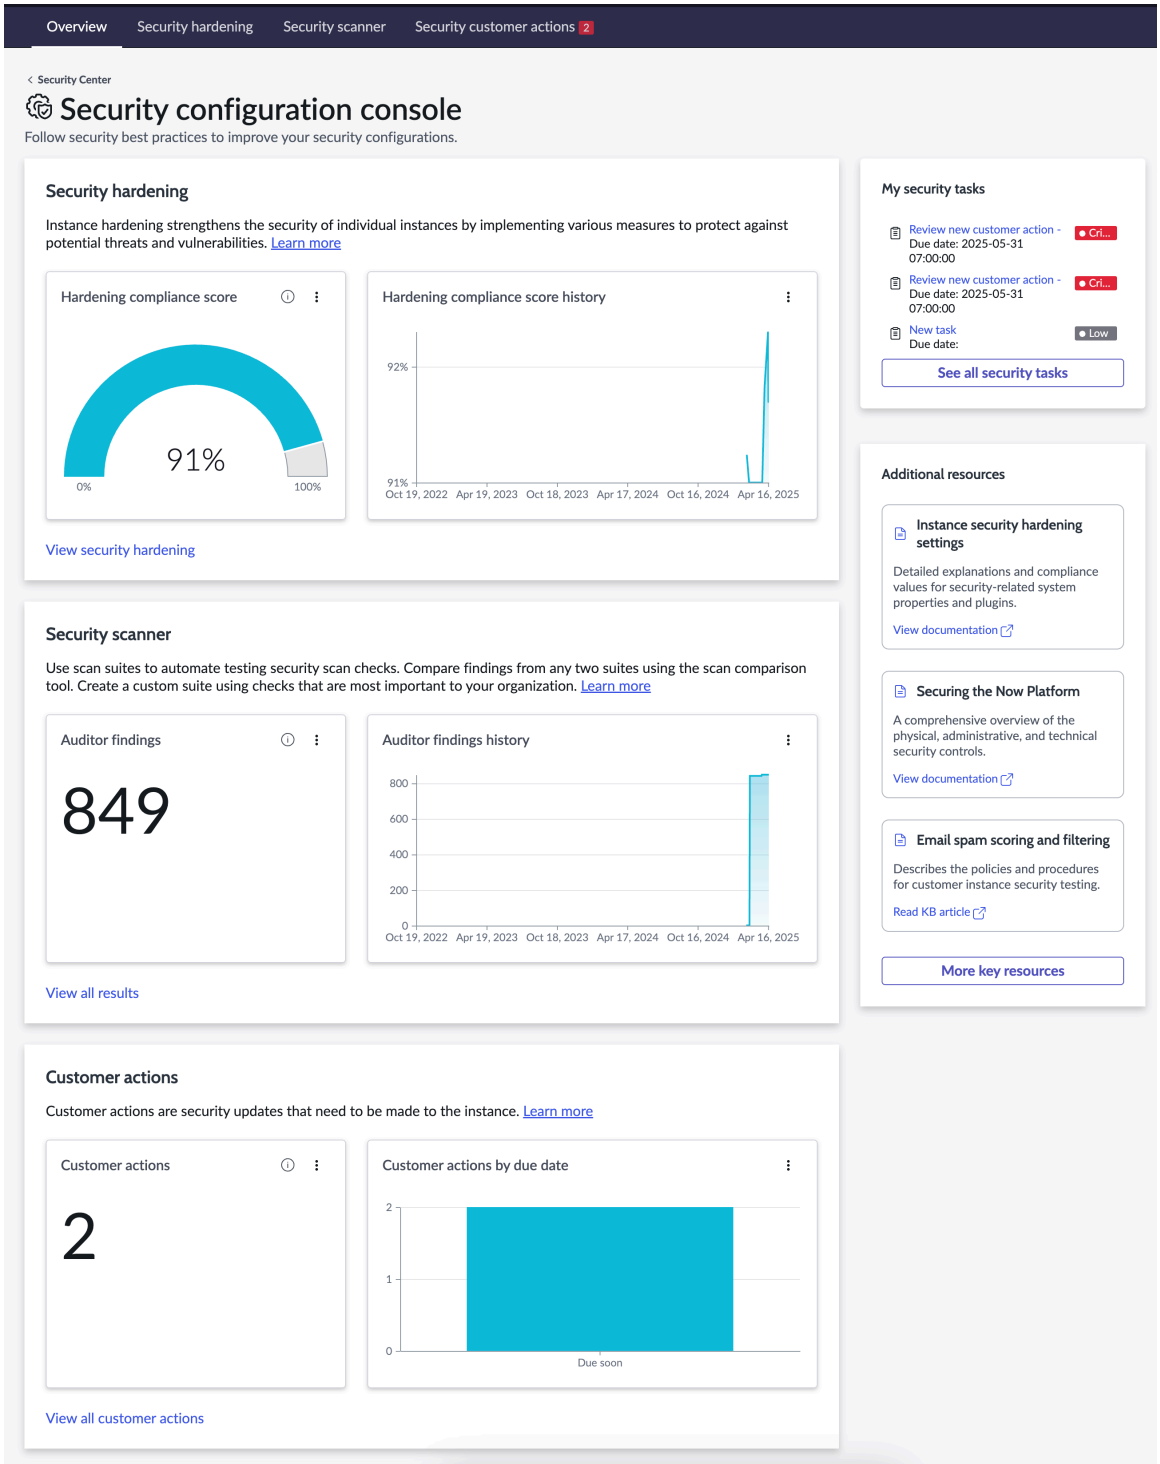

Navigation bar

Use the bar at the top of the page to navigate between this page and the security hardening, security scanner, or security Customer Actions section of Security Center.

Security hardening

Find details about your hardening compliance in the security hardening section. Select the cards on in this section to navigate to the Security hardening page.

The hardening compliance score is a percentage that shows how closely your security settings align with the recommended settings. The security hardening section shows this percentage, as well as the history of the score on your instance over time.

Security scanner

Review the results of automated security scan checks using the security scanner section. These checks and help you find security vulnerabilities, inefficiencies, and issues with compliance.

Select the cards on in this section to navigate to the security scanner page.

Customer actions

Find security updates that can be made to the instance in the Customer Actions section. This section shows the number of available actions and a chart indicating the due date for these actions. Select the cards on in this section to navigate to the Customer Actions page.

My Security Tasks

View the most urgent Security Tasks assigned to you. Select a task to view its details, or select **See all Security Tasks** to view a complete list of Security Tasks.

Additional resources

Use the additional resources section on the edge of the screen to navigate to documentation and videos relevant to learning Security Center as well as information on managing instance security.

Security hardening

View your hardening compliance score, compare it with previous scores, and change settings to improve your compliance score and security posture in the security hardening page.

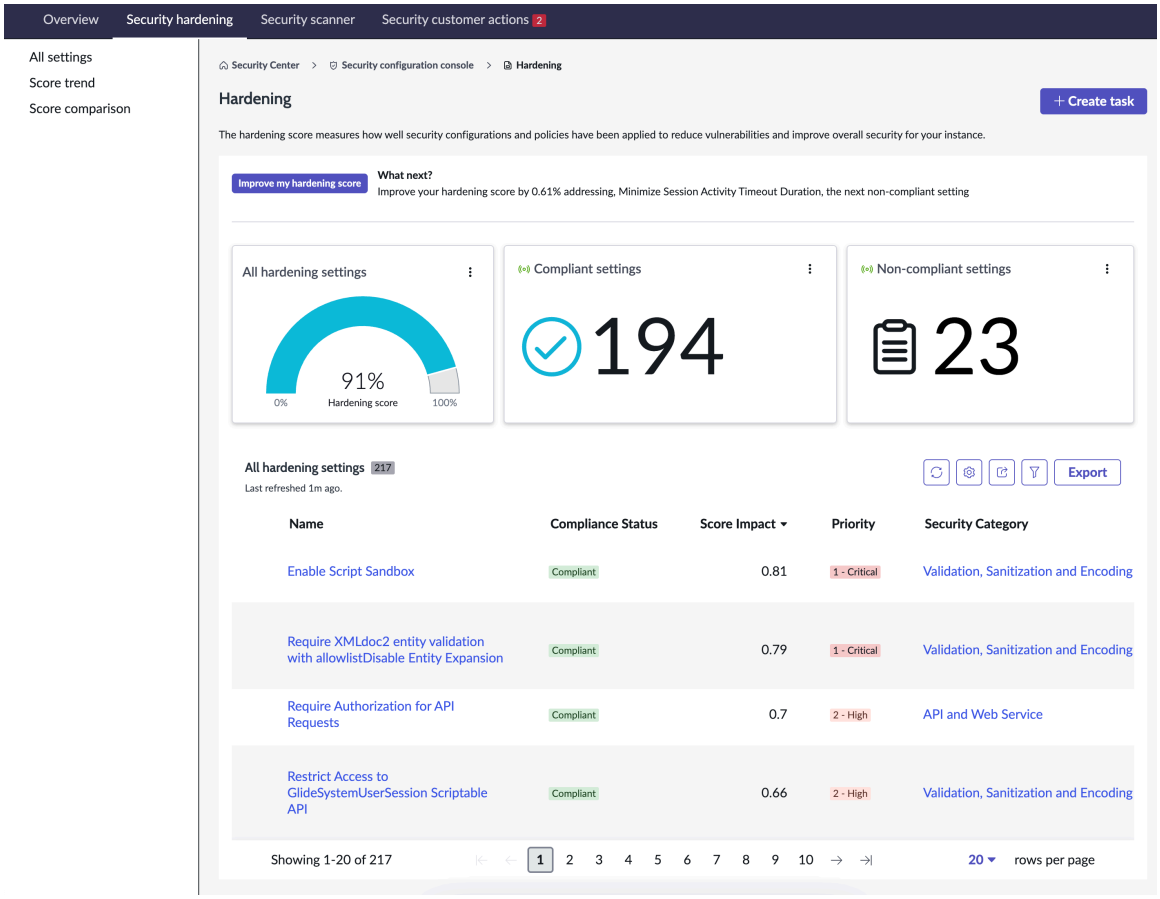

Hardening settings specify recommended values for the security-related properties and plugins in the ServiceNow AI Platform. The hardening tool calculates the hardening setting compliance score as a percentage. This number indicates how compliant your instance is with the Security Center hardening settings.

The formula for calculating the hardening compliance score:

- Each hardening setting has a risk score between 0-10. You can see the values of individual settings in the **All settings** section.
- The score equals the sum of all the compliant risk scores divided by the sum of all risk scores.

For example, the sum of all the compliant risk scores is 25.4, and the total for all risk scores is 34.9. For this sum, the compliance score is (25.4 /34.9) × 100 which equals 72.7. This decimal gets rounded up to the nearest whole number and will therefore be equal to 73.

This calculation is automatically performed on the first of the month, or after an installation or reinstallation of ServiceNow Security Center. You can trigger a recalculation at any time using the **Update score** button on this page.

Security hardening subsections

The security hardening section contains three subsections that you can select on the left edge of the screen.

- All settings
- Hardening compliance score trend
- Hardening score comparison

All settings

Review all of your instance hardening settings available from a single page.

The **All Settings** page displays information about all instance hardening settings. Use the buttons in the upper right to refresh, filter, and export this information.\

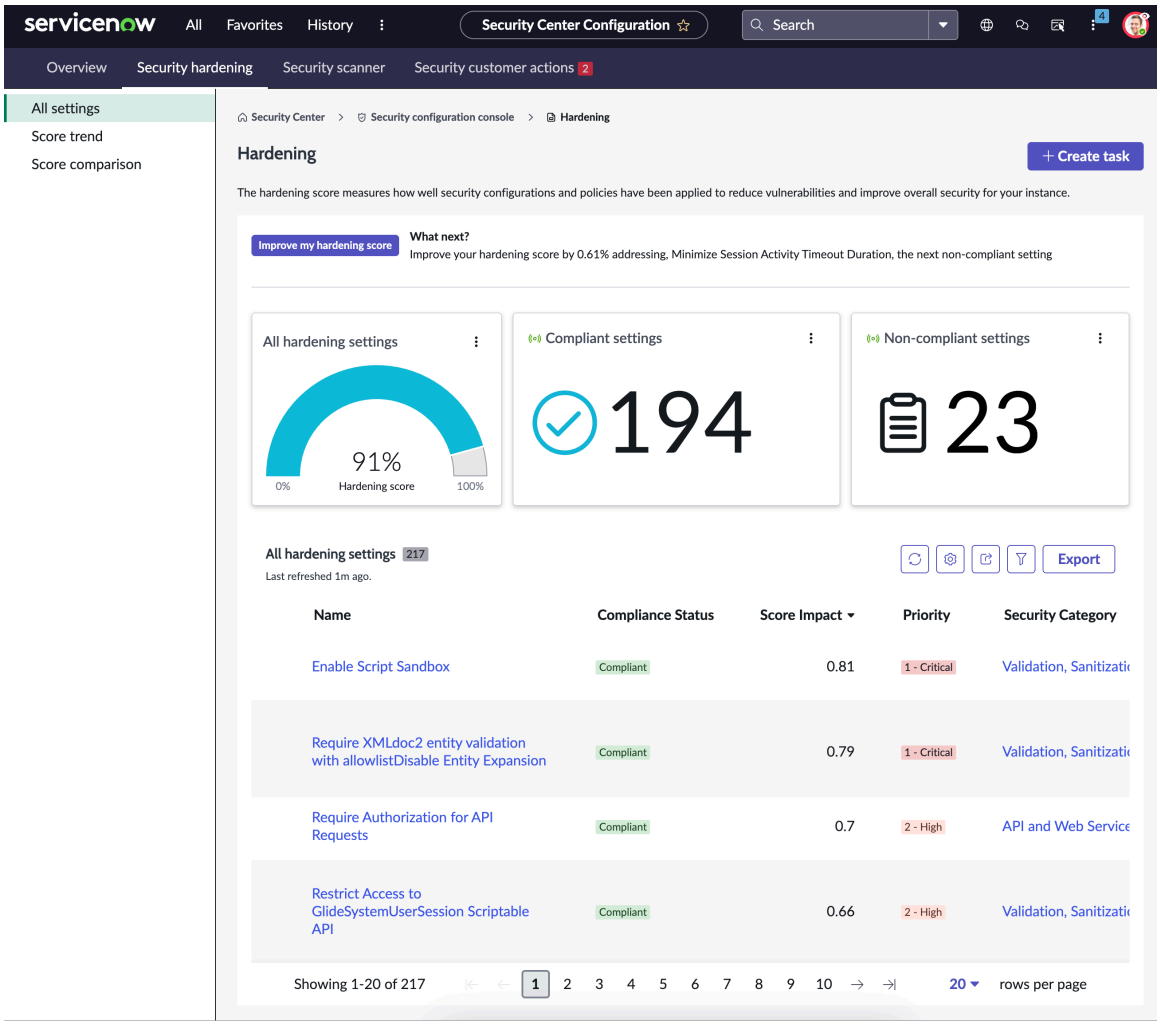

Use the **+Create task** button to create a Security Task to complete a Customer Action. This button appears both in the **All Settings** page as well as the pages for individual hardening settings. For details, see Security Tasks.

The following information about hardening settings can be found on the list.

Name

Name of the hardening setting.

Compliance Status

Whether the setting is properly configured according to the system's recommendation (compliant) or must be configured (non-compliant)

Score Impact

Impact this hardening setting has on your security posture, expressed as a percentage. All score impacts sum to 100%.

Priority

The criticality of the setting: Critical, High, Moderate, and Low. A higher score indicates a greater impact and priority.

Security category

The security category of the property. Select the category to view details.

Resolution Details

Description of the steps to remediate the security vulnerability of the hardening setting.

To learn how to configure a property hardening setting see Increase hardening compliance score.

Hardening settings details

Analyze the details of a hardening setting by selecting its link within the Security Center app.

Navigate to **Hardening** > **All settings**, and then select a hardening setting to be redirected to a page within the hardening setting tool that displays its security-related information.

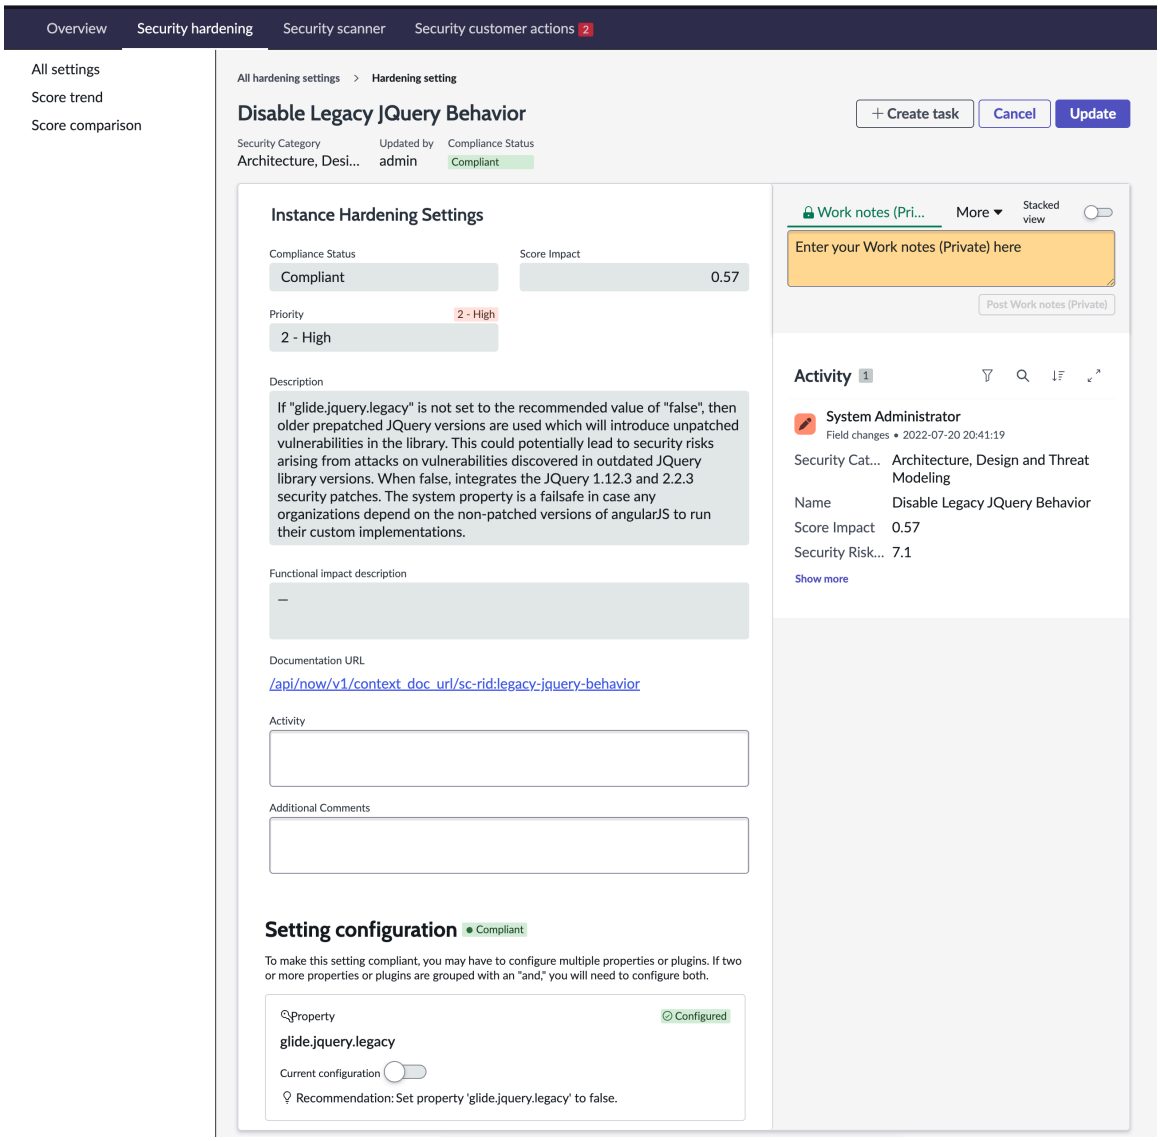

Use the **+Create task** button to create a Security Task to complete a Customer Action. For details, see Security Tasks.

Hardening settings configuration details

Configuration attribute	Description
Compliance status	Indicates if the hardening setting is compliant nor not.
Score impact	A percentage that indicates how much this hardening setting impacts your security posture.
Priority	A number within the range of 1-4, with 1 having the most weight that indicates the strength of impact this hardening setting has on your security posture.
Functional impact	The impact this hardening setting has on the operation of your instance.
Description	A general overview about the hardening setting.
Documentation Url	Link to documentation for the hardening setting.
Activity	Notifications of updates related to the hardening setting.
Setting configuration	Details related to the compliance status of your hardening setting along with instructions on how to make them compliant. **Note:** Some hardening settings may require you to configure multiple properties and plugins to make them compliant.

Filter hardening settings

Simplify your hardening review process using filters. These filters can create a working list of hardening settings for review, which restored for later use and shared with other users.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Security** > **Security Center**.
2. In the **Security Consoles** section, select the **Security configuration console**.
3. Select the **Security Hardening** tab.
4. Select **All settings**.
5. Select the **Show filter panel** button ( ).
6. In the **Filter** window, select **Advanced view**.
7. Create a filter showing only the hardening settings you want to review.
8. Once your filter is complete, select **Save Filter**.

9. Enter a name for your filter and select **Save**.

> **Note:** You may also make this list available to other users by setting permissions to **Everyone**, or **Group** to share with a specific group. These users must also have permission to use Security Center.

What to do next

After saving a filter you can load it in the **Advanced View** using the **Use existing filter** button.

Hardening compliance score trend

View the trends of your hardening compliance score over time in a chart or table.

See trends in your compliance score over time. Use the date picker to select a time range to analyze to gain insights into the data by applying performance analytics functionality such as targets, thresholds, and KPIs.

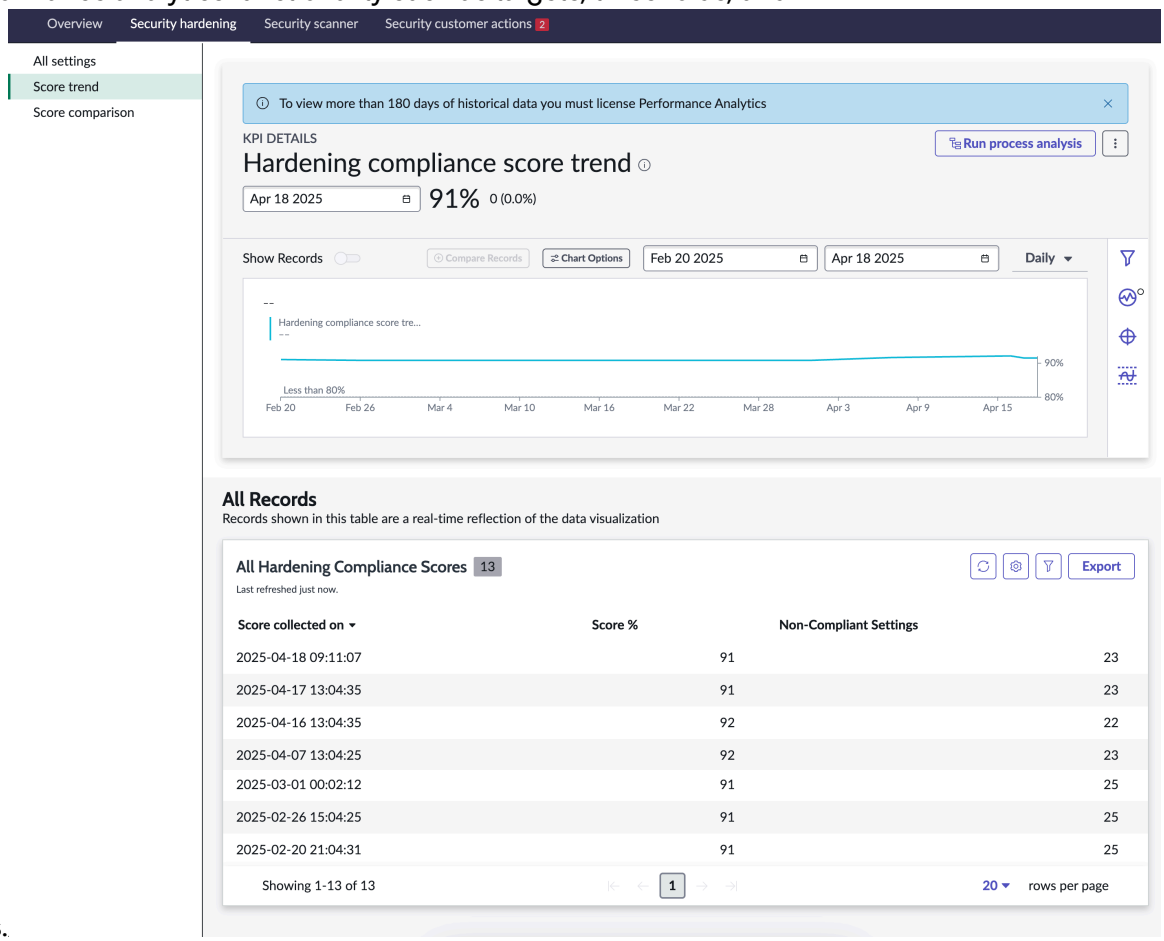

Compliance score chart

Use the **Chart options** button to specify what information to analyze and how to present the information.

Analysis

Select from the list of options. Some options may be unavailable depending on the selected chart type.

Target

Goals your organization wants to achieve. See Indicator targets.

Threshold

Defines a normal range of scores for an indicator and alert you when a certain event occurs. See Indicator thresholds.

Forecast

Describes the ability to forecast future scores based on past behavior. See Performance Analytics scores forecasts.

Trend

Shows how the value of one or more items change over time.

Comments

Displays annotations on individual data points.

Labels

Displays scores related to visualizations.

Statistics

Displays statistics related to your compliance score.

Time series

Select which metric to display on the chart.

Score

Score for the Key Performance Indicator (KPI).

Change

Change of score for this indicator.

Change percentage

Change as a percentage of scores.

Chart type

Select a chart type to control how the information your selected is visualized. See Use cases for different time series visualization types.

All Records table

The All Records table displays the date on which the compliance scores were collected, the score percentages, and the number of non-compliant settings. You can use this table as another option to analyze the security posture of your instance over time.

Increase hardening compliance score

Increase your hardening compliance score by ensuring that the hardening settings are compliant with the system's recommendations.

Before you begin

Role required: admin

Identify non-compliant hardening settings with the highest score impact on your instance. Review them to see whether you can comply with system recommendations so that you can increase your overall compliance score.

Procedure

1. Navigate to **Hardening** > **All settings**.
2. Filter Compliance Status column to view only non-compliant hardening settings.
3. Select **Score Impact** to sort from largest to smallest.

4. Select settings and review the setting details to decide whether you want to comply with recommendations.

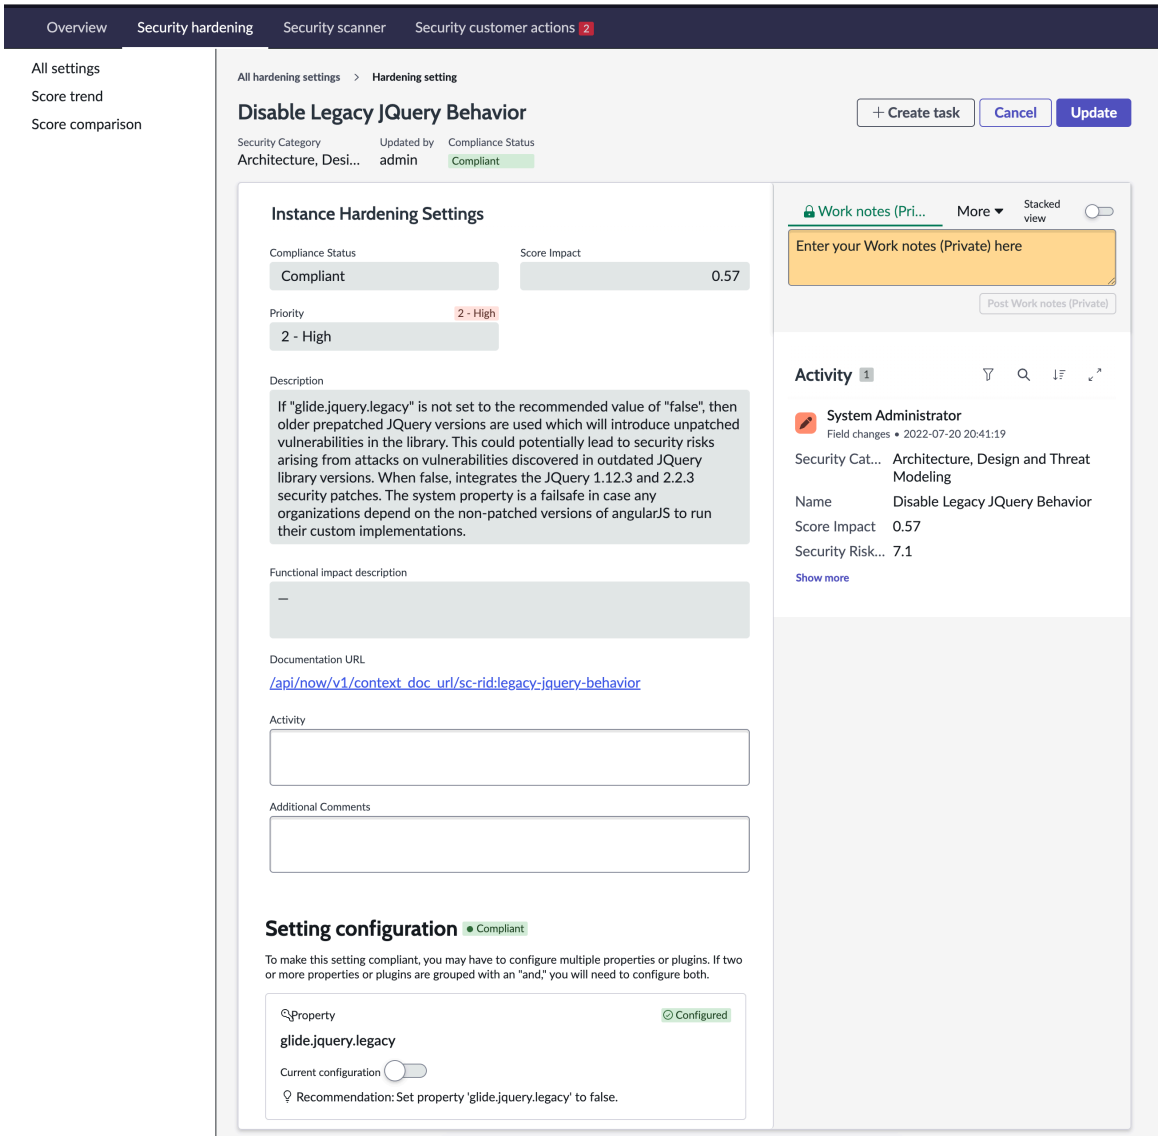

5. Make the hardening setting compliant.
6. If you update a non-compliant hardening score to make it compliant, go to the Homepage and select **Update score** to view the most up-to-date score.
   The hardening score is rounded up. A score of 86.75% will be rounded up to 87.

Hardening score comparison

Gain visibility to the health of your hardening settings and use this data to improve the security posture of your instance.

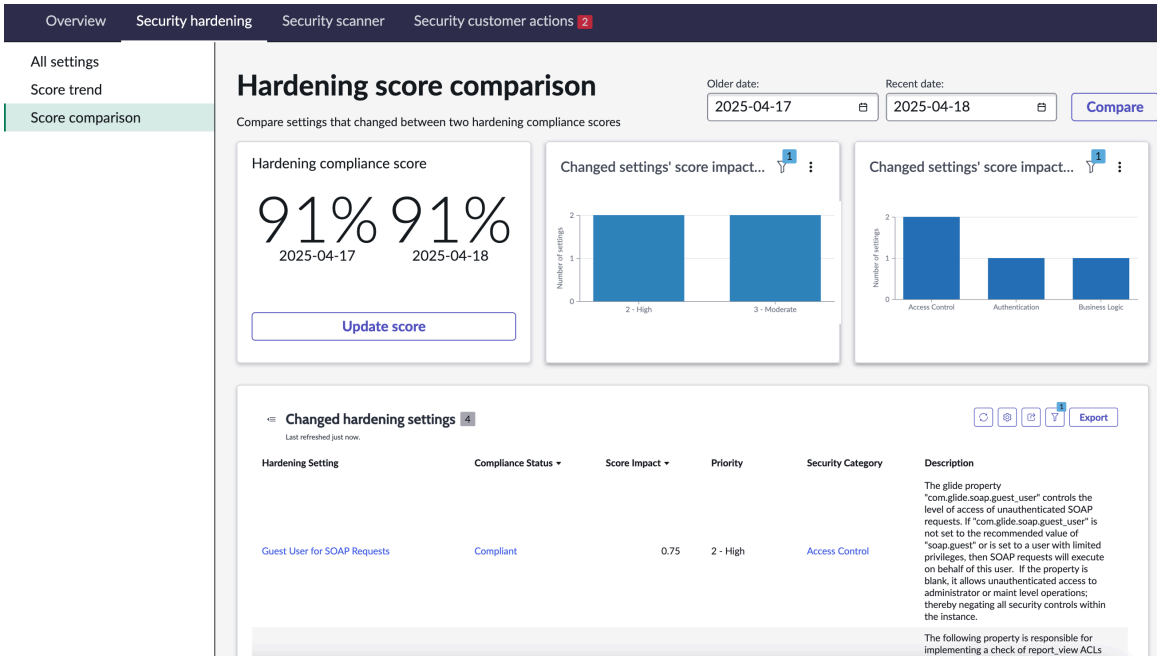

This page displays a summary of your hardening compliance and the changes made to your instance that impact this score.

The hardening score comparison page displays hardening settings with changes in compliance status between the dates selected in the **Older date** and **Recent date** fields. With this information you can see how they've impacted your hardening compliance score.

Hardening compliance score

The card compares the compliance score of your instance on the selected dates as a percentage.

Changed settings' score impact by priority

Displays the number of hardening settings that have changed their compliance status since the last score update, organized by priority value.

Changed settings' score impact by security area

Displays the number of hardening settings that changed compliance status between the two selected dates, organized by security area.

Changed hardening settings

Displays the list of hardening settings that have changed compliance statues between the selected dates. Review settings that became non-compliant and decreased your hardening comparison score for opportunities to make them compliant to increase your score. See Increase hardening compliance score.

Security scanner

Scan your instance against a set of security checks to identify misconfiguration. The scanner tool simplifies the process of creating different suites of checks for different use cases so that you can analyze the results over time.

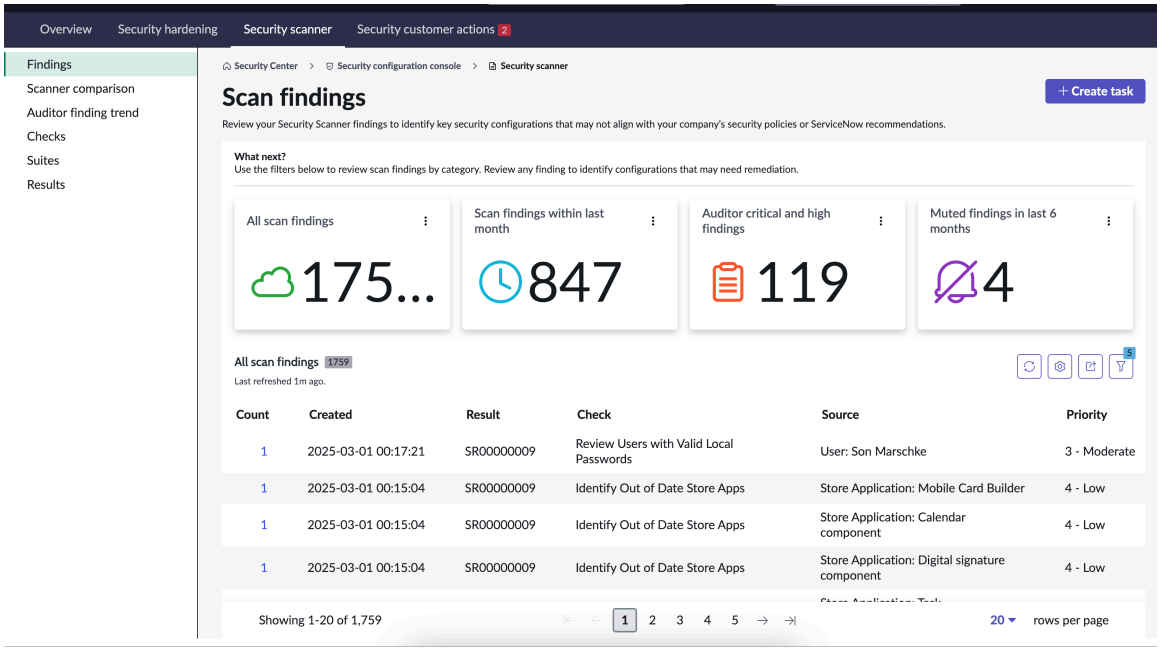

Security scanning is a method to investigate your instance for configurations that indicate security health issues. This method enables you to identify opportunities to implement security recommendations for your organization.

When accessing the scanner tool, there's no comparison available until you select the suite to be compared with at least two scan results of the suite. You can use the default suite and checks, or you can create your own custom checks and suites.

Scan findings

A finding is a reference to a record that has violated a rule from a check on the instance. You can find the source record and the number of times the record triggered the rules of a given check.

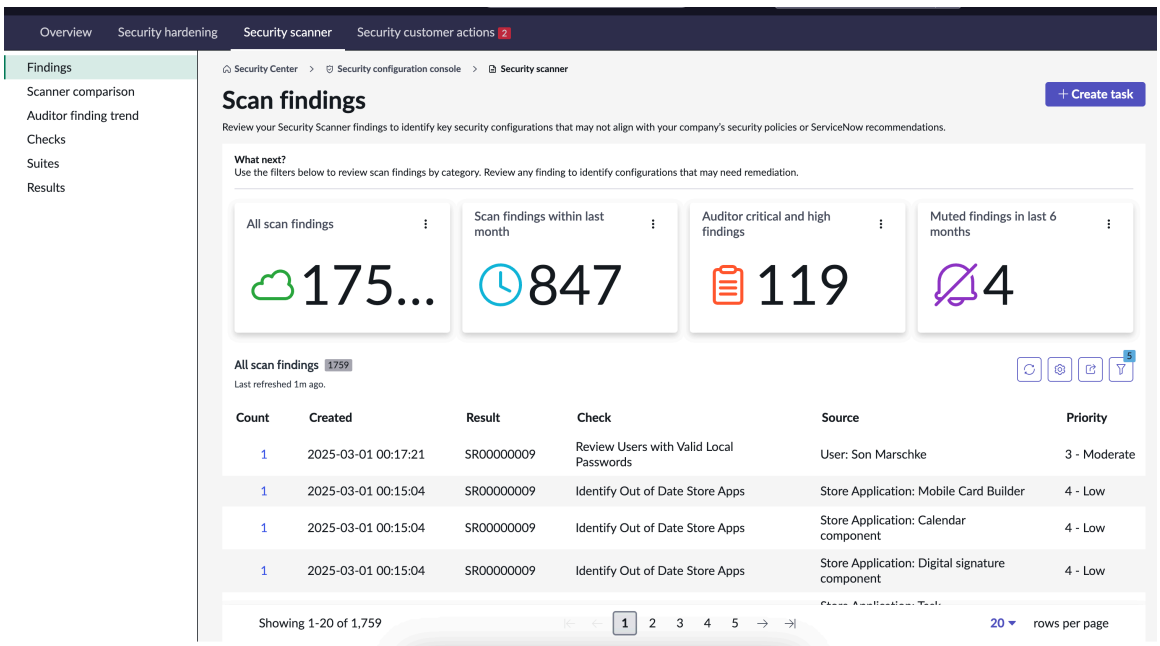

Navigate to the **Findings** tab to view scan findings in a list. The cards above the list provide a count of the findings that match specific criteria listed on the card. Select any of these cards to filter the list to show only those that match the criteria.

Select the **+Create task** button to create a Security Task to resolve a finding. This button appears both on the list and within the finding record. For details on Security Tasks, see Security Tasks.

Scan findings

Select a link under the **Count** column to view a finding record, which displays granular details related to a specific finding.

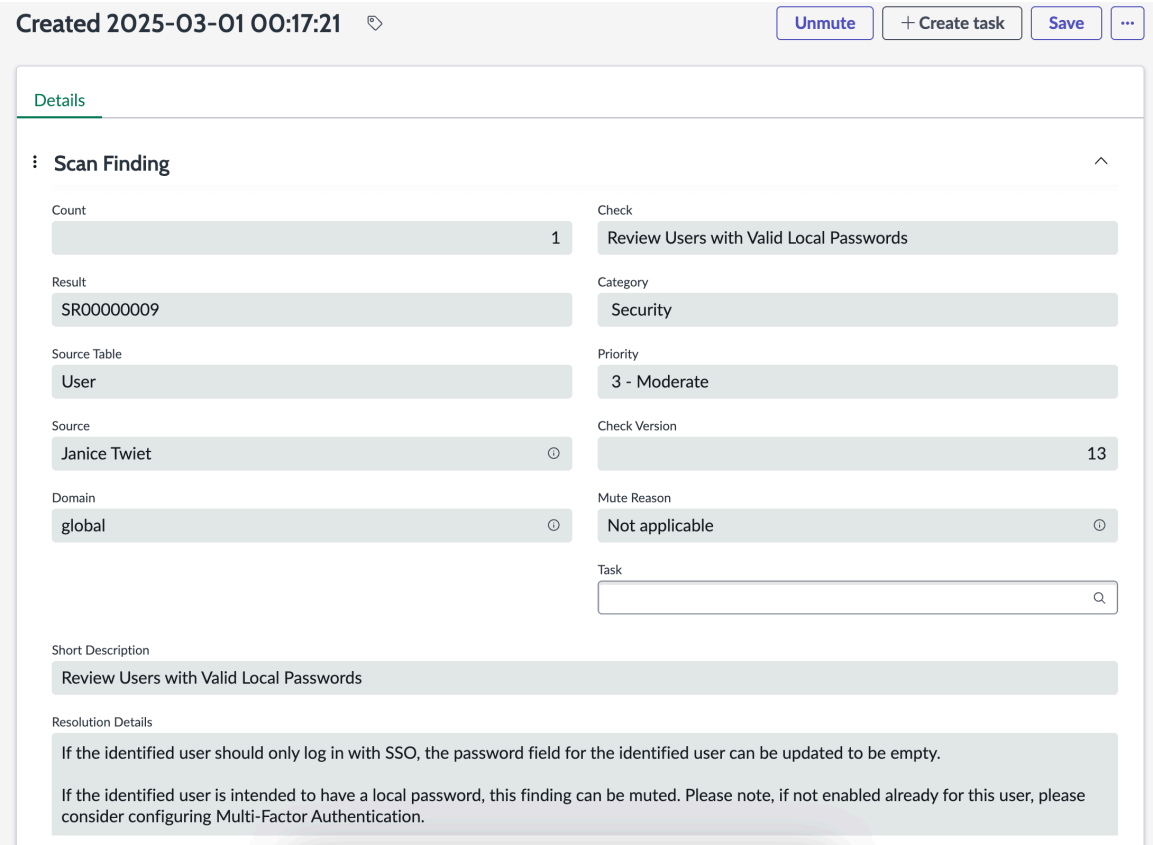

Check

List of checks associated with the scan.

Category

Security category associated with the scan. For example, access control or malicious code.

Count

Number of times the record violated check rules.

Priority

Severity of the security risk: 1 is highest priority while 4 is lowest.

Result

Status and type of scan.

Check Version

Version of the check.

Source Table

Record that has violated a rule from the check.

Mute Reason

Reason for muting the finding.

Source

Date the finding was created.

Task

Task record associated with the scan. Used to facilitate task assignments from the finding of a record.

Domain

Which domain the scan is applied to.

Short Description

Brief explanation of the scan.

Resolution Details

Instructions on how to resolve issues reported by the scan.

Findings can be muted by selecting the **Mute / Unmute** button. When muting a scan finding, you're prompted a reason for muting the finding. Findings muted in the last six months are available in the muted findings card in the **Scan Findings** page.

Security scan comparison

Compare two scans of the same security suite to gain visibility to the health of your hardening settings and improve the security posture of your instance.

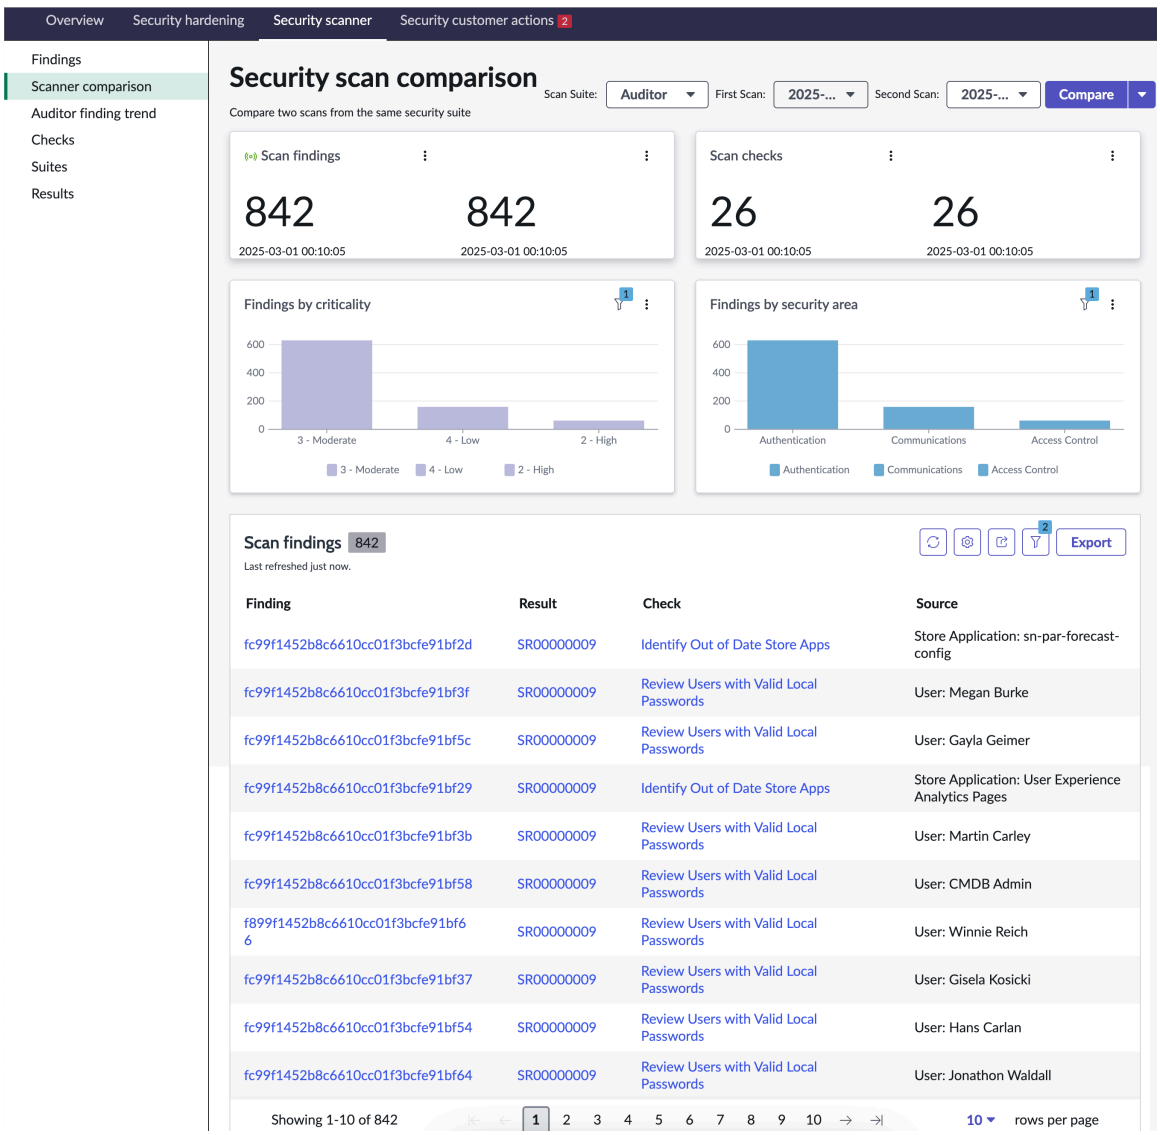

Important: This page appears empty if you don't have two or more instances of the same scan to compare. Run a scan at least twice to view a comparison on this page.

Select a scan in the **Scan Suite** list, then select a **First Scan** and **Second Scan** to begin a comparison.

The security scan comparison page displays the security changes in your hardening settings between the first and second scans. Below is an explanation of each card:

Scan findings

Number of findings in the selected scan at each selected date.

Scan checks

Number of scan checks performed in the selected scan at each selected date.

Changed findings by criticality

Chart displaying all findings organized by criticality.

Changed findings by security area

Chart displaying all findings organized security area.

Scan findings

List of all scan findings. Select a finding, scan result, or check to view the associated record and find additional details. Use the bottoms in the upper right to filter, refresh, or export the list.

Auditor suite findings

Review Security Center Auditor suite findings over time.

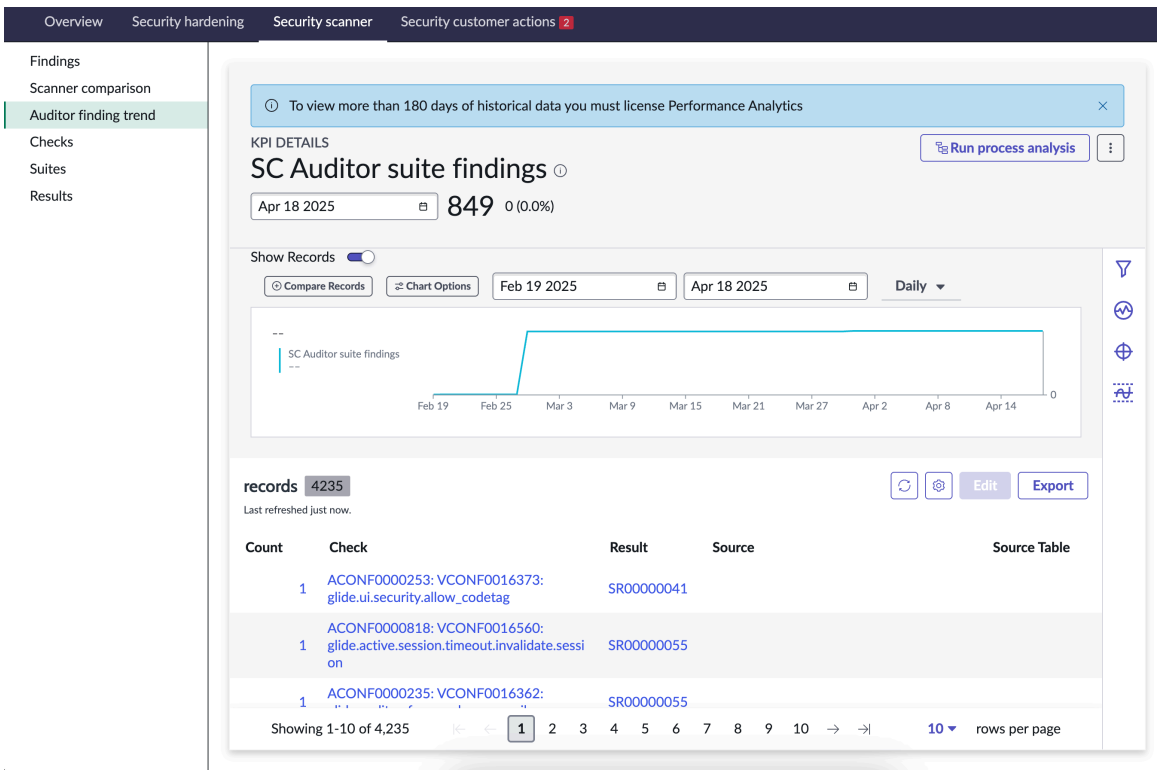

SC Auditor suite findings chart

The page shows SC Auditor suite findings compared between the two selected dates.Use the **Chart options** button to specify what information to analyze and how to present the information.

Analysis

Select from the list of options. Some options may be unavailable depending on the selected chart type.

Target

Goals your organization wants to achieve. See Indicator targets.

Threshold

Defines a normal range of scores for an indicator and alert you when a certain event occurs. See Indicator thresholds.

Forecast

Describes the ability to forecast future scores based on past behavior. See Performance Analytics scores forecasts.

Trend

Shows how the value of one or more items change over time.

Comments

Displays annotations on individual data points.

Labels

Displays scores related to visualizations.

Statistics

Displays statistics related to your compliance score.

Time series

Select which metric to display on the chart.

Score

Score for the Key Performance Indicator (KPI).

Change

Change of score for this indicator.

Change percentage

Change as a percentage of scores.

Chart type

Select a chart type to control how the information your selected is visualized. See Use cases for different time series visualization types.

Records list

Lists the checks performed as part of the suite.

Scan checks

Use checks to detect anomalies within an instance, running against tables, records, or metadata.

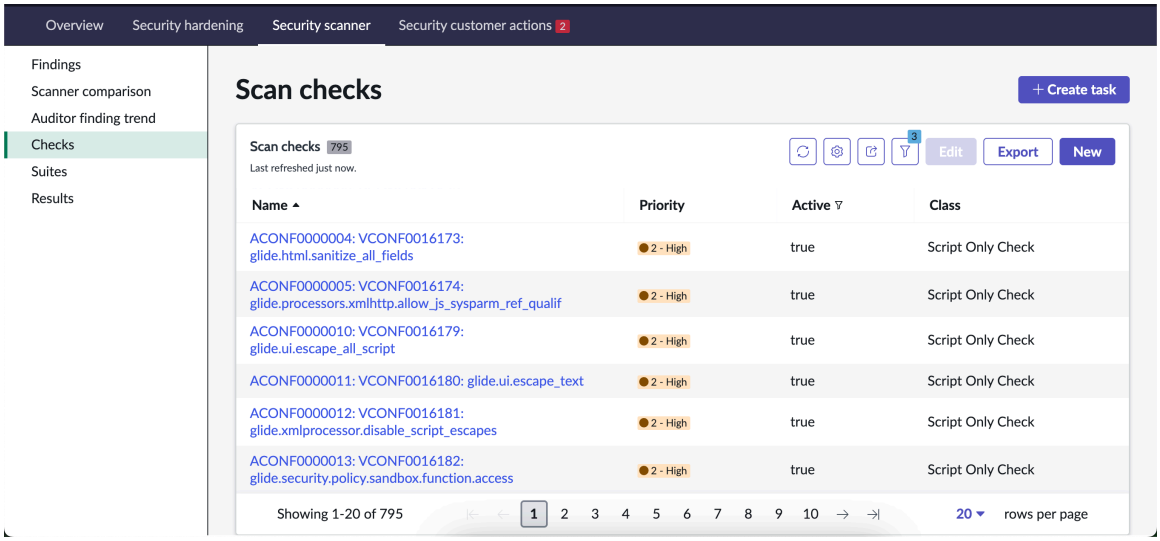

Checks are rules designed to detect anomalies within an instance. Select a check on the list to view details including what the check evaluates, and possible steps to correct the issue if the check returns any findings.

Select the **+Create task** button to create a Security Task related to a scan check. For details on Security Tasks, see Security Tasks.

Checks are divided into four classes to identify issues and implement security recommendations for your organization: table checks, column type checks, script only checks, and linter checks.

Check classes

Check class	Description
Table checks	Use this check class when you know the specific tables and checks you want to test.
Column checks	Use this check class to implement the rule you created to iterate all records matching the target column field type.
Script checks	Use this check class to verify meta data, configurations, and execute complex checks by writing your own checks.
Linter checks	Use this check class to identify any issues in a script. When a linter check is run on a record, an abstract syntax tree for its code is generated which can be used to analyze issues with the code.

Scan suites

Review details on the scan suites available on your instance.

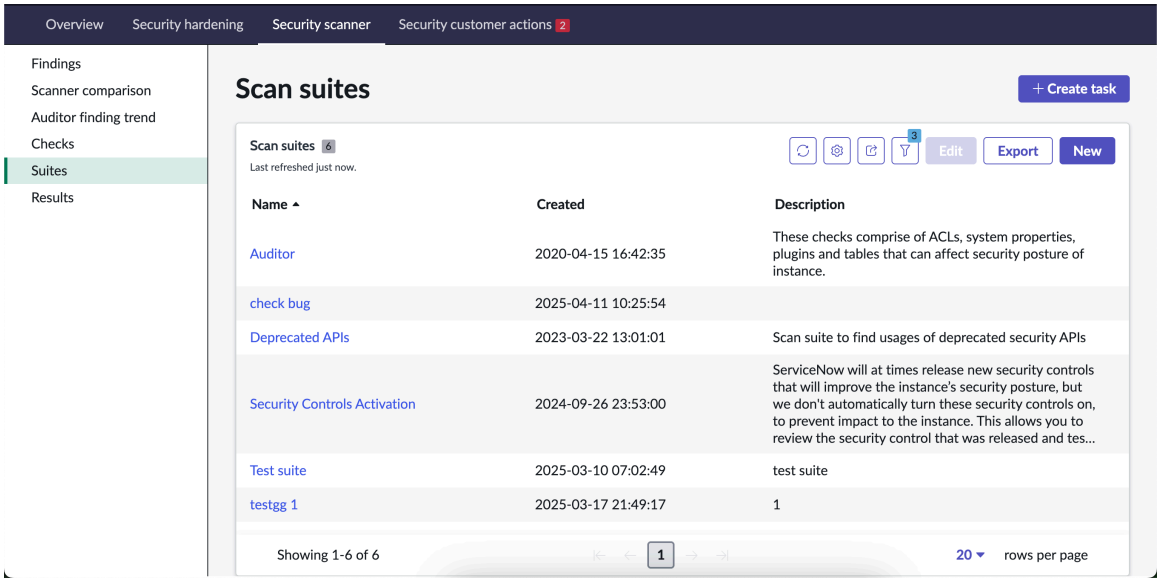

Scan suites are collections of security center checks that execute together. You can use base system suites or create your own by cloning an existing suite and updating the checks made in the clone. For details, see Create a scan suite.

Select the **+Create task** button to create a Security Task related to a scan suite. For details on Security Tasks, see Security Tasks.

Scan suite details

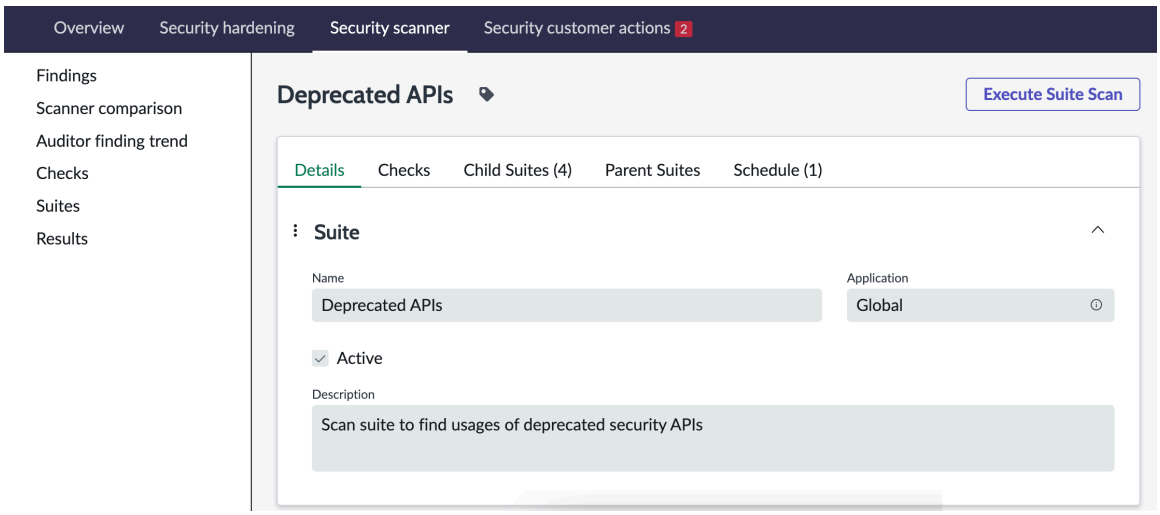

Select the **Name** field of a suite to view the suite details. This page provides details on the scan suites in a tabbed interface. In the tabbed page, you can find this information:

Details

Details on the suite include its name, application, and description.

Checks

List of the checks included in this suite. Details on individual checks are the same as are found in Scan checks.

Child Suites

List of child suites associated with this suite. When you execute this suite, all child suites also execute.

Parent Suites

List of suites in which this suite is a child suite. When you execute a parent suite, this suite is also executed.

Schedule

Details of the scheduled execution of this suite.

Access Controls Auditor checks

Learn about the checks available in the default Access Controls Auditor Suites, what criteria they evaluate, and how they can be used to improve the security of your instance.

Access Control List rules (ACLs) restrict access to data by requiring users to pass a set of requirements before they can interact with it. Access Controls Auditor checks evaluate your instance according to the eight criteria listed in the following table. Use the findings on these checks to improve the security of your instance.

Access Controls Auditor checks

Check Name	Check Criteria	Description
All Processors of type - SCRIPT must be protected with CSRF Token	Checks for Processors with the **SCRIPT** type that aren't protected with a CSRF token.	All Processors with the **SCRIPT** type should be protected with a Cross-site Request Forgery (CSRF) token. These processors should have the CSRF option checked, which prohibits the processor from running unless the instance uses a CSRF token.
Can Contribute / Cannot Contribute user criteria to be defined on each knowledge	Checks for knowledge base records that don't have **Can Contribute** or **Cannot Contribute** user criteria defined.	Each knowledge base should have either **Can Contribute** or **Cannot Contribute** user criteria defined. Otherwise, any user can contribute content to a knowledge base with no Contribute criteria defined.
Empty ACLs	Checks for Access Control List (ACL) records that have no security attribute, no role, or the **public** role.	Leaving ACLs empty or using the **public** role may provide open access to any content protected by this ACL.
Access Controls on Client callable Script Includes	Checks for client-callable script includes that aren't secured by ACLs.	All client callable script includes should be secured with an ACL using required roles.
Access controls on UI Pages	Checks for UI Pages that aren't secured by ACLs	Without an ACL securing access to a UI Page, that UI Page is accessible to all logged-in internal users. Without any restrictions logged-in users can potentially make unauthorized changes.

Access Controls Auditor checks (continued)

Check Name	Check Criteria	Description
Access controls on Tables	Checks for tables without ACLs	Tables should be secured with ACLs. Access to data stored in tables should be limited only to users that need it.
User Account shouldn't have both Internal and External roles	Checks for user records with both **Internal** and **External** roles assigned	Internal user roles are intended for users within your company. External user roles are intended for external personnel, such as customers and partners.
Publicly accessible knowledge base and articles	Checks for publicly accessible knowledge bases and knowledge base articles	Publicly accessible knowledge bases and articles are visible to all users in the instance. Increase security by limiting knowledge bases and articles to the specific audience that needs them.

Auditor checks

Use the Auditor suite to SecureCheck to detect misconfiguration that can impact the security posture of your instance.

Check information

Access Controls Auditor checks

Check Name	Description	Scan finding type
Identify out of date store apps	Identifies apps activated on your instance that have updated versions available. Verify you're running to the most up-to-date versions of store applications, which can include fixes to potential security issues.	Resolution Recommended
Insecure GlideRecord calls	Identifies scripts that are directly invokable by end users (such as Client-Callable Script Includes, Widgets, Processors, REST Endpoints) These scripts should respect ACLs and use GlideRecordSecure or GlideRecord with canRead, canWrite, canCreate, canDelete.	Resolution Recommended
Review allowed JavaScript libraries	Identifies scripts where JavaScript Content Access Control is used to allow or deny specific third-party JavaScript libraries.	Resolution Recommended

Access Controls Auditor checks (continued)

Check Name	Description	Scan finding type
	Review instance customizations to verify that libraries aren't in use before blocking access. The JavaScript Content provider Access Tracking [sys_js_content_provider_access_tracking] table can be reviewed to see the last date that the library was accessed. **Note:** This check can be ignored in instances initially provisioned on Tokyo or later. Records on the associated table have deny rules set by default. In instances initially provisioned prior to Tokyo, there may be allow rules in the JavaScript Access Control tables.	
Review client callable script includes with no corresponding ACL	Identifies client callable script includes that don't have a corresponding ACL. These scripts use the default ("*") client callable script include ACL. For these scripts, create ACLs that defines the appropriate criteria for access to verify that only expected users can interact with the functionality provided.	Resolution Recommended
Review custom tables with record producers and no business rule	Identifies record producers that don't have additional server-side validation. This check identifies custom tables with a Record Producer but without an associated business rule. These may enable users to submit unexpected data into the associated table.	Resolution Recommended
Review empty ACLs	Identifies ACL records which have no script, condition, security attribute, or role, or ACLs with the **public** role. Leaving ACLs empty or using the **public** role provides open access to any content protected by this ACL.	Resolution Recommended
Review fields with HTML Sanitization disabled	Identifies HTML fields where HTML Sanitization is inactive. HTML sanitization removes or replaces potentially harmful elements and attributes within HTML code. Review HTML fields where	Resolution Recommended

Access Controls Auditor checks (continued)

Check Name	Description	Scan finding type
	sanitization is inactive to confirm whether this configuration is necessary.	
Review inactive security feature plugins	Identifies plugins that aren't activated that provide additional, configurable security controls. The findings produced by this check are provided for informational purposes. Before enabling one of the identified plugins, verify that the plugin meets your use cases or requirements. You can mute these findings if you don't have a use case for the identified.	Inform
Review large allowed IP address ranges	Identifies IP Address Access Control Ranges that contain a large number of IP Addresses. **Note:** If you're seeing a large number of false positives, consider adjusting the largestExpectedCIDRBlock variable for your specific business needs. Classless Inter-Domain Routing (CIDR) blocks contain a larger amount of IP addresses as the number decreases. For example, the CIDR block size 8 is larger (contains more IP addresses) than the CIDR block size 16. Review and confirm that the current configuration aligns with your business needs.	Review and Decide
Review public GraphQL schemas	Identifies public GraphQL schemas in the GraphQL API [sys_graphql_schema] table. These schemas can be configured to be available without authentication. Depending on the endpoint's functionality, this may allow unauthenticated users to perform unexpected actions or interact with unexpected data.	Review and Decide
Review public knowledge base articles	Identifies knowledge bases and knowledge base Articles configured to enable access to unauthenticated users. Review and confirm that the current configuration aligns with your business needs.	Review and Decide

Access Controls Auditor checks (continued)

Check Name	Description	Scan finding type
Review public REST API endpoints	Identifies Rest API Endpoints in the Scripted REST Resource [sys_ws_operation] table that are configured to be available without authentication. Depending on the endpoint's functionality, this may allow unauthenticated users to perform unexpected actions or interact with unexpected data.	Review and Decide
Review public Service Portal pages	Identifies Service Portal pages that are made public. Service Portal pages are made available to unauthenticated users by setting the "public" field to "true." Review and confirm that the current configuration aligns with your business needs.	Review and Decide
Review public UI Pages	Identifies UI Pages that are made public. UI Pages can be made available to unauthenticated users using the [sys_public] page. Review and confirm that the current configuration aligns with your business needs.	Review and Decide
Review roles that contain the 'admin' role	Identifies any roles (Roles [sys_user_role] table) that contains the **admin** role. The **admin** role grants users administrative privileges and should be used only when necessary. Review and confirm that the current configuration aligns with your business needs. If this is an intentional configuration, this check can be muted.	Review and Decide
Review UI Pages without corresponding ACLs	Identifies UI Pages that don’t have an ACL for that UI Page. UI Pages that don’t have a specific ACL default to a generic UI Page ACL, which may grant access to unintended users.	Resolution Recommended
Review users with valid local passwords	Identifies users with locally set passwords. Users with local passwords may interact with the instance via APIs using the local credentials, even if local logins are disallowed.	Review and Decide

Access Controls Auditor checks (continued)

Check Name	Description	Scan finding type
	This password configuration is needed for integration user accounts to function correctly. Review these user accounts to verify that only intended users (such as integration accounts) can authenticate with local authentication.	
Rotate passwords stored with outdated hashing algorithms	Identifies user accounts with passwords created in previous versions of the ServiceNow AI Platform, which may have used what is now considered a legacy or outdated hashing algorithm. Accounts created on old platform versions that haven't rotated their passwords may still have passwords stored with a legacy hashing algorithm. Review the identified accounts created consider password resets.	Resolution Recommended
Securing record producers	Identifies insecure record producers. If not assigned to appropriate roles unauthorized users can access them, potentially revealing sensitive information. Assign appropriate roles to record producers to verify that they're accessible only to users that need them.	Resolution Recommended
UI action visibility	Identifies UI actions that can be accessed by a user with no roles who doesn't have read access to the table. These users may be able to alter data on a table they don't have access to via these UI actions. Verify that UI actions are only available to users with access to the table they affect.	Resolution Recommended

Create a scan suite

Create and schedule a custom suite so that you can analyze the security of your instance for your organization.

Before you begin

Role required: admin or sn_vsc.security_center_viewer.

A suite is a collection of checks that can be used for a scan. View a list of scan suites organized in a table by navigating to **Scanner** > **Suites**. Create your own suites or to use the default suite, Auditor. Auditor is a default base system suite that contains checks for security best practices. These checks consist of system properties, plugins, and tables that can affect the security posture of your instance. The following steps show how to create a suite along with the several options available for configuring them.

Procedure

1. In Security Center, select the **Scanner** tab, then select **Suites** in the panel on the left side of the screen.
2. In the **Scan Suites** page, select the **New** button.
3. Enter a suite **name** and **description**, and then select **Save**

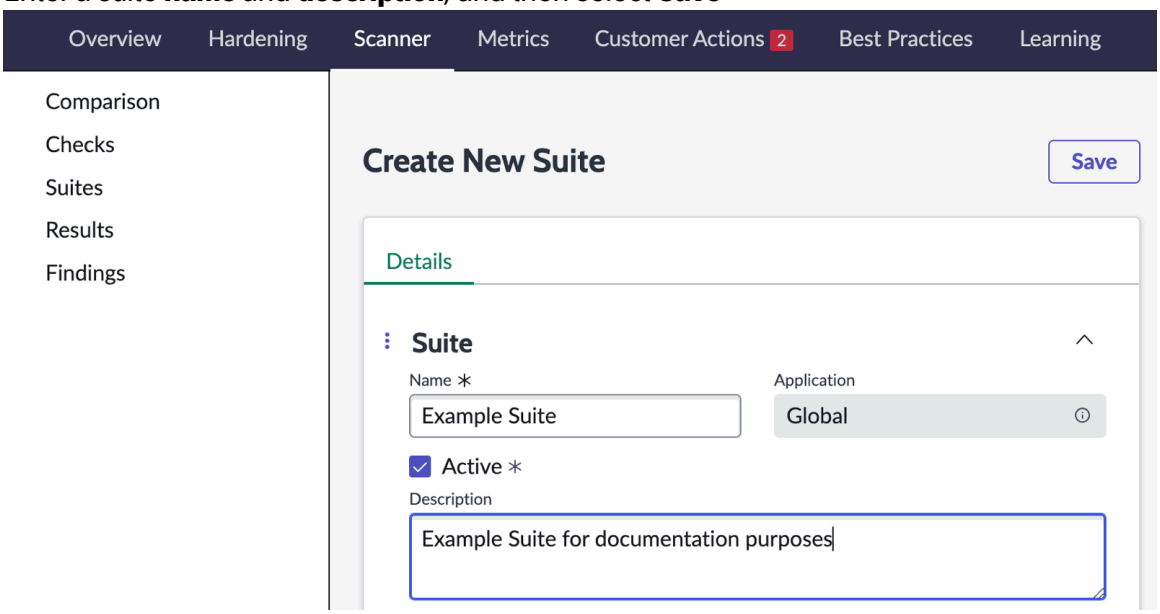

After selecting **Save**, configuration options for the suite display in a tabbed interface.

4. Select the **Checks** tab.
   Use this tab to add checks to your suite.
   a. Select **Edit**.
   b. Select the checks that you want to add, and then select Add (**>**) to place it in your suite.
   c. **Save**.
5. Select the **Child Suites** tab.
   Use this tab to add child suites. Suites added as child suites are also executed.when this suite is used in a scan.
   a. Select **Edit**.
   b. Select the child suites that you want to add, and then select the Add (**>**) to place it in your child suite.
   c. **Save**.
6. Select the **Parent Suites** tab.
   Use this tab to add parent suites. Parent suites that are executed in a scan will also execute this suite.
   a. Select **Edit**.
   b. Select the child suites that you want to add, and then select Add (**>**) to place it in your parent suite.
   c. **Save**.
7. Select the **Schedule** tab.
   Use this tab to set a time for your suite to execute.

a. Select **New**.

b. Enter details of the scheduled scan.
The time fields are in the format: hour:minutes:seconds.

c. **Save**.

Clone the access controls auditor suite

Clone and customize the default access controls auditor suite in your instance to create a new suite tailored to your organization's security practices.

Before you begin

Role required: admin

About this task

The default access controls auditor suite provided with your instance can't be modified. However, if you want to add, remove, or edit checks to align with your organization's security practices, you can duplicate the access controls auditor suites. Copying the access controls auditor suites enables you to customize it and create a new suite based on the default one. The access controls auditor suites includes checks related to security best practices, covering system properties, plugins, and tables that impact the instance's security posture. The following steps demonstrate how to duplicate the default access controls auditor suites to tailor it to your organization's requirements.

Procedure

1. Navigate to **All** > **Instance Scan** > **Suites**.
2. Select **New** and enter a name for your suite, along with an optional description.
3. Right-click the form header and select **Save**.
4. Add the checks needed for your scan.

   a. In the **Checks** tab, select **edit**.

   b. Add the conditions.
   For example, to add scan checks apply the following fields, operators, values, and conditions:

   [Category][is][Security] AND [Application][is][Global]

   For example, it should look like the following: `Category is Security And Application is Global`.

5. Select **run filter**.
6. Select the scan checks that you want to add from the collection list to the check list, and then select the add (>) button.
7. Select **save**.
   The suite with your custom checks added have been created.
8. Select **Execute Suite Scan**.

View the Access Controls Auditor Suite

View the checks available in the default Access Controls Auditor Suites to understand which checks are executed when this suite runs.

Before you begin

Role required: admin

About this task

The steps you need to complete to access the default access controls auditor suites within your ServiceNow instance.

Procedure

1. In Security Center, select the **Scanner** tab, then select **Suites** from the panel on the left.
2. In the **Scan Suites** page, select **Auditor** from the list.
3. To review the suite checks, select the **Checks** tab.
   You should see the list of 8 checks available for the suite.
4. Select the name of a check you want to analyze.
5. Analyze the fields associated with the check.

Name	Description
Name	Name of the check.
Application	The application to which the check belongs (Security Center).
Category	Category associated with the check.
Priority	The level of urgency.
Version	Version number of the check.
Active	The status of the check, for example, active or inactive.
Short Description	A brief summary of the check.
Description	A more comprehensive summary of the check.
Resolution Details	Mitigated potential security incidents.
Documentation URL	Links to related documentation in the product documentation or Knowledge Base (KB) articles.
Run Condition	Conditions that trigger the start of the checks.
Table	The table to which the check belongs.
Conditions	Conditional logic applied to the checks.
Advanced	Advanced configuration options.

Reschedule a scan suite

Change the schedule of your scan suites to suit your needs.

Before you begin

Role required: admin, sn_vsc.security_center_viewer

Procedure

1. Navigate to **All** > **System Security** > **Security Center**.
2. In the **Tools** section, select **Security scanner**.
3. Select **Suites** in the menu on the left edge of the screen.
4. Open the suite you want to reschedule.
5. Select the **Schedule** tab.
6. Select an existing schedule, or select **New** to create a new schedule.
7. In the **Scheduled Scan** form, fill in the fields as needed.

Field	Description
Name	Name of the scan schedule
Run	Select the how often your suite will run. **Note:** Additional fields will appear to help define your schedule. These fields vary depending on your selection.
Active	Whether the schedule is active. Inactive schedules are never executed.
Conditional	Select to enable the **Condition** field, where you can define a scripted condition for your schedule.
Condition	A scripted condition for your schedule. The suite will only run when the script evaluates to `true`.
Run as tz	Which timezone the schedule uses when determining when to run.

Scan results

View data related to your scan results from a single view.

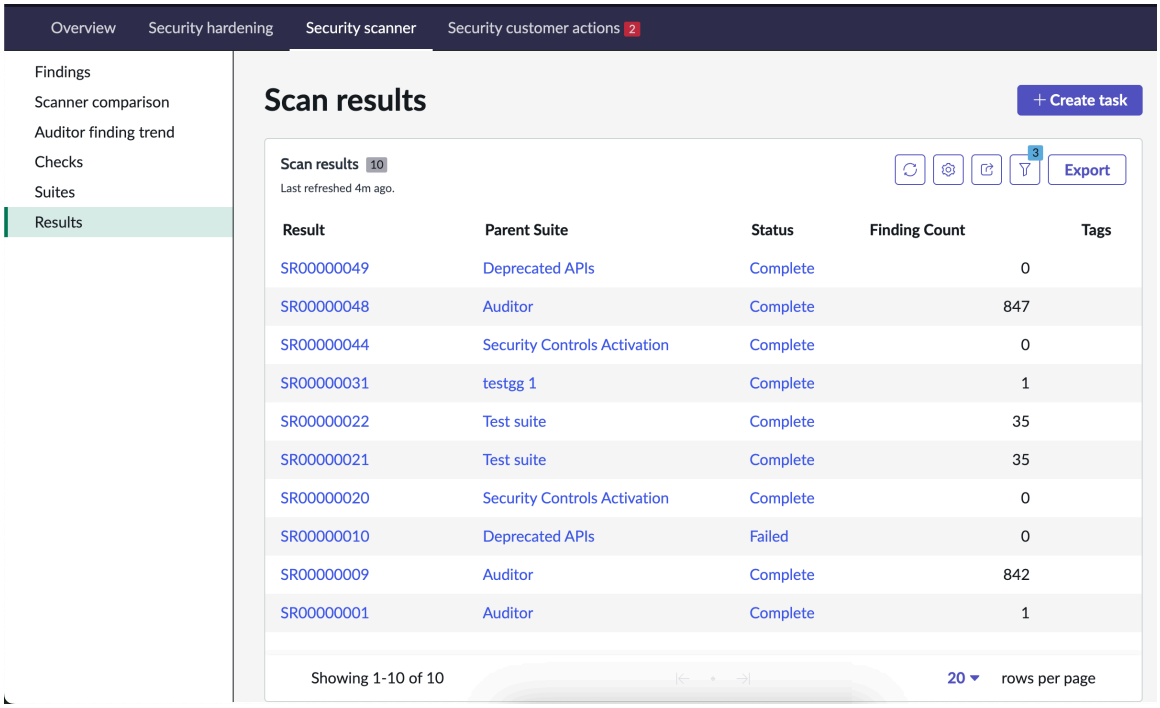

A scan result reports the status and type of the scan. You can also see all the checks that ran as part of the scan and all other information related to the scan such as errors and scan logs.

You can find your scan results dashboard by selecting the **Scanner** tab in Security Center, then selecting **Results** in the left panel. The scan results are listed under the **Scan results** list. Select a result to drill in and view the details, which are separated into tabbed sections.

Details

Overview of the scan results including the status, type, and execution time in seconds.

Scan Findings

Findings encountered during the execution of the checks.

Suites

Displays the suites that ran as part of this scan.

Checks

Displays the checks that ran as part of this scan.

Failures

Displays the checks that failed during the scan and the reason of its failure in the form of an error message.

Scan Log

Displays messages output during the scan.

Scan Statistics

Displays statistics related to the scan.

Targets

Displays all the targets against which the checks have executed.

Customer Actions

Use the Customer Actions tool to implement important security updates based on your instance and the configuration of plugins.

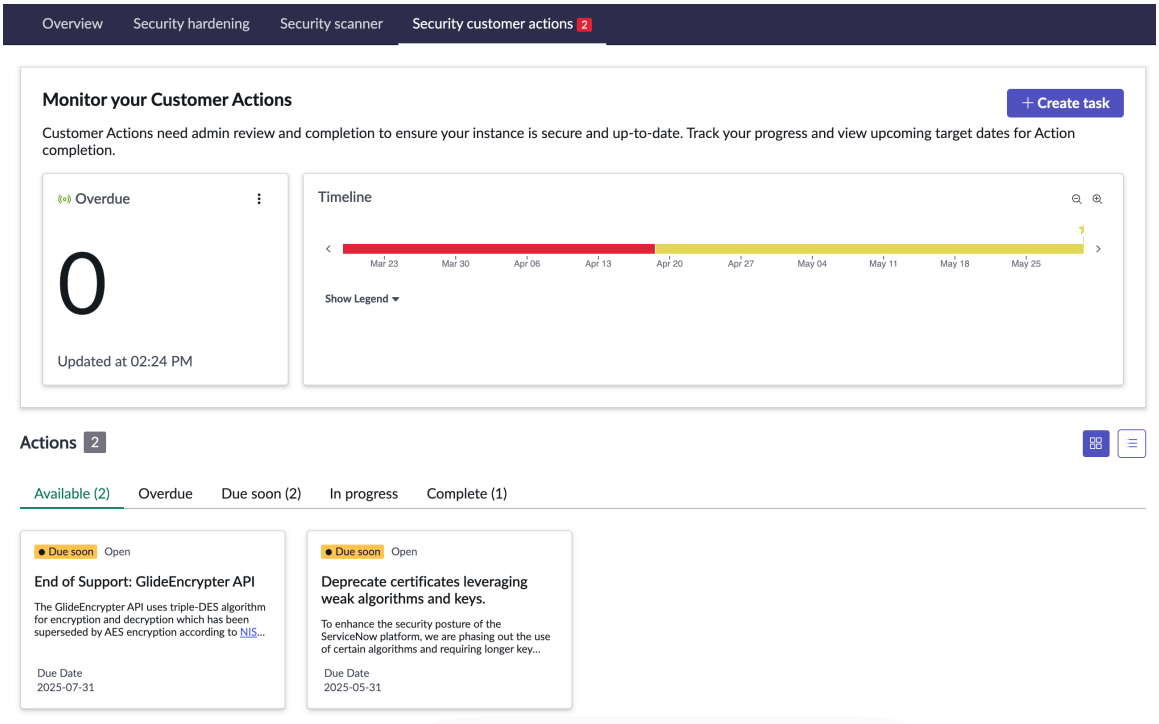

ServiceNow updates your instance with family releases and patches to keep base system platform code and functionality current and secure. However can be additional actions you can take to secure your instance that are unique to your custom instance configuration. These recommended updates are found in the Customer Actions tool.

Customer Actions vs. Family releases and patches

Feature	Customer Actions	Family releases and patches
Definition	Manual, guided tasks for admins to implement important security changes.	System-driven updates applied automatically during family releases or patches.
Purpose	Address security risks that can't be resolved by ServiceNow upgrades alone. Designed to avoid disrupting custom configurations.	Keep platform code and functionality current and secure.
Examples	Deprecating weak certificates, enforcing multi-factor authentication (MFA), disabling insecure protocols.	Platform version upgrades, patch-level fixes, backend security enhancements.

Access this tool by navigating to **Customer Actions** within Security Center.

Elements of the Customer Actions tool

Create Task button

Use the **+Create task** button to create a Security Task to complete a Customer Action. For details, see Security Tasks.

Overdue actions

A count of the number of Customer Actions that are part their due date. Select the overdue actions card to see a list of past due Customer Actions.

Customer Actions timeline

A timeline of Customer Actions so that you can prioritize when to implement the steps. The timeline in the preceding image shows the Customer Actions that are overdue, and the ones that are due soon.

Actions tabs

A tabbed interface showing Customer Actions by category:

- Available
- Overdue
- Due soon
- In progress
- Complete

Select a tab to view a list of actions in each category.

For details on how to review individual Customer Actions, see View Customer Actions.

For information on how to apply the changes recommended by Customer Actions, see Implement Customer Actions.

Implement Customer Actions

Learn how to implement Customer Actions on your instance to increase its security posture.

Before you begin

Role required: admin

About this task

Customer Actions provides you step by step instructions for how to implement the changes on your instance to increase your security posture. Customer Actions pulls relevant content from the product documentation to guide you through each step of the implementation process. Follow the steps below to learn how to implement a Customer Action.

Procedure

1. Navigate to the **Customer Actions** app.
2. Select a Customer Action to implement.
   Under the **Updates** section, the Available tab lists all the Customer Actions on your instance. Click on a Customer Action under this tab to view the instructions.
3. Read the overview to learn why the Customer Action is needed, and to prepare for implementing the steps.
   Click on **Go to next step** to progress to the implementation details.
4. Implement the update steps.

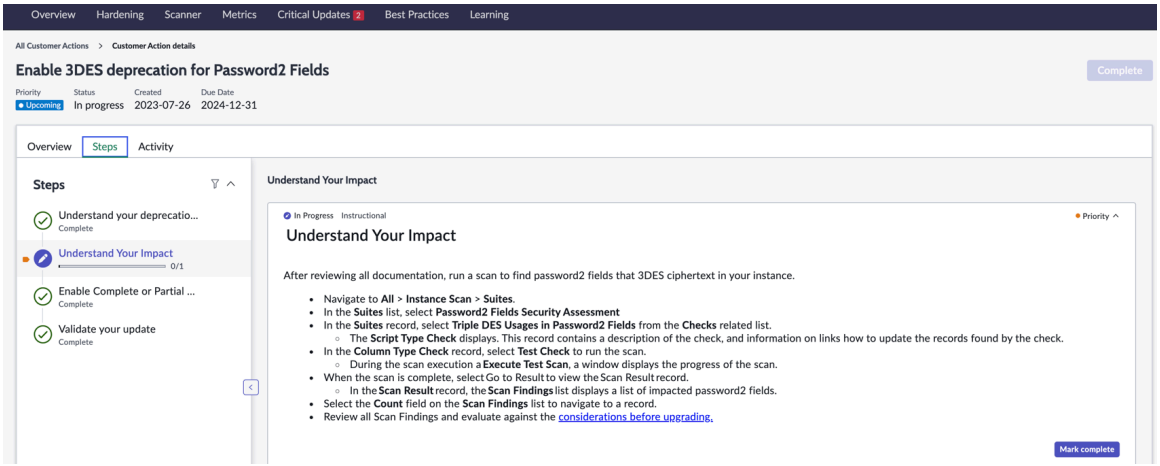

a. Read each group of instructions and implement them within your instance.

b. Click **Mark complete** when you are done with a set of instructions.

c. Repeat steps a-b until you have completed all of the update steps.

d. Click **Complete**.
The Customer Action should appear in the **Complete** tab.

View Customer Actions

See details of all of the activity related to your Customer Actions.

Before you begin

Role required: admin

About this task

The activities will be listed from newest to oldest so that you can analyze the most recent activity first. Each update to an activity is timestamped, and you can use search and filter to query for specific information.

Procedure

1. Navigate to the **Customer Actions** app within Security Center.

2. Navigate to the **Updates** label of Customer Actions.

3. Click a state, and then select a Customer Action from that state.

In the example below, the **Complete** state is selected and then the Customer Action for **End of Support: GlideEncrypter API** is selected.

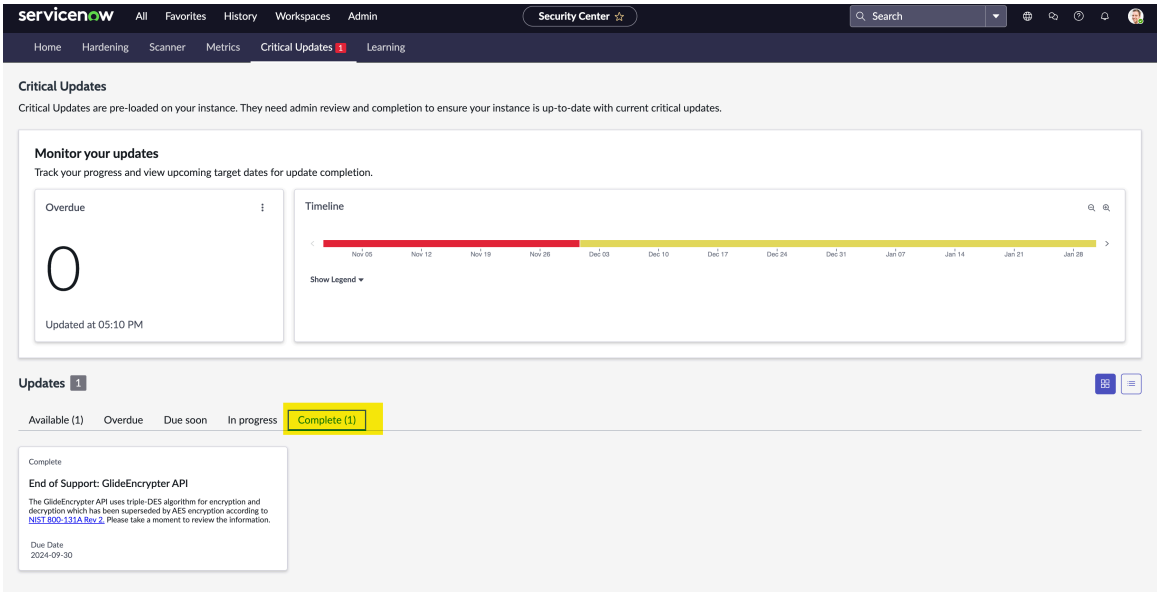

4. View the activity of your Customer Action.

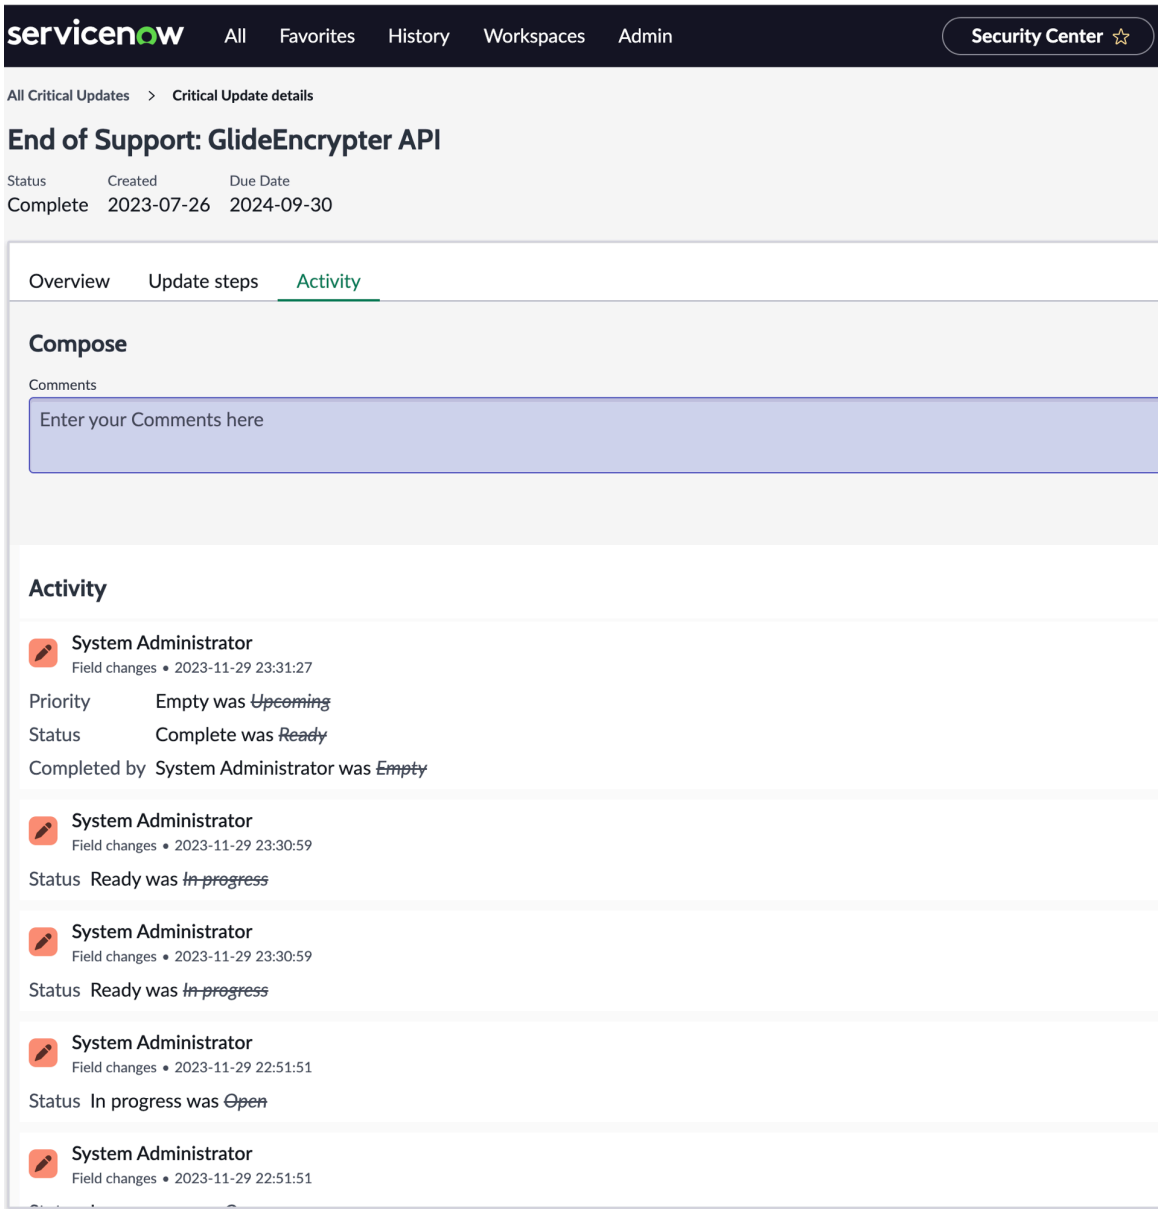

All of the activity related to a Customer Action is automatically recorded. In addition, you have the ability to add additional comments to activities.

Security monitoring console

Supervise security notifications and metrics to stay informed about potential security risks on your instance.

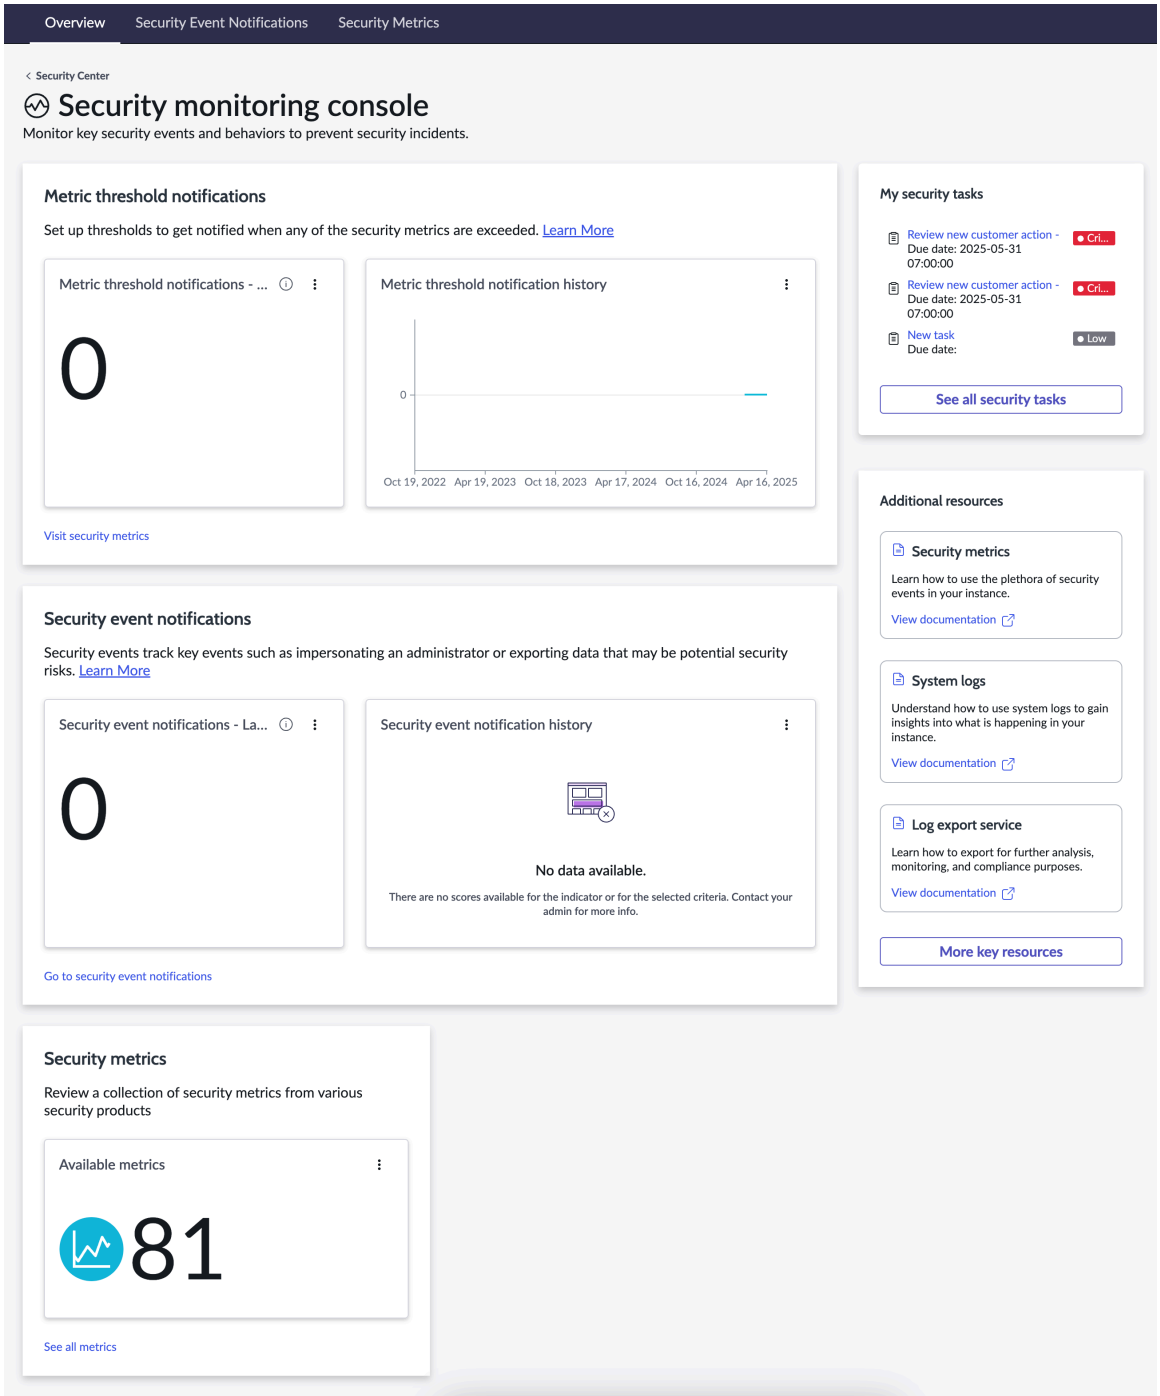

The Security monitoring console is organized into sections, each with cards containing information on a specific aspect of security monitoring. Select the info (ⓘ) icon on any card for details on what each represents.

Navigation bar

Use the bar at the top of the page to navigate between the **Overview**, **Security Event Notifications**, and **Security Metrics** sections.

Metric threshold notifications

Learn about the metric threshold notifications on your instance in the Metric threshold notifications panel. Thresholds define a normal range of scores for an indicator and alert you

when certain events occur, like when a score reaches an all-time high. This panel displays the count of threshold notifications in the last seven days, and a chart showing the threshold count over time.

Select either card to open the Security metrics tab to see more detail on these notifications.

Security event notifications

Track security events such as impersonation or data exports in the Security event notifications section. This panel displays the count of security events in the last seven days, and a chart showing the count over time.

Select either card to open the Security metrics tab to see more detail on these notifications.

Security metrics

This section displays a count of available Security Metrics on your instance. Review a collection of Security Metrics from various security products.

My Security Tasks

View the most urgent Security Tasks assigned to you. Select a task to view its details, or select **See all Security Tasks** to view a complete list of Security Tasks. For details on Security Tasks, see Security Tasks.

Additional resources

Use the additional resources section on the edge of the screen to navigate to documentation and videos relevant to learning Security Center as well as information on managing instance security.

Security Event Notifications

View, manage, and analyze the default security event notification policies on your ServiceNow instance, as well as access the functionality to create custom policies.

Access Security Event Notifications from the Security Center homepage by opening the **Security Monitoring Console** in the **Security Consoles** section, then selecting the **Security Event Notifications** tab.

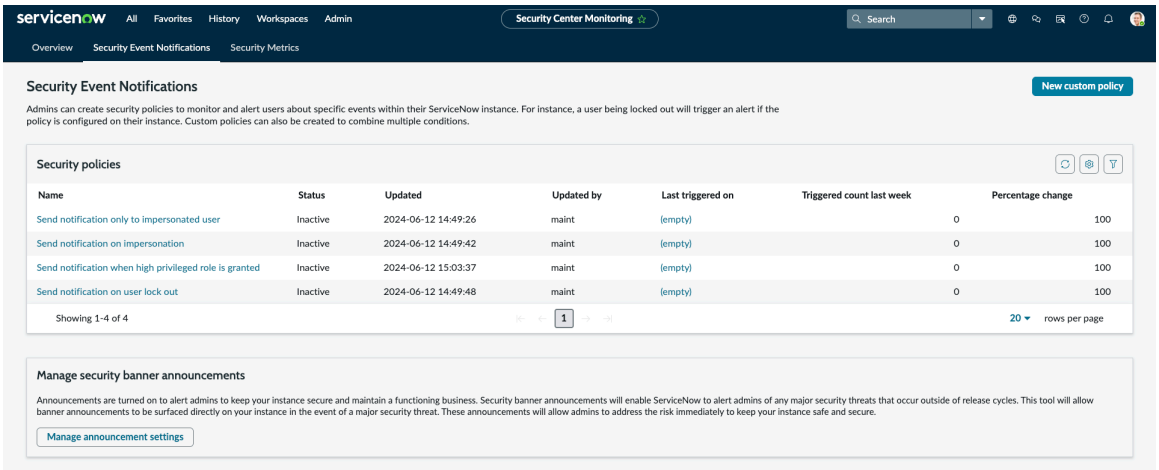

Customize and create notification policies that trigger email notifications when users perform actions in the ServiceNow AI Platform that might be insecure or require admin monitoring.

Admins can customize the values for the default policies, clone, and edit default policies or create their own custom ones.

Security policies

The Security Event Notifications console displays important information regarding the security policies on your instance. This table includes columns displaying:

Label	Description
Name	The name of the security event notification policy.
Status	The status of a policy can be either active or inactive.
Updated	The date and time of when the policy was last updated.
Updated by	The user that updated a policy.
Last triggered on	The date and time the policy last ran.
Triggered count last week	The number of triggers a policy generated the past 7 days.
Percentage change	The rate of change for weekly triggers, calculated as a percentage. **Note:** For example, if this policy triggers an average of 5 notifications a week and it triggered 10 times last week, that would be a percentage increase of 100%.

Select a policy on the table to view more detail, including an overview, settings, and notification history. Select the **New custom policy** button to create a custom policy tailored to your organization's use cases.

Manage security banner announcements

Select **Manage announcement settings** to control whether admins receive banner notifications on this instance. For information on these announcements, see Security banner announcements.

Create custom security event notification policies

Learn how to create custom security event notifications that are specific to your organization's needs. This enables you to monitor actions taken by users and groups on your instance and generate notifications for potential security risks.

Before you begin

Role required: admin

Procedure

1. Navigate to the Security Center app and select **Notifications**.
2. Select **New custom policy**.
3. Select one of the three event types that will trigger the start of the policy, and select **Create**.
4. Configure the policy.
   See Configure preferences for security event notification policies to learn how to configure your new custom security event notification policy.

Modify security event notification policies

Learn how to modify the settings of your security event notification policies.

Before you begin
Role required: admin

Procedure

1. Navigate to the Security Event **Notification** homepage and select any of the policies you want to modify.

   Note: You cannot modify the trigger or condition logic for default policies. To modify these settings, you must first duplicate the policy which can be done by selecting the down arrow next to **Update**, and then selecting **Duplicate**.

2. Select **policy settings** to change the settings.
3. Select **Update** to save your changes.
4. **Optional:** Select **deactivate** until you are ready to start triggering notifications.
5. **Optional:** Start triggering notifications by selecting **Activate**.
   See Configure preferences for security event notification policies to learn how to configure your new custom security event notification policy.

Configure preferences for security event notification policies

Discover how to customize security event notification policies in Security Center to align with your organization's specific needs.

Before you begin
Role required: admin

About this task
The **Policy settings** page is where you can customize the settings for a security event notification policy. Here, you can adjust options related to when the policy runs, the conditional logic, and the notifications that are sent.

Procedure

1. In Security Center, select **Notifications**.
2. Under the Security policies table, select a security event notification policy that you want to configure.
   For example, **Send notification when high privileged role is granted**.
3. Select **Policy settings**.
4. Configure the policy settings.

Policy settings configuration options

Label	What to configure
When to run	**Trigger**: The event that starts the policy.
Conditions	The conditional logic and conditions to apply to your policies.

Label	What to configure
	○ **Condition logic**: two options ▪ **All conditions must be met to trigger this policy**: all conditions need to be valid for the policy to send notifications. ▪ **Any of the conditions must be met to trigger this policy**: notifications will be sent if any one of the conditions is valid. ○ **Add Condition**: Select Add Condition and provide the appropriate values for it. You can remove a condition by selecting **Remove condition**.
Notifications	The email body and email recipients: ○ **Notification**: Select one of the predefined notification templates. See Create a custom email for security event notifications for more information. ○ **Add Notification**: Select Add Notification. ○ **Groups**: Select the group who should receive the notification. ○ **Users**: Select the users who should receive the notification. You can remove a notification by selecting **Remove notification**.

Note: If you want to make a policy inactive, select **Deactivate**. If you want to replicate a policy, select the down arrow next to the **Update** () button.

5. Select **Update** to save your settings.

Create a custom email for security event notifications

Learn how for creating a custom email for security event notifications by configuring new notifications, setting triggers, defining recipients, and crafting email content with dynamic event fields.

Before you begin
Role required: admin

Procedure

1. Select **system notifications** > **notification**.
2. Select **new**.
3. Enter a **name** for the notification field.
4. Select the field next to **table** and enter security policy notification (*sn_vsc_security_policy*).

5. Configure the new record.
   There are three tabs within the record that must be configured: **When to send**, **Who will receive**, and **What it will contain**.

Tab	Description
When to send	Select the field next to **Send when** label and select **Triggered**.
Who will receive	○ Select the local icon next to **Users/Groups in fields**, and select both **Users** and **Groups**. Select the **Add Item** (>) icon to move the fields into the **Selected** text field. ○ Right click on the record header banner and select **Save**.)
What it will contain	Create an email template. See Create an email template for more information. Enter fields of the event into **Message HTML**. This can be done by using the `${event_id.FIELD_NAME}`. Here's an example: `Role: ${event_id.role.URI_REF} Granted to: ${event_id.granted_to.URI_REF} Granted by: ${event_id.user.URI_REF} Logged at: ${event_id.sys_created_on} Security Center Notification: ${execution.policy.name}`

6. Select **Submit**.
7. Add a custom notification to email template in security event notifications.
   a. In Security Center, select **Notifications** > **Policy settings**.
   b. Navigate to the **Notifications** label.
   c. Enter the name of the custom notification you created in the **Notification** field.

Security Event Notifications history

Explore the complete history of security event notifications on your instance.

The **Notifications** page provides an overview of the notification history for all security event notification policies on your instance. You can view the data in charts, presenting it graphically and applying performance analytics functionalities like targets, thresholds, and KPI signals. Alternatively, you can view the data in a table format, enabling you to utilize standard functionalities such as search, sort, filter, and querying.

Security metrics

Monitor over 50 different Security Metrics to identify potential security threats or insecure behaviors. Set thresholds for email notifications, visualize, and analyze the data in multiple ways. Export the data, or create dashboards with the metrics that are most important to your organization.

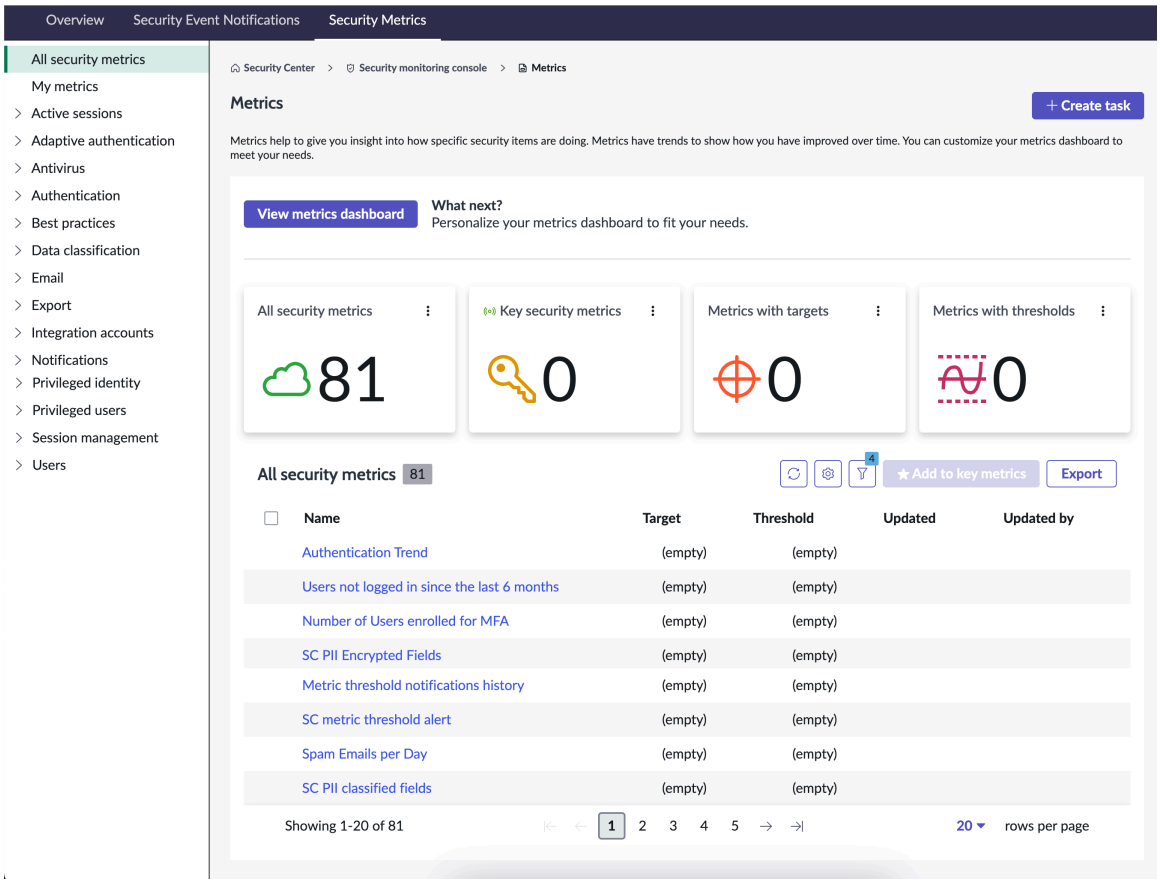

Access Security Metrics within Security Center by selecting **Security Monitoring** in the tools section. Then, select the **Security metrics** card.

All Security Metrics dashboard

View metrics such as failed logins, exports, impersonations, and more using this dashboard. The cards above the list provide a count of the metrics that match specific criteria listed on the card. Select any of these cards to filter the list to show only those metrics that match the criteria.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

My metrics dashboard

This is a customizable dashboard that displays information about metrics you care about most. Select **Edit** to customize your dashboard by adding visualizations, filters, headings, images, rich text, and lists. To learn about the various ways you can customize your dashboard, see Dashboards in Platform Analytics.

Metrics navigation pane

The navigation pane on the edge of the screen displays your instance metrics. Select **All Security Metrics** to view all metrics in a single list. Select any of the items under **All Security Metrics** the entry to browse the metrics organized by category.

Each category displays a dashboard similar to the security **My metrics dashboard**, with cards for the metrics in that category.

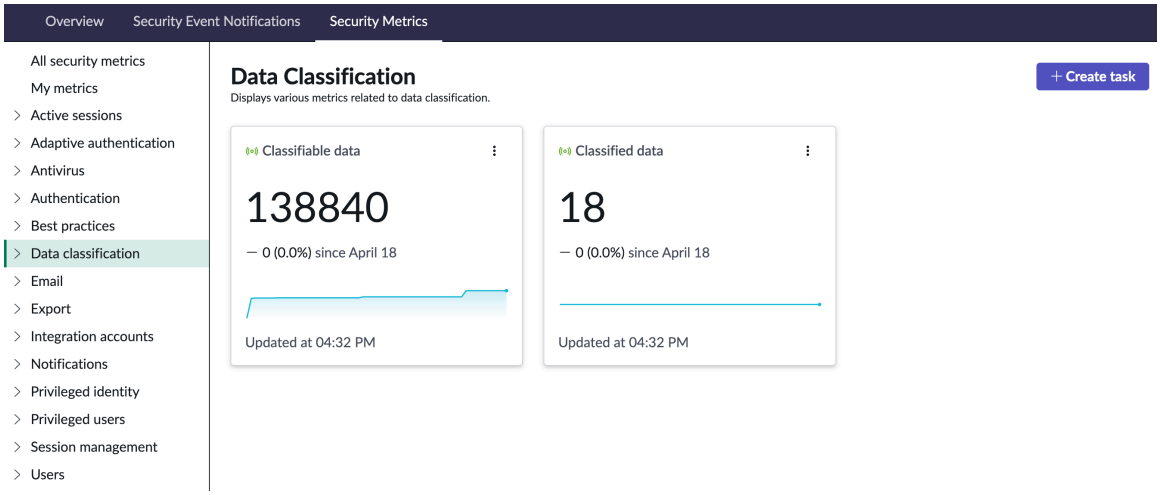

Customize the My security metrics dashboard

Discover the flexibility of the My security metrics dashboard, which can be customized with metrics from various sources like graphs and charts. Tailor the dashboard to suit your organization's specific requirements.

Before you begin

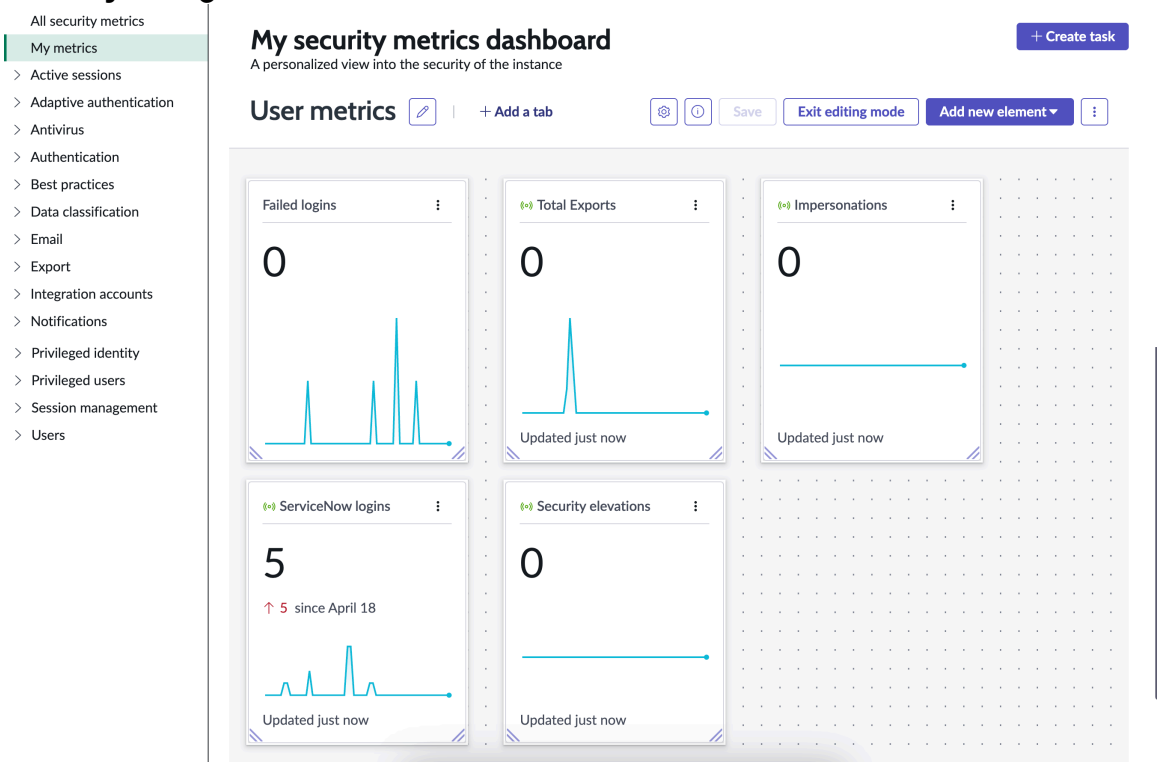

To customize the Security Metrics dashboard, confirm you are within the Security Center application scope. To switch to this scope, select **application picker** > **Application scope**, and type security center.

Role required: admin

About this task

Implement and customize My security metrics dashboard. Add metrics and content to your dashboard from a variety of sources, such as graphs and charts generated from a report. Learn

the steps required to add a tab to your dashboard so that you can customize it to meet your organization's needs.

Procedure

1. Access Security Metrics within Security Center by selecting **Security Monitoring** in the tools section, and selecting the **Security metrics** card.
2. Enter edit mode by selecting **Edit**.
3. Create a tab by selecting **+Add a tab**.

   Note: You can also add a new tab by selecting the **plus (+)** icon next to an existing tab.

4. Edit the title of your newly created tab by selecting the pencil icon and entering a name.
5. Add an element to your tab by selecting **Add a new element**.
6. Select **Save**.
7. **Optional:** To share the dashboard with users, groups, and roles in your organization, select **More actions menu** > **Share**.
   a. Enter the name of a user to grant access to.
   b. Select **Confirm**.

Configure Security Metrics to send email when thresholds are triggered

Learn how to configure Security Metrics so that your instance generates an email notification when a threshold is triggered.

Before you begin

Role required: admin

Procedure

1. In the Security Center, navigate to **metrics**.
2. Select a metric that you want to set a threshold for in your organization.

   Example
   To manage failed logins, you can set a threshold. When the number of failed logins reach this threshold, admins or security experts in your organization are notified. The appropriate Security Metric to target in this case is **users** > **failed logins**.
3. Select the **Thresholds** icon.
4. Select the **plus (+)** icon on the thresholds panel to configure the form.
   The options available are to set the visibility, threshold type, and threshold value.

   a. For visibility, select **for everyone** if you want everyone to receive the notification, or **for me** if you only want to receive a notification.

   b. In **threshold type**, select **all time high**.
   This field triggers the threshold notification.

   Note: A threshold value must only be entered when the **Threshold Type** is set to **Less than** or **More than**.

5. Select **save** to create the threshold.
   Next, the notification must be set up.

6. Select **all**, and enter **system notification**.
7. Navigate to **email** > **notifications**.
8. In the **name** column, enter `*pa thre`.
   The threshold notifications are event-based, and they're set up in the PA Thresholds Notification table.
9. Select **PA Thresholds Notification**.
   There are three tabs in the results:
   - When to send: Specifies what must occur to send the email notification. In this example, once the threshold you previously set is reached, the email is triggered.
   - Who will receive: Specifies which users receive the notification. You can add users and groups.
   - What it will contain: Specifies the script action that a user can modify to customize the message that users and the group will receive. To learn more about creating script actions, see Script actions.

All Security Metrics

Navigate to **All Security Metrics** to view a table with the data related to the Security Metrics of your instance.

Access the list of all Security Metrics on a filterable list by accessing the Security Metrics Dashboard, and selecting the **All Security Metrics** option on from the list on the left.

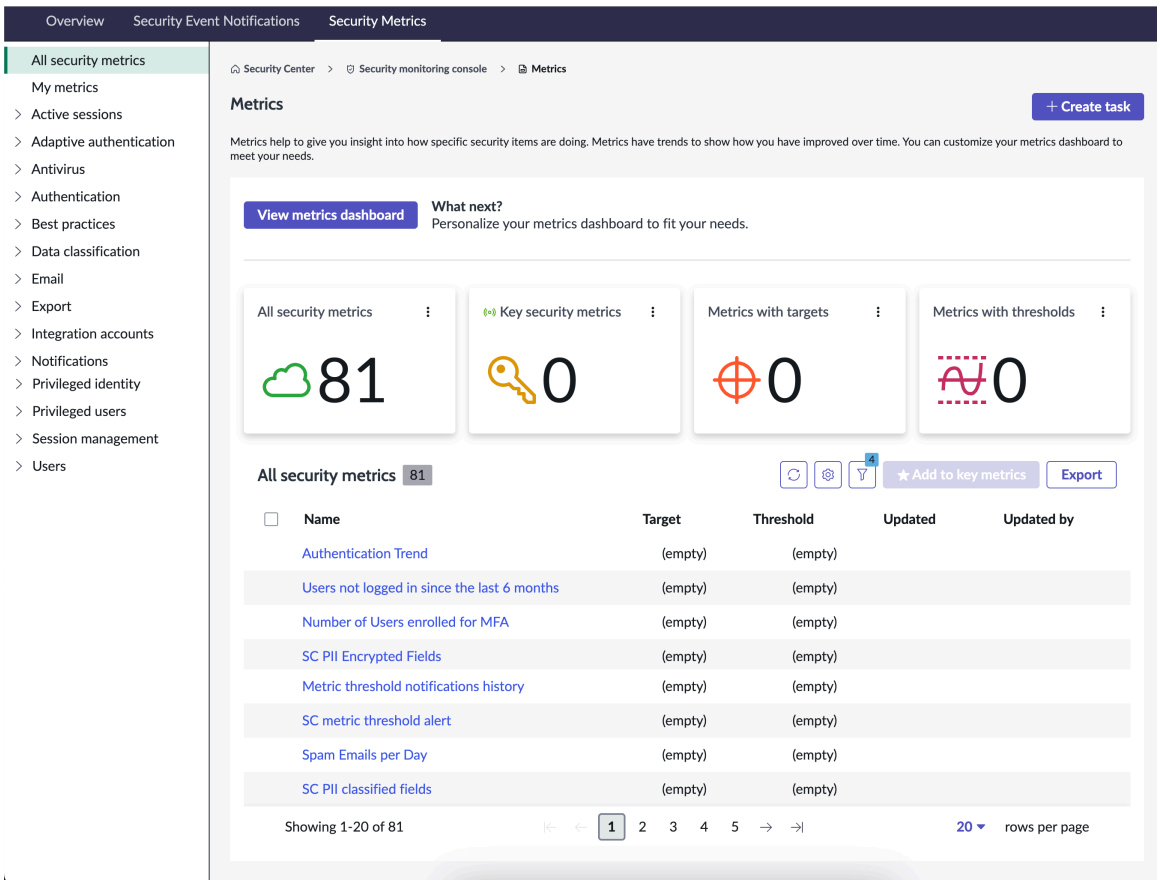

Select any item on the list to view details on each metric.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Active Sessions

View the trend line of the active users on the ServiceNow AI Platform.

This page displays cards with information on metrics related to active sessions. Each card displays a trend line for the following metrics:

- User sessions: Metrics related to active user sessions on your instance.
- Privileged user sessions: Metrics related to privileged users or those who have been assigned additional roles.

Select the cards to view the individual metrics page with additional details.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Adaptive authentication Security Metrics

Use authentication policies to evaluate authentication requests and deny or allow access to your instance based on the specified policy conditions.

Adaptive authentication metrics enable you to monitor how adaptive authentication is being used on your instance. View a summary of all of your metrics on Adaptive authentication, individual metrics such as policy result rates or denied IP addresses. This page requires the Adaptive Authentication (*com.snc.adaptive_authentication*) plugin for adaptive authentication to be available in your instance. You must enable Authentication Policy to see the metrics. See Activate adaptive authentication and Configure adaptive authentication properties for more details.

Note: This feature was released with version 1.2.

- Policy result rates: All the successful and failed adaptive authentication events.
- Event failure distribution: All the failed events for each event type.
- Event success distribution: The successful events associated with every event type.
- Denied IP addresses: The number of IP addresses blocked by the instance, along with their associated data.
- Authenticated user logins: The number of events counted for each event type, excluding the pre-login event.
- API user logins: The number of events associated with API Authentication policies for each event type.
- Authentication trend: The total number of events recorded.
- Authenticated users: The number of users counted for each event type, excluding the pre-login event.

Select the cards to view the individual metrics page with additional details.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Antivirus

Displays the trend of when events occur on potentially infected files. See when they are discovered, placed into quarantine, restored, or deleted.

This page displays cards with information on metrics related to antivirus activity on your instance. Each card displays a trend line for the following metrics:

- Quarantined files: The number of files that potentially contain malware.
- Downloaded files: The number of downloaded files.
- Restored files: The number of files that have been restored.
- Deleted files: The number of deleted files.

Select the cards to view the individual metrics page with additional details.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Authentication

View trends for metrics related to authentication schemes, such as Multi-Factor Authentication(MFA) use, web service accounts, and biometric scanner usage.

This page displays cards with information on metrics related to authentication. Each card displays a trend line for the following metrics:

- Users enrolled for MFA
- Users using MFA
- High privileged non-MFA users
- Active MFA users
- Locked out MFA users
- Web service account only
- X509 certificates expiring in 30 days
- Biometric scanner/hardware key users
- REST API's Without Access Policy
- Integration Accounts without Web Service Access Only Flag enabled

Select the cards to view the individual metrics page with additional details.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Data Classification

Access the Security Metrics for data classification on your ServiceNow instance.

Data classification enables you to apply data classification labels to existing dictionary entries in any table. These labels help to provide visibility into the types of data hosted in your instances. Use these classifications to comply with privacy laws, and meet industry regulation requirements.

See data classification for more information.

- **Classifiable data**: Tables or columns that can be classified.
- **Classified data**: Dictionary entries, tables, or columns that are classified.

The data is graphically displayed in a customizable chart that provides detailed analysis such as setting targets, thresholds, trends, statistics, or forecasts. The data is also organized as records in a table. See Analytics Hub for more information.

Select the cards to view the individual metrics page with additional details.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Authentication metrics

View metrics related to authentication on your instance from a dashboard.

This page displays cards with information on metrics related to authentication. Each card displays a trend line for the following metrics:

- Users enrolled for MFA: Total count of users enrolled to use MFA.
- Users using MFA bypass: Total count of users who are circumventing multi factor authentication.
- High privileged non-MFA user: Total count of high-privileged users that are not using MFA.
- Active MFA users: Total count of MFA users that are active on your instance.
- Locked out MFA users: Total count of MFA users who are locked out on your instance.
- Web service account user: Total count of users with web service account only.
- X509 certificates expiring: Total count of X509 certificates that are expiring in the next 30 days.
- Biometric scanner/Hardware key users: Total count of users login with biometric scanner or Hardware key.
- REST API's Without Access Policy: Total count of users login without REST API's Access Policy.
- Integration Accounts without Web Service Access Only Flag enabled: Total count of users login without Integration Accounts Web Service Access Only Flag enabled.

Select the cards to view the individual metrics page with additional details.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Email

Displays data related to spam emails that are being received externally.

The page displays a card displays a trend line for the spam emails, representing the number of junk mail received per day.

Select the card to view additional details.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Export

Discover the data that is commonly exported, and which users do the exporting.

This page displays cards with information on metrics related to export. Each card displays a trend line for the following metrics:

- Total exports: Total count of table records exported by user.
- Classified exports: Total count of table records exported, summarized by the data classification assigned to them.

Select the cards to view the individual metrics page with additional details.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Integration Accounts

View the trends about the integration accounts that are created on the ServiceNow AI Platform.

Security Center identifies integration accounts as user accounts with the **Internal integration user** field selected on the User [sys_user] record. These accounts are intended to manage integrations between your instance and external applications or systems. To view all integration accounts on your instance, navigate to **All** > **User Administration** > **Users** and filter by **Internal integration user = true**.

This page displays cards with information on metrics related to integration accounts that are created on the ServiceNow AI Platform. Each card displays a trend line for the following metrics:

- Total integration accounts: Trends related to integration accounts on your instance.
- Active integration accounts: Trends related to active integration accounts on your instance.
- Inactive integration accounts: Trends for the total count of integration accounts that are inactive on your instance.

Select the cards to view the individual metrics page with additional details.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Privileged Identity

Analyze data related to metrics for users with privileged identity.

This page displays cards with information on metrics related to privileged identity users and their activity on your instance. Each card displays a trend line for the following metrics:

- Admin logins: Total count of logins with users of the admin role.
- Impersonation: Total count of impersonations performed by users with the impersonator role.
- Elevation: Total count of security elevations performed by users with the security_admin role.
- ServiceNow logins: Total count of logins by ServiceNow employees.
- Admin users added: Total number of users given the admin role.

Select the cards to view the individual metrics page with additional details.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Privileged Users

View the trend line for the privileged users (active and inactive) and their activity on the ServiceNow AI Platform.

The Privileged users overview section displays the trend line for the privileged users (active and inactive) and their activity on the ServiceNow AI Platform. Privileged users are users who have been assigned additional roles by admins to access the features like High Security Settings, Import, and Portal users.

Below is an explanation of the privileged users:

- Total users: Total count of users on your instance.
- Active users: Total number of users who initiated sessions on your instance.
- Inactive users: Users who have not recently logged into your instance.

- Inactive users who are not locked out: Users who have not recently logged in but still have access to their account.
- Locked out users: Users who are not permitted to authenticate into their account.
- New users: Users who have recently been added to your instance.
- Successful logins: Users that have successfully logged in.
- Failed logins: Unsuccessful login attempts.
- Local logins not protected by MFA: Users that logged in without MFA.
- Users never logged in: Users that have never logged into your instance.
- Users not logged in since last month: Users that have not logged in the last 30 days.
- Users not logged in since the last 6 months: Users that have not logged in the last 6 months.
- Users not logged in since the last 1 year: Users that have not logged in the past year.
- Need to reset password: Users that need to reset their password.
- Password reset failures of the users: Number of password failures per user.

Select the cards to view the individual metrics page with additional details.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Session management

View metrics related to user sessions and the frequency of lockouts of the sessions.

The page displays a card displays a trend line for active sessions, representing the number of active sessions on your instance.

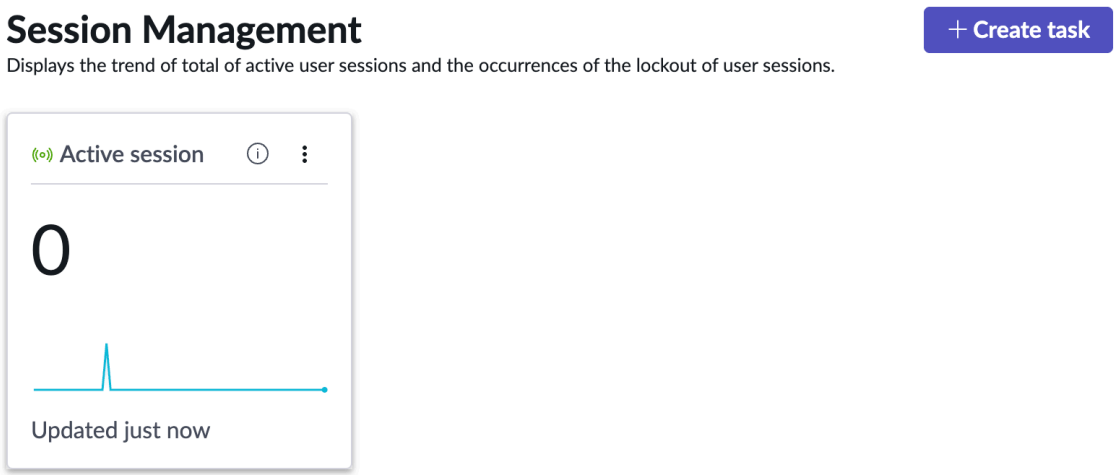

Select the card to view additional details.

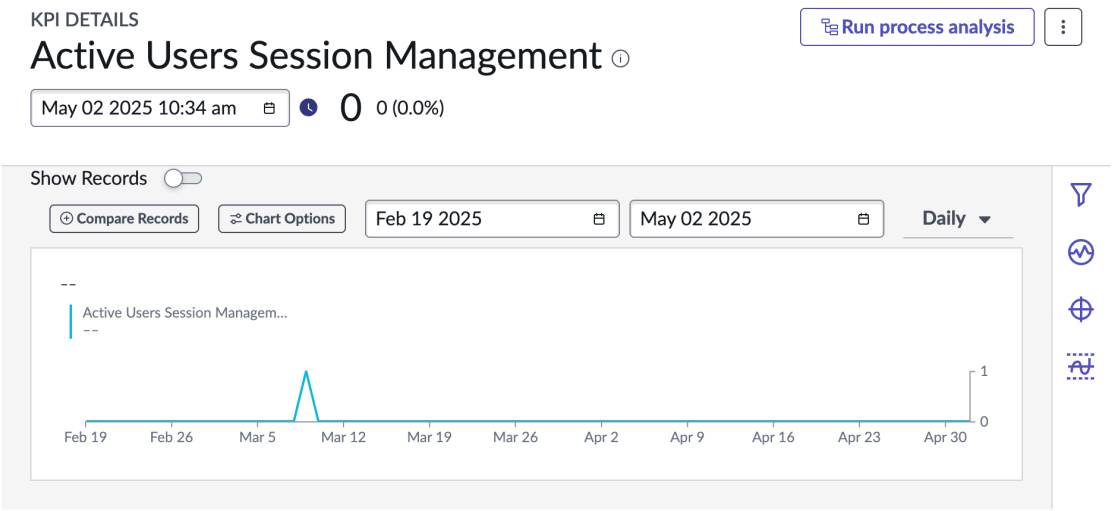

From the details page, you can use the buttons on the edge of the screen to configure filters, KPI signals, targets, and thresholds.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Users

View the trend line for total users (active and inactive) and their activity on the ServiceNow AI Platform.

The Users overview section displays the trend line for total users (active and inactive) and their activity on the ServiceNow AI Platform.

Below is an explanation of the type of users:

- Total users: Total count of users on your instance.
- Active users: Total number of users who initiated sessions on your instance.
- Inactive users: Users who have not recently logged into your instance.
- Inactive users who are not locked out: Users who have not recently logged in but still have access to their account.
- New users: Users who have recently been added to your instance.
- Successful logins: Users that have successfully logged in.
- Failed logins: Unsuccessful login attempts.
- External logins: Authentications that uses a 3rd party.
- Local logins not protected by MFA: Logins that do not use MFA.
- Users not logged in since last month: Users that have not logged in the last 30 days.
- Users not logged in since the last 6 months: Users that have not logged in the last 6 months.
- Users not logged in since the last 1 year: Users that have not logged in the past year.
- Need to reset password: Users that need to reset their password.
- Password reset failures of the users: Number of password failures per user.

Select the card to view additional details.

Select the **+Create task** button to create a Security Task related to a metric. For details on Security Tasks, see Security Tasks.

Security posture console

Improve your ability to identify, respond to, and recover from security threats with comprehensive visibility and step-by-step instructions.

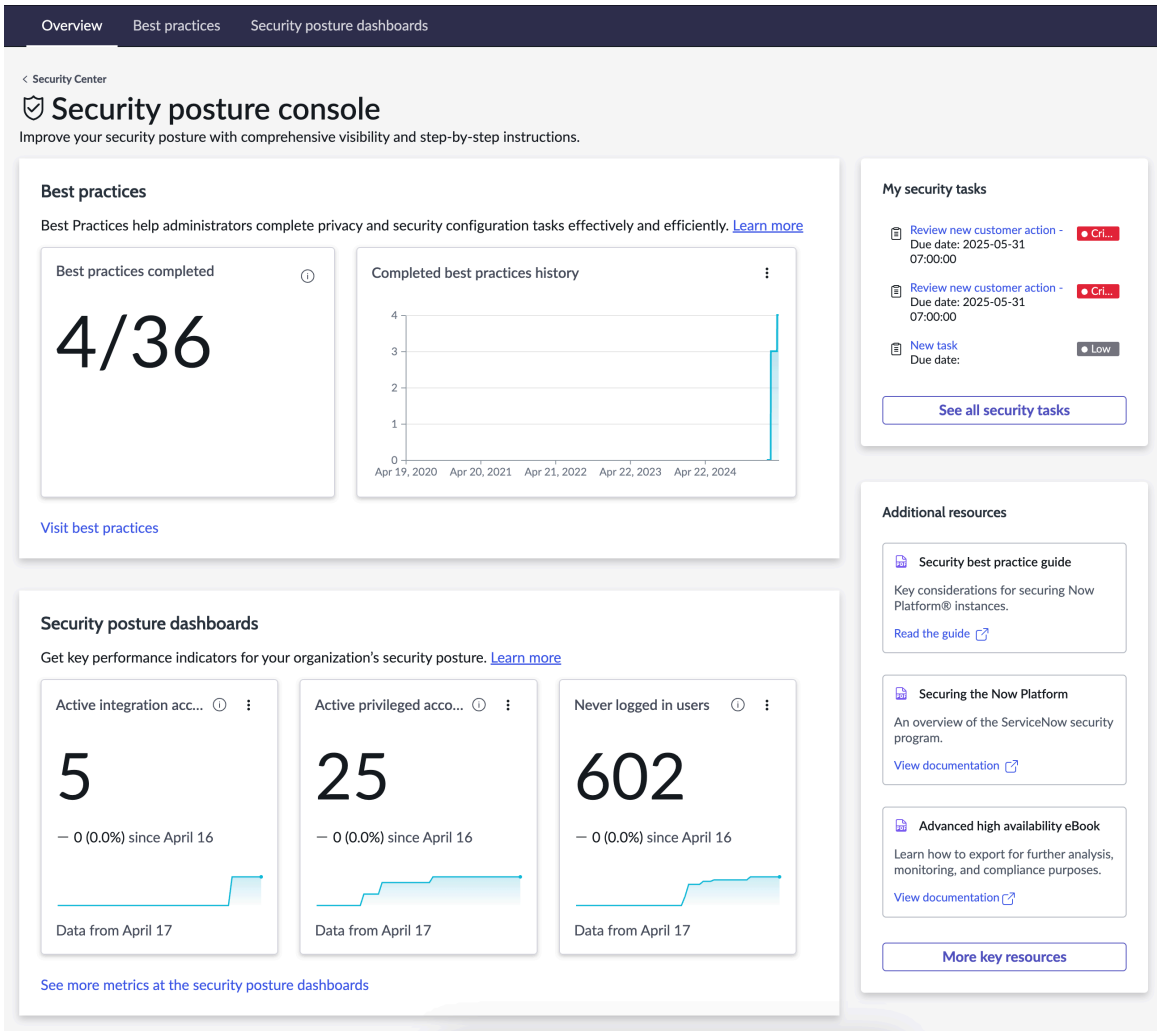

The security posture console is divided into sections to provide information about best practices and security posture dashboards. Find additional details by selecting any of the cards in these sections. You can also use the bar at the top of the page to navigate between this page and the best practices and security posture dashboards section of Security Center.

Best Practices

This section displays the count of best practices, and how many you have applied to your instance. You can see a chart showing your completed best practices over time.

Security posture dashboards

This section displays the number of active integration accounts, active privileged accounts, and users that have never logged in.

Active integration accounts

Integration accounts are user accounts that manage integrations between your instance and third party applications.

Active privileged accounts

Privileged users are users who have been assigned additional roles by admins to access the features like High Security Settings, Import, and Portal users.

Never logged-in users

User accounts that have never logged in to your instance.

My Security Tasks

View the most urgent Security Tasks assigned to you. Select a task to view its details, or select **See all Security Tasks** to view a complete list of Security Tasks.

Additional resources

Use the additional resources section on the edge of the screen to navigate to documentation and videos relevant to learning Security Center as well as information on managing instance security.

Security Best Practices

Use Security Best Practices to implement privacy and security configuration tasks on your ServiceNow instance.

Identify best practices to improve your security posture, and follow step-by-step instructions on how to implement them. Security Best Practices provide the following:

- The home page shows an overview of your progress on implementing security best practices. You can also organize and manage lists of security best practices according to your organization's goals.
- The overview page provides details of each security best practice, the steps to implement them, and a record of all activities and comments.
- The task steps page provides you with instructions on how to implement security best practices.
- The activity page tracks the history of the user and system actions related to your security best practices.

Security Best Practices home page

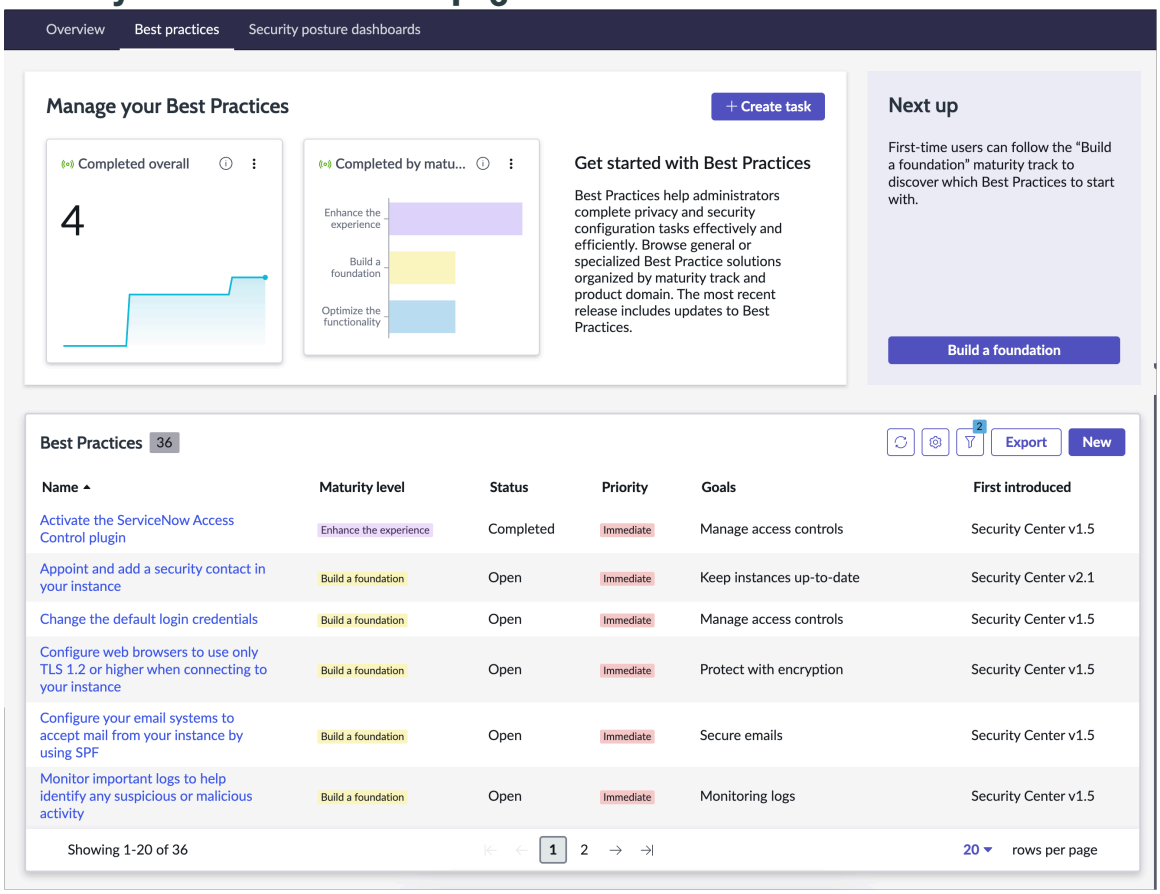

The home page displays a **Manage your best practices** section, which includes graphs provide an overview of your progress.

Completed overall

Displays a count and trend line of best practices you have completed. Select the card to view the **Completed Overall** metric page in Security metrics.

Completed by maturity level

Displays a chart of completed best practices organized by maturity level (see a description of maturity levels in the proceeding table). Select the card to view the **Completed by Maturity Level** metric page in Security metrics.

Build a foundation

Select the **Build a foundation** button to filter the table on this page to display only best practices in the **Build a foundation** maturity level. These are lower impact changes you can make to start improving instance security.

Create a task

Use the **+Create task** button to create a Security Task to track or delegate best practice work. For details on Security Tasks, see Security Tasks.

The table enables you to apply filters so that you can sort and save filtered lists, which you can use as work lists for different use cases or roles. See save a filtered list for more information.

The following are explanations of the fields related to the security best practices table.

Summary of fields used in the security best practices table

Name	Description
Name	Word used to identify a security best practice.
Maturity level	Applications and features that have been arranged by the order of impact to provide you measurable results. The values for maturity levels are: • Build a foundation • Enhance the experience • Optimize the functionality • Add advanced features These can also be thought of as crawl, walk, run, and fly phases.
Status	Current state of a best practice: • Open • In progress • Completed
Priority	Order of importance for implementing a best practice in your organization: • Immediate • Later • Not applicable
Goals	Security category that a best practice addresses: • Address initial security configurations • Secure emails • Monitoring logs • Manage access controls • Protect with encryption • Keep instances up to date
First introduced	Which Security Center version the best practice was introduced.
Changed	Which Security Center version the best practice was changed.
Removed	Which Security Center version the best practice was removed.

Security Best Practices details page

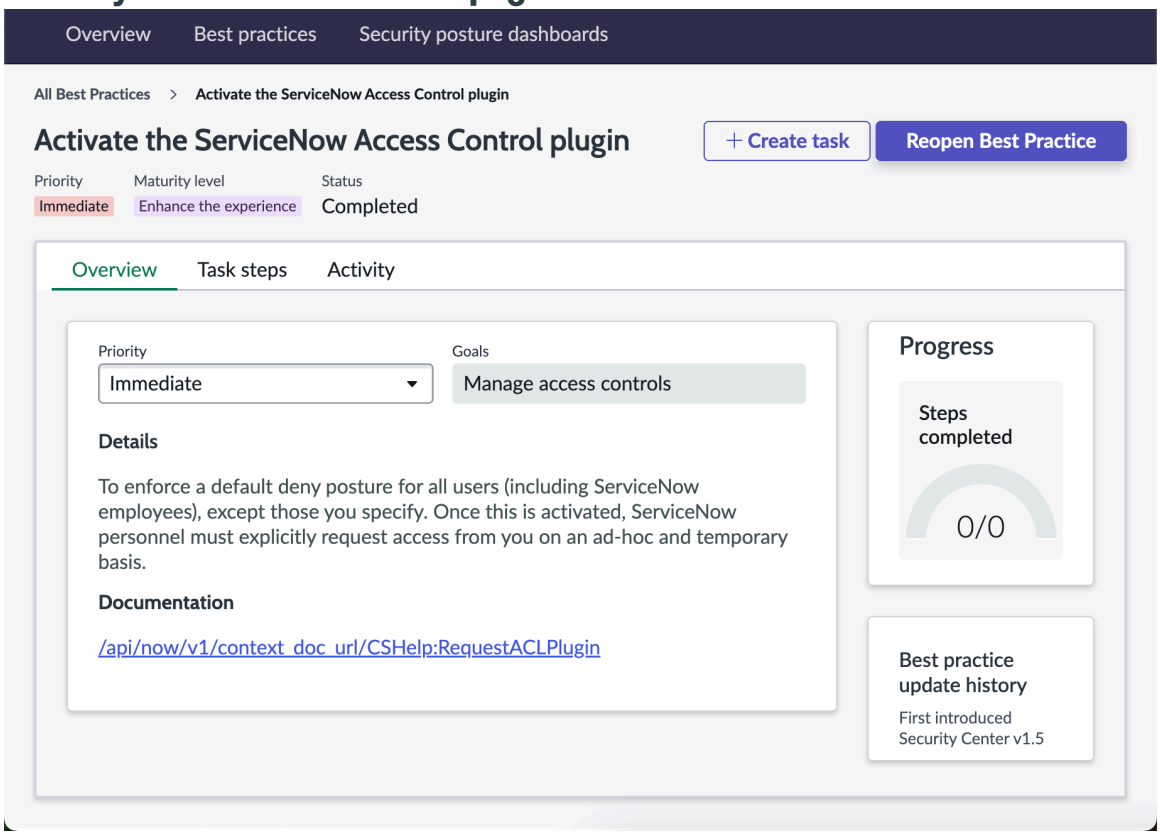

Select a best practice from the table to view its page. At the top of the details page, you can view general information about the security best practice including priority, maturity level, and status. Use **Complete Best Practice** button to mark a the practice as complete, or **Reopen Best Practice** button to mark the practice complete. Use the **+Create Task** to create a Security Task to track and delegate this task. For details on Security Tasks, see Security Tasks.

This page provides more information on the best practice, divided into tabs:

Overview

This tab contains the **Priority** drop-down menu, which enables you to specify the security best practices that are important to you at this time and which are not applicable.

The details section provides content about the features associated with the security best practice, and the documentation section provides one or more links where you can find additional information.

The **Progress** card on the right shows the number of steps completed versus the total number of steps included. Select **go to next step** to navigate to the next incomplete step.

The **best practice update history** card provides a snapshot of the release information for the best practice. You can track which ServiceNow Security Center version the security best practice was released in, and which versions it was subsequently last updated in.

Task Steps

This tab provides step-by-step instructions for how to implement this security best practice. See complete a security best practice for more information.

Activity

This tabe displays timestamped activities listed from newest to oldest. Use search and filter to query for information. See apply filters to the security best practices table for more information.

Complete a security best practice

Learn how to complete a security best practice on your ServiceNow instance.

Before you begin

Role required: admin

About this task

Complete the following steps to implement security best practices on your ServiceNow instance.

Procedure

1. Access Security Best Practices by selecting **All** on your ServiceNow instance and then entering **best practice**.
2. Apply a filter to view security best practices to implement.
   a. Navigate to the **Best Practices** table in the UI.
   b. Select the kebab menu, or the item with three vertical dots (⋮) in the **Status** column.
   c. Select the **Open** and **In progress** options, and then select **apply**.
3. Select the name of a security best practice to implement.
4. Select **Go to next step**.
5. Select a task step from the **task steps** list.
6. Select the caret down (v) icon to view the implementation instructions for that step.
7. Read the instructions and then implement them on your ServiceNow instance.
8. Select **Mark step complete** when you're done with the task.

   Note: You can skip the task if you want to come back to it later, or if it doesn't apply to your organization's goals. In addition, you can reopen a completed task step by selecting the task step, and then selecting **Restart step**.

9. Select **Complete Best Practice**.

   You can complete a security best practice at any time. You don't have to complete all the steps and it's ok if you skip them. However, consider inserting a comment in the Activity subtab when completing a security best practice when you have not completed all the steps for an audit reference.

View activity of a best practice

Track the chronological and timestamped history related to a security best practice that you're completing on your ServiceNow instance, along with the user that initiated the activity.

Before you begin

Role required: admin

Procedure

1. Access Security Best Practices Manager by selecting **All** on your ServiceNow instance and then enter best practice.
2. On the Best Practices table, select a security best practice you want to view the activity for in the **Name** column.
3. Select the **Activity** tab.

   Examples of activity that you can view includes field changes such as if the status of a security best practice has changed from Open to Complete, or when a task step is completed or skipped. In addition, you can enter additional information related to an activity by entering in a comment in the text box.

View data of completed best practices

See visual representations of the total amount of best practices completed, or segmented by maturity level.

Before you begin
Role required: admin

About this task
Complete the actions to see charts that provide visual trends of security best practices that were completed overall, and by maturity level.

Procedure

1. Access the Security Center app and navigate to the **Best Practices** tab.
2. Under the **Manage your Best Practices** card select the **Completed overall** chart.
3. View the chart to see trends regarding security best practices configurations that were completed.
4. Select **Completed by maturity level** to see a chart, which displays trends of best practices completed by maturity level.

Filters the security best practices table

Apply filters to return more relevant results for your security best practices.

Before you begin
Role required: admin

Procedure

1. Access Security Best Practices Manager by selecting **All** on your ServiceNow instance and entering best practice.
2. Navigate to the **Best Practices** table and select a field you want to apply a filter to by selecting its kebab menu, or the one with three vertical dots (⋮).

Fields	Filter
Name	Combine a conditional operator with a string of text to return security best practices that match the name you entered.

Fields	Filter
Maturity level	Apply a filter to return security best practices that match one or more of the following maturi ty levels: ○ Build a foundation ○ Enhance the experience ○ Optimize the functionality ○ Add advanced features
Status	Apply a filter to return security best practices that match one or more of the following status es: ○ Empty ○ Open ○ In progress ○ Completed
Priority	Apply a filter to return security best practices that match the following priorities: ○ Empty ○ Immediate ○ Later ○ Not applicable
Goals	Apply a filter to return security best practices that match one or more of the following goals: ○ Address initial security configurations ○ Keep instances up to date ○ Manage access controls ○ Monitoring logs ○ Protect with encryption ○ Secure emails **Note**: Advanced filters are applied for goals by default. To modify them, select **Make changes** > **Advanced view** > **Select field goals** and then build the rest of the condi tions and select **Update**.
First introduced	Combine a conditional operator with a string of text to return best practices that were first introduced with the version you entered.

Fields	Filter
Changed	Combine a conditional operator with a string of text to return security best practices that changed in the version you entered.
Removed	Combine a conditional operator with a string of text to return security best practices that were removed in the version you entered.

Save a filter for future use

Discover the steps for saving filters in your security best practices for later reuse.

Before you begin
Role required: none

About this task
After you create filters in the security best practices table, you can save them for use the next time you authenticate into your instance. If you want more flexibility on the type of filters you can create, you can apply an advanced filter.

Procedure

1. Select the Show filter ( ) panel icon.
2. In the **Filter** dialog panel, scroll down and select **Advanced view**.
3. Build a filter by adding conditions that contain a field, operator, and values.
4. Select **Save filter**
5. Enter the Filter name.
6. Select a permission.
7. Select **Save**.
8. Select **Update** to apply the saved filter to the security best practices table.

Use a saved filter

Discover the steps for using the filters that you saved for your security best practices.

Before you begin
Role required: none

About this task
After you log out of your instance your filters are reset to the default settings. The following steps show how to retrieve a saved filter so that you can reapply it on the security best practices table.

Procedure

1. Select the Show filter panel ( ) icon.
2. In the Filter dialog box, scroll down and select **Advanced view**.
3. Select **Use existing filter**, and then select a saved filter from the dropdown list.
4. Select **Update** to apply the saved filter.

Best Practices

Learn details about the Best Practices in the Security Posture Console.

Best Practice	Description
Activate the ServiceNow Access Control plugin	Use the ServiceNow Access Control plugin to control which employees can access your instance, and when. Enforcing a default deny posture for all users except those users you specify. This can include ServiceNow employees. Using this plugin helps prevent unwanted access to your instance. **Note:** After ServiceNow Access Control is activated, personnel must explicitly request access from you on an ad-hoc and temporary basis. For details on this plugin, see ServiceNow access control.
Appoint and add a security contact in your instance	Select an information security contact within your organization who receives security-related information from the Security Team. This contact is in addition to your admin, who also receives these updates. This information could be security issues, security alerts, or details about important software updates. For details on adding a security contact, see KB0621516.
Change the default login credentials	Change the passwords on your instance's built-in user accounts, such as admin, ITIL, and employee. These accounts are provisioned with default passwords, unique to your instance, but should be changed as soon as possible. For details on how to change the passwords for user accounts on your instance, see Configure password for a user.
Configure web browsers to use only TLS 1.2 or higher when connecting to your instance	Verify that browsers that connect to your instance are using the more secure Transport Layer Security (TLS) 1.2. This change can be made on the browser or enforced by your web proxy or other gateway. Consult the documentation for your browser, web proxy, or gateway for steps on how to configure those products to use only TLS 1.2.
Configure your email systems to accept mail from your instance by using SPF	If your organization uses Sender Policy Framework (SPF) to control inbound email with anti-spam technology, you must configure it to accept email originating from your instance. Configure SPF to query the SPF records dynamically. If SPF isn't an option, another approach is to add the mail server IP addresses to your allow list. This configuration must be monitored as the addresses could be subject to change. For steps and additional details on these solutions, see KB0535456

Best Practice	Description
Consider limiting file attachments, uploads, and downloads	Restrict attachment uploads by role, file extension, MIME type, or size to help prevent potentially malicious files being stored and then delivered from your instance. You can also control which file types can be downloaded, including by MIME type, and prevent image access by unauthenticated users. These attachment restrictions are controlled by system properties on your instance. For details on their configuration, see Configure attachment system properties.
Disable browser SQL messages	Prevent SQL error messages from being presented in the web browser. Though useful to users and developers, these messages can be used by attackers to learn information about your system or to help guide their attempts to access your data. These messages can be turned off using a system property. For details on this system property, see Disable SQL Error Messages [Updated in Security Center 1.3 and 1.5].
Disable password-less authentication	Help ensure strong authentication by disabling password-less authentication when possible. Without disabling password-less authentication, potential attackers could gain access to your instance by correctly guessing a user name (such as firstname.lastname or a role title). You can disable password-less authentication on your instance using a system property. For details on this property see Disable password-less authentication.
Enable table auditing for important or sensitive data	Track changes to your data using table auditing. Auditing tracks the creation, update, and deletion of all records in the table where it's enabled, enabling admins to track changes to important or sensitive data. Admins can also choose to select specific fields in a table for auditing to see more targeted results or to reduce performance impact. For details on auditing on instances, see Auditing. For specific instructions on enabling auditing on a table, see Configuring auditing for a table.
Encrypt data at rest within the instance	Encrypt your data to maintain its confidentiality and integrity. Data on your instance can be within the database. You can also elect to subscribe to functionality to encrypt the data volume transparently on the backend. The physical disks on which the instance runs can also be encrypted in their entirety to guard data in case of their loss or theft. You can use different methods of encryption simultaneously for data stored in your instance, depending on your use case and the risks you wish to mitigate. For example, you

Best Practice	Description
	can choose to transparently encrypt your data at rest using database encryption on most tables, cloud encryption on the entire data volume. You could also use full disk hardware encryption, which also requires a dedicated environment to protect against drive or server theft. Review the encryption options available with in Key Management Framework.
Enforce the use of strong passphrases	Use password policies to enforce the length, complexity, expiration, uniqueness, lockout, and more for native and local accounts on your instance. Use these policies to maximize security, encourage the adoption of long passphrases and help to eliminate the use of simple passwords. You can retain your existing policies for any external authentication services you have integrated, such as LDAP or SAML. For details on password policy configuration, see Configure your password policy.
Ensure automatic account creation	Use this feature to create user accounts by email dynamically. Activate this feature only if necessary for your business needs, only after you have defined a list of trusted domains from which accounts can be created. You can also control how passwords are assigned to new accounts created this way. For details on automatic user creation, see Enable automatic user creation.
Ensure Secure Access to knowledge bases	Manage access to knowledge bases and articles to help ensure secure and efficient information sharing. You can determine whether certain users or categories of users can access knowledge bases and knowledge articles by controlling contribute and read access. The specific configuration depends on your business needs. Learn about your options for configuring knowledge access at Managing access to knowledge bases and knowledge articles.
Ensure that the High Security plugin is installed and activated	Use the High Security plugin (HSP) to enhance security management and applying appropriate settings. High Security Settings provides a central location for security settings, creates a distinct security administrator role, a default deny property, and other important security features. HSP is installed and enabled by default on all new instances. You can request HSP activation for older instances, including instances that have had upgrades from an older version. Enabling HSP should be done only after careful testing in a

Best Practice	Description
	non-production environment, as activation changes some fundamental properties and behaviors. For more details on the High Security plugin, see Enable High Security Plugin [Updated in Security Center 1.3].
Familiarize yourself with NOW security resources	Security information constantly evolves, so it's crucial to stay updated with security resources to keep your information security strong. Use the following resources to stay informed about security resources: • CORE Directory: ServiceNow CORE Compliance Portal • Securing the ServiceNow AI Platform: How ServiceNow protects customer data • Secure your instance
Harden your Instance	Use the Security Center Hardening tool to reduce risk by limiting weaknesses that could be exploited, and implement recommended settings to secure your instance further. Learn more about Security Center at Security Center. Review the available hardening settings at Hardening settings.
Install patches as soon as possible	Install patches and platform updates as soon as possible help ensure the highest levels of security for both your instance and those of other customers. Keeping current with updates also enables you to maintain continuous support by conforming to the EOL policy. Use Upgrade Center to help manage the process. Security fixes are routinely released for the Now Platform via the patches and hot fixes that accompany product feature updates. Upgrading when new patches and hot fixes are available helps reduce the risk of potential vulnerabilities. Information about Now Platform releases, patches, and hot fixes can be found in the Release Notes section of the product documentation. For more information, see Phase 1 - .
Integrate with MFA	Integrate third-party multi-factor authentication (MFA) with your existing SAML IdP to provide additional login security. MFA provides a high level of security because authentication requires multiple authentication factors. Something the user knows (the password) as well as something they own (a one-time code, mobile phone, or biometric attributes, such as a fingerprint).

Best Practice	Description
	Learn more about MFA integration at Multi-factor authentication.
Limit accepted email sender domains	Control which domains and users your instance can communicate with via email by using system address filters. These filters can be customized to your requirements. Learn how to configure trusted domains at Designate email domains as untrusted or trusted.
Monitor important logs to help identify any suspicious or malicious activity	The system logs module provides a variety of logs that you can use to troubleshoot and debug transactions and events that take place within the instance. **Event logs** Event logs reveal much about system activity, including login events (successful or otherwise), and privilege escalation. **System logs** System logs contain extensive information about general activity, including configuration changes, system errors, workflows, and inbound / outbound data connections. **Note:** The Event and System logs can also be used to provide an audit trail of any activity by personnel. **Transaction logs** Transaction logs record all web browser-related activity for an instance and can provide details of every request made. Transaction logs can be useful for identifying unusual or malicious activity. **Table auditing and record history** Enable auditing for database tables. Record history is perpetual and enables you to track and view details of any changes made to the data since its creation. By default, only the incident, problem, and change tables are tracked. For other tables, auditing must be enabled manually. **Import logs** You can view detailed information related to data import activity into your instance by checking the import logs. These logs include information about source and status, time, and so on **Outbound web services logs** These show REST and SOAP request activity and can help you to track the volume and destination of connections to external services. Learn more about system logs in System logs.

Best Practice	Description
Monitor login failure rates and create alerts	Monitor for unusual activity such as high numbers of failed logins, especially within short time frames. You can create alerts to send emails when a threshold you define is exceeded. Learn how to configure these thresholds at Indicator thresholds.
Monitor security events	Review the **My security metrics dashboard** to see the available security metrics for your instance, and set thresholds to generate email notifications for notable activity. Examples of notable activity can include: **Privilege escalation** Unexpected modifications made to privileged roles, such as Admin, ITIL_Admin, or any other roles with higher privileges could indicate suspicious actions. **Failed logins** Unusual numbers or patterns of failed logins can reveal potential brute force attempts or password spray attacks. **Admins and high privilege users added** New admin account creation should be checked for validity to help to prevent attempts at unauthorized privileged access.
Monitor your instance's Hardening Compliance level	Ensure that your instances are in compliance with the latest security hardening metrics using the Hardening tool in Security Center. Access this tool in a non-production instance assess impact to your environment. Ideally, the score should be as close to 100% as possible with a minimum score of 83%, without affecting product functionality. Learn more about Security Center's hardening settings tool at Hardening settings.
Refer your developers to the Secure Coding Guide	Help to ensure that your instance is secure and resistant to unauthorized access as possible using secure coding practices. The ServiceNow Secure Coding guide for Instance developers provides an overview of application security-related GlideScriptable classes and methods offered by ServiceNow. This guide is designed to assist and educate developers while creating and modifying the code on the target Instance. Review the guide at ServiceNow Secure Coding guide for Instance developers.
Remove the 'Remember Me' check box	Help prevent unwanted access to your instance by deactivating the **Remember Me** feature. When active, this feature stores a cookie is on the user's computer, which automatically authenticates the user on subsequent visits.

Best Practice	Description
	This can present security issues if users access your instance from an insecure endpoint, such as a shared computer. Learn more about this feature, and how to deactivate it in Remember me.
Restrict access to your instance from unknown IP addresses	Help prevent unauthorized access to your instance by restricting access from IP addresses unrelated to your organization. Anyone trying to access the instance from an unauthorized IP address are denied. If using this approach, consider allowing only your gateway or web proxy external addresses, as well as addresses from which your users access the instance from, including remote users. You can restrict both outbound and inbound access by IP address. Learn how to restrict access to your instance by IP address in Restrict access to specific IP ranges plugin [Updated in Security Center 1.3].
Review ServiceNow's guidance on password spray attacks	Protect your instance against password spray attacks. These attacks attempt to gain access by testing a commonly used password against multiple accounts in succession. Learn more about spray attacks, and how to protect your instances from them in Password Spray Attack Mitigating Strategies.
Review the Shared Security Model	Understand your shared role as a customer in maintaining the security of your instances by reviewing the Shared Responsibility Model. The Shared Responsibility Model defines the partnership between and the customer, both with specific responsibilities. Learn more at ServiceNow Shared Responsibility Model.
Transfer log data from the instance for archival and reference	Archive your log data to retain it beyond the default 21-day log rotation period. This archival can be achieved using web services requests, the data export feature, the MID Server, or the Log Export Service from the Vault package. Use the following resources to learn more about these methods: • Web services • Exploring Log Export Service (LES)
Use encryption modules with RBAC to further enhance data access control	Learn how to use the Key Management Framework (KMF) to protect the data on your instance using Role-Based Access Control (RBAC). KMF uses cryptographic modules, which enable you to define what data on your instance is encrypted, and what method of encryption to use. Using

Best Practice	Description
	multiple modules, you can encrypt different areas of your instance with different specifications. Learn how KMF and its components are used to encrypt your data at Exploring the Key Management Framework. Learn about cryptographic modules in Cryptographic module overview.
Use of certificate-based authentication with integration providers	Configure traffic to your integration providers using REST/SOAP connections to use certificate-based authentication. Secure Socket Layer (SSL) certificate authentication encrypts data in transit, helping to prevent it from being read as it is sent. Learn more about this configuration in Configure mutual authentication.
Use SAML authentication	Integrate third-party multi-factor authentication (MFA) with your existing SAML IdP to provide additional login security. MFA provides a high level of security because authentication requires multiple authentication factors. Something the user knows (the password) as well as something they own (a one-time code produced by an MFA token or mobile phone, or biometric attributes, such as a fingerprint). ServiceNow supports direct MFA integration with local accounts, LDAP, SSO with SAML, OIDC, or Digest. Adaptive Authentication is a prerequisite for SSO with MFA. MFA can be enabled for specified users and specified roles, and configured for ease of use. For example you can exempt recognized devices for a number of hours. You can view metrics for MFA use in the Security Center. Learn more about SAML authentication using these resources: • SAML 2.0 concepts • SAML 2.0 configuration using Multi-Provider SSO
Use the email filters feature set to deal with suspect inbound messages	Create email filters to filter out messages marked as suspicious by ServiceNow Antivirus Protection. In addition to virus protection, Antivirus Protection analyzes email for malware and SPAM, scoring and the results adding this information to the message in x-headers. You can use these headers as criteria for the Email Filters Plugin to act on if desired. Learn more about ServiceNow's antivirus feature at Antivirus Scanning.

Best Practice	Description
	Learn how to configure email filters on your instance at Email filters.
Use the Syslog Probe to send logs to your SIEM	Use the ServiceNow syslog probe to send log messages from your instance to a Security Information and Event Manager (SIEM). An SIEM is third-party software or service that can be used for activity monitoring and identifying security events. Learn more about ServiceNow syslog probe configuration at Syslog probe.
Use your own mature email security environment	Consider using your own (or third-party) infrastructure to send and receive instance-related email and benefit from more precise perimeter email control. By using your own SMTP, POP3, or IMAP servers, you can control how mail is filtered and received before being sent to your instance. **Note:** Configuration of your own email infrastructure is considered an advanced email configuration, and can optionally use a third-party email infrastructure via OAuth 2.0 email authentication. See your own email vendor documentation and instructions for details.
Validate access using Access Analyzer	Use the ServiceNow Access Analyzer tool to help you compare and analyze permissions for selected users, roles, or groups. You can use this information to troubleshoot access issues, identify who has access to your sensitive data, and determine the correct level of access for users on your instance. Learn more about the benefits of Access Analyzer at Explore Access analyzer.

Security posture dashboards

Use the customizable single and multi-instance security posture dashboards to monitor your security KPIs. These dashboards consolidate the important information regarding the security of your instances in a single location and include a number of base system dashboard widgets.

Accessing the Security posture dashboards

To access the Security posture dashboard, open Security Center by navigating to **All** > **Security Center**. Select **Security posture console** in the **Security consoles** section. On the **Security posture console** page, select **Security posture dashboards** at the top.

Use the down arrow next to **Security posture dashboard** to switch between instance dashboards.

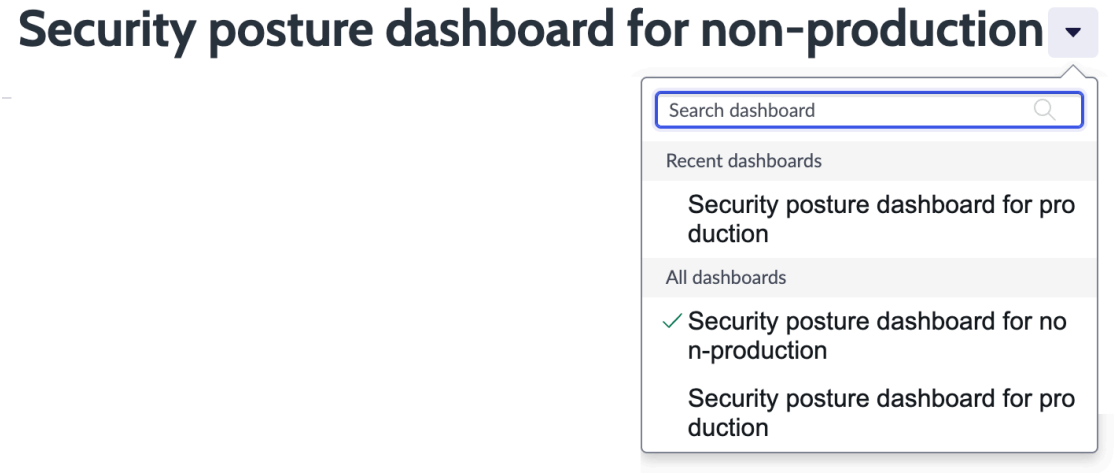

The dashboard is divided into multiple sections containing widgets related to an aspect of instance security. Select any widget on the dashboard to view more detail on this aspect of your instance's security.

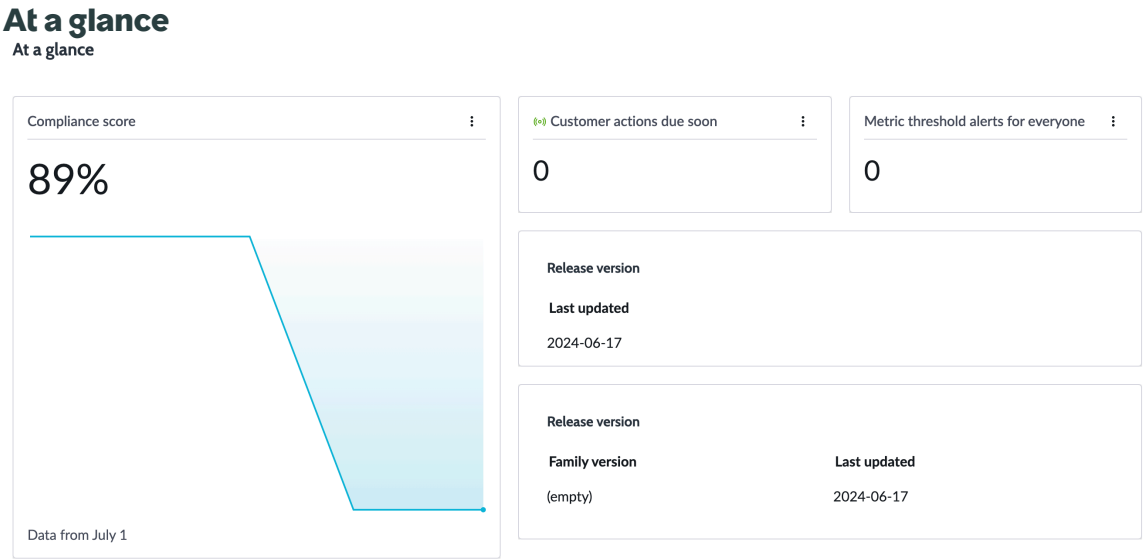

The **At a glance** section displays an overview of security on an instance, such as a compliance score, Customer Actions due, and release information for the instance.

Compliance score

Displays your instance compliance score percentage over time, beginning with the date sown at the bottom of the widget. Select this widget to navigate to the Hardening compliance score trend.

Metric threshold alerts for me

Displays a count of Metric threshold event [n_vsc_metric_threshold_event] records assigned to the current user. Select this widget to navigate to a list of these records.

Metric threshold alerts for everyone.

Displays a count of unassigned Metric threshold event [n_vsc_metric_threshold_event] records. Select this widget to navigate to a list of these records.

Customer actions due soon

Displays a count of customer actions due soon. Select this widget to navigate to Customer Actions.

Antivirus downloaded files

Displays a count of files uploaded to your instance that have been quarantined. Select this widget to view these quarantined files in the Antivirus section of Security metrics.

Release version

Displays the family version of the instance and the date of the last instance update.

Users

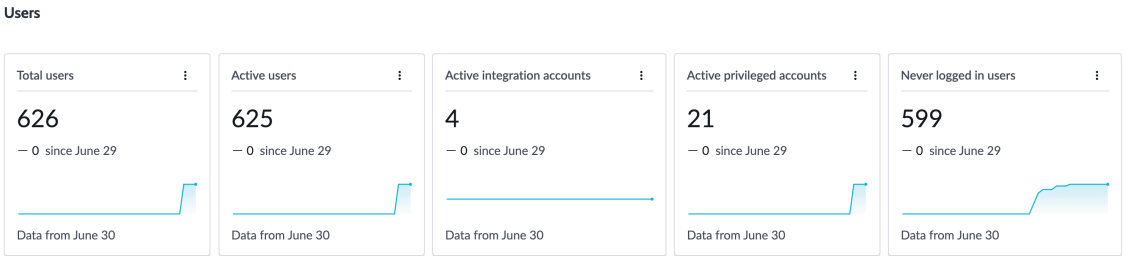

The **Users** section provides information on the users in your instance. The widgets on this section show user information, and a line graph showing changes to this information over time. Select a widget to view more detail.

Total users

Displays a count of users on your Users [sys_user] table. Select this widget to see details on these records in the Active Sessions section of Security metrics.

Active users

Displays a count of active users on your Users [sys_user] table. Active users are user records where the **Active** field is selected. Select this widget to see details on these records in the Active Sessions section of Security metrics.

Active integration accounts

Displays a count of integration accounts on your Users [sys_user] table. Integration accounts are user records where the **Web service access only** field is selected. Select this widget to see details on these records in the Active Sessions section of Security metrics.

Active privileged accounts

Displays a count of active privileged users on your Users [sys_user] table. Privileged users are user records where the **Active** field is true, the **Internal Integration User** is false, and **Internal Integration User** is inactive. Select this widget to see details on these records in the Active Sessions section of Security metrics.

Never logged-in users

Displays a count of users on your Users [sys_user] table created in the last 60 days, and have no value in the **Last login time** field. Select this widget to see details on these records in the Active Sessions section of Security metrics.

Login protection

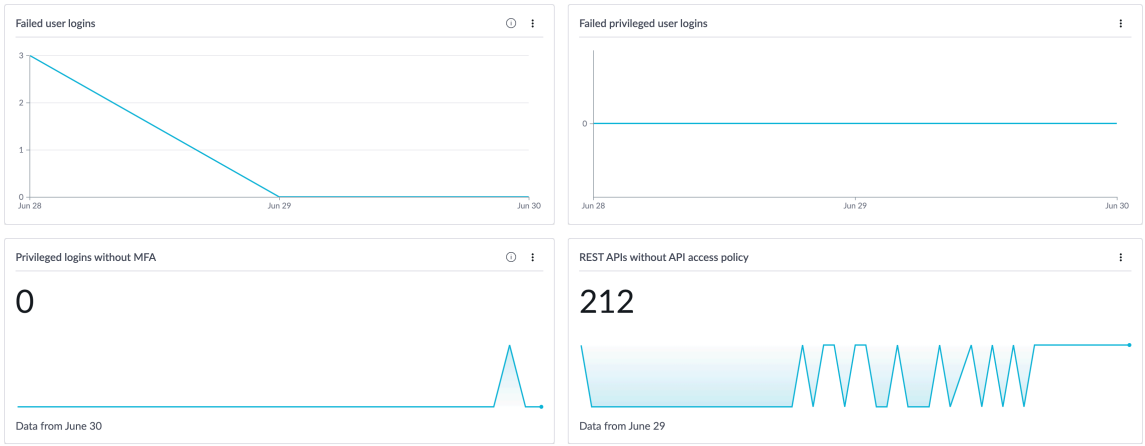

The **Login protection** section includes information on failed logins, including failed login attempts for privileged users. These widgets include a line graph showing changes to this information over time. Select a widget to view more detail.

Failed user logins

Displays a count of failed user login attempts. Select this widget to navigate to see details on failed logins in Security metrics.

Failed privileged user logins

Displays a count of failed user login attempts from privileged accounts. Privileged users are user records where the **Active** field is true, the **Internal Integration User** is false, and **Internal Integration User** is inactive. Select this widget to navigate to see details on failed logins in Security metrics.

Privileged logins without MFA

Displays a count of privileged accounts that aren't configured for multi-factor authentication (MFA). Select this widget to navigate to see a list these accounts in Security metrics.

REST APIs without API access policy

Displays a list of REST APIs that aren't restricted with an API access policy. Select this widget to navigate to see a list these REST APIs in Security metrics.

Instance hardening

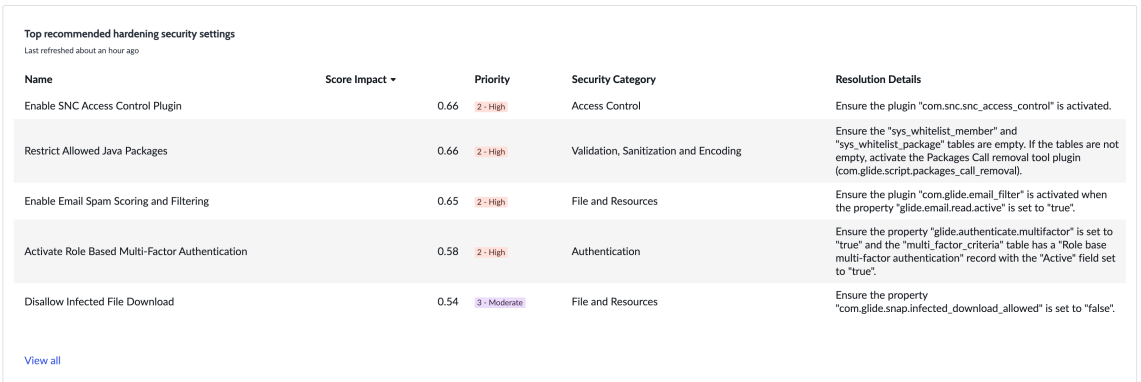

The **Instance hardening** section contains recommended hardening security settings that you can change to improve instance security. Use this section to see the priority and potential impact of these changes.

Instance trends

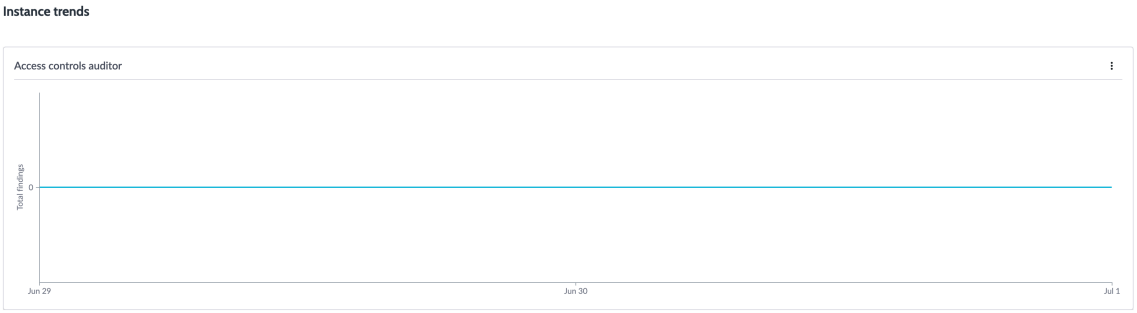

The **Instance trends** dashboard displays the results of the access controls auditor scan suite.

Data protection

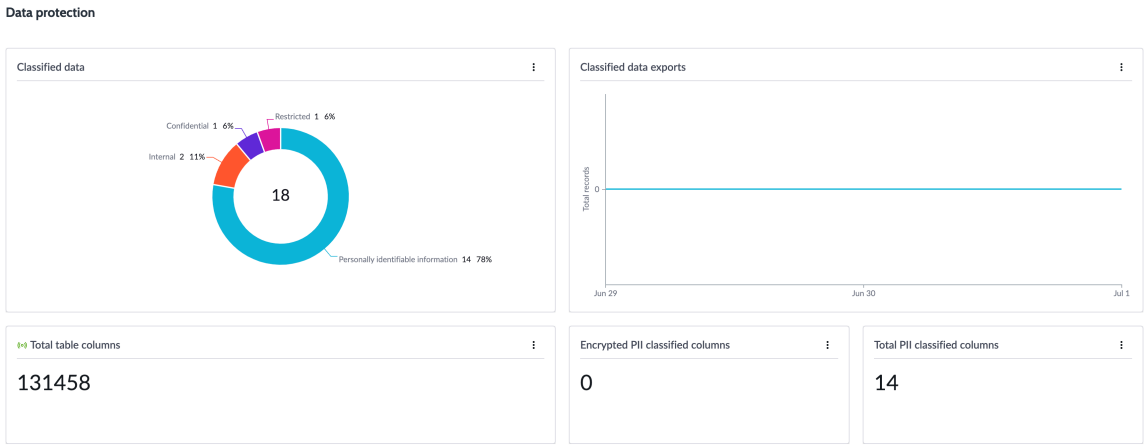

Use the **Data protection** section to see an overview of classified data, such as personally identifiable information (PII). The dashboard also tracks exports of classified data.

Classified data

Displays a pie chart of classified data on your instance, separated by type. Select a section of the chart to view details on these records. For details on data classification, see Data classification.

Classified data exports

Displays a count of classified records that have been exported from your instance. Select this widget to see a list of classified exports in the Export section of Security metrics.

Total table columns

Displays a count of all columns(fields) in tables on your instance. Select this widget to review these columns in the Data Classification section of Security metrics.

Encrypted PII classified columns

Displays a count of encrypted records classified as Personally Identifiable Information (PII). Select this widget to review these records in the Data Classification section of Security metrics.

Total PII classified columns

Displays a count of all records classified as Personally Identifiable Information (PII). Select this widget to review these records in the Data Classification section of Security metrics.

Review multiple instances

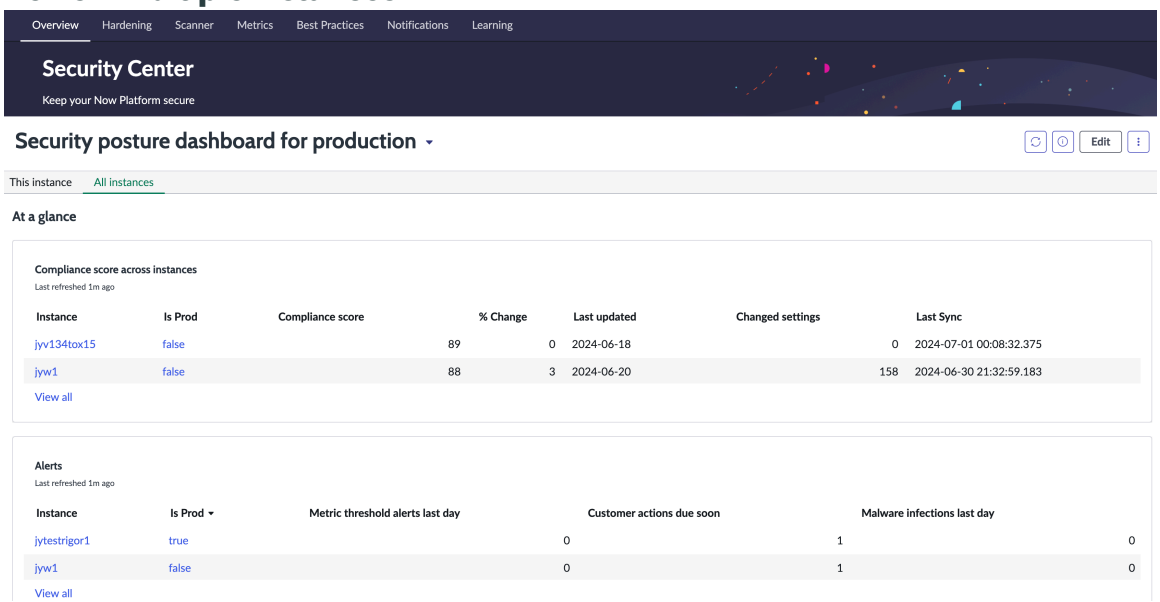

View the security posture of your non-production instances without leaving your production instance using the **All instances** tab at the top of the dashboard. The **All instances** tab displays a condensed version of the same information as the **This instance** tab, but also includes data from all your non-production instances.

By default, the **All instances** tab displays information on the production instance you're logged in to, and all non-production instances across all your production environments.

You may add or remove instances that appear on this dashboard by modifying your trust configuration. Providing data visibility between instances allows them to appear within your dashboard. For details on this process, see Basic trust configuration for data sync applications.

Dashboard customization

The instance security posture dashboard can't be customized, but you can duplicate the dashboard by selecting the **More Actions** (⋮) icon and selecting **Duplicate**. You can change the duplicate dashboard.

Security Tasks

Use Security Tasks to monitor, prioritize, and assign all your security-related tasks in one place.

Save time tracking and organizing all the tasks needed to improve and maintain your instance security posture using Security Tasks.

Review and manage Security Tasks

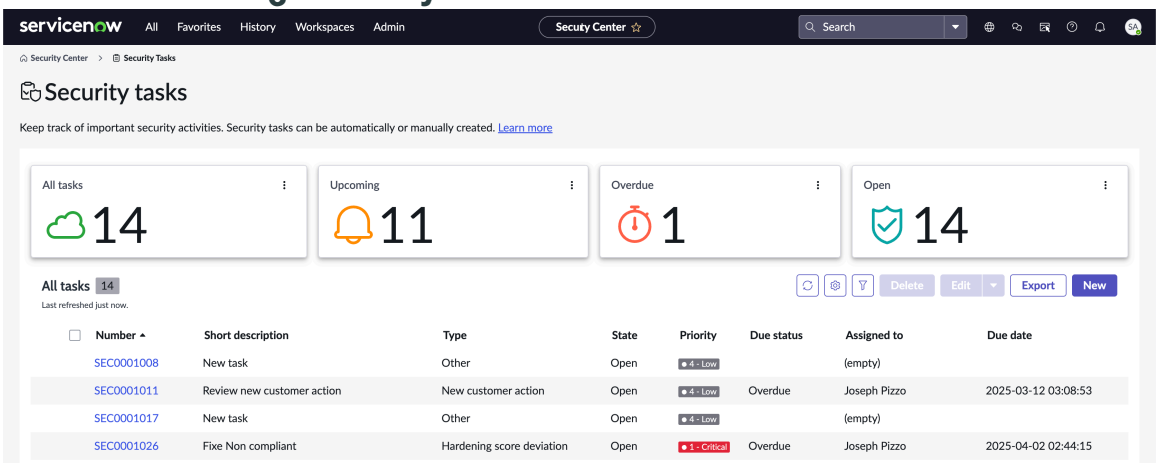

Access the Security Tasks by navigating to **All** > **Security Center** > **Security Tasks**. By default, all tasks are displayed in the list. You can easily filter upcoming, overdue, or open tasks by selecting the cards above the list. These cards also display the current count of tasks in each of these categories. You can further refine your filter by using the normal table field filtering options and then selecting the filter button ( ), then selecting the Advanced view button to create and save a custom filters to use later.

User roles

The Security Tasks use the following roles:

User	Required role	Description
System Administrator	admin	System Administrators can view, create, assign, and delete Security Tasks.
Security Center Viewer	sn_vsc_security_center_viewer	Viewers can view, but not create, delete or edit Security Tasks.
Security Task Manager	sn_vsc_task_manager	Security Task Managers can work on assigned tasks, view task manager-related pages, and create and manage tasks within Security center. They can view tasks assigned to other users but can't edit tasks assigned to other users. Task managers can't access non-task manager pages in Security Center.
ITIL	itil	ITIL users can interact with Security Tasks assigned to them, but can't view all tasks as the admin and Security Task manager roles can.

Quickly assign to yourself Security Tasks

Use the check boxes to the left of the items in the list to select multiple items. When one of more tasks is selected, you can use the **Delete** button to delete the selected tasks, or the **Edit** button to assign the selected tasks to yourself.

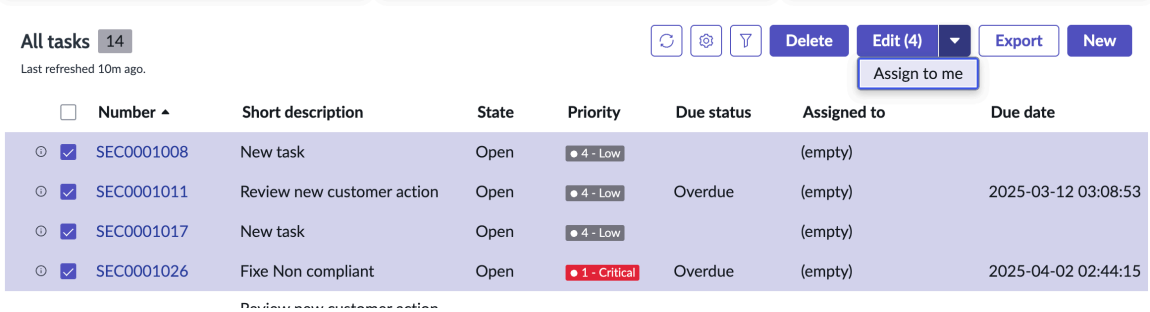

Create Security Tasks

Create a Security Task with the **+Create Task** button, which is available various platform security pages.

For example, in Security Center, on the Best Practices tool details page, the admin can select the **+Create Task** button to create and assign a task to complete this best practice.

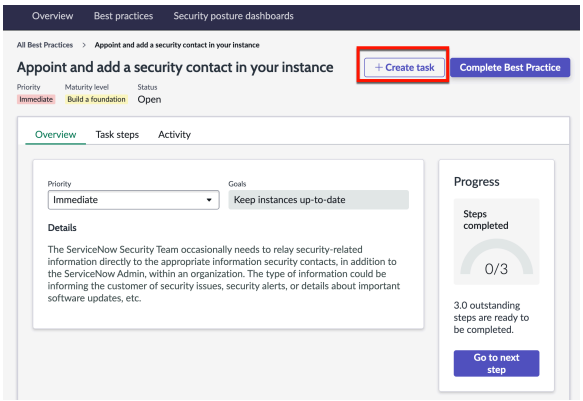

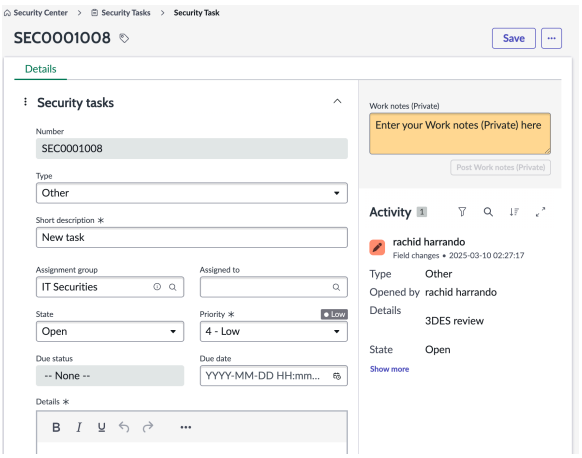

Automatically generated tasks

Security Tasks can be automatically generated. Automatic Security Task generation is triggered by an associated event that occurred on the platform. Learn more about generated Security Tasks in Automatic Security Task generation.

Edit and assign Security Tasks

Edit Security Tasks to assign them to users, define due dates, and provide additional details to the users who will complete these tasks. For details, see Edit Security Tasks.

Export tasks

You can export Security Tasks into the format of your choice. For details, see Export Security Tasks.

Automatic Security Task generation

Learn about how and when your instance generates Security Tasks.

Automatically generated Security Tasks

Security Tasks can be automatically generated. Automatic Security Task generation is triggered by an associated event that occurred on the platform. For example:

Metrics threshold breached

A Security Task is generated for a metric when it's threshold is breached. There's only one open task for a metric, even if the same metric has multiple breaches. If a Security Task for a breached metric is closed, a new task is generated is the threshold is breached again.

Event notification

A Security Task is generated when a security event notification triggers (meets conditions of it's policy). As with metrics, there's only one open task for a policy, with a new task being generated if the previous one is closed, and the notification triggers again. For details on security event notifications, see Security Event Notifications.

Hardening score deviation

When the hardening score degrades below the configured threshold value (default 3), a Security Task is generated. For example, if the hardening score is 97, and is 94 the next day, a Security Task is created just after the score of 94 is calculated. Only a single open task is created. The platform will not generate another task if the score lowers again the next day from 94 to 91. For details on your hardening score, see Hardening compliance score trend.

Customer Action

A Security Task is generated whenever a Customer Action is installed. For details on Customer Actions, see Customer Actions.

Banner announcement

A Security Task is generated for each new banner announcement. For details on banner announcements, see Security banner announcements.

Automatically generated security settings

Configuration options for automatically generated Security Tasks are found in the security center properties page. Find these settings by navigating to **All** > **Security Center** > **Security Center Properties**.

Enable/Disable automated Security Tasks

When the **Yes/No** field is selected, automated Security Tasks are enabled on your instance.

Hardening score degradation threshold

The value in the field represents the amount by which the hardening score must degrade (since the last daily score) to generate a Security Task. This value must be a positive integer. The default value is 3.

Edit Security Tasks

Learn how to create, edit, delete, or export Security Tasks in Security Center

Before you begin

Role required: admin or sn_vsc_task_manager

Procedure

1. Access your Security Task list by navigating to **All** > **Security Center** > **Security Tasks**.
2. From the Security Task list, you can edit your Security Tasks in two ways:

Edit in the Security Task form

Select a Security Task number from the list to open the Security Task record and see its details. Here you can edit, assign to another user, and add work notes.

Important: Remember that users who are assigned tasks must be admins or have the sn_vsc_task_manager role.

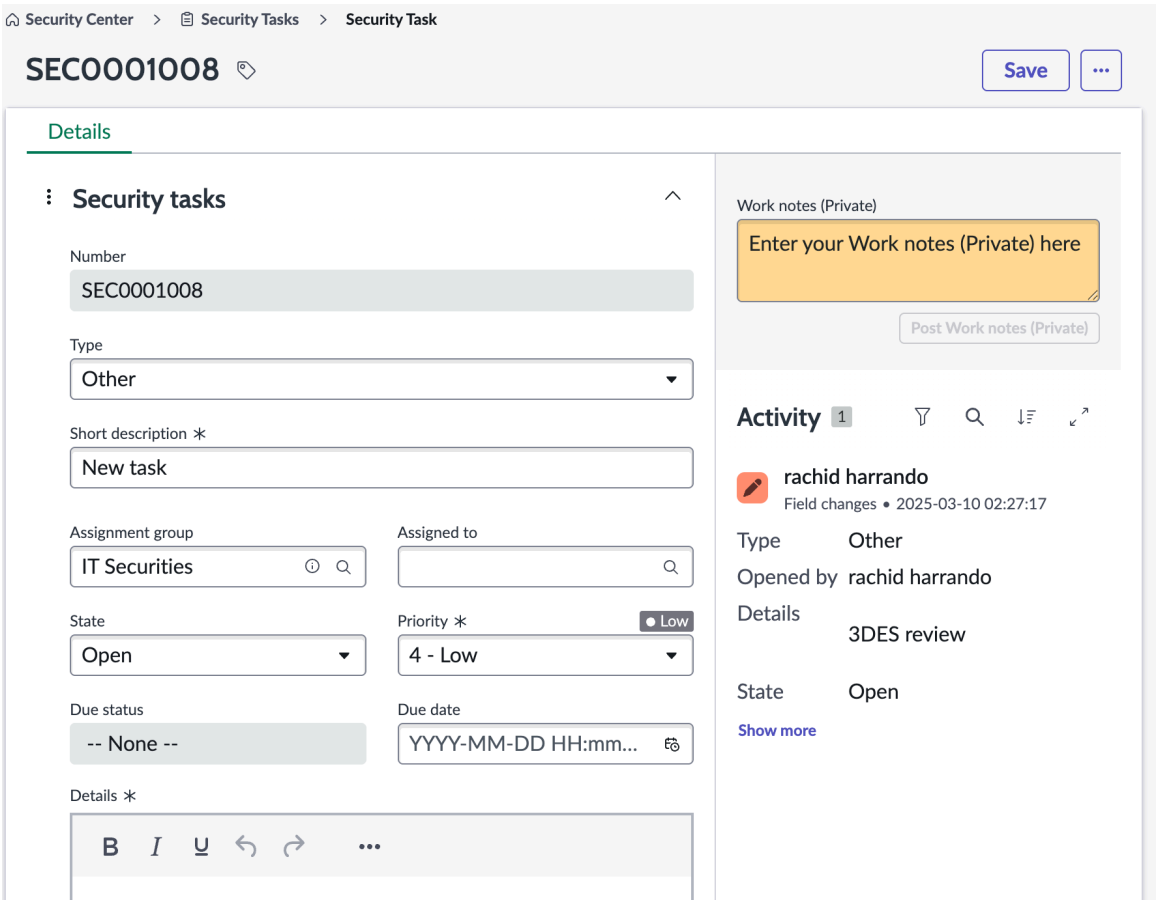

Edit from the Security Task list

Select one or more Security Tasks by selecting on its check box on the left and then select the **Edit** button.

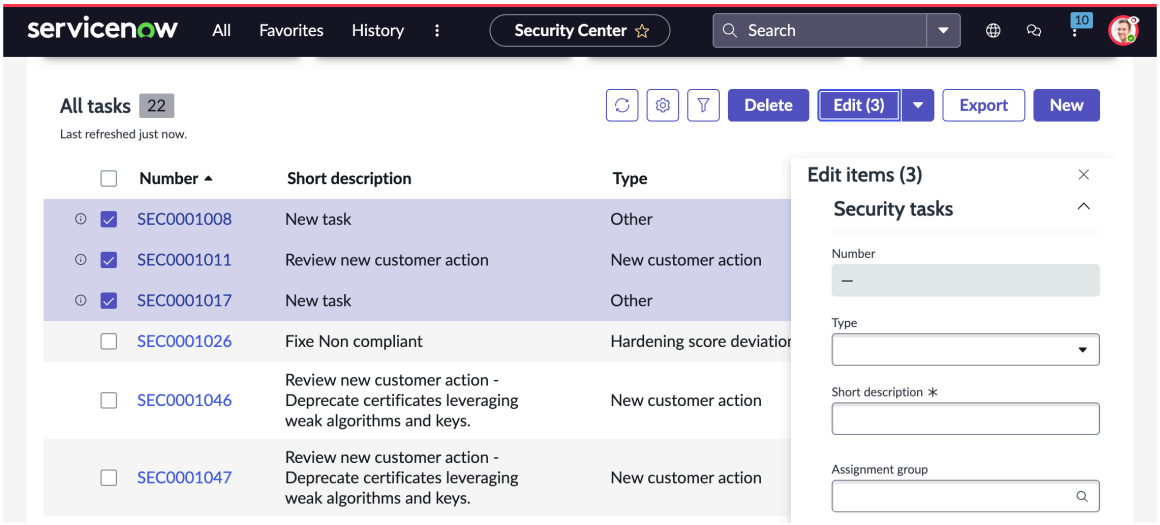

3. Fill in the fields as needed:

Field	Description
Number	Automatically generated number used to identity the task
Type	Select a task type: ○ Metrics threshold breached ○ New Customer Action ○ Event notification ○ Hardening score deviation ○ New banner announcement
Short description	Brief description of the task
Assignment group	Assignment group for the task
Assigned to	Person assigned to the task. This user must be in the assignment group selected in the **Assignment Group** field, and have the *admin* or *sn_vsc_task_manager* role.
State	State of the task: ○ Pending ○ Open ○ In Progress ○ Closed ○ Canceled
Priority	Priority of the task: ○ 1- Critical ○ 2- High ○ 3- Moderate ○ 4- Low

Field	Description
Due status	Whether the task is **Upcoming** or **Overdue**, depending on the date in the **Due date** field.
Due date	The date and time by which the task needs to be completed.
Details	Used to provide details, links, or other information users need to know about this task.
Work notes	Notes about this task. Notes entered into this field are timestamped and displayed on the Activity list on the right of the screen.
Work notes (private)	Work notes can be privately added by entering them into this field and selecting the **Post work notes (private)** button.

4. Select **Save** to save the Security Task record, if you are editing a record from the list, select **Update**.
5. **Optional:** If needed, you can delete a task from the list by selecting on its check box on the right, and selecting **Delete**, or in the form by selecting the more icon (...), then selecting **Delete**.

Export Security Tasks

Learn how to export Security Tasks into files you can download and use in other software.

Before you begin

Role required: admin

Procedure

1. Access your Security Task list by navigating to **All** > **Security Center** > **Security Tasks**.
2. Select one or more Security Tasks by selecting on its check box on the left
3. Select the **Export** button.
4. Under **File type**, select the file format you want the Security Tasks exported to. You can select from Excel, CSV, JSON, or PDF formats.
5. Under **Delivery Type**, select either **Download** or **Email**.

   Download

   The export file is downloaded using your browser.

   Email

   The **Email** field displays. Enter a email address in the **Email** field, and the export file is delivered to the specified email address.

Security learning

Access security learning materials from a single page.

Read security white papers, ebooks, knowledge base (KB) articles, product documentation, and community discussions from a consolidated view by navigating to **Learning**. The content is organized by headings and descriptive UI cards to make it quick for you to pinpoint the correct resource.

Security banner announcements

Enable security banner announcements to stay informed about urgent and critical security alerts using high visibility banners visible to administrators within the instance UI.

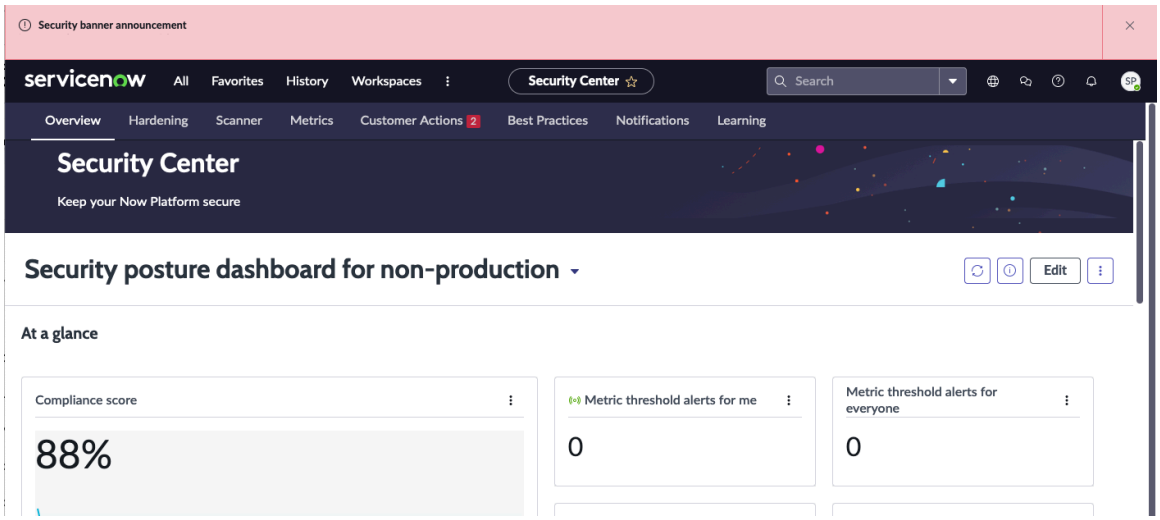

Security banner announcements are announcements displayed to customer administrators, sent by ServiceNow, to keep you informed about fixes for potential security threats that were discovered recently. These alerts contain a summary of the security risk and include a link where you can learn more and act to secure your instance.

Administrators can dismiss banner by selecting the close (x) button, but the banner re-appears with each new session until the banner expiration date is passed. Administrators can disable banner announcements by setting the *`sn_vsc.configure_customer_push_action`* system property value to `false`. The duration for which the banner appears is controlled by **Start** and **End** field values for the banner. These fields are found in banner's record in the Banner Announcement [sys_ux_banner_announcement] table.

Enable or disable security banner announcements

The security banner feature is enabled by default. To enable or disable security banner announcements, navigate to **System Security** > **Security Center** > **Notifications**. From this page, select the **Manage announcement settings** button.

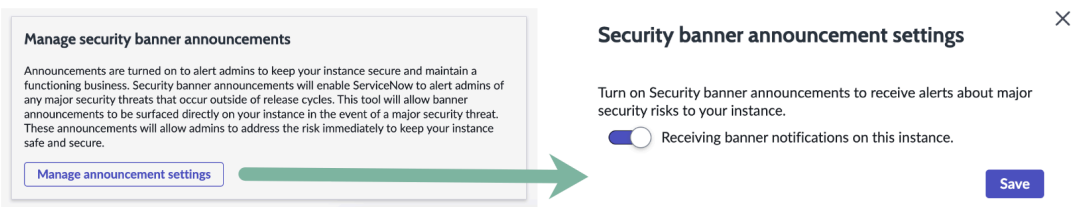

Instance Security Center

Monitor the compliance level of instance security controls, view security event monitoring metrics, and configure and maintain instance security settings all from within the Instance Security Center. The Instance Security Center consolidates several key security components into a single control console that helps you detect, protect, and respond to instance-based security events.

Important:

Instance Security Center (ISC) has reached the end of sales as of September 2024, and is no longer supported or available for new activation.

ServiceNow Security Center (SSC) is the recommended solution going forward. For more information, see Instance Security Center to ServiceNow Security Center migration.

Instance Security Center components

To access the Instance Security Center, navigate to **System Security** > **Instance Security Center** or the System Administration homepage.

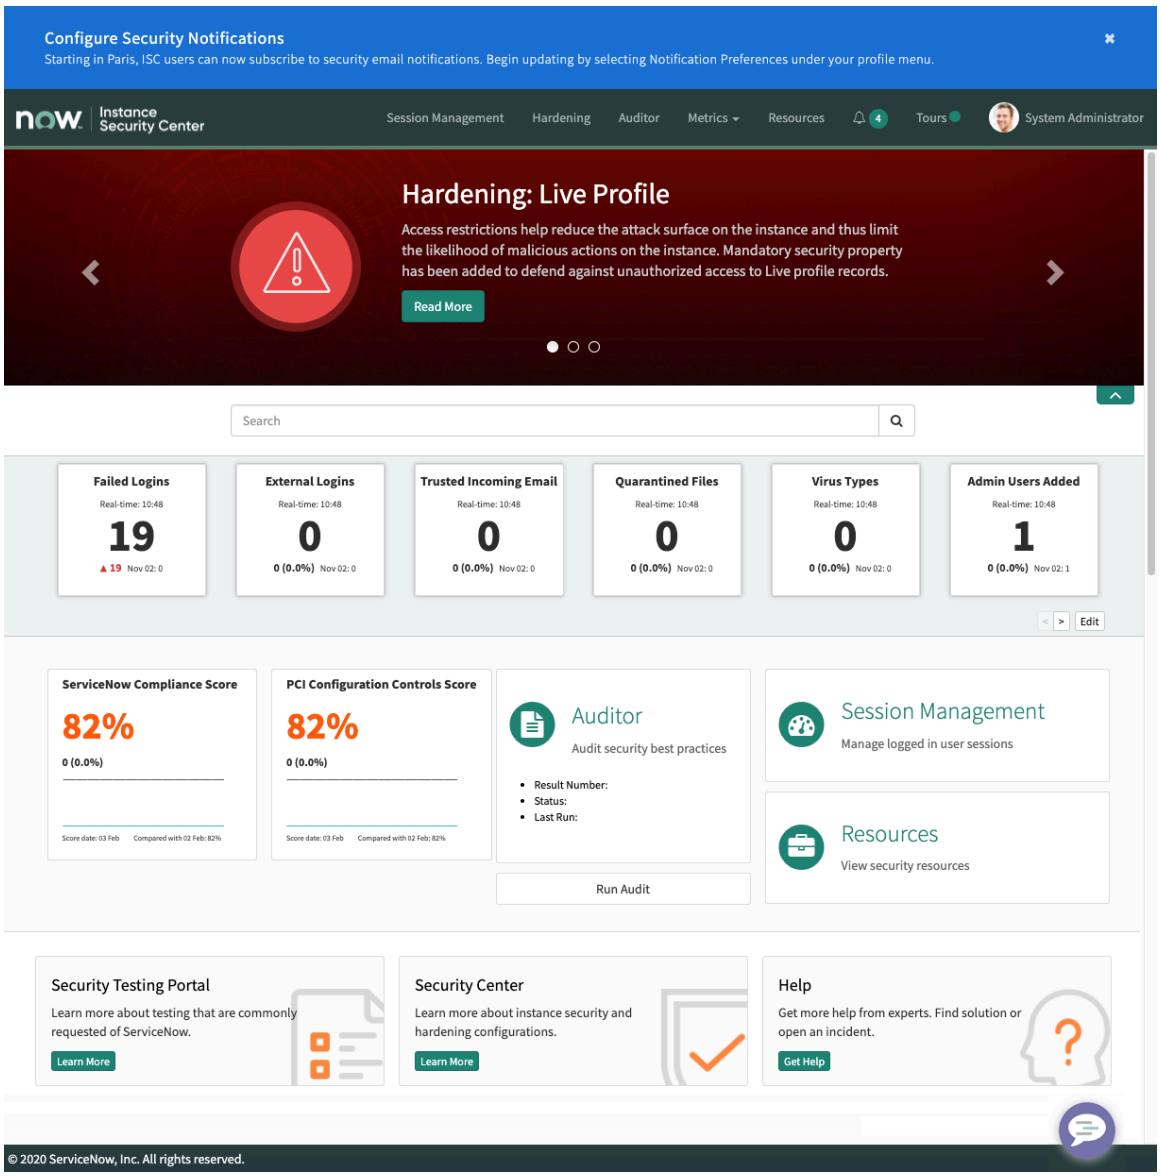

The Instance Security Center homepage contains the following security components:

- Administrator messages
- Rotating security banner
- Search
- Security event ribbon
- Daily compliance score

- PCI Configuration Controls Score
- Session Management
- Hardening
- Auditor
- Metrics (user, email, and antivirus)
- Resources
- Security notifications
- Tours
- Security testing portal
- Security center
- Help
- Virtual Agent access

From the Instance Security Center homepage, you can view the security compliance score for your instance and monitor its overall security health. You can then configure or update system properties that are related to your instance security so that they comply with security requirements.

Note: The Instance Security Center does not support domain separation.

User roles

To use the Instance Security Center, you must have the admin or security_dashboard_user role.

To learn more about managing per-user subscriptions, see Managing per-user subscriptions in Subscription Management and contact your account representative.

roles

Required role	User	Benefits
admin	This role has access to all system features, functions, and data because administrators can override access control list (ACL) rules and pass all role checks. Avoid assigning this role to your users when more targeted roles are available.	Leverage Security Center tools to improve instance security posture and monitor security related behaviors.
sn_vsc.security_center_viewer	This role allows users who are not Admins to view the information in Security Center but not make changes to the Security Center tools or make instance configurations change leveraging the Security Center tools.	Gain visibility into Security Center tools to monitor instance security posture and monitor security behaviors.

roles (continued)

Required role	User	Benefits
	For example, platform owners, security operations analysts or compliance stakeholder might want or need to view some of the security KPIs, security insights and security learning material available in Security Center.	

Warning: To ensure that the Instance Security Center receives up-to-date security information with every upgrade, do not customize this module. If you change any security settings on your instance, make sure that you test them in a non-production environment first.

Administrator messages

Messages and reminders, intended mostly for administrators, appear above the rotating security banner.

For example, a `Configure Security Notifications` message appears to remind administrators to configure preferences for security notifications if they have not done so. It also points them to the proper page to do so.

Note: The administrator messages banner does not appear for non-admin users, or if there are no actionable items for admin users.

Rotating security banner

To assist you in monitoring the security health of your instance, critical instance security messages appear in the rotating banner.

- Two to three security messages normally rotate at a regular interval.
- The dots at the bottom of the banner show you the total number of current security messages.
- To navigate through them, select the dots, or select the arrows that appear on either side of the messages.

The banner background colors indicate the relative severity of the messages.

Color	Description
Red	Critical security situation requiring a timely response, or a recommendation on how to protect or respond to critical security events.
Dark gray	Non-critical warning message.
Blue	General information message.

To collapse or minimize the text content in the banner, select . To maximize the text content, select it again.

- When you use the Instance Security Center again, the text content appeared as collapsed or expanded, depending on how you used it during your previous session.
- If the text content itself changes, it appears as maximized for all users.

Search

Use the search bar to search the entire Instance Security Center for security resources that assist you with understanding and resolving security issues. You can search the following security-related resources:

- Now Support Knowledge Base articles
- Instance Security Center pages
- External Now Support links
- PA security widgets, such as the Daily Compliance Score and External Incoming Emails
- Banner content

Event ribbon

Use the event ribbon to view key security event monitoring metrics for the current instance.

- To manually scroll through the metrics, select the right or left arrow keys.
- To configure the event ribbon, select **Edit**.

To learn more about the event ribbon and how to configure it, see Monitor security events and Configure the security event ribbon.

Daily compliance score

The Daily Compliance Score section contains the **Daily Compliance Score**, **Session Management**, **Hardening**, **Auditor**, and **Resources** tiles.

You use the Daily Compliance Score to gauge how healthy your instance is from a security standpoint.

The Daily Compliance Score is a percentage score. It is based on how compliant the current settings of your instance security properties are with the compliance values published in the Hardening settings.

- To learn more about Daily Compliance Score calculations, and how hardening settings impact it, see Check the daily compliance score and configure security property settings.
- The **Refresh** button enables an administrator to instantly recalculate the Daily Compliance Score. To learn more, see How Daily Compliance score, trend, and graph data is refreshed.

Hardening

Use this process to adjust the specific security configuration properties that affect the Daily Compliance Score:

1. To access the Hardening Compliance Configurations page and perform instance security hardening, select the **Daily Compliance Score** tile or the **Hardening** link.
2. Specify whether you want to view all or only recommended security controls. Then, select the category you want to work in.
3. Set each security configuration property in the selected category. Click **More Info** to view detailed information for a property.

To learn more about hardening and optimizing security configuration properties to further increase compliance, see Adjust instance security settings to increase compliance.

To learn more about how trend and graph data is refreshed, see How Daily Compliance score, trend, and graph data is refreshed.

Auditor

Run the Auditor to scan your instance and find incorrect security definitions. It provides findings you can correct to help improve the security posture of your instance.

To access the Auditor page, select the **Auditor** tile or the **Auditor** link. To learn more, see Scan for incorrect security definitions.

Session management

Use Session Management to:

- View and manage user login sessions.
- See the user login session of the current node that you are connected to.
- See detailed information about each session, such as the user name and IP address.
- Isolate and lock out specific user sessions that pose security risks.

To access the Session Management page, select the **Session Management** tile or link.

Field	Description
User	Name of the user associated with this login session. • To locate a specific user session, select the spotlight search icon () to search by user, user agent keyword, or IP address. For example, if you want to find all current logins from a specific type of browser, enter the browser name as a keyword into the **User Agent** field. • Click a user name to access the user profile record. You can modify the user profile only if you have an assigned admin role. **Note:** To learn more about user profiles, see Create a user.
MFA	Check box indicating if Multi-factor Authentication (MFA) is enabled for the logged in user. To learn more about MFA, see Multi-factor authentication.
Active	Check box indicating if the logged in user is active or inactive.
User Agent	Type of browser and the device operating system for the user login session.
IP Address	IP address of the logged in user.
Last Accessed	Date and time this user session last accessed the instance.

Field	Description
	Note: To view detailed information for a particular login session, or to lock out the session itself, select the **User Agent**, **IP Address**, or **Last Accessed** fields.

Metrics

View detail for the following types of metrics:

Type of metric	Description
User	Security metrics that are associated with user activity in the instance. To access the User Metrics page, select the **Metrics** link, and then select **User Metrics**.
Export	Security metrics that are associated with data exported from the instance. To access the Export Metrics page, select the **Metrics** link, and then select **Export Metrics**.
Authentication	Security metrics that are associated with authentication, such as infrequently used IP addresses, failed logins, and types of authentication schemes used by your users. To access the Export Metrics page, select the **Metrics** link, and then select **Authentication Metrics**.
Email	Security metrics associated with anomalous behaviors related to the incoming emails to your instance. To access the Email Metrics page, select the **Metrics** link, and then select **Email Metrics**.
Antivirus	Security metrics that are associated with antivirus event activity in the instance. To access the Antivirus Metrics page, select the **Antivirus** tile or select the **Metrics** link, and then select **Antivirus**.

Note: To learn more about monitoring each type of metric, see Monitor instance metrics.

Resources

Access Now Support Knowledge Base articles, resources, and blogs that are related to instance security. These resources include security settings, coding, compliance, fixes, and related topics. To access the Resources page:

1. Click the **Resources** tile or link.
2. In the Resources page, select a category:

Category	Description
Recommended Guidelines	Access to recommended security guidelines, including the Hardening settings and Secure Coding Guide [KB0623354] articles.
Security Resources	Access to security-related resources in the Knowledge Base, including: ○ Customer Instance Security Testing ○ Cloud Security, Trust, and Compliance Center KB articles

Security notifications

A notifications bell icon () appears in the Instance Security Center, with a count total of unread security notifications. Notifications persist and are included in this count until you mark them as read.

1. Click the bell icon to view the first five unread security notifications.

   A notification appears when **Admin Login**, **Admin Unlocked**, **Failed Login**, **High Privilege Role**, **Impersonation**, **Security Elevation**, and **Weekly Digest** events take place in your instance. To learn more about these security events, see Monitor security events.

2. To view detailed information for a specific security event, select the notification.

   For example, if you select a High Privilege Role notification, you can view the Roles (sys_user_role) table. Use this table to see which users were assigned privileged roles during the calendar day. Using this history helps you to determine if roles have been properly assigned.

3. If there are more than five unread notifications, select **View All Notifications** to access an All Notifications page with a listing of all unread notifications.
   - To view detailed information for a specific security event, select the notification.
   - To mark all listed notifications as read, select **Mark All as Read**.

Note: As an administrator, you can also configure preferences for sending specific types of notifications for each type of security event. To learn more, see Set preferences for security event notifications.

Tours

Click the **Tours** link to view a guided visual tour of the Instance Security Center.

- The guided tour includes only the security monitoring functions that are listed on the homepage.
- It does not include the security functions that you access when you select the tiles or links on the homepage.

Security testing portal, security center, and help

The Now Support Service Portal is a central resource that you use to manage instances, tasks, and accounts. You can also access useful resources you can use to diagnose and resolve security and technical issues in your instance.

To access these resources, select **Learn More** or **Get Help** in the following tiles:

Tile	Description
Security Testing Portal	Access to the Security Dashboard in the Now Support Security Testing Portal.
Security Center	Access to Security Compliance in the Now Support Security Portal.
Help	Access to the following help resources in the Now Support Security Portal: • Ask an expert to find answers to common questions. • Report an issue or outage to ServiceNow Global Technical Support by opening a case. • Self-Service Support Resources, including: ○ Videos ○ Documentation ○ ServiceNow Community ○ Knowledge Base ○ Known Error Portal ○ Security RFX Database • ServiceNow Community questions that are recommended for your use.

Virtual Agent access

The Virtual Agent is a platform for providing user assistance through conversations within a messaging interface.

With the associated plugins installed, administrators can access the Virtual Agent and Natural Language Understanding (NLU) functions by selecting the Virtual Agent icon:

It enables you to perform the following tasks:

- Ask security-related questions, then get quick summary answers and reference links to learn more.
- Get answers related to processes such as:
  - Instance Security Center
  - Platform Security
  - ServiceNow security policies
  - Trust, Governance, and Risk
- Search for security-related resources, such as Knowledge Base topics in the Now Support Security Portal.

Note: To learn more about how to use and activate the Virtual Agent, see:

- Virtual Agent
- Activate the ISC Virtual Agent interface

Instance Security Center to ServiceNow Security Center migration

Learn the key differences when migrating from Instance Security Center (ISC) to ServiceNow Security Center (SSC).

Important:

Instance Security Center is a legacy product that will reach the end of sales by September 2024. ServiceNow Security Center is the recommended solution for customers to adopt going forward. For more information, see the deprecation process article (KB0867184) in the Now Support knowledge base.

ServiceNow Security Center (SSC) is a security application that consists of a set of purpose-built tools designed to help organizations maintain the security of their ServiceNow deployments. Using SSC, organizations can improve their security posture, strengthen their compliance levels, and do so with a seamless user experience.

Use this document to learn about SSC enhancements in functionality and how to execute tasks in SSC that were previously done in ISC.

Security hardening

Feature enhancements	ISC vs SSC
User interaction updates	• Review latest score by navigating to the Overview or hardening comparison page. • Manually update latest score by navigating to Overview and selecting update score. • Change setting configuration by selecting a setting name link on any of these pages and then make edits on the **Settings detail** page: ○ Overview ○ **Hardening** > **All settings** ○ **Hardening** > **Comparison**
Set a schedule for to recalculate your compliance score	Using SSC, you can set when the system triggers a refresh of your compliance score. **1.** Navigate to the Scheduled Script Executions [sysauto_script] table. **2.** Find and open the `SC - Calculate Compliance Monthly` record. **3.** Use the **Run**, **Day**, **Time zone**, and **Time** fields to set your preferred schedule. 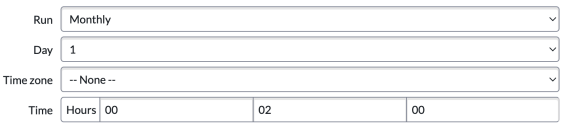 **4.** Select **Update** to save your changes.

Security KPIs and metrics

Feature enhancements	ISC vs SSC
Enhancements to KPIs	Using SSC, you can: • See the trend of the score over time. • Monitor all KPIs and metrics using the same interface and analytical capabilities.
Enhancements to metrics	SSC includes these metrics enhancements: • Over 65 metrics have been added. • Monitor and analyze KPIs and metrics from a single user interface. • Create targets and thresholds for metrics.
User interaction updates	Access Metrics in SSC by navigating to **Overview** > **Metrics**, and then selecting metric from the navigation menu. Alternatively, you can navigate to **All metrics**.

Security scanner

Feature enhancements	ISC vs SSC
New functionality	SSC includes these functions: • Manually execute scans or schedule them to run at specific times. • Create your own scan checks. • Create your own scan suite. • Compare the results of two scans.
Enhanced email functionality	Set up email notifications to stay informed about security events on your instance.
User interaction updates	• Run the Auditor scan by navigating to **Overview** > **Scan** > **Suites** and selecting **Auditor suite** from the list. • See scan results by navigating to **Overview** > **Scan** > **Results**.

Additional security features

Feature enhancements	ISC vs SSC
Monitor sessions functionality	• Visualize session activities using the Analytics Hub. • Monitor session activity by navigating to **Metrics** > **Sessions**.
Find resources	Access all learning resources in a single page at **Overview** > **Learning**.

Opt out of Security Center features

To disable Security Center features on your instance, follow the steps provided in KB1702514: Guidelines to disable Instance Security Center components (while using Security Center).

Monitor security events

Analyze the event metrics in your instance so that you can identify and prevent potential security events.

Important:

Instance Security Center (ISC) has reached the end of sales as of September 2024, and is no longer supported or available for new activation.

ServiceNow Security Center (SSC) is the recommended solution going forward. For more information, see Instance Security Center to ServiceNow Security Center migration.

In the event ribbon, which is on the Instance Security homepage, you can analyze these metrics and accompanying detail to identify potential security events in the instance.

- For each event metric, a real-time single score count appears, indicating how many times that the event occurred during the day in this instance. These single score reports are updated automatically as the corresponding events take place.
- Each event metric also contains compliance trend and graph information over a range of dates. This information updates on a daily basis when you run the performance analytics job. To learn more, see the **Analyzing event trend detail** section.

Event types

You can monitor at least six of the following types of events. For more than six events, use the left or right arrows below the event ribbon to scroll through them. To learn how to configure the event ribbon, see Configure the security event ribbon.

Notification preference	Description
Admin Logins	Number of login attempts in this instance, during the calendar day, by users who have an assigned admin role.
Admin Users Added	Number of users with an admin role that were added in this instance during the calendar day. For example, your instance may have a security issue if the count is 10, but 4 users are known to have an assigned admin role.
External Incoming Email	To learn more, see Email metrics.
External Logins	Number of users with an assigned snc_external role who logged into this instance during the calendar day. These logins typically occur for maintenance, support, consulting, or audit purposes. Monitoring this metric enables you to verify that the external login attempts are legitimate and not potential security issues.

Notification preference	Description
	To learn more about assigning external user roles, see Explicit Roles.
Failed Logins	Number of attempted logins that failed in this instance during the calendar day. This metric may indicate that attempts are being made to log in and compromise your instance security.
Impersonations	Number of impersonation logins in this instance during the calendar day. To learn about impersonating users, see Impersonate a user.
Quarantined Files	Number of files that were quarantined when you ran Antivirus Scanning in this instance during the calendar day. To learn more about quarantined files and Antivirus Scanning, see Antivirus metrics and Antivirus Scanning.
Security Elevations	Number of times that a security administrator has elevated security for standard users by changing their assigned user role to a high privilege security role during the calendar day. These high privilege security roles include oauth_admin, admin, security_admin, and impersonator. • This metric indicates that someone might have tried to elevate the security of an unauthorized user. Do not use this metric by itself to detect a specific security compromise. Instead, treat this metric as an indication that you should check another metric to see if a security compromise has occurred. • To learn more about elevating user security, see Elevate to a privileged role and Elevated privilege roles.
SNC Logins	Number of Customer Service and Support personnel who logged into this instance using the hi-hopping technique during the calendar day. These logins typically occur for maintenance, support, consulting, or audit purposes. For information on how to control ServiceNow corporate employee access, see ServiceNow access control.
Spam	To learn more, see Email metrics.

Notification preference	Description
Trusted Incoming Email	To learn more, see Email metrics.
Untrusted Incoming Email	To learn more, see Email metrics.
Virus Types	Number of different types of antivirus events that occurred in this instance during the calendar day. To learn more about antivirus event types, see Antivirus metrics.

Analyzing event trend detail

To view trend details for an event metric, click the event count to access the Analytics Hub page. The details that appear for the instance depend on the type of metric.

For example, to view a listing of each failed attempt on the Security Dashboard Event Logs page:

- Select the **Failed Logins** metric.
- In the Analytics Hub page, click **Show Records**.
- Click one of the failed login attempts.
- The detail includes the name of the user who attempted to log in, their IP address, and the table name that they tried to access.

You can set up event threshold triggers in the Core UI Analytics Hub or Platform Analytics KPI Details to provide alerts when a certain event occurs within a range of scores for an indicator. You can also set targets that enable you to visualize the difference between the desired score and the actual score of an event.

For example, you can set a threshold of `10` for the **Failed Logins** metric. When ten or more failed login attempts occur during the day, an alert is sent to specific security personnel. You can also set a similar target that provides a visual highlight in the Analytics Hub when 10 failed logins occur during a day.

Trend data and graphs that appear in Event ribbon tile and the Analytics Hub are updated after the performance analytics job executes at 02:00 local time. To learn more, see How Daily Compliance score, trend, and graph data is refreshed.

Related topics

Instance Security Center

Now Intelligence

Analytics Hub

Performance Analytics targets and thresholds

Configure the security event ribbon

Configure the security event ribbon on the Instance Security Center homepage to include only those events that are relevant for tracking instance security in your operations. You can also change the order in which the security event tiles appear on the ribbon.

Before you begin

Role required: security_dashboard_user or admin

About this task

The security event ribbon is initially populated with a full set of standard security events. You can customize it by removing the events that are not relevant to your organization.

- For example, if you suspect that security issues are due to the actions of internal personnel, include the Admin Logins, Admin Users added, and Security Elevations event indicators.
- These indicators monitor how many times users with admin roles attempted to log in and, if admin users were added, what attempts were made to elevate security roles.

Note: To learn about the types of security events that appear in the event ribbon, see Monitor security events.

Procedure

1. Navigate to **All** > **System Security** > **Instance Security Center**.
2. In the event ribbon, click **Edit** (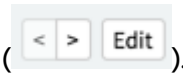

).
On the Edit Events form, the **Selected** column contains the events that are already listed.
3. To add security events to the event ribbon, move them from the **Available** column to the **Selected** column.
To change the order the events appear on the ribbon, select an event, then click the up or down arrow to move it to its correct position.
   - Place the events in the same sequential order they should appear on the Instance Security Center event ribbon.
   - The events you place at the top of the **Available** column appear sequentially, beginning on the left side of the Instance Security Center event ribbon. The events placed towards the bottom of the column appear on to the right on the event ribbon.
4. To remove security events from the event ribbon:
   - In the **Selected** column, select the security events that you want to remove from the event ribbon.
   - Move them from the **Selected** column to the **Available** column.
5. Click **Save**.

Related topics

Instance Security Center

Set preferences for security event notifications

Configure preferences for the types of notifications you want to receive for occurrences of specific security events. For each type, you designate whether to receive notifications by email, by push notification in Now Mobile, or in third party messaging applications such as Slack or Microsoft Teams.

Before you begin

To enable third party messaging applications to send security event notifications, you must activate the Messaging Notification (com.glide.notification.messaging) plugin. Individual users must configure their own settings. For details, see Notifications in messaging applications.

Role required: admin.

Procedure

1. In the Instance Security Center home page, click the profile menu, then click **Notification Preferences**.

Notification preferences

Notification preference	Description
Admin Login	Send the selected type of notification whenever other users with assigned admin roles log into this instance from a different IP address.
Admin Unlock	Send the selected type of notification whenever an account for a high privilege user has been unlocked.
Failed Login	Send the selected type of notification whenever other users fail to log in into this instance in less than the number of attempts defined in the *glide.user.max_unlock_attempts* property. If you don't configure this property, the default value is 5. To learn more about this property, see Specify lockout for failed login attempts.
HP Role Added	Send the selected type of notification whenever a high privilege security role (including oauth_admin, admin, security_admin, and impersonator roles) is granted to another user. To learn more about elevating user security, see Elevate to a privileged role and Elevated privilege roles.
Impersonation	Send the selected type of notification whenever another user is impersonating you. To learn more about impersonating users, see Impersonate a user.
Security Elevation	Send the selected type of notification whenever other users are elevated to a security admin role in this instance.
Weekly Digest	Send a weekly digest on the selected type of notification. It includes: ○ A summary all security activities that took place in this instance throughout the week. ○ The current daily compliance score for the instance.

2. For each type of security event, select the appropriate check boxes to designate the type of notifications to send you.

You can select multiple notification methods for each.

Check box	Description
Email	Send an email for this type of security event.
Slack	Send notifications for this type of security event through Slack. **Note:** This column appears only if you have set up Slack integration to the ServiceNow AI Platform.
Teams	Send notifications for this security event through Microsoft Teams. **Note:** This column appears only if you have set up Microsoft Teams integration to the ServiceNow AI Platform.
Push	Send push notifications on Now Mobile for this type of security event. **Note:** This column appears only if you first log into Now Mobile.
Select All	Select a specific type of notification for all types security events. For example, click **Select All** located above **Email** if you want to receive email notifications for each security event type.

3. Click **Save**.

Check the daily compliance score and configure security property settings

Review the Daily Compliance Score metric and security configuration properties to see if your instance complies with the suggested security requirements. You can affect the daily compliance score by updating non-compliant security properties in the Hardening Compliance Configurations page.

Important:

Instance Security Center (ISC) has reached the end of sales as of September 2024, and is no longer supported or available for new activation.

ServiceNow Security Center (SSC) is the recommended solution going forward. For more information, see Instance Security Center to ServiceNow Security Center migration.

The Daily Compliance Score is a percentage score. It is based on how compliant the current settings of your instance security properties are with the compliance values published in the Hardening settings.

Review the daily compliance score of your instance regularly. Follow these guidelines when you are evaluating your daily compliance score:

- Greater than or equal to 90% indicates that the instance is compliant with critical security controls.
- Greater than or equal to 50% and less than 90% indicates a moderate level of security compliance.
- Less than 50% indicates a low level of security compliance.

Adjust instance security settings to increase compliance

Using the Hardening Compliance Configuration page, harden and optimize non-compliant security properties that affect the daily compliance score of your instance. Its use ensures that your instance complies with the published security hardening standards, while fulfilling your company's security requirements.

Before you begin

Role required: security_dashboard_user or admin.

Refer to the Hardening settings content for detailed descriptions, and compliance values, for the security-related system properties and plugins in the ServiceNow AI Platform.

- Consult the Instance Security Hardening Settings whenever you set or update security-related properties, even if some of the compliance values may not be suitable for your instance.
- When you are updating these properties, ensure that the instance continues to behave as expected. Consult with the appropriate internal personnel who have the expertise to determine the security impacts.

Note: If you have an admin role, you can view and edit security controls. If you have a security_dashboard_user role, you can view security controls, but you cannot edit them.

Procedure

1. Navigate to **All** > **System Security** > **Instance Security Center**.
2. Click the **Daily Compliance Score** tile or the **Hardening** link to access the Hardening Compliance Configuration page.
3. In the Hardening Compliance chart, view the statistics for compliant and non-compliant security configuration properties.

Option	Description
Compliant	Number of security configuration properties that comply with the compliance values in the Instance Security Hardening Settings. **Note:** You cannot change the settings for compliant security properties in the Hardening Compliance Configuration. If you want to do so, you must update them in System Properties. To learn more, see
Non-Compliant	Number of security configuration properties that do not comply with the compliance val ues in the Instance Security Hardening Set

Option	Description
	tings. You can update settings for non-compli ant properties.

Note: To view the number of compliant or non-compliant security scores over a range of dates, move the blue dot on the slider below the Daily Compliance Score.

4. In the **Show** list below the chart, specify whether you want to access all security configuration properties, or only recommended ones.

Option	Description
All	(Default) All compliant and non-compliant se curity configuration properties in each select ed category.
Recommended	Only recommended security configuration properties appear in each selected category. These security configuration properties are a selected subset of the most critical ones used to secure the ServiceNow AI Platform. Consider these security configuration prop erties to be the bare minimum number of set tings you must set to secure the ServiceNow AI Platform. **Note:** To fully secure your instance, use the **All** option. It includes all recom mended security configuration proper ties too.

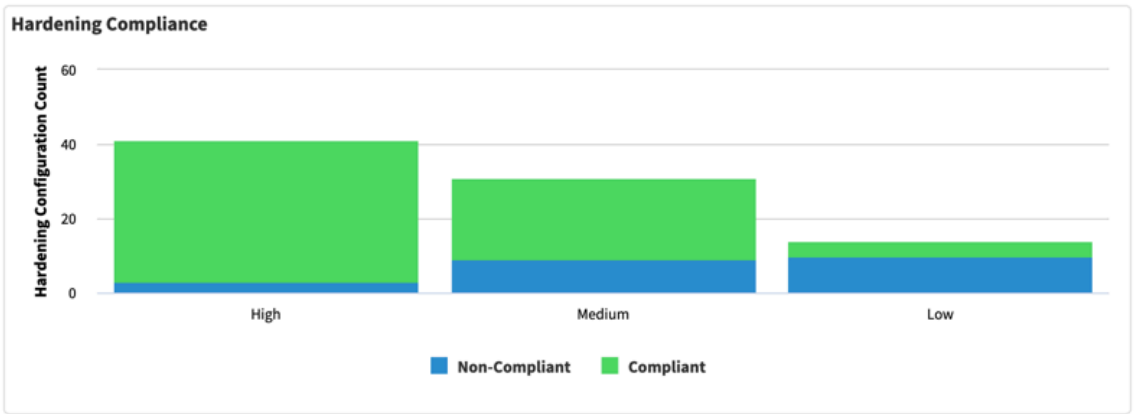

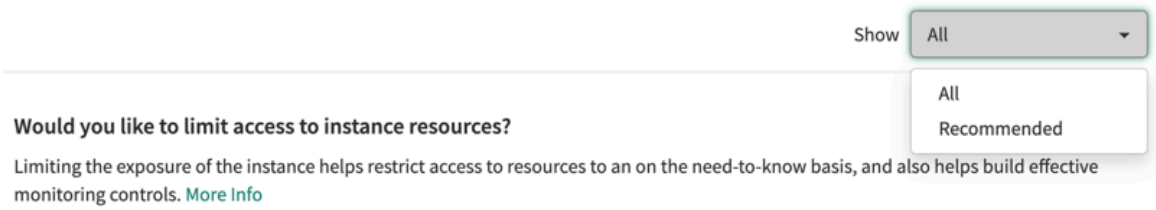

5. In **Categories**, select the category that contains the security configuration properties you would like to access:

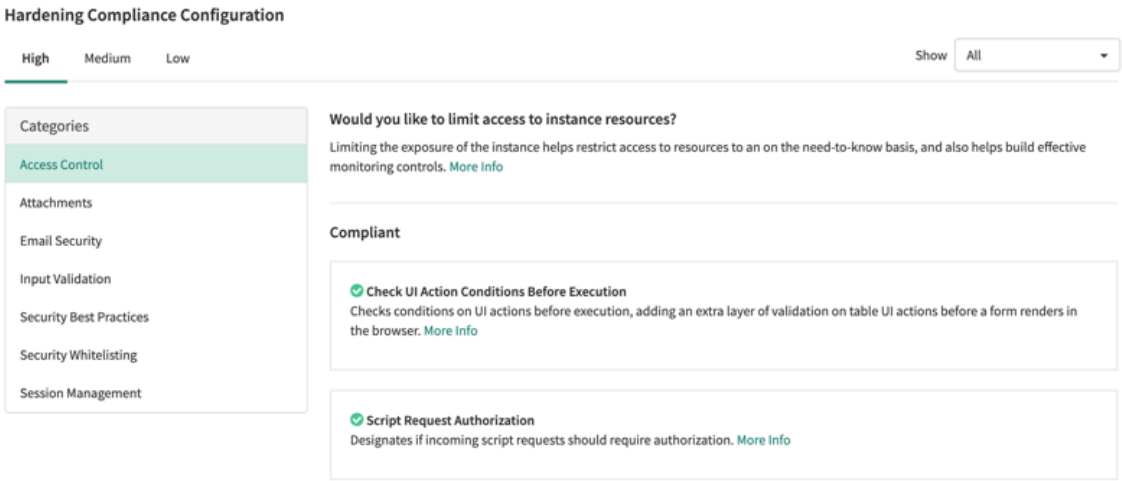

Option	Description
Access Control	Access controls determine whether to grant or deny user access to a particular resource based on who is permitted to use those re sources. To learn more, see Access control in the Instance Security Hardening Settings.
Attachments	Attachment security controls enable valida tion of incoming attachments to protect your instance against malicious files sent by attack ers. To learn more, see Validate file mime type in AttachmentCreator soap web service [New in Security Center 1.3 and updated in 1.5] in the Instance Security Hardening Settings.
Email Security	Email security encompasses security config uration properties an administrator can con figure to ensure that proper security policies are in place for all inbound emails. To learn more, see Enable email spam scoring and fil tering [Updated in Security Center 1.3] in the Instance Security Hardening Settings.
Input Validation	Input validation includes security-related properties that an administrator can config ure to minimize entry of malformed data, re gardless of source. To learn more, see Vali dation, sanitization, and encoding in the In stance Security Hardening Settings.
Secure communications	Secure communications properties are those that an administrator can configure to secure the transportation of HTTP traffic. To learn more, see Communications in the Instance Security Hardening Settings.
Security Best Practices	Security best practices encompass Security Tasks that an administrator should perform pe riodically, within a certain interval of time, and

Option	Description
	include related configuration properties. To learn more, see Security Best Practices in the Instance Security Hardening Settings.
Security Inclusion Listing	Security inclusion listing includes security-related properties that an administrator can configure to restrict behavior to known inclusion listings.
Session Management	Session management includes security-related properties that an administrator can configure to ensure secure session management in the ServiceNow AI Platform. To learn more, see Session management in the Instance Security Hardening Settings

6. Configure the non-compliant security properties in the selected category.
   - Unless otherwise specified, sliding the switch on sets a security property to its recommended setting. For example, you set most controls to true or false, but some require entry of a value, or values, such as a comma-separated value list.
   - To access the dedicated Instance Security Hardening Settings topic for the security control, and learn more about it, click **More Info**.

Result

The Daily Compliance score increases or decreases depending on the changes that you make to the non-compliant security control settings.

Related topics

Instance Security Center

How Daily Compliance score, trend, and graph data is refreshed

Trend and graph data in the Instance Security Center is updated after the performance analytics job executes at 02:00 local time. It appears in the Daily Compliance Score tile, in the Event ribbon tiles, and in the Analytics Hub page detail when you click one of the event tiles.

The [AppSec] Daily Data Management job is a regularly scheduled job that runs nightly and performs the following tasks:

1. Verifies if valid users have been assigned to the [AppSec] Daily Data Management and [PA AppSec] Daily Data Collection jobs when you first schedule them.
   - If you entered a valid user into the **Run As** field, the job continues processing. A valid user is one that is not locked out of the instance and has an assigned admin role.
   - If you entered an invalid user, an error message appears above the rotating security banner in the Instance Security Center.

   Note: To learn more about updating the assigned user when running scheduled jobs, see Create or schedule a data collection job.

2. Executes business logic to set the compliance state for the security properties you configure in the Hardening Compliance Configuration page. To learn more, see Check the daily compliance score and configure security property settings.
3. Runs the [PA AppSec] Daily Data Collection performance analytics job to collect compliance data and update the Daily Compliance Score.

Manually refreshing the Daily Compliance Score

Alternately, if you have an assigned admin role, you can refresh and recalculate the Daily Compliance Score at any time by clicking **Refresh**.

Note: The **Refresh** button does not appear for users with an assigned security_dashboard_user role.

- The Refresh function performs the same tasks as the Daily Data Collection performance analytics job but does it in real time, rather than in a batch process.
- You typically use it when you want to perform updates to the Daily Compliance Score to immediately view the impact of instance security activities.
- There may be a slight delay before the updated score appears.

Note: When you perform an upgrade (for example, from London to Zurich), the Instance Security Center (ISC) plugin is automatically activated. A supplied fix script automatically assigns a custom user without any assigned roles.

PCI compliance score dashboard

The PCI compliance score dashboard shows how your instance conforms to payment card industry (PCI) security standards. Use the dashboard to view your compliance score and modify your configuration to improve security.

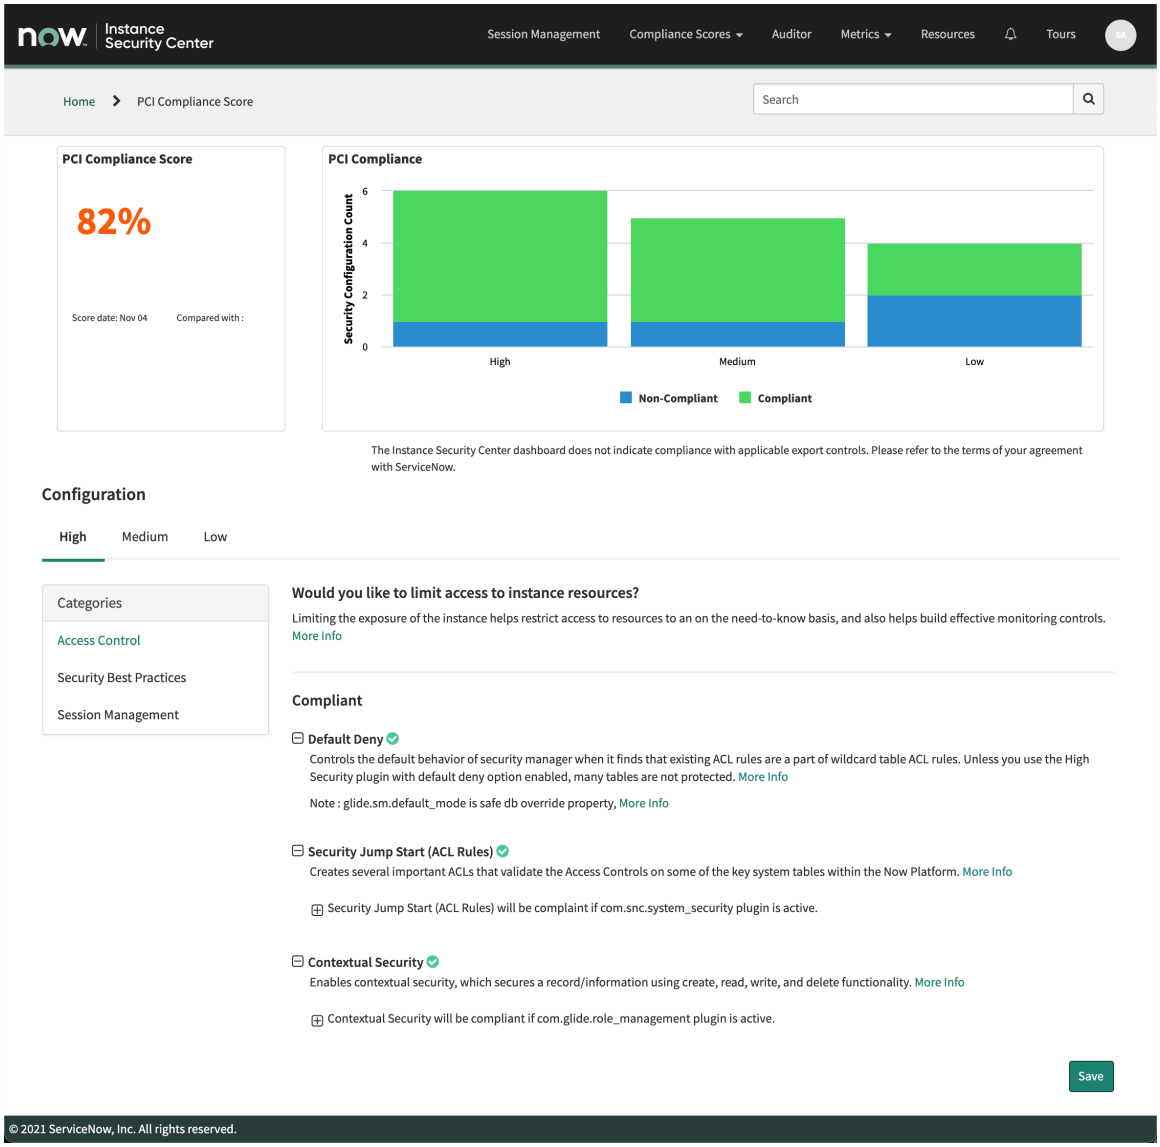

Required ServiceNow AI Platform roles

security_dashboard_user or admin, needed to view the PCI compliance score dashboard.

Access the PCI compliance score dashboard

To open the dashboard, navigate to **System Security** > **Instance Security Center**. From the instance security center, click **Compliance Scores** in the header, then select **PCI Compliance Score**.

Use cases

For examples of how different people in your organization would use this dashboard, see these use cases.

User	Dashboard use
Security dashboard user (security_dashboard_user)	Continually monitor and manage instance security compliance.

User	Dashboard use
Admin (admin)	Continually monitor instance security compliance to detect and respond to security threats.

Indicators

Indicator	Description
PCI Compliance Score	Displays your instance's compliance score as a percentage. This percentage of represents the percentage of security configurations on your instance that meet compliance standards. The indicator also displays the date the compliance score was calculated, and a comparison with the previous calculated score.

Data visualization

Title	Type	Source table	Description
PCI Compliance	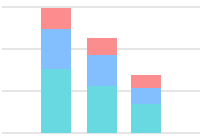 Stacked bar chart	Security Configurations [isc_security_configurations]	Displays compliant and non-compliant security configurations in high, medium, and low categories. Click an area of the report to display the matching security configurations.

PCI configuration controls score dashboard

Use the PCI configuration controls score dashboard to review your PCI configuration and determine which security checks are non-compliant. You can change the configuration of the non-compliant security checks from the instance security center.

Required ServiceNow AI Platform roles

The security_dashboard_user or admin role is required to view the PCI compliance score dashboard.

Access the PCI configuration controls score dashboard

To open the dashboard, navigate to **System Security** > **Instance Security Center**. From the instance security center, click **Compliance Scores** in the header, then select **PCI Configuration Controls Score**.

Use cases

For examples of how different people in your organization would use this dashboard, see these use cases.

User	Dashboard use
Security dashboard user (security_dashboard_user)	Continually monitor and manage PCI configuration controls compliance on your instance. Modify PCI configuration to ensure compliance and improve instance security.
Admin (admin)	Continually monitor PCI configuration controls compliance to detect and respond to potential security threats.

Indicators

PCI Configuration Controls Score

Displays your instance's PCI configuration controls score as a percentage. This percentage of represents the percentage of security PCI control configurations on your instance that meet compliance standards. The indicator also displays the date the compliance score was calculated, and a comparison with the previous calculated score.

Data visualization

Title	Type	Source table	Description
PCI Configuration Controls	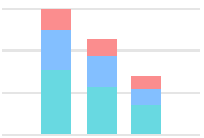 Stacked bar chart	Security Configurations [isc_security_configurations]	Displays compliant and non-compliant PCI control configurations in high, medium, and low categories. Click an area of the report to display the matching security configurations.

Scan for incorrect security definitions

Run the Auditor to scan your instance and find incorrect security definitions. It provides findings you can correct to help improve the security posture of your instance.

Important:

Instance Security Center (ISC) has reached the end of sales as of September 2024, and is no longer supported or available for new activation.

ServiceNow Security Center (SSC) is the recommended solution going forward. For more information, see Instance Security Center to ServiceNow Security Center migration.

The Auditor performs a “full-body” assessment of your instance health that analyzes your system configuration. For security scans, it compares your current security configuration to best practice definitions, and to security property compliance values.

From an instance security standpoint, it provides insights and recommendations into what you should continue doing, and where you might be able to improve. These insights and recommendations help you answer the following questions:

- Are the appropriate security-related properties set?
- Is the High Security plugin enabled?
- Do the right access control rules exist?

Run the Auditor and analyze scan results

1. To run the Auditor, click **Audit** on the Instance Security Center home page.
2. When it completes, open Scan Results to review and analyze the security findings.
3. To review the detail for a specific scan result, double-click the result number. This information includes its status, scan type, execution time, and error messages.
4. Each of the Auditor findings contains resolution details, and a URL to product content about how to address them. Follow the documented guidelines to resolve the issues in each of the findings.

Related topics

Hardening settings

Enable High Security Plugin [Updated in Security Center 1.3]

Access Control List Rules

Monitor instance metrics

Monitor user, export, authentication, email, and antivirus metrics for your instance. For example, you can monitor your email security by checking metrics for spam, external emails, and inbound emails from untrusted and trusted domains for your instance. Analyze these metrics to look for anomalous security behaviors that are related to activities that take place in your instance.

Important:

Instance Security Center (ISC) has reached the end of sales as of September 2024, and is no longer supported or available for new activation.

ServiceNow Security Center (SSC) is the recommended solution going forward. For more information, see Instance Security Center to ServiceNow Security Center migration.

User metrics

Analyze user metrics to look for anomalous behaviors that are related to specific types of user activity in your instance.

Not Logged in Last Month / Last Six Months / Last Year

Indicates the number of users who have not logged into the instance within the last month, within the last six months, and within the last calendar year. To view user detail for a specific metric:

- Click the metric to view a listing of users that have not logged in to the instance during the indicated time period.
- Click a user name to view more details about that user.

Users with High Privilege Roles

Indicates the number of users with the following high privilege role types:

User role	Description
admin	Primary administrator role that has access to all system features, functions, and data, regardless of security constraints.
ais_high_security_admin	Elevated privilege role that enables a user to access High Security settings for AI Search. To learn more, see Assign roles to AI Search administrators and users.
password_reset_admin	Administrator role that enables a user to view the status of password reset activities, identify potential security threats, and monitor for compliance with password security policies. To learn more, see Password Reset and Password Change reports and logs.
script_include_admin	Administrator role that also has access to script includes.
security_admin	Elevated privilege role that enables a user to create and change access controls and High Security Settings. To learn more, see Security_admin role
user_admin	Administrator role that can also manage users, roles, user groups, roles, and department assignments.

Note: To learn more about these administrative role types, see Special administrative roles.

To view user detail for a specific user role metric:

- Click the user count role metric to view a listing of users with that high privilege role type.
- Click a user name to view more details about that user. You can then determine if these security-critical roles are assigned to the proper personnel.

Users Trend

Shows count trend information over a time period for the following types of users:

Count type	Description
Active Users	Number of users who are marked as Active in the instance.
Inactive Users	Number of users who are marked as Inactive in the instance.
Locked Out	Number of users who are locked out of the instance.

To view user detail for a specific user count (for example, Locked Out Users):

- Click the **Locked out users** metric.
- In the Analytics Hub page, click **Show Records**.
- Click a user name to view more details about that user. You can then determine if there is a reason this person is locked out and remedy the situation.

Events Trend

Shows count trend information for specific types of events, over a time period:

Event type	Description
Admin login	Number of users with high privilege administrator user roles who logged in on a specific day.
External login	Number of users with an assigned snc_external role who logged into this instance during the calendar day. These logins typically occur for maintenance, support, consulting, or audit purposes. Monitoring this metric enables you to verify that the external login attempts are legitimate and not potential security issues.
Failed login	Number of failed login attempts on a specific day.
Impersonation	Number of users logged in on a specific day who are impersonating other users.
Security elevation	Number of times that a security administrator has elevated security for standard users by changing their assigned user role to a high privilege security role during the calendar day. These high privilege security roles include oauth_admin, admin, security_admin, and impersonator.
SNC login	Number of Customer Service and Support who logged in to this instance using the hi-hopping technique during a specific day.

To view user detail for a specific event count (for example, Impersonation):

- Click the user count metric. The Security Dashboard Event Logs page lists event logs for that type of event.
- Click a user name to view more details about that event.

Related topics

Analytics Hub

Export metrics

Analyze export metrics to see what data is most commonly exported and which users export the most data.

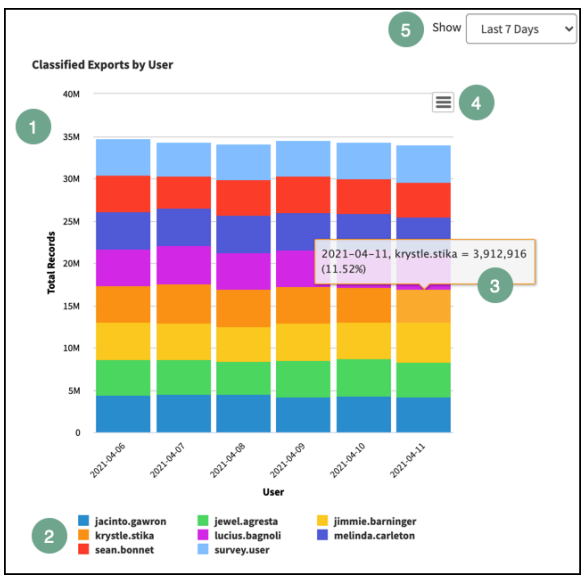

1. Export chart

Each report displays the number of export events by date, using a color coded key to indicate which users performed the export. Click a colored section of a column to see the list of **Export Events** [isc_export_event] records matching that entry.

2. Report key

The key at the bottom of the report indicates which colors identify which users or tables.

3. Preview pop-up

Point to an entry in the chart to see a pop-up preview. This

preview shows the user or table name, as well as a count of exports and a percentage of the total on that column.

4. **Image export**

   Click the icon to save the report as an image. You can save the report in PNG or JPEG format.

5. **Report date range**

   Use the **Show** list to display exports within the last 24 hours or within the last 7 days.

Export metrics reports

The export metrics page displays four reports.

Exports by user

Use the **Exports by User** report to see which of your users are exporting the most data.

Classified exports by user

Use the **Classified Exports by User** report to see which of your users are exporting the most data matching classifications like confidential, restricted, or personal information. Administrators can define which classifications this report uses in the **Settings** tab.

Exports by table

Use the **Exports by Table** report to see which tables are exported from most frequently.

Classified exports by table

Use the **Classified Exports by Table** report to see which tables are exported from most frequently that match classifications like confidential, restricted, or personal information. Administrators can define which classifications this report uses in the **Settings** tab.

Note: Export metric reports only track export events. Exports from other sources, such as rest APIs or workflows are not tracked as part of this feature.

Export metrics settings

Use the configuration options in the Settings tab to narrow down reporting results.

Access the settings for your export metrics by selecting the **Settings** tab.

Settings configuration fields

Export metrics configurations

Configuration	Description
Classifications for Metrics	Add or remove classifications to this field to determine which exports are included in the **Classified Exports by User** and **Classified Exports by Table** reports. These reports support the following classifications:

Export metrics configurations (continued)

Configuration	Description
	• Personally identifiable information • Confidential • Restricted • Internal • Public For more detail on data classifications, see Data classifications
Classifications for Alerts	Add or remove classifications to this field to determine which exports trigger instance security notifications. The classifications supported in the **Classifications for Metrics** field are supported here. For more detail on these alerts, see the security notifications section on the Instance Security Center page. The **Record Threshold** field defines the number of records exported before your instance triggers and alert.
Record Threshold	Number of record a user must export to trigger an alert. To trigger an alert, these records must also match the classifications listed in the **Classifications for Alerts** field.

Save your settings by entering **Control + s** (Windows) or ⌘ **+ s** (macOS).

Authentication Metrics

Analyze authentication metrics to see information related to authentication, such as infrequently used IP addresses, failed logins, and types of authentication schemes used by your users.

Use the **Authentication Metrics** page to view reports relating to your authentication configuration. The following reports are displayed in this tab.

Note: The authentication metrics page requires the **REST API Access Policy plugin** (com.glide.rest.policy) plugin. For more details about this feature, see REST API access policies.

APIs without authentication policies

Displays a real-time count of all API's without an access policy

Hardening: Account recovery flow

Hardening: Role-based MFA feature related setting

Web service access only accounts

Shows the count of all user records with web service access option enabled in *User* [sys_user] table.

X509 certificates expiring in 30 days

Displays a count of all X.509 certificates from *X.509 Certificates* [sys_certificate] table which are due expire in 30 days.

Adaptive authentication metrics

Analyze adaptive authentication metrics to monitor and add insights on how adaptive authentication is being used on your instance.

View reports, settings, and policies associated with adaptive authentication in on place using the adaptive authentication metrics page. Security administrators can use reports to monitor the results of their adaptive authentication policies. Use this data to gain insights and adapt your policies to improve their performance.

Note: The adaptive authentication metrics page requires the **Adaptive Authentication** (com.snc.adaptive_authentication) plugin. For more details about this feature, see Adaptive authentication.

Metrics

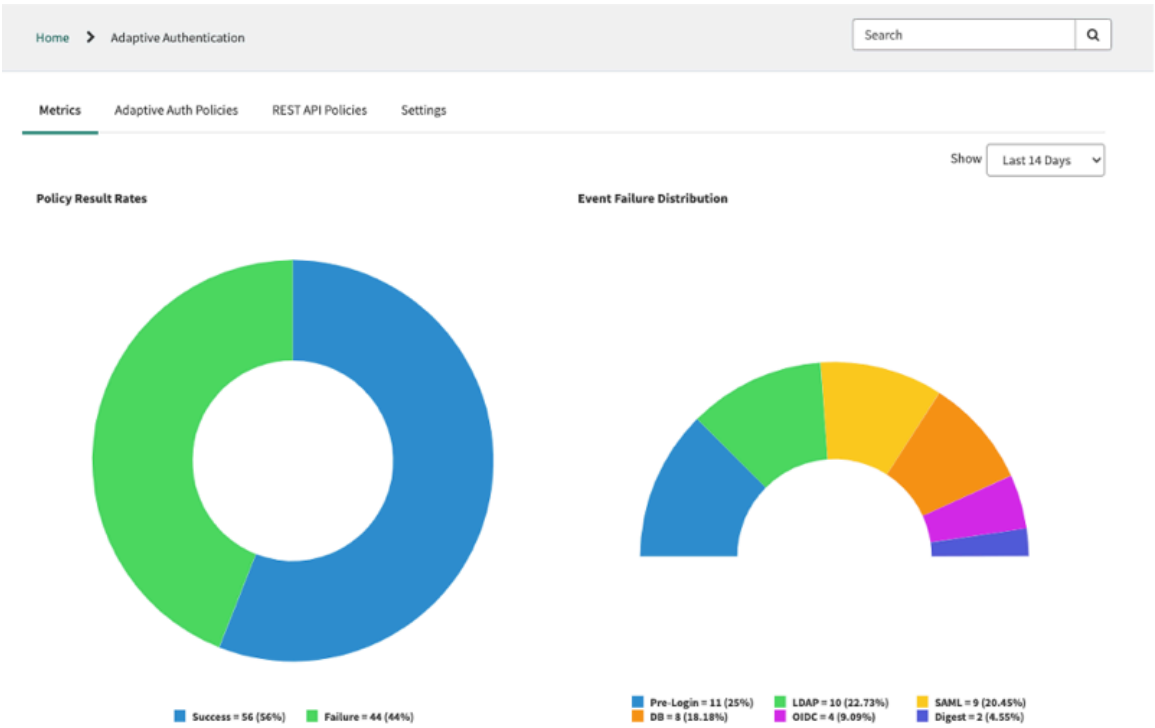

Use the **Metrics** tab to view reports relating to your adaptive authentication configuration. The following reports are displayed in this tab.

- Policy Results Rates
- Event Failure Distribution
- Event Success Distribution
- Denied IP Addresses
- Authentication User Logins
- API User Logins
- Authentication Trend

Use the **Show** list to select a time span for the displayed reports

Adaptive Auth Policies

Use the **Adaptive Auth Policies** tab to view the adaptive authentication policies and policy contexts on your instance. Click on any entry on these lists to view the associated record. For more details on these records, see Adaptive authentication

Settings

Use the settings tab to view and configure adaptive authentication system properties. For more information on these properties, see Configure adaptive authentication properties

Email metrics

Analyze your email metrics to look for anomalous behaviors that are related to the incoming emails to your instance. For example, if the metrics indicate a spike in spam emails from specific domains, you can define inbound actions that prevent their delivery to the instance.

For each email metric, a count appears for each type of email that is delivered or sent to the instance during the calendar day.

Notification preference	Description
External Incoming Email	Number of incoming emails for the calendar day that were delivered to the instance from external email domains. **Note:** The external email domains are those domains that are not listed in the *security.list.internal.domains* system property, because this property tracks only your internal email domains. To learn more about this property, see Available system properties.
Spam	Number of incoming emails for the calendar day that were delivered to the instance and marked as spam. A number that is out of line with historical trends may indicate that attempts are being made to compromise your instance security.
Trusted Incoming Email	Number of incoming emails to the instance for the calendar day from email domains designated as trusted.
Untrusted Incoming Email	Number of incoming emails to the instance for the calendar day from email domains designated as untrusted. You can designate untrusted or trusted email domains in the Untrusted And Trusted Domain form so that you can track your inbound emails that are sent from them. To learn how to designate untrusted or trusted email domains, see Designate email domains as untrusted or trusted.

After you click an email metric, you can learn about the possible email security issues in your instance by clicking one of the following:

Command	Description
Chart	Incoming email counts and count trends over time for the selected email type (spam, incoming external, untrusted, or trusted).
Records	Individual email records that compromise the daily count for the selected email type.
More Info	Additional information for the selected email type.

Note: The email metrics apply only to your incoming emails to the instance. The metrics do not apply to the normal traffic that is processed through your enterprise-wide email servers. To learn about defining inbound actions and how they impact the processing of your inbound emails, see Inbound email actions.

Designate email domains as untrusted or trusted

Designate specific email domains as untrusted or trusted so that you can monitor the metrics for incoming emails from these sources in your instance.

Before you begin

Role required: security_dashboard_user or admin

About this task

Important: Instance Security Center (ISC) reached end of sales in September 2024 and is a legacy product. ServiceNow Security Center is the recommended replacement. For more information, see Migrating to Security Center.

This procedure applies to instances still running ISC. If you have migrated to ServiceNow Security Center, email monitoring is available by navigating to **Security Center** > **Metrics** > **Email**. The trusted and untrusted email domain designation feature described in this topic is not available in Security Center.

When untrusted or trusted domains send emails to your instance, their daily counts appear on the **Untrusted Incoming Email** or **Trusted Incoming Email** metrics on the Email page. You can then track email activity from these domains and use email logs to view specific incoming emails. You can also specify a user, usually a manager, or a security analyst, to notify whenever activity occurs from the untrusted or trusted domain.

Note: Designating an email domain as untrusted is for security tracking purposes only. Administrators can also set up a system address filter to ignore emails from untrusted domains. To learn about filtering emails to block their delivery, see System address filters.

Procedure

1. Navigate to **All** > **System Security** > **Instance Security Center**.
2. On the Instance Security Center homepage, select **Email** from the **Metrics** menu.

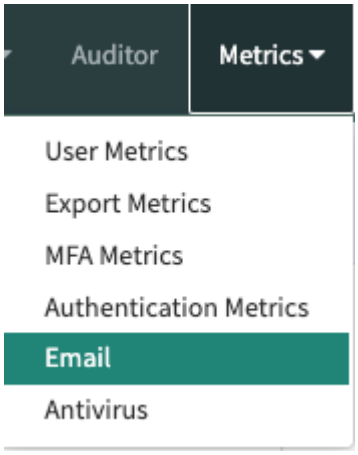

3. On the Email page, in the Untrusted And Trusted Domains section, click **New**.
4. On the form, fill in the fields.

Untrusted And Trusted Domain form

Field	Description
Domain	Name of the email domain that you are designating as untrusted or trusted. For example, enter `servicenow.com` to designate ServiceNow employees can send trusted emails to the instance.
Category	Category that indicates if the email domain is untrusted or trusted: **Untrusted** Designates that the email domain as untrusted. You use it to identify domains that send suspicious or emails that pose a potential security threat to the instance. **Trusted** Designates that the email domain as trusted. You use it to identify domains when your metrics indicate that the incoming emails from it pose no security threats. Designating the domain as trusted enables you to track its inbound email activity over time.
Active	Check box for enabling or disabling the designated untrusted or trusted status for the specified email domain.
Notify	Name of the user to notify by email when activity occurs in the untrusted or trusted domain. Click the spotlight search icon ( ) to search for the name of the user. Leave the **Notify** field blank if you do not want notifications sent.

5. Select **Save**.

Result

Untrusted or trusted email domain information is also added to the **Untrusted And Trusted Domains** listing on the Email page.

Related topics

Instance Security Center

Email metrics

Antivirus metrics

If the Antivirus Scanning plugin is activated, Antivirus Scanning runs in your instance to help protect it against virus infections from attachments.

The following metrics appear for the last 60 days of activity, and enable you to assess the effectiveness of the Antivirus Scanning functions.

Antivirus Events

Antivirus Events indicate the number of antivirus events in your instance, by date. To access the antivirus events, navigate to **System Security** > **Instance Security Center** and select the Metrics tab. Color coded graph lines represent the following types of antivirus events:

Color	Description
Blue	Number of files quarantined by Antivirus Scanning in this instance for the indicated date.
Green	Number of infected files downloaded to the instance, and then quarantined for the indicated date. These files are primarily email attachments that contain a virus or rouge code.
Yellow	Number of quarantined files in the instance that were deleted for the indicated date.
Orange	Number of quarantined files in the instance that were restored for the indicated date. **Note:** After Antivirus Scanning runs and finds any false positives, you can restore a quarantined file and make it accessible in the instance.

- To access the Analytics Hub page and view the detailed score card and analytics information for a specific date, click a colored line in the Antivirus Events graph. For example, click the blue graphics line to view analytics information for files quarantined for a specific date.
- To view the following breakdowns in the Analytics Hub page, click , then click:

Breakdown	Description
AppSec - Antivirus Event Source	Source of the antivirus event. ○ **On Upload:** Occurred due to an upload of an infected file, usually an attachment. ○ **From Quarantine:** Occurred due to the quarantine of an infected file, usually an attachment. ○ **On Download:** Occurred due to a download of an infected file, usually an attachment. ○ **From Record:** Occurred due to an infected record in a table.
AppSec - Antivirus Event Type	Type of antivirus event.

Breakdown	Description
	○ **Quarantined:** Occurred due to the quarantine of a file, usually an attachment. ○ **Downloaded:** Occurred due to a download of a file, usually an attachment. ○ **Restored:** Occurred due to the restoration of a quarantined file. ○ **Deleted:** Occurred due to the deletion of a quarantined file.
AppSec - Antivirus Uploader	Name of the logged in user who uploaded the files that were the source of virus infections detected by the Antivirus Scanning application.

Quarantined Files

Lists the infected files in the instance quarantined by Antivirus Scanning:

Field	Description
File name	Name of the infected file.
Content type	Type of content that was infected in the file. For example, **application/x-dosexec** is an infected application or DOS executable file, while **text/plain** is an infected .txt file.
Table	Name of the table that contains the infected file. For example, **incident** appears for an incident file record.
Virus	Name of the file quarantined by Antivirus Scanning.
Detected	Date and time the infected file was detected.
Created By	Name of the user who quarantined the infected file.
Created	Date and time the quarantine file record was created.
Table sys ID	Table system identifier assigned to the quarantine file record.

Note: You can also add **Quarantined Files** and **Virus Types** tiles to the Event ribbon. To learn more, see Monitor security events and Configure the security event ribbon.

Related topics

- Antivirus Scanning
- Configuring Antivirus Scanning
- Reviewing quarantined files
- Review antivirus activity
- Analytics Hub
- Performance Analytics breakdowns
- Analytics, Intelligence, and Reporting

MFA metrics dashboard

The MFA metrics dashboard shows information on your instances multi-factor authentication configuration. Use the dashboard to ensure your MFA configuration meets your security standards.

Required ServiceNow AI Platform roles

security_dashboard_user or admin, needed to view the PCI compliance score dashboard.

Access the MFA metrics dashboard

To open the dashboard, navigate to **System Security** > **Instance Security Center**. From the instance security center, click **Metrics** in the header, then select **MFA Metrics**.

Use cases

For examples of how different people in your organization would use this dashboard, see these use cases.

User	Dashboard use
Security dashboard user (security_dashboard_user)	Continually monitor and manage instance security compliance.
Admin (admin)	Continually monitor instance security compliance to detect and respond to security threats.

Indicators

Users Enrolled for MFA

Displays the total number of users on the instance enrolled in MFA . Click to open the Analytics hub for more detail.

Users using MFA Bypass

Displays the total number of users using MFA bypass . Click to open the Analytics hub for more detail.

Data visualizations

Title	Type	Source table	Description
High Privilege MFA Users	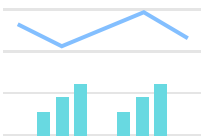 Bar		
MFA User Trend	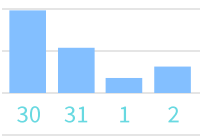 Trend		

Activate the ISC Virtual Agent interface

If you have the admin role, you can activate the ISC Virtual Agent Conversations plugin (com.glide.isc_virtualagent). Activating this plugin installs the Virtual Agent and Natural Language Understanding (NLU content packs, providing Virtual Agent access from the Instance Security Center.

Before you begin

Important:

Instance Security Center (ISC) has reached the end of sales as of September 2024, and is no longer supported or available for new activation.

ServiceNow Security Center (SSC) is the recommended solution going forward. For more information, see Instance Security Center to ServiceNow Security Center migration.

The Virtual Agent interface in the Instance Security Center is only available for users with paid Virtual Agent subscriptions, and who have activated the Glide Virtual Agent (com.glide.cs.chatbot) plugin. To learn more, see Activate Virtual Agent.

Role required: admin.

About this task

The ISC Virtual Agent interface enables you to perform the following tasks:

- Ask security-related questions, then get quick summary answers and reference links to learn more.
- Get answers related to the Instance Security Center, Platform Security, ServiceNow security policies, trust, governance, risk, and other processes.
- Search for security-related resources, such as Knowledge Base topics in the Now Support Security Portal.

Plugins for ISC Virtual Agent

Plugin	Description
ISC Virtual Agent Conversations [com.glide.isc_virtualagent]	Activates the ISC Virtual Agent Conversations content pack for the Instance Security Center.
ISC NLU Model for Virtual Agent Conversations [com.glide.isc_nlu]	Activates the Natural Language Understanding (NLU) content pack for the Instance Security Center.

Note: Activating the com.glide.isc_virtualagent plugin automatically activates com.glide.isc_nlu. However, if you activate the com.glide.isc_nlu plugin first, you must manually activate com.glide.isc_virtualagent as well.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the plugin using the filter criteria and search bar.

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

3. Select **Install** to start the installation process.

> **Note:** When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Hardening settings

The ServiceNow Security Center (SSC) hardening settings content contains detailed descriptions and compliance values for the security-related system properties and plugins in the ServiceNow AI Platform. You can set these properties using the hardening settings app in the Security Center.

Overview and purpose

The Security Center calculates a daily compliance score, expressed as a percentage that is based on how compliant your current instance security settings are with the compliance values in Security Center hardening settings.

You can manage the specific security configuration settings that may affect the score for your instance directly from the Security Center.

The hardening settings configurations are explained with several attributes described in the table.

Hardening settings configuration details

Configuration attribute	Description
Overview	Provides a high level overview of the recommendation.
Configuration name	The property or plugin name.
Configuration type	Describes where the property can be configured outside of the Security Center, such as in system properties (*sys_properties_list.do*).
Data type	Describes the type of value required for the configuration. Examples are true/false boolean, installation, plugin, string, etc.
Recommended value	The value that is recommended by the Security Center to enhance security compliance in your instance.
Default value	The value that the configuration is set to in the base system.
Category	The name and link to the category for the hardening setting.
Security risk	**Severity score**: The score indicates the potential security risk to your instance as per the likelihood of the vulnerability to be exploited. The security vulnerability is considered and scored individually using the CVSS (Common Vulnerability Scoring System) score on a scale ranging from **0.0** to **10.0**. See https://nvd.nist.gov/vuln-metrics/cvss/v3-calculator for additional information.
	Severity rating per CVSS score:

Hardening settings configuration details (continued)

Configuration attribute	Description
	• Critical: 9.0-10.0 • High: 7.0-8.9 • Medium: 4.0-6.9 • Low: .01-3.9 • None: 0.0
	Security risk details: Describes the importance of the setting configuration and the risk of not utilizing the recommended configuration.
Dependencies and prerequisites	Related settings or configurations that are required before or in conjunction with the hardening configuration.
Functional impact	The impact this hardening setting has on the operation of your instance.
References	Links to configuration documentation or other helpful information.

Note: Some of the configurations can only be completed by Customer Service and Support and will be indicated as such.

To learn more about ensuring your instances meet hardening requirements, see Security hardening.

Other resources

For user reference, the ServiceNow AI Platform maintains extensive configuration capabilities information in the product documentation. You access most of the security content using the links found in Secure your instance. Also, see the following:

- Available system properties
- General security settings properties
- High Security Settings

Hardening settings baseline versions

Explore how baseline versions for hardening settings align with family and store releases.

Security Center works by ingesting a subset of system properties from an instance and displaying its configuration details, along with the security impact of non-compliance within the app. The baseline serves as the reference point for the system properties that are ingested with each release of the Security Center app.

Hardening settings baseline overview

Security Center version	Hardening settings baseline version	Supported families	Store release date	Installed by default with
SSC v1.1	Baseline v1.0	Utah, Vancouver	May 2023	Vancouver family
SSC v1.2	Baseline v1.0	Utah, Vancouver	Aug 2023	Store only
SSC v1.3	Baseline v2.0	Vancouver, Washington DC	Nov 2023	Washington family
SSC v1.5	Baseline v4.0	Washington DC, Xanadu	May 2024	Xanadu
SCC v1.6	Baseline v4.0	Washington DC, Xanadu	August 2024	Store only
SCC v 2.0	Baseline v5.0	Xanadu, Yokohama	November 2024	Yokohama
SCC v 2.2	Baseline v6.0	Xanadu, Yokohama	May 2025	Store Only
SCC v 3.2	Baseline v7.0	Zurich	August 2025	Zurich

New hardening settings

New hardening settings are broken out by baseline release.

New hardening settings for baseline 7.0

New hardening settings have been released with Security Center baseline version 7.0.

- Enable Cross Scope Privilege Checks on Service Portal Form [New in Security Center 7.0]
- Validate query ACLs on Glide DB functions [New in Security Center 7.0]
- Use Document Classification to limit publicly accessible documents [New in Security Center 7.0]
- Restrict write access on system fields to admin users [New in Security Center 7.0]
- Require approval for agent-based Office 365 group membership changes [New in Security Center 7.0]
- Exclude Sensitive Tables and Fields from Data Generation [New in Security Center 7.0]
- Enforce Read Roles for Catalog Variable Search [New in Security Center 7.0]
- Enforce valid query string choice [New in Security Center 7.0]
- Restricted Binding functionality in case Bearer Authorization [New in Security Center 7.0]
- Disable resource owner password credentials (ROPC) in OAuth 2 token grants [New in Security Center 7.0]
- Enforce certificate trust [Updated in Security Center 1.3, removed in 2.0, added in 7.0]
- Prevent usage of 3DES keys [New in Security Center 7.0]
- Allow HTML Links to Trusted Domains in the Description Fields of the Impact Workspace Module [New in Security Center 7.0]
- Ensure Contextual Search Do Not Contain An Unvalidated Redirect [New in Security Center 7.0]
- Sanitize HTML in the Description Fields of the Impact Workspace Module [New in Security Center 7.0]

New hardening settings for baseline version 6.0

New hardening settings have been released with Security Center baseline version 6.0.

- Enforce application specific ACLs only for application data
- Disable legacy JQuery UI usage
- Display recommendations for high risk UI pages
- Enforce current password policy compliance requirements on login
- Prevent Users From Accepting Warning To Bypass CSRF Validation [Updated in Security Center 1.3 and 1.5]
- Set minimal password length [Updated in Security center 2.2]
- Limit session length for high assurance sessions
- Disable deprecated TLS versions
- Disable local login for users with Single Sign-On (SSO) enabled
- Reduce allowed bypasses for multifactor setup
- Prevent impersonating user from viewing application data
- Prevent verbose HTTP request logging
- Enable relay state in SAML requests to prevent replay attacks
- Minimize failed login attempts for high assurance sessions
- Apply continuous authentication policies to mobile sessions
- Disable use of TripleDES/3DES encryption algorithm

New hardening settings for baseline version 5.0

New hardening settings have been released with Security Center baseline version 5.0.

- Enforce ACL on HR Lifecycle Events Data [New in Security Center 2.0]
- Sanitize All Translated HTML Fields [New in Security Center 2.0]
- Configure Service Portal Widgets Allow List [New in Security Center 2.0]
- Enforce ACL on HR Core Data [New in Security Center 2.0]
- Enforce ACL on HR Virtual Agent Data [New in Security Center 2.0]
- Configure Service Portal Widgets Table Allow List [New in Security Center 2.0]
- Enforce Security Scope for Service Application Information [New in Security Center 2.0]
- Prevent Empty ACL Creation [New in Security Center 2.0]
- Prevent Unauthenticated Access to Virtual Agent Embedded Web Client
- Set Automatic Token Cleanup for Token Credentials [New in Security Center 2.0]
- Restrict Global App Development by Role [New in Security Center 2.0]
- Enable ACLs for Encoded Query in Simple List Widget [New in Security Center 2.0]
- Invalidate Session After OAuth Token Expiration [New in Security Center 2.0]
- Set Allowed MIME Child Types [New in Security Center 2.0]
- Restrict Impersonation to Admin [New in Security Center 2.0]

New hardening settings for baseline version 4.0

New hardening settings have been released with Security Center baseline version 4.0.

- Review extraneous explicit role access control conditions [Removed in Security Center 1.5]
- Enable work order management query rules for service organizations [New in Security Center 1.5 and updated in 2.0]
- Restrict flow context read access [New in Security Center 1.5]
- Limit policy based session access mobile refresh token interval [New in Security Center 1.5]
- Enable policy based session access for mobile [New in Security Center 1.5]
- Prevent inactive users from logging in [New in Security Center 1.5]
- Configure event management assignment group admin roles [New in Security Center 1.5]
- Enforce Security Scope for Agent Workspace for HR Case Management [New in Security Center 1.5 and updated in 2.0]
- Enforce security scope license and permit playbook [New in Security Center 1.5 and updated in 2.0]

New hardening settings for baseline version 2.0

New hardening settings have been released with Security Center baseline version 2.0.

- Ensure archive table ACLs are checked [New in Security Center 1.3 and updated in 1.5]
- Enforce application scope restrictions [New in Security Center 1.3 and removed in 1.5]
- Enable the hardened java security manager [New in Security Center 1.3]
- Verify certificate revocation [New in Security Center 1.3]
- Require clearing pasteboard when backgrounding mobile application [New in Security Center 1.3 and updated in 1.5]
- Enable protected tables plugin [New in Security Center 1.3]
- Enforce strict elevate privilege [New in Security Center 1.3]
- Limit integrations' active session life span [New in Security Center 1.3]
- Proactively invalidate inactive sessions [New in Security Center 1.3 and updated in 1.5 and 2.0]
- Enable MID audit log [New in Security Center 1.3 and updated in 1.5]
- Use of secure insert multiple operation within import set API [New in Security Center 1.3]
- Enforce OCSP check on network error [New in Security Center 1.3 and updated in 2.0]
- Enforce security rules to sharing dashboards [New in Security Center 1.3]
- Restrict oauth parameters to POST body [New in Security Center 1.3]
- Limit attachment size in training and prediction flows for GraphQL endpoints [New in Security Center 1.3 and updated in 1.5]
- Disable GlideRecord Scope Fencing Legacy Behavior [New in Security Center 1.3 and updated in 1.5 and 2.0]
-
- Required jms connection factories [New in Security Center 1.3 and updated in 1.5 and 2.0]
- Limit attachment size in training and prediction flows [New in Security Center 1.3 and updated in 1.5]
- Log session audit events [New in Security Center 1.3 and updated in 1.5]
- Require write access to access service catalog add item page [New in Security Center 1.3]
- Define active session timeout exception roles [New in Security Center 1.3]

- Certificate based authentication not enforced [New in Security Center 1.3]
- Enforce scoped ACL access for information request playbooks [New in Security Center 1.3 and updated in 1.5]
- Hide user comments on articles [New in Security Center 1.3]
- Ensure dashboards creation/deletion requires access check [New in Security Center 1.3 and updated in 2.0]
- Enforce device encryption and passcode requirements [New in Security Center 1.3]
- Validate file mime type in AttachmentCreator soap web service [New in Security Center 1.3 and updated in 1.5]
- Verify certificate revocation [New in Security Center 1.3]
- Check impersonation on ACL evaluation in HR App [New in Security Center 1.3 and updated in 1.5]
- Require captcha for guest walk-up experience in customer service application [New in Security Center 1.3 and updated in 1.5]
- Require Authentication on Event Management HTTP Processor [New in Security Center 1.3, Updated in 1.5, and removed in 2.0]
- Limit guest's active session life span [New in Security Center 1.3]
- Disallow target cloning [New in Security Center 1.3]
- Set safe content security policy for svg files [New in Security Center 1.3]
- Anti-CSRF token validation time [New in Security Center 1.3]
- Restrict knowledge bases access [New in Security Center 1.3]
- Enforce scope security for public sector digital services [New in Security Center 1.3]
- Restrict HR case updates from personal emails [New in Security Center 1.3 and updated in 1.5]
- Limit UI active session life span [New in Security Center 1.3]
- Enforce secure referrer policy [New in Security Center 1.3]

Updated hardening settings

Updated hardening settings are broken out for each baseline release.

Updated hardening settings for baseline version 7.0

Some hardening settings have been updated with the release of Security Center baseline version 7.0.

Documentation	Updates
Enable relay state in SAML requests to prevent replay attacks	• Description: ○ (Old) When "glide.authenticate.sso.saml2.enable_relay_state_with_id" is set to "true", the relay state parameter will contain the sys id of a record in the table multisso_request_parameter which the relay state url to redirect to. This relay state protects against SAML Replay attacks which could be possible in some ServiceNow

Documentation	Updates
	instance versions if the property glide.authenticate.sso.saml2.enable_relay_state_with_id is set to false. Replay attacks allow a well-positioned attacker who has gained access to the SAML request to resubmit a valid request in orger to gain unauthorized access to the platform. ○ (New) Protect against SAML replay attacks using the glide.authenticate.sso.saml2.enable_relay_state_with_id system property. When this property is set to true, the relay state parameter contains the sys_id of a record in the MultiSSO Request Parameters [multisso_request_parameter] table, which the relay state URL redirects to. Set the system property glide.authenticate.sso.saml2.enable_relay_state_with_id to true. This helps prevent attackers who have gained access to a SAML request from accessing your instance by resubmitting a valid request. The relay state enabled by this system property helps protect your instance against replay attacks. Enabling the property helps prevent attackers who have gained access to a SAML request from accessing your instance by resubmitting a valid request. • Remediation: ○ (Old) Set the property glide.authenticate.sso.saml2.enable_relay_state_with_id to true. ○ (New) Set the property glide.authenticate.sso.saml2.enable_relay_state_with_id to true. If the property does not exist in the System Properties [sys_properties] table, the default value is false. • security risk: ○ (Old) This relay state protects against SAML Replay attacks which could be possible in some ServiceNow instance versions if the property glide.authenticate.sso.saml2.enable_relay_state_with_id is set to false. Replay attacks allow a well-positioned attacker who has gained access to the SAML request to resubmit a valid request in orger to gain unauthorized access to the platform.

Documentation	Updates
	○ (New) The relay state enabled by this system property helps protect your instance against replay attacks. Enabling the property helps prevent attackers who have gained access to a SAML request from accessing your instance by resubmitting a valid request. • Functional Impact: ○ (Old) When this property is set to true, relay state in the SAML request will contain sys id of the record in the table multisso_request_parameter which contains relay state url to redirect to. ○ (New) When this property is set to true, the relay state in a SAML request contains the sys_id of a record in the MultiSSO Request Parameters [multisso_request_parameter] table, which contains relay state URL to redirect to.
Disable local login for users with Single Sign-On (SSO) enabled	• Technical Configuration Name ○ (Old) <blank> ○ (New) glide.sso.acr.enabled,glide.authenticate.multisso.enabled • Description ○ (Old) ServiceNow instance owners are responsible for provisioning user accounts for their instance and ensuring that users can access the instance in expected ways. When a user, that is not locked out, has a valid password hash on the "sys_user.user_password" field, that user can perform local database authentication for both interactive and non-interactive access types. If a user is also configured to use SSO authentication, local database authentication can still occur without additional controls in place. That is, SSO users with valid local credentials may access the instance, or parts of the instance, with those local credentials. When SSO authentication is enabled for a user, it is best practice to prevent that user from logging in locally. This reduces the chance that the valid local

Documentation	Updates
	login credentials are stolen and used to login by a malicious user. ○ (New) Users configured to use SSO authentication may be able to access the instance, or parts of the instance, with local credentials stored in the user_password field of their User [sys_user] record. This access applies to both interactive and non-interactive access for users who aren't locked out. Help prevent SSO-configured users from using local credentials to reduce the chance that valid local login credentials are stolen and used by malicious users. Review Now Support Knowledge Base article KB1649420 for instructions on identifying and addressing accounts with local login still enabled on an instance with SSO enabled. • CVSS score ○ (Old) 5.9 ○ (New) 4.2
Disable legacy JQuery UI usage	Fallback Value • (Old) true • (New) false
Enforce application specific ACLs only for application data	• Description ○ (Old) Properties in the format 'glide.enforce_security_scope.' property, such as 'glide.enforce_security_scope.sn_hr_core', control the behavior of application data residing in primary tables outside the application, such as 'sys_attachment' or 'sys_email'. When properties in the format have a value of 'true', only the application specific ACLs are evaluated for access to the application data residing in these tables. When the properties do not have a value of 'true', the ACLs on the primary table will still be evaluated for access, potentially allowing unauthorized or undesired access to application

Documentation	Updates
	data. Not all applications are designed to work in this configuration or ship a 'sys_properties' record for this purpose. The following application scopes contain this property: ▪ sn_uni_task ▪ sn_uni_req ▪ sn_svc_appl_info ▪ sn_professional ▪ sn_opp_market ▪ sn_lg_ops ▪ sn_lg_matter ▪ sn_lg_contracts ▪ sn_jny ▪ sn_ja ▪ sn_imt_vaccine ▪ sn_imt_tracing ▪ sn_imt_health_test ▪ sn_hr_ws ▪ sn_hr_va ▪ sn_hr_sp ▪ sn_hr_pj ▪ sn_hr_pad ▪ sn_hr_mii_base ▪ sn_hr_le ▪ sn_hr_hc ▪ sn_hr_gen_ai ▪ sn_hr_er ▪ sn_hr_ef ▪ sn_hr_core ▪ sn_hr_awa ▪ sn_hr_agent_ws ▪ sn_hc_professional ▪ sn_gsm_soc_bnfts ▪ sn_gsm_lic_prmt_ex ▪ sn_gsm_lic_prmt ▪ sn_gsm_info_req ▪ sn_gsm

Documentation	Updates
	▪ sn_em ▪ sn_egd_goals ▪ sn_egd_core ▪ sn_egd_act ▪ sn_doc ○ (New) Control the behavior of application data residing in primary tables outside the application. When these properties have a primary value of true, only the application-specific ACLs are evaluated for access to the application data residing in these tables. Not all applications are designed to work in this configuration or use a System Property [sys_properties] record for this purpose. These system properties use the glide.enforce_security_scope. naming format. For example, use the glide.enforce_security_scope.sn_hr_sp property for the Employee Center Core (sn_hr_sp) scope. The following application scopes contain this property: ▪ sn_doc ▪ sn_egd_act ▪ sn_egd_core ▪ sn_egd_goals ▪ sn_em ▪ sn_gsm ▪ sn_gsm_info_req ▪ sn_gsm_lic_prmt ▪ sn_gsm_lic_prmt_ex ▪ sn_gsm_soc_bnfts ▪ sn_hc_professional ▪ sn_hr_agent_ws ▪ sn_hr_ai_agents ▪ sn_hr_awa ▪ sn_hr_core ▪ sn_hr_ef ▪ sn_hr_er ▪ sn_hr_gen_ai ▪ sn_hr_hc

Documentation	Updates
	▪ sn_hr_le ▪ sn_hr_le_ent ▪ sn_hr_mii_base ▪ sn_hr_na_galileo ▪ sn_hr_pad ▪ sn_hr_pj ▪ sn_hr_sp ▪ sn_hr_va ▪ sn_hr_ws ▪ sn_imt_health_test ▪ sn_imt_tracing ▪ sn_imt_vaccine ▪ sn_ja ▪ sn_jny ▪ sn_lg_contracts ▪ sn_lg_matter ▪ sn_lg_ops ▪ sn_opp_market ▪ sn_professional ▪ sn_svc_appl_info ▪ sn_svc_appl_pgm_mg ▪ sn_talent_aia ▪ sn_uni_req ▪ sn_uni_task When the properties do not have a value of 'true', the ACLs on the primary table will still be evaluated for access, potentially allowing unauthorized or undesired access to application data. • Remediation ○ (Old) For any applications installed with the 'glide.enforce_security_scope.' property, such as 'glide.enforce_security_scope.sn_hr_core', in the 'sys_properties' table, ensure the property value is set to 'true'. These properties can only be modified by the scoped administrator for the specific application.

Documentation	Updates
	○ (New) For each application installed with the glide.enforce_security_scope property in the System Properties [sys_properties] table, (for example, glide.enforce_security_scope.sn_hr_core), ensure the property value is set to true. These properties can only be modified by the scoped administrator for the specific application. If a sys_properties record does not exist for the given application and respective property, it must be created. Use this script can to find which properties need to be updated or created on the instance: `var properties = [` `'glide.enforce_security_scope.sn_uni_task',` `'glide.enforce_security_scope.sn_uni_req',` `'glide.enforce_security_scope.sn_svc_appl_info',` `'glide.enforce_security_scope.sn_professional',` `'glide.enforce_security_scope.sn_opp_market',` `'glide.enforce_security_scope.sn_lg_ops',` `'glide.enforce_security_scope.sn_lg_matter',` `'glide.enforce_security_scope.sn_lg_contracts',` `'glide.enforce_security_scope.sn_jny',` `'glide.enforce_security_scope.sn_ja',` `'glide.enforce_security_scope.sn_imt_vaccine',`

Documentation	Updates
	'glide.enforce_security_scope.sn_imt_tracing', 'glide.enforce_security_scope.sn_imt_health_test', 'glide.enforce_security_scope.sn_hr_ws', 'glide.enforce_security_scope.sn_hr_va', 'glide.enforce_security_scope.sn_hr_sp', 'glide.enforce_security_scope.sn_hr_pj', 'glide.enforce_security_scope.sn_hr_pad', 'glide.enforce_security_scope.sn_hr_mii_base', 'glide.enforce_security_scope.sn_hr_le', 'glide.enforce_security_scope.sn_hr_le_ent', 'glide.enforce_security_scope.sn_hr_hc', 'glide.enforce_security_scope.sn_hr_gen_ai', 'glide.enforce_security_scope.sn_hr_er', 'glide.enforce_security_scope.sn_hr_ef', 'glide.enforce_security_scope.sn_hr_core', 'glide.enforce_security_scope.sn_hr_awa', 'glide.enforce_security_scope.sn_hr_agent_ws', 'glide.enforce_security_scope.sn_hc_professional',

Documentation	Updates
	'glide.enforce_security_scope.sn_gsm_soc_bnfts', 'glide.enforce_security_scope.sn_gsm_lic_prmt_ex', 'glide.enforce_security_scope.sn_gsm_lic_prmt', 'glide.enforce_security_scope.sn_gsm_info_req', 'glide.enforce_security_scope.sn_gsm', 'glide.enforce_security_scope.sn_em', 'glide.enforce_security_scope.sn_egd_goals', 'glide.enforce_security_scope.sn_egd_core', 'glide.enforce_security_scope.sn_egd_act', 'glide.enforce_security_scope.sn_doc', 'glide.enforce_security_scope.sn_talent_aia', 'glide.enforce_security_scope.sn_hr_na_galileo', 'glide.enforce_security_scope.sn_svc_appl_pgm_mg', 'glide.enforce_security_scope.sn_hr_ai_agents', 'glide.enforce_security_scope.sn_hr_mii_base']; var pm = new GlidePluginManager(); for (var i = 0; i < properties.length; i++) { var property = properties[i];

Documentation	Updates
	`` var application = property.split('.')[2]; var propertyValue = gs.getProperty(property, 'false'); if (pm.isActive(application) && propertyValue.toLowerCase () != 'true') { gs.print(property); } } ``` • CVSS score ○ (Old) 5.5 ○ (New) 4.1
Prevent impersonating user from viewing application data	• Description ○ (Old) Properties in the format '.impersonateCheck', such as 'sn_hr_core.impersonateCheck', control how an impersonating user can access specific application data belonging to another user. When properties in this form are set to 'true', an admin level user impersonating another user will not be able to access the application specific data belonging to that user. Not all applications are designed to work in this configuration or ship a 'sys_properties' record for this purpose. A value of "false" for these properties allows an admin level user to impersonate another user and access application data with the impersonated user's access. This may be undesirable or allow for unauthorized data access in specific application contexts. The following scopes contain this property: ▪ sn_opp_market ▪ sn_jny ▪ sn_imt_vaccine ▪ sn_imt_health_test

Documentation	Updates
	▪ sn_hr_core ▪ sn_egd_goals ▪ sn_egd_core ▪ sn_egd_act ○ (New) Use system properties to prevent an impersonating user from viewing application data. Prevent admin level from accessing the application specific data belonging to that user when impersonating an account. This permission can be set at the application level by creating a system property specific to the application. These system properties use the .impersonateCheck naming format (for example sn_hr_core.impersonateCheck). Create a system property with a value of true to prevent users from accessing the application-specific data belonging to another user when impersonating an account. NOTE: Not all applications are designed to work in this configuration or have a System Properties [sys_properties] record for this purpose. The following scopes are configured to work with this property. ▪ sn_opp_market ▪ sn_jny ▪ sn_imt_vaccine ▪ sn_imt_health_test ▪ sn_hr_core ▪ sn_egd_goals ▪ sn_egd_core ▪ sn_egd_act ▪ sn_em ▪ sn_talent_aia • Remediation ○ (Old) For any application with the '.impersonateCheck' property, such as 'sn_hr_core.impersonateCheck', in the 'sys_properties' table, ensure the property

Documentation	Updates
	value is set to 'true'. These properties can only be modified by the scoped administrator for the specific application. ○ (New) For each application with the .impersonateCheck property in the System Properties [sys_properties] table, ensure the property value is set to true. These properties can only be modified by the scoped administrator for the specific application. Use this script to find which properties need to be updated or created on the instance: `var properties = ['sn_opp_market.impersonateCheck', 'sn_jny.impersonateCheck', 'sn_imt_vaccine.impersonateCheck', 'sn_imt_health_test.impersonateCheck', 'sn_hr_core.impersonateCheck', 'sn_egd_goals.impersonateCheck', 'sn_egd_core.impersonateCheck', 'sn_egd_act.impersonateCheck', 'sn_em.impersonateCheck', 'sn_talent_aia.impersonateCheck'];` `var pm = new GlidePluginManager();` `for (var i = 0; i < properties.length; i++) { var property = properties[i];`

Documentation	Updates
	`` var application = property.split('.')[0]; var propertyValue = gs.getProperty(property, 'false'); if (pm.isActive(application) && propertyValue.toLowerCase() != 'true') { gs.print(property); } } ``` • CVSS score ○ (Old) 4.4 ○ (New) 3.8
Escape JavaScript [Updated in Security Center 1.3]	• Description ○ (Old) <blank> ○ (New) The property "glide.html.escape_script" helps sanitize HTML fields. If "glide.html.escape_script" is not set to the recommended value of "true", then inputs will not be sanitized for HTML fields (output encoding) from a backend Java context by removing embedded JavaScript. Javascript in HTML fields can lead to stored and reflected XSS. • security risk ○ (Old) <blank> ○ (New) The ability to have XSS can lead to easily attained privilege escalation to higher roles such as admin.
Enable HTML Sanitizer [Updated in Security Center 1.3]	• Description ○ (Old) <blank> ○ (New) This property controls sanitization behavior of HTML fields on a global level from a backend Java context. If "glide.html.sanitize_all_fields" is not set to the recommended value of "true", then, the ServiceNow instance is open to XSS in HTML fields.

Documentation	Updates
	• security risk ○ (Old) <blank> ○ (New) The ability to have XSS can lead to easily attained privilege escalation to higher roles such as admin.
Escape jelly script [Updated in Security Center 1.3 and 1.5]	• Description ○ (Old) <blank> ○ (New) This property escapes all the JS and HTML strings included in <j:jelly> ... </j:jelly> before they are written to the output stream, preventing several XSS issues from occurring. If "glide.ui.escape_all_script" is not set to the recommended value of "true", then escaping of scripts injected into Jelly is disabled. • security risk ○ (Old) <blank> ○ (New) Without this mitigation, the platform becomes widely open to a variety of script injection attacks. An attacker could execute arbitrary Rhino scripts on the instance.
	• Description ○ (Old) <blank> ○ (New) This property enables the "script sandbox" feature. The script sandbox is mainly utilized when executing client-generated scripts (such as query conditions and GlideAjax expressions). If "glide.script.use.sandbox" is not set to the recommended value of "true", then the script sandbox feature will be disabled. • security risk ○ (Old) <blank> ○ (New) Without the script sandbox, unauthorized/unauthenticated users can execute

Documentation	Updates
	arbitrary privileged script on a ServiceNow instance. This would lead to complete security impact across all areas, including, but not limited to potentially malicious access to all data on the ServiceNow instance.
Disallow target cloning [New in Security Center 1.3]	• Description ○ (Old) If glide.db.clone.allow_clone_target is not set to the recommended value of false, then the instance can be used as a clone target, or a record that specifies the instance URL and credentials used for cloning. A system clone is when everything in a database is copied from one instance to another. This is a security risk because the instance database can be overwritten in the cloning process, leading to data loss and lack of data integrity. As a remediation, ensure that glide.db.clone.allow_clone_target is set to false. ○ (New) Protect your instance from being used as a clone target by setting the glide.db.clone.allow_clone_target system property to false. A system clone copies everything in a database from a source instance to the target instance. This is a security risk because the instance database on the target instance is overwritten in the cloning process, leading to data loss and lack of data integrity. • Remediation ○ (Old) Ensure the property "glide.db.clone.allow_clone_target" is set to "false". ○ (New) Set the glide.db.clone.allow_clone_target system property to false on production instances to disallow your instance from being selected as a clone target. • CVSS score

Documentation	Updates
	○ (Old) 5.9 ○ (New) 4.4 • Data type ○ (Old) <blank> ○ (New) Boolean • Out of box value ○ (Old) <blank> ○ (New) true
Deny internal access to explicit external roles [Updated in Security Center 1.3 and 1.5]	• Description ○ (New) Use system properties to determine whether external users can be assigned the snc_internal role. Use the glide.security.explicit_roles.enable_internal_user_blacklist system property to prevent external users from being assigned the snc_internal role. When this property is set to true, it enforces the parameters of the maint-protected glide.security.explicit_roles.internal_user_blacklist property. This property assigns the snc_external role to a list of untrusted user classes. If glide.security.explicit_roles.enable_internal_user_blacklist is set to false, the glide.security.explicit_roles.internal_user_blacklist property is ignored. ○ (Old) This property prevents external users from being assigned the snc_internal role. When "glide.security.explicit_roles.enable_internal_user_blacklist is set to the recommended value of "true", then it enforces the parameters of the maint-protected "glide.security.explicit_roles.internal_user_blacklist" property which assigns the 'snc_external' role to a list of untrusted user classes. If the value is set to false, the "glide.security.explicit_roles.internal_user_blacklist" property is ignored. Misconfiguration of this property increases the risk that an external user account gains access to internal information. • Plugin applicability

Documentation	Updates
	○ (New) Explicit Roles Plugin, Customer Service Base Extension Entities ○ (Old) Explicit Roles Plugin • security risk ○ (New) Misconfiguration of this property increases the risk that an external user account gains access to internal information. ○ (Old) <blank> • Data type ○ (New) Boolean ○ (Old) <blank> • Out of box value ○ (New) true ○ (Old) <blank> • Fallback value ○ (New) false ○ (Old) true
Restrict oauth parameters to POST body [New in Security Center 1.3]	• Description ○ (Old) This property controls the inbound OAuth authentication's acceptance of access tokens. Access tokens are sensitive and should only be accepted when located within a POST request body. ○ (New) Use the glide.oauth.allow.parameters.in.post.body.only property to control the inbound OAuth authentication's acceptance of access tokens. Access tokens are sensitive and should only be accepted when located within a POST request body. • Remediation

Documentation	Updates
	○ (Old) Ensure the property "glide.oauth.allow.parameters.in.post.body.only" is set to "true". ○ (New) Ensure the property "glide.oauth.allow.parameters.in.post.body.only" is set to "true". If the property does not exist in the "sys_properties" table, the default is "false". • Plugin applicability ○ (Old) <blank> ○ (New) OAuth 2.0 • security risk ○ (Old) If "glide.oauth.allow.parameters.in.post.body.only" is not set to the recommended value of "true", then access tokens could be present in the GET request parameter which could linger in client and infrastructure logs and potentially lead to account takeover if those logs are leaked. ○ (New) If glide.oauth.allow.parameters.in.post.body.only isn't set to the recommended value of true, access tokens could be present in the GET request parameter. These access tokens could linger in client and infrastructure logs and potentially lead to account takeover if those logs are leaked. • Dependencies and prerequisites ○ (Old) <blank> ○ (New) Plugin OAuth 2.0 • Data type ○ (Old) <blank> ○ (New) Boolean • Out of box value ○ (Old) <blank> ○ (New) true

Documentation	Updates
Enforce GroupBy ACLs	• Description ○ (Old) If "glide.security.groupby_acl_check" is not set to the recommended value of "True", then make sure that a table has "groupby_acl_check" attribute set in order to honor groupby ACLs. In other case there will be no ACLs check on groupby columns of a table. This could lead to information disclosure. ○ (New) Use the glide.security.groupby_acl_check system property to configure your instance to conduct ACL checks on groupby columns. If this property is set to the recommended value of true, then ACLs on groupby columns are honored by default. A table's groupby_acl_check attribute takes precedent over the glide.security.groupby_acl_check property. If the property is set to false, then ensure that any table which should have ACL checks on groupby columns has the groupby_acl_check attribute set to true. • security risk ○ (Old) <blank> ○ (New) If "glide.security.groupby_acl_check" is set to false and there is no "groupby_acl_check" attribute on the individual table, then ACLs on groupby columns will not be honored which could lead to information leakage. • Functional impact ○ (Old) <blank> ○ (New) ACLs on groupby columns will be enforced by default for tables. • Data type ○ (Old) <blank> ○ (New) Boolean • Out of box value

Documentation	Updates
	○ (Old) <blank> ○ (New) true
Require XMLdoc2 entity validation with allowlistDisable entity expansion [Updated in Security Center 1.3]	• Description ○ (Old) If the Glide Property "glide.stax.whitelist_enabled" does not exist in the sys_properties table or is not set to the recommended value of "true", then all external entities are allowed when the Glide Property "glide.stax.allow_entity_resolution" is set to the value of "true". If customizations do not require entity expansion, use the "glide.stax.allow_entity_resolution" property to completely disable external entity expansion. The XML completes parsing but doesn't include any internal or external entities. ▪ If you set "glide.stax.allow_entity_resolution" to true, all external entities attempt to resolve or expand subject entities, subject to the setting of the "glide.stax.whitelist_enabled" property. ▪ If you set "glide.stax.allow_entity_resolution" to false, all entity resolution and expansion is blocked. To learn more about this property, see https://www.servicenow.com/docs/csh?topicname=sc-disable-entity-expansion.html&version=latest. When "glide.stax.whitelist_enabled" is set to true, define a listing of comma-delimited FQDN in the "glide.xml.entity.whitelist" property, which is the only URLs that can be reached using XML Entity processing property. To learn more, see https://www.servicenow.com/docs/csh?topicname=sc-xml-entity-validation-url-allowlist.html&version=latest. ○ (New) If the Glide Property "glide.stax.whitelist_enabled" does not exist in the sys_properties table or is not set to the recommended

Documentation	Updates
	value of "true", then all external entities are allowed when the Glide Property "glide.stax.allow_entity_resolution" is set to the value of "true". If customizations do not require entity expansion, use the "glide.stax.allow_entity_resolution" property to completely disable external entity expansion. The XML completes parsing but doesn't include any internal or external entities. ▪ If you set "glide.stax.allow_entity_resolution" to true, all external entities attempt to resolve or expand subject entities, subject to the setting of the "glide.stax.whitelist_enabled" property. ▪ If you set "glide.stax.allow_entity_resolution" to false, all entity resolution and expansion is blocked. To learn more about this property, see https://www.servicenow.com/docs/csh?topicname=sc-disable-entity-expansion.html&version=latest. When "glide.stax.whitelist_enabled" is set to true, define a listing of comma-delimited FQDN in the "glide.xml.entity.whitelist" property, which is the only URLs that can be reached using XML Entity processing property. To learn more, see https://www.servicenow.com/docs/csh?topicname=sc-xml-entity-validation-url-allowlist.html&version=latest. • Remediation ○ (Old) Ensure the property "glide.stax.whitelist_enabled" is set to "true". ○ (New) Ensure the property "glide.stax.whitelist_enabled" is set to "true" when the Glide Property "glide.stax.allow_entity_resolution" is set to the value of "true".

Documentation	Updates
Define restricted downloadable MIME types [Updated in Security Center 1.3, 1.5, and 2.0]	• Description ◦ (Old) If "glide.ui.attachment.download_mime_types" does include dangerous items such as "text/html,image/svg,image/svg +xml,application/xml", then dangerous files could be rendered inline in the browser which could lead to Cross Sitte Scripting attacks (XSS). This property is the list of comma separated attachment mime types which will not render inline in the browser. For example, including text/html will force HTML files to be downloaded to the client as attachments rather than viewed inline in the browser. Maintaining this list properly will prevent cross site scripting attacks. ◦ (New) If glide.ui.attachment.download_mime_types does include dangerous MIME types such as text/html, image/svg ,image/svg +xml,application/xml, then dangerous files could be rendered inline in the browser, which could lead to Cross Site Scripting attacks (XSS). This property is the list of comma-separated attachment mime types, which won't render inline in the browser. For example, including text/html forces HTML files to be downloaded to the client as attachments rather than viewed inline in the browser. Maintaining this list properly prevents cross-site scripting attacks. If the glide.ui.attachment.download_mime_types system property doesn't include dangerous MIME types such as "text/html, image/svg,image/svg+xml,application/ xml", then dangerous files could be rendered inline in the browser. This can lead to Cross Site Scripting (XSS) attacks. This check is only relevant when glide.ui.attachment.force_download_all_mime_types is set to false. This property is a list of comma-separated attachment MIME types, which don't render inline in the browser. For example,

Documentation	Updates
	including text/html forces HTML files to be downloaded to the client as attachments rather than viewed inline in the browser. • Remediation ○ (Old) If glide.ui.attachment.force_download_all_mime_types is set to false, verify that the glide.ui.attachment.download_mime_types system property includes the dangerous MIME types "text/html,image/svg,image/svg+xml,application/xml". ○ (New) Ensure the property "glide.ui.attachment.download_mime_types" includes the dangerous items "text/html,image/svg,image/svg +xml,application/xml". • Security Risk ○ (Old) <blank> ○ (New) Maintaining this list properly can prevent cross site scripting attacks. • Functional Impact ○ (Old) <blank> ○ (New) Attachments with the MIME type listed in this property cannot be viewed inline in the browser. • Dependencies and prerequisites ○ (Old) <blank> ○ (New) This check is only relevant when glide.ui.attachment.force_download_all_mime_types is set to false or does not exist in the System Properties [sys_properties] table. • Data Type ○ (Old) <blank> ○ (New) Comma separated list of MIME types • Out of box value

Documentation	Updates
	○ (Old) <blank> ○ (New) text/html,image/svg,image/svg +xml,application/xml
Escape HTML in list views [Updated in Security Center 1.3 and 1.5]	• Description ○ (Old) This property helps sanitize list view displaying of HTML fields. If "glide.ui.escape_html_list_field" is not set to the recommended value of "true", then a malicious user can inject HTML code within the form field to execute unwanted scripts on different client/user sessions. This could potentially be leveraged by attackers to steal session information and sensitive data. ○ (New) Set glide.ui.escape_html_list_field to true to prevent HTML from being rendered in HTML fields in list view. Leaving HTML sanitization inactive platform wide (via system property) or by field (via a schema attribute), may lead to XSS style attacks. XSS attacks may allow a low privileged user to hijack the session of a high privileged user or interfere in standard web application behaviors, including redirects or defacement. • CVSS score ○ (Old) 8.8 ○ (New) 3.1 • Security risk ○ (Old) <blank> ○ (New) When HTML sanitization is disabled platform wide (via Glide Properties) or per field (schema attribute), this may lead to XSS style attacks if low privileged users have access to write to an HTML field. XSS attacks may allow a low privileged user to hijack the session of a high privileged user or interfere in standard web application behaviors (redirects or defacement).

Documentation	Updates
	• Functional impact ○ (Old) <blank> ○ (New) By default, HTML can be rendered (and is sanitized) in form view. This same behavior may be desired in list view, in which case setting this property to "false" may be the preferred experience. There will be minimal security impact as long as HTML sanitizaiton is not disabled platform wide or on the field. • Data type ○ (Old) <blank> ○ (New) Boolean • Out of box value ○ (Old) <blank> ○ (New) true
Minimize reset password max SMS per day [Updated in Security Center 1.3]	• Data type ○ (Old) <blank> ○ (New) Integer • Out of box value ○ (Old) <blank> ○ (New) 10 • Fallback value ○ (Old) 10 ○ (New) 5
Disallow infected file download [Updated in Security Center 1.5 and 2.0]	Out of box value • (Old) <blank> • (New) true
Restrict unauthenticated access to attachments	Description • (Old) Description (Old): If "glide.image_provider.security_enabled" is not set to the recommended value of "True", then all images are accessible via urls that end in ".iix". This would

Documentation	Updates
	allow unauthenticated access to images leading to sensitive information leak. This property is not honored for the images from attachment table where origin table is [sysevent_email_style, sys_home, sys_properties]. Restriction should be applied for unauthenticated users as some attachments might contain sensitive information. • (New) Secure the images on your instance to prevent sensitive information leak. Images on your instance are accessible via urls that end in `.iix`. Set the **glide.image_provider.security_enabled** system property to **true** to prevent access to your images via these URLs. **Note:** This property is not honored for images from the attachment table if the origin table is one of: ○ Stationeries [sysevent_email_style] ○ Welcome Page Sections [sys_home] ○ System Properties [sys_properties] Restriction should be applied for unauthenticated users as some attachments might contain sensitive information.
Activate role based multi-factor authentication [Updated in Security Center 1.3]	CVSS Score • (Old) 8.8 • (New) 3.1
Maximize failed login unlock timeout duration [Updated in Security Center 1.3]	• Description ○ (Old) If "glide.user.unlock_timeout_in_mins" is not set to the recommended value of "15", then it may be easier to brute force accounts in a faster timeframe. This property unlocks the user account after the time period that is specified for the glide.user.unlock_timeout_in_mins

Documentation	Updates
	property. If no value is specified, then the system unlocks the user account after the default period of 15 minutes. If the property is not configured to a secure value and the lockout duration is not enabled, then it may be easier to brute force account logins in a faster timeframe. This may allow a malicious user to eventually obtain unauthorized access to the instance. Impact on the instance will be limited to the privileged of the affected user login brute-forced. ○ (New) If "glide.user.unlock_timeout_in_mins" is not set to at least the minimum value of "15", then it may be easier to brute force accounts in a faster timeframe. This property unlocks the user account after the time period that is specified for the glide.user.unlock_timeout_in_mins property. If no value is specified, then the system unlocks the user account after the default period of 15 minutes. If the property is not configured to a secure value and the lockout duration is not enabled, then it may be easier to brute force account logins in a faster timeframe. This may allow a malicious user to eventually obtain unauthorized access to the instance. Impact on the instance will be limited to the privileged of the affected user login brute-forced. • Remediation ○ Set the glide.user.unlock_timeout_in_mins system property value to a minimum of 15. If glide.user.unlock_timeout_in_mins does not exist, the default lockout time is set to 15 minutes. Ensure that the SNC User Lockout Check with Auto Unlock script action (found on the Script Action [sysevent_script_action] table) is present and active. The SNC User Lockout Check with Auto Unlock script action is installed with the High Security Settings (com.glide.high_security) plugin. Ensure the property "glide.user.unlock_timeout_in_mins" is set to "15" or more and that the Script

Documentation	Updates
	Action (sysevent_script_action) "SNC User Lockout Check with Auto Unlock" is present and active. If the Glide Property "glide.user.unlock_timeout_in_mins" does not exist, it will default to a secure value of "15". The "SNC User Lockout Check with Auto Unlock" is installed with the High Security Plugin. ○ (New) 10
Maximize failed login unlock timeout duration [Updated in Security Center 1.3]	Remediation • (Old) Ensure at least one of the script actions: "SNC User Lockout Check" or "SNC User Lockout Check with Auto Unlock" are enabled to manage failed login attempts. Additionally, ensure the property "glide.user.max_unlock_attempts" is set to "5" or less. • (New) Ensure at least one of the script actions: "SNC User Lockout Check" or "SNC User Lockout Check with Auto Unlock" are enabled to manage failed login attempts. These script actions are stored in the "sysevent_script_action" table. Additionally, ensure the property "glide.user.max_unlock_attempts" is set to "5" or less.
Set OTP lifetime for password reset to 1 hour [Updated in Security Center 2.0]	• Short Description ○ (Old) Set OTP Lifetime for Password Reset to 12 Hours or Less ○ (New) Set OTP lifetime for password reset to 1 hour • Description ○ (Old) This property "glide.pwd_reset.onetime.token.validity" allows the link in the password reset email to expire after the number of hours specified in that "glide.pwd_reset.onetime.token.validity"

Documentation	Updates
	property. Validity time of password reset token should be keeped as short as possible in according of normal user experience. Have long validity time for password reset token can help malicious actors to perform account takeover. ◦ (New) Control the time duration of the link in the password reset email. The property glide.pwd_reset.onetime.token.validity makes the link in the password reset email expire after the number of hours specified in the property. The validity time of a password reset token should be kept as short as possible while not disrupting normal user experience. A longer validity time for password reset token gives malicious actors a wider window to perform account takeover if the email with the reset token is leaked or otherwise compromised. • Remediation ◦ (Old) Set the property value to 12 hours or a shorter validity time. ◦ (New) Set the property value to 1 (in hours) • CVSS ◦ (Old) 5.6 ◦ (New) 4.6 • Security Risk ◦ (Old) <blank> ◦ (New) A longer validity time for password reset token gives malicious actors a wider window to perform account takeover if the email with the reset token is leaked or compromised. • Functional Impact ◦ (Old) <blank> ◦ (New) A user must reset their password within the number of hours specified in this property. Otherwise they will need to request a new link. • Data type

Documentation	Updates
	○ (Old) <blank> ○ (New) Integer representing number of hours • Out of box value ○ (Old) <blank> ○ (New) 1 • Fallback value ○ (Old) 12 ○ (New) 1
Limit concurrent interactive sessions [Updated in Security Center 1.3]	• Short Description ○ (Old) Limit Concurrent Interactive Sessions ○ (New) Limit Concurrent Interactive Sessions When Limit Concurrent Sessions Plugin Is Installed • Plugin Applicability ○ (Old) <blank> ○ (New) com.glide.limit.concurrent.sessions
Limit maximum number of attachments in email	• Data type ○ (Old) <blank> ○ (New) Integer • Out of box value ○ (Old) <blank> ○ (New) 30 • Fallback value ○ (Old) 30 ○ (New) 50
Enable protected tables plugin [New in Security Center 1.3]	• Description ○ (Old) When the Glide Property 'com.glide.security.protected_table.enabled' is set to 'true', The Protected Tables plugin will be utilized to prevent higher privilege users on an instance from tampering with log tables. The following logging tables will have special protections when this property is set to 'true':

Documentation	Updates
	▪ syslog (config not modifiable) ▪ syslog_transaction ▪ sys_outbound_http_log ▪ sysevent ▪ sys_audit ▪ sys_push_notification ▪ protected_table_configuration (config not modifiable) ○ (New) When the Glide Property 'com.glide.security.protected_table.enabled' is set to 'true', The Protected Tables plugin will be utilized to prevent higher privilege users on an instance from tampering with log tables. The following logging tables will have special protections when this property is set to 'true': ▪ syslog (config not modifiable) ▪ syslog_transaction ▪ sys_outbound_http_log ▪ sysevent ▪ sys_audit ▪ sys_push_notification ▪ protected_table_configuration (config not modifiable) ▪ syslog_app_scope The integrity of logs is important for determining malicious activity on an instance by a customer admin. • Remediation ○ (Old) Set the Glide Property 'com.glide.security.protected_table.enabled' to 'true'. ○ (New) Set the Glide Property 'com.glide.security.protected_table.enabled' to 'true'. This requires a user with the "security_admin" role. • CVSS Score ○ (Old) 4.5 ○ (New) 4

Documentation	Updates
Ensure dashboards creation/deletion requires access check [New in Security Center 1.3 and updated in 2.0]	• Description ○ (Old) The property 'glide.processors.check_access_before_process' enables ACL enforcement for creating or deleting dashboards so long as a user is logged in. Disabling this property (i.e., setting it to false) effectively allows for an ACL bypass on dashboards which allows all authenticated low privileged users to arbitrarily delete and add dashboards. This property should always be set to true. ○ (New) The property 'glide.processors.check_access_before_process' enables ACL enforcement for creating or deleting dashboards. When the property is set to "true", access control checks are performed on the dashboards. When this property is set to "false", authenticated users can arbitrarily delete and add dashboards. • Remediation ○ (Old) Ensure the Glide Property 'glide.processors.check_access_before_process' exists and is set to the value 'true'. If the property does not appear in the sys_properties table, add a new record. ○ (New) Set the glide.processors.check_access_before_process system property to true. If the property does not appear in the System Properties [sys_properties] table, the fallback value is true. • Security Risk ○ (Old) <blank> ○ (New) When this property is set to "false", authenticated users can arbitrarily delete and add dashboards. • Functional impact

Documentation	Updates
	○ (Old) <blank> ○ (New) Users may not have access to delete dashboards as before. To mitigate this functional impact, users can be granted standard access to dashboards. • Fallback Value ○ (Old) false ○ (New) true
Define active session timeout exception roles [New in Security Center 1.3]	• Description ○ (Old)Active session timeouts are a feature to ensure that a hijacked session cannot be used indefinitely without providing authentication information. This property controls the roles that are exempted from an active session timeout limit. It is best practice to only consider an active session timeout limit exception for internal integration account roles. If a role is given an exception to the session timeout, and that role is given to a user who is the victim of a session hijacking attack, an attacker can continue to authenticate to that session indefinitely. This may increase the affected scope of a security incident by allowing an attacker more time to make use of a hijacked account. ○ (New) Use the glide.active.session.timeout.exception.roles system property to exempt roles from an active session timeout limit. The active session timeout feature helps ensure that a hijacked session can't be used indefinitely without providing authentication information. It is best practice to only consider an active session timeout limit exception for internal integration account roles. Consider an active session timeout limit exception only for internal integration account roles. If a user is a victim of a session hijacking attempt, and has a role with an exception, attackers using that session can continue to authenticate to that session indefinitely. This may increase the impact of a security incident by enabling an attacker more time to make use of a hijacked account.

Documentation	Updates
	• Remediation ◦ (Old) Ensure Glide Property 'glide.active.session.timeout.exception.roles' is set to value 'edge_encryption,mid_server'. ◦ (New) Configure the glide.active.session.timeout.exception.roles property to roles which should be exempt from active session timeouts. This property value is a comma separated list of roles. The default value is edge_encryption,mid_server,maint. • Security Risk ◦ (Old) Only consider an active session timeout limit exception for internal integration account roles. If a role is given an exception to the session timeout, and that role is given to a user who is the victim of a session hijacking attack, then an attacker can continue to authenticate to that session indefinitely. This may increase the impact of a security incident by enabling an attacker more time to make use of a hijacked account. ◦ (New) Consider an active session timeout limit exception only for internal integration account roles. If a user is a victim of a session hijacking attempt, and has a role with an exception, attackers using that session can continue to authenticate to that session indefinitely. This may increase the impact of a security incident by enabling an attacker more time to make use of a hijacked account. • Functional impact ◦ (Old) <blank> ◦ (New) Roles added to this list will be exempt from active session timeout limit. • Data type ◦ (Old) <blank> ◦ (New) Comma-separated list of roles • Out of Box Value ◦ (Old) <blank> ◦ (New) edge_encryption,mid_server,maint

Documentation	Updates
	• Fallback Value ○ (Old) edge_encryption,mid_server ○ (New) edge_encryption,mid_server,maint
Limit HTTP response body size [New in Security Center 1.3 and updated in 1.5]	• Data type ○ (Old) <blank> ○ (New) Boolean, Integer • Out of box value ○ (Old) <blank> ○ (New) true,524288000

Updated hardening settings for baseline version 6.0

Some hardening settings have been updated with the release of Security Center baseline version 6.0.

Documentation	Updates
Prevent Unauthenticated Access to Virtual Agent Embedded Web Client	• New Short Description: Prevent Unauthenticated Access to Virtual Agent Embedded Web Client • Old Short Description:Publicly Exposed Virtual Agent Embedded Web Client sn_va_web_client_app_embed
Prevent Empty ACL Creation [New in Security Center 2.0]	Rule: Script: Script has been updated to improve detection accuracy.
Enable ACLs for Encoded Query in Simple List Widget [New in Security Center 2.0]	• CVSS Score (New): 4.3 • CVSS Score (Old): 5.3
Sanitize All Translated HTML Fields [New in Security Center 2.0]	• CVSS Score (New): 4.6 • CVSS Score (Old): 8.8
Enable HTML Sanitizer [Updated in Security Center 1.3]	Rule Script: Script has been updated to improve detection accuracy
Implement the x-frame-options: SAMEORIGIN security header [Updated in Security Center 1.3]	• CVSS Score (New): 5.9 • CVSS Score (Old): 7.1
Restrict access to GlideSystemUserSession scriptable API [Updated in Security Center 1.3 and 2.0]	• Remediation (New): Ensure the property *glide.sandbox.usersession.allow_unsanitize* is set to **false**. If a System Property

Documentation	Updates
	[sys_properties] record does not exist for this property, create one. • Remediation (Old): Ensure the property *glide.sandbox.usersession.allow_unsanitize* is set to **false**.
Disallow target cloning [New in Security Center 1.3]	• Description (New): If *glide.db.clone.allow_clone_target* is not set to the recommended value of **false**, then the instance can be used as a clone target, or a record that specifies the instance URL and credentials used for cloning. A system clone is when everything in a database is copied from one instance to another. This is a security risk because the instance database can be overwritten in the cloning process, leading to data loss and lack of data integrity. As a remediation, ensure that *glide.db.clone.allow_clone_target* is set to **false**. Not setting this property to the recommended value of false enables the instance to be used as a clone target. This is a security risk because the instance database can be overwritten in the cloning process. • Description (Old): If *glide.db.clone.allow_clone_target* is not set to the recommended value of **false**, then the Instance can be used as a clone target. This risks the instance database being overwritten by cloning process, leading to integrity and availability loss.
Restrict oauth parameters to POST body [New in Security Center 1.3]	• CVSS Score (New): 4.2 • CVSS Score (Old): 7.4
Enforce URL allowlist check [Updated in Security Center 1.3, 1.5, and 2.0]	• Description (New): If *glide.security.url.whitelist.strict_check* is not set to the recommended value of **true** then all external URLs are allowed for redirection when *glide.security.url.whitelist* is empty. If *glide.security.url.whitelist* is not empty, then only external URLs in the whitelist are allowed. Thus either setting *glide.security.url.whitelist.strict_check* to true OR ensuring

Documentation	Updates
	glide.security.url.whitelist is set to a non-empty value with the allowed external URLs leaves the instance in a secure state. If all external URLs are allowed for redirection, this could allow an attacker to redirect a user to a malicious website. • Description (Old): If *glide.security.url.whitelist.strict_check* is not set to the recommended value of **true**, and if *glide.security.url.whitelist* is not set to an organization's approved URL, then all external URLs are allowed for redirection. This could allow an attacker to redirect a user to a malicious website. • CVSS Score (New): 6.3 • CVSS Score (Old): 8.3
Require XMLdoc2 entity validation with allowlistDisable entity expansion [Updated in Security Center 1.3]	• If the **glide.stax.whitelist_enabled** system property does not exist in the System Properties [sys_properties] table, or it is not set to the recommended value of **true**, then all external entities are allowed when the **glide.stax.allow_entity_resolution** system property is set to the value of **true**. If customizations do not require entity expansion, use the **glide.stax.allow_entity_resolution** system property to disable external entity expansion. The XML completes parsing but doesn't include any internal or external entities. ○ If you set **glide.stax.allow_entity_resolution** to **true**, all external entities attempt to resolve or expand subject entities, subject to the setting of the **glide.stax.whitelist_enabled** property. ○ If you set **glide.stax.allow_entity_resolution** to **false**, all entity resolution and expansion is blocked. To learn more about this property, see Disable Entity Expansion within the XMLDocument2 Streaming Parser [Updated in Security Center 1.5]. When **glide.stax.whitelist_enabled** is set to **true**, define a listing of comma-delimited FQDN in the **glide.xml.entity.whitelist** property, which are the only URLs that can be reached using the XML entity processing

Documentation	Updates
	property. To learn more, see Restrict XML external entities [Updated in Security Center 1.3 and 2.0]. Attackers can use this vulnerability to expand data exponentially in an External Entities Expansion (XXE) attack, quickly consuming all system resources. • Description (Old): If *glide.stax.whitelist_enabled* is not set to the recommended value of **true**, then all external entities are allowed. This could lead to External Entities Expansion (XXE) attacks. • Rule Script: Script has been updated to improve detection accuracy
Restrict XML external entities [Updated in Security Center 1.3 and 2.0]	• Protect against XXE attacks by using an allow list to prevent attackers from including arbitrary HTTP requests that the server may execute. This could lead to additional attacks using the server's trust relationship with other entities. Add `http://java.sun.com/j2ee/dtds/` to the value of the **glide.xml.entity.whitelist** system property, then set the **glide.xml.entity.whitelist.enabled** system property to **true**. Values other than `http://java.sun.com/j2ee/dtds/` can be included in the in the **glide.xml.entity.whitelist** property, but are unnecessary for the out of the box platform state. Review any additional values to determine if they are safe. ⚠ **Warning:** This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible. • Description (Old): If "glide.xml.entity.whitelist" is not set to the recommended value of "http://java.sun.com/j2ee/dtds/" and "glide.xml.entity.whitelist.enabled" is not set to "true", then a malicious external entity can be allowed which could cause an XXE attack. An attacker can use the DTD to include arbitrary HTTP requests that the server may execute. This could lead to

Documentation	Updates
	additional attacks using the server's trust relationship with other entities. • Rule Script: Script has been updated to improve detection accuracy
Restrict email domains for external user registration [Updated in Security Center 1.3, 1.5, and 2.0]	Rule Script: Script has been updated to improve detection accuracy
Validate file mime type in AttachmentCreator soap web service [New in Security Center 1.3 and updated in 1.5]	• Ensure that MIME-types are validated for attachments to prevent dangerous files from being uploaded on your instance using wrong file extensions. Set the **glide.attachment.enforce_security_validation** system property to **true**. When set to **true**, files are uploaded with the correct file type extension. • Description (Old): If "glide.attachment.enforce_security_validation" is not set to the recommended value of "true", then there will be no validation for attachment mime-type and dangerous files could be uploaded on the system using wrong file extensions. When this property is set to 'true', files are uploaded with the correct file type extension. It is a security best practice to validate file uploads at least with MIME type validation. • CVSS Score (New): 6.7 • CVSS Score (Old): 7.5
Define allowed ServiceNow internal IP addresses [Updated in Security Center 1.3 and 1.5]	• Description (New) Prevent unnecessary exposure of instance access to wider group of people using the **glide.ip.authenticate.strict** and **glide.ip.authenticate.allow.secured** system properties. When the **glide.ip.authenticate.strict** system property is set to **true**, internal ServiceNow personnel and systems can only make inbound connections to your instance from essential IP ranges. This limits ServiceNow's visibility to essential internal infrastructure on your instance, and prevents access by broader ServiceNow personnel such as support

Documentation	Updates
	and sales staff via corporate networks. The **glide.ip.authenticate.allow.secured** system property grants internal ServiceNow inbound connections, including regular authenticated access and unauthenticated diagnostic pages. If not set to **true**, then a broader ServiceNow internal IP range defined in the **glide.ip.authenticate.allow** property is used to grant these internal ServiceNow inbound connections. Ensure the **glide.ip.authenticate.allow.secured** system property contains only trusted values and that the property **glide.ip.authenticate.strict** is set to **true**. • Description (Old): If "glide.ip.authenticate.strict" is set to "true", then internal ServiceNow personelle and systems can only make inbound connections to the instance from essential IP ranges. This limit's ServiceNow's visibility into the instance to essential internal infrastructure, and prevents access by broader ServiceNow personelle such as support and sales staff via corporate networks. When set to "true", the "glide.ip.authenticate.allow" property is used to grant internal ServiceNow inbound connections. If not set to "true", then a broader ServiceNow internal IP range as defined in "glide.ip.authenticate.allow" is used to grant internal ServiceNow inbound connections.
Disable Entity Expansion within the XMLDocument2 Streaming Parser [Updated in Security Center 1.5]	• Description (New): Disable entity expansion on your instance to secure your instance from attacks such as ability to read system files, and Denial of Service. Use the system property to disallow XML entities to be expanded during parsing by the streaming parser (XMLDocument2). Set the **glide.stax.allow_entity_resolution** system property to **false** to disable entity expansion on your instance. If this property does not appear in the System Properties [sys_properties] table, the default value is

Documentation	Updates
	true. Create the property record and set the value to **false** to change it's value. • Description (Old): If "glide.stax.allow_entity_resolution" is not set to the recommended value of "False", then this property allow XML entities to be expanded during parsing by the streaming parser (XMLDocument2). XML entity expansion can lead to attacks such as ability to read system files, and Denial of Service. • Remediation (New): Ensure the property *`glide.stax.allow_entity_resolution`* exists in the sys_properties table and is set to **false**. If the property does not appear in the sys_properties list the default value is **true**. • Remediation (Old): Ensure the property "glide.stax.allow_entity_resolution" is set to "false".
Deny by default with empty ACLs [Updated in Security Center 1.3]	• Description (New): Prevent your instance's legacy security manager from allowing access to resources when there are no ACLs defined for that resource, or if there are only wildcard table-level ACLs (for example, `incident.*`). When allowed access by default, anything that does not have explicit ACLs set is susceptible to manipulation. Set the **glide.sm.default_mode** system property value to **deny** to disallow access when there are no define ACL rules, or there are only wildcard table-level ACLs. • Description (Old): If "glide.sm.default_mode" is not set to the recommended value of "deny", then the instance's legacy security manager allows access to a resource when there are no ACLs defined for that resource; or only wildcards table-level ACLs. By setting this to "allow" anything that does not have explicit ACLs set is susceptible to manipulation. • CVSS Score (New): 6.3 • CVSS Score (Old): 8.8 • Rule Script: Script has been updated to improve detection accuracy

Documentation	Updates
Restrict unauthenticated access to attachments	• Description (New): Secure the images on your instance to prevent sensitive information leak. Images on your instance are accessible via urls that end in `.iix`. Set the **glide.image_provider.security_enabled** system property to **true** to prevent access to your images via these URLs. **Note:** This property is not honored for images from the attachment table if the origin table is one of: ○ Stationeries [sysevent_email_style] ○ Welcome Page Sections [sys_home] ○ System Properties [sys_properties] Restriction should be applied for unauthenticated users as some attachments might contain sensitive information. • Description (Old): If "glide.image_provider.security_enabled" is not set to the recommended value of "True", then all images are accessible via urls that end in ".iix". This would allow unauthenticated access to images leading to sensitive information leak.
Disable embedded HTML code [Updated in Security Center 1.3]	• Description (New): Disable support for displaying HTML code embedded using the `[code]` tag. This tag allows rendered HTML to display in journal fields and may lead to cross-site scripting (XSS) attacks. These attacks can enable foreign scripts to execute on a user session in the logged in browser's context. Attackers can use these scripts to steal session information and sensitive data. The HTML language was not designed to separate script from formatting, so allowing user-controlled HTML in any system has inherent risk. Setting the **glide.ui.security.codetag.allow_script**

Documentation	Updates
	to **false** is compliant, and significantly reduces this risk, however some small risk remains. It disables only the script portion of a `code` tag, and relies on sanitizing all known conventions of script in the HTML. Set the **glide.ui.security.allow_codetag** system property to **false** to completely prohibit journal fields and forms from displaying rendered HTML. The ServiceNow AI Platform mitigates many injection and cross-site attacks by implementing escaping and encoding techniques. As a result, users can't write/submit HTML formatted inputs for journal fields. But journal fields can render text enclosed within code tags as HTML. ○ However, there is an associated security risk. If set to `true`, malicious users can write harmful HTML JS code that may be executed on a different client browser after rendering of journal fields. ○ Set this property to `false` so that administrators can prevent journal fields from rendering HTML code by disabling support for the `[code]` tag. • Description (Old): Disables support for embedding HTML code created using the [code] tag. If "glide.ui.security.allow_codetag" is not set to the value of "false", then rendered HTML in journal fields and forms will not be displayed. Setting "glide.ui.security.allow_codetag" to "true" with display embedded HTML code with may lead to cross-site scripting (XSS) attacks.
Enable password reset policy checks [Updated in Security Center 2.0]	Rule Script: Script has been updated to improve detection accuracy
Enable email spam scoring and filtering [Updated in Security Center 1.3]	• Technical Configuration Name (New): com.glide.email_filter,glide.email.read.active • Technical Configuration Name (Old): com.glide.email_filter • Rule Script: Script has been updated to improve detection accuracy

Documentation	Updates
Escape Excel Formulas [Updated in Security Center 1.3]	• Description (New): Prevent potentially malicious formulas in programs such as Excel from being executed after exporting and opening the file by escaping formulas in these files. Excel injection occurs when websites embed untrusted entries inside Excel files. When you use a spreadsheet application such as Microsoft Excel, or LibreOffice Call, to open a file, any cells starting with +, -, =, or @ are interpreted as a formula unless properly escaped. Malicious formulas pose a risk even when the spreadsheet doesn't contain any sensitive information, as they can be used to compromise the viewer's computer through code execution. Set the **glide.export.escape_formulas** system property to **true** to escape these formulas from executing. • Description (Old): Setting the property "glide.export.escape_formulas" to the recommended value of "true" prevents potentially malicious formulas in programs such as Excel from being executed after exporting and opening the file. Cell values for CSV, Xls, and XLSX can be interpreted as formulas by spreadsheet applications unless properly escaped which can lead to malicious code execution. • CVSS Score (New): 6.4 • CVSS Score (Old): 6.5
Restrict JSONP Requests to Trusted URLs [Updated in Security Center 1.3]	• Description (New): Increase security on your instance by ensuring that only trusted URLs for the AngularJS $http service can allow/reject JSONP requests. JSONP requests are allowed to any URL if these properties are not configured and enabled. Use the value of the **angular.jsonp.inclusion_list.urls** system property to define a list of URLs that are trusted and allow for this purpose. Set the value of the **angular.jsonp.inclusion_list.enabled** system property to **true** to limit allowed JSONP to only the URLs listed in **angular.jsonp.inclusion_list.urls**.

Documentation	Updates
	• Description (Old): This property specifies trusted URLs for the angularJS $http service to allow/reject JSONP requests. Property is necessary because this is a potentially breaking change for customers, so they need a way to add their trusted URLs. If "angular.jsonp.inclusion_list.enabled" is not set to the recommended value of "true", then JSONP requests are allowed to any URL.
Enable SNC access control plugin [Updated in Security Center 1.3]	• Description (New): Prevent ServiceNow Customer Service and Support personnel from accessing the instances without your express permission by enabling the SNC Access Control (com.snc.snc_access_control) plugin. Although all access to your instance is audited, you may prefer to control this access. This access method is fully auditable and tracked. **Note:** Other authorized ServiceNow Operations personnel, in their capacity to support and manage the product, are required to perform administrative actions on the underlying infrastructure. Enabling this plugin may affect support service levels and the Availability SLA. Availability SLA is then measured from the time that Support staff personnel are granted access to your instance. Enable the SNC Access Control (com.snc.snc_access_control) plugin to restrict access to your instance without your express permission. For more details on this feature, see ServiceNow access control. For activation information, see Activate ServiceNow access control • Description (Old): The SNC Access Control (com.snc.snc_access_control) plugin prevents Customer Service and Support personnel from accessing the instances without your express permission. However, other authorized ServiceNow Operations personnel, in their capacity to support and manage the product, are required to perform administrative actions on the underlying infrastructure. This infrastructure includes servers and databases, among other infrastructure components that make

Documentation	Updates
	up the SaaS solution. This access method is fully auditable and tracked.This plugin enables you to restrict access to your instance without your express permission, so it may affect support service levels and the Availability SLA. Availability SLA is then measured from the time that Support staff personnel are granted access to your instance. • Remediation (New): Ensure the plugin "com.snc.snc_access_control" is activated. Read the documentation on activating at https://www.servicenow.com/docs/csh?topicname=t_ActivateSNCAccessControl.html&version=lates • Remediation (Old): Ensure the plugin "com.snc.snc_access_control" is activated. • CVSS Score (New): 3.3 • CVSS Score (Old): 8.2
Maximize failed login unlock timeout duration [Updated in Security Center 1.3]	• Technical Configuration Name (New): glide.user.unlock_timeout_in_mins, sysevent_script_action • Technical Configuration Name (Old): glide.user.unlock_timeout_in_mins • Description (New): Help secure your instance against brute force attacks by defining a time period during which a user cannot attempt to log in after being locked out. The **glide.user.unlock_timeout_in_mins** system property unlocks the user account after the time period that is specified in it's value. If no value is specified, your instance unlocks the user account after the default period of 15 minutes. • Description (Old): If "glide.user.unlock_timeout_in_mins" is not set to the recommended value of "15", then it may be easier to brute force accounts in a faster timeframe. This property unlocks the user account after the time period that is specified for the glide.user.unlock_timeout_in_mins property. If no value is specified, then the system unlocks the user account after the default period of 15 minutes. • Remediation (New):

Documentation	Updates
	Set the **glide.user.unlock_timeout_in_mins** system property value to a minimum of **15**. If **glide.user.unlock_timeout_in_mins** does not exist, the default lockout time is set to 15 minutes. Ensure that the **SNC User Lockout Check with Auto Unlock** script action (found on the Script Action [sysevent_script_action] table) is present and active. The **SNC User Lockout Check with Auto Unlock** script action is installed with the High Security Settings (com.glide.high_security) plugin. • Remediation (Old): Ensure the property "glide.user.unlock_timeout_in_mins" is set to "15" or more. • Rule Script: Script has been updated to improve detection accuracy
Restrict access to specific IP ranges plugin [Updated in Security Center 1.3]	• Technical Configuration Name (New): com.snc.ipauthenticator,ip_access • Technical Configuration Name (Old): com.snc.ipauthenticator
Restrict access to emails with empty target table	• CVSS Score (New): 6.5 • CVSS Score (Old): 5.4
Restrict downloadable MIME types [Updated in Security Center 1.3 and 2.0]	• CVSS Score (New): 6.4 • CVSS Score (Old): 8
Enable updated version of MultiSSO plugin [Updated in Security Center 1.3 and 1.5]	• Technical Configuration Name (New): glide.authenticate.multissov2_feature.enabled • Technical Configuration Name (Old): glide.authenticate.multisso.enabled,glide.authenticate.multiss • Description (New): If the Multi SSO plugin is enabled on an instance, reduce security vulnerabilities by confirming that the v2 version is enabled. The latest version enhances security and has more features, such as Assertion encryption support, and IDP-initiated Single Logout (SLO). If the latest version is not enabled, the new security features cannot be used and the instance is at risk of using an plugin which is deprecated.

Documentation	Updates
	Follow the steps in KB0756504 to upgrade to the latest version. This process includes checking for and migrating any customization-related changes, then upgrading the version. When complete, the **glide.authenticate.multissov2_feature.enabled** system property is automatically set **true**. • Description (Old): If the Multi SSO plugin is enabled on an instance, the v2 version should be enabled. Versions prior to MultiSSOv2, including SAML 1.1 and SAML 2.0, do not follow best practice and use opensaml library versions with known CVEs. If the known CVEs were exploitable in outdated opensaml libraries, this could allow the attacker to forge messages and bypass authentication through XML Signature Wrapping attacks, impersonate entities or allow person-in-the-middle attackers to gain unauthorized access to the platform. • CVSS Score (New): 0 • CVSS Score (Old): 7.1
Prevent inactive users from logging in [New in Security Center 1.5]	Rule Script: Script has been updated to improve detection accuracy
Enable MID audit log [New in Security Center 1.3 and updated in 1.5]	Rule Script: Script has been updated to improve detection accuracy
Ensure archive table ACLs are checked [New in Security Center 1.3 and updated in 1.5]	Rule Script: Script has been updated to improve detection accuracy
Define active session timeout exception roles [New in Security Center 1.3]	• CVSS Score (New): 6.4 • CVSS Score (Old): 7.1
Limit HTTP response body size [New in Security Center 1.3 and updated in 1.5]	• Description (New): Prevent `OutOfMemoryExceptions` that can result from a request response body being too large using the **glide.http.response.get_body.limit.enabled** and **glide.http.response.get_body.limit** system properties. These exceptions can cause denial of service (DOS) attacks as well as other issues that may aid attackers in compromising an instance. Not setting these properties to the recommended values could make your instance vulnerable to OutOfMemoryExceptions and denial of service attacks.

Documentation	Updates
	To protect your instance against these security vulnerabilities: ○ Set the **glide.http.response.get_body.limit.enabled** system property to **true**. ○ Ensure that the **glide.http.response.get_body.limit** system property set to no more than 524,288,000 megabytes (500 MB). • Description (Old): The properties glide.http.response.get_body.limit.enabled and glide.http.response.get_body.limit were introduced to enable new functionality that prevents Out of memory exceptions being thrown as a result of a requests response body being too big. Out of memory exceptions can cause denial of service attacks as well as other issues that may aid attackers in compromising an instance. • CVSS Score (New): 3.1 • CVSS Score (Old): 6.4
Limit UI active session life span [New in Security Center 1.3]	• Description (New): Reduce the scope of potential security incidents by decreasing the lifespan of active HTTP sessions. The **glide.ui.active.session.life_span** system property enforces a maximum lifespan on active HTTP sessions irrespective of inactive timeout. Longer maximum lifespans can allow an attacker to use a stolen session a for longer time, increasing the scope of a security incident. The default value of **0** disables timeout of active sessions Set the **glide.ui.active.session.life_span** to a value between 1 and 720. This value represents the time in minutes that HTTP sessions can remain active. • Description (Old): This configuration will enforce max lifespan on active guest HTTP sessions irrespective of inactive timeout. The configured value is in minutes and the value of zero will disable timing out the active sessions. A larger maximum lifespan could allow an attacker to persist a stolen

Documentation	Updates
	session for longer, increasing the scope of a security incident. This particular property is limited to UI session timeout.

Updated hardening settings for baseline version 5.0

Some hardening settings have been updated with the release of Security Center baseline version 5.0.

Baseline version 5 includes several updates to short descriptions for style and consistency across records. In addition, many property related scripts were updated to improve the accuracy of default values in cases where the property has been removed from the sys_property table.

Documentation	Updates
Require authorization for SOAP requests [Updated in Security Center 1.3, 1.5, and 2.0]	• New remediation: Ensure the Glide Property *glide.basicauth.required.soap* exists and is set to the value true. Alternatively, configure the instance for WS Security by setting the property glide.soap.require_ws_security to true and following the product documentation to configure WS Security Profiles. If the property does not appear in the sys_properties table, add a new record. • Old remediation: Ensure the property *glide.basicauth.required.soap* is set to the value true. Alternatively, configure the instance for WS Security by setting the property *glide.soap.require_ws_security* to true and following the product documentation to configure WS Security Profiles.
Enforce OCSP check on network error [New in Security Center 1.3 and updated in 2.0]	• New remediation: Ensure the property *com.glide.communications.httpclient.ocsp_a* exists and is set to false. If the property does not appear in the sys_properties table, add a new record. • Old Remediation: Ensure the property *com.glide.communications.httpclient.ocsp_a* is set to false.
Disable external content url [Updated in Security Center 2.0]	• New remediation: Ensure the Glide Property *glide.ui.url.external.content* exists and is set to the value false. If the property does not appear in the sys_properties table, add a new record.

Documentation	Updates
	• Old Remediation: Ensure the property *glide.ui.url.external.content* is set to false. • New CVSS Score: 7.2 • Old CVSS Score: 8.1 • Rule Script: Script has been updated to improve detection accuracy.
Restrict XML external entities [Updated in Security Center 1.3 and 2.0]	• New remediation: Ensure the Glide Property *glide.xml.entity.whitelist* exists and is set to "http://java.sun.com/j2ee/dtds/" and the Glide Property *glide.xml.entity.whitelist.enabled* exists and is set to the value true. If the properties do not appear in the sys_properties table, add new records. • Old Remediation: Ensure the property *glide.xml.entity.whitelist* is set to "http://java.sun.com/j2ee/dtds/" and the property *glide.xml.entity.whitelist.enabled* is set to true.
Disable unauthenticated published reports [Updated in Security Center 2.0]	• New remediation: Ensure the Glide Property *glide.report.published_reports.enabled* exists and is set to the value false. If the property does not appear in the sys_properties table, add a new record. • Old Remediation: Ensure the property *glide.report.published_reports.enabled* is set to false.
Enable password reset policy checks [Updated in Security Center 2.0]	• New remediation: Ensure the Glide Property *glide.enable.password_policy* exists and is set to the value true. If the property does not appear in the sys_properties table, add a new record. • Old Remediation: Ensure the property *glide.enable.password_policy* is set to true.
Minimize Entity Expansion Threshold for GlideXMLUtil Scriptable [Updated in Security Center 1.3, 1.5, and 2.0]	• New remediation: Ensure the property *glide.xmlutil.max_entity_expansion* is set to 3000 or less. If the instance is on Washington or later, the default implied value is 3000 if the sys_properties record does not exist. If the instance is not on

Documentation	Updates
	Washington or later, the recommendaiton is for the instance admin to create a sys_properties record with name *glide.xmlutil.max_entity_expansion* and the value 3000. • Old Remediation: Ensure the property *glide.xmlutil.max_entity_expansion* is set to 3000 or less.
Disable outbound SSLv2/SSLv3 connections [Updated in Security Center 1.3]	• New remediation: Ensure the Glide Property *glide.outbound.sslv3.disabled* exists and is set to the value true. If the property does not appear in the sys_properties table, add a new record. • Old Remediation: Ensure the property *glide.outbound.sslv3.disabled* is set to true. **Important:** The value for the *glide.outbound.sslv3.disabled* property is a safe override and cannot be altered once changed.
Disable GlideRecord Scope Fencing Legacy Behavior [New in Security Center 1.3 and updated in 1.5 and 2.0]	• New short description: Disable GlideRecord Scope Fencing Legacy Behavior • Old short description: Enable GlideRecord Scope Fencing Legacy Behavior
Restrict uploaded MIME types [Updated in Security Center 1.3 and 2.0]	• New remediation: Ensure the property *glide.security.file.mime_type.validation* exists and is set to true. If the property does not appear in the sys_properties table, add a new record. • Old remediation: Ensure the property *glide.security.file.mime_type.validation* is set to true.
Enable Jelly JS interpolation protection for nested expressions [Updated in Security Center 2.0]	• New remediation: Ensure the Glide Property *glide.ui.jelly.js_interpolation.protect_ne* exists and is set to the value true. If the property does not appear in the sys_properties table, add a new record. • Old remediation: Ensure the property *glide.ui.jelly.js_interpolation.protect_ne* is set to true.
Enable SSL in LDAP authentication [Updated in Security Center 1.5 and 2.0]	Rule Script: Script has been updated to improve detection accuracy.

Documentation	Updates
Enable UserCookie version 3.1 [Updated in Security Center 2.0]	• New description: UserCookie v3 is generated only when property *glide.ui.secure.cookies.use_kmf is disabled*. UserCookie v3 is not secure due to storing secret key for HMAC in source code and identical for all customers. That can support malicious actors to use this one secret key for attempts to hijacking user sessions. By setting the property *glide.ui.secure.cookies.use_kmf* to true UserCookie v3.1 will be used and secret key will be stored in security storage such as KMF. • Old description: UserCookie v3 is generated only when property *glide.ui.secure.cookies.use_kmf* is disabled. UserCookie v3 is not secure due to storing secret key for HMAC in source code and identical for all customers. That can support malicious actors to use this one secret key for attempts to hijacking user sessions. • New remediation: Ensure the property *glide.ui.secure.cookies.use_kmf* exists and is set to true. If the property does not appear in the sys_properties table, add a new record. • Old remediation: Ensure the property *glide.ui.secure.cookies.use_kmf* is set to true. Which means UserCookie v3.1 will be used and secret key will be stored in security storage such as KMF.
Set OTP lifetime for password reset to 1 hour [Updated in Security Center 2.0]	Rule Script: Script has been updated to improve detection accuracy.
Log user impersonation [Updated in Security Center 1.3 and 2.0]	• New remediation: Ensure the property *glide.sys.log_impersonation* exists and is set to true. If the property does not appear in the sys_properties table, add a new record. • Old remediation: Ensure the property *glide.sys.log_impersonation* is set to true.
Required jms connection factories [New in Security Center 1.3 and updated in 1.5 and 2.0]	Rule Script: Script has been updated to improve detection accuracy.
Ensure dashboards creation/deletion requires access check [New in Security Center 1.3 and updated in 2.0]	• New remediation: Ensure the Glide Property *glide.processors.check_access_before_proce* exists and is set to the value true. If

Documentation	Updates
	the property does not appear in the sys_properties table, add a new record. • Old remediation: Ensure the value of *glide.processors.check_access_before_proce* is always true.
Proactively invalidate inactive sessions [New in Security Center 1.3 and updated in 1.5 and 2.0]	• New remediation: Ensure the Glide Property *glide.active.session.timeout.invalidate.se* exists and is set to the value true. If the property does not appear in the sys_properties table, add a new record. • Old remediation: Set the Glide Property *glide.active.session.timeout.invalidate.se* to true.
Enforce Security Scope for Agent Workspace for HR Case Management [New in Security Center 1.5 and updated in 2.0]	Rule Script: Script has been updated to improve detection accuracy.
Enforce security scope license and permit playbook [New in Security Center 1.5 and updated in 2.0]	Rule Script: Script has been updated to improve detection accuracy.
Restrict downloadable MIME types [Updated in Security Center 1.3 and 2.0]	• New description: If the property *glide.ui.attachment.force_download_all_mi* is set to true, then the *glide.ui.attachment.download_mime_types* property will be overridden so that all MIME types will be downloaded rather than rendered by the browser. For example, downloading text/html forces an HTML file to be downloaded to the client as a file rather than viewed inline in the browser, preventing a XSS attack. XSS can lead to easily attained privilege escalation to higher roles such as admin where more lateral movement can be taken. • Old description: If the property *glide.ui.attachment.force_download_all_mi* is not set to true, then the *glide.ui.attachment.download_mime_types* property will be overridden so that all MIME types will be downloaded rather than rendered by the browser. For example, downloading text/html forces an HTML file to be downloaded to the client as a file rather than viewed inline in the browser, preventing a XSS attack. The ability to have XSS can lead to easily attained privilege escalation to higher roles such as admin where more lateral movement can be taken.

Documentation	Updates
	• New remediation: Ensure the property *glide.ui.attachment.force_download_all_mi* is set to true. If the property does not exist in the sys_properties table, the default value is false. • Old remediation: Ensure the property *glide.ui.attachment.force_download_all_mi* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Define restricted downloadable MIME types [Updated in Security Center 1.3, 1.5, and 2.0]	Rule Script: Script has been updated to improve detection accuracy.
Disallow infected file download [Updated in Security Center 1.5 and 2.0]	• New description: When the property *com.glide.snap.infected_download_allowed* is set to true, users can still download non-scanned attachments in the case that the antivirus service is down or unreachable. This means it is possible that a user downloads a malicious file and risks infecting the user's desktop (in the case there is no other endpoint protection on the device). • Old description: If *com.glide.snap.infected_download_allowed* is not set to the recommended value of False, then it is possible to download a malicious file that has not been scanned leading to a risk of infecting the user's desktop. • New remediation: Ensure the property *com.glide.snap.infected_download_allowed* is set to false. • Old remediation: Ensure the property *com.glide.snap.infected_download_allowed* is set to False.
Restrict access to GlideSystemUserSession scriptable API [Updated in Security Center 1.3 and 2.0]	• New description: *gs.addErrorMessageNoSanitizationMessaging* and *gs.addInfoMessageNoSanitization()* are used within the scripting environment for logging and notifications. Both of these are available in the sandbox if this property is not set to the recommended value of false. The sandbox is a low privileged scripting environment available to unauthenticated and no role users. Both of these methods can be used to display unsanitized input to a user. Displaying

Documentation	Updates
	unsanitized input to the user is dangerous, as unsanitized input may contain dangerous code that runs in the user's browser. This can be utilized for traditional reflected XSS attacks. Reflected XSS attacks can be used in multiple scenarios, including session hijacking. • Old description: Messaging within the glide scripting sandbox is used for logging purposes. Calling this unsanitized error function exposes the platform to reflected XSS attacks. XSS attacks can allow for easy privilege escalation by stealing someone's session cookies. If `glide.sandbox.usersession.allow_unsanitize` is not set to the recommended value of false, then the unsanitized error messaging functions `addErrorMessageNoSanitization` and `addInfoMessageNoSanitization` are available to script.
Enable work order management query rules for service organizations [New in Security Center 1.5 and updated in 2.0]	• New description: When set to true, rules/ filters from sn_query_rule table will be used to determine read access to Field Service Management-related tables (Work Order and Work Order Task) to the logged in user through query business rules and read ACLs. When false, the records won't be filtered based on query rules. Query business rules add additional security validations. Specifically, this property will filter records for agents, qualifiers, and dispatchers based on their assigned territory or territory membership. It is best practice to follow the principle of least privilege when reading records. When this property is not set to true, there may be increased risk of data exposure from Field Service Management tables. • Old description: When set to true, rules/ filters from sn_query_rule table will be used to determine read access to Field Service Management-related tables (Work Order and Work Order Task) to the logged in user through query business rules and read ACLs. When false, the records won't be filtered based on query rules. Query business rules add additional security validations. Specifically, this property will filter records for agents, qualifiers, and dispatchers based on their assigned

Documentation	Updates
	territory or territory membership. It is best practice to follow the principle of least privilege when reading records.
Restrict email domains for external user registration [Updated in Security Center 1.3, 1.5, and 2.0]	• New description: The *sn_ext_usr_reg.allowed_email_domains* property defines which email addresses are allowed to self-register to a ServiceNow instance. The format should be a comma separated list of acceptable email domains such as domain1.com,domain2.com where emails such as example@domain2.com will be accepted. If *sn_ext_usr_reg.allowed_email_domains* is not set with a list of acceptable domains, then users with any email address are allowed to register accounts on the instances. If not defined, malicious actors could perform registration using emails addresses from unwanted domains to gain authenticated access to the instance. • Old description: The *sn_ext_usr_reg.allowed_email_domains* property defines which email addresses are allowed to self-register to a ServiceNow instance. If *sn_ext_usr_reg.allowed_email_domains* is not set with a list of acceptable domains, then users with any email address are allowed to register accounts on the instances. If not defined, malicious actors could perform registration using emails addresses from unwanted domains to gain authenticated access to the instance.
Apply domain separation on dot walked fields [Updated in Security Center 1.3, 1.5, and 2.0]	• New description: This property controls whether join queries are given domain separated conditions or not, in order to ensure they apply domain separation functionality for dot walked fields. If *glide.sys.domain.include_domain_condition_* is not set to the recommended value of true on an instance using domain separation, then sensitive information could be disclosed that is not to be shared with a specific domain. There may be moderate functional impact to the instance if components are reliant on the unsafe cross domain queries. Instances should be tested in subproduction environments before enabling.

Documentation	Updates
	• Old description: This property controls whether join queries are given domain separated conditions or not, in order to ensure they apply domain separation functionality for dot walked fields. If *glide.sys.domain.include_domain_condition_* is not set to the recommended value of true on an instance using domain separation, then sensitive information could be disclosed that is not to be shared with a specific domain.
Enforce URL allowlist check [Updated in Security Center 1.3, 1.5, and 2.0]	• New remediation: Ensure the property *glide.security.url.whitelist.strict_check* is set to true or the property *glide.security.url.whitelist* is set to a value. • Old remediation: Ensure the property *glide.security.url.whitelist.strict_check* is set to "true" and the property *glide.security.url.whitelist* is set to a value.
Set guest user for soap requests [Updated in Security Center 1.3 and 2.0]	Rule Script: Script has been updated to improve detection accuracy.
Restrict access to background script [Updated in Security Center 1.3 and 2.0]	• New description: This property holds the required role to access Script Background module. If *glide.script_processor.admin* is not set to the recommended and default value of admin, then users having a lower privileged role will be able to run background scripts on the instance. This will lead to a complete bypass of the ACL system allowing full access to tables. • Old description: This property holds the required role to access Script Background module. If *glide.script_processor.admin* is not set to the recommended value of admin, security_admin, or maint, then users having a lower privileged role will be able to run background scripts on the instance. This will lead to a complete bypass of the ACL system allowing full access to tables. • New remediation: Ensure the property *glide.script_processor.admin* is set to the admin. This is the default value on instances.

Documentation	Updates
	• Old remediation: Ensure the property *glide.script_processor.admin* is set to the admin, security_admin, or maint role.
Verify certificate chain and hostname [New in Security Center 1.3 and updated in 2.0]	• New description: When the Glide Property *com.glide.communications.httpclient.verify* is not set to the secure value of true, the hostname and certificate chain presented by remote hosts during a TLS connection initiated from the ServiceNow instance are not validated. This could compromise the security of the TLS connection and allow person-in-the-middle attacks, where communications between two parties are intercepted. This may lead to sensitive data disclosure. • Old description: If *com.glide.communications.httpclient.verify* is not set to true this could allow person-in-the-middle attacks where communications between two parties are intercepted. Setting this property to an insecure value disables the certificate verification process which evaluates all certifications in the certificate chain through checking revocation status. Set this property to true to prevent the http client from connecting to a potentially harmful hostname.
Control Lockout Time for Invalid Password Reset Attempts [Updated in Security Center 1.3 and 2.0]	• New short description: *Control Lockout Time for Invalid Password Reset Attempts* • Old short description: *Minimize Reset Password Request Max Attempts Window Duration* • New description: The *password_reset.request.max_attempt_window* property defines the number of minutes a user must wait to reset or change their password after exceeding the maximum number of unsuccessful attempts that is set with the *password_reset.request.max_attempt property*. A small number of minutes for the *password_reset.request.max_attempt_window* property increases the risk of successfully brute forcing a password as a greater number of password reset attempts can

Documentation	Updates
	be made. The default of 1440 minutes is recommended. • Old description: If *password_reset.request.max_attempt_window* is not set to the recommended value of 1440 or less, then it could be possible to perform account bruteforce as the account will not be locked after a maximum number of wrong authentication attempts. • New remediation: Ensure the property *password_reset.request.max_attempt_window* is set to 1440 or greater. • Old remediation: Ensure the property *password_reset.request.max_attempt_window* is set to 1440 or less. • Rule Script: Script has been updated to improve detection accuracy.
Disable GlideRecord Scope Fencing Legacy Behavior [New in Security Center 1.3 and updated in 1.5 and 2.0]	• New short description: *Disable GlideRecord Scope Fencing Legacy Behavior* • Old short description: *Enable GlideRecord Scope Fencing Legacy Behavior* • New remediation: Set the Glide Property *glide.record.legacy_cross_scope_access_pol* to false. When not present in the sys_properties table, the default value is true. • Old remediation: Set the Glide Property *glide.record.legacy_cross_scope_access_pol* to false.
Limit Invalid Password Reset Attempts [Updated in Security Center 1.3 and updated in 2.0]	• New short description: *Limit Invalid Password Reset Attempts* • Old short description: *Minimize Reset Password Request Max Attempt Allowance*

Updated hardening settings for baseline version 4.0

Some hardening settings have been updated with the release of Security Center baseline version 4.0.

Baseline version 4.0 includes several updates to short descriptions for style and consistency between records. In addition, many property related scripts were also updated to improve accuracy of the default value for cases where the property has been removed from the sys_property table.

Documentation	Updates
Require authorization for SOAP requests [Updated in Security Center 1.3, 1.5, and 2.0]	• New technical configuration name: *glide.basicauth.required.soap, glide.soap.require_ws_security* • Old technical configuration name: *glide.basicauth.required.soap* • New description: The glide property *glide.basicauth.required.soap* controls whether basic authentication is required to make a SOAP request to an instance. If *glide.basicauth.required.soap* is not set to the recommended value of true, then unauthenticated users performing SOAP operations will be mapped to the soap.guest user. This may enable an unauthenticated user to perform operations on the instance as if a logged in user to the instance. There may be additional impact if the user define within *com.glide.soap.guest_user* is assigned additional roles. • Old description: The glide property *glide.basicauth.required.soap* controls whether authentication is required to make a SOAP request to an instance. If *glide.basicauth.required.soap* is not set to the recommended value of true, then authentication is disable for SOAP requests on the instance. It allows unauthenticated access to administrator or maint level operations; thereby negating security controls within the instance. • New remediation: Ensure the property *glide.basicauth.required.soap* is set to the value true. Alternatively, configure the instance for WS Security by setting the property *glide.soap.require_ws_security* to true and following the product documentation to configure WS Security Profiles. • Old remediation: Ensure the property *glide.basicauth.required.soap* exists in the sys_properties table and is set to true. • Rule Script: Script has been updated to improve detection accuracy.

Documentation	Updates
Escape jelly script [Updated in Security Center 1.3 and 1.5]	• New description: This property escapes all the JS and HTML strings included in before they are written to the output stream, preventing several XSS issues from occurring. If *glide.ui.escape_all_script* is not set to the recommended value of true, then escaping of scripts injected into Jelly is disabled. Without this mitigation, the platform becomes widely open to a variety of script injection attacks. An attacker could execute arbitrary Rhino scripts on the instance. • Old description: The following property escapes all the JS and HTML strings included in <j:jelly> ... </j:jelly> before they are written to the output stream, preventing several XSS issues from occurring. If *glide.ui.escape_all_script* is not set to the recommended value of "true", then escaping of scripts injected into Jelly is disabled. Without this mitigation, the platform becomes widely open to a variety of script injection attacks. An attacker could execute arbitrary Rhino scripts on the instance.
Prevent Users From Accepting Warning To Bypass CSRF Validation [Updated in Security Center 1.3 and 1.5]	• New short description: Prevent Users From Accepting Warning To Bypass CSRF Validation • Old short description: Enforce CSRF Token Strict Validation • New description: This property prevents users from being able to accept a warning which allows a potentially malicious request to be sent to the instance. This warning appears when a POST request fails due to having a mis-matched anti-CSRF token belonging to one of the victim's other active sessions. If *glide.security.csrf.strict.validation.mode* is not set to the recommended value of true, then an attacker can formulate a CSRF attack utilizing a leaked anti-CSRF token from a different active session belonging to the victim.A POST request to an instance contains an anti-CSRF token within sysparm_ck or X-UserToken which matches the user's current session. If the anti-CSRF token is instead tied to one of the user's other active sessions, the POST request will return a 302 redirection to

Documentation	Updates
	security_interceptor.do with a Continue button available to the user when this property is set to false. Clicking this button will re-submit the request to the instance, except it will now having a valid anti-CSRF token. When this property is set to true, the 302 redirection to the security_interceptor.do page will not display a Continue button and the user will not be allowed to resubmit the request.A successful CSRF attack will allow an attacker to effectively perform any operation that the victim is able to perform. • Old description: This property enables CSRF token strict validation which prevents the reuse of CSRF tokens. If *glide.security.csrf.strict.validation.mode* is not set to the recommended value of true, then CSRF tokens could be reused which opens a door to CSRF attacks. • New CVSS Score: 3.7 • Old CVSS Score: 3.1
Require Authentication on Event Management HTTP Processor [New in Security Center 1.3, Updated in 1.5, and removed in 2.0]	• New short description: Require Authentication on Event Management HTTP Processor • Old short description: Require Authentication on Event Management HTTP Processor
Enable Anti-CSRF token [New in Security Center 1.3, updated in 1.5, and removed in 2.0]	• New description: Cross-Site Request Forgery (CSRF) is an attack that forces authenticated users to submit a request to a Web application against which they are currently authenticated. CSRF attacks exploit the trust a Web application has in an authenticated user. This property enables usage of a secure token to identify and validate incoming requests. This token is used to prevent cross site request forgery attacks. If *glide.security.use_csrf_token* is not set to the recommended value of true, then CSRF is possible. • Old description: Cross-Site Request Forgery (CSRF) is an attack that forces authenticated users to submit a request to a Web application against which they are currently authenticated. CSRF attacks exploit the trust a Web application has in an authenticated user. This property enables usage of

Documentation	Updates
	a secure token to identify and validate incoming requests. This token is used to prevent cross site request forgery attacks. If *glide.security.use_csrf_token* is not set to the recommended value of true, then CSRF is possible.
Enable HTML Sanitizer within Virtual Agent [Updated in Security Center 1.3 and 1.5]	• New short description: Enable HTML Sanitizer within Virtual Agent • Old short description: Enable HTML Sanitizer • New description: This property controls the whether the HtmlSanitizerService is enabled. If *com.glide.cs.html.sanitizer.enabled* is not set to true, then a Stored Cross-Site Scripting (XSS) attack is possible in the VA web client. • Old description: This property controls the whether the HTMLSanitezerService is enabled. If *com.glide.cs.html.sanitizer.enabled* is not set to true, then a Stored Cross-Site Scripting (XSS) attack is possible in the VA web client.
Deny internal access to explicit external roles [Updated in Security Center 1.3 and 1.5]	• New technical configuration name: *glide.security.explicit_roles.enable_intei* • Old technical configuration name: *glide.security.explicit_roles.enable_intei* • New description: This property prevents external users from being assigned the snc_internal role. When *glide.security.explicit_roles.enable_intei* is set to the recommended value of true, then it enforces the parameters of the maint-protected *glide.security.explicit_roles.internal_use property* which assigns the snc_external role to a list of untrusted user classes. If the value is set to false, the *glide.security.explicit_roles.internal_use* property is ignored. Misconfiguration of this propery increases the risk that an external user account gains access to internal information. • Old description: This prevents external users from being assigned the snc_internal role. If *glide.security.explicit_roles.enable_intei*

Documentation	Updates
	is not set to the recommended value of true, and the *glide.security.explicit_roles.internal_use* property is not set to a list of untrusted user classes, then the specified roles can be assigned the snc_internal role instead of the snc_external role. If the list is empty, then all users will be assigned the snc_internal role by default. The property should contain at least the default roles csm_consumer_user, customer_contact. Misconfiguration of these properties increases the risk that an external user account gains access to internal information. • New remediation: Ensure the property *glide.security.explicit_roles.enable_inte* is set to true. • Old Remediation: Ensure the property *glide.security.explicit_roles.enable_inte* is set to true and that the property *glide.security.explicit_roles.internal_use* includes the dangerous items csm_consumer_user, customer_contact. • Rule Script: Script has been updated to improve detection accuracy.
Require authorization for WSDL request [Updated in Security Center 1.3 and 1.5]	• New description: If *glide.basicauth.required.wsdl* is not set to the recommended value of true, then this will disable Basic Authentication for WSDL requests. WSDL is a protocol that is used to describe web services such as instance table schemas, and is not a mechanism for sharing the data within tables. Setting this property to true allows for disclosure of table schemas to unauthenticated users. • Old description: If *glide.basicauth.required.wsdl* is not set to the recommended value of true, then this will disable Basic Authentication for WSDL requests. This could lead to information disclosure to unauthenticated users. • New CVSS Score: 5.3 • Old CVSS Score: 4.3

Documentation	Updates
Enforce URL allowlist check [Updated in Security Center 1.3, 1.5, and 2.0]	Rule Script: Script has been updated to improve detection accuracy.
Define restricted downloadable MIME types [Updated in Security Center 1.3, 1.5, and 2.0]	• New short description: Define Restricted Downloadable MIME Types • Old short description: Restrict Downloadable MIME Types • New description: If *glide.ui.attachment.download_mime_types* does include dangerous items such as text/html,image/svg,image/svg+xml,application/xml, then dangerous files could be rendered inline in the browser which could lead to Cross Sitte Scripting attacks (XSS). This property is the list of comma separated attachment mime types which will not render inline in the browser. For example, including text/html will force HTML files to be downloaded to the client as attachments rather than viewed inline in the browser. Maintaining this list properly will prevent cross site scripting attacks. • Old description: If *glide.ui.attachment.download_mime_types* does include dangerous items such as text/html,image/svg,image/svg+xml,application/xml, then dangerous files could be rendered inline in the browser which could lead to Cross Sitte Scripting attacks (XSS). This property is the list of comma separated attachment mime types which will not render inline in the browser. For example, including text/html will force html files to be downloaded to the client as attachments rather than viewed inline in the browser. Maintaining this list properly will prevent cross site scripting attacks.
Escape HTML in list views [Updated in Security Center 1.3 and 1.5]	• New description: This property helps sanitize list view displaying of HTML fields. If *glide.ui.escape_html_list_field* is not set to the recommended value of true, then a malicious user can inject HTML code within the form field to execute unwanted scripts on different client/user sessions. This could potentially be leveraged by attackers to steal session information and sensitive data. • Old description: The following property helps sanitize list view displaying of HTML fields. If *glide.ui.escape_html_list_field*

Documentation	Updates
	is not set to the recommended value of true, then a malicious user can inject HTML code within the form field to execute unwanted scripts on different client/user sessions. This could potentially be leveraged by attackers to steal session information and sensitive data.
Restrict email domains for external user registration [Updated in Security Center 1.3, 1.5, and 2.0]	• New short description: Restrict Email Domains for External User Registration • Old short description: Restrict Email Domains for External User Registration (Plugin Applicability: External User Registration) • New Description: The *sn_ext_usr_reg.allowed_email_domains* property defines which email addresses are allowed to self-register to a ServiceNow instance. If *sn_ext_usr_reg.allowed_email_domains* is not set with a list of acceptable domains, then users with any email address are allowed to register accounts on the instances. If not defined, malicious actors could perform registration using emails addresses from unwanted domains to gain authenticated access to the instance. • Old Description: If *sn_ext_usr_reg.allowed_email_domains* is not set with a whitelist of acceptable domains, then malicious actors could perform registration using emails addresses from unwanted domains. • Rule Script: Script has been updated to improve detection accuracy.
Enable Captcha for External User Registration [Updated in Security Center 1.3 and 1.5]	• New short description: Enable Captcha for External User Registration • Old short description: Enable Captcha for External User Registration (Plugin Applicability: External User Registration) • Rule Script: Script has been updated to improve detection accuracy
Minimize external user registration link expiration duration [Updated in Security Center 1.3 and 1.5]	• New short description: Minimize External User Registration Link Expiration Duration • Old short description: Minimize External User Registration Link Expiration Duration

Documentation	Updates
	(Plugin Applicability: External User Registration) • Rule Script: Script has been updated to improve detection accuracy
Disallow infected file download [Updated in Security Center 1.5 and 2.0]	• New short description: Disallow Infected File Download • Old short description: Disallow Infected Files Download • New remediation: Ensure the property *com.glide.snap.infected_download_allowed* is set to False. • Old Remediation: Ensure the property *com.glide.snap.infected_download_allowed* is set to True. • Rule Script: Script has been updated to improve detection accuracy.
Validate file mime type in AttachmentCreator soap web service [New in Security Center 1.3 and updated in 1.5]	• New description: If *glide.attachment.enforce_security_validati* is not set to the recommended value of true, then there will be no validation for attachment mime-type and dangerous files could be uploaded on the system using wrong file extensions. When this property is set to true, files are uploaded with the correct file type extension. It is a security best practice to validate file uploads at least with MIME type validation. • Old description: If *glide.attachment.enforce_security_validati* is not set to the recommended value of True, then there will be no validation for attachment mime-type and dangerous files could be uploaded on the system using wrong file extensions. When this property is set to true, files are uploaded with the correct file type extension. It is a security best practice to validate file uploads at least with MIME type validation. • New remediation: Ensure the property *glide.attachment.enforce_security_validati* is set to true. • Old Remediation: Ensure the property *glide.attachment.enforce_security_validati* is set to True.

Documentation	Updates
Disable MultiSSO Debugging [Updated in Security Center 1.3 and 1.5]	• New short description: Disable MultiSSO Debugging • Old short description: Disable MultiSSO Debugging (Plugin Applicability: Multiple Provider Single Sign-On)
Define allowed ServiceNow internal IP addresses [Updated in Security Center 1.3 and 1.5]	• New technical configuration name: *glide.ip.authenticate.strict* • Old technical configuration name: *glide.ip.authenticate.strict,glide.ip.auth* • New description: If *glide.ip.authenticate.strict* is set to true, then internal ServiceNow personelle and systems can only make inbound connections to the instance from essential IP ranges. This limit's ServiceNow's visibility into the instance to essential internal infrastructure, and prevents access by broader ServiceNow personelle such as support and sales staff via corporate networks. When set to "true", the *glide.ip.authenticate.allow* property is used to grant internal ServiceNow inbound connections. If not set to true, then a broader ServiceNow internal IP range as defined in *glide.ip.authenticate.allow* is used to grant internal ServiceNow inbound connections. • Old description: If *glide.ip.authenticate.strict* is set to true, then only IP ranges specified in *glide.ip.authenticate.allow.secured* can make inbound connections to the instance. This property contains a list of only essential ServiceNow internal IP ranges (Secure VPN, DC). If *glide.ip.authenticate.allow.secured* is not set to the recommended value or permutation of "10.0.0.0/8, 37.98.232.0/21, 103.23.64.0/22, 149.96.0.0/17, 149.96.0.0/16, 199.91.136.0/21, 148.139.0.0/16, 127.0.0.1" or the newer value list "10.0.0.0/8, 37.98.232.0/21, 103.23.64.0/22, 149.96.0.0/17, 149.96.0.0/16, 199.91.136.0/21, 148.139.0.0/16, 127.0.0.1, 0:0:0:0:0:0:0:1, ::1" which adds IPv6 localhost to Utah, then it may allow untrusted sources outside of SN DataCenter and secure VPN to access sensitive monitoring endpoints on instances.

Documentation	Updates
	• New remediation: Ensure the property *glide.ip.authenticate.allow.secured* contains only trusted values and that the property *glide.ip.authenticate.strict* is set to true. • Old remediation: Ensure the property *glide.ip.authenticate.allow.secured* contains only values in "10.0.0.0/8, 37.98.232.0/21, 103.23.64.0/22, 149.96.0.0/17, 149.96.0.0/16, 199.91.136.0/21, 148.139.0.0/16, 127.0.0.1, 0:0:0:0:0:0:0:1, ::1" and that the property *glide.ip.authenticate.strict* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Disable Entity Expansion within the XMLDocument2 Streaming Parser [Updated in Security Center 1.5]	• New short description: Disable Entity Expansion within the XMLDocument2 Streaming Parser • Old short description: Disable Entity Expansion • Rule Script: Script has been updated to improve detection accuracy.
Apply domain separation on dot walked fields [Updated in Security Center 1.3, 1.5, and 2.0]	• New short description: Apply Domain Separation on Dot Walked Fields • Old short description: Apply Domain Separation on Dot Walked Fields (Plugin Applicability: Domain Separation) • Rule Script: Script has been updated to improve detection accuracy.
Restrict permissions for CMDB model [Updated in Security Center 1.3 and 1.5]	Rule Script: Script has been updated to improve detection accuracy.
Require clearing pasteboard when backgrounding mobile application [New in Security Center 1.3 and updated in 1.5]	• New description: The *glide.sg.clear_pasteboard_when_backgrounde* property controls if text copied from ServiceNow mobile app is kept in the clipboard and pasteboard after the app is in background mode. If it is not set to the recommended value of true, then sensitive information may be disclosed to the Android or iOS clipboard where it can be exposed to other applications on the device. • Old description: The property *glide.sg.clear_pasteboard_when_backgrounde*

Documentation	Updates
	controls if text copied from ServiceNow mobile app is kept in the clipboard/ pasteboard after the app is no longer in focus. If it is not set to the recommended value of true, then sensitive information may be disclosed to the Android or iOS clipboard where it can be exposed to other applications on the device.
Enable account recovery [Updated in Security Center 1.3 and 1.5]	• New short description: Enable Account Recovery • Old short description: Enable Account Recovery (Plugin Applicability: Multiple Provider Single Sign-On)
Disable SQL Error Messages [Updated in Security Center 1.3 and 1.5]	• New description: If *glide.db.loguser* is not set to the recommended value of false, then sensitive server-side error messages could be displayed to end-users. Error messages can include stack traces and information about the structure of the database that could provide an attacker the knowledge needed to perform successful SQL Injection should the preconditions exist. As defense in depth, these error messages should not be displayed to the end user. • Old description: If *glide.db.loguser* is not set to the recommended value of false, then sensitive server-side error messages could be displayed to end-users.
Enforce relative links [Updated in Security Center 1.3 and 1.5]	• New description: The *glide.cms.catalog_uri_relative* property enforces relative links from the URI parameter on /ess/catalog.do. If *glide.cms.catalog_uri_relative* is not set to the recommended value of true, then the URL will not be sanitized with the enforceRelativeURL(url) function. Absolute URLs can pose a security risk when used as a part of parameter or a field value, thus redirecting the source page to an adversary-controlled website. This property impacts the legacy Content Management System (CMS) which has been replaced with Service Portal. • Old description: The *glide.cms.catalog_uri_relative* property enforces relative links from the

Documentation	Updates
	URI parameter on /ess/catalog.do. If *glide.cms.catalog_uri_relative* is not set to the recommended value of true, then the URL will not be sanitized with the enforceRelativeURL(url) function. Absolute URLs can pose a security risk when used as a part of parameter or a field value, thus redirecting the source page to an adversary-controlled website.
Minimize Entity Expansion Threshold for GlideXMLUtil Scriptable [Updated in Security Center 1.3, 1.5, and 2.0]	• New short description: Minimize Entity Expansion Threshold for GlideXMLUtil Scriptable • Old short description: Minimize Entity Expansion Threshold • New description: This property controls the maximum amount of entity expansion within an XML Parser. If *glide.xmlutil.max_entity_expansion* is not set to the recommended value of 3000 or less, then the GlideXMLUtil parsing scriptable may be vulnerable to denial of service attacks. • Old description: This property controls the maximum amount of entity expansion within an XML Parser. If *glide.xmlutil.max_entity_expansion* is not set to the recommended value of 3000 or less, then XML parser may be vulnerable to denial of service attacks. • Rule Script: Script has been updated to improve detection accuracy.
Disable GlideRecord Scope Fencing Legacy Behavior [New in Security Center 1.3 and updated in 1.5 and 2.0]	• New description: GlideRecord provided cross scope create/update access to tables that were not configured with that level of access. In order to prevent customers from having applications broken when this scoped access behavior was patched, the property *glide.record.legacy_cross_scope_access_pol* was created. When true, cross scope access falls back onto legacy behavior (insecure). This property disables scope fencing, allowing scoped apps to access global script interfaces. It is best security practice to have scope fencing restrictions in place. Scoping ensures applications can only access resources with explicit access or within their scope, following the principle of least privilege. Disabiling this feature

Documentation	Updates
	could lead to confidentiality, availability, and integrity impacts. • Old description: Legacy behavior provided create/update access to tables that did not allow so. In order to prevent legacy customers from having applications broken when this scoped access behavior was patched, the property *glide.record.legacy_cross_scope_access_pol* was created. When true, cross scope access falls back onto legacy behavior (insecure). This property disables scope fencing, allowing scoped apps to access global script interfaces. It is best security practice to have scope fencing restrictions in place. Scoping ensures applications can only access resources with explicit access or within their scope, following the principle of least privilege. Disabiling this feature could lead to confidentiality, availability, and integrity impacts.
Enable updated version of MultiSSO plugin [Updated in Security Center 1.3 and 1.5]	• New short description: Enable Updated Version of Multi SSO Plugin • Old short description: Enable Updated Version of Multi SSO Plugin (Plugin Applicability: Multiple Provider Single Sign-On)
Enable SSL in LDAP authentication [Updated in Security Center 1.5 and 2.0]	Script has been updated to improve detection accuracy.
Enforce password reset on api requests [Updated in Security Center 1.5]	Script has been updated to improve detection accuracy.
Do not apply password policy at login [Updated in Security Center 1.5 and removed in 2.0]	• New description: By setting the property *glide.apply.password_policy.on_login* to False there will be no password complexity enforcement at login time. Setting the property to True will enforce password complexity and lead to organization policy compliance issues. As per ASVS 4.03 v2.1.9 recommendations : Verify that there are no password composition rules limiting the type of characters permitted. There should be no requirement for upper or lower case or numbers or special characters. (C6) Instead of password complexity enforcement, ASVS recommendations are to

Documentation	Updates
	enforce a minimum length of 12 characters for password length. Ref: OWASP ASVS v4.0 Authentication • Old description: By setting the property *glide.apply.password_policy.on_login* to False there will be no password complexity enforcement at login time. Setting the property to True will enforce password complexity and lead to organisation policy compliance issues. As per ASVS 4.03 v2.1.9 recommendations : Verify that there are no password composition rules limiting the type of characters permitted. There should be no requirement for upper or lower case or numbers or special characters. (C6) Instead of password complexity enforcement, ASVS recommendations are to enforce a minimum length of 12 characters for password length. Ref: OWASP ASVS v4.0 Authentication
Do not use demo certificates for active saml configurations [Updated in Security Center 1.5]	• New short description: Do Not Use Demo Certificates for Active SAML Configurations • Old short description: Do Not Use Demo Certificates for Active SAML Configurations (Plugin Applicability: Multiple Provider Single Sign-On)
Minimize SAML notBefore or notOnOrAfter constraint duration [Updated in Security Center 1.3 and 1.5]	• New short description: Minimize SAML "notBefore" or "notOnOrAfter" Constraint Duration • Old short description: Minimize SAML "notBefore" or "notOnOrAfter" Constraint Duration (Plugin Applicability: Multiple Provider Single Sign-On)
Block Expired Anti-CSRF Tokens [Updated in Security Center 1.5]	• New short description: Block Expired Anti-CSRF Tokens • Old short description: Block Expired CSRF Tokens

Documentation	Updates
Require captcha for guest walk-up experience in customer service application [New in Security Center 1.3 and updated in 1.5]	• New short description: Require Captcha for Guest Walk-up Experience in Customer Service Application • Old short description: Require Captcha for Guest Walk-up Experience in Customer Service Application (Plugin Applicability: Guest Walk-up Experience for Customer Service)
Check impersonation on ACL evaluation in HR App [New in Security Center 1.3 and updated in 1.5]	• New short description: Check Impersonation on ACL Evaluation in HR App • Old short description: Check Impersonation on ACL Evaluation in HR App (Plugin Applicability: Human Resources Scoped App)
Restrict HR case updates from personal emails [New in Security Center 1.3 and updated in 1.5]	• New short description: Restrict HR Case Updates from Personal Emails • Old short description: Restrict HR Case Updates from Personal Emails (Plugin Applicability: Human Resources Scoped App) • Rule Script: Script has been updated to improve detection accuracy.
Enable MID audit log [New in Security Center 1.3 and updated in 1.5]	• New short description: Enable MID Audit Log • Old short description: Enable MID Audit Log (Plugin Applicability: MID Server)
	• New short description: Enforce Credential Alias Usage • Old short description: Enforce Credential Alias Usage (Plugin Applicability: MID Server)
Required jms connection factories [New in Security Center 1.3 and updated in 1.5 and 2.0]	• New short description: Required JMS Connection Factories • Old short description: Required JMS Connection Factories (Plugin Applicability: MID Server) • Rule Script: Script has been updated to improve detection accuracy.

Documentation	Updates
Limit attachment size in training and prediction flows [New in Security Center 1.3 and updated in 1.5]	• New short description: Limit Attachment Size in Training and Prediction Flows • Old short description: Limit Attachment Size in Training and Prediction Flows (Plugin Applicability: Platform Document Intelligence)
Ensure archive table ACLs are checked [New in Security Center 1.3 and updated in 1.5]	Rule Script: Script has been updated to improve detection accuracy.
Log session audit events [New in Security Center 1.3 and updated in 1.5]	• New description: When the Glide Property *glide.authenticate.session_access.log_audi* is set to true, session audit events will be created in the sys_session_access_audit table. It is best practice to log information about who accessed a session to assist in malicious actor investigations. Information logged will include user, session ID (non-sensitive), IP address, roles, and policies. • Old description: When the Glide Property *glide.authenticate.session_access.log_audi* is set to true, session audit events will be created in the sys_session_access_audit table. It is best practice to log general information about session access to assist in malicious actor investigations. Information logged will include user, session ID (non-sensitive), IP address, roles, and policies.
Enforce scoped ACL access for information request playbooks [New in Security Center 1.3 and updated in 1.5]	• New short description: Enforce Scoped ACL Access for Information Request Playbooks • Old short description: Enforce Scoped ACL Access for Information Request Playbooks • Rule Script: Script has been updated to improve detection accuracy.
Proactively invalidate inactive sessions [New in Security Center 1.3 and updated in 1.5 and 2.0]	• New description: The Glide Property *glide.active.session.timeout.invalidate.se* controls if a timed out session is proactively invalidated before the Tomcat container invalidates the session. When this property is not set to true, there can be a small interval of time where a timed out session is not invalidated (60+ seconds, depending on queue size). If a session is hijacked, an attacker may be able to utilize a session during this small period of time. • Old description: The Glide Property *glide.active.session.timeout.invalidate.se*

Documentation	Updates
	controls if a timeout session is proactively invalidated before the Tomcat container. When this property is not set to true, there can be a small interval of time where a timed out session is not invalidated (60+ seconds, depending on queue size). If a session is hijacked, an attacker may be able to utilize a session during this small period of time.
Limit HTTP response body size [New in Security Center 1.3 and updated in 1.5]	• New short description: Limit HTTP Response Body Size • Old short description: Ensure HTTP Responses Do Not Trigger a OutofMemory Exception Due to Response Body Size

Updated hardening settings for baseline version 2.0

Some hardening settings have been updated with the release of Security Center baseline version 2.0.

Documentation	Updates
Minimize concurrent interactive session quantity [Updated in Security Center 1.3]	• New short description: Minimize Concurrent Interactive Session Quantity • Old short description: Glide Authenticate Max Concurrent Interactive Sessions
Enforce certificate trust [Updated in Security Center 1.3, removed in 2.0, added in 7.0]	• New short description: Enforce Certificate Trust • Old short description: Certificate Trust
Maximize reset password SMS complexity [Updated in Security Center 1.3]	• New short description: Maximize Reset Password SMS Complexity • Old short description: Reset Password SMS Complexity
Enable High Security Plugin [Updated in Security Center 1.3]	• New short description: Enable High Security Plugin • Old short description: High Security Plugin
Enforce strict security of session cookies [Updated in Security Center 1.3]	• New short description: Enforce Strict Security of Session Cookies • Old short description: Secure Session Cookies

Documentation	Updates
Do not use demo certificates for active saml configurations [Updated in Security Center 1.5]	• New short description: Do Not Use Demo Certificates for Active SAML Configurations (Plugin Applicability: Multiple Provider Single Sign-On) • Old short description: Do Not Use Demo Certificates for Active SAML Configurations
Disable Entity Expansion within the XMLDocument2 Streaming Parser [Updated in Security Center 1.5]	Rule Script: Script has been updated to improve detection accuracy.
Restrict allowed Java packages [Updated in Security Center 1.3]	• New short description: Restrict Allowed Java Packages • Old short description: Java Packages Allowlist
Require obfuscation of mobile app UI [Updated in Security Center 1.3]	• New short description: Require Obfuscation of Mobile App UI • Old short description: Mobile App UI Obfuscation
Disable public access to favorites [Updated in Security Center 1.3 and 2.0]	• New short description: Disable Public Access to Favorites • Old short description: Public Access to Favorites
Escape JavaScript [Updated in Security Center 1.3]	• New description: The glide property *glide.html.escape_script* helps sanitize HTML fields. If *glide.html.escape_script* is not set to the recommended value of true, then inputs will not be sanitized for HTML fields (output encoding) from a backend Java context by removing embedded JavaScript. Javascript in HTML fields can lead to stored and reflected XSS. The ability to have XSS can lead to easily attained privilege escalation to higher roles such as admin where more lateral movement can be taken. • Old description: The glide property *glide.html.escape_script* helps sanitize html fields. If *glide.html.escape_script* is not set to the recommended value of true, then inputs will not be sanitized for HTML fields (output encoding) from a backend Java context by removing embedded JavaScript. Javascript in HTML fields can lead to stored

Documentation	Updates
	and reflected XSS. The ability to have XSS can lead to easily attained privilege escalation to higher roles such as admin where more lateral movement can be taken.
Set Xframe options to prevent embedding third-party websites [Updated in Security Center 1.3]	• New short description: Set Xframe Options to Prevent Embedding Third-Party Websites • Old short description: Xframe Options • New description: If *com.glide.cs.embed.xframe_options* is not set to the recommended value of DENY or SAMEORIGIN, then content of the web application could be embedded in a third-party site using an ALLOW-FROM uri. Allowing untrusted third-party sites could enable attacks such as clickjacking. • Old description: If *com.glide.cs.embed.xframe_options* is not set to the recommended value of DENY or SAMEORIGIN, then content of a the web application could be embedded in a third-party site using an ALLOW-FROM uri. Allowing untrusted third-party sites could enable attacks such as clickjacking. • Rule Script: Script has been updated to improve detection accuracy.
Escape HTML in list views [Updated in Security Center 1.3 and 1.5]	• New short description: Escape HTML in List Views • Old short description: Escape HTML
Require obfuscation of classic mobile app UI [Updated in Security Center 1.3]	• New short description: Require Obfuscation of Classic Mobile App UI • Old short description: Classic Mobile App UI Obfuscation
Deny by default with empty ACLs [Updated in Security Center 1.3]	• New short description: Deny by Default with Empty ACLs • Old short description: Security Manager Default Deny • New description: If *glide.sm.default_mode* is not set to the recommended value of deny, then the instance's legacy security manager allows access to a resource when there are no ACLs defined for that resource; or only wildcards table-level ACLs. By setting this

Documentation	Updates
	to allow anything that does not have explicit ACLs set is susceptible to manipulation. • Old description: If *glide.sm.default_mode* is not set to the recommended value of deny, then it allows access by the legacy security manager to a resource when there are no ACLs defined for that resource; or only wildcards table-level ACLs. By setting this to allow anything that does not have explicit ACLs set is susceptible to manipulation.
Maximize reset password request retry window duration [Updated in Security Center 1.3]	• New short description: Maximize Reset Password Request Retry Window Duration • Old short description: Reset Password Request Retry Window
Require Authorization for XSD Requests [Updated in Security Center 1.3]	• New short description: Require Authorization for XSD Requests • Old short description: XSD Request Authorization • New remediation: Ensure the property *glide.basicauth.required.xsd* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.xsd* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Escape jelly script [Updated in Security Center 1.3 and 1.5]	• New short description: Escape Jelly Script • Old short description: Escape Jelly
Double check inbound transactions [Updated in Security Center 1.3]	• New remediation: Ensure the property *glide.security.strict.updates* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.security.strict.updates* is set to true. • Rule Script: Script has been updated to improve detection accuracy.

Documentation	Updates
Restrict downloadable files types in static content [Updated in Security Center 1.3]	• New short description: Restrict Downloadable Files Types in Static Content • Old short description: Files Types Download Restrictions from Static Content
Require authorization for pdf requests [Updated in Security Center 1.3]	• New short description: Require Authorization for PDF Requests • Old short description: PDF Request Authorization • New remediation: Ensure the property *glide.basicauth.required.pdf* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.pdf* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Restrict uploaded MIME types [Updated in Security Center 1.3 and 2.0]	• New short description: Restrict Uploaded MIME Types • Old short description: Upload MIME Type Restriction • Rule Script: Script has been updated to improve detection accuracy.
Disable legacy JQuery behavior [Updated in Securty Center 1.3]	• New short description: Disable Legacy JQuery Behavior • Old short description: Legacy JQuery Behavior
Maximize reset password request unlock window duration [Updated in Security Center 1.3]	• New short description: Maximize Reset Password Request Unlock Window Duration • Old short description: Reset Password Request Unlock Window
Disable MultiSSO Debugging [Updated in Security Center 1.3 and 1.5]	• New short description: Disable MultiSSO Debugging (Plugin Applicability: Multiple Provider Single Sign-On) • Old short description: Disable MultiSSO Debugging • Rule Script: Script has been updated to improve detection accuracy.

Documentation	Updates
Enforce production instance behavior [Updated in Security Center 1.3 and 1.5]	• New short description: Enforce Production Instance Behavior • Old short description: Production Instance Behavior
Limit Invalid Password Reset Attempts [Updated in Security Center 1.3 and updated in 2.0]	• New short description: Minimize Reset Password Request Max Attempt Allowance • Old short description: Reset Password Request Max Attempts
Require authorization for csv requests [Updated in Security Center 1.3]	• New short description: Require Authorization for CSV Requests • Old short description: CSV Request Authorization • New remediation: Ensure the property *glide.basicauth.required.csv* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.csv* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Minimize reset password request success window duration [Updated in Securty Center 1.3]	• New short description: Minimize Reset Password Request Success Window Duration • Old short description: Reset Password Request Success Window
Enforce SOAP request strict security [Updated in Security Center 1.3]	• New short description: Enforce SOAP Request Strict Security • Old short description: SOAP Request Strict Security
Require authorization for SOAP requests [Updated in Security Center 1.3, 1.5, and 2.0]	• New short description: Require Authorization for SOAP Requests • Old short description: SOAP Request Authorization • New description: The glide property *glide.basicauth.required.soap* controls whether authentication is required to make a SOAP request to an instance. If *glide.basicauth.required.soap*

Documentation	Updates
	is not set to the recommended value of true, then authentication is disable for SOAP requests on the instance. It allows unauthenticated access to administrator or maint level operations; thereby negating security controls within the instance. • Old description: The glide property *glide.basicauth.required.soap* controls whether authentication is required i order to make a SOAP request to an instance. If glide.basicauth.required.soap is not set to the recommended value of true, then authentication is disable for SOAP requests on the instance. It allows unauthenticated access to administrator or maint level operations; thereby negating all security controls within the instance. • New remediation: Ensure the property *glide.basicauth.required.soap* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.soap* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Require XMLdoc2 entity validation with allowlistDisable entity expansion [Updated in Security Center 1.3]	• New short description: Require XMLdoc2 entity validation with allowlistDisable Entity Expansion • Old short description: XMLdoc2 entity validation with allowlistDisable Entity Expansion
Apply domain separation on dot walked fields [Updated in Security Center 1.3, 1.5, and 2.0]	• New short description: Apply Domain Separation on Dot Walked Fields (Plugin Applicability: Domain Separation • Old short description: Apply Domain Separation • New description: This property controls whether join queries are given domain separated conditions or not, in order to ensure they apply domain separation functionality for dot walked fields. If *glide.sys.domain.include_domain_condition_* is not set to the recommended value of true on an instance using domain separation, then sensitive information could be disclosed that is not to be shared with a specific domain.

Documentation	Updates
	• Old description: This property controls whether join queries are given domain separated conditions or not, in order to ensure they apply domain separation functionality for dot walked fields. If *glide.sys.domain.include_domain_condition_* is not set to the recommended value of true, then sensitive information could be disclosed that is not to be shared with a specific domain. • New remediation: Ensure the property *glide.sys.domain.include_domain_condition_* is set to true when the Domain Separation plugin is active. • Old remediation: Ensure the property *glide.sys.domain.include_domain_condition_* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Restrict JSONP Requests to Trusted URLs [Updated in Security Center 1.3]	• New short description: Restrict JSONP Requests to Trusted URLs • Old short description: JSONP Request Inclusion List • New description: This property specifies trusted URLs for the angularJS $http service to allow/reject JSONP requests. Property is necessary because this is a potentially breaking change for customers, so they need a way to add their trusted URLs. If *angular.jsonp.inclusion_list.enabled* is not set to the recommended value of "true", then JSONP requests are allowed to any URL. • Old description: This property specifies trusted URLs for the angularJS $http service to allow/reject JSONP requests. Property is necessary because this is a potentially breaking change for customers, so they need a way to add their trusted URLs. If *angular.jsonp.inclusion_list.enabled*) is not set to the recommended value of true, then jsonp requests are allowed to any url.
Minimize reset password max SMS per day [Updated in Security Center 1.3]	• New short description: Minimize Reset Password Max SMS Per Day • Old short description: Reset Password Max SMS Per Day

Documentation	Updates
Maximize reset password verification delay duration [Updated in Security Center 1.3]	• New short description: Maximize Reset Password Verification Delay Duration • Old short description: Reset Password Verification Delay • New description: If *password_reset.verification.delay* is not set to the recommended value of 1000 or more, then it will lead the login more susceptable to brute force attacks. This number of milliseconds delay limits the ability of a malicious actor to attempt to guess users identification or verification details, by using automation tools ("bots"). • Old description: If *password_reset.verification.delay* is not set to the recommended value of 1000 or more, then it will lead the login more susceptable to brute force attacks. This number of milliseconds delay limits the ability of a hacker to attempt to guess users identification or verification details, by using automation tools ("bots").
Require authorization for data broker rest API [Updated in Security Center 1.3]	• New short description: Require Authorization for Data Broker Rest API • Old short description: Data Broker Rest API Authorization • New description: If *glide.basicauth.required.databrokerrestapi* is not set to the recommended value of true, then basic authorization is not required for all inbound Data Broker Rest API requests. This could lead to unauthenticated information disclosure from the instance. • Old description: Starting in Utah release, if *glide.basicauth.required.databrokerrestapi* is not set to the recommended value of "true", then basic authorization is not required for all inbound Data Broker Rest API requests. This could lead to unauthenticated information disclosure from the instance. • New remediation: Ensure the property *glide.basicauth.required.databrokerrestapi* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.databrokerrestapi*

Documentation	Updates
	is set to true on instances running Utah release and later. • Rule Script: Script has been updated to improve detection accuracy.
Require authorization for JSONv2 request [Updated in Security Center 1.3]	• New short description: Require Authorization for JSONv2 Request • Old short description: JSONv2 Request Authorization • New remediation: Ensure the property *glide.basicauth.required.jsonv2* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.jsonv2* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Disable JavaScript tags in embedded HTML [Updated in Security Center 1.3]	• New short description: Disable JavaScript Tags in Embedded HTML • Old short description: Allow Javascript tags in Embedded HTML • New remediation: Ensure the property *glide.ui.security.codetag.allow_script* exists in the sys_properties table and is set to false. • Old remediation: Ensure the property *glide.ui.security.codetag.allow_script* is set to false. • Rule Script: Script has been updated to improve detection accuracy.
Enable security jump start plugin (ACL Rules) [Updated in Security Center 1.3]	• New short description: Enable Security Jump Start Plugin (ACL Rules) • Old short description: Security Jump Start Plugin (ACL Rules)
Set guest user for soap requests [Updated in Security Center 1.3 and 2.0]	• New short description: Set Guest User for SOAP Requests • Old short description: Guest User for SOAP Requests

Documentation	Updates
Restrict XML external entities [Updated in Security Center 1.3 and 2.0]	• New short description: Restrict XML External Entities • Old short description: XML Entity Validation URL Allowlist
Enable ACLs to Control Live Profile Details [Updated in Security Center 1.3]	• New short description: Enable ACLs to Control Live Profile Details • Old short description: Enable ACLs to control Live Profile Details
Restrict access to custom journal entries [Updated in Security Center 1.3 and removed in 2.0]	• New short description: Restrict Access to Custom Journal Entries • Old short description: Secure Custom Journal Entries • New description: If *glide.live_feed.custom_journal.acl_check_e* is not set to the recommended value of true, then all users to be able to see all journal entries within the life feed feature. Setting property to true will respect ACL's on custom journal fields which is a good to have feature. • Old description: If *glide.live_feed.custom_journal.acl_check_e* is not set to the recommended value of true, then all users to be able to see all journal entries. Setting property to true will respect ACL's on custom journal fields which is a good to have feature.
Set OTP lifetime for password reset to 1 hour [Updated in Security Center 2.0]	• New description: This property *glide.pwd_reset.onetime.token.validity* allows the link in the password reset email to expire after the number of hours specified in that *glide.pwd_reset.onetime.token.validity property*. Validity time of password reset token should be keeped as short as possible in according of normal user experience. Have long validity time for password reset token can help malicious actors to perform account takeover. • Old description: This property *glide.pwd_reset.onetime.token.validity* allows the link in the password reset email to expire after the number of hours specified in that *glide.pwd_reset.onetime.token.validity*

Documentation	Updates
	property. Validity time of password reset token should be keeped as short as possible in according of normal user experience. Have long validity time for password reset token can help hackers to perform account takeover.
Restrict delegated developers read access [Updated in Security Center 1.3]	• New short description: Restrict Delegated Developers Read Access • Old short description: Delegated Developers Read Access Allowlist
Define allowed ServiceNow internal IP addresses [Updated in Security Center 1.3 and 1.5]	• New short description: Define Allowed ServiceNow Internal IP Addresses • Old short description: IP Addresses Access Allowlist
Validate SOAP content type [Updated in Security Center 1.3]	• New short description: Validate SOAP Content Type • Old short description: SOAP Content Type Checking • Rule Script: Script has been updated to improve detection accuracy.
Require authorization for excel requests [Updated in Security Center 1.3]	• New short description: Require Authorization for Excel Requests • Old short description: Excel Request Authorization • New remediation: Ensure the property *glide.basicauth.required.excel* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.excel* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Require authorization for API requests [Updated in Security Center 1.3]	• New short description: Require Authorization for API Requests • Old short description: API Request Authorization • New remediation: Ensure the property *glide.basicauth.required.api*

Documentation	Updates
	exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.api* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Minimize Entity Expansion Threshold for GlideXMLUtil Scriptable [Updated in Security Center 1.3, 1.5, and 2.0]	• New short description: Minimize Entity Expansion Threshold • Old short description: Setting Entity Expansion Threshold
Notify users during password reset/change process [Removed in Security Center 1.5]	• New short description: Notify Users During Password Reset/Change Process • Old short description: Password Reset/ Change Notification Process • New remediation: Ensure Password Reset process notifies users upon password change or reset. • Old remediation: Ensure Password reset process notifies users upon password change or reset.
Disable legacy AngularJS behavior [Removed in Security Center 2.2]	• New short description: Disable Legacy AngularJS Behavior • Old short description: Legacy AngularJS Behavior
Maximize failed login unlock timeout duration [Updated in Security Center 1.3]	• New short description: Maximize Failed Login Unlock Timeout Duration • Old short description: Managing Unlock Timeout after Failed Logins
Enable HTTP Only Cookie Flag [Updated in Security Center 1.3]	• New short description: Enable HTTP Only Cookie Flag • Old short description: HTTP Only Cookie Flag • Rule Script: Script has been updated to improve detection accuracy.

Documentation	Updates
Enable scoped admin application ACLs [Updated in Security Center 1.3]	• New short description: Enable Scoped Admin Application ACLs • Old short description: Administer Scoped Admin Application ACLs • Rule Script: Script has been updated to improve detection accuracy.
Enable UserCookie version 3.1 [Updated in Security Center 2.0]	• New description: UserCookie v3 is generated only when property *glide.ui.secure.cookies.use_kmf* is disabled. UserCookie v3 is not secure due to storing secret key for HMAC in source code and identical for all customers. That can support malicious actors to use this one secret key for attempts to hijacking user sessions. • Old description: UserCookie v3 is generated only when property glide.ui.secure.cookies.use_kmf is disabled. UserCookie v3 is not secure due to storing secret key for HMAC in source code and identical for all customers. That can support hackers to use this one secret key for attempts to hijacking user sessions.
Require authorization for XML requests [Updated in Security Center 1.3]	• New short description: Require Authorization for XML Requests • Old short description: XML Request Authorization • New remediation: Ensure the property *glide.basicauth.required.xml* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.xml* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Minimize external user registration link expiration duration [Updated in Security Center 1.3 and 1.5]	• New short description: Minimize External User Registration Link Expiration Duration • Old short description: External User Registration Link Expiration

Documentation	Updates
	• New short description: Convert Inbound Email Images to Attachments • Old short description: Convert Inbound Email HTML
Minimize SMTP Recipient Quantity [Updated in Security Center 1.3]	• New short description: Minimize SMTP Recipient Quantity • Old short description: Max SMTP Recipients
Enable updated version of MultiSSO plugin [Updated in Security Center 1.3 and 1.5]	• New short description: Enable Updated Version of Multi SSO Plugin (Plugin Applicability: Multiple Provider Single Sign-On) • Old short description: Updated Version of Multi SSO Plugin is Enabled • New CVSS Score: 7.1 • Old CVSS Score: 5
Disable raw database query execution [Updated in Security Center 1.3 and removed in 2.0]	• New short description: Disable Raw Database Query Execution • Old short description: Operation Level Access Control Requirements • New description: This property allows a user to perform raw SQL queries on the database which can give access to tables and data outside of GlideRecord restrictions. If `glide.db.allow_unsafe_dbi_execute_sql` is not set to the recommended value of false, then this allows calling `dbi.executeStatement()` from a Glide Scriptable. • Old description: This property allows a user to perform raw SQL queries on the database which can give access to tables and data out of GlideRecord restrictions. If `glide.db.allow_unsafe_dbi_execute_sql` is not set to the recommended value of false, then this allows calling `dbi.executeStatement()` from a Glide Scriptable.
Escape XML markup [Updated in Security Center 1.3]	• New short description: Escape XML Markup • Old short description: Escape XML

Documentation	Updates
	• New remediation: Ensure the property *glide.ui.escape_text* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.ui.escape_text* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Require authorization for RSS requests [Updated in Security Center 1.3]	• New short description: Require Authorization for RSS Requests • Old short description: RSS Request Authorization • New remediation: Ensure the property *glide.basicauth.required.rss* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.rss* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Maximum allowed attachment size [Updated in Security Center 1.3]	• New short description: Minimize Allowed Attachment Size • Old short description: Max Allowed Attachment Size
Enforce relative links [Updated in Security Center 1.3 and 1.5]	• New description: The *glide.cms.catalog_uri_relative* property enforces relative links from the URI parameter on /ess/catalog.do. If *glide.cms.catalog_uri_relative* is not set to the recommended value of true, then the URL will not be sanitized with the enforceRelativeURL(url) function. Absolute URLs can pose a security risk when used as a part of parameter or a field value, thus redirecting the source page to an adversary-controlled website. • Old description: Use the *glide.cms.catalog_uri_relative* property to enforce relative links from the URI parameter on /ess/catalog.do. If *glide.cms.catalog_uri_relative*

Documentation	Updates
	is not set to the recommended value of true, then it may not sanitize URL with the enforceRelativeURL(url) function.
Enable SMS code notification for enrollment and verification [Updated in Security Center 1.3]	• New short description: Enable SMS Code Notification for Enrollment and Verification • Old short description: SMS Code Notification for Enrollment and Verification
Cache-Control HTTP Header Value [Updated in Security Center 1.3 and removed in 1.5]	• New short description: Cache-Control HTTP Header Value • Old short description: Cache-Control HTTP header value • Rule Script: Script has been updated to improve detection accuracy.
Deny internal access to explicit external roles [Updated in Security Center 1.3 and 1.5]	• New short description: Deny Internal Access to Explicit External Roles • Old short description: Enable Explicit Roles Internal Denylist • New technical configuration name: glide.security.explicit_roles.enable_internal_user_blacklist,gli • Old technical configuration name: glide.security.explicit_roles.enable_internal_user_blacklist • New description: This prevents external users from being assigned the snc_internal role. If *glide.security.explicit_roles.enable_intei* is not set to the recommended value of true, and the *glide.security.explicit_roles.internal_use* property is not set to a list of untrusted user classes, then the specified roles can be assigned the snc_internal role instead of the snc_external role. If the list is empty, then all users will be assigned the snc_internal role by default. The property should contain at least the default roles csm_consumer_user,customer_contact. Misconfiguration of these properties increases the risk that an external user account gains access to internal information. • Old description: This property prevents external users from being assigned the snc_internal role. If *glide.security.explicit_roles.enable_intei* is set to the recommended

Documentation	Updates
	value of true, then it enables *glide.security.explicit_roles.internal_use property* which allows to assign snc_external role. If the value is set to false, it disables *glide.security.explicit_roles.internal_use* property. • New remediation: Ensure the property *glide.security.explicit_roles.enable_inte* is set to true and that the property *glide.security.explicit_roles.internal_use* includes the dangerous items csm_consumer_user, customer_contact. • Old remediation: Ensure the property *glide.security.explicit_roles.enable_inte* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Minimize one-time out of band verifier lifetime duration [Updated in Security Center 1.3]	• New short description: Minimize One-Time Out of Band Verifier Lifetime Duration • Old short description: Short One-Time Out of Band Verifier Lifetime
Require authorization for script requests [Updated in Security Center 1.3]	• New short description: Require Authorization for Script Requests • Old short description: Script Request Authorization • New remediation: Ensure the property *glide.basicauth.required.scriptedprocesso* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.scriptedprocesso* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Limit concurrent interactive sessions [Updated in Security Center 1.3]	• New short description: Limit Concurrent Interactive Sessions • Old short description: Glide Authenticate Limit Concurrent Interactive Sessions • New description: This property is meant to be used with the Limit Concurrent Sessions (*com.glide.limit.concurrent.sessions*) plugin. When the plugin is active and the property is set to false, a user can have any

Documentation	Updates
	number of concurrent interactive sessions on an instance. A greater number of open sessions means there is a great possibility for session hijacking to occur. • Old description: This propert is meant to be used with the Limit Concurrent Sessions (*com.glide.limit.concurrent.sessions*) plugin. When the plugin is active and the property is set to false, a user can have any number of concurrent interactive sessions on an instance. A greater number of open sessions means there is a great possibility for session hijacking to occur.
Prevent Users From Accepting Warning To Bypass CSRF Validation [Updated in Security Center 1.3 and 1.5]	• New short description: Enforce CSRF Token Strict Validation • Old short description: CSRF Strict Validation • New description: This property enables CSRF token strict validation which prevents the reuse of CSRF tokens. If *glide.security.csrf.strict.validation.mode* is not set to the recommended value of true, then CSRF tokens could be reused which opens a door to CSRF attacks. • Old description: This property enables CSRF token strict validation which prevents the reuse of CSRF tokens. If *glide.security.csrf.strict.validation.mode* is not set to the recommended value of true, then CSRF token could be reused which opens a door tot CSRF attacks.
Minimize session activity timeout duration [Updated in Security Center 1.3]	• New short description: Minimize Session Activity Timeout Duration • Old short description: Session Activity Timeout
Enable HTML Sanitizer [Updated in Security Center 1.3]	• New short description: Enable HTML Sanitizer • Old short description: HTML Sanitizer
Restrict access to background script [Updated in Security Center 1.3 and 2.0]	• New description: This property holds the required role to access Script Background module. If *glide.script_processor.admin* is not set to the recommended value of admin, security_admin, or maint, then users having a lower privileged role will be able to run

Documentation	Updates
	background scripts on the instance. This will lead to a complete bypass of the ACL system allowing full access to tables. • Old description: This property holds the required role to access Script Background module. If *glide.script_processor.admin* is not set to the recommended value of Admin, then any user having a low privileged role will be able to run background scripts on the instance. This will lead to a complete bypass of the ACL system allowing full access to tables • New remediation: Ensure the property *glide.script_processor.admin* is set to the admin, security_admin, or maint role. • Old remediation: Ensure the property *glide.script_processor.admin* is set to Admin. • Rule Script: Script has been updated to improve detection accuracy.
Disable embedded HTML code [Updated in Security Center 1.3]	• New short description: Disable Embedded HTML Code • Old short description: Embedded HTML Code
Minimize absolute session timeout duration [Updated in Security Center 1.3]	• New short description: Minimize Absolute Session Timeout Duration • Old short description: Absolute Session Timeout
Require authentication by default for client-callable script includes [Updated in Security Center 1.3]	• New short description: Require Authentication by Default for Client-Callable Script Includes • Old short description: Privacy on Client-Callable Script Includes
Restrict access to GlideSystemUserSession scriptable API [Updated in Security Center 1.3 and 2.0]	• New short description: Restrict Access to GlideSystemUserSession Scriptable API • Old short description: Access to GlideSystemUserSession scriptable API

Documentation	Updates
Enforce HTML Sanitization [Updated in Security Center 1.3]	• New short description: Enforce HTML Sanitization • Old short description: Check Unsanitized HTML • Rule Script: Script has been updated to improve detection accuracy.
Minimize absolute session timeout duration [Updated in Security Center 1.3]	• New short description: Minimize Absolute Session Timeout Duration • Old short description: Absolute Session Timeout
Activate role based multi-factor authentication [Updated in Security Center 1.3]	• New short description: Activate Role Based Multi-Factor Authentication • Old short description: Role Based Multi-Factor Authentication
Minimize SAML notBefore or notOnOrAfter constraint duration [Updated in Security Center 1.3 and 1.5]	• New short description: Minimize SAML "notBefore" or "notOnOrAfter" Constraint Duration (Plugin Applicability: Multiple Provider Single Sign-On) • Old short description: SAML "notBefore" or "notOnOrAfter" Constraint • Rule Script: Script has been updated to improve detection accuracy.
Restrict email domains for external user registration [Updated in Security Center 1.3, 1.5, and 2.0]	• New short description: Restrict Email Domains for External User Registration (Plugin Applicability: External User Registration) • Old short description: External User Registraiton Email Domain Allowlist • New remediation: Ensure the property *sn_ext_usr_reg.allowed_email_domains* is not set to an empty value. • Old remediation: Ensure the property *sn_ext_usr_reg.allowed_email_domains* is not set to an empty value.
Maximize reset password SMS pause window duration [Updated in Security Center 1.3]	• New short description: Maximize Reset Password SMS Pause Window Duration • Old short description: Reset Password SMS Pause Window

Documentation	Updates
	• New remediation: Ensure the property *password_reset.sms.pause_window* is set to 2 or greater. • Old remediation: Ensure the property *password_reset.sms.pause_window* is set to 2. • Rule Script: Script has been updated to improve detection accuracy.
Disable outbound SSLv2/SSLv3 connections [Updated in Security Center 1.3]	• New short description: Disable Outbound SSLv2/SSLv3 Connections • Old short description: Disabling SSLv2/ SSLv3
Require authorization for unload requests [Updated in Security Center 1.3]	• New short description: Require Authorization for Unload Requests • Old short description: Unload Request Authorization • New remediation: Ensure the property *glide.basicauth.required.unl* exists in the sys_properties_table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.unl* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Enable email spam scoring and filtering [Updated in Security Center 1.3]	• New short description: Enable Email Spam Scoring and Filtering • Old short description: Email Spam Scoring and Filtering
Unset LDAP Initial distinguished name [Updated in Security Center 1.3 and removed in 2.0]	• New short description: Unset LDAP Initial Distinguished Name • Old short description: LDAP Initial Distinguished Name
Enable Anti-CSRF token [New in Security Center 1.3, updated in 1.5, and removed in 2.0]	• New short description: Enable Anti-CSRF Token • Old short description: Anti-CSRF Token

Documentation	Updates
Require AJAXGlideRecord ACL checking [Updated in Security Center 1.3]	• New short description: Require AJAXGlideRecord ACL Checking • Old short description: Enabling AJAXGlideRecord ACL Checking
Log user impersonation [Updated in Security Center 1.3 and 2.0]	Rule Script: Script has been updated to improve detection accuracy.Script has been updated to improve detection accuracy.
Disallow infected file download [Updated in Security Center 1.5 and 2.0]	Rule Script: Script has been updated to improve detection accuracy.
Enable Captcha for External User Registration [Updated in Security Center 1.3 and 1.5]	• New short description: Enable Captcha for External User Registration (Plugin Applicability: External User Registration) • Old short description: Enable Captcha for External User Registration
Disable SQL Error Messages [Updated in Security Center 1.3 and 1.5]	• New short description: Disable SQL Error Messages • Old short description: Disabling SQL error messages
Minimize reset password request expiration duration [Updated in Security Center 1.3]	• New short description: Minimize Reset Password Request Expiration Duration • Old short description: Reset Password Request Expiration • Rule Script: Script has been updated to improve detection accuracy.
Control Lockout Time for Invalid Password Reset Attempts [Updated in Security Center 1.3 and 2.0]	• New short description: Minimize Reset Password Request Max Attempts Window Duration • Old short description: Reset Password Request Max Attempts Window
Restrict downloadable MIME types [Updated in Security Center 1.3 and 2.0]	• New short description: Restrict Downloadable MIME Types • Old short description: Downloadable Mime Type Denylist
Escape Excel Formulas [Updated in Security Center 1.3]	• New short description: Escape Excel Formulas • Old short description: Escape Excel Formula

Documentation	Updates
Enable contextual security plugin [Updated in Security Center 1.3]	• New short description: Enable Contextual Security Plugin • Old short description: Contextual Security Plugin
Enable account recovery [Updated in Security Center 1.3 and 1.5]	• New short description: Enable Account Recovery (Plugin Applicability: Multiple Provider Single Sign-On) • Old short description: Account Recovery • New description: This property controls the account recovery feature which binds the ability to bypass single sign-on to specifically designated administrators. If *glide.sso.acr.enabled* is not set to the recommended value of true, then the local interactive log-ins (username or password based) will be remain enabled when single sign-on is enabled on the instance. Eliminating local interactive log-ins reduces the potential for unauthorized access to the instance. • Old description: This property controls the account recovery feature. If *glide.sso.acr.enabled* is not set to the recommended value of true, then Account recovery by userId will not be possible. • New CVSS Score: 6.5 • Old CVSS Score: 9.1 • Rule Script: Script has been updated to improve detection accuracy.
Require authorization for import requests [Updated in Security Center 1.3]	• New short description: Require Authorization for Import Requests • Old short description: Import Request Authorization • New remediation: Ensure the property *glide.basicauth.required.importprocessor* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.importprocessor* is set to true. • Rule Script: Script has been updated to improve detection accuracy.

Documentation	Updates
Enable SNC access control plugin [Updated in Security Center 1.3]	• New short description: Enable SNC Access Control Plugin • Old short description: SNC Access Control Plugin
Limit concurrent sessions across all nodes [Updated in Security Center 1.3]	• New short description: Limit Concurrent Sessions Across All Nodes • Old short description: Glide Authenticate Limit Concurrent Sessions Across All Nodes
Require authorization for XML output requests [Updated in Security Center 1.3]	• New short description: Require Authorization for XML Output Requests • Old short description: XML Output Authorization • New remediation: Ensure the property *glide.basicauth.required.xmloutputprocesso* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.xmloutputprocesso* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Escape scripts in scratchpad [Updated in Security Center 1.3]	• New short description: Escape Scripts in Scratchpad • Old short description: Escape Scratchpad • New description: The scratchpad is an easy way to set information on the server that can be accessed in the browser. An admin can script anything to be on it, including arbitrary data from arbitrary records. If *glide.ui.escape_scratchpad* is not set to the recommended value of true, then it is possible to execute malicious script like a cross-site scripting vulnerability. • Old description: The scratchpad is an easy way to set information on the server that can be accessed in the browser. An admin can script anything to be on it, including arbitrary data from arbitrary records.If *glide.ui.escape_scratchpad* is not set to the recommended value of true, then it is possible to execute malicious script like a cross-site scripting vulnerability.

Documentation	Updates
Require authorization for WSDL request [Updated in Security Center 1.3 and 1.5]	• New short description: Require Authorization for WSDL Request • Old short description: WSDL Request Authorization • New remediation: Ensure the property *glide.basicauth.required.wsdl* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.wsdl* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Require authorization for SCHEMA requests [Updated in Security Center 1.3]	• New short description: Require Authorization for SCHEMA Requests • Old short description: SCHEMA Request Authorization • New remediation: Ensure the property *glide.basicauth.required.schema* exists in the sys_properties table and is set to true. • Old remediation: Ensure the property *glide.basicauth.required.schema* is set to true. • Rule Script: Script has been updated to improve detection accuracy.
Restrict downloadable MIME types [Updated in Security Center 1.3 and 2.0]	• New short description: Restrict Downloadable MIME Types • Old short description: Downloadable Mime Types
Disable logger for low privilege users in script sandbox [Updated in Security Center 1.3]	• New short description: Disable Logger for Low Privilege Users in Script Sandbox • Old short description: Glide Security Logger No Loggining for Sandbox • Rule Script: Script has been updated to improve detection accuracy.

Documentation	Updates
Implement the x-frame-options: SAMEORIGIN security header [Updated in Security Center 1.3]	• New short description: Implement the X-Frame-Options: SAMEORIGIN Security Header • Old short description: X-Frame-Options: SAMEORIGIN • Rule Script: Script has been updated to improve detection accuracy.
Restrict performance monitoring access [Updated in Security Center 1.3]	• New short description: Restrict Performance Monitoring Access • Old short description: Performance Monitoring ACL
Turn off verbose SQL error messages for import processor [Updated in Security Center 1.3]	Rule Script: Script has been updated to improve detection accuracy.
Minimize reset password SMS expiracy duration [Updated in Security Center 1.3]	• New short description: Minimize Reset Password SMS Expiracy Duration • Old short description: Reset Password SMS Expiracy
Disable creating users from incoming emails [Updated in Securty Center 1.3]	• New short description: Disable Creating Users from Incoming Emails • Old short description: Restrict Emails by Domain • New description: An administrator can set an email property to automatically create users from incoming emails. If set this property to the insecure value, the instance will automatically create users from incoming email. Each user created will have the same hardcoded default password which makes bypassing authentication through brute force easier. • Old description: An administrator can set an email property to automatically create users from incoming emails. If set this property to the insecure value, the instance will automatically create users from incoming email. Each user created will have the same hardcoded default password which makes bypassing authentication through brute force easier.

Documentation	Updates
	• New remediation: Ensure the property *glide.pop3readerjob.create_caller* is set to false. • Old remediation: Ensure the property *glide.pop3readerjob.create_caller* is set to false

Deleted hardening settings

Some hardening settings are removed from Security Center baselines.

Deleted hardening settings for baseline version 7.0

Some hardening settings have been removed with the release of Security Center baseline version 7.0.

- Disable Password-Less Authentication
- Escape XML Response
- Minimize One-Time Out of Band Verifier Lifetime Duration

Deleted hardening settings for baseline version 6.0

Some hardening settings have been removed with the release of Security Center baseline version 6.0.

- Enforce ACL on HR Lifecycle Events Data
- Enforce Security Scope License and Permit Playbook
- Enforce ACL on HR Core Data
- Restrict Global Administrators from Bypassing Scoped App Access Restrictions
- Disable Legacy AngularJS Behavior
- Enforce Field Level ACLs in GlideRecordSandbox
- Enforce Scope Security for Public Sector Digital Services
- Restrict Downloadable Files Types in Static Content
- Enforce Scoped ACL Access for Information Request Playbooks
- Minimize SAML "notBefore" or "notOnOrAfter" Constraint Duration
- Restrict Delegated Developers Read Access
- Enforce Password Policy on All Methods
- Enforce Security Scope for Service Application Information
- Require Obfuscation of Classic Mobile App UI
- Enable SSL in LDAP Authentication
- Enforce Security Scope for Agent Workspace for HR Case Management
- Limit Attachment Size in Training and Prediction Flows
- Remove Credentials From Welcome Page
- Enforce ACL on HR Virtual Agent Data
- Check Impersonation on ACL Evaluation in HR App

- Restrict Allowed Java Packages
- Enforce Strict Code Signing Checks

Deleted hardening settings for baseline version 5.0

Some hardening settings have been removed with the release of Security Center baseline version 5.0.

- Require Authentication on Event Management HTTP Processor
- Restrict Access to Custom Journal Entries
- Do Not Apply Password Policy at Login
- Disable Raw Database Query Execution
- Log HTML Sanitization
- Enable Anti-CSRF Token
- Enforce Certificate Trust
- Unset LDAP Initial Distinguished Name
- Set Minimal Password Length

Deleted hardening settings for baseline version 4.0

Some hardening settings have been removed with the release of Security Center baseline version 4.0.

- LDAP Initial Password
- Mobile Offline Roles
- Restrict Access to Critical Data through Zero Trust Access Policies (Plugin Applicability: Adaptive Authentication)
- Read Only Tables Allowlist For Write
- Allowed JDBC Probe Operations (Plugin Applicability: MID Server)
- Set Allowed Domains which Create Users from Incoming Emails
- Set Complex "Default" Password
- Accessible Properties in GraphQL Allowlist
- Cross Origin Messaging Allowlist
- Limit Attachment Size in Training and Prediction Flows for GraphQL Enpoints (Plugin Applicability: Platform Document Intelligence)
- Edit Content Roles Allowlist (Plugin Applicability: Communities)
- Ensure Database Queries Do Not Trigger a OutofMemory Exception Due to Query Size
- Role Allowlist for Script Execution
- Downloadable File Type Allowlist
- Read Only Tables Allowlist For Delete
- Enforce Authentication for Roleless ACL
- Minimize LDAP One-Time Token Expiry Time
- Allow Only Trusted IP Addresses for Authentication
- Set Mobile Password Reset URL
- Password Complexity of Service Accounts (Plugin Applicability: Service Bridge)

- Notify Users During Password Reset/Change Process
- Enforce Application Scope Restrictions
- Access Control Requirements (Plugin Applicability: Communities)
- Record History Access Role Allowlist
- Restrict Role Access for Attachments
- Only Allow PDFs from Predefined List of Trusted URLs
- Prevent Emailing One-Time Password During LDAP Server Outage
- Convert Inbound Email Images to Attachments
- Review Extraneous Explicit Role Access Control Condition

Deleted hardening settings for baseline version 2.0

Some hardening settings have been removed with the release of Security Center baseline version 2.0.

- Enable Code Signing for Application Configuration Data and Scripts
- Enable Glide KMF encrypter
- Disable the Use of Instance Level Encrypter
- Enable Explicit Roles Internal Denylist
- Block Direct Inserts to MID ECC Queue for Code Signing Requirement
- Log All Outbound HTTP Request Fields

Access control

The access control category audits the process of protecting resources from unauthorized access through granting and denying requests based on a permission model. This includes ensuring an entity accessing a resource holds valid credentials to do so, creating and protecting a well-defined set of roles or permissions and ensuring role or permission controls are protected from replay and tampering.

Access controls determine whether access to a particular resource should be granted or denied. It only allows access to resources to those users permitted to use them.

Anti-CSRF token validation time [New in Security Center 1.3]

The *glide.security.csrf_previous.time_limit* property specifies the time in seconds for a secure token to expire.

When a user session expires, the secure token expires with it unless the **allowing reuse of expired tokens are allowed** property is enabled, and it is within the time frame described by this property. This token is used to prevent cross site request forgery attacks.

More information

Attribute	Description
Configuration name	*glide.security.csrf_previous.time_limit*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer

Attribute	Description
Recommended value	86400 seconds or 1 day
Default value	86400 seconds or 1 day
Category	Access control
Security risk	• Severity score: 5.3 • CVSS score: Medium • Security risk details: Not setting this property to the recommended value disables the token validation used to prevent cross site request forgery attacks.
Dependencies and prerequisites	None
Functional impact	This property determines the duration in seconds for a secure token to remain valid. The secure token expires when the user session expires unless the **allowing reuse of expired tokens** property is enable, and the token is within the time frame specified in this property. This token prevents cross-site request forgery attacks. It has a default value of 86400 seconds or 1 day.

Apply domain separation on dot walked fields [Updated in Security Center 1.3, 1.5, and 2.0]

The *glide.sys.domain.include_domain_condition_on_join* property controls whether join queries are given domain separated conditions or not in order to ensure they apply domain separation functionality for dot walked fields.

This property controls whether join queries are given domain separated conditions or not, in order to ensure they apply domain separation functionality for dot walked fields. If *glide.sys.domain.include_domain_condition_on_join* is not set to the recommended value of true on an instance using domain separation, then sensitive information could be disclosed that is not to be shared with a specific domain. There may be moderate functional impact to the instance if components are reliant on the unsafe cross domain queries. Instances should be tested in subproduction environments before enabling.

⚠ **Warning:** This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.sys.domain.include_domain_condition_on_join*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	True, when domain separation is installed, otherwise the property won't exist.
Default value	false
Category	Access control

Attribute	Description
Purpose	Controls whether join queries are given domain separated conditions or not, in order to ensure they apply domain separation functionality for dot walked fields.
Security risk	• Severity score: 6.5 • CVSS score: Medium • Security risk details: If *glide.sys.domain.include_domain_condition_on_joi* is not set to the recommended value of **true**, then sensitive information could be disclosed that is not to be shared with a specific domain.
References	Domain separation for service providers

Block access for delegated developers

This configuration affects access for delegated developers that are updating user roles through script. When the configuration is compliant, the developer will not be able to update or insert records into the sys_user_has_role table without also having the user_admin role.

The value of this property affects whether a delegated developer is allowed to grant or receive unexpected access to functionality in the instance. When the property contains roles, only those roles may execute script modules.

More information

Attribute	Description
Property name	*com.glide.sys.security.delegateddev.block_grant_ro*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	The value of this property affects whether a delegated developer is allowed to grant or receive unexpected access to functionality in the instance.
Type	toggle switch
Default value	true
Recommended value	true
Security Dependencies	none
Security risk rating	6.7
Functional impact	When a user with the delegated_developer role is attempting to modify a record in the sys_user_has_role table, this property enables additional security checks against the operation. The additional security checks validate that the user has been granted the user_admin role if they're trying to create or update the sys_user_has_role table. If they do not have the

Attribute	Description
	user_admin role, the access will be denied. When the property is false, these additional checks are not validated.
Security risk	(Moderate) Without appropriate authorization, unauthorized users may access sensitive content/data on the instance.
References	Access control

To learn more about adding or creating a system property, see Add a system property.

Block Expired Anti-CSRF Tokens [Updated in Security Center 1.5]

Block expired CSRF tokens to prevent cross-site request forgery attacks.

Overview

Cross-site request forgeries are a type of malicious exploit whereby unauthorized commands are performed on behalf of an authenticated user.

Configuration details

Attribute	Description
Overview	Controls the usage of an expired secure token to identify and validate incoming requests. Set to **false** to prevent a previously expired token to validate an incoming request.
Configuration name	*glide.security.csrf_previous.allow*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	**false**
Default value	**true**
Category	Access control
Security risk	**Severity score**: 6.5
	Severity rating per CVSS score: Medium
	Security risk details: Enforces a strong anti-CSRF mechanism to protect authenticated functionality, and effective anti-automation or anti-CSRF protects unauthenticated functionality.
Dependencies and prerequisites	None
References	Enable Anti-CSRF token [New in Security Center 1.3, updated in 1.5, and removed in 2.0], cross-site request forgery.

Check UI action conditions before execution

Use the *glide.security.strict.actions* property to enable checking of UI actions conditions in forms and lists before they execute. When you set this property to true, it adds an extra layer of validation on the table UI actions before they are executed.

More information

Attribute	Description
Property name	*glide.security.strict.actions*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	To ensure an extra layer of validation on the table UI actions before they are executed.
Data type	Boolean
Recommended value	true
Default value	true
Security risk rating	3.3
Functional impact	This remediation only adds an extra layer of validation to check for UI actions on the target table/page on the instance. As long as the access controls are set appropriately on the customer instance, there should not be an impact here.
Security risk	(Low) Access request is always checked when transactions happen between two zones. This operation validates any UI actions before the form renders for the end user.
References	High security plugin

To learn more about adding or creating a system property, see Add a system property.

Configure event management assignment group admin roles [New in Security Center 1.5]

Use the *evt_mgmt.connector_assignment_group_admin_roles* property to set which roles are authorized for admin access over the assignment group field in connector instances.

The *evt_mgmt.connector_assignment_group_admin_roles* property contains a comma separated string which indicates the role names that have admin access over the assignment group field in connector instances. Changing the default roles in this list may enable unauthorized users to alter event integrations on the instance. To prevent unauthorized access to roles, set *evt_mgmt.connector_assignment_group_admin_roles* to the value of admin,evt_mgmt_admin,sn_sow_srm.srm_admin. Review any additional roles in the recommended value string to ensure that the role should be included.

More information

Attribute	Description
Configuration name	*evt_mgmt.connector_assignment_group_admin_roles*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	admin,evt_mgmt_admin,sn_sow_srm.srm_admin
Default value	admin,evt_mgmt_admin,sn_sow_srm.srm_admin
Category	Access control
Security risk	• Severity score: 3.1 • CVSS score: Low • Security risk details: Changing the default roles could enable unauthorized users to alter event integrations on the instance.
Dependencies and prerequisites	None
References	Creating groups

Configure Service Portal Widgets Allow List [New in Security Center 2.0]

Learn how to configure the glide.service_portal.widget.allow_list property securely so that the access control lists (ACLs) for the tables do not expose sensitive information.

The *glide.service_portal.widget.allow_list* property identifies the widgets that can access any table within the instance. However, the access control lists (ACLs) for these tables will continue to apply. If the ACLs are incorrectly configured or absent, widgets on this list might enable access to these tables, potentially exposing sensitive information. This property is effective only if the widget uses SNCACLWidgetUtil and the *glide.service_portal.widget.enforce_public_check* property is enabled (set to true).

More information

Attribute	Description
Configuration name	*glide.service_portal.widget.allow_list*
Configuration type	System Properties (/sys_properties_list.do)
Data type	array
Recommended value	Empty
Default value	Empty - in some customer's cases there might be some values.
Category	Access control

Attribute	Description
Security risk	• Severity score: 3.7 • CVSS score: Low • Security risk details: Not configuring this property to the recommended values could enable widgets to access any table within the instance.
Dependencies and prerequisites	For the *glide.service_portal.widget.allow_list* setting to be applicable, the *glide.service_portal.widget.enforce_public_check* property must be set to true.
Functional impact	This property enables customers to access any table information if the widget is set to public and is included in the property's value.

Configure Service Portal Widgets Table Allow List [New in Security Center 2.0]

Learn how the *glide.service_portal.widget.table_allow_list* property enhances security by listing tables accessible to unauthenticated users through Service Portal widgets, dependent on additional checks and specific glide property settings.

The *glide.service_portal.widget.table_allow_list* property contains a list of tables that unauthenticated users can access through Service Portal widgets, which utilize the additional security checks in the SNCACLWidgetUtil script. This property is enforced only if the glide property *glide.service_portal.widget.enforce_public_check* is set to true. Including unnecessary tables in this property may lead to the disclosure of sensitive information. Nonetheless, Table ACLs will continue to be evaluated as they were previously.

More information

Attribute	Description
Configuration name	*glide.service_portal.widget.table_allow_list*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Access control
Security risk	• Severity score: 3.7 • CVSS score: Low • Security risk details: If this property is not set to the secure value, unnecessary tables may be included, potentially leading to the disclosure of sensitive information.

Attribute	Description
Dependencies and prerequisites	The *glide.service_portal.widget.enforce_public_check* property must be set to true for the *glide.service_portal.widget.table_allow_list* setting to take effect.
Functional impact	The table list controls access to the tables from which the widget is allowed to retrieve data.

Deny internal access to explicit external roles [Updated in Security Center 1.3 and 1.5]

Use system properties to determine whether external users can be assigned the snc_internal role.

Use the *glide.security.explicit_roles.enable_internal_user_blacklist* system property to prevent external users from being assigned the *snc_internal* role. When this property is set to **true**, it enforces the parameters of the maint-protected *glide.security.explicit_roles.internal_user_blacklist* property. This property assigns the *snc_external* role to a list of untrusted user classes. If *glide.security.explicit_roles.enable_internal_user_blacklist* is set to **false**, the *glide.security.explicit_roles.internal_user_blacklist* property is ignored.

Note: Instances without Explicit Roles installed are not affected. As of the Paris release, new installations of Explicit Roles get the property with a default value of true.

More information

Attribute	Description
Configuration name	*glide.security.explicit_roles.enable_internal_user*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Fallback value	false
Category	Session management
Security risk	• Severity score: 5.4 • CVSS score: Medium • Misconfiguration of this property increases the risk that an external user account gains access to internal information.
Dependencies and prerequisites	None

Deny unauthorized access to request items [Updated in Security Center 1.3]

The *glide.sc.req_for.roles.default* property defines a default behavior for the *retrieveAddress* API.

This property is functional only when *glide.sc.req_for.roles* has no values. If *glide.sc.req_for.roles* has values, then this property has no significance and users with only defined roles are given access to the API.

More information

Attribute	Description
Property name	*glide.sc.req_for.roles.default*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	When there are no roles given in the property, the Client Callable Script Include *ScriptServiceCatalogGetLocation* can be called by any unprivileged logged-in user and can retrieve the address of any other users in the system. This property protects this API to be exposed to unprivileged users.
Recommended value	deny
Default value	deny
Configuration type	Choicelist (allow | deny)
Security risk	(Moderate) If *glide.sc.req_for.roles.default* is not set to the recommended value of **deny** (allow) and the value of *glide.sc.req_for.roles* is empty, then any user can request items for other users allowing unauthorized resource access.
References	Client-callable script includes

To learn more about adding or creating a system property, see Add a system property.

Display recommendations for high risk UI pages

Decrease the likelihood of authorization errors, and unintended information disclosure by displaying recommendations for high risk UI pages.

Use the **glide.script.ui_page.customer_scoped.security_msgs_enabled** system property to determine whether to display security recommendations to users configuring UI pages when:

- An ACL is missing
- GlideRecord/GlideDBQuery APIs are used instead of GlideRecordSecure
- the page is configured as public in the "sys_public" table

When the property is enabled, recommendations when the preceding conditions are met, decreasing the likelihood of authorization errors, and helping prevent unintended information disclosure.

Set the system property **glide.script.ui_page.customer_scoped.security_msgs_enabled** to **true** to display these recommendations.

More information

Attribute	Description
Technical configuration name	glide.script.ui_page.customer_scoped.security_msgs_enabled
Plugin applicability	None
Security risk	Setting this system property to **false** increase the likelihood of authorization errors, and unintended information disclosure.
Common Vulnerability Scoring System (CVSS) score	5.3
Common Vulnerability Scoring System (CVSS) rating	Medium
Functional impact	Displays security recommendations to users configuring UI pages.
Dependencies and prerequisites	None
Data type	Boolean
Base system value	true
Fallback value	true
Recommended value	true

To learn more about adding or creating a system property, see Add a system property.

Disable inbound emails for locked out users

Use the *glide.pop3.process_locked_out* property to control inbound email actions for locked out, active users.

Set this property to **false** to disable inbound emails for locked out users.

Note: Consider the security implications of allowing users from untrusted domains, and why they were locked out, before allowing emails from them to trigger inbound email actions.

More information

Attribute	Description
Property name	*glide.pop3.process_locked_out*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	This property controls inbound email actions for locked out users.

Attribute	Description
Data type	Boolean
Recommended value	false
Default value	false
Security risk	(High) When you set this property to **true**, there may be an information disclosure as inbound emails would be received by users with locked accounts.
Security risk rating	7.5

To learn more about adding or creating a system property, see Add a system property.

Double check inbound transactions [Updated in Security Center 1.3]

Use the *glide.security.strict.updates* property to enable double-checking of security on inbound transactions during form submission. When you set this property to **true**, it adds an extra layer of table validation before a form renders in the browser.

Ensure the property *glide.security.strict.updates* exists in the sys_properties table and is set to true.

More information

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

Attribute	Description
Property name	*glide.security.strict.updates*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	To ensure an added layer of verification of user permissions before presenting the form in the browser.
Data type	Boolean
Recommended value	true
Default value	true
Security risk rating	8.1
Functional impact	This remediation adds an extra layer of validation to check for user permissions on the target table/page on the instance. As long as the access controls are set appropriately on the customer instance, there should be no impact.
Security risk	(High) You should always check access request when transactions happen between two zones. This operation checks for permissions when the form is requested and before form rendering happens.
References	

To learn more about adding or creating a system property, see Add a system property.

Enable scoped admin application ACLs [Updated in Security Center 1.3]

The *glide.security.scoped_administration.honor_global_acl* determines whether an application administration app can inherit global access control list (ACL) rules.

This property is especially useful when there are no scoped admin application ACLs defined for the record scope.

Set *glide.security.scoped_administration.honor_global_acl* to true to prevent a low privileged user with permissions to the application to potentially access sensitive records.

More information

Attribute	Description
Property name	*glide.security.scoped_administration.honor_global_*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	Controls ACL access rule in scoped admin application.
Recommended value	True
Default value	True
Configuration type	Boolean
Security risk	(Low) When the property value is true and there are no scoped admin application ACLs defined for the record scope, the global ACLs will be honored. If set to false, with no scoped admin application ACLs defined for the record scope, ACL checks will be ignored.
Security risk rating	3.8
References	Access control rules in application administration apps

To learn more about activating a plugin, see Activate a plugin

Enable work order management query rules for service organizations [New in Security Center 1.5 and updated in 2.0]

Use the *sn_fsm.use_query_rules* property to apply rules and filters to the Field Service Management tables.

When set to **true**, rules/filters from sn_query_rule table will be used to determine read access to Field Service Management-related tables (Work Order and Work Order Task) to the logged in user through query business rules and read ACLs. When **false**, the records won't be filtered based on query rules. Query business rules add additional security validations. Specifically, this property will filter records for agents, qualifiers, and dispatchers based on their assigned territory or territory membership. It is best practice to follow the principle of least privilege when reading records.

Set the *sn_fsm.use_query_rules* system property to **true**.

More information

Attribute	Description
Configuration name	*sn_fsm.use_query_rules*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Access control
Security risk	• Severity score: 4.3 • CVSS score: Medium • Security risk details: There may be increased risk of data exposure from Field Service Management tables.
Dependencies and prerequisites	None
Functional impact	When set to true, rules/filters from sn_query_rule table are used to determine read access to Field service management related tables. For example, Work Order (WO) and Work Order Table (WOT) to the logged-in user through query business rules and READ ACLs. When false, the records aren't filtered based on query rules. Enabling this property secures the data, and all data (wm_task and wm_order) won't be visible to their users.
References	

Enable ACLs to Control Live Profile Details [Updated in Security Center 1.3]

Use the *glide.live_profile.details* property to designate whether a user should be able to view all detail fields, such as company name and phone numbers, in a live profile.

Depending on the setting of the *glide.live_profile.details* property, the following occur:

- If the value is set to Show, access to the live profile information is granted, regardless of the ACLs created for the user profile.
- If the value is set to ACL, access to the live profile information is restricted, as per the ACLs created for the user profile.
- If the value is set to Hide, access to the live profile information is restricted, regardless of the ACLs created for the user profile.

More information

Attribute	Description
Property name	*glide.live_profile.details*

Attribute	Description
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	The purpose is to enable only authorized users to access the details of a Live Profile (such as Company name, Phone numbers)
Data type	choicelist
Recommended value	ACL
Default value	ACL
Security risk rating	4.3
Functional impact	If property is not enabled, unauthorized users can access the Live profile details of all other users.
Security risk	(Moderate) API requests should always honor table ACLs. Restriction must be applied to prevent unauthorized users accessing details of a Live Profile.

To learn more about adding or creating a system property, see Add a system property.

Enable ACLs for Encoded Query in Simple List Widget [New in Security Center 2.0]

Learn how to set the *glide.service_portal.enable_acls_for_encoded_query_in_list* property to the secure value to prevent users from bypassing access control list (ACL) evaluations on a query condition in the Simple List Widget.

When the *glide.service_portal.enable_acls_for_encoded_query_in_list* property is not set to the secure value of true, a user may be able to bypass access control list (ACL) evaluations on a query condition in the Simple List Widget. If the property is set to false, it reverts to previous behavior, enforcing ACL checks for an encoded query based on the enforce_acl checkbox value.

It is a best practice to evaluate ACLs within queries to ensure that a user has access to the fields being queried, thereby preventing unauthorized data leakage.

Ensure that the glide property *glide.service_portal.enable_acls_for_encoded_query_in_list* is set to true. If the property does not exist in the sys_properties table, the default value is true.

More information

Attribute	Description
Configuration name	*com.glide.script.fencing.cross_scope_access.shared*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	true
Default value	true

Attribute	Description
Category	Access control
Security risk	• Severity score: 4.3 • CVSS score: Medium • Security risk details: Not setting this property to the recommended value of true.
Dependencies and prerequisites	None
Functional impact	The Simple List Widget may not display any data depending on the user's role and the underlying ACLs. Additionally, users might encounter security warnings if the Simple List query contains filter conditions with properties that are not accessible to the current user.

Enable URL allowlist for cross-origin iframe communication

Use the *glide.ui.concourse.onmessage_enforce_same_origin* property to enable cross-origin communication between `iframes`.

`Openframe` can only process messages from trusted domains that are specified in the *glide.ui.concourse.onmessage_enforce_same_origin_whitelist* property. To learn more, see Enable URL allowlist for cross-origin iframe communication.

More information

Attribute	Description
Property name	*glide.ui.concourse.onmessage_enforce_same_origin*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Category	Access control
Purpose	To enable inclusion listing of trusted domains, so they can communicate between iframes for openframe.
Recommended value	true
Default value	true
Security risk rating	4.2
Functional impact	If you do not inclusion list intended domains, the ability to embed other pages within ServiceNow AI Platform instances may be limited.
Security risk	(High) If a web page contains event handlers that do not perform proper origin validation, a web page, or script from any origin, can communicate with it. It can also initiate any functionality performed by the event handler.
References	OpenFrame overview

To learn more about adding or creating a system property, see Add a system property.

Enable Anti-CSRF token [New in Security Center 1.3, updated in 1.5, and removed in 2.0]

Use the *glide.security.use_csrf_token* property to ensure the use of a secure token to identify and validates incoming requests, which in turn are used to prevent these attacks.

Cross-Site Request Forgery (CSRF) is an attack that forces authenticated users to submit a request to a Web application against which they are currently authenticated. CSRF attacks exploit the trust a Web application has in an authenticated user.This property enables usage of a secure token to identify and validate incoming requests. This token is used to prevent cross site request forgery attacks. If *glide.security.use_csrf_token* is not set to the recommended value of true, then CSRF is possible.

More information

Attribute	Description
Property name	*glide.security.use_csrf_token*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	To protect the application from potential CSRF attack.
Security risk rating	8.1
Recommended value	true
Default value	true
Functional impact	This remediation enables an extra validation step before the instance user submits a write request to the instance. Every write request contains a CSRF token (i.e a validation/CSRF ID tied to the user session). When the user session expires, the secure token expires with it.
Security risk	(High) Cross Site Request Forgery is a significant security risk that violates the integrity of the instance data. An attacker can launch the CSRF attack by abusing the trust of an instance user. With the help of social engineering attacks, a user can submit a malformed request on behalf of the attacker on the instance.

To learn more about adding or creating a system property, see Add a system property.

Enable contextual security plugin [Updated in Security Center 1.3]

Activate the Contextual Security Plugin (*com.glide.role_management*) plugin to enable contextual security, which secures a record/information using create, read, write, and delete functionality.

After it is installed and activated, the dictionary roles (created by simple security manager) are no longer tested. Instead, the ServiceNow AI Platform looks for ACL rules on fields and tables. It secures the data with the help of ACL rules instead of traditional, role-based dictionary rules implemented by simple security manager. Even if you configure the dictionary form and add roles to a dictionary entry, no change in rights occurs.

More information

Attribute	Description
Plugin ID	*com.glide.role_management*
Configuration type	System Definition > Plugins
Category	Access control
Purpose	Unlike the simple security manager, the contextual security manager is aware of the system table hierarchy. You can potentially have different security rules for a field based on where in the hierarchy it appears.
Recommended value	Active
Default value	There is no default value as this is a plugin, not a Glide property.
Security risk rating	8.1
Functional impact	This remediation enforces functional level of access controls, which would let application determine the access restrictions based on ACL table alone.
Security risk	(High) Functional level access controls must be enforced from the server side prior to executing CRUD operations, ensuring the appropriate level of access to instance users.
References	Contextual Security Manager

Enable Cross Scope Privilege Checks on Service Portal Form [New in Security Center 7.0]

Use a system property to enforce cross scope privilege checks on the Service Portal form widget and prevent unauthorized retrieval of forms and table data between scopes.

In Yokohama and later releases, queries enforce cross scope privilege checks on the table before reading the given sys_id information.

When the *glide.service_portal.enforce_cross_scope_check_in_form* property is set to the recommended value of **true**, cross scope privilege checks are enforced on the table. When set to **false**, the cross-scope privilege check isn't enforced.

Set the *glide.service_portal.enforce_cross_scope_check_in_form* system property to **true** or confirm that the property doesn't exist in the System Properties [sys_properties] table. If a record doesn't exist in the sys_properties table, the value defaults to **true**.

More information

Attribute	Description
Configuration name	*glide.service_portal.enforce_cross_scope_check_in*
Configuration type	System Properties (/sys_properties_list.do)

Attribute	Description
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: 3.1 • CVSS score: Low • Security risk details: A lack of cross scope privilege checks could lead to unauthorized retrieval of forms and table data between scopes.
Dependencies and prerequisites	None

Enforce ACL on HR Lifecycle Events Data [New in Security Center 2.0]

Learn how to prevent unauthorized access to data in the Human Resources Lifecycle Events application by verifying that the *glide.enforce_security_scope.sn_hr_le property* is set to the secured value.

The *glide.enforce_security_scope.sn_hr_le* property limits the access control lists (ACLs) of several HR tables so that only the "sn_hr_le" scope is considered. If *glide.enforce_security_scope.sn_hr_le* isn't set to the recommended value of true, then the data from the Human Resources: Lifecycle Events application will be exposed to ACLs from all other scopes which could lead to unauthorized users accessing sensitive data. For example, an IT administrator gaining access to HR data.

⚠ **Warning:** This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Configuration name	*glide.enforce_security_scope.sn_hr_le*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Access control
Security risk	• Severity score: 6.5 • CVSS score: Medium • Security risk details: Not setting this property to the value of true, could cause the Human Resources:

Attribute	Description
	Lifecycle Events application data to be exposed to ACLs from all other scopes. This could lead to unauthorized users accessing sensitive data.
Dependencies and prerequisites	None

Enforce ACL on HR Core Data [New in Security Center 2.0]

Learn how to configure the *glide.enforce_security_scope.sn_hr_core* property so that the Human Resources Scoped App: Core (com.sn_hr_core) plugin does not expose sensitive data to access control lists (ACLs) from all other scopes.

The *glide.enforce_security_scope.sn_hr_core* property restricts the access control lists (ACLs) of several global data tables like sys_attachment and sys_email to only consider the sn_hr_core scope. If this property is not set to the recommended value of true, then data from the Human Resources Scoped App: Core plugin will be exposed to ACLs from all other scopes. For instance, this could result in the IT administrator gaining access to human resources data.

⚠ **Warning:** This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Configuration name	*glide.enforce_security_scope.sn_hr_core*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: 6.5 • CVSS score: Medium • Security risk details: If this property is not set to the secure value of true, then data from the Human Resources Scoped App: Core plugin will be exposed to ACLs from all other scopes.
Dependencies and prerequisites	None
References	Activate Case and Knowledge Management

Enforce ACL on HR Virtual Agent Data [New in Security Center 2.0]

Discover how to set the *glide.enforce_security_scope.sn_hr_va* property to a secure value, preventing data leakage from the Virtual Agent Conversations scoped application.

The *glide.enforce_security_scope.sn_hr_va* property restricts the access control lists (ACLs) of several human resources (HR) tables to only consider the sn_hr_va scope. If this property is not set to the recommended value of true, then data from the Human Resources: Virtual Agent Conversations scoped application will be exposed to ACLs from all other scopes. For example, this could allow the IT Administrator to access HR data.

More information

Attribute	Description
Configuration name	*glide.enforce_security_scope.sn_hr_va*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Access control
Security risk	• Severity score: 6.5 • CVSS score: Medium • Security risk details: Failing to set this property to the secure value true could expose data from the Human Resources: Virtual Agent Conversations scoped application to ACLs from all other scopes.
Dependencies and prerequisites	None

Enforce application specific ACLs only for application data

Avoid unauthorized or undesired access to application data by enforcing application-specific access control lists (ACLs) only for application data.

Control the behavior of application data residing in primary tables outside the application. When these properties have a value of **true**, only the application-specific ACLs are evaluated for access to the application data residing in these tables. Not all applications are designed to work in this configuration or use a System Property [sys_properties] record for this purpose.

These system properties use the **glide.enforce_security_scope.<scope>** naming format. For example, use the **glide.enforce_security_scope.sn_hr_sp** property for the Employee Center Core (sn_hr_sp) scope. The following application scopes contain this property:

- sn_doc
- sn_egd_act
- sn_egd_core
- sn_egd_goals
- sn_em
- sn_gsm
- sn_gsm_info_req
- sn_hc_professional
- sn_hr_agent_ws
- sn_hr_ai_agents
- sn_hr_awa
- sn_hr_core
- sn_hr_ef
- sn_hr_er
- sn_hr_le_ent
- sn_hr_mii_base
- sn_hr_na_galileo
- sn_hr_pad
- sn_hr_pj
- sn_hr_sp
- sn_hr_va
- sn_ja
- sn_jny
- sn_lg_contracts
- sn_lg_matter
- sn_lg_ops
- sn_opp_market
- sn_professional

- sn_gsm_lic_prmt
- sn_gsm_lic_prmt_ex
- sn_gsm_soc_bnfts
- sn_hr_gen_ai
- sn_hr_hc
- sn_hr_le
- sn_hr_ws
- sn_imt_health_test
- sn_imt_tracing
- sn_imt_vaccine
- sn_svc_appl_info
- sn_svc_appl_pgm_mg
- sn_talent_aia
- sn_uni_req
- sn_uni_task

For each application installed with the *glide.enforce_security_scope* property in the System Properties [sys_properties] table, (for example, *glide.enforce_security_scope.sn_hr_core*), ensure the property value is set to **true**.

Note: These properties can only be modified by the scoped administrator for the specific application. If a sys_properties record does not exist for the given application and respective property, it must be created.

Use this script can to find which properties need to be updated or created on the instance:

```
var properties = [
       'glide.enforce_security_scope.sn_uni_task',
       'glide.enforce_security_scope.sn_uni_req',
       'glide.enforce_security_scope.sn_svc_appl_info',
       'glide.enforce_security_scope.sn_professional',
       'glide.enforce_security_scope.sn_opp_market',
       'glide.enforce_security_scope.sn_lg_ops',
       'glide.enforce_security_scope.sn_lg_matter',
       'glide.enforce_security_scope.sn_lg_contracts',
       'glide.enforce_security_scope.sn_jny',
       'glide.enforce_security_scope.sn_ja',
       'glide.enforce_security_scope.sn_imt_vaccine',
       'glide.enforce_security_scope.sn_imt_tracing',
       'glide.enforce_security_scope.sn_imt_health_test',
       'glide.enforce_security_scope.sn_hr_ws',
       'glide.enforce_security_scope.sn_hr_va',
       'glide.enforce_security_scope.sn_hr_sp',
       'glide.enforce_security_scope.sn_hr_pj',
       'glide.enforce_security_scope.sn_hr_pad',
       'glide.enforce_security_scope.sn_hr_mii_base',
       'glide.enforce_security_scope.sn_hr_le',
       'glide.enforce_security_scope.sn_hr_le_ent',
       'glide.enforce_security_scope.sn_hr_hc',
       'glide.enforce_security_scope.sn_hr_gen_ai',
       'glide.enforce_security_scope.sn_hr_er',
       'glide.enforce_security_scope.sn_hr_ef',
       'glide.enforce_security_scope.sn_hr_core',
       'glide.enforce_security_scope.sn_hr_awa',
       'glide.enforce_security_scope.sn_hr_agent_ws',
       'glide.enforce_security_scope.sn_hc_professional',
       'glide.enforce_security_scope.sn_gsm_soc_bnfts',
       'glide.enforce_security_scope.sn_gsm_lic_prmt_ex',
       'glide.enforce_security_scope.sn_gsm_lic_prmt',
       'glide.enforce_security_scope.sn_gsm_info_req',
       'glide.enforce_security_scope.sn_gsm',
       'glide.enforce_security_scope.sn_em',
       'glide.enforce_security_scope.sn_egd_goals',
```

```
        'glide.enforce_security_scope.sn_egd_core',
        'glide.enforce_security_scope.sn_egd_act',
        'glide.enforce_security_scope.sn_doc',
        'glide.enforce_security_scope.sn_talent_aia',
        'glide.enforce_security_scope.sn_hr_na_galileo',
        'glide.enforce_security_scope.sn_svc_appl_pgm_mg',
        'glide.enforce_security_scope.sn_hr_ai_agents',
        'glide.enforce_security_scope.sn_hr_mii_base'
    ];

    var pm = new GlidePluginManager();

    for (var i = 0; i < properties.length; i++) {
        var property = properties[i];
        var application = property.split('.')[2];
        var propertyValue = gs.getProperty(property, 'false');

        if (pm.isActive(application) &&
propertyValue.toLowerCase() != 'true') {
            gs.print(property);
        }
    }
```

More information

Attribute	Description
Configuration name	*glide.enforce_security_scope.<scope>*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: 4.1 • CVSS score: Medium • Security risk details: When these properties don't have a value of **true**, the ACLs on the primary table are still evaluated for access, potentially allowing unauthorized or undesired access to application data.
Dependencies and prerequisites	None

To learn more about adding or creating a system property, see Add a system property.

Enforce application scope restrictions [New in Security Center 1.3 and removed in 1.5]

Use the *glide.record.legacy_cross_scope_access_policy_in_script* property to control the permissions of scoped apps.

If the *glide.record.legacy_cross_scope_access_policy_in_script* property is set to true, scoped apps can call APIs which should only be available to global apps. This property bypasses the intended access controls for creating and updating developers for those scoped apps.

More information

Attribute	Description
Configuration name	*glide.record.legacy_cross_scope_access_policy_in_s*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	true (when the property does not exist in the sys_properties table.)
Category	Access control
Security risk	• Severity score: 3.5 • CVSS score: Low • Security risk details: If this property is not set to the recommended value, then scoped apps and delegated developers for those scoped apps can create and update records in global tables such as Incident.
Dependencies and prerequisites	None

Enforce Read Roles for Catalog Variable Search [New in Security Center 7.0]

Use system properties to ensure that only catalog variables with an empty read role are indexed for search.

When the property *glide.ais.ingestion.ignore_catalog_variables_read_roles* is set to the recommended value of **false**, only catalog variables with an empty read role are indexed for search. If this property is set to **true** then all variables are indexed for search regardless of any read roles specified on the variable.

Verify that the *glide.ais.ingestion.ignore_catalog_variables_read_roles* system property does not exist in the System Properties [sys_properties] table, or exists and is set to **false**.

More information

Attribute	Description
Configuration name	*glide.ais.ingestion.ignore_catalog_variables_read_*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false

Attribute	Description
Default value	false
Category	Access control
Functional impact	This property affects the indexing of searchable content from catalog variables that require specific roles for read access.
Security risk	• Severity score: 2.6 • CVSS score: Low • Security risk details: Users will be able to search for variables that they do not have read roles to access to causing information disclosure.
Dependencies	None
References	Service catalog variables

Enforce security rules to sharing dashboards [New in Security Center 1.3]

Use the *glide.cms.dashboards.sharing_with_secure_search* property to control whether users can share dashboards.

When the *glide.cms.dashboards.sharing_with_secure_search* property is not set to **true**, users can share dashboard groups and roles that they do not have access to. Enabling this property enforces access control lists (ACLs) when searching the sys_user, sys_user_role, and sys_user_group tables during the dashboard sharing process. Sharing a dashboard excessively can lead to users, groups, or roles accessing data that they should not have permission to view, potentially compromising sensitive information. Therefore, it is recommended to set *glide.cms.dashboards.sharing_with_secure_search* to true so that dashboards are shared only with those who have the appropriate permissions.

More information

Attribute	Description
Configuration name	*glide.cms.dashboards.sharing_with_secure_search*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Access control
Security risk	• Severity score: 3.5 • CVSS score: Low • Security risk details: Not setting this property to the recommended value of true causes access control lists to not be enforced when searching the sys_user,

Attribute	Description
	sys_user_role, and sys_user_group tables. This could lead to dashboards being shared with unauthorized users, exposing sensitive information.
Dependencies and prerequisites	None
References	Access control
Functional impact	This property applies security rules to the list of users, user groups, and roles that are visible when sharing dashboards.

Enforce scope security for public sector digital services [New in Security Center 1.3]

Use the *glide.enforce_security_scope.sn_gsm* property to control how the application data from the Public Sector Digital Services application is accessed.

The ServiceNow Public Sector Digital Services application lets you develop public sector applications that deliver digital services to constituents, businesses, and agencies.

When *glide.enforce_security_scope.sn_gsm* is set to false, access to the application data within the global tables of the Public Sector Digital Services app may be accessible based on the access control lists (ACLs) of those global tables. When this property is set to true, access to data residing in global tables are only evaluated based off the ACLs shipped directly in the Public Sector Digital Services app. Setting this property to false may lead to information disclosure from over permissive ACLs.

To remediate this security risk, set *glide.enforce_security_scope.sn_gsm* to true.

More information

Attribute	Description
Configuration name	*glide.enforce_security_scope.sn_gsm*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: 4.2 • CVSS score: Medium • Security risk details: Not setting this property to the recommended value could lead to information disclosure from over permissive ACLs.
Dependencies and prerequisites	None
References	Configuring Public Sector Digital Services

Enforce scoped ACL access for information request playbooks [New in Security Center 1.3 and updated in 1.5]

Use the *glide.enforce_security_scope.sn_gsm_info_req* property to control access to playbook data for the Information Request playbooks feature.

The Information Request Playbook application enables public sector end users to submit and track public record requests and provides government agents with a pre-defined process for handling and resolving these requests. If *glide.enforce_security_scope.sn_gsm_info_req* is not set to true, unexpected access could be granted to playbook data for the Information Request playbooks application. Set this property to true to only consider ACLs from the sn_gsm_info_req scope when granting access.

More information

Attribute	Description
Configuration name	*glide.enforce_security_scope.sn_gsm_info_req*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: 4.3 • CVSS score: Medium • Security risk details: If this property is set to false, then ACLs from all scopes are considered when granting access to playbook data in the scope master table. This would expose information request playbook data.
Dependencies and prerequisites	None
References	• Using Information Request Playbook • Configure Information Requests service channel

Enforce strict elevate privilege [New in Security Center 1.3]

Use the *glide.security.strict_elevate_privilege* property to control whether roles marked as privileged must be manually elevated for the user to be granted the role's capabilities.

Set this property to true to add an extra layer of security validation when a privileged user elevates their role.

More information

Attribute	Description
Configuration name	*glide.security.strict_elevate_privilege*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: 6.7 • CVSS score: Medium • Security risk details: When *glide.security.strict_elevate_privilege* is set to false, roles marked as privileged are automatically elevated upon an admin user new session, and do not need to be manually elevated (with the exception of security_admin).
Dependencies and prerequisites	None
Functional impact	This property strictly requires admin role users to elevate privileges when needed.

Enforce security scope license and permit playbook [New in Security Center 1.5 and updated in 2.0]

Use this property to determine if only the access control lists (ACLs) within the License and Permit plugin will be used in determining access to the scope, or if ACLs from all scopes will be considered.

When the *glide.enforce_security_scope.sn_gsm_lic_prmt* property is set to the recommended value of true, then only ACLs within the License and Permit plugin are used to determine access to the scope. When this setting is configured to false, then License and Permit Playbooks data in scope master tables are exposed because ACLs from all scopes are granted access. To reduce data exposure, set *glide.enforce_security_scope.sn_gsm_lic_prmt* to the recommended value of true.

More information

Attribute	Description
Configuration name	*glide.enforce_security_scope.sn_gsm_lic_prmt*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true

Attribute	Description
Category	Access control
Security risk	• Severity score: 2.7 • CVSS score: Low • Security risk details: Configuring this setting to the recommended value of true secures the License and Permit Playbooks data in scope master tables by considering only ACLs from *sn_gsm_lic_prmt* scope for granting access. Setting this to false exposes the License and Permit Playbooks data in scope master tables by considering the ACLs from all scopes for granting access. For example, the IT Administrator can access License and Permit Playbooks data when this setting is false.
Dependencies and prerequisites	None
References	• Using License and Permit Playbook • Application scope

Enforce Security Scope for Agent Workspace for HR Case Management [New in Security Center 1.5 and updated in 2.0]

Configure the Agent Workspace for HR Case Management plugin so that data in scope master tables can only be accessed by users with the correct permissions, enforcing the principle of least privilege.

When the *glide.enforce_security_scope.sn_hr_agent_ws* plugin is configured to the recommended value of true, then only the access control lists (ACLs) within the Agent Workspace for HR Case Management plugin are used to determine access to a resource. When this setting is set to false, then Agent Workspace for HR Case Management data in scope master tables are exposed because the ACLs from all scopes are granted access.For example, an IT Administrator can access Agent Workspace for HR Case Management data when this setting is set to false. To prevent this from happening, set *glide.enforce_security_scope.sn_hr_agent_ws* to the recommended value of true which ensures that the principle of least privilege exists as users can only access resources they have permission to.

More information

Attribute	Description
Configuration name	*glide.enforce_security_scope.sn_hr_agent_ws*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control

Attribute	Description
Security risk	• Severity score: 2.7 • CVSS score: Low • Security risk details: Configuring this setting to false causes the Agent Workspace for HR Case Management data in scope master tables to be exposed because the ACLs from all scopes are granted access.
Dependencies and prerequisites	Agent Workspace for HR Case Management
Functional impact	Configuring this setting to true will enforce global ACLs to be executed for a table, if scoped ACLs do not exist for it.
References	• https://owasp.org/www-project-proactive-controls/#div-numbering • Add a component to Agent Workspace

Enforce Security Scope for Service Application Information [New in Security Center 2.0]

Use the *glide.enforce_security_scope.sn_svc_appl* property to ensure that the data in master scope tables is secured.

When the *glide.enforce_security_scope.sn_svc_appl_info* property is set to true, access to resources within the scope is determined solely by the access control lists (ACLs) from the Service Application Information plugin (sn_svc_appl_info). This ensures the security of data in master scope tables by restricting access permissions to those defined within the sn_svc_appl_info scope.

If set to the insecure value of false, ACLs from all scopes are considered when granting access to data in master scope tables such as sys_attachment. This could lead to unauthorized access to sensitive information by users who do not have permissions for the Service Application Information data.

More information

Attribute	Description
Configuration name	*glide.enforce_security_scope.sn_svc_appl_info*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Access control
Security risk	• Severity score: 4.3 • CVSS score: Medium

Attribute	Description
	• Security risk details: If this property is set to the insecure value of false, it can lead to unauthorized access to sensitive data by users who do not have permissions for the Service Application Information data.
Dependencies and prerequisites	The Service Applicant Information plugin (*com.sn_svc_appl_info*) must be activated for the *glide.enforce_security_scope.sn_svc_appl_info* property to be effective.

Enforce field level ACLs in GlideRecordSandbox

Manage field level ACLs in GlideRecordSandbox on your instance.

Use the *glide.sandbox.fields.check_acl* property to enforce field level ACLs in GlideRecordSandbox. An example in which this property is applied is when a user can provide a script, like in *sysparm_query*. If this property is not set to the recommended value of **true**, ACL restrictions can be bypassed, which enables sensitive data to be compromised, such as a *sys_user.user_password* from an unauthorized user.

Warning: The value for this property is a no DB override. It can't be altered or overridden.

More information

Attribute	Description
Configuration name	*glide.sandbox.fields.check_acl*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: 7.5 • CVSS score: High • Security risk details: Setting this property to **false** enables ACL restrictions to be bypassed which could expose sensitive data.
Dependencies and prerequisites	None

Enforce GroupBy ACLs

Configure your instance to conduct ACL checks on groupby columns.

Use the *glide.security.groupby_acl_check* system property to configure your instance to conduct ACL checks on groupby columns. If this property is set to the recommended value of **true**, then ACLs on groupby columns are honored by default. A table's *groupby_acl_check* attribute takes precedent over the

glide.security.groupby_acl_check property. If the property is set to **false**, then ensure that any table which should have ACL checks on groupby columns has the *groupby_acl_check* attribute set to **true**.

Ensure the property *glide.security.groupby_acl_check* is set to **true**.

More information

Attribute	Description
Configuration name	*glide.security.groupby_acl_check*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: 3.7 • CVSS score: Low • Security risk details: Setting this property to **false** will disable ACLs check on groupby columns which could lead to information leakage.
Dependencies and prerequisites	None

Ensure archive table ACLs are checked [New in Security Center 1.3 and updated in 1.5]

The *glide.security.enable_archive_table_acls* property controls whether access control lists (ACLs) of the original table, the table the archive table was created from, are evaluated to false.

The *glide.security.enable_archive_table_acls* property should not be set to false since the original table's ACLs will be evaluated regardless of its value. You can avoid additional ACLs for an archive table by not adding them.

More information

Attribute	Description
Configuration name	*glide.security.enable_archive_table_acls*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control

Attribute	Description
Security risk	• Severity score: 3 • CVSS score: Low • Security risk details: If the property is set to false, ACLs added to archived tables will be ignored, an action that is counter intuitive and therefore may lead to authorization bypass.
Dependencies and prerequisites	None
Functional impact	When this property is set to true, any active read ACLs on archive tables will be honored. If no active read ACLs exist or the property is set to false, the original table's (table from which data was archived) will apply to the archive table. **Note:** Only read ACLs are supported on archive tables. Other operations on archive tables are governed internally through an Access Handler.

Ensure dashboards creation/deletion requires access check [New in Security Center 1.3 and updated in 2.0]

The *glide.processors.check_access_before_process* system property enables access control list (ACL) enforcement for creating or deleting dashboards when a user is logged in.

Set the *glide.processors.check_access_before_process* system property to **true**. If the property does not appear in the System Properties [sys_properties] table, the fallback value is **true**.

More information

Attribute	Description
Configuration name	*glide.processors.check_access_before_process*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Fallback value	true
Category	Access control
Security risk	• Severity score: 6.3 • CVSS score: Medium

Attribute	Description
	• Security risk details: Disabling this property by setting it to false, enables an ACL bypass on dashboards. This allows all authenticated users with low privileges to delete and add dashboards.
Dependencies and prerequisites	None
Functional impact	This property controls the ability to create new sys_dashboards and delete existing dashboards when a user lacks the necessary access rights. When the value is set to false, users with inappropriate roles can add and delete sys_dashboard entries (though the GlideRecord layer should recheck the existing ACLs). A value of true restricts add and delete operations for users without the required access rights.

Exclude Sensitive Tables and Fields from Data Generation [New in Security Center 7.0]

Use system properties to exclude tables and fields from Data Generation, which is used to generate fake data sets based on existing data. Tables and fields that are added to these exclusion lists can't be used for Data Generation feature.

Tables included in the *glide.data.generation.excluded.tables* system property are excluded from Data Generation in addition to metadata tables.

Fields included in a comma separated list for *glide.data.generation.excluded.fields.<TABLE-NAME>* are excluded from Data Generation in addition to any of these fields, if applicable:

- number
- roles
- sys_class_name
- sys_created_by
- sys_id
- sys_mod_count
- sys_tags
- sys_updated_by

Review the list of tables included in the property *glide.data.generation.excluded.tables*. Add any tables which should be excluded from Data Generation to the comma separated list of tables. In addition, metadata tables will be ignored for data generation.

Review the list of fields for each table by looking at properties with the format: `glide.data.generation.excluded.fields.<TABLE-NAME>`. Add any sensitive fields for the specified table as a comma-separated list of values.

More information

Attribute	Description
Configuration name	• *glide.data.generation.excluded.tables* • *glide.data.generation.excluded.fields.**
Configuration type	System Properties (/sys_properties_list.do)
Data type	• Comma-separated list of table names • Comma-separated list of fields for a given table
Recommended value	• Comma-separated list of table names which should be excluded from Data Generation • Comma-separated list of field names which should be excluded from Data Generation for each applicable table
Default value	"","
Category	Access control
Security risk	• Severity score: 2.6 • CVSS score: Low • Security risk details: Data in tables used in the Data Generation feature can have real or fake values populated as intended by the feature. Sensitive values in tables which should not be visible or replicated could be revealed to other instance users.
Dependencies and prerequisites	Data Generation plugin is in use

Prevent Users From Accepting Warning To Bypass CSRF Validation [Updated in Security Center 1.3 and 1.5]

Use the *glide.security.csrf.strict.validation.mode* property to enable CSRF token strict validation. If the CSRF token doesn't match, it prevents resubmission of the request.

This property prevents users from being able to accept a warning which allows a potentially malicious request to be sent to the instance. This warning appears when a POST request fails due to having a mis-matched anti-CSRF token belonging to one of the victim's other active sessions. If *glide.security.csrf.strict.validation.mode* is not set to the recommended value of true, then an attacker can formulate a CSRF attack utilizing a leaked anti-CSRF token from a different active session belonging to the victim. A POST request to an instance contains an anti-CSRF token within "sysparm_ck" or "X-UserToken" which matches the user's current session.

If the anti-CSRF token is instead tied to one of the user's other active sessions, the POST request will return a 302 redirection to security_interceptor.do with a Continue button available to the user when this property is set to false. Clicking this button will re-submit the request to the instance, except it will now having a valid anti-CSRF token. When this property is set to true, the

302 redirection to the security_interceptor.do page will not display a Continue button and the user will not be allowed to resubmit the request.A successful CSRF attack will allow an attacker to effectively perform any operation that the victim is able to perform.

More information

Attribute	Description
Property name	*glide.security.csrf.strict.validation.mode*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	To enforce strict validation of CSRF token, and prevents its reuse.
Data type	Boolean
Recommended value	true
Default value	true
Security risk rating	(Medium) Cross site Request Forgery is a significant security risk that violates the integrity of the instance data. An attacker can launch the CSRF attack on any instance user by abusing the trust of the instance user. With the help of social engineering attacks, a user can submit a malformed request to the instance on behalf of the attacker.
Security risk rating	3.7
Functional impact	This remediation enables an extra validation step before the instance user submits a write request to the instance. It checks whether the current CSRF token has been used previously. If Yes, it prevents submission of further write requests.
Security risk	(Medium) Cross site Request Forgery is a significant security risk that violates the integrity of the instance data. An attacker can launch the CSRF attack on any instance user by abusing the trust of the instance user. With the help of social engineering attacks, a user can submit a malformed request to the instance on behalf of the attacker.

Return to Configure and upload your customer supplied key to upload your wrapped key.

Restrict delegated developers read access [Updated in Security Center 1.3]

If *com.glide.dd_allow_global_access_tables* does not contain the recommended value of wf_activity, wf_activity_definition, wf_workflow, wf_workflow_version, sp_portal, sp_widget, and sp_page, then those tables could be read by a delegated developer. This could provide the delegated developer read access to sensitive information.

More information

Attribute	Description
Property name	*com.glide.dd_allow_global_access_tables*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	wf_activity, wf_activity_definition, wf_workflow, wf_workflow_version, sp_portal, sp_widget, sp_page
Default value	wf_activity, wf_activity_definition, wf_workflow, wf_workflow_version, sp_portal, sp_widget, sp_page
Category	Access control
Security risk	• Severity score: 2.7 • CVSS score: Low • Security risk details: Ensure that *com.glide.dd_allow_global_access_tables* is set to wf_activity, wf_activity_definition, wf_workflow, wf_workflow_version, sp_portal, sp_widget, sp_page.

Require AJAXGlideRecord ACL checking [Updated in Security Center 1.3]

Use the *glide.script.secure.ajaxgliderecord* property to perform access control rule (ACL) validation when server-side records, such as tables, are accessed using GlideAjax APIs within a client script.

From client scripts, it is possible to query arbitrary data from the server using the AJAXGlideRecord (GlideAjax - Client) API, by using a syntax such as a server-side glide record. It is a powerful and useful tool in many deployments.

If you choose to apply Access Control Lists (ACL) to GlideAjax API calls, you can only query data to which the currently connected user has access. For example, if an ESS user who has no rights to read the cmn_location table is logged in, any GlideAjax API call to that table would fail.

If the ServiceNow AI Platform is running without GlideAjax ACL call checking, an API can return information that the currently logged in user could not otherwise access.

Use GlideRecordSecure when querying data to ensure the highest level of security. GlideRecord relies on ACL enforcement through configurations whereas GlideRecordSecure applies stricter security controls. GlideRecordSecure offers a more secure, out-of-the-box solution for handling sensitive data.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.script.secure.ajaxgliderecord*

Attribute	Description
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	Ensure security ACLs are checked and validated even when the records are accessed through Client Side APIs.
Recommended value	true
Default value	true
Security risk rating	8.1
Functional impact	This remediation enforces the ACL relationship with server-side records when the requests are made using the AJAXGlideRecord API calls. If the ACL configuration is not properly configured, then there is potential impact. For more details on its impact, and how to identify it, see Refer to the Audit and review client-side GlideRecord (AJAXGlideRecord) transactions [KB0550828] article in the HI Knowledge Base .
Security risk	(High) Through client scripts, it is possible to query arbitrary data from the server through the GlideAjax API. Server-side resources can be accessed without proper authorization, so using ACL validation helps the application validate the request based on the configured authorization.
Workaround	Ensure that proper ACLs are created for script includes, processors, and other entities used by a GlideAjax (AJAXGlideRecord) API so that it executes under proper authorization. Implement methods like `canRead ()`, `canWrite()`, `canCreate ()`, and `canDelete ()` to perform user authorization before accessing table records using GlideRecord. Another method is to use GlideRecordSecure. The class is inherited from the GlideRecord Server that performs the same functions as GlideRecord, and also enforces ACLs.
References	Apply ACLs to AJAXGlideRecord (client-side Glide record) This property belongs to the same family of properties that secure and restrict execution of scripts originating from the client, such as *`glide.script.allow.ajaxevaluate`*. For more information, see Enable AJAXEvaluate.

To learn more about adding or creating a system property, see Add a system property .

Restrict write access on system fields to admin users [New in Security Center 7.0]

Use the *glide.rest.table_api.admin_only_sys_fields* system property to control write access the fields generated by the system.

The *glide.rest.table_api.admin_only_sys_fields* property controls write access to these fields:

- sys_id
- sys_created_by
- sys_created_on
- sys_updated_by
- sys_updated_on

When this property is set to **true**, only an admin can write to these fields that are system-generated values. When the property is set to **false** or it doesn't exist in the System Properties [sys_properties] table, users who have create or write access to a table can write to these system values using the Table API.

Set the property *glide.rest.table_api.admin_only_sys_fields* to **true** to prevent non-admin users from updating system fields on records.

More information

Attribute	Description
Configuration name	*glide.rest.table_api.admin_only_sys_fields*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Fallback value	false
Category	Access control
Security risk	• Severity score: 2.7 • CVSS score: Low • Users without the admin role could update fields such as sys_created_by or sys_updated_on. This access can impact the integrity of the created or updated record metadata, making records appear as if they were created by another user or updated at a different time.
Functional impact	Any integrations or users without the admin role that use the Table API to create or update records and change these fields are impacted. These requests to the Table API set the query parameter

Attribute	Description
	sysparm_suppress_auto_sys_field to **false** and set those fields in the request body.
Dependencies and prerequisites	None

To learn more about adding or creating a system property, see Add a system property.

Require approval for agent-based Office 365 group membership changes [New in Security Center 7.0]

Enable the approval flow for adding or removing Office 365 group members through the Microsoft 365 group membership AI Agent using a system property.

Use the *sn_itsm_aia.office_365_group_member_approval.required* system property to control whether the approval flow for adding or removing Office 365 group members through the AI agent is on or off. When the approval workflow is enabled, an approval record must be set to approved by a member of the group specified in the *sn_itsm_aia.office_365_group_member_approval.group_id* system property. If the *sn_itsm_aia.office_365_group_member_approval.group_id* property isn't configured, the *Microsoft 365 group member approvers* group is used.

More information

Attribute	Description
Configuration name	*sn_itsm_aia.office_365_group_member_approval.requi*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Fallback value	true
Category	Access control
Security risk	• Severity score: 4.9 • CVSS score: Medium • Any user who can access the Microsoft 365 group membership AI Agent can and add and remove Office 365 group members from the Azure AD group if explicit approval from the specified approvers group isn't required. Ensure that these changes are approved to reduce the risk of elevation of privilege by adding/removing members without proper authorization.
Functional impact	When *sn_itsm_aia.office_365_group_member_approval.requi* is set to **true**, a member of the group specified in *sn_itsm_aia.office_365_group_member_approval.<grou* must approve the Incident requesting addition

Attribute	Description
	or removal of Office 365 group members. If *sn_itsm_aia.office_365_group_member_approval.requi* is set to **false** then no approval is required and the AI Agent can autonomously handle the process of adding or removing members from an Office 365 group.
Dependencies and prerequisites	None

To learn more about adding or creating a system property, see Add a system property.

Prevent impersonating user from viewing application data

Use system properties to prevent an impersonating user from viewing application data.

Prevent admin level from accessing the application specific data belonging to that user when impersonating an account. This permission can be set at the application level by creating a system property specific to the application.

These system properties use the *<scope>.impersonateCheck* naming format (for example *sn_hr_core.impersonateCheck*). Create a system property with a value of **true** to prevent users from accessing the application-specific data belonging to another user when impersonating an account.

Note: Not all applications are designed to work in this configuration or have a System Properties [sys_properties] record for this purpose. The following scopes are configured to work with this property.

- sn_opp_market
- sn_jny
- sn_imt_vaccine
- sn_imt_health_test
- sn_hr_core
- sn_egd_goals
- sn_egd_core
- sn_egd_act
- sn_em
- sn_talent_aia

For each application with the *<scope>.impersonateCheck* property in the System Properties [sys_properties] table, ensure the property value is set to **true**.

Note: These properties can only be modified by the scoped administrator for the specific application.

Use this script to find which properties need to be updated or created on the instance:

```
var properties = [
    'sn_opp_market.impersonateCheck',
    'sn_jny.impersonateCheck',
    'sn_imt_vaccine.impersonateCheck',
    'sn_imt_health_test.impersonateCheck',
    'sn_hr_core.impersonateCheck',
```

```
    'sn_egd_goals.impersonateCheck',
    'sn_egd_core.impersonateCheck',
    'sn_egd_act.impersonateCheck',
'sn_em.impersonateCheck',
'sn_talent_aia.impersonateCheck'
];

var pm = new GlidePluginManager();

for (var i = 0; i < properties.length; i++) {
    var property = properties[i];
    var application = property.split('.')[0];
    var propertyValue = gs.getProperty(property, 'false');

    if (pm.isActive(application) &&
 propertyValue.toLowerCase() != 'true') {
        gs.print(property);
    }
}
```

More information

Attribute	Description
Configuration name	*<scope>.impersonateCheck*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Access control
Security risk	• Severity score: 3.8 • CVSS score: Low • A value of **false** for these properties allows an admin level user to impersonate another user and access application data with the impersonated user's access. This may be undesirable or allow for unauthorized data access in specific application contexts.
Dependencies and prerequisites	None

To learn more about adding or creating a system property, see Add a system property.

Enforce oauth state parameter validation

Configure the *glide.oauth.state.parameter.required* property to prevent your instance from cross-site request forgery (CSRF) attacks.

The *glide.oauth.state.parameter.required* property enables the State parameter to be required in an OAuth request for authorization code flow. The State parameter is a string value that should not contain special characters or be empty. Setting this property to **true**

ensures that an attacker cannot perform Cross-site request forgery (CSRF) attacks during authentication, which protects your instance from attacks from an impersonated user.

More information

Attribute	Description
Configuration name	*glide.oauth.state.parameter.required*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: 4.2 • CVSS score: Medium • Security risk details: Set this property to **true** to ensure that CSRF attacks are prevented.
Dependencies and prerequisites	None

Enforce Strict User Image Upload

Use the *glide.security.strict.user_image_upload* property to enable Access Control for the upload/update of a profile picture when performed on a user record.

This setting opens the possibility of an unauthorized user uploading an image to another user's profile.

- When you set this property to **true**, the table ACLs are enforced when uploading photos, only allowing authorized users to upload an image.
- When you set it to **false**, ACLs are not enforced on image uploads to the Photo field.

More information

Attribute	Description
Property name	*glide.security.strict.user_image_upload*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	To restrict uploading of user image only to authorized users.
Recommended value	true
Security risk rating	3.7
Functional impact	No functionality impact as authorized users are still able to upload images to their user profile.

Attribute	Description
Security risk	(Low) When you set this property to **false**, an authenticated user could upload an image to another user's account without authorization.

To learn more about adding or creating a system property, see Add a system property.

Restrict email domains for external user registration [Updated in Security Center 1.3, 1.5, and 2.0]

Use the *sn_ext_usr_reg.allowed_email_domains* property to list acceptable external email domains.

The **sn_ext_usr_reg.allowed_email_domains** system property defines which email addresses are allowed to self-register to a ServiceNow instance. The format should be a comma separated list of acceptable email domains such as `domain1.com,domain2.com` where emails such as `example@domain2.com` will be accepted. If **sn_ext_usr_reg.allowed_email_domains** is not set with a list of acceptable domains, then users with any email address are allowed to register accounts on the instances. If not defined, malicious actors could perform registration using emails addresses from unwanted domains to gain authenticated access to the instance.

More information

Attribute	Description
Property name	*sn_ext_usr_reg.allowed_email_domains*
Configuration type	System Properties (/sys_properties_list.do), Communities Properties
Category	Access control
Purpose	List email domains to allow user email for registration.
Recommended value	Enter a list of domains in a comma-separated format, for example `domain1.com, domain2.com, domain3.com`. This format works with or without spaces between elements.
Configuration type	String
Security risk	(High) Malicious actors could perform registration using emails addresses from unwanted domains. Ensure that **sn_ext_usr_reg.allowed_email_domains** is not set to an empty value.
Functional impact	Email addresses from domains that are not included in the comma separated list defined in the property aren't allowed to self-register to a ServiceNow instance.
Security risk rating	7.5
References	Communities

To learn more about adding or creating a system property, see Add a system property.

Enable High Security Plugin [Updated in Security Center 1.3]

When you activate the High Security plugin, it creates or updates hundreds of different configurations to control the level of security on your instance. These configurations mitigate many of the top OWASP attacks by enabling strict access control, input validation, and output encoding.

These configurations include:

- Access Control
- Business rules
- System properties
- UI policy action
- script actions
- script includes

Example

Refer to the examples for the following properties:

Property	Topic
glide.ui.escape_all_script	Escape jelly script [Updated in Security Center 1.3 and 1.5]
glide.security.strict.actions	Check UI action conditions before execution
glide.security.csrf_previous.allow	Enable Anti-CSRF token [New in Security Center 1.3, updated in 1.5, and removed in 2.0]
glide.security.csrf.strict.validation.mode	Prevent Users From Accepting Warning To Bypass CSRF Validation [Updated in Security Center 1.3 and 1.5]

More information

Attribute	Description
Plugin Name	com.glide.high_security
Configuration type	System Definition > Plugins - Development
Category	Access control
Purpose	It is mandatory to activate this plugin. It increases the security level of an instance, which reduces the attack surface by mitigating owasp top 10 attacks, including CSRF, XSS, Securing Session Cookies, and File uploads.
Recommended value	Active
Security risk rating	9.8
Functional impact	This plugin enables several system security configurations, which may impact UI and functionality as well.

Attribute	Description
Security risk	(High) Many security configurations are unintentionally left open, which may open the door for some of the critical vulnerabilities.
References	Activating High Security Settings High Security Settings

To learn more about activating a plugin, see Activate a plugin

Honor Admin Override ACLs

The *glide.security.admin.override.accessterm* property controls admins to be unable to override ACL evaluation even where the override should be in effect.

More information

Attribute	Description
Property name	*glide.security.admin.override.accessterm*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	Controls admins to be unable to override ACL evaluation.
Data type	Boolean
Recommended value	True
Default value	True
Security risk	(Low) ACLs are evaluated cumulatively. If there are number of ACLs on any given field and the Admin Overrides option is false (not selected) on one of them, then the effective admin overrides for all the ACLs are considered to be false.
Security risk rating	3.8
References	Access Control List Rules

To learn more about adding or creating a system property, see Add a system property.

Prevent inactive users from logging in [New in Security Center 1.5]

Configure this property to control if inactive users can authenticate on your instance.

When the system property *glide.authenticate.only.allow.active.user.login* is not set to a value of **true**, users in the User [sys_user] table marked inactive can still login to the instance. Users may be marked inactive if they no longer have permission to login (such as during termination from a company). If the property is not set to **true**, then these users may still access the instance and any data they could previously access.

Set the *glide.authenticate.only.allow.active.user.login* property value to the recommended value of **true** (recommended), to ensures the users in the User [sys_user] table marked inactive cannot log in to the instance and are locked out.

More information

Attribute	Description
Configuration name	*glide.authenticate.only.allow.active.user.login*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: 7.5 • CVSS score: High • Security risk details: Not setting this property to the recommended value of true, could enable inactive users like a terminated employee to still be able to access the instance and any data.
Dependencies and prerequisites	None

Prevent Unauthenticated Access to Virtual Agent Embedded Web Client

Learn how to configure the *sn_va_web_client_app_embed* table to block unauthenticated users from accessing embedded web clients.

The UI page, *sn_va_web_client_app_embed*, which is an embedded web client for Virtual Agent, contains the access control lists (ACLs) marked true in the sys_public table out of the box. It has been confirmed that there are use cases where public accessibility is needed however this is not a standard to set it to default publicly accessible.

If the embedded web client is not needed for unauthenticated users, open the *sn_va_web_client_app_embed* record (sys_id `04b1905473222300e985658b4cf6a7ef`) in the Public Pages [sys_public] table and deselect the **Active** field to deactivate the page.

More information

Attribute	Description
Configuration name	**sn_va_web_client_app_embed**
Configuration type	UI Page(sys_ui_page_list.do)
Data type	table
Recommended value	The **sn_va_web_client_app_embed** public page [sys_public] (sys_id `04b1905473222300e985658b4cf6a7ef`) does not exist or is not active.

Attribute	Description
Default value	Not available (this is a table value)
Category	Access control
Security risk	• Severity score: 7.5 • CVSS score: High • Security risk details: It is recommended to deactivate the UI page, *sn_va_web_client_app_embed*, if an embedded web client is not needed for unauthenticated users.
Dependencies and prerequisites	None

Restrict JSONP Requests to Trusted URLs [Updated in Security Center 1.3]

Specify trusted URLs for the AngularJS $http service to allow or reject JSONP requests.

Increase security on your instance by ensuring that only trusted URLs for the AngularJS $http service can allow/reject JSONP requests. JSONP requests are allowed to any URL if these properties are not configured and enabled.

Use the value of the **angular.jsonp.inclusion_list.urls** system property to define a list of URLs that are trusted and allow for this purpose. Set the value of the **angular.jsonp.inclusion_list.enabled** system property to **true** to limit allowed JSONP to only the URLs listed in **angular.jsonp.inclusion_list.urls**.

More information

Attribute	Description
Configuration name	*angular.jsonp.inclusion_list.enabled*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: Medium • CVSS score: 5.4 • Security risk details: Setting this property to **false** enables JSONP requests to any URL.
Dependencies and prerequisites	None

Prevent users from accepting warning to bypass CSRF validation

Reduce the risk of Cross-Site Request Forgery (CSRF) by preventing users from accepting warning to bypass CSRF validation.

Enable CSRF token strict validation to prevent Cross-Site Request Forgery (CSRF) tokens from being reused, which may allow CSRF attacks.

Set the **glide.security.csrf.strict.validation.mode** system property value to **true** to enable CSRF token strict validation. If this property doesn't exist on your System Properties [sys_properties] table, the default value is **true** starting in Xanadu.

More information

Attribute	Description
Technical configuration name	glide.security.csrf.strict.validation.mode
Plugin applicability	None
Security risk	Cross-site request forgery is a significant security risk that violates the integrity of the instance data. An attacker can launch the CSRF attack on any instance user by abusing the trust of the instance user. With the help of social engineering attacks, a user can submit a malformed request to the instance on behalf of the attacker.
Common Vulnerability Scoring System (CVSS) score	3.7
Common Vulnerability Scoring System (CVSS) rating	Low
Functional impact	This remediation enables an extra validation step before the user submits a write request to the instance. It checks whether the current CSRF token has been used previously. If it has, then it prevents submission of further write requests.
Dependencies and prerequisites	None
Data type	Boolean
Base system value	true
Fallback value	true
Recommended value	true

To learn more about adding or creating a system property, see Add a system property.

Disable raw database query execution [Updated in Security Center 1.3 and removed in 2.0]

Control whether a user can perform raw SQL queries on the database.

The *glide.db.allow_unsafe_dbi_execute_sql* property enables users to perform raw SQL queries on the database, which can give access to tables and data outside of GlideRecord restrictions. If this property is not set to the recommended value of *false*, this allows for the calling of *dbi.executeStatement()* from a Glide Scriptable which can lead to malicious SQL statements being executed.

Warning: This property is both safe and no db override.

More information

Attribute	Description
Configuration name	*glide.db.allow_unsafe_dbi_execute_sql*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	false
Category	Access control
Security risk	• Severity score: 7.2 • CVSS score: High • Security risk details: Not setting this property to **false** enables calling of *dbi.executeStatement()* from a Glide Scriptable.
Dependencies and prerequisites	None
References	Access Control List Rules

Hide user comments on articles [New in Security Center 1.3]

Use the *glide.knowman.show_user_feedback* property to control whether feedback comments are visible.

When *glide.knowman.show_user_feedback* is not set to never, feedback comments will be visible on knowledge base (KB) articles to users with roles defined in the Glide property, *glide.knowman.show_user_feedback.roles*. As feedback comments may contain sensitive information, you may not want the feedback to be visible. If this property is not set to never, there could be confidentiality impacts if sensitive information is disclosed in feedback.

More information

Attribute	Description
Configuration name	*glide.knowman.show_user_feedback*
Configuration type	System Properties (/sys_properties_list.do)
Data type	choicelist
Recommended value	never
Default value	onload
Category	Access control

Attribute	Description
Security risk	• Severity score: 3.5 • CVSS score: Low • Security risk details: Not setting this property to never could lead to sensitive information being disclosed in feedback comments.
Dependencies and prerequisites	None
Functional impact	Shows user comments on KB articles based on choices mentioned in the configuration.

Require authentication by default for client-callable script includes [Updated in Security Center 1.3]

By default, client-callable script includes that do not explicitly set visibility, are public. If needed, add the *glide.script.ccsi.ispublic* property to enable privacy control over all client-callable script includes accessed by public pages.

When you add this property, you must set its value to **false**, which designates that all client-callable script includes are private, and changes their visibility in public pages.

Note: You cannot add the property with a value of **true**, or change its value from **false** to **true**. If you attempt to do so, an error message appears.

If needed, you can change the privacy setting for an individual client-callable script include by adding the `isPublic()` function.

- The `isPublic()` setting takes precedence over the *glide.script.ccsi.ispublic* property.
- For example, if you set `isPublic()` to **true** in an individual script, it makes it public, which overrides the *glide.script.ccsi.ispublic* property that makes all other client-callable script includes private.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.script.ccsi.ispublic*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	Making client-callable script includes private means that guests who access public pages can't access the client-callable script include. A non-logged-in user can't execute a private script.
Recommended value	false
Default value	false

Attribute	Description
Security risk rating	7.5
Functional impact	If the client-callable script includes are designated as public (that is, this property is missing), then unauthenticated users can execute client scripts. Add the property restricts the execution of scripts by a non-logged-in user.
Security risk	(High) If you do not add this property, client-side script includes circumvent ACLs, which may result in unintended public functionality. If the client script provides confidential information, it could have an adverse potential security risk.
Workaround	Setting the *glide.script.ccsi.ispublic* property to **false** makes all client-callable script includes private. You can change the privacy setting for an individual client-callable script include by adding the `isPublic()` function. The `isPublic` function takes precedence over the *glide.script.ccsi.ispublic* property. Add the following syntax to the script include: `isPublic:function(){return[true/false];},`

Enforce production instance behavior [Updated in Security Center 1.3 and 1.5]

Configure whether your instance should be handled like a production or non-production instance.

If the *glide.installation.production* property is not set to the recommended value of **true**, the instance is not treated as a production instance, therefore permitting zboot and other potentially dangerous scripts to run. Enabling a production instance to be evaluated as a non-production instance can lead to information leakages or Denial of Service (DoS) attacks.

More information

Attribute	Description
Configuration name	*glide.installation.production*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: 6.3 • CVSS score: Medium

Attribute	Description
	• Security risk details: Setting this property value to **false** causes the instance to be treated as a non-production instance which enables zboot and other potentially dangerous scripts to run.
Dependencies and prerequisites	None

Restrict access to background script [Updated in Security Center 1.3 and 2.0]

Configure the *glide.script_processor.admin* property to set the role required for accessing the Script Background module.

This property holds the required role to access Script Background module. If *glide.script_processor.admin* is not set to the recommended and default value of admin, then users having a lower privileged role will be able to run background scripts on the instance. This will lead to a complete bypass of the ACL system allowing full access to tables.

Ensure the property *glide.script_processor.admin* is set to the admin. This is the default value on instances.

⚠ **Warning:** This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Configuration name	*glide.script_processor.admin*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	admin
Default value	admin
Category	Access control
Security risk	• Severity score: 8.8 • CVSS score: High • Security risk details: Not setting this property to the recommended value of **admin** lets any user run background scripts on the instance.
Dependencies and prerequisites	None

Restrict access to emails with empty target table

Activate the *glide.email.email_with_no_target_visible_to_all* property to restrict user access to emails, unless they were the one who sent the email or have an admin role.

Unauthorized users are able to access emails in the sys_email_list table that are missing a target record. Instead of enforcing ACLs on email entries, this property restricts access only to the email sender and users with the admin role.

Note: Emails sent to and received by the instance appear in the sys_email_list table. However, only received emails that were marked with an Error and Ignored state should have an empty target table.

More information

Attribute	Description
Property name	*glide.email.email_with_no_target_visible_to_all*
Configuration type	System Properties (/sys_properties_list.do)
Configure	Access control
Purpose	To block email client from showing emails when user doesn't authorize access.
Recommended value	true
Default value	true
Security risk rating	6.5
Functional impact	Users are no longer able to see emails where target table is empty unless they are an admin or were the sender of the email.
Security risk	(Moderate) If the property is not enabled, unauthorized users are able to access any email where the target_table field is empty.
References	Advanced email properties https://support.servicenow.com/kb_view.do?sysparm_article=KB0690043

To learn more about adding or creating a system property, see Add a system property.

Restrict access to specific IP ranges plugin [Updated in Security Center 1.3]

Use the *com.snc.ipauthenticator* plugin to restrict access to specific IP ranges. Unless public access is intended for the instance, administrators should limit access to their assigned IP net blocks.

Prerequisites

This plugin when set to true restricts access to specific IP ranges. Unless public access is intended for the instance, administrators should limit access to their assigned IP net blocks. An exclusion list (Deny) or an inclusion list (Allow) of IP addresses can be created through IP Address Access Control (ip_access_list.do).

Before setting this property, you must activate the IP Range Based Authentication (com.snc.ipauthenticator) *com.snc.ipauthenticator* plugin. To learn more, see IP range based authentication and in the Steps to configure section (below).

Ensure the plugin *com.snc.ipauthenticator* is activated and there is at least one active IP access policy in the table ip_access.

More information

Attribute	Description
Plugin Name	• com.snc.ipauthenticator • ip_access
Configuration type	System Security > IP Address Access Control
Category	Access control
Purpose	To add the range of IP address that can or can't access the instance to the trusted and untrusted domain lists.
Recommended value	Active
Default value	None. This is a plugin, not a Glide property; therefore, there is no default value.
Security risk rating	5.3
Functional impact	Customer-denied IP ranges are used for this remediation item. No impact as customer defines the target list.
Security risk	(Low) Unnecessary exposure to the target instance on the internet should be restricted with the help of IP access controls functionality.
References	IP range based authentication

Steps to configure

1. Ensure that the com.snc.ipauthenticator plugin is active.
2. Navigate to **System Security** > **IP Address Access Control**.
3. Click **New** to create an exclusion list (Deny) or an inclusion list (Allow) of IP addresses.
4. Click **Submit**.

Restrict knowledge bases access [New in Security Center 1.3]

The *glide.knowman.block_access_with_no_user_criteria* property is used to control the read/write access of users on knowledge based articles.

More information

Attribute	Description
Configuration name	*glide.knowman.block_access_with_no_user_criteria*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false

Attribute	Description
Category	Access control
Security risk	• Severity score: 9.1 • CVSS score: Critical • Security risk details: If this property is not set to the recommended value of true, then any user can read and contribute to a KB.
Dependencies and prerequisites	None
Functional impact	Denies access to a knowledge base when either Can Read or Can Contribute isn't specified.

Restrict permissions for CMDB model [Updated in Security Center 1.3 and 1.5]

Use the *csm_cmdb_model.customer_visible_flag* system property to limit customer access to data in the Product Models table as an additional access control to the CMDB model.

Set the *csm_cmdb_model.customer_visible_flag* property to **true** to enable the Customer Visible field for the tables listed below:

- Product Models table [cmdb_model]
- Software Models table [cmdb_software_product_model]
- Application Models table [cmdb_application_product_model]
- Consumable Models table [cmdb_consumable_product_model]
- Facility Models table [cmdb_facility_product_model]
- Hardware Models table [cmdb_hardware_product_model]

Setting this property as **true** hides all cmdb_model values by default.

Set the property to **false** to not consider the customer_visible column/atrribute on the cmdb_model table and to rely on the bases cmdb_model ACLs which are accessible to sn_esm_user.

More information

Attribute	Description
Property name	*csm_cmdb_model.customer_visible_flag*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	When set to tue**true**, the system uses the setting in the Customer Visible field to determine access to product model data on the Customer Service Portal.
Recommended value	true
Default value	false
Configuration type	Boolean

Attribute	Description
Security risk	(Moderate) Any user with the sn_esm_user role and out of the box ACLs could have permissions to the CMDB model. **Note:** this role tends to be granted to external users. External users could unwillingly be given permissions to the CMDB model.
References	Limit access to product model data on the Customer Service Portal

To learn more about adding or creating a system property, see Add a system property.

Restrict unauthenticated access to attachments

Use the *glide.image_provider.security_enabled* property to control the security settings for images. If set to **true**, images are visible only to authenticated and authorized users. If set to **false**, images are visible to anyone with a URL to the attachment.

Secure the images on your instance to prevent sensitive information leak. Images on your instance are accessible via urls that end in `.iix`.

Set the **glide.image_provider.security_enabled** system property to **true** to prevent access to your images via these URLs.

Note:

This property is not honored for images from the attachment table if the origin table is one of:

- Stationeries [sysevent_email_style]
- Welcome Page Sections [sys_home]
- System Properties [sys_properties]

Restriction should be applied for unauthenticated users as some attachments might contain sensitive information.

More information

Attribute	Description
Property name	*glide.image_provider.security_enabled*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	To prevent unauthenticated access of attachment when rendered using the *.iix* format.
Recommended value	true
Default value	false
Functional impact	No significant impact on the functionality. User experience might be affected a bit because the user

Attribute	Description
	who formerly directly accessed .iix must go through authentication.
Security risk	• Severity Score: 6.5 • Security Risk Details: Restriction must be applied for unauthenticated users as some attachment might contain sensitive information.
References	Administering attachments Available system properties

To learn more about adding or creating a system property, see Add a system property.

Restrict access to custom journal entries [Updated in Security Center 1.3 and removed in 2.0]

Use the *glide.live_feed.custom_journal.acl_check_enabled* property to respect ACL's on custom journal fields.

If *glide.live_feed.custom_journal.acl_check_enabled* is not set to the recommended value of true, then all users to be able to see all journal entries within the life feed feature. Setting property to true will respect ACL's on custom journal fields which is a good to have feature.

More information

Attribute	Description
Property name	*glide.live_feed.custom_journal.acl_check_enabled*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	To control which users see which journal entries based on ACLs.
Recommended value	true
Default value	true
Configuration type	Boolean.
Security risk	(Moderate) When set to true, only the custom journal entries which pass ACL will be shown in Live Feed, otherwise, all users will be able to see all journal entries.

To learn more about adding or creating a system property, see Add a system property.

Restrict flow context read access [New in Security Center 1.5]

Use the *com.snc.process_flow.reporting.require_flow_access* property to enforce if an additional access check is required for a user to read a flow check.

When the *com.snc.process_flow.reporting.require_flow_access* property is set to the recommended value of true, there is an additional access check for a user trying to read a flow context. There may be minor information disclosure if this property is set to false.

More information

Attribute	Description
Configuration name	*com.snc.process_flow.reporting.require_flow_access*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: 2.7 • CVSS score: Low • Security risk details: Setting this property to false retains its existing behavior. Setting this property to true enforces the added security layer of read access.
Dependencies and prerequisites	None
Functional impact	When this property is enabled, the security for reading flow context records are increased. The instance enforces that a user trying to read the flow context has read access to the parent flow as well.

Restrict Impersonation to Admin [New in Security Center 2.0]

The *glide.sys.permissive.impersonate* property can be used to prevent non-admin roles from impersonating other users.

When the *glide.sys.permissive.impersonate* property is set to false, only users with the admin role can impersonate other users. When this property is set to true, users may be able to make use of application components that expose impersonation APIs to impersonate a user of higher privilege. This could result in unauthorized access if these application components are misconfigured because non-admin users can access the Impersonation functionality.

You may want to set the property to the non-default value when you need non-admin users to have the capability to impersonate other users.

⚠ **Warning:** This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Configuration name	*glide.sys.permissive.impersonate*
Configuration type	System Properties (/sys_properties_list.do)

Attribute	Description
Data type	Boolean
Recommended value	false
Default value	false
Category	Access control
Security risk	• Severity score: 6.7 • CVSS score: Medium • Security risk details: Failing to set this property to the recommended value of false may allow a non-admin user to utilize application components that expose APIs, enabling them to impersonate a user with higher privileges.
Dependencies and prerequisites	None
Functional impact	Non-admin users can access Impersonation features with some customizations to other scripts and UI pages. However, it is essential to ensure that only the correct users are granted access to these features. **Note:** When *glide.sys.permissive.impersonate* is set to true, Non-admin users with the *impersonate* role can still impersonate.

Enable security jump start plugin (ACL Rules) [Updated in Security Center 1.3]

Activate the Security Jump Start (ACL Rules) (com.snc.system_security *com.snc.system_security*) plugin to create several important ACLs that validate the Access Controls on some of the key system tables within the ServiceNow AI Platform.

These rules provide a jump-start on securing many system tables, making it easier for an organization to get an instance into production. The Security Jump Start (ACL Rules) plugin is installed automatically on all new instances.

More information

Attribute	Description
Plugin ID	*com.snc.system_security*
Configuration type	System Definition > Plugins
Category	Access control
Purpose	Activate the Security Jump Start (ACL Rules) plugin to achieve proper security compliance. It provides some basic ACLs that secure system tables in the first place instead of creating manually for each system table that comes with default provisioning of

Attribute	Description
	an instance. These ACLs are helpful when the newly created instance must quickly move into production.
Recommended value	Active
Default value	None. This is a plugin, not a Glide property, so there is no default value. The plugin is installed by default on zBoot (resets).
Security risk rating	8.1
Functional impact	There is significant functional impact if this plugin is installed without auditing of the existing ACLs on the instance. Customer outreach and definitions are required before the remediation can occur.
Security risk	(High) Access control should be enforced to lock down the unintended access to the instance. ACL jumpstart rules were created to provide a starting point on securing many system tables to make it easier for an organization to quickly get into production.
References	Security jump start - ACL rules

Steps to configure

If this plugin is not activated on your instance, contact ServiceNow Support. Activating the plugin at this point might modify security access to tables already in use in a production environment. If an administrator is interested in the new ACL rules the plugin provides, you can manually create one or more of them in an existing instance if needed. This list of ACLs may be used as a guideline in that case.

Use of secure insert multiple operation within import set API [New in Security Center 1.3]

Use the *com.glide.import_set_api.insert_multiple_optimize* property to control whether *GlideRecordSecure* or *GlideRecord* is used for the Insert Multiple operation within the Import Set API.

If *com.glide.import_set_api.insert_multiple_optimize* is set to the recommended value of false, then *GlideRecordSecure* will be used to insert records, and table-level access control lists (ACLs) will be evaluated. If this property is set to true, *GlideRecord* will be used to insert records, and table-level ACLs will not be evaluated; In addition, you must ensure that Import Set API Insert Multiple REST Endpoint ACL (sys_id: 3101b770ff2211105cf343d0653bf182) is active and that users have the role, import_transformer.

More information

Attribute	Description
Configuration name	*com.glide.import_set_api.insert_multiple_optimize*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	false

Attribute	Description
Category	Access control
Security risk	• Severity score: 6.5 • CVSS score: Medium • Security risk details: If this property is not set to false, a low-privileged user could insert data into tables outside the scope of their privileged roles.
Revertible behavior	Both safe override and No DB Override.
Dependencies and prerequisites	None
Functional impact	This property optimizes the performance of Import Set API by using GlideRecord to save data. When the parameter is set, it requires an integration user to have the import_transformer role to access the API.
References	https://developer.servicenow.com/blog.do?p=/post/gliderecord-vs-gliderecordsecure

Enforce SOAP request strict security [Updated in Security Center 1.3]

Use the *glide.soap.strict_security* property to enforces web service security.

This property uses a combination of:

- Basic authentication challenge/response over the HTTP protocol and
- System level access controls in the Enable security jump start plugin (ACL Rules) [Updated in Security Center 1.3].

If you set this property to **true**, it performs the following actions:

- If the user has appropriate role to perform the operation, it checks incoming SOAP request for role authorization to validate. It occurs during SOAP web service calls/requests made against ServiceNow AI Platform tables when performing CREATE, READ, UPDATE or DELETE operations.
- Checks the system-level ACLs while retrieving data in the form of SOAP data on the table.
- Checks the field-level ACLs for any CRUD operation performed against a field of table.

ACL checks are only complete for standard Table API calls and not web services.

More information

Attribute	Description
Property name	*glide.soap.strict_security*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Default value	true
Recommended value	true

Attribute	Description
Functional impact	This remediation enforces the system-level access control while retrieving data from tables/pages in the form of SOAP data on the instance. If there are users currently accessing this data, they are restricted/allowed to access the data based on the ACL rules. For the default roles that have access to the SOAP data, see SOAP web service.
Security risk	(Moderate) Without appropriate authorization configured on the incoming SOAP requests, an unauthorized user can get access to sensitive content/data on the target instance.
References	Enforce strict security for inbound SOAP SOAP web service

To learn more about adding or creating a system property, see Add a system property.

Required jms connection factories [New in Security Center 1.3 and updated in 1.5 and 2.0]

The *mid.property.jms.command.allowed_factory_names* property controls the Java Messaging Service (JMS) connection factories that the MID Server can use.

It is intended for a few select factories needed by plugins for JMS activity or action. Including additional factories could be a step in a chain of attack for vulnerabilities such as JDNI insertion that rely on capabilities an attacker can leverage in allowed factories. To prevent the possibility of any leveraged vulnerability, do not include factories beyond the necessary defaults.

To remediate this security risk review the list of names provided to the mid property, *mid.property.jms.command.allowed_factory_names*. Ensure any additional Java factory names beyond the default of connectionFactory, queueConnectionFactory, and topicConnectionFactory are necessary.

More information

Attribute	Description
Configuration name	*mid.property.jms.command.allowed_factory_names*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Default value	*connectionFactory*, *queueConnectionFactory*, *topicConnectionFactory*
Recommended value	*connectionFactory*, *queueConnectionFactory*, *topicConnectionFactory*
Category	Access control

Attribute	Description
Security risk	• Severity score: 4.1 • CVSS score: Medium • Security risk details: If the MID Server (*com.glideapp.agent*) plugin is active, review the list of names provided to the mid property *mid.property.jms.command.allowed_factory_names*. Ensure any additional factory names beyond the default of connectionFactory, queueConnectionFactory, and topicConnectionFactory are necessary.
Dependencies and prerequisites	None

Restrict Global App Development by Role [New in Security Center 2.0]

Use the *sn_g_app_creator.allow_global* property to control which users can create applications in the global scope using the Guided Application Creator.

If *sn_g_app_creator.allow_global* is set to the recommended value of false, users require the *sn_g_app_creator.global* role to create applications in the global scope. Conversely, if set to the insecure value of true, any user with the basic sn_g_app_creator.app_creator role can create global applications. Global applications lack scope protection, allowing developers access to extensive features and functions beyond specific scopes. Restricting global application development to users with the additional role adheres to the principle of least privilege.

Note: This property does not come pre-configured in your instance. You must manually create and configure this property according to your organization's needs.

More information

Attribute	Description
Configuration name	*sn_g_app_creator.allow_global*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	false
Category	Access control
Security risk	• Severity score: 3.3 • CVSS score: Low • Security risk details: Failing to set this property to the recommended value could allow any user with the s*n_g_app_creator.app_creator* role to create global applications, which does not adhere to the principle of least privilege.

Attribute	Description
Dependencies and prerequisites	None
Functional impact	Enhanced the API (/api/now/templates) to validate the create global application ACL and property.

Review extraneous explicit role access control conditions [Removed in Security Center 1.5]

The Explicit Roles plugin is recommended to mandate that all users have either the snc_internal role to access internal resources, or the snc_external role to access external resources.

After the installation of this plugin, all existing users are assigned the snc_internal role, and existing access control lists (ACLs) are populated with the role conditions. Due to automation logic or intervention by an instance admin, the snc_internal or snc_external roles may be incorrectly added to an ACL that already contains a more strict role requirement. Since ACL role evaluation will pass for any user containing any role mapped to an ACL, the addition of snc_internal or snc_external may be too broad for the intended purpose of an ACL. This could lead to data leakage if a low privileged user is granted access through the ACL.

For example, it would be unnecessary for both the snc_internal and the admin roles to be mapped to the same ACL within a table. The ACL is meant to grant access to admins, in which case the snc_internal role is a mistake. Or, the ACL is meant to grant access to all snc_internal users which makes the admin role unnecessary. When the Explicit Roles plugin is installed, review the ACLs which contain a role condition for snc_internal or snc_external while also containing a condition for another role. If the roles are able to function for a specific use case, then the finding should be periodically reviewed.

Important: This hardening setting will be removed in the next Security Center v1.5 store patch release and future versions. An Instance Scan suite called "Explicit Roles ACL Config Check Suite" is available in the Washington release. We recommend that you review the findings of this new Instance Scan.

Set guest user for soap requests [Updated in Security Center 1.3 and 2.0]

Configure this property to control the level of access of unauthenticated SOAP requests.

This property controls the level of access of unauthenticated SOAP requests. If it is not set to the recommended value of *`soap.guest`*, or is set to a user with limited privileges, then SOAP requests will execute on behalf of the user. If this property is blank, then it enables unauthenticated access to administrator or maintenance level operations which negates all security controls within the instance.

More information

Attribute	Description
Configuration name	*`com.glide.soap.guest_user`*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	soap.guest
Default value	soap.guest
Category	Access control

Attribute	Description
Security risk	• Severity score: 8.1 • CVSS score: High • Security risk details: Setting this property to blank enables unauthenticated access to administrator or maintenance level operations.
Dependencies and prerequisites	None

Disable public access to favorites [Updated in Security Center 1.3 and 2.0]

Use the *glide.ui.magellan.favorites.allow_public* to specify whether unauthenticated users are allowed to see **Favorites** in the navigator.

Public Access to Favorites is compliant when *glide.ui.magellan.favorites.allow_public* is set to **false**.

More information

Attribute	Description
Property name	*glide.ui.magellan.favorites.allow_public*
Configuration type	System Properties (/sys_properties_list.do)
Configure in Instance Security Center	Yes
Purpose	Control if unauthenticated users are allowed to see **Favorites** in the navigator.
Type	true/false
Recommended value	*false*
Security Dependencies	Set *glide.ui.magellan.favorites.allow_public* to **false**.
Functional impact	(Medium) Enabling this property acts as a layer of protection from unauthorized users.
Security risk	(Medium) If this property isn't enabled, there's a risk of unauthorized access to sensitive data.
References	Configure favorites option

Enable SNC access control plugin [Updated in Security Center 1.3]

Activate the SNC Access Control (com.snc.snc_access_control) plugin to control access to your instances by Customer Service and Support personnel.

Prevent ServiceNow Customer Service and Support personnel from accessing the instances without your express permission by enabling the SNC Access Control (com.snc.snc_access_control) plugin. Although all access to your instance is audited, you may prefer to control this access. This access method is fully auditable and tracked.

Note: Other authorized ServiceNow Operations personnel, in their capacity to support and manage the product, are required to perform administrative actions on the underlying infrastructure. Enabling this plugin may affect support service levels and the Availability SLA. Availability SLA is then measured from the time that Support staff personnel are granted access to your instance.

Enable the SNC Access Control (com.snc.snc_access_control) plugin to restrict access to your instance without your express permission. For more details on this feature, see ServiceNow access control. For activation information, see Activate ServiceNow access control

More information

Attribute	Description
Plugin Name	*com.snc.snc_access_control*
Configuration type	System Definition > Plugins
Category	Access control
Purpose	Prohibits Customer Service and Support employees from accessing the instance
Recommended value	Active
Default value	None. This is a plugin, not a Glide property, so there is no default value. The plugin is not installed by default.
Role required	The customer administrator can't activate this plugin. It must be explicitly requested because it requires elevated privileges to activate the plugin.
Security risk rating	3.3
Functional impact	If this plugin is inactive, all Customer Service and Support employees can access the customer's instance. Enabling the plugin enables the customer to restrict access to authorized Customer Service and Support employees only.
Security risk	(High) Unnecessary exposure of instance access to wider group of people.
References	ServiceNow access control

Steps to configure

1. To request the plugin, follow the steps in Activate ServiceNow access control. Customers must request the SNC Access Control plugin (com.snc.snc_access_control) from HI.
2. To enable SNC access control, follow the steps in Configure ServiceNow access control. Configure an access control record to specify one or more Customer Service and Support employees that have permission to log in your instance.

Use Document Classification to limit publicly accessible documents [New in Security Center 7.0]

Control public access to permalinked documents using system properties.

Note: This hardening setting is not a part of the hardening baseline. It does not appear in Security Center hardening pages and affect your hardening score.

By default, the Document Management plugin includes these classifications:

- **public**
- **restricted**
- **confidential**
- **none** (no classification provided)

Use the *com.snc.documents.permalink.allowed_classifications* to create a list of document classifications. Documents under these classifications (and documents without classification) are publicly accessible to any unauthenticated users with an appropriate link.

Use this property to control public access to document permalinks. Previously, these links were publicly accessible to anyone with the link. The value of this property depends on your specific needs. You may have additional custom document categories that may need to be added to this property to enable public access.

Set the *com.snc.documents.permalink.allowed_classifications* property value to a comma-separated list of document classifications. Documents with these classifications are publicly accessible by unauthenticated users.

If this property isn't present in the System Properties [sys_properties] table, it defaults to an empty list. In this case, only documents with no classification are publicly accessible by unauthenticated users. The default value of this property is **public**, meaning that only documents classified as **public** or documents with no classification are accessible to unauthenticated users.

More information

Attribute	Description
Configuration name	*com.snc.documents.permalink.allowed_classificatio*
Configuration type	System Properties (/sys_properties_list.do)
Data type	String list
Recommended value	public
Default value	public
Fallback value	<empty>
Category	Access control
Security risk	• Severity score: 5.3 • CVSS score: Medium • If a document category containing documents which shouldn't be publicly accessible is added to this property, it may expose sensitive information to unauthenticated users.
Functional impact	If a document category containing documents which should be publicly accessible isn't added to this property, it prevents those documents from being accessed.
Dependencies and prerequisites	None

To learn more about adding or creating a system property, see Add a system property.

Validate query ACLs on Glide DB functions [New in Security Center 7.0]

Control whether query ACLs are applied to Glide DB functions using system properties.

When the *glide.db.encoded_query.check_function_field_query_acls* system property is set to **true**, query ACLs (query_range and query_match) are validated by default on these Glide DB functions:

- glidefunction:position
- glidefunction:substring
- glidefunction:concat
- glidefunction:coalesce
- glidefunction:length

To validate these query ACLS on additional functions, add those functions to the *glide.db.encoded_query.force_query_range_on_functions* system property.

Set the *glide.db.encoded_query.check_function_field_query_acls* system property to **true**, or confirm that the property isn't on the System Properties [sys_properties] list.

More information

Attribute	Description
Configuration name	• *glide.db.encoded_query.check_function_field_quer* • *glide.db.encoded_query.force_query_range_on_func*
Configuration type	System Properties (/sys_properties_list.do)
Data type	• Boolean • String list
Recommended value	• true • glidefunction:position, glidefunction:substring, glidefunction:concat, glidefunction:coalesce, glidefunction:length
Default value	• true • glidefunction:position, glidefunction:substring, glidefunction:concat, glidefunction:coalesce, glidefunction:length

Attribute	Description
Fallback value	• true • glidefunction:position, glidefunction:substring, glidefunction:concat, glidefunction:coalesce, glidefunction:length
Category	Access control
Security risk	• Severity score:5.3 • CVSS score: Medium • If the *glide.db.encoded_query.check_function_field_quer* system property exists and isn't set to a value of **true**, a logged-in user may be able to infer data blindly, leading to sensitive information disclosure.
Functional impact	Users may expect to see values from function fields but get blocked due to ACLs enforced on the function fields.
Dependencies and prerequisites	None

To learn more about adding or creating a system property, see Add a system property.

API and web service

The API and Web Service category ensures that applications have appropriate authentication, authorization and session management, validate all input that traverses a trust boundary and include security controls for all API types.

Specific controls in this category address input validation by service type such as XDS schema validation for SOAP web services or Denial of Service protection for graphQL APIs.

Validate SOAP content type [Updated in Security Center 1.3]

Use the *glide.soap.require_content_type_xml* property to enable validation of a content type as text/xml and protect against invalid SOAP requests.

- When set to **true**, the ServiceNow AI Platform validates the content type as text/xml and protects against invalid SOAP requests.
- If set to **false**, any content-type value is allowed.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.soap.require_content_type_xml*
Configuration type	System Properties (/sys_properties_list.do)

Attribute	Description
Category	API and web service
Purpose	Protect against invalid SOAP requests
Recommended value	true
Default value	true
Security risk rating	8.8
Functional impact	This remediation enables validation of SOAP content type for all the inbound SOAP requests. • If you are using a content type other than text/xml for inbound requests, it may cause potential failure of SOAP transactions. • If you are not using the correct MIME type, it could disrupt third-party integrations.
Security risk	(Moderate) When accepting inbound SOAP requests, the appropriate validation is performed to ensure that the relevant content type is being defined as a part of the request. It restricts the invalid SOAP responses that can be viewed as a security risk.
Reference	Content types

To learn more about adding or creating a system property, see Add a system property.

Require authorization for pdf requests [Updated in Security Center 1.3]

Use the *glide.basicauth.required.pdf* property to designate if incoming PDF requests should require basic authentication.

Ensure the property *glide.basicauth.required.pdf* exists in the sys_properties table and is set to true.

More information

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

Attribute	Description
Property name	*glide.basicauth.required.pdf*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	To enforce basic authentication on PDF requests.
Recommended value	true
Default value	true
Security risk rating	7.5

Attribute	Description
Functional impact	This remediation enforces a combination of authentication methods, in the form of basic authentication and system level access control. • It performs this authentication while retrieving data from tables/pages in the form of PDF data on the instance. • It restricts any guest users who are currently accessing this data. If applicable, you may need to create a new account for users who need access to this content, with necessary access control permissions. To learn more, see Web service import sets.
Security risk	(High) Without appropriate authorization configured on the incoming PDF requests, an unauthorized user can get access to sensitive content/data on the target instance.

To learn more about adding or creating a system property, see Add a system property.

Require Authentication on Event Management HTTP Processor [New in Security Center 1.3, Updated in 1.5, and removed in 2.0]

Learn how to establish secure basic authentication for inbound Amazon Simple Notification Service (SNS) requests when the Event Management plugin (*com.glideapp.itom.snac*) is enabled.

If the *glide.basicauth.required.evtmgmthttpprocessor* property isn't set to the recommended value of **true**, and the Event Management plugin (*com.glideapp.itom.snac*) is active, then basic authentication is not required for all inbound Amazon Simple Notification Service (SNS) requests. This can lead to unauthenticated access to instance data.

To remediate this security risk, ensure that *glide.basicauth.required.evtmgmthttpprocessor* is set to**true** and that *com.glideapp.itom.snac* is active.

More information

Attribute	Description
Configuration name	*glide.basicauth.required.evtmgmthttpprocessor*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	API and web service

Attribute	Description
Security risk	• Severity score: 7 • CVSS score: High • Security risk details: Not setting *glide.basicauth.required.evtmgmthttpprocesso* to the recommended value of true, and not activating the *com.glideapp.itom.snac* plugin causes basic authentication not to be enabled for inbound SNS requests. This could lead to unauthenticated access to instance data.
Dependencies and prerequisites	None
References	• https://docs.aws.amazon.com/sns/latest/dg/welcome.html • Access control
Functional impact	If *glide.basicauth.required.evtmgmthttpprocessor* is not set to the recommended value of True, and if the Event Management plugin (com.glideapp.itom.snac) is active, then basic authentication is not required for all Inbound Amazon Web Services SNS requests. This can lead to unauthenticated access to instance data.

Require authorization for SOAP requests [Updated in Security Center 1.3, 1.5, and 2.0]

Use the *glide.basicauth.required.soap* property to designate if incoming SOAP requests should require basic authorization.

More information

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

Attribute	Description
Property name	*glide.basicauth.required.soapglide.soap.require_ws*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	To enforce soap requests authorization.
Recommended value	true
Security risk rating	8.1
Functional impact	This remediation enforces a combination of authentication methods, in the form of basic authentication and system level access control.

Attribute	Description
	• It performs this authentication while retrieving data from tables/pages in the form of SOAP data on the instance. • It restricts any guest users who are currently accessing this data. • Create an account for a user who needs access to this content, with the necessary access control permissions. To learn more, see SOAP web service and MID Server authentication credentials and SOAP requests.
Security risk	(Medium) Without appropriate authorization configured on the data source SOAP requests, an unauthorized user can access sensitive content/data on the target instance.
References	Authentication

To learn more about adding or creating a system property, see Add a system property.

Require authorization for unload requests [Updated in Security Center 1.3]

Use the *glide.basicauth.required.unl* (useUnloadFormat) property to designate if incoming unload requests should require basic authentication.

More information

Attribute	Description
Property name	*glide.basicauth.required.unl*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	To enforce basic authentication on unload requests.
Recommended value	true
Security risk rating	7.5
Functional impact	This remediation enforces a combination of authentication methods, in the form of basic authentication and system level access control. It performs this authentication while retrieving data from tables/pages in the form of unload data on the instance.
Security risk	(High) Without appropriate authorization configured on the datasource unload requests, an unauthorized user can get access to sensitive content/data on the target instance. Ensure that *glide.basicauth.required.unl* exists in the sys_properties_table and is set to true.
References	Authentication

Require authorization for csv requests [Updated in Security Center 1.3]

Use the *glide.basicauth.required.csv* property to designate if incoming CSV (Comma-Separated Values) requests should require basic authentication.

More information

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

Attribute	Description
Property name	*glide.basicauth.required.csv*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	To enforce basic authentication on CSV requests.
Recommended value	true
Security risk rating	7.5
Functional impact	This remediation enforces a combination of authentication methods, in the form of basic authentication and system level access control. • It performs this authentication while retrieving data from tables/pages in the form of CSV data on the instance. • It restricts any guest users who are currently accessing this data. If applicable, you may need to create a new account for users who need access to this content, with necessary access control permissions. To learn more, see Retrieving data from a CSV formatted file.
Security risk	(High) Without appropriate authorization configured on the incoming CSV requests, an unauthorized user can get access to sensitive content and data on the target instance. Ensure that *glide.basicauth.required.csv* exists in the sys_properties table and is set to true.
References	Web service security

Require authorization for excel requests [Updated in Security Center 1.3]

Use the *glide.basicauth.required.excel* property to designate if incoming Excel requests should require basic authentication.

More information

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

Attribute	Description
Property name	*glide.basicauth.required.excel*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	To enforce basic authentication on Excel requests.
Recommended value	true
Security risk rating	7.5
Functional impact	This remediation enforces a combination of authentication methods, in the form of basic authentication and system level access control. • It performs this authentication while retrieving data from tables/pages in the form of Excel data on the instance. • It restricts any guest users who are currently accessing this data. If applicable, you may need to create a new account for users who need access to this content, with necessary access control permissions.
Security risk	(High) Without appropriate authorization configured on the incoming Excel requests, an unauthorized user can get access to sensitive content/data on the target instance.

To learn more about adding or creating a system property, see Add a system property.

Require authorization for import requests [Updated in Security Center 1.3]

Use the *glide.basicauth.required.importprocessor* property to designate if incoming import requests should require basic authentication.

More information

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

Attribute	Description
Property name	*glide.basicauth.required.importprocessor*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	To enforce basic authentication on import requests.
Recommended value	true
Security risk rating	5.3

Attribute	Description
Functional impact	This remediation enforces a combination of authentication methods, in the form of basic authentication and system level access control. • It performs this authentication while importing data sources into the instance tables/pages. • It restricts any guest users who are currently accessing this data. If applicable, you may need to create a new account for users who need access to this content, with necessary access control permissions. To learn more, see Retrieving data from a CSV formatted file.
Security risk	(Moderate) Without appropriate authorization configured on the datasource import requests, an unauthorized user can get access to sensitive content/data on the target instance.
References	SOAP web services security SOAP web service

Require authorization for JSONv2 request [Updated in Security Center 1.3]

Use the *glide.basicauth.required.jsonv2* property to designate if incoming JSONv2 requests should require basic authorization.

More information

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

Attribute	Description
Property name	*glide.basicauth.required.jsonv2*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	To enforce JSONv2 requests authorization.
Security risk rating	7.5
Recommended value	true
Functional impact	This remediation enforces a combination of authentication methods, in the form of basic authentication and system level access control.

Attribute	Description
	• It performs this authentication while retrieving data from tables/pages in the form of JSON data on the instance. • It restricts any guest users who are currently accessing this data. • Create an account for a user who needs access to this content, with the necessary access control permissions. To learn more, see JSONv2 Web Service JSONv2 Web Service.
Security risk	(High) Without appropriate authorization configured on the data source JSON requests, an unauthorized user can access sensitive content/data on the target instance.
References	Authentication Requiring basic authentication for incoming JSONv2 requests

Require authorization for WSDL request [Updated in Security Center 1.3 and 1.5]

Use the *glide.basicauth.required.wsdl* property to designate if incoming WSDL (Web Services Description Language) requests should require basic authentication.

If *glide.basicauth.required.wsdl* is not set to the recommended value of true, then , then this will disable Basic Authentication for WSDL requests. WSDL is a protocol that is used to describe web services such as instance table schemas, and is not a mechanism for sharing the data within tables. Setting this property to true allows for disclosure of table schemas to unauthenticated users.

Note: If you choose not to require basic authentication for incoming WSDL requests, you must modify Access Control (ACL) rules to enable guest users to access the WSDL content.

More information

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

Attribute	Description
Property name	*glide.basicauth.required.wsdl*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	To enforce basic authentication on WSDL requests.
Recommended value	true
Security risk rating	5.3

Attribute	Description
Functional impact	This remediation enforces a combination of authentication methods, in the form of basic authentication and system level access control. • It performs this authentication while retrieving data from tables/pages in the form of WSDL data on the instance. • It restricts any guest users who are currently accessing this data. If applicable, you may need to create a new account for users who need access to this content, with necessary access control permissions.
Security risk	(Medium) Without appropriate authorization configured on the WSDL web services, an unauthorized user can get access to sensitive WSDL content/data on the target instance.
References	Web service security

Require authorization for XML requests [Updated in Security Center 1.3]

Use the *glide.basicauth.required.xml* property to designate if incoming XML requests should require basic authentication.

More information

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

Attribute	Description
Property name	*glide.basicauth.required.xml*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	To enforce basic authentication on XML requests.
Security risk rating	7.5
Recommended value	true
Functional impact	This remediation enforces a combination of authentication methods, in the form of basic authentication and system level access control. • It performs this authentication while retrieving data from tables/pages in the form of XML data on the instance. • It restricts any guest users who are currently accessing this data. If applicable, you may need to create a new account for users who need access to this content, with necessary access control permissions.

Attribute	Description
	To learn more, see XML parser step.
Security risk	(High) Without appropriate authorization configured on the incoming XML requests, an unauthorized user can get access to sensitive content/data on the target instance.
References	Authentication

Require authorization for XML output requests [Updated in Security Center 1.3]

Configure this property so that basic authorization is required for all inbound XMLOutputProcessor requests.

If the *glide.basicauth.required.xmloutputprocessor* property is not set to the recommended value of **true**, then basic authorization is not required for inbound XMLOutputProcessor requests which could lead to unauthenticated information disclosure from the instance.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Configuration name	*glide.basicauth.required.xmloutputprocessor*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	API and web service
Security risk	• Severity score: 7.5 • CVSS score: High • Security risk details: Not setting the property to the recommended value of **true** could lead to sensitive information being leaked from the instance.
Dependencies and prerequisites	None

Require Authorization for XSD Requests [Updated in Security Center 1.3]

Use the *glide.basicauth.required.xsd* property to designate if incoming XSD (XML Schema Definition) requests should require basic authentication.

Ensure the property *glide.basicauth.required.xsd* exists in the sys_properties table and is set to true.

More information

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

Attribute	Description
Property name	glide.basicauth.required.xsd*glide.basicauth.required.xsd*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	To enforce basic authentication on XSD requests.
Recommended value	true
Security risk rating	5.3
Functional impact	This remediation enforces a combination of authentication methods, in the form of basic authentication and system level access control. • It performs this authentication while retrieving data from tables/pages in the form of XSD data on the instance. • It restricts any guest users who are currently accessing this data. If applicable, you may need to create a new account for users who need access to this content, with necessary access control permissions. To learn more, see Non-interactive sessions.
Security risk	(Moderate) Without appropriate authorization configured on the incoming XSD requests, an unauthorized user can get access to sensitive content/data on the target instance.
References	Authentication

Require authorization for script requests [Updated in Security Center 1.3]

Use the *glide.basicauth.required.scriptedprocessor* property to designate if incoming script requests should require basic authentication.

More information

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

Attribute	Description
Property name	*glide.basicauth.required.scriptedprocessor*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	To enforce basic authentication on scripts requests.

Attribute	Description
Recommended value	true
Security risk rating	7.2
Functional impact	This remediation enforces the authentication in the form of Basic authorization. • It performs this authentication while processing script requests on the instance. • It restricts any guest users who are currently accessing this data. If applicable, you may need to create a new account for users who need access to this content, with necessary access control permissions.
Security risk	(High) Without appropriate authorization configured on the incoming script requests, an unauthorized user access sensitive content/data on the target instance.
References	Authentication

To learn more about adding or creating a system property, see Add a system property.

Require authorization for SCHEMA requests [Updated in Security Center 1.3]

Use the *glide.basicauth.required.schema* property to require basic authorization for all Inbound Table Schema Processor requests.

The Inbound Table Schebma Processor handles incoming schema requests for the platform.

Set *glide.basicauth.required.schema* to the recommended value of **true** to require basic authorization for all Inbound Table Schema Processor requests. Set the value to **false** to not require basic authorization for all Inbound Table Schema Processor requests.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.basicauth.required.schema*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	To require basic authorization for all Inbound Table Schema Processor requests.
Recommended value	True (default value).
Configuration type	Boolean
Security risk	(Moderate) Omitting authentication from this processor will lead to unauthenticated access to instance data.

Require authorization for RSS requests [Updated in Security Center 1.3]

Use the *glide.basicauth.required.rss* property to designate if incoming RSS requests should require basic authentication.

More information

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

Attribute	Description
Property name	*glide.basicauth.required.rss*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	To enforce basic authentication on RSS requests.
Recommended value	true
Security risk rating	7.5
Functional impact	This remediation enforces a combination of authentication methods, in the form of basic authentication and system level access control. • It performs this authentication while processing RSS requests on the instance. • It restricts any guest users who are currently accessing this data. If applicable, you may need to create a new account for users who need access to this content, with necessary access control permissions. To learn more, see RSS feed generator.
Security risk	(High) Without appropriate authorization configured on the incoming RSS requests, an unauthorized user can get access to sensitive content/data on the target instance.
References	RSS basic authentication

Require authorization for API requests [Updated in Security Center 1.3]

Use the *glide.basicauth.required.api* property to enhance security for basic authorization for incoming REST requests.

Set the *glide.basicauth.required.api* property to **true** to require authorization for all REST requests. Set the property to **false** to bypass authorization for all REST requests.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.basicauth.required.api*
Configuration type	System Properties (/sys_properties_list.do)
Category	API and web service
Purpose	Basic authorization for incoming REST requests.
Recommended value	True (default)
Configuration type	String
Security risk	(High) If "glide.basicauth.required.api" is not set to the recommended value of "true", then this will disable Basic Authentication on API request and will lead to unauthenticated access to instance data.
Security risk rating	8.6

Architecture, design, and threat modeling

This broad control addresses high level design considerations and key elements to implement a secure application. This covers the tenants of availability, confidentiality processing integrity, non-repudiation and privacy. Additionally, elements of a secure software development lifecycle are included.

Certificate based authentication not enforced [New in Security Center 1.3]

The *glide.authenticate.mutual.enabled* property enables certificate based authentication, a type of mutual authentication for inbound REST connections to REST and SOAP APIs in the ServiceNow AI Platform.

Mutual authentication establishes trust between server and client by exchanging secure socket layer (SSL) certificates to validate the certificate with a trusted Certificate Authority. This allows verification that a trusted source is connecting to the ServiceNow AI Platform. If this instance is not set to the recommended value of true, an instance could be vulnerable to man-in-the-middle attacks (MitM).

To remediate this security threat, enable mutual authentication for inbound web services. If it's your first time installing the certificate-based authentication plugin (*com.glide.auth.mutual)* for the ServiceNow AI Platform, then follow the Set up Certificate-based authentication instructions. In addition, ensure that the *glide.authenticate.mutual.enabled* property is set to true to activate the plugin.

More information

Attribute	Description
Configuration name	*glide.authenticate.mutual.enabled*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true

Attribute	Description
Default value	true
Category	Architecture, design, and threat modeling
Security risk	• Severity score: 5.3 • CVSS score: Medium • Security risk details: If this property is not set to the recommended value of true, then certificate based authentication does not validate certificates with a trusted Certificate Authority. This increases the chances of a bad actor attacking an instance using MitM attacks.
Dependencies and prerequisites	None
References	• https://csrc.nist.gov/glossary/term/man_in_the_middle_attack • Certificate-based authentication • Configure mutual authentication

Check impersonation on ACL evaluation in HR App [New in Security Center 1.3 and updated in 1.5]

Use the *sn_hr_core.impersonateCheck* property to prevent a user from impersonating another user and accessing their HR information.

A secure setting prevents an admin from seeing another user's HR information while using impersonation. An insecure setting for this property allows an admin to impersonate a user and access HR data such as survey results or audit records with the impersonated user's access. Due to the nature of this type of data, such as information which should be available only to the user themselves like email, this is not recommended. Setting *sn_hr_core.impersonateCheck* to true only allows access to HR information when the user is not impersonating any others.

More information

Attribute	Description
Configuration name	*sn_hr_core.impersonateCheck*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Architecture, design, and threat modeling
Security risk	• Severity score: 2.7 • CVSS score: Low

Attribute	Description
	• Security risk details: An insecure setting for this property allows an admin to impersonate a user and access HR data such as survey results or audit records with the impersonated user's access.
Dependencies and prerequisites	None
Functional impact	When this property set to true, it prevents an admin from seeing another user's HR information while using impersonation. When set to false, it allows an admin to impersonate a user and access HR data such as survey results or audit records with the impersonated user's access. Due to the nature of this type of data, such as information which should available only to the user themselves like an email, this is not recommended. Setting sn_hr_core.impersonateCheck to true only allows access to HR information when the user is not impersonating any others.

Disable local login for users with Single Sign-On (SSO) enabled

Update user records to disable local login for users with Single Sign-On (SSO) enabled.

Users configured to use SSO authentication may be able to access the instance, or parts of the instance, with local credentials stored in the **user_password** field of their User [sys_user] record. This access applies to both interactive and non-interactive access for users who aren't locked out. Help prevent SSO-configured users from using local credentials to reduce the chance that valid local login credentials are stolen and used by malicious users.

Review Now Support Knowledge Base article KB1649420 for instructions on identifying and addressing accounts with local login still enabled on an instance with SSO enabled.

More information

Attribute	Description
Security risk	When SSO authentication is enabled for a user, it's best practice to prevent that user from logging in locally. This reduces the chance that the valid local login credentials are stolen and used to login by a malicious user.
Common Vulnerability Scoring System (CVSS) score	4.2
Common Vulnerability Scoring System (CVSS) rating	Medium
Functional impact	SSO configured users are able to log in with local credentials.
Dependencies and prerequisites	Single Sign-On must be enabled (the **glide.authenticate.multisso.enabled** system property set to **true**.)
Data type	N/A
Base system value	N/A

Attribute	Description
Fallback value	N/A
Recommended value	N/A

Disable unauthenticated published reports [Updated in Security Center 2.0]

Deactivate this property to prevent the user from publishing or accessing reports. This property disables the published reports feature in reporting.

Enable publishing reports by setting the *glide.report.published_reports.enabled* to **true**.

Ensure the Glide Property *glide.report.published_reports.enabled* exists and is set to the value false. If the property does not appear in the sys_properties table, add a new record.

More information

Attribute	Description
Property name	*glide.report.published_reports.enabled*
Configuration type	System Properties (/sys_properties_list.do)
Category	Architecture, design, and threat modeling
Purpose	Disables the published reports feature in reporting.
Type	true | false
Recommended value	**false**
Security risk rating	6.5
Functional impact	The user cannot publish reports.
Security risk	(Moderate) If this property is not enabled, users may be able to access or publish reports exposing sensitive data. Publishing a report creates a URL that anyone can use to access the report, including people who are not users. When anyone navigates to the URL, the report is generated with current data from the instance.
References	Publish a report

To learn more about adding or creating a system property, see Add a system property.

Enforce field ACLs for inbound query requests

Manage how incoming queries are validated on your instance.

Use the *glide.export.query.enforce_field_acl* property to control whether field-level ACLs are enforced on the fields referenced in an inbound query requests. When set to **true**, field ACLs are checked against fields used in the incoming query, and the query is rejected if the user is unauthorized to access those fields. When set to **false**, field ACLs are not checked on query conditions, and the query executes regardless of field-level access restrictions.

This property applies only to field ACL enforcement on query conditions. Setting this property to **false** does not affect whether users can read field values they are not otherwise authorized to view. Field-level read ACLs remain enforced regardless of this setting.

Set the property *glide.export.query.enforce_field_acl* to **true**.

More information

Attribute	Description
Configuration name	*glide.export.query.enforce_field_acl*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Architecture, design, and threat modeling
Security risk	• Severity score: 4.4 • CVSS score: Medium • Security risk details: If this property is set to **false**, ACLs are not checked against incoming queries which can lead to information disclosure.
Dependencies and prerequisites	None

Enforce read ACLs on report views

Manage how Read ACLs are enforced on your instance.

Use the *glide.report.report_view.read_acl* property to enforce the Read ACL (table level) on reporting functions when there is no Report View ACL on the table or field. If this property is not set to **enforce**, ACLs could be bypassed leading to sensitive information leakage.

More information

Attribute	Description
Configuration name	*glide.report.report_view.read_acl*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	enforce
Default value	enforce
Category	Architecture, design, and threat modeling
Security risk	• Severity score: 7.1 • CVSS score: High • Security risk details: Not setting this property to **enforce** could cause ACLs to be bypassed.
Dependencies and prerequisites	None

Enforce Query ACLs for Knowledge Quick Links

Enforce query ACLs for Knowledge Quick Links using a system property.

When the *com.glide.security.query_acl.enabled.knowledge_quick_links* system property is set to **true**, query ACLs are enforced for Knowledge Quick Links. If this property is set to **false**, an attacker can use blind queries to enumerate and exfiltrate data due to the default behavior of `GlideRecord.addEncodedQuery`.

If the property doesn't exist in the System properties [sys_properties] table, the secure default of **true** is used. A third option, **external_and_guests** enforces query ACLs only for external users and guests.

Verify that the *com.glide.security.query_acl.enabled.knowledge_quick_links* system property is set to **true**.

More information

Attribute	Description
Configuration name	*com.glide.security.query_acl.enabled.knowledge_qui*
Configuration type	System Properties (/sys_properties_list.do)
Data type	String - Used as a ternary operator
Recommended value	true
Default value	true
Fallback value	true
Category	Architecture, design, and threat modeling
Security risk	• Severity score: 5.3 • CVSS score: Medium • Security Risk: ACLs can be bypassed, disclosing field data to users who don't have permissions to see it. This disclosure could include sensitive data depending on the table exploited.
Functional impact	Users aren't able to perform specific queries on fields where they don't have field level access.
Dependencies and prerequisites	None

To learn more about adding or creating a system property, see Add a system property.

Enforce Query ACLs for SubLists, List Counts and Widget Data Tables

Enforce query ACLs on sublist, list count, and widget data table queries using system properties.

Set *com.glide.security.query_acl.enabled.sub_lists* to **true** to enforce query ACLs on sublist queries, such as grouped lists and related lists.

Set *com.glide.security.query_acl.enabled.list_count* to **true** to enforce query ACLs on list count queries.

Set *glide.security.query_acl.enabled.data_table* to **true** to enforce query ACLs on widget data tables.

If any of these system properties are set to **false**, an attacker can use blind queries to enumerate and exfiltrate data due to the default behavior of `GlideRecord.addEncodedQuery`. If these properties don't exist in the System Properties [sys_properties] table, the secure default of true is used. A third option, *external_and_guests*, enforces ACLs only for external users and guests.

Ensure these system properties do not appear in the System Properties [sys_properties] table or are set to **true**.

More information

Attribute	Description
Configuration name	• *com.glide.security.query_acl.enabled.sub_lists* • *com.glide.security.query_acl.enabled.list_count* • *glide.security.query_acl.enabled.data_table*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	• true • true • true
Default value	• true • true • true
Fallback value	• true • true • true
Category	Architecture, design, and threat modeling
Security risk	• Severity score: 5.3 • CVSS score: Medium • Security Risk: ACLs can be bypassed disclosing field data to users who do not have permissions to see it. This could include sensitive data depending on the table exploited.
Dependencies and prerequisites	None

Enforce valid query string choice [New in Security Center 7.0]

Use a system property to ensure that any choice field value, when passed via a URL query string, is a valid active choice when a record is created.

When the *glide.ui.query_string.enforce_valid_choice_on_create* system property is set to **true**, the platform validates that any choice field value passed via a URL query string (for example, from a list filter) is a valid active choice when the record is created.

If invalid, the value is ignored and the field falls back to its default value. When the property is **false**, validation is inactive, and the system accepts any value, even invalid or inactive ones. This acceptance can potentially result in incorrect or unexpected data being stored on records.

Ensure *glide.ui.query_string.enforce_valid_choice_on_create* exists in the System Properties [sys_properties] table and is set to **true**. If the property doesn't exist in the table, the fallback value is **false**.

More information

Attribute	Description
Configuration name	*glide.ui.query_string.enforce_valid_choice_on_crea*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	false
Fallback value	true
Category	Architecture, design, and threat modeling
Security risk	• Severity score: 2.2 • CVSS score: Low • Security Risk: Set this property to **true** to ensure that only valid choices are used for new record creation. Invalid choices may lead to minor and unexpected data integrity issues, where a user is able to select an unintended value.
Functional impact	Existing instances and workflows may rely on new records being created based on a filter that contains an invalid or inactive choice. Use the following process to observe this behavior: **1.** Log in to an instance as an admin user. **2.** Create a String field on any table. For example, an incident with 2 choices: **Test1** and **Test2**. **3.** Create a list filter on incident table selecting the string field with value set to **Test2**. **4.** Go to dictionary field and deactivate the choice **Test2**.

Attribute	Description
	5. Go to the filter selected in step 3 and select **New** button. 6. Check the value of the String type field in the newly opened record. When the property is set to **true**, the string field shouldn't have a value or should show the default value. When the property is set to **false** (default), the string field has the value set to **Test2**.
Dependencies and prerequisites	None

To learn more about adding or creating a system property, see Add a system property.

Define allowed ServiceNow internal IP addresses [Updated in Security Center 1.3 and 1.5]

Use the *glide.ip.authenticate.strict* property to specify IP ranges that can make inbound connections on an instance.

Use the *glide.ip.authenticate.strict* property to reduce the scope of the IP allow list of an instance and restrict the IP addresses that ServiceNow personnel can use to make inbound/outbound connections to an instance. The exact IP ranges removed from the IP allow list by this property may be adjusted over time as the ServiceNow internal network changes. When set to true, *glide.ip.authenticate.strict* always guarantees an IP allow list that is equal to or more restrictive than the default.

When glide.ip.authenticate.strict is set to true:

A strict list of ServiceNow IP ranges replaces the default IP allow lists for inbound and outbound requests. This IP allow list, which starts with a more restrictive predefined set of allowed IP ranges, is replaced by the property glide.ip.authenticate.allow.secured.self_hosted_list if the instance is self-hosted.

When glide.ip.authenticate.strict is set to false:

The default IP allow list is used which contains a wider set of ServiceNow IP ranges. The default IP allow list is replaced by the contents of glide.ip.authenticate.allow.self_hosted_list if the instance is self-hosted.

Note: The *glide.ip.authenticate.allow.secured* property contains the ServiceNow data center and secure VPN IP ranges used when *glide.ip.authenticate.strict* is set to true. This property is maintained and deployed automatically by ServiceNow. Customers do not need to configure it. The only customer action required by this hardening setting is to set *glide.ip.authenticate.strict* to `true`.

Note: Regardless of the value of *glide.ip.authenticate.strict* or if the instance is self-hosted, the allow list includes IP addresses in the *glide.custom.ip.authenticate.allow* and *glide.custom.ip.outbound.authenticate.allow* system properties, if defined.

All IP list properties share the same format, which is a comma separated range of IP addresses in IPv4 or IPv6 format. IP ranges are specified using a hyphen (10.0.10.14-10.0.10.19), using CIDR notation (10.0.10.0/24), or consist of a single IP address (10.0.10.5). At runtime, you can make additions to the IP allow list by adding entries to the IP Address Access Controls [ip_access]

table. To restrict all inbound access to a defined set of IP addresses, add a deny-all entry to the [ip_access] table and then add the specific IP addresses or ranges that you want to allow.

Warning: The value for this property is a no DB override. It can't be altered or overridden.

More information

Attribute	Description
Configuration name	• *glide.ip.authenticate.strict* • *glide.ip.authenticate.allow.secured*
Configuration type	System Properties (/sys_properties_list.do)
Data type	• Boolean • String
Recommended value	• true • Managed by ServiceNow. Contains ServiceNow data center and secure VPN IP ranges. This value is deployed automatically and cannot be modified by customers.
Default value	• false • <empty>
Fallback value	• false • <empty>
Category	Architecture, design, and threat modeling
Security risk	• Severity score: 4.3 • CVSS score: Medium • Security Risk Details: Broader access increases the risk of unauthorized or unnecessary access to the instance by non-essential internal users, such as support or sales staff, and reduces control over privileged access. Enforcing strict IP authentication limits connectivity to essential infrastructure, strengthening security and reducing exposure to insider threats or misconfiguration.
Functional impact	This should have no functional impact. It may restrict non-essential ServiceNow personal from accessing an instance. However, these are individuals that do not generally require access to such instances. If access is

Attribute	Description
	required, you can grant it on a case by case basis using the IP Address Access Controls [ip_access] table.
Dependencies and prerequisites	None

To learn more about adding or creating a system property, see Add a system property.

Disable legacy JQuery behavior [Updated in Securty Center 1.3]

The *glide.jquery.legacy* is used to prevent older prepatched JQuery versions from being used which will introduce unpatched vulnerabilities in the library.

Set *glide.jquery.legacy* to the recommended value of **false** to prevent older prepatched JQuery versions from being used which introduce unpatched vulnerabilities in the library. Set the value to **true** to allow prepatched JQuery versions.

More information

Attribute	Description
Property name	*glide.jquery.legacy*
Configuration type	System Properties (/sys_properties_list.do)
Category	Architecture, design, and threat modeling
Purpose	To prevent potential security risks arising from attacks on vulnerabilities discovered in outdated JQuery library versions.
Recommended value	False
Configuration type	Boolean
Security risk	(High) Prevent older prepatched JQuery versions from being used which introduce unpatched vulnerabilities in the library. The system property is a failsafe in case any organizations depend on the non-patched versions of angularJS to run their custom implementations.
Security risk rating	7.1

To learn more about adding or creating a system property, see Add a system property.

Disable GlideRecord Scope Fencing Legacy Behavior [New in Security Center 1.3 and updated in 1.5 and 2.0]

The *glide.record.legacy_cross_scope_access_policy_in_script* property disables scope fencing allowing scoped apps to access global script interfaces. It was created as a patch to GlideRecord's cross scope access.

GlideRecord provided cross scope create/update access to tables that were not configured with that level of access. In order to prevent customers from having applications broken when this scoped access behavior was patched, the property *glide.record.legacy_cross_scope_access_policy_in_script* was created. When true, cross scope access falls back onto legacy behavior (insecure). This property disables scope fencing, allowing scoped apps to access global script interfaces.

It is best security practice to have scope fencing restrictions in place. Scoping ensures applications can only access resources with explicit access or within their scope, following the principle of least privilege. Disabiling this feature could lead to confidentiality, availability, and integrity impacts.

Set the Glide Property *glide.record.legacy_cross_scope_access_policy_in_script* to false. When not present in the sys_properties table, the default value is true.

More information

Attribute	Description
Configuration name	*glide.record.legacy_cross_scope_access_policy_in_s*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	true (when the property does not exist in the sys_properties table.)
Category	Architecture, design, and threat modeling
Security risk	• Severity score: 5 • CVSS score: Medium • Security risk details: Scoping ensures applications can only access resources with explicit access or within their scope following the principle of least privilege. Disabling this feature could lead to confidentiality, availability, and integrity impacts.
Dependencies and prerequisites	None

Disable legacy AngularJS behavior [Removed in Security Center 2.2]

Use the *glide.angular.legacy* property to protect from potential security risks arising from attacks on vulnerabilities discovered in outdated AngularJS library versions.

Set *glide.angular.legacy* to the recommended value of **false** to prevent older prepatched angularJS versions from being used. Set the property to **true** to use older prepatched angularJS versions.

More information

Attribute	Description
Property name	*glide.angular.legacy*
Configuration type	System Properties (/sys_properties_list.do)
Category	Architecture, design, and threat modeling
Purpose	The system property is a failsafe in case any organizations depend on the non-patched versions of angularJS to run their custom implementations.

Attribute	Description
Recommended value	False
Configuration type	Boolean
Security risk	(High) Using older versions of angularJS could potentially lead to security risks arising from attacks on vulnerabilities discovered in outdated AngularJS library versions.
Security risk rating	7.1

Require authorization for data broker rest API [Updated in Security Center 1.3]

Use the *glide.basicauth.required.databrokerrestapiprocessor* property to require basic authorization for all inbound Data Broker Rest API requests.

If this property is set to **true**, then authorization is applied. If it is set to **false**, then no authorization is used which could lead to sensitive information being leaked from your instance.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Configuration name	*glide.basicauth.required.databrokerrestapiprocesso*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Architecture, design, and threat modeling
Security risk	• Severity score: 8.6 • CVSS score: High • Security risk details: If property is set to **false**, then API authentication will not be applied enabling a bad-actor to access sensitive data.
Dependencies and prerequisites	None

Restricted Binding functionality in case Bearer Authorization [New in Security Center 7.0]

Use a system property and restricted binding to ensure that an access token generated using that entity can't be used for UI calls.

Use the *glide.oauth.enforce_restricted_binding_for_ui* system property and enable restricted binding for an OAuth entity to prevent the access tokens generated by that entity from being used for UI calls (For example, `incident_list.do`).

When restricted binding is turned off, the access token generated can be used for UI calls regardless of the value of the system property.

Ensure that *glide.oauth.enforce_restricted_binding_for_ui* is set to **true** and **Enforce Token Restrictions** is set to **true** in all OAuth entity entries. For details on OAuth entity entries, see OAuth inbound.

More information

Attribute	Description
Configuration name	*glide.oauth.enforce_restricted_binding_for_ui*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Fallback value	false
Category	Architecture, design, and threat modeling
Security risk	• Severity score: 5.0 • CVSS score: Medium • Security risk details: When the *glide.oauth.enforce_restricted_binding_for_ui* system property isn't set to **true** or restricted binding is turned off, then a user with the access token to access an API (for example, mobile API) can get a session issued and use it to access other restricted resources on the instance (for example, `incident_list.do`).
Dependencies and prerequisites	None

Deny by default with empty ACLs [Updated in Security Center 1.3]

Use the *glide.sm.default_mode* property to control the default behavior of security manager when it finds that existing Access Control List (ACL) rules are a part of wildcard table ACL rules.

Prevent your instance's legacy security manager from allowing access to resources when there are no ACLs defined for that resource, or if there are only wildcard table-level ACLs (for example, `incident.*`). When allowed access by default, anything that does not have explicit ACLs set is susceptible to manipulation.

Set the **glide.sm.default_mode** system property value to **deny** to disallow access when there are no define ACL rules, or there are only wildcard table-level ACLs.

⚠ **Warning:** This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.sm.default_mode*
Configuration type	System Properties (/sys_properties_list.do)
Category	Architecture, design, and threat modeling
Purpose	Best security practice would be to restrict an Access to the tables by an unauthorized user. • If there are no ACL rules in place for tables, this property ensures that at least wildcard ACLs are validated for any CRUD operation performed on the table/field. • These rules restrict the read, write, create, and delete operations on all tables, unless the user has the admin role or meets the requirements of another table ACL rule.
Recommended value	deny
Functional impact	If you set this property to **Allow**, the wildcard table ACL rules allow CRUD operations on all tables unless there are specific table ACL rules in place to restrict such operations. **Note:** This plugin is not intended for existing instances, as it might modify security access to tables that are already in use in a production environment.
Security risk	6.3
References	Default deny property

To learn more about adding or creating a system property, see Add a system property.

Set Automatic Token Cleanup for Token Credentials [New in Security Center 2.0]

Use the *com.snc.platform.security.token.auth.cleanup* property to ensure that expired API keys and HMAC secrets are deleted, thereby limiting the potential for token reuse.

If the *com.snc.platform.security.token.auth.cleanup* property is set to the insecure value of false, expired API keys and HMAC secrets will not be deleted, creating a potential for token reuse. If a token was expired due to leakage or compromise, its reuse could expose the instance to anyone possessing the leaked token.

Expired tokens are retained for the number of days defined by *com.snc.platform.security.token.auth.days.expired.hmac_secret.is.kept* and *com.snc.platform.security.token.auth.days.expired.api_key.is.kept*. Valid values for these settings are integers of 0 or greater. A value of 0 results in the expired tokens being deleted on the same day, while a higher number of days increases the exposure period. A default value of 7 days or fewer is recommended.

More information

Attribute	Description
Configuration name	*com.snc.platform.security.token.auth.cleanup*, *com.snc.platform.security.token.auth.days.expired*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	The recommended values are true, and any integer less than or equal to 7.
Default value	7
Category	Architecture, design, and threat modeling
Security risk	• Severity score: 5.1 • CVSS score: Medium • Security risk details: Not configuring this property to the recommend value of true could prevent expired API keys and HMAC secrets from being deleted which increases the likelihood for token reuse.
Dependencies and prerequisites	None

Authentication

The authentication category covers the main elements of modern authentication to confirm an entity and its claims are authentic and correct, resistant to impersonation and prevent interception of passwords.

The ASVS standard builds on the NIST 800-63b (https://pages.nist.gov/800-63-3/sp800-63b.html) specification for this section.

Authentication includes password policy, controls and storage, proper implementation of authenticators and proper implementation of out of band or one time verifiers.

Activate role-based multi-factor authentication [Updated in Security Center 1.3]

Use the *glide.authenticate.multifactor* property to enforce role-based multi-factor authentication (MFA) for all users assigned to specific roles.

Enforce multi-factor authentication based on the roles assigned to the user. If a user has been assigned "admin", "security_admin" or "user_admin" roles in the multi-factor roles list, MFA will be enforced.

- Set this property to **true** to enforce role-based multi-factor authentication for all users assigned to specific roles.
- Set this property to **false** to disable role-based multi-factor authentication for all users assigned to specific roles.

More information

Attribute	Description
Property name	*glide.authenticate.multifactor*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	Enforce role-based multi-factor authentication for all users assigned to specific roles.
Type	true/false
Recommended value	*true*
Security Dependencies	Activate **Role based multi-factor authentication** within the **Multi-factor Criteria table**.
Security risk rating	7.2
Functional impact	Enabling this property improves the experience of the user. It acts as an extra layer of protection and security against compromised credentials.
Security risk	(Moderate) If this property is not enabled, there is a risk of unauthorized access to sensitive data.
References	Configure role-based multi-factor criteria

To learn more about adding or creating a system property, see Add a system property.

Activate role based multi-factor authentication [Updated in Security Center 1.3]

Use the *glide.authenticate.multifactor* property to enforce role-based multi-factor authentication (MFA) for all users assigned to specific roles.

Enforce multi-factor authentication based on the roles assigned to the user. If a user has been assigned "admin", "security_admin" or "user_admin" roles in the multi-factor roles list, MFA will be enforced.

- Set this property to **true** to enforce role-based multi-factor authentication for all users assigned to specific roles.
- Set this property to **false** to disable role-based multi-factor authentication for all users assigned to specific roles.

More information

Attribute	Description
Property name	*glide.authenticate.multifactor*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	Enforce role-based multi-factor authentication for all users assigned to specific roles.

Attribute	Description
Type	Boolean
Recommended value	*true*
Default value	false
Security Dependencies	Activate **Role based multi-factor authentication** within the **Multi-factor Criteria table**.
Security risk rating	7.2
Functional impact	Enabling this property improves the experience of the user. It acts as an extra layer of protection and security against compromised credentials.
Security risk	(Moderate) If this property is not enabled, there is a risk of unauthorized access to sensitive data.
References	Configure role-based multi-factor criteria

Anti-CSRF token (instance security hardening)

Use the *glide.security.use_csrf_token* property to ensure the use of a secure token to identify and validates incoming requests, which in turn are used to prevent these attacks.

Cross-Site Request Forgery (CSRF) is an attack that forces an end user to execute unwanted actions on a web application in which they're currently authenticated. CSRF attacks specifically target state-changing requests, not theft of data, since the attacker has no way to see the response to the forged request.

Note: By default, the *glide.security.use_csrf_token* property is set to `true` for zBoot instance.

The following properties can be enabled for added controls over CSRF token:

- *glide.security.csrf_previous.time_limit*
- *glide.security.csrf_previous.allow*
- *glide.security.csrf.strict.validation.mode*

More information

Attribute	Description
Property name	glide.security.use_csrf_token
Configuration type	System Properties (/sys_properties_list.do)
Configure in Instance Security Center	Yes
Purpose	To protect the application from potential CSRF attack.
Recommended value	true
Functional impact	(Low) This remediation enables an extra validation step before the instance user submits a write request to the instance. Every write request contains a CSRF token (i.e a validation/CSRF ID tied to the user session). When the user session expires, the secure token expires with it.

Attribute	Description
Security risk	(High) Cross Site Request Forgery is a significant security risk that violates the integrity of the instance data. An attacker can launch the CSRF attack by abusing the trust of an instance user. With the help of social engineering attacks, a user can submit a malformed request on behalf of the attacker on the instance.

To learn more about adding or creating a system property, see Add a system property.

Control Lockout Time for Invalid Password Reset Attempts [Updated in Security Center 1.3 and 2.0]

The *password_reset.request.max_attempt_window* property controls the number of minutes a user must wait to reset or change their password after exceeding the maximum number of unsuccessful attempts that is set with the *password_reset.request.max_attempt* property.

The *password_reset.request.max_attempt_window* property defines the number of minutes a user must wait to reset or change their password after exceeding the maximum number of unsuccessful attempts that is set with the *password_reset.request.max_attempt property*. A small number of minutes for the *password_reset.request.max_attempt_window* property increases the risk of successfully brute forcing a password as a greater number of password reset attempts can be made. The default of 1440 minutes is recommended.

Ensure the property *password_reset.request.max_attempt_window* is set to 1440 or greater.

More information

Attribute	Description
Property name	*password_reset.request.max_attempt_window*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	Denotes the lockout period in minutes after the maximum number of unsuccessful password reset attempts has been met.
Recommended value	1440
Default value	1440
Configuration type	Positive integer values
Security risk	(High) If the property is not set to the recommended value of 1440 or less, then it could be possible to perform account brute force as the account will not be locked after a maximum number of wrong authentication attempts.
Security risk rating	7.5
References	Configure Password Reset properties

To learn more about adding or creating a system property, see Add a system property.

Disable creating users from incoming emails [Updated in Securty Center 1.3]

Use the *glide.user.trusted_domain* property to specify the comma-separated list of trusted domains used in the creation of users from incoming emails.

An administrator can set an email property to automatically create users from incoming emails. If set this property to the insecure value, the instance will automatically create users from incoming email. Each user created will have the same hard coded default password which makes bypassing authentication through brute force easier.

More information

Attribute	Description
Property name	*glide.pop3readerjob.create_caller*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Recommended value	false
Default value	false
Security risk rating	5.4
Functional impact	Once this property is configured, the instance only accepts emails from trusted domains. If you do not include the domain in the trusted list, there is an impact to guest users because accounts are created automatically.
Security risk	(Moderate) If the property is not enabled, an attacker might use an email spoofing/spamming campaign to send multiple emails resulting in the creation of more unnecessary guest users.
References	Inbound mail configuration

To learn more about adding or creating a system property, see Add a system property.

Disable password-less authentication

Use the *glide.login.no_blank_password* property to prevent users from logging into the NOW platform with blank passwords, or by leaving the **Password** field empty.

Even if the admin purposefully assigns an empty value or blank password in user records, a user can't log in without providing a value in the **Password** field.

More information

Attribute	Description
Property name	*glide.login.no_blank_password*
Configuration type	System Properties (/sys_properties_list.do)

Attribute	Description
Configure in Instance Security Center	Yes
Purpose	To ensure strong authentication as sometimes the username are easy to guess within an organization.
Recommended value	true
Functional impact	Operations should not use blank passwords because it is viewed as a critical security risk. However, if there is a valid case for such usage, there is a possibility of an outage. Users with blank passwords wouldn't be able to log in to the instance.
Security risk	(High) An attacker is able to log in to the instance with the default usernames, or by specific individual/group (usually firstname.lastname) without any password. Doing so is viewed as a critical security risk, because it would enable a public user to violate the confidentiality and integrity of the instance data.

Disable resource owner password credentials (ROPC) in OAuth 2 token grants [New in Security Center 7.0]

Prevent Resource Owner Password Credentials (ROPC) from granting OAuth 2 tokens.

By default, Resource Owner Password Credentials (ROPC) are allowed to grant OAuth 2 tokens on your instances when a client application directly requests an access token using a user name and password. When the *glide.oauth.inbound.ropc.grant_type.disabled* is set to **true**, ROPC is inactive and can't be used to grant OAuth 2 tokens.

Ensure that the *glide.oauth.inbound.ropc.grant_type.disabled* system property is set to **true**. If the property doesn't exist on the System Properties [sys_properties] table, the default value is **false**. If this property exists on that table, it defaults to **false**.

More information

Attribute	Description
Configuration name	*glide.oauth.inbound.ropc.grant_type.disabled*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Fallback value	false
Category	Authentication
Security risk	• Severity score: 3.3 • CVSS score: Low • When the property is set to **false**, using ROPC to grant OAuth 2 tokens is allowed. ROPC is considered less secure that other authentication flows because the

Attribute	Description
	user's credentials are exposed to the application. This can lead to vulnerabilities in situations where the client is compromised and suffers from weaknesses similar to those of basic auth. OAuth 2.1 has deprecated ROPC.
Functional impact	When the property is set to **true**, ROPC is inactive and cannot be used to grant OAuth 2 tokens. This prevents any applications that are accessing the platform by granting OAuth 2 token using ROPC.
Dependencies and prerequisites	The OAuth 2.0 (com.snc.platform.security.oauth) plugin must be active.

To learn more about adding or creating a system property, see Add a system property.

Do not apply password policy at login [Updated in Security Center 1.5 and removed in 2.0]

Manage how password complexity is handled in your instance.

By setting the property *glide.apply.password_policy.on_login* to false there will be no password complexity enforcement at login time. Setting the property to true will enforce password complexity and lead to organization policy compliance issues.

As per ASVS 4.03 v2.1.9 recommendations:

"Verify that there are no password composition rules limiting the type of characters permitted. There should be no requirement for upper or lower case or numbers or special characters. (C6)"

Instead of password complexity enforcement, ASVS recommendations are to enforce a minimum length of 12 characters for password length.

Refer to OWASP ASVS v4.0 Authentication.

More information

Attribute	Description
Configuration name	*glide.apply.password_policy.on_login*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	false
Category	Authentication

Attribute	Description
Security risk	• Severity score: 4.4 • CVSS score: Medium • Security risk details: Setting this property to **true** could enforce password complexity and lead to organization compliance issues.
Dependencies and prerequisites	None

Enable account recovery [Updated in Security Center 1.3 and 1.5]

The *glide.sso.acr.enabled* property controls the account recovery feature.

Set *glide.sso.acr.enabled* to the recommended value of **true** to allow account recovery by userid possible. Set the value to **false** to disallow account recovery by userid.

More information

Attribute	Description
Property name	*glide.sso.acr.enabled*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	Controls the account recovery by userid feature.
Recommended value	True (default)
Configuration type	Boolean
Security risk	Critical (Without this property enabled, users will not be allowed to recover their account by userid.
Security risk rating	9.1
References	See Account recovery (ACR) for additional information.

Enable CAPTCHA for customer registration

Reduce the risk of requests by malicious bots by enabling CAPTCHA for customer registration.

Use the **sn_customerservice.captchaEnabled** system property to determine whether CAPTCHA validation is enabled for customer registration on the Customer Service Management Portal. Use CAPTCHA validation to help prevent potentially malicious bots from automatically submitting requests against an application.

Set the system property **sn_customerservice.captchaEnabled** to **true** to enable CAPTCHA validation. If the property isn't on the System Properties [sys_properties] table, the default value is **true**.

More information

Attribute	Description
Technical configuration name	sn_customerservice.captchaEnabled

Attribute	Description
Plugin applicability	Customer Service Management
Security risk	CAPTCHA validation helps prevent potentially malicious bots from automatically submitting requests against an application.
Common Vulnerability Scoring System (CVSS) score	3.7
Common Vulnerability Scoring System (CVSS) rating	Low
Functional impact	Registering users may have a negative experience from having to pass the CAPTCHA validation.
Dependencies and prerequisites	None
Data type	Boolean
Base system value	**true**
Fallback value	**true**
Recommended value	**true**

To learn more about adding or creating a system property, see Add a system property.

Enable a deny-list password validation check

Manage the deny-list passwords in the Excluded Password table.

Use the *glide.enable.blacklist_password* property to monitor deny-list passwords. When the property is set to **True**, the user's password is checked against a list of the deny-list passwords to prevent them from using a password from a set of breached passwords. The administrator can maintain the list by inserting passwords into the Excluded Password table.

More information

Attribute	Description
Configuration name	*glide.enable.blacklist_password*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Authentication
Dependencies and prerequisites	None
References	Exclude passwords through password policies on your instance

Enable Captcha for External User Registration [Updated in Security Center 1.3 and 1.5]

The *sn_ext_usr_reg.captchaEnabled* controls if CAPTCHA will be validated for external user registration.

Set *sn_ext_usr_reg.captchaEnabled* to the recommended value of **true** to help prevent automatic account creation attacks with requiring CAPTCHA for external user registration. Set the value to **false** to not require CAPTCHA for external user registration.

More information

Attribute	Description
Property name	*sn_ext_usr_reg.captchaEnabled*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	This property is used to enable or disable CAPTCHA validation while doing external user registration on portals like CSP, Community. This is also used in store apps like VAM and CSM Guest Walkup to enable/ disable captcha.
Recommended value	True
Configuration type	Boolean
Security risk	(Low) The property controls CAPTCHA enablement in external user registration. Unideal value may result in security vulnerability.
Security risk rating	3.7

To learn more about adding or creating a system property, see Add a system property.

Enable CAPTCHA in password reset

Use the *password_reset.captcha.ignore* property to enable or disable requiring a CAPTCHA challenge when a user resets their password.

Set *password_reset.captcha.ignore* to the recommended value of **false** to require a CAPTCHA challenge for a user to reset their password. Set the value to **true** to ignore the CAPTCHA option for a password reset.

CAPTCHAs help prevent automation attacks by prompting the user for a challenge-response that is not easily answered by automated systems. If CAPTCHA is disabled, an attacker may be more successful during automated attacks against the password reset feature.

Note: This property is used for password reset automation only.

More information

Attribute	Description
Property name	*password_reset.captcha.ignore*

Attribute	Description
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	This property is used to enable or disable CAPTCHA validation during password reset.
Recommended value	false
Configuration type	Boolean
Security risk	(Moderate) Unideal value may result in security vulnerability.
Security risk rating	5.5

To learn more about adding or creating a system property, see Add a system property.

Enable email OTP for multi-factor authentication

Manage how two-factor authentication is applied on your instance.

Use the *glide.authenticate.multifactor.email.otp.enabled* property to control whether a token for two-factor authentication is sent using email. Email is considered a weak MFA factor which an attacker is more likely to gain access into for bypassing MFA. By setting this property to **false**, the risk of an attacker bypassing MFA when they compromised a user's password is reduced.

More information

Attribute	Description
Configuration name	*glide.authenticate.multifactor.email.otp.enabled*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	true
Category	Authentication
Security risk	• Severity score: 3.1 • CVSS score: Low • Security risk details: Setting this property to **false** reduces the risk of a bad actor bypassing two-factor authentication.
Dependencies and prerequisites	None
References	Email as an MFA factor

Enable password reset policy checks [Updated in Security Center 2.0]

Use the **glide.enable.password_policy** property to enable password policy checks whenever a user changes their password using the user interface.

To define which password policy to use once this property is enabled, see Enable password policies on your instance. Ensure the Glide Property *glide.enable.password_policy* exists and is set to the value true. If the property does not appear in the sys_properties table, add a new record.

Note: The **glide.enable.password_policy** does not apply when an administrator changes a password or adds a user through script.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.enable.password_policy*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	To apply password policy at time of password change.
Recommended value	true (for higher strength passwords)
Security risk rating	7.4
Functional impact	Setting the property to true turns on password policy checks when a user resets their password.
Security risk	(Moderate) Without a password policy, a user can create a weak password which increases the likelihood of an adversary gaining access to the instance.

Steps to configure

If you configure this setting in the Hardening Compliance Configuration page in the Instance Security Center:

1. Under **Medium**, Select **Session Management**.
2. In the **Enable Password Reset Policy Check** setting, select **Medium** for medium strength passwords, or **Strong** for more robust, higher strength passwords. Selecting one of these options sets the **glide.enable.password_policy** property to true and starts a workflow that automatically updates your password policy.

Additionally, you can set the *glide.apply.password_policy.on_login* system property to enable password policy checks at the time of log in.

Enable policy based session access for mobile [New in Security Center 1.5]

Use the The Zero Trust- Policy Based Session Access plugin to control if users authenticating through a mobile app will have their roles reduced.

The Zero Trust- Policy Based Session Access plugin enables security admins to reduce user access in a session based on parameters such as IP address, location, identify provider attributes, and user attributes with adaptive authentication policies. When this plugin is enabled or set to true, users authenticating through a mobile device will have their roles restricted according to the plugin's policies. Instance admins may wish to restrict high privileged access

when users authenticate through a mobile device as it could indicate an unsafe environment for sensitive operations.

More information

Attribute	Description
Configuration name	*glide.authenticate.session_access.mobile.enabled*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Access control
Security risk	• Severity score: 4.7 • CVSS score: Medium • Security risk details: If this hardening setting is set to true, then policy based session access is enforced on the instance for mobile logins, and users that are not coming from a trusted environment, or are not using trusted devices will have their roles with reduced privileges. True is the secure setting. If this setting is configured to false, then policy based session access is disabled, and users will continue to have full roles including the high privileged roles like admin all the time.
Dependencies and prerequisites	None
Functional impact	If admin has configured the session access policy on the instance, then users will have their roles reduced after mobile login if they are not coming from a trusted environment or using a trusted device.
References	Adaptive authentication

Enable relay state in SAML requests to prevent replay attacks

Reduce the risk of replay attacks by enabling relay state in SAML requests to help prevent replay attacks.

Protect against SAML replay attacks using the **glide.authenticate.sso.saml2.enable_relay_state_with_id** system property. When this property is set to **true**, the relay state parameter contains the sys_id of a record in the MultiSSO Request Parameters [multisso_request_parameter] table, which the relay state URL redirects to.

Set the system property **glide.authenticate.sso.saml2.enable_relay_state_with_id** to **true**. This helps prevent attackers who have gained access to a SAML request from accessing your instance by resubmitting a valid request.

More information

Attribute	Description
Technical configuration name	glide.authenticate.sso.saml2.enable_relay_state_with_id
Plugin applicability	Multi-Provider SSO plugin (com.snc.integration.sso.multi.installer)
Security risk	The relay state enabled by this system property helps protect your instance against replay attacks. Enabling the property helps prevent attackers who have gained access to a SAML request from accessing your instance by resubmitting a valid request.
Common Vulnerability Scoring System (CVSS) score	3.8
Common Vulnerability Scoring System (CVSS) rating	Low
Functional impact	When this property is set to true, the relay state in a SAML request contains the sys_id of a record in the MultiSSO Request Parameters [multisso_request_parameter] table, which contains relay state URL to redirect to.
Dependencies and prerequisites	None
Data type	Boolean
Base system value	true
Fallback value	false
Recommended value	true

To learn more about adding or creating a system property, see Add a system property.

Enable SMS code notification for enrollment and verification [Updated in Security Center 1.3]

The *password_reset.sms.use_notify* property controls the usage of SMS code notifications for password reset.

If the *password_reset.sms.use_notify* property is set to the recommended value of **true**, the user is notified to do a password reset using SMS verification and new device enrollment which is more secure than email.

More information

Attribute	Description
Configuration name	*password_reset.sms.use_notify*
Configuration type	System Properties (/sys_properties_list.do)
Data type	booelan
Recommended value	true

Attribute	Description
Default value	true
Category	Authentication
Security risk	• Severity score: 3.7 • CVSS score: Low • Security risk details: Setting this property to**false** makes email the default method for password recovery which is less secure than SMS.
Dependencies and prerequisites	None

Enable SSL in LDAP authentication [Updated in Security Center 1.5 and 2.0]

Manage the encryption of LDAP authentication requests on your instance.

Use the *glide.ldap.use.ssl* property to enable or disable TLS encryption for LDAP authentication requests sent over the network. If this property is not set to the recommended value of **true**, LDAP authentication is susceptible to a man-in-the-middle attack.

More information

Attribute	Description
Configuration name	*glide.ldap.use.ssl*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Authentication
Security risk	• Severity score: 8.1 • CVSS score: High • Security risk details: Setting this property to **false**, makes LDAP authentication vulnerable to a man-in-the-middle attack.
Dependencies and prerequisites	None
References	Lightweight Directory Access Protocol integration

Enforce current password policy compliance requirements on login

Reduce the risk of brute force account login by enforcing current password policy compliance requirements on login.

Help prevent interactive users from logging in to the instance with passwords that don't meet current administrator requirements using the **glide.apply.password_policy.on_login** system property.

Set the **glide.apply.password_policy.on_login** system property to **true** to enforce current password policy compliance requirements. If this property doesn't exist on the System Properties [sys_properties] table, the default value is **false**.

More information

Attribute	Description
Technical configuration name	glide.apply.password_policy.on_login
Plugin applicability	None
Security risk	Interactive users may continue to log in to the instance with passwords that don't meet current administrator requirements. This may mean that the users have weak passwords that don't meet up-to-date security requirements, potentially leading to an increased risk of brute-force account logins by a malicious user.
Common Vulnerability Scoring System (CVSS) score	4.4
Common Vulnerability Scoring System (CVSS) rating	Medium
Functional impact	If the existing passwords aren't in compliance with your current password policy, enabling this property forces users to change the passwords during their next login. This property is automatically set to **false**. Setting the value to **true** enforces a password policy during login. **Note:** Enabling this property might force a significant number of users who aren't in compliance with the new password policy to change their passwords.
Dependencies and prerequisites	None
Data type	Boolean
Base system value	false
Fallback value	false
Recommended value	true

To learn more about adding or creating a system property, see Add a system property.

Enforce device encryption and passcode requirements [New in Security Center 1.3]

The *glide.sg.device_encryption_enabled* property enforces the Federal Information Processing Standard (FIPS 140-2) Encryption. Mobile device encryption and passcode ensure that an unauthorized user cannot access the content of a device even if the device is physically obtained.

When *glide.sg.device_encryption_enabled* is set to true, the ServiceNow mobile app will check that device encryption and device passcode are enabled.

More information

Attribute	Description
Configuration name	*glide.sg.device_encryption_enabled*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Authentication
Security risk	• Severity score: 4.2 • CVSS score: Medium • Security risk details: If encryption or passcode is not enabled, then the user will not be allowed to log into the instance on mobile
Dependencies and prerequisites	None
Functional impact	When this property is set to true, the mobile app will verify if device encryption is enabled. If encryption is not enabled, users will not be allowed to log into the current instance on mobile. Users are logged out and see the following warning message, suggesting that they set a device pin or encrypt the device and to try to login again. `You were logged out` `You need a passcode in order to use this instance on this device. Go to your device's settings to set one up`

Limit Invalid Password Reset Attempts [Updated in Security Center 1.3 and updated in 2.0]

The *password_reset.request.max_attempt* is used to control the maximum number of unsuccessful attempts that a user can reset or change their password before being locked out for a specified period of time.

More information

Attribute	Description
Property name	*password_reset.request.max_attempt*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	Denotes the maximum number of unsuccessful password reset attempts that can be taken before

Attribute	Description
	the user is locked out of password reset process. The lockout period is determined by the value in *password_reset.request.max_attempt_window*.
Recommended value	Set to a positive integer value less than three. The default value is **3**. When you determine the limit for the upper range of the property, consider the task that the user is performing.
Configuration type	Positive integer values
Security risk	(High) If the property is not set to the recommended value of "**3**" or other reasonable small value, then it could be possible to perform a brute force attack against the password reset process.
Security risk rating	7.5
References	Configure Password Reset properties

Managing unlock timeout after failed logins [Updated in Security Center 1.3]

Two script actions are available that enable a site administrator to manage the number of times a user can provide an incorrect password before being locked out from the ServiceNow AI Platform. You can enable either of these script actions to manage failed login attempts.

More information

Attribute	Description
Property/Plugin Name	N/A
Configuration type	System Policy > Script Actions
Category	Authentication
Purpose	To enforce strict policy for failed login attempts to avoid brute forcing of credentials.
Recommended value	Active
Security risk rating	7.3
Functional impact	This remediation would enable administrator of the instance to monitor and report any malicious user access. No functionality impact, only User experience change.
Security risk	(Moderate) Apply a defined logging and auditing strategy so that you can identify and act on suspicious activity in a timely manner.

Steps to configure

1. Navigate to **System Policy** > **Script Actions.**
2. Search for the name **SNC User*.

3. To enable management of failed login attempts, change the Active state of either the *SNC User Lockout Check with Auto Unlock* or *SNC User Lockout Check* scripts actions from **false** to **true**.
4. To reset the failed login counter after a successful login, you can activate the *SNC User Clear* script action.

Maximize failed login unlock timeout duration [Updated in Security Center 1.3]

A script action is available that enables site administrators to manage the number of times a user can provide an incorrect password before being locked out from the ServiceNow AI Platform. You can enable this script action to manage failed login attempts.

Help secure your instance against brute force attacks by defining a time period during which a user cannot attempt to log in after being locked out. The **glide.user.unlock_timeout_in_mins** system property unlocks the user account after the time period that is specified in it's value. If no value is specified, your instance unlocks the user account after the default period of 15 minutes.

Set the **glide.user.unlock_timeout_in_mins** system property value to a minimum of **15**. If **glide.user.unlock_timeout_in_mins** does not exist, the default lockout time is set to 15 minutes.

Ensure that the **SNC User Lockout Check with Auto Unlock** script action (found on the Script Action [sysevent_script_action] table) is present and active. The **SNC User Lockout Check with Auto Unlock** script action is installed with the High Security Settings (com.glide.high_security) plugin.

More information

Attribute	Description
Configuration name	• glide.user.unlock_timeout_in_mins (System Property) • sysevent_script_action (Script Action)
Configuration type	System Policy > Script Actions
Category	Authentication
Purpose	To enforce strict policy for failed login attempts to avoid brute forcing of credentials.
Recommended value	• **15** for the **glide.user.unlock_timeout_in_mins** system property • **Active** for the **SNC User Lockout Check with Auto Unlock** script action.
Functional impact	This remediation would enable administrator of the instance to monitor and report any malicious user access. No functionality impact, only User experience change.
Security risk	• Severity Score: 6.8 • Security Risk Details: If the property is not configured to a secure value and the lockout duration is not enabled, then it may be easier to brute force account logins in a faster time frame. This may allow a

Attribute	Description
	malicious user to eventually obtain unauthorized access to the instance. Impact on the instance will be limited to the privileged of the affected user login brute-forced.

Steps to configure

1. Navigate to **System Policy** > **Script Actions.**
2. Search for the name **SNC User*.
3. To enable management of failed login attempts, change the Active state of either the *SNC User Lockout Check with Auto Unlock* or *SNC User Lockout Check* scripts actions from **false** to **true**.
4. To reset the failed login counter after a successful login, you can activate the *SNC User Clear* script action.

Maximize reset password request retry window duration [Updated in Security Center 1.3]

The *password_reset.request.retry_window* property controls the number of minutes before the count for password reset attempts refreshes.

More information

Attribute	Description
Property name	*password_reset.request.retry_window*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	Denotes the length of time in minutes before the count for password attempts refreshes from the last request before the retry count is reset to zero.
Recommended value	Set to a positive integer value of **1440** or more. The default value is **1440** minutes.
Configuration type	Positive integer values.
Security risk	(High) If the property is not set to the recommended value of 1440 or more, then it could be possible to perform account brute force against password reset process.
Security risk rating	7.5
References	Configure Password Reset properties

To learn more about adding or creating a system property, see Add a system property.

Maximize reset password request unlock window duration [Updated in Security Center 1.3]

The *password_reset.request.unlock_window* property controls the number of minutes a user must wait to start a reset request after the last successful unlock account action.

This property controls the number of minutes a user must wait to start a reset request after the last successful unlock account. If *password_reset.request.unlock_window* is not set to the recommended value of 1440 or more, it increases the opportunity for a malicious actor from brute forcing the user's password using automated tools.

More information

Attribute	Description
Property name	*password_reset.request.unlock_window*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	It denotes the time period in minutes that a user must wait after successfully resetting the password to reset the password again.
Recommended value	1440
Default value	1440
Configuration type	Positive integer values
Security risk	(High) If the property is not set to the recommended value of 1440 or greater, then it increases the opportunity of a malicious actor to brute force password access using automatic tools.
Security risk rating	5.9
References	Configure Password Reset properties

To learn more about adding or creating a system property, see Add a system property.

Maximize reset password SMS complexity [Updated in Security Center 1.3]

The *password_reset.sms.default_complexity* property controls the minimum required SMS code verification size required during password reset.

More information

Attribute	Description
Property name	*password_reset.sms.default_complexity*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	Denotes the SMS code verification size required during password reset.
Recommended value	6
Default value	4
Configuration type	Integer value greater than zero
Security risk	(Low) If the property is not set to the recommended value, then a weak SMS validation token is used. This

Attribute	Description
	increases the possibility of token guessing which could lead to account takeover.
Security risk rating	3.8

To learn more about adding or creating a system property, see Add a system property.

Maximize reset password SMS pause window duration [Updated in Security Center 1.3]

Manage the time duration in minutes that a user must wait before they can request a new password reset code.

If this property is not set to the recommended value of **2** minutes or more, then a malicious user could initiate many passwords reset codes in a brief window of time. This increases the chance of a bad actor predicting the SMS reset code.

More information

Attribute	Description
Configuration name	*password_reset.sms.pause_window*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	2
Default value	2
Category	Authentication
Security risk	• Severity score: 4.8 • CVSS score: Medium • Security risk details: Ensure that *password_reset.sms.pause_window* is set to a value of 2 or more.
Dependencies and prerequisites	None

Maximize reset password verification delay duration [Updated in Security Center 1.3]

Configure the delay, in milliseconds, that a user must wait before submitting a new password reset request.

A bad actor could attempt to brute force login credentials by using automation tools like bots which the **reset password verification delay** property helps defend against. The property value represents the delay, in milliseconds, that a user must wait before they can place a request to reset the password. If this property is not set to the recommended value of **1000** or more, the login is more vulnerable to brute force attacks.

More information

Attribute	Description
Configuration name	*password_reset.verification.delay*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	1000
Default value	1000
Category	Authentication
Security risk	• Severity score: 5.9 • CVSS score: Medium • Security risk details: Setting the property value to less than 1000 makes your login more vulnerable to brute force attacks.
Dependencies and prerequisites	None
References	Configure password for a user

Minimize external user registration link expiration duration [Updated in Security Center 1.3 and 1.5]

Manage the number of days that a registration link can be accessed.

Use the *sn_ext_usr_reg.Reg_link_expiration_days* property to manage who can access a registration link. If the link is set to the recommended value of **3**, a registration link could be used by someone other than the intended user if the link is discovered at a later date.

More information

Attribute	Description
Configuration name	*sn_ext_usr_reg.Reg_link_expiration_days*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Integer
Recommended value	3
Default value	3
Category	Authentication
Security risk	• Severity score: Medium • CVSS score: 6.6 • Security risk details: Not setting this property to the integer 3 could lead to the registration link being used by an unintended user.

Attribute	Description
Dependencies and prerequisites	None

Minimize reset password max SMS per day [Updated in Security Center 1.3]

Manage the maximum number of SMS codes sent for verification per day by user.

The *password_reset.sms.max_per_day* property represents the maximum number of SMS codes that can be sent for verification daily by a user.

More information

Attribute	Description
Configuration name	*password_reset.sms.max_per_day*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	10
Default value	10
Category	Authentication
Security risk	• Severity score: 5.9 • CVSS score: Medium • Security risk details: If this property is not set to the recommended value of **10** or less, it's easier to brute force the SMS code.
Dependencies and prerequisites	None

Minimize failed login attempts for high assurance sessions

Decrease the likelihood of a brute force attack by minimizing failed login attempts for high assurance sessions.

Use the **glide.zta.high_assurance.session.max.login.failed_attempts** system property to limit the number of failed authentication attempts allowed before users are logged out when re-authenticating through Continuous Authentication.

Set the value of this system property to a low value (such as 5) to decrease the likelihood of a brute force attack.

More information

Attribute	Description
Technical configuration name	glide.zta.high_assurance.session.max.login.failed_attempts
Plugin applicability	Zero Trust - Continuous Authentication (com.snc.zero_trust_continuous_authentication)
Security risk	A high number of allowed authentication attempts increases the likelihood of a brute force attack.

Attribute	Description
Common Vulnerability Scoring System (CVSS) score	3.3
Common Vulnerability Scoring System (CVSS) rating	Low
Functional impact	Users are logged out of their sessions after the number of authentication failures selected in the property.
Dependencies and prerequisites	None
Data type	Integer
Base system value	5
Fallback value	5
Recommended value	5

To learn more about adding or creating a system property, see Add a system property.

Minimize reset password request expiration duration [Updated in Security Center 1.3]

The *password_reset.request.expiry* denotes the time period in minutes during which a user must perform the password reset process.

Note: The setting for the *password_reset.request.expiry* property takes precedence over the setting for *glide.pwd_reset.onetime.token.validity*property that has a 12 hour default.

More information

Attribute	Description
Property name	*password_reset.request.expiry*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	Denotes the time period in minutes during which a user must perform the password reset process.
Recommended value	Set to an integer of **10** or less. The default value is 10.
Configuration type	Integer values
Security risk	(Moderate) If the property is not set to the recommended value of 10 or less, then it increases the opportunity for someone else to guess and use the request and attempt to reset the password.
Security risk rating	4.2
References	Configure Password Reset properties

To learn more about adding or creating a system property, see Add a system property.

Minimize reset password request success window duration [Updated in Securty Center 1.3]

The *password_reset.request.success_window* property controls the number of minutes a user must wait to reset or change their password again after successfully resetting the password. The user will be blocked to reset the password again for the specified duration.

More information

Attribute	Description
Property name	*password_reset.request.success_window*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	It denotes the time period in minutes that a user must wait after successfully resetting the password to reset the password again.
Recommended value	1440
Default value	1440
Configuration type	Positive integer values
Security risk	(High) If the property is not set to the recommended value of 1440 or less,then it increases the opportunity of someone else abusing the password reset functionality to gain unauthorized access to a user account.
Security risk rating	4.9
References	Configure Password Reset properties

To learn more about adding or creating a system property, see Add a system property.

Minimize reset password SMS expiracy duration [Updated in Security Center 1.3]

Control the number of minutes remaining before the SMS code expires.

The *password_reset.sms.expiry* property represents the number of minutes remaining before the SMS code expires.

More information

Attribute	Description
Configuration name	*password_reset.sms.expiry*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	5
Default value	5
Category	Authentication

Attribute	Description
Security risk	• Severity score: 5.6 • CVSS score: Medium • Security risk details: If this property is not set to the recommended value of **5** or less, the chances of a bad actor guessing and using the SMS code to reset the password increases.
Dependencies and prerequisites	None

Minimize one-time out of band verifier lifetime duration [Updated in Security Center 1.3]

Manage the time duration for out-of-band verifiers.

An out-of-band verifier is an alternative delivery method for one-time code situations. For example, resetting a multi-factor token. If this method is enabled by administrators in the Multi-factor authentication plugin, a one-time code is delivered by email. Set one-time out-of-band verifiers to expire after 10 minutes to limit the validity window. A larger time window allows more time for the code to be compromised through illicit means such as phishing, social engineering, or shoulder-surfing attacks.

More information

Attribute	Description
Configuration name	*glide.multifactor.onetime.code.validity*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	10
Default value	10
Category	Authentication
Security risk	• Severity score: 3.9 • CVSS score: Low • Security risk details: Set one-time out-of-band verifiers to expire after 10 minutes. Anything longer increases the risk of the code being compromised by a bad actor.
Dependencies and prerequisites	Multi-factor authentication
References	Multi-factor Authentication criteria

Minimize SAML notBefore or notOnOrAfter constraint duration [Updated in Security Center 1.3 and 1.5]

Configure this property to add a grace period in which SAML requests and responses are considered valid.

This property adds a grace period during which SAML requests and responses are considered valid. The property value represents the number of seconds to add to the *NotBefore* and *NotOnOrAfter* constraints to account for time differences between the Identity Provider (IdP) clock, and Service Provider (SP) clock. These constraints defend against replay attacks by denying requests that aren't made within the specified time frame. If the IdP and SP clocks are significantly different, then the network latency may result in the SAML request being unauthorized.

More information

Attribute	Description
Configuration name	*glide.authenticate.sso.saml2.clockskew*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	less than 60
Default value	180
Category	Authentication
Security risk	• Severity score: 7.5 • CVSS score: High • Security risk details: Setting the property to a value of 60 or higher may prevent the constraints from defending against replay attacks.
Dependencies and prerequisites	None

Notify users during password reset/change process [Removed in Security Center 1.5]

Use this property to enable end users to reset or change passwords using a self-service process.

This property enables an end user to reset or change a password using a self-service process. Alternatively, your organization could implement a process that requires a service desk agent to reset passwords for end users. If a password change and or reset process doesn't notify users on password update, a bad actor may be able to lock that user out of their account without their knowledge. This would provide the bad actor more time to perform malicious activities. Ensure password reset process notifies users upon password change or reset.

More information

Attribute	Description
Configuration name	*pwd_process.change, pwd_process.reset*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true

Attribute	Description
Category	Authentication
Security risk	• Severity score: 8.1 • CVSS score: High • Security risk details: A bad actor may be able to lock a user out of their account without their knowledge if no notification is sent to them when a password change is reset.
Dependencies and prerequisites	None

Reduce allowed bypasses for multifactor setup

Decrease the window of time an account is at risk of compromise by reducing allowed bypasses for multifactor setup.

Reduce the number of times a user can skip the multifactor passcode setup requirement to decrease the window of time an account is at risk of compromise. Multifactor authentication (MFA) protects against password-related attacks or weak passwords by requiring an additional form of verification. Reducing the length of time a user can skip this setup reduces this vulnerability.

Set the **glide.authenticate.multifactor.setup.bypass.count** system property to **0** to prevent users from skipping multifactor passcode setup.

More information

Attribute	Description
Technical configuration name	glide.authenticate.multifactor.setup.bypass.count
Plugin applicability	None.
Security risk	Multifactor authentication protects against password-related attacks or weak passwords by requiring an additional form of verification. A high number of allowed multifactor setup bypasses increases the risk of account compromise as the account isn't protected by multifactor. A small number of allowed bypasses reduces this window of time.
Common Vulnerability Scoring System (CVSS) score	3.9
Common Vulnerability Scoring System (CVSS) rating	Low
Functional impact	Users aren't allowed to login to an instance without setting up MFA if they've surpassed the number of logins specified in the property.
Dependencies and prerequisites	None
Data type	Integer
Base system value	0

Attribute	Description
Fallback value	0
Recommended value	0

To learn more about adding or creating a system property, see Add a system property.

Remove credentials from Welcome page

Modify the default content on the Welcome page to remove the default credentials.

Two **How to Login** records are installed as part of the demo data for the CMS plugin.

Note: If you do not install the demo data for an instance, these records do not exist. In that case, the configuration is considered as Security Compliant as per the recommended security practices.

More information

Attribute	Description
Name	*How to login*
Configuration type	Table: sys_home
Category	Authentication
Purpose	To remove default credentials from Welcome page that were added with demo data.
Recommended value	False or null if no demo data was utilized.
Default value	None. This is a table configuration, not a Glide property. Therefore, there is no value for default value.
Configuration type	Boolean
Security risk	(Moderate) Demo data is provided for the CMS plugin, which includes two default passwords included on the welcome page. If this is not removed, an unauthorized attacker could gain access to the instance.
References	https://support.servicenow.com/kb_view.do?sysparm_article=KB0550107 Welcome pages

To learn more about adding or creating a system property, see Add a system property.

Require captcha for guest walk-up experience in customer service application [New in Security Center 1.3 and updated in 1.5]

The captcha for the Guest Walk-up experience prevents unauthenticated guest users to create bookings by requiring users to complete a captcha verification.

More information

Attribute	Description
Configuration name	*sn_guest_walkup_cs.captcha.enabled*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Authentication
Security risk	• Severity score: 3.7 • CVSS score: Low • Security risk details: If captcha is not enabled, this could lead to automated creation of spam appointments to overwhelm the system or fill up all available booking spots creating a Denial of Service (DoS) attack.
Dependencies and prerequisites	None
Functional impact	This property enables or disables the captcha on the CSM Guest Walkup Check-in widgets. By default, it is set to true.

Require obfuscation of classic mobile app UI [Updated in Security Center 1.3]

Use the *glide.ui.m.blur_ui_when_backgrounded* property to obfuscate all fields from the snapshot as the image is saved during the backgrounding process.

On Android devices, the Android operating system takes a screenshot for usage in the recent task menu when the application is sent to background. Users can also take manual screenshots of the application, which are stored publicly on the device.

On iOS devices, the iOS operating system also allows applications to save an image file. This file represents the last screen seen by the user when the application is sent to the background. While the intent is to provide a better user experience, it also creates a security risk because the images are saved as PNG image files.

Note: This setting or configuration is per instance basis so the user must connect to the instance with the property configured.

To obfuscate all fields from the snapshot in the ServiceNow Classic app, see Configure the blur app option to improve security.

Example

When you set this property to true, the background application is obfuscated for iOS devices, and blacked out for Android iOS devices.

More information

Attribute	Description
Property/Plugin Name	*glide.ui.m.blur_ui_when_backgrounded*
Configuration type	System Properties (/sys_properties_list.do)
Category	Authentication
Purpose	To obfuscate all fields from the snapshot as the image is saved during the backgrounding process.
Recommended value	true
Security risk rating	2.4
Functional impact	If the *glide.ui.m.blur_ui_when_backgrounded* property is set to true, the native apps use the parameter defined on the server to blur the screen when the app enters the background. • It blurs the screenshots taken by iOS and Android when the app enters the background. **Note:** Enabling this setting can prevent users from taking screenshots of the classic mobile app, which may be desirable for support purposes. • The user experience may be adversely affected because they would not be able to see the content when the app is sent to background.
Security risk	(Medium) A compromised (jailbroken) device would enable an attacker to have full access to the file system, with access those files/snapshots with sensitive information embedded in them.

To learn more about adding or creating a system property, see Add a system property.

Require obfuscation of mobile app UI [Updated in Security Center 1.3]

Configure the *glide.sg.blur_ui_when_backgrounded* property so that the UI of the app is blurred when the app is running in the background.

If this property is not set to the recommended value of **true**, the mobile app's user interface is visible when viewed from the app switcher. The UI is still visible even when the app is running in background which provides a lower level of security and confidentiality to end users.

More information

Attribute	Description
Configuration name	*glide.sg.blur_ui_when_backgrounded*
Configuration type	System Properties (/sys_properties_list.do)

Attribute	Description
Data type	Boolean
Recommended value	true
Default value	false
Category	File and resources
Security risk	• Severity score: 2.4 • CVSS score: Low • Security risk details: Setting the value to **true** provides a higher level of confidentiality and privacy on the local device by blurring the UI when the app is running in the background.
Dependencies and prerequisites	None
References	Require obfuscation of classic mobile app UI [Updated in Security Center 1.3]

Set minimal password length [Updated in Security center 2.2]

Set minimal password length to avoid compliance issues and reduce the risk of a successful brute force attack

Enforce a minimum password length of at least 12 characters to avoid compliance issues and reduce the risk of a successful brute force attack.

For every utilized Password credential store in use on your instance, ensure that a Password Policy is being enforced by selecting the **Enable password policy** field on the associated record in on Password Reset Credential Store [pwd_cred_store] table.

Next, open the record on the Password Policy [password_policy] record table and set the **Minimum Password Length** field to at least **12**. You can find the associated Password Policy record in the **Password policy** field of the Password Reset Credential Store [pwd_cred_store] record.

For more information configuring a password policy, see Enable password policies on your instance.

More information

Attribute	Description
Technical configuration name	Records on the Password Reset Credential Store [pwd_cred_store] and Password Policy [password_policy] tables.
Plugin applicability	None
Security risk	Setting the **Minimum Password Length** Policy [password_policy] records to value of less than 12 could lead to compliance issues and increases the risk of an attacker successfully brute forcing passwords.

Attribute	Description
Common Vulnerability Scoring System (CVSS) score	5.9
Common Vulnerability Scoring System (CVSS) rating	Medium
Functional impact	The instance will not suffer any impact from a minimum password length of 12.
Dependencies and prerequisites	None
Data type	Integer
Base system value	8
Fallback value	8
Recommended value	12

Set OTP lifetime for password reset to 1 hour [Updated in Security Center 2.0]

Control the time duration of the link in the password reset email.

The *glide.pwd_reset.onetime.token.validity* system property makes the link in the password reset email expire after the number of hours specified in the property. The validity time of a password reset token should be kept as short as possible while not disrupting normal user experience

Set the property value to 1 (in hours).

More information

Attribute	Description
Configuration name	*glide.pwd_reset.onetime.token.validity*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	1
Default value	1
Fallback value	1
Category	Authentication
Security risk	• Severity score: 4.6 • CVSS score: Medium • Security risk details: A longer validity time for password reset token gives malicious actors a wider window to perform account takeover if the email with the reset token is leaked or compromised.
Dependencies and prerequisites	None

Business Logic

This category looks at the logic and flow unique to each application with general secure principles. Specifically ensure that the intended sequence of business logic flow cannot by bypassed, that limits exist to detect and prevent automated attacks, and that protections against spoofing, tampering, information disclosure and elevation of privilege attacks exist.

The following are some security controls that an administrator can configure to restrict unauthorized access to sensitive entities within the ServiceNow AI Platform.

Limit max comments per user per day

Configure the *sn_kb_social_qa.max_comments_per_user_daily* property to restrict the number of QA comments per day.

If this property is not set to the recommended value of **500** or less, there is no restriction on the number of QA comments per day which could lead to resource exhaustion.

More information

Attribute	Description
Configuration name	*sn_kb_social_qa.max_comments_per_user_daily*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	500
Default value	500
Category	Business Logic
Security risk	• Severity score: 3.7 • CVSS score: Low • Security risk details: Not setting the recommended value of this property to 500 or less places no limit on QA comments per day which could lead to resource exhaustion.
Dependencies and prerequisites	None

Limit max subscriptions per user per day

Configure the *sn_kb_social_qa.max_subscriptions_per_user_daily* property to limit the max number subscriptions a user can subscribe to in a day.

If this property is not set to the recommended value of **500** or less, there is no restriction on the maximum number of Q&A questions that a user can subscribe to in a day. This no limitation could lead to resource exhaustion and impact the availability of your instance.

More information

Attribute	Description
Configuration name	*sn_kb_social_qa.max_subscriptions_per_user_daily*

Attribute	Description
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	500
Default value	500
Category	Business Logic
Security risk	• Severity score: 3.7 • CVSS score: Low • Security risk details: Set the property value to 500 or less to prevent resource exhaustion.
Dependencies and prerequisites	None

Minimize SMTP Recipient Quantity [Updated in Security Center 1.3]

The *glide.email.smtp.max_recipients* specifies the maximum number of recipients the instance can list in the **To:** line for a single email notification.

Set *glide.email.smtp.max_recipients* to the recommended value of **100 or less**. Notifications that would exceed this limit instead create duplicate email notifications addressed to a subset of the recipient list. Each email notification has the same maximum number of recipients.

More information

Attribute	Description
Property name	*glide.email.smtp.max_recipients*
Configuration type	System Properties (/sys_properties_list.do)
Category	Business Logic
Purpose	If this property is set to an insecure value above the default value of 100, then its possible a denial of service could happen on the instance.
Recommended value	100
Default value	100
Configuration type	Integer
Security risk	(Moderate) Notifications that would exceed this limit instead create duplicate email notifications addressed to a subset of the recipient list.
Security risk rating	4.9

To learn more about adding or creating a system property, see Add a system property.

Timeout Guest Sessions

Use a system property to control the inactive session timeout for unauthenticated users.

Use the *glide.guest.session_timeout* system property to set the inactive session timeout duration (in minutes) for unauthenticated users. Raise the value of this property to extend the time your instance persists sessions beyond the default of 30 minutes. Avoid large timeout values, which can increase the number of sessions persisted by the instance, and cause minor availability concerns.

More information

Attribute	Description
Configuration name	*glide.guest.session_timeout*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Integer (in minutes)
Recommended value	30
Default value	30
Fallback value	0
Category	Business Logic
Security risk	• Severity score:4.3 • CVSS score: Medium • Security risk details: Large timeout values can increase the number of concurrent sessions on your instance, causing minor availability concerns.
Dependencies and prerequisites	None

Validate remote host

Set the property to true to prevent bad actors from using internal port scanning in your network.

If the *glide.update_set.remote.check_host* property is not set to the recommended value of **true**, then the remote instance test feature will allow internal port scanning which is a method bad actors can use to discover vulnerabilities in a network. It is then possible to enumerate all open ports on a given host, and in some cases pull response data which could lead to information leakage or unauthorized data access.

⚠ **Warning:** This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Configuration name	*glide.update_set.remote.check_host*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true

Attribute	Description
Category	Business Logic
Security risk	• Severity score: 6.3 • CVSS score: Medium • Security risk details: Not setting property to the recommended value of **true** could enable bad actors to use internal port scanning to gain access to unauthorized data.
Dependencies and prerequisites	None
References	https://support.servicenow.com/kb?id=kb_article_view&sysparm_article=KB0755132 Define a remote instance

Communications

This control ensures proper encryption using strong algorithms and ciphers. This includes ensuring the recommended version of TLS is used for client connectivity, use of strong cipher suites, use of trusted and signed certificates, ensuring connections are encrypted between components and logging of connection failures.

Enforce certificate trust [Updated in Security Center 1.3, removed in 2.0, added in 7.0]

Use system properties to ensure that certificate expiration and trust are checked for certificates received from outbound HTTPS call endpoints when host verification is not performed.

When *com.glide.communications.trustmanager_trust_all* is set to **true**, then certificate expiration and trust are not checked for the certificate received from an outbound HTTPS call endpoint when host verification is not performed.

Verify that the *com.glide.communications.trustmanager_trust_all* system property is set to the recommended value of **false**. This ensures that your instance only trusts certificates that it can verify against the JVM certificate store. Self-signed and enterprise-signed certificates are not trusted. This property only applies when *com.glide.communications.httpclient.verify_hostname* is set to **false**.

Note: The values for these properties are and cannot be altered once changed (they are non-revertible). For security purposes, do not change this property value. If you have further questions, contact Customer Service and Support.

More information

Attribute	Description
Property name	*com.glide.communications.trustmanager_trust_all*
Configuration type	System Properties (/sys_properties_list.do)
Category	Communications

Attribute	Description
Purpose	To enforce certificate validation for outgoing requests.
Recommended value	false
Security risk rating	5.7
Functional impact	This remediation enforces strict validation on certificate CA (certificate authority) field. If a trusted entity (CA) issued the certificate, the instance accepts it for further use.
Security risk	(Medium) For confidentiality and integrity reasons, application should validate the certificate's CA before using the certificate for any transactional operations.
References	Certificates Verify certificate chain and hostname [New in Security Center 1.3 and updated in 2.0]

Disable outbound SSLv2/SSLv3 connections [Updated in Security Center 1.3]

Use the *glide.outbound.sslv3.disabled* property to force the MID Server to use TLS when making outbound connections, such as REST and SOAP requests. Normally, outbound connections from an instance are forced to use TLS instead of SSL.

More information

Attribute	Description
Property name	*glide.outbound.sslv3.disabled*
Configuration type	System Properties (/sys_properties_list.do)
Category	Communications
Purpose	To enforce the use if TLS during all outbound connections from ServiceNow instance.
Recommended value	true
Default value	false **Important:** The value for the *glide.outbound.sslv3.disabled* property is a safe override and cannot be altered once changed.
Security risk rating	6.5
Functional impact	This remediation enforces the usage of TLS protocol version when communicating on HTTPS. If there are devices that customer/users of the instance are using that do not support TLS communication, there may be a potential outage.
Security risk	(Moderate) Legacy versions of SSL were proven to be insecure when utilized for HTTP secure shell

Attribute	Description
	implementation, due to client-side attacks, including BEAST and SSL heart-bleed.

To learn more about adding or creating a system property, see Add a system property.

Do not use demo certificates for active saml configurations [Updated in Security Center 1.5]

Control whether demo certificates are used in production SAML configurations.

The demo certificates provided by ServiceNow should not be used in production SAML configurations because they are common among all instances with a known passphrase. If one of the SAML properties using a certificate keystore is active (*require_signed_authnrequest*, *require_signed_logoutrequest*, or *encrypt_assertion*), then the demo data shouldn't be used. Since demo data is shared among all instances, there is no integrity guarantee of requests signed with shared certificates. Therefore, any message encrypted by the IDP could be decrypted by a bad actor if intercepted.

More information

Attribute	Description
Configuration name	*glide.authenticate.sso.saml2.keystore*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	sys_id of a custom keystore
Default value	empty string
Category	Communications
Security risk	• Severity score: 3.9 • CVSS score: Low
Dependencies and prerequisites	None

Disable deprecated TLS versions

Avoid loss or leakage of sensitive data by disabling deprecated TLS versions.

Confirm that your instance only negotiates for communication via Transport Layer Security (TLS) versions 1.2 when communicating with other servers to help prevent data transfer over weaker TLS 1 and TLS 1.1 protocol versions.

Set the **com.glide.communications.disable.deprecated.tls** to **true** to use only TLS versions 1.2 and up.

More information

Attribute	Description
Technical configuration name	com.glide.communications.disable.deprecated.tls

Attribute	Description
Plugin applicability	None
Security risk	Using outdated and unsupported TLS versions could result in loss and leakage of sensitive data.
Common Vulnerability Scoring System (CVSS) score	4.4
Common Vulnerability Scoring System (CVSS) rating	Medium
Functional impact	When this property is set to **true**, older insecure servers that require a weaker TLS 1 or 1.1 protocol aren't able to communicate with your instance.
Dependencies and prerequisites	None
Data type	Boolean
Base system value	true
Fallback value	true
Recommended value	true

To learn more about adding or creating a system property, see Add a system property.

Enforce OCSP check on network error [New in Security Center 1.3 and updated in 2.0]

Learn how to configure the *com.glide.communications.httpclient.ocsp_allow_network_error* property to prevent bad actors from bypassing Online Certificate Status Protocol (OCSP) checks.

If *com.glide.communications.httpclient.ocsp_allow_network_error* is not set to the recommended value of false, and the Online Certificate Status Protocol (OCSP) check encounters a network error (for example, a timeout or problem fetching the revocation information), it will bypass the OCSP security check and consider it successful. This could allow an attacker with a revoked certificate to break the Public Key infrastructure (PKI) and digital certificate trust that is foundational to the web. The use of revoked certificates is often an indicator of malicious activity unless the servers are out of sync.

Ensure the property *com.glide.communications.httpclient.ocsp_allow_network_error* exists and is set to false. If the property does not appear in the sys_properties table, add a new record.

More information

Attribute	Description
Configuration name	*com.glide.communications.httpclient.ocsp_allow_net*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	true

Attribute	Description
Category	Communications
Security risk	• Severity score: 5.9 • CVSS score: Medium • Security risk details: Not setting this property to false could enable a bad actor to bypass the OCSP security check.
Dependencies and prerequisites	None
Functional impact	This property determines whether a request against the Authority Information Access (AIA) Online Certificate Status Protocol (OCSP) uri results in a pass or fail outcome in the event of a connection or timeout error. When set to false, the revocation status of the presented server certificate can't be validated and will lead to a communication failure with that endpoint. If a network error occurs when the property is set to its default value of true, the certificate is treated as valid from a revocation standpoint.

Verify certificate chain and hostname [New in Security Center 1.3 and updated in 2.0]

Configure the *com.glide.communications.httpclient.verify_hostname* property to prevent man-in-the-middle-attacks by ensuring that the certification verification process is executed.

When the Glide Property *com.glide.communications.httpclient.verify_hostname* is not set to the secure value of true, the hostname and certificate chain presented by remote hosts during a TLS connection initiated from the ServiceNow instance are not validated. This could compromise the security of the TLS connection and allow person-in-the-middle attacks, where communications between two parties are intercepted. This may lead to sensitive data disclosure.

More information

Attribute	Description
Configuration name	*com.glide.communications.httpclient.verify_hostnam*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Communications
Security risk	• Severity score: High • CVSS score: 7.4

Attribute	Description
	• Security risk details: Not setting *com.glide.communications.httpclient.verify_hostn* to the recommended value of true could make your instance vulnerable to man-in-the-middle-attacks.
Dependencies and prerequisites	None
Functional impact	Verifies hostname and certificate chain presented by remote secure socket layer (SSL) hosts. Set this property to true to secure against Man-in-the-middle (MITM) attacks. **Note:** This property overrides the *com.glide.communications.trustmanager_trust_ali* property.

Verify certificate revocation [New in Security Center 1.3]

The *com.glide.communications.httpclient.verify_revoked_certificate* property checks certificate revocation during the Transport Layer Security (TLS) handshake to ensure that security checks are not bypassed.

If *com.glide.communications.httpclient.verify_revoked_certificate* is not set to the recommended value of **true**, then certificate revocation will not be checked during the TLS handshake. TLS encrypts data sent over the Internet to ensure that bad actors are unable to see sensitive information such as passwords or credit card numbers. Bypassing the TLS handshake is a security risk because an attacker with a revoked certificate can neglect to provide a valid certificate and break public key infrastructure (PKI) and digital certificate trust.

More information

Attribute	Description
Configuration name	*com.glide.communications.httpclient.verify_revoked*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Communications
Security risk	• Severity score: 6.5 • CVSS score: Medium • Security risk details: Not setting *com.glide.communications.httpclient.verify_revok* to the recommended value of true will cause certificate revocation to not be checked during the TLS handshake.
Dependencies and prerequisites	None

Attribute	Description
Functional impact	This property should be set to true to ensure that a Transport Layer Security (TLS) session is started with an authentic endpoint. If this property is set to false, then the certificate is not checked, which could compromise the security of the instance.

Configuration

The Configuration category ensures applications have a secure build environment and hardened third party library components. Specifically, ensuring a build and deploy pipeline is repeatable and includes automated testing and prevents known security issues from being deployed. This includes keeping dependencies up to date and free from known vulnerabilities.

Auto set content type options [Removed in Security Center 1.3.3]

Configure the Auto set content type options property on your instance to prevent MIME confusion attacks.

Use this property to control the X-Content-Type-Options response HTTP header. The X-Content-Type-Options response HTTP header is used by the server to indicate that the MIME types advertised in the Content-Type headers should be followed. If this property is set to false, then it is possible for an attacker to conduct MIME confusion attacks; if set to true then this header will prevent the browser from interpreting files as anything but the content type in the HTTP headers.

Warning: The value for this property is a no DB override. It can't be altered or overridden.

More information

Attribute	Description
Configuration name	*glide.security.header.auto_set_x_content_type_opti*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Configuration
Security risk	• Severity score: 7.3 • CVSS score: High • Security risk details: Setting this property to **false** could make it possible for an attacker to conduct MIME confusion attacks.
Dependencies and prerequisites	None
References	Add a system property

Cache-Control HTTP Header Value [Updated in Security Center 1.3 and removed in 1.5]

Use the *glide.http.cache_control* property to set the default cache-control value in the HTTP response headers that the ServiceNow AI Platform sends when requesting static content data for a page. Examples of static content include images, CSS, and JavaScript rendered from within, for a page.

The *glide.http.cache_control* property sets the default Cache-Control value in HTTP response headers to **private** or **public**. The default is **private**.

Value	Description
private	Static content can be cached at the browser (client) level, but not at the proxy server level.
public	Static content can be cached at the browser (client) level, and also at the proxy server level.

The Expires value in the HTTP response headers control when the static content expires, and has a default value of 369 days. To manually override the default value, use the *glide.http.expire.days* property.

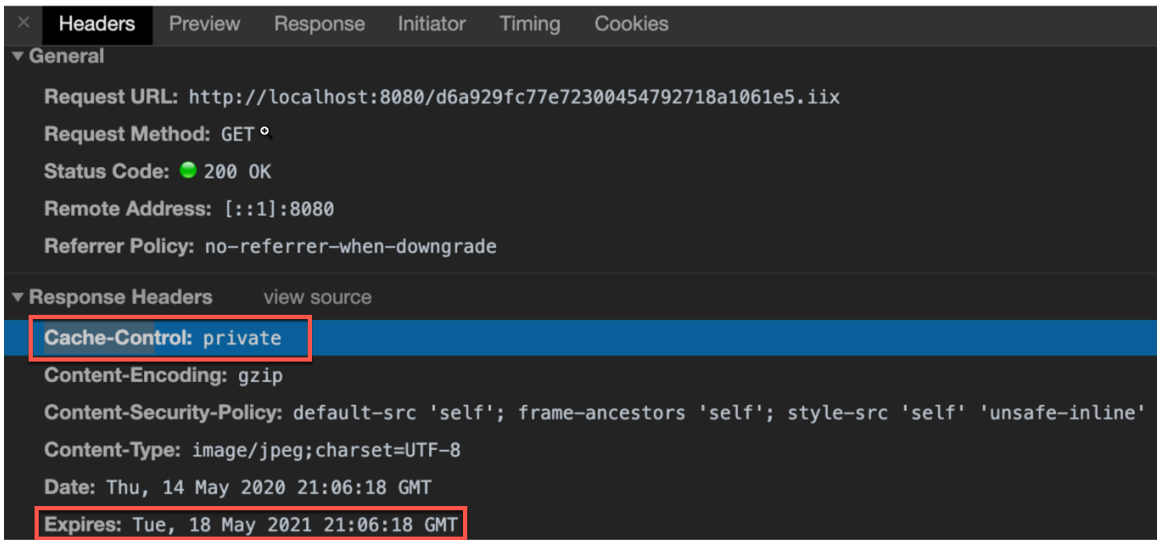

Note: You can use the *glide.http.cache* property designate whether to enable or disable setting of the Cache-Control and Expires values in HTTP response headers. Its default is **true**, which allows you to set the:

- Cache-Control value default using the *glide.http.cache_control* property.
- Expires value default using the *glide.http.expire.days* property.

More information

Attribute	Description
Property name	*glide.http.cache_control*
Configuration type	System Properties (/sys_properties_list.do)
Category	Configuration

Attribute	Description
Purpose	To configure the Cache-Control HTTP response header value for static content.
Recommended value	private
Default value	private
Security risk rating	4.3
Functional impact	Sets the default Cache-Control value in an HTTP response header.
Security risk	(High) If you set this property to **public**, instances with CDN/proxies may cache static content and render without authentication. • **private** is a more appropriate setting for instances with CDN/proxy setup. • If an the instance does not have a CDN/proxy setup, either value should be fine.

To learn more about adding or creating a system property, see Add a system property.

Enable HTTP response headers configuration

Reduce the risk of cookie/session-related hijacking of web apps using a system property.

If *glide.http.headers_config.enabled* isn't set to **true**, then response header configurations defined in the HTTP Response Headers [sys_response_header] table aren't used. Security related HTTP response headers include Content Security Policy, which aids in XSS-related protections. For details on HTTP response headers, see HTTP Response Headers.

Verify that the property *glide.http.headers_config.enabled* is set to **true**.

More information

Attribute	Description
Configuration name	*glide.http.headers_config.enabled*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Fallback value	true
Category	Session management
Security risk	• Severity score: 5.5 • CVSS score: Medium

Attribute	Description
	• Security Risk: The security risks of missing, incorrect, or weak HTTP response headers may allow for XSS, CSRF, and cookie/session related hijacking of web apps.
Dependencies and prerequisites	None

Disable legacy JQuery UI usage

Avoid the introduction of unpatched vulnerabilities in the library by disabling legacy JQuery UI usage.

Prevent the use of older prepatched JQuery UI versions, which introduce unpatched vulnerabilities in the library. Using old versions can potentially lead to security risks arising from attacks on vulnerabilities discovered in outdated JQuery UI library versions.

Confirm that the **glide.jquery_ui.legacy** system property is set to **false** to prevent older prepatched JQuery UI versions from being used. This system property is a failsafe in case organizations depend on the non-patched versions to run their custom implementations.

More information

Attribute	Description
Configuration name	*glide.jquery_ui.legacy*
Configuration type	System Properties (/sys_properties_list.do)
Data type	String
Recommended value	false
Default value	false
Fallback value	true
Category	Configuration
Security risk	• Severity score: 3.9 • CVSS score: Low • Preventing older prepatched JQuery UI versions from being used can introduce unpatched vulnerabilities in the library.
Functional impact	This system property is a failsafe in case organizations depend on the non-patched versions to run their custom implementations.
Dependencies and prerequisites	None

To learn more about adding or creating a system property, see Add a system property.

Disable locked form elements debugging

Here's the description for *glide.security.explain.write.locks*.

Set *glide.security.explain.write.locks* to the recommended value of **false** to prevent the display of explanation of locked form elements. Set the value to **true** to display the explanation of locked form elements.

More information

Attribute	Description
Property name	*glide.security.explain.write.locks*
Configuration type	System Properties (/sys_properties_list.do)
Category	Configuration
Purpose	Limits display behavior of SecurityDebugger without dependence on other properties.
Recommended value	false
Default value	false
Configuration type	Boolean
Security risk	(Low) Will prevent the display of the explanation on locked form elements. This makes the application slightly more secure, as less information is being provided by the security debugger.
Security risk rating	3.5

To learn more about adding or creating a system property, see Add a system property.

Disable MultiSSO Debugging [Updated in Security Center 1.3 and 1.5]

The *glide.authenticate.multisso.debug* property controls debug logging for Multi-SSO.

More information

Attribute	Description
Property name	*glide.authenticate.multisso.debug*
Configuration type	System Properties (/sys_properties_list.do)
Category	Configuration
Purpose	Disables Multi-SSO debug.
Recommended value	false
Default value	false
Configuration type	Boolean
Security risk	(High) Set the property to the recommended value of "False", otherwise, MultiSSO debug is enabled which could lead to unintended sensitive information leak.
Security risk rating	4.0
References	Multi-Provider SSO properties, tables, and scripts

To learn more about adding or creating a system property, see Add a system property.

Disallow target cloning [New in Security Center 1.3]

Configure the *glide.db.clone.allow_clone_target* property to prevent your instance from being used as a clone target.

Protect your instance from being used as a clone target by setting the **glide.db.clone.allow_clone_target** system property to **false**. A system clone copies everything in a database from a source instance to the target instance. This is a security risk because the instance database on the target instance is overwritten in the cloning process, leading to data loss and lack of data integrity.

Set the **glide.db.clone.allow_clone_target** system property to **false** on production instances to disallow your instance from being selected as a clone target.

More information

Attribute	Description
Configuration name	*glide.db.clone.allow_clone_target*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	true
Category	Configuration
Security risk	• Severity score: 4.4 • CVSS score: Medium • Security risk details: Not setting this property to the recommended value of false enables the instance to be used as a clone target. This is a security risk because the instance database can be overwritten in the cloning process.
Dependencies and prerequisites	None
References	• • Register target instance (legacy)
Functional impact	This property provides an additional safeguard which prevents a production instance from being cloned to. The default value is false for production instances and true for sub production instances such as dev or qa. To enable an instance to be used as a clone target set this property to true.

Disable soap fault stack trace display

Manage how stack traces are displayed in your instance.

Use the *glide.soapfault.display_stack_trace*property to manage stack traces in your instance. If this property is configured to **false**, then sensitive information could be leaked. If it is set to **true**, then no stack trace will display.

More information

Attribute	Description
Configuration name	*glide.soapfault.display_stack_trace*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	false
Category	Configuration
Security risk	• Severity score: 4.3 • CVSS score: Medium • Security risk details: Setting this property to **false** could expose sensitive information from the stack trace.
Dependencies and prerequisites	None

Restrict performance monitoring access [Updated in Security Center 1.3]

Use the *glide.security.diag_txns_acl* property to control stats.do, threads.do, thread_pool_stats, and replication.do access from an unauthenticated connection.

When you set this property to **true**, the *glide.security.diag_txns_acl* property only allows access to the following by the administrator account:

- https://<instancename>.service-now.com/stats.do
- https://<instancename>.service-now.com/threads.do
- https://<instancename>.service-now.com/replication.do
- https://<instancename>.service-now.com/thread_pool_stats.do

Without enabling this setting, it is still possible to access these resources from an unauthenticated connection.

More information

Attribute	Description
Property name	*glide.security.diag_txns_acl*
Configuration type	System Properties (/sys_properties_list.do)
Category	Configuration
Purpose	Restrict the access to configuration pages to administrator account only

Attribute	Description
Recommended value	true
Default value	true
Security risk rating	5.3
Functional impact	This remediation enforces only administrator account to get access to the application sensitive data for logging and troubleshooting purposes.
Security risk	(Moderate) Sensitive data such as server details, threads, and processes executed on the server should never be visible or accessible to the end user without appropriate privileges.

To learn more about adding or creating a system property, see Add a system property.

Enable updated version of MultiSSO plugin [Updated in Security Center 1.3 and 1.5]

Verify that you're using v2 of the MultiSSO plugin and that it's set to true to reduce security vulnerabilities.

If the Multi SSO plugin is enabled on an instance, reduce security vulnerabilities by confirming that the v2 version is enabled. The latest version enhances security and has more features, such as Assertion encryption support, and IDP-initiated Single Logout (SLO). If the latest version is not enabled, the new security features cannot be used and the instance is at risk of using an plugin which is deprecated.

Follow the steps in KB0756504 to upgrade to the latest version. This process includes checking for and migrating any customization-related changes, then upgrading the version. When complete, the **glide.authenticate.multissov2_feature.enabled** system property is automatically set **true**.

More information

Attribute	Description
Configuration name	glide.authenticate.multissov2_feature.enabled
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Configuration
Security risk	If the latest version is not enabled, the new security features cannot be used and the instance is at risk of using an plugin which is deprecated
Dependencies and prerequisites	None
References	https://support.servicenow.com/kb?id=kb_article_view&sysparm_article=KB0756504

Enforce secure referrer policy [New in Security Center 1.3]

Use the *com.glide.security.referrerpolicy* property to ensure that the Referrer-Policy HTTP header sends the appropriate level of data to each ServiceNow® page to prevent data leaks.

When the *com.glide.security.referrerpolicy* property is set to default, it ensures that the Referrer-Policy HTTP header is managed with the appropriate level of information sent, specifically tailored for the ServiceNow AI Platform® request page. This prevents unauthorized data leaks that could be accessible from other parts of the full URL, such as the path and query string.

More information

Attribute	Description
Configuration name	*com.glide.security.referrerpolicy*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	default
Default value	default
Category	Configuration
Security risk	• Severity score: 4.3 • CVSS score: Medium • Security risk details: Ensure that the *com.glide.security.referrerpolicy* property is set to default to prevent leaks of unauthorized data.
Dependencies and prerequisites	None
References	Referrer-Policy
Functional impact	This property controls how much information is sent via the referrer header when a request is sent from a page: • default: Instance will take care of the referrer headers • same-origin: Send full referrer URL within the instance/same domain and no referrer to outside origin • origin: Send only the origin as a referrer inside and outside the origin • origin-when-cross-origin: Send full referrer URL within the instance/same domain and only the origin outside the origin

Ensure minimum private key size

Use a system property to determine the minimum size of the private key used for Certificate Signing Request (CSR) generation with the Certificate Inventory Management application.

The *sn_disco_certmgmt.private_key_size* system property determines the minimum size of the private key used for CSR generation with the Certificate Inventory Management application. Valid choices are 512, 1024, 2048, or 4096.

Verify that this property is set to a value of 2048 or higher. Valid choices for this property are 512, 1024, 2048, or 4096. If the property doesn't exist in the System Properties [sys_properties] table, or the value is invalid, the value is 2048 by default.

More information

Attribute	Description
Configuration name	*sn_disco_certmgmt.private_key_size*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Integer, representing the size of the private key generated. Use values of 512, 1024, 2048, or 4096.
Recommended value	2048
Default value	2048
Fallback value	2048
Category	Communications
Security risk	• Severity score: 3.1 • CVSS score: Low • Security Risk: Using keys smaller than 2048 can result in future information disclosure in the event the key is brute forced. Use a valid value greater than or equal to 2048 to future proof the key for a longer time period.
Functional impact	Legacy systems and applications may not handle keys greater than or equal to 2048.
Dependencies and prerequisites	None

To learn more about adding or creating a system property, see Add a system property.

Implement the x-frame-options: SAMEORIGIN security header [Updated in Security Center 1.3]

Use the *glide.set_x_frame_options* property to set the X-Frame-Options response header to SAMEORIGIN for all UI pages.

Use the X-Frame-Options HTTP response header to indicate whether browser should be allowed to render a page in a `<frame>` or `<iframe>`. Sites can use this function to avoid clickjacking attacks by ensuring that their content is not embedded into other sites. An attacker could embed your page into their own page and make your page elements perform maliciously. The end user may think the page is legitimate because it resembles your page. The end user may click on elements like usual only to have malicious scripts or elements run.

More information

Attribute	Description
Property name	*glide.set_x_frame_options*
Configuration type	System Properties (/sys_properties_list.do)
Category	Configuration
Purpose	To mitigate against ClickJacking attacks.
Recommended value	true
Default value	true
Security risk rating	5.9
Functional impact	This remediation enforces the restriction for rendering a ServiceNow AI Platform application in a third-party application in the form of an iFrame. If you have such an integration, the application wouldn't render in the customized third-party app.
Security risk	(Medium) The Same Origin policy enables you to restrict a domain from retrieving a script or a resource from another domains. All modern browsers support this functionality. The policy validates the connection based on protocol, port, and host. CORS (Cross Origin Request) is a modification to Same Origin Policy that enables access to resources/scripts from another domain when explicitly stated as a part of header value. • In this case, the X-Frame-Options header controls whether the ServiceNow AI Platform application can be rendered on the third-party website. • It reduces the sensitive exposure, because the property value, when set to **SAMEORIGIN** doesn't enable the rendering to happen.
References	Available system properties Configure iFrames

To learn more about adding or creating a system property, see Add a system property.

Require write access to access service catalog add item page [New in Security Center 1.3]

Use the *glide.sc.request.add_item_write_access* property to prevent unauthorized operations from being performed on catalog items.

More information

Attribute	Description
Configuration name	*glide.sc.request.add_item_write_access*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Configuration
Security risk	• Severity score: 4.3 • CVSS score: Medium • Security risk details: When the *glide.sc.request.add_item_write_access* property is set to false, any logged in user can access the Add Catalog Item UI page. This could result in unauthorized operations performed on catalog items. To remediate this security risk, set this property to true.
Dependencies and prerequisites	None
Functional impact	When the property is true, the user must have write access to the record in the context of the UI page.

Set Xframe options to prevent embedding third-party websites [Updated in Security Center 1.3]

Configure this property to prevent the content of a web-application from being embedded in a third-party site.

If *com.glide.cs.embed.xframe_options* is not set to the recommended value of DENY or SAMEORIGIN, then content of the web application could be embedded in a third-party site using an ALLOW-FROM uri. Allowing untrusted third-party sites could enable attacks such as clickjacking.

More information

Attribute	Description
Configuration name	*com.glide.cs.embed.xframe_options*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	sameorigin
Default value	sameorigin
Category	Configuration

Attribute	Description
Security risk	• Severity score: 3.1 • CVSS score: Low • Security risk details: Not setting this property to the recommended value could enable the content of a web application to be embedded in a third-party site enabling attacks such as click-jacking.
Dependencies and prerequisites	None

Data protection

The data protection category addresses the elements of confidentiality, integrity and availability (CIA) of data.

The CIA components are:

- Confidentiality: Data is protected from unauthorized access in transit and at rest.
- Integrity: Data is protected from unauthorized creation, deletion or change.
- Availability: Data is accessible when required.

Remove remember me

Use the *glide.ui.forgetme* property to remove the **Remember Me** check box from the login page to prevent login information from being cached.

More information

Attribute	Description
Property name	*glide.ui.forgetme*
Configuration type	System Properties (/sys_properties_list.do)
Category	Data protection
Purpose	To ensure that no authentication information is cached.
Security risk rating	3.5
Recommended value	true
Default value	true
Functional impact	This remediation would change the user experience by automatically logging them out of the instance when their session expires. The session expiration would solely depend on the value set in the system property as detailed in Managing user sessions.
Security risk	(Low) When you select the **Remember me** check box at login, an extra cookie is stored on the user's computer.

Attribute	Description
	• Its purpose is to automatically re-establish the session for the subsequent visits of the logged-in user. • It poses a security risk as it allows the user session to be active until they deliberately log out. The likelihood of an attack for this scenario increases when the end user has left the browser unattended, or if it is compromised from a different attack.
References	Remove the Remember me check box

To learn more about adding or creating a system property, see Add a system property

Require clearing pasteboard when backgrounding mobile application [New in Security Center 1.3 and updated in 1.5]

The *glide.sg.clear_pasteboard_when_backgrounded* property controls if text copied from ServiceNow mobile app is kept in the clipboard and pasteboard after the app is in background mode. If it is not set to the recommended value of true, then sensitive information may be disclosed to the Android or iOS clipboard where it can be exposed to other applications on the device.

More information

Attribute	Description
Configuration name	*glide.sg.clear_pasteboard_when_backgrounded*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Data protection
Security risk	• Severity score: 3.5 • CVSS score: Low • Security risk details: If this property is not set to the recommended value of true, then sensitive information may be disclosed to the Android or iOS clipboard where it can be exposed to other applications on the device.
Dependencies and prerequisites	None
Functional impact	This property clears the copy and paste clipboard when the ServiceNow app enters the background.

Restrict HR case updates from personal emails [New in Security Center 1.3 and updated in 1.5]

Use the *sn_hr_core.restrict_guest_email* property to control whether a user can respond back to a HR case with their personal email.

When the *sn_hr_core.restrict_guest_email* property is not set to true, a user can send an email from a personal account referencing the HR case to be included in the work notes. This could result in minor confidentiality or integrity issues if the personal email is compromised or communicating insecurely. An admin may want to restrict the ability of users to respond to HR cases from their personal email, since they cannot be confident of the user accessing the personal email account.

More information

Attribute	Description
Configuration name	*sn_hr_core.restrict_guest_email*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Data protection
Security risk	• Severity score: 3.5 • CVSS score: Low • Security risk details: Not having this property set to true could result in minor confidentiality or integrity issues if the personal email is compromised or communicating insecurely.
Dependencies and prerequisites	None
Functional impact	This property controls whether or not a reply from a personal email address will update an HR Case. Set to true, any reply from personal email will be added to the case notes. If false, the case and notes will not be updated.

Restrict oauth parameters to POST body [New in Security Center 1.3]

Use the *glide.oauth.allow.parameters.in.post.body.only* property to control the inbound OAuth authentication's acceptance of access tokens. Access tokens are sensitive and should only be accepted when located within a POST request body.

More information

Attribute	Description
Configuration name	*glide.oauth.allow.parameters.in.post.body.only*
Configuration type	System Properties (/sys_properties_list.do)

Attribute	Description
Data type	Boolean
Recommended value	true
Default value	true
Category	Data protection
Security risk	• Severity score: 4.2 • CVSS score: Medium • Security risk details: If *glide.oauth.allow.parameters.in.post.body.only* isn't set to the recommended value of **true**, access tokens could be present in the GET request parameter. These access tokens could linger in client and infrastructure logs and potentially lead to account takeover if those logs are leaked.
Dependencies and prerequisites	None
References	• OAuth 2.0 • Manage OAuth tokens
Functional impact	Ensures that oauth_token.do processor accepts only POST body parameters as input for all supported grant types.

Error handling and logging

The error handling and logging category addresses the quality and verbosity of logged information exposed to stakeholders.

This includes ensuring logs and error messages do not collect sensitive information, correctly protect data according to classification and have an appropriate lifetime. Additionally, this category relates to appropriate error handling and not revealing sensitive errors to end users, such as verbose stack traces for unhandled exceptions with security implications.

Disable logger for low privilege users in script sandbox [Updated in Security Center 1.3]

Manage Glide System's ability to log scripts being executed in the sandbox environment.

Use the *glide.security.sandbox_no_logging* property to control Glide System's ability to log scripts being executed in the sandbox environment. If *glide.security.sandbox_no_logging* is set to **false**, logging is available for lower-privileged users using sandboxed scripts. This is a potential security vulnerability because low privileged users can inject logs allowing a bad actor to potentially obfuscate an attack. Configure the property to **true** to prevent lower-privileged users that are using a sandboxed script from having logging functionality.

More information

Attribute	Description
Configuration name	*glide.security.sandbox_no_logging*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Error handling and logging
Security risk	• Severity score: 2.2 • CVSS score: Low • Security risk details: Setting this property to **false** enables logging for lower-privileged users which could allow a bad actor to obfuscate an attack.
Dependencies and prerequisites	None

Disable secure cookie debugging

Manage the log messages related to cookies in your instance.

Use the *glide.secure_cookie.debug* property to manage your log messages related to cookies. If this property is set to **false**, no log messages are displayed. If it is set to **true**, messages in the `SecureUserCookie` and `Cookie` classes are logged. This could lead to sensitive information being exposed in your instance.

More information

Attribute	Description
Configuration name	*glide.secure_cookie.debug*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	false
Category	Error handling and logging
Security risk	• Severity score: 4.2 • CVSS score: Medium • Security risk details: Configuring this property to **true** could lead to sensitive information being exposed.
Dependencies and prerequisites	None

Disable SQL Error Messages [Updated in Security Center 1.3 and 1.5]

Use the *glide.db.loguser* property to disable SQL error messages from rendering in a browser.

If *glide.db.loguser* is not set to the recommended value of false, then sensitive server-side error messages could be displayed to end-users. Error messages can include stack traces and information about the structure of the database that could provide an attacker the knowledge needed to perform successful SQL Injection should the preconditions exist. As defense in depth, these error messages should not be displayed to the end user.

More information

Attribute	Description
Property name	*glide.db.loguser*
Configuration type	System Properties (/sys_properties_list.do)
Category	Error handling and logging
Purpose	To disable SQL error messages from displaying within the browser.
Type	Boolean
Recommended value	false
Default value	true
Security risk rating	3.1
Functional impact	This remediation disables rendering of SQL error messages. There is no impact to any functionality.
Security risk	(Medium) No sensitive SQL information that could help an attacker should appear as a part of error message on a web page.

To learn more about adding or creating a system property, see Add a system property.

Enable MID audit log [New in Security Center 1.3 and updated in 1.5]

The MID Server command audit log records details such as the command name, command hash, name of credential used, and execution status.

The MID Server command audit log tracks details such as the command name, command hash, name of credential used and execution status. When enabled, users with the *agent_security_admin* role can view these logs in the MID Server Command Audit Logs [ecc_agent_command_audit_log] table. Navigate to **All** > **MID Server** > **Audit Logs** > **Command Audit Logs** to see this table.

Set *mid.log.command_audit.enable* property to **true** in the MID Server Properties [ecc_agent_property] table for each MID Server to turn on auditing for commands run by the MID Server.

For more details on setting this property, see MID Server command audit log.

For information about MID Servers and how they work, see MID Server.

More information

Attribute	Description
Configuration name	*mid.log.command_audit.enable*
Configuration type	MID Server Property [ecc_agent_property] record
Data type	Boolean
Recommended value	true
Default value	false
Category	Error handling and logging
Security risk	• Severity score: 2.2 • CVSS score: Low • Security risk details: In the event of security investigation, this table can be used by incident response teams to audit the commands run on the MID server. Without this log, there might not be sufficient details to respond to situations such as unauthorized account use.
Dependencies and prerequisites	None

Enable protected tables plugin [New in Security Center 1.3]

Use the *com.glide.security.protected_table.enabled* property to prevent higher privilege users from tampering with log tables.

When the *com.glide.security.protected_table.enabled* property is set to **true**, the protected tables plugin will be used to prevent higher privilege users on an instance from tampering with log tables. The following log tables have special protections when this property is set to **true**:

- syslog (No DB Override)
- syslog_transaction
- sys_outbound_http_log
- sysevent
- sys_audit
- sys_push_notification
- protected_table_configuration (No DB Override)

The integrity of logs is important for determining malicious activity on an instance by a customer admin.

More information

Attribute	Description
Configuration name	*com.glide.security.protected_table.enabled*

Attribute	Description
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Error handling and logging
Security risk	• Severity score: 4.5 • CVSS score: Medium • Security risk details: Not setting *com.glide.security.protected_table.enabled* to the recommended value of true enables higher privilege users on an instance to tamper with log tables.
Dependencies and prerequisites	None
References	System logs

Log all outbound http request fields [Removed in Security Center v1.3.2]

Configure the *glide.outbound_http.security.log.allow.all.fields* property to false to prevent sensitive Outbound HTTP fields from being logged in plain text.

If this property is not set to the recommended value of **false**, sensitive Outbound HTTP fields might be logged in plaintext. This can decrease the security posture of your enterprise network because outbound requests with sensitive data and credentials can be logged in plaintext which is unencrypted, and can be viewed by lower-privileged users.

More information

Attribute	Description
Configuration name	*glide.outbound_http.security.log.allow.all.fields*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	false
Category	Error handling and logging
Security risk	• Severity score: 6.8 • CVSS score: Medium • Security risk details: Not setting this property to **false** increases the chance of sensitive outbound fields being logged in plaintext which is a security risk.
Dependencies and prerequisites	None

Log html sanitization [Removed in Security Center 2.0]

Configure *glide.html_sanitize.discarded_log.enable* property to determine if HTML sanitization events will be logged in your instance.

If this property is not set to the recommended value of **true**, HTML sanitization events are not logged in the *sys_log* table. Lack of logging could negatively impact your instance automated security detection and investigation capabilities.

More information

Attribute	Description
Configuration name	*glide.html_sanitize.discarded_log.enable*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Error handling and logging
Security risk	• Severity score: 2.4 • CVSS score: Low • Security risk details: If this property is not set to the recommended value of **true**, HTML sanitization events are not logged in the *sys_log* table which could impact automated security detection and investigation capabilities.
Dependencies and prerequisites	None
References	Enabling HTML sanitizer

Log session audit events [New in Security Center 1.3 and updated in 1.5]

Set the *glide.authenticate.session_access.log_audit_event* property to **true**, so that session audit events will be created in the *sys_session_access_audit* table.

When the Glide Property *glide.authenticate.session_access.log_audit_event* is set to true, session audit events will be created in the sys_session_access_audit table. It is best practice to log information about who accessed a session to assist in malicious actor investigations. Information logged will include user, session ID (non-sensitive), IP address, roles, and policies.

Note: The *glide.authenticate.session_access.log_audit_event* system property is specific to Zero trust access. For more information, see Zero Trust Access (ZTA).

More information

Attribute	Description
Configuration name	*glide.authenticate.session_access.log_audit_event*

Attribute	Description
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Error handling and logging
Security risk	• Severity score: 6.3 • CVSS score: Medium • Security risk details: Not setting this property to the recommended value of true prevents events from being logged. This could prevent you from finding bad actors in the event of a cyber attack.
Dependencies and prerequisites	None

Log user impersonation [Updated in Security Center 1.3 and 2.0]

Configure *glide.sys.log_impersonation* to control if user-impersonating events are logged in your instance.

If this property is not set to the recommended value of **true**, user-impersonating events are no longer logged. The absence of logging could impact the automated security detection and investigation capabilities on your instance. Ensure the property *glide.sys.log_impersonation* exists and is set to true. If the property does not appear in the sys_properties table, add a new record.

More information

Attribute	Description
Configuration name	*glide.sys.log_impersonation*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Error handling and logging
Security risk	• Severity score: 6.4 • CVSS score: Medium • Security risk details: If this property is not set to **true**, user-impersonating events are no longer logged which could impact your instance's security detection and investigation capabilities.
Dependencies and prerequisites	None

Attribute	Description
References	Impersonating users

Prevent verbose HTTP request logging

Help prevent access to sensitive information by reducing verbose HTTP request logging.

Control the level of logging for outbound HTTP requests to help prevent access to sensitive information, such as authorization headers or cookies. This information can be used like credentials to access the requested resource.

The logging level for these requests is controlled by the **glide.outbound_http_log.override** and **glide.outbound_http_log.override.level** properties. When **glide.outbound_http_log.override** is set to **true**, the log level for requests and responses is controlled by the **glide.outbound_http_log.override.level** property. If **glide.outbound_http_log.override.level** is set to **all** or **elevated**, then request and response headers are logged.

Set **glide.outbound_http_log.override** to **false** and **glide.outbound_http_log.override.level** to **basic**. If these properties don't appear in the System Properties [sys_properties] table, they are in a secure state by default.

More information

Attribute	Description
Technical configuration name	• glide.outbound_http_log.override • glide.outbound_http_log.override.level
Plugin applicability	None
Security risk	Outbound HTTP request headers logged with verbose settings can include sensitive information such as Authorization headers or cookies. This information can be used like credentials to access the requested resource. Users with access to the Outbound HTTP Logs [sys_outbound_http_log] table can view this information. The severity depends on what type of outbound requests are made.
Common Vulnerability Scoring System (CVSS) score	5
Common Vulnerability Scoring System (CVSS) rating	Medium
Functional impact	The **glide.outbound_http_log.override** system property enables you to override outbound http request log level. A value of **false** defaults the log level to basic. If **glide.outbound_http_log.override** is set to **true**, the level of logging is determined by the value of the **glide.outbound_http_log.override.level** property. This value can be **basic**, **elevated**, or **all**. All 3 are string/text based values. Any value other than these is interpreted as **basic**.

Attribute	Description
	For additional details, see Configure outbound logging.
Dependencies and prerequisites	None
Data type	• Boolean • String
Base system value	• false • <blank>
Fallback value	• false • <blank>
Recommended value	• false • basic

To learn more about adding or creating a system property, see Add a system property.

Turn off verbose SQL error messages for import processor [Updated in Security Center 1.3]

Configure this property to control whether verbose SQL error messages are displayed.

If the property is set to **false**, then a verbose SQL error message will display which can leak sensitive information. To avoid this, set the property to **true** to display a generic message.

More information

Attribute	Description
Configuration name	*glide.import.error_message.generic*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Error handling and logging
Security risk	• Severity score: 3.1 • CVSS score: Low • Security risk details: Setting the property to **false** enables verbose SQL messages which can lead to information leakage.

Attribute	Description
Dependencies and prerequisites	None

File and resources

The file and resources category ensures applications handle untrusted file data securely and store untrusted data from untrusted sources with limited permissions in an appropriate location.

This includes controls such as avoiding denial of service through large or unexpected file types, validating file type and preventing against path traversal.

Disallow infected file download [Updated in Security Center 1.5 and 2.0]

Control whether users can download non-scanned attachments if the antivirus service is down or unreachable.

When the *com.glide.snap.infected_download_allowed* property is set to true, users can still download non-scanned attachments in the event that the antivirus service is down or unreachable. This situation potentially exposes users to the risk of downloading a malicious file, thereby risking the infection of their desktop, especially if there is no other endpoint protection installed on the device.

Ensure the property *com.glide.snap.infected_download_allowed* is set to false.

More information

Attribute	Description
Configuration name	*com.glide.snap.infected_download_allowed*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	false
Category	File and resources
Security risk	• Severity score: 6.7 • CVSS score: Medium • Security risk details: Not setting this property to the recommended value of false could expose your instance to downloading malicious files.
Dependencies and prerequisites	None

Enable email spam scoring and filtering [Updated in Security Center 1.3]

Install the Email Filter (*com.glide.email_filter*) plugin to install email filtering within the instance. This filtering identifies existing headers, which enables you to decide what to do with the email based on the associated header. Alternatively, set *com.glide.email_filter* to false.

Every message sent through ServiceNow AI Platform email servers is assessed for the likelihood of being spam.

Note: If an instance uses a private email server, this topic is not applicable. For more information, see Email Spam Scoring and Filtering.

Prerequisites

Before setting this property:

Set the *glide.email.read.active* property to true. To learn more, see Enable using your own POP3 server.

More information

Attribute	Description
Plugin Name	com.glide.email_filter, glide.email.read.active
Configuration type	System Definition > Plugins
Category	File and resources
Purpose	To enforce filtering to avoid spamming of emails.
Recommended value	Either: • The **glide.email.read.active** property to **false** • The **glide.email.read.active** property to **true** and activate the **Email Filter** (com.glide.email_filter) plugin. Active
Default value	None. This is a plugin, not a Glide property so there is no default value.
Security risk rating	8.1
Functional impact	Email is never filtered, blocked, or quarantined from the instance as part of spam scoring. It is only scored and then sent on to the instance. All filtering is done within the instance with the Email Filters plugin.
Security risk	(Moderate) Email filters enable administrators to use a condition builder or conditional script to specify when to ignore malicious incoming emails from known/unknown sender.
References	Email filters https://support.servicenow.com/kb_view.do?sysparm_article=KB0549426

To learn more about activating a plugin, see Activate a plugin.

Enable antivirus scan

The *com.glide.snap.enable_scan* property activates the antivirus scan functionality.

Set *com.glide.snap.enable_scan* to the recommended value of **true** to enable antivirus scanning.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*com.glide.snap.enable_scan*
Configuration type	System Properties (/sys_properties_list.do)
Category	File and resources
Purpose	To enable or disable antivirus scanning on the specific instance.
Recommended value	true
Default value	true
Data type	Boolean
Security risk	(High) Antivirus scanning helps protect your instance against virus infections that can be introduced by file attachments to your system records, such as incidents, problems, and stories.
Security risk rating	7.7

Restrict downloadable files types in static content [Updated in Security Center 1.3]

Use the *glide.ui.strict_customer_uploaded_static_content* property to enable restrictions on the file types that can be downloaded when they have been uploaded using the Upload File functionality.

You use this property with the *glide.ui.strict_customer_uploaded_content_types* property, which creates a comma-delimited list of restricted downloadable file types.

Warning: The value for this property is a no DB override. It can't be altered or overridden.

More information

Attribute	Description
Property name	*glide.ui.strict_customer_uploaded_static_content*
Configuration type	System Properties (/sys_properties_list.do)
Category	File and resources
Purpose	To ensure that safe file types are permitted to be downloaded from the application.
Recommended value	true
Default value	true
Security risk rating	3.1

Attribute	Description
Functional impact	This remediation enforces restriction of file downloads based on the values specified in the *glide.ui.strict_customer_uploaded_content_types* property.
Security risk	(Low) File download restrictions should be applied to any untrusted user input sources.

Limit attachment size in training and prediction flows for GraphQL endpoints [New in Security Center 1.3 and updated in 1.5]

The *glide.platform_ml_di.max_attachment_size_graphql* property controls the maximum allowed size limit for returning attachments in GraphQL endpoints of training or prediction flows.

More information

Attribute	Description
Configuration name	*glide.platform_ml_di.max_attachment_size_graphql*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	5,242,880
Default value	5,242,880
Category	File and resources
Security risk	• Severity score: 4.3 • CVSS score: Medium • Security risk details: If this property is not set to the recommended value of 5,242,880 or less, then returning large files could cause a denial of service (DoS).
Dependencies and prerequisites	None

Limit attachment size in training and prediction flows [New in Security Center 1.3 and updated in 1.5]

The *glide.platform_ml_di.max_attachment_size* property controls the maximum allowed size limit for returning attachments in training and prediction flows.

More information

Attribute	Description
Configuration name	*glide.platform_ml_di.max_attachment_size*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer

Attribute	Description
Recommended value	4,000,000
Default value	4,000,000
Category	File and resources
Security risk	• Severity score: 4.3 • CVSS score: Medium • Security risk details: If *glide.platform_ml_di.max_attachment_size* is not set to the recommended value of 4,000,000 or less, then returning large files could cause a denial of service (DoS) attack.
Dependencies and prerequisites	None

Limit HTTP response body size [New in Security Center 1.3 and updated in 1.5]

Configure the *glide.http.response.get_body.limit.enabled* and *glide.http.response.get_body.limit* properties to protect your instance against *OutOfMemoryExceptions*.

Prevent `OutOfMemoryExceptions` that can result from a request response body being too large using the **glide.http.response.get_body.limit.enabled** and **glide.http.response.get_body.limit** system properties. These exceptions can cause denial of service (DOS) attacks as well as other issues that may aid attackers in compromising an instance. Not setting these properties to the recommended values could make your instance vulnerable to OutOfMemoryExceptions and denial of service attacks.

To protect your instance against these security vulnerabilities:

- Set the **glide.http.response.get_body.limit.enabled** system property to **true**.
- Ensure that the **glide.http.response.get_body.limit** system property set to no more than 524,288,000 megabytes (500 MB).

More information

Attribute	Description
Configuration name	• **glide.http.response.get_body.limit.enabled** • **glide.http.response.get_body.limit**
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	File and resources

Attribute	Description
Security risk	• Severity score: 3.1 • Security risk details: Not setting these properties to the recommended values could make your instance vulnerable to *OutOfMemoryExceptions* and denial of service attacks.
Dependencies and prerequisites	None
Functional impact	This property reduces the chances of an *OutOfMemoryException* due to a customer accidentally loading a large file into memory.

Limit maximum number of attachments in email

Configure the number of inbound email attachments allowed per Email [sys_email] record on your instance.

Use the *glide.email.inbound.max_attachment_count* system property to control the maximum number of inbound email attachments allowed per Email [sys_email] record on your instance.

Set the value of this property to **30** or less to help avoid instance performance degradation.

More information

Attribute	Description
Configuration name	*glide.email.inbound.max_attachment_count*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	30 or less
Default value	30
Fallback value	50
Category	File and resources
Security risk	• Severity score: 5.3 • CVSS score: Medium • Security risk details: Setting the property value to greater than 30 can lead to degradation issues.
Dependencies and prerequisites	None
References	Email properties

Maximum allowed attachment size [Updated in Security Center 1.3]

Configure the *com.glide.attachment.max_size* property to control the maximum size (in megabytes) permitted for an uploaded attachment.

This property controls the maximum size (in megabytes) of an uploaded attachment. If this property is not set to the recommended value of **1024** (1 gigabyte) or less, the platform can accept large files that could fill up storage and lead to a Denial of Service (DoS) attack.

More information

Attribute	Description
Configuration name	*com.glide.attachment.max_size*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	1024
Default value	1024
Category	File and resources
Security risk	• Severity score: 6.5 • CVSS score: Medium • Security risk details: If this property is not set to **1024** (1 gigabyte) or less, the platform could accept large files which may lead to a DoS attack.
Dependencies and prerequisites	None
References	Attachment limit properties

Set Allowed MIME Child Types [New in Security Center 2.0]

Learn how to configure the *glide.security.mime.type.allowed_child_types* property to a secure setting so that file types will not pass the Multipurpose Internet Mail Extensions (MIME) type checking. This reduces the risk of remote code execution on an uploaded file.

The *glide.security.mime.type.allowed_child_types* property defines the MIME file types that may have a file extension not matching the data within an uploaded file. This allows such file types to bypass MIME type checking. The property accepts a comma-separated list of file type pairs, such as *application/zip=application/java-archive*. In this example, if the property is set to such a value, files with a `.zip` extension that are technically `.jar` files are allowed to pass MIME type checking despite the inconsistency. If not set properly, this bypass can lead to remote code execution of an uploaded file. Therefore, it should always be set to an empty string ("") unless a valid use case arises. For instance, if a certain MIME type must be allowed under a different file extension and is valid as per the Tika configuration, then those key-value pairs will be updated as part of this property value.

More information

Attribute	Description
Configuration name	*glide.security.mime.type.allowed_child_types*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string

Attribute	Description
Recommended value	""
Default value	""
Category	File and resources
Security risk	• Severity score: 4.6 • CVSS score: Medium • Security risk details: Not setting this property to the secure value could cause files with incorrect configurations to lead to remote code execution of an uploaded file.
Dependencies and prerequisites	Yes, when *glide.security.mime.type.detection.allow_child_typ* is set to true, the values of this property will be used to validate against the configured list of allowed MIME child types.
Functional impact	To support MIME types whose file extensions do not match the content of the files but are valid according to the Tika sub-type configurations in *tika-mimetypes.xml*.

Validate file mime type in AttachmentCreator soap web service [New in Security Center 1.3 and updated in 1.5]

The *glide.attachment.enforce_security_validation* property determines whether Multipurpose internet Mail Extensions (MIME) files undergo validation.

Ensure that MIME-types are validated for attachments to prevent dangerous files from being uploaded on your instance using wrong file extensions.

Set the **glide.attachment.enforce_security_validation** system property to **true**. When set to **true**, files are uploaded with the correct file type extension.

More information

Attribute	Description
Configuration name	*glide.attachment.enforce_security_validation*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	File and resources
Security risk	• Severity score: 6.7 • CVSS score: Medium

Attribute	Description
	• Security risk details: If the property is set to false, there's no validation for MIME files during uploads. This could enable malicious files to be disguised by changing their file extension.
Dependencies and prerequisites	None
References	https://developer.mozilla.org/en-US/docs/Web/HTTP/Basics_of_HTTP/MIME_types/Common_types
Functional impact	Set this hardening setting to true to run mime-type and file extension validations on uploaded file attachments. No validations are run if this property is set to false. This property is set to true by default.

Malicious code

The Malicious Code category ensures that best efforts are made to confirm that your code is free of vulnerabilities and unwanted functionality.

This includes ensuring secure and proper handling for malicious activity, no time based attacks, no outbound communications to untrusted destinations, and that no unauthorized or attacker-controlled code is included. This category includes audit or third party libraries from the application codebase.

Block rooted or jailbroken mobile devices

Secure your instance by preventing unauthorized access from jailbroken devices.

Use the *glide.sg.allow_rooted_jailbroken_device* property to secure your instance from unauthorized access by jailbroken devices. If a user tries to authenticate into an instance using a mobile app while this property is set to **false**, they receive the following alert: `This device appears to be jailbroken and cannot be used to access this instance. Please contact your ServiceNow Administrator.` The app is frozen while the alert message is displayed, and the only way to dismiss this message is to select **Log out**. If this property is set to **true** users authenticate into an instance using a jailbroken device.

More information

Attribute	Description
Configuration name	*glide.sg.allow_rooted_jailbroken_device*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	false
Category	Malicious code
Security risk	• Severity score: 4.5 • CVSS score: Medium

Attribute	Description
	• Security risk details: The lack of security on jailbroken devices makes them a prime target by bad actors. Having unauthorized entities accessing corporate data can weaken the security posture of a corporate network.
Dependencies and prerequisites	None
References	Access control

Enable Code Signing for application configuration data and scripts [Removed in Security Center 1.3]

Manage Code Signing for application configuration data and scripts on your instance.

Code Signing can help improve security by validating sensitive application configuration data and scripts before they are used. Code Signing creates digital signatures for the data which later are checked to confirm the authenticity and integrity of the data. This verification prevents malicious data or scripts from being used on the instance which may lead to full compromise of the instance.

Enable Code Signing on your instance by following the steps in Configuring Code Signing. This results in the *com.snc.kmf.signature.validation.flag* system property being set to true.

⚠ **Warning:** This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Configuration name	*com.snc.kmf.signature.validation.flag*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Malicious code
Security risk	• Severity score: 6 • CVSS score: Medium • Security risk details: Setting this property to **true** enables Code Signing, which helps improve security by validating sensitive application configuration data and scripts.
Dependencies and prerequisites	None

Session management

This category looks at the security of the application state for a user. Sessions should be unique to each individual, unable to be guessed or shared, and invalidated after periods of inactivity or when not required. This includes factors such as cookie attributes for cookie-based sessions, session token generation, and storage and requirements for federated re-authentication.

Apply continuous authentication policies to mobile sessions

Reduce the risk of session hijacking by applying continuous authentication policies to mobile sessions.

Ensure that mobile users are using the high assurance session feature of the Zero Trust - Continuous Authentication plugin by setting the **glide.zta.high_assurance.mobile.session.allowed** system property to **false**. To bypass continuous authentication policies for mobile sessions, set this property to **true**.

Set the **glide.zta.high_assurance.mobile.session.allowed** system property to **false**. This action decreases the risk of session hijacking and persistent access if a session is compromised. If this property doesn't appear in the System Properties [sys_properties] table, the default value is **false**.

More information

Attribute	Description
Technical configuration name	glide.zta.high_assurance.mobile.session.allowed
Plugin applicability	Zero Trust - Continuous Authentication (com.snc.zero_trust_continuous_authentication)
Security risk	If this system property set to **true**, the high assurance session feature isn't applied to mobile sessions. Reauthentication doesn't happen as defined in policies for high role sessions. This increases the risk of session hijacking and persistent access if a session is compromised.
Common Vulnerability Scoring System (CVSS) score	3.9
Common Vulnerability Scoring System (CVSS) rating	Low
Functional impact	If set to **true**, the high assurance session feature doesn't apply to mobile sessions, which increases the risk of session hijacking and persistent access if a session is compromised.
Dependencies and prerequisites	None
Data type	Boolean
Base system value	false
Fallback value	false
Recommended value	false

To learn more about adding or creating a system property, see Add a system property.

Minimize absolute session timeout duration [Updated in Security Center 1.3]

Use the *glide.ui.user_cookie.max_life_span_in_days* property to set a maximum life span for user cookies created when users log in with the **Remember Me** checkbox selected. When the cookie expires, users who have selected the **Remember Me** checkbox are forced to reauthenticate into the instance.

It enables the user cookie to be valid for the duration of specified days, starting when the cookie was first issued. The default value is 30 days, and the maximum cap is at 365 days.

Note: To enforce a maximum session time for any active user sessions, see Managing user sessions.

More information

Attribute	Description
Property name	*glide.ui.user_cookie.max_life_span_in_days*
Configuration type	System Properties (/sys_properties_list.do)
Category	Session management
Purpose	To force users who have selected the **Remember Me** checkbox to reauthenticate after specific days.
Recommended value	Less than or equal to 30
Default value	30 days
Functional impact	This property enforces mandatory relogin by avoiding any sort of cookie rotation after a given timeframe.
Security risk rating	4.2
Security risk	(Medium) The user cookies being active for an indefinite amount of time is a security risk and should expire on a time-based configuration.
References	Available system properties Remember me

To learn more about adding or creating a system property, see Add a system property.

Define active session timeout exception roles [New in Security Center 1.3]

Use a system property to exempt roles from active session timeout limits.

Use the *glide.active.session.timeout.exception.roles* system property to exempt roles from an active session timeout limit. The active session timeout feature helps ensure that a hijacked session cannot be used indefinitely without providing authentication information. It is best practice to only consider an active session timeout limit exception for internal integration account roles.

Configure the *glide.active.session.timeout.exception.roles* property to roles which should be exempt from active session timeouts. This property value is a comma separated list of roles. The default value is *edge_encryption,mid_server,maint*.

More information

Attribute	Description
Configuration name	*glide.active.session.timeout.exception.roles*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	edge_encryption,mid_server,maint
Default value	edge_encryption,mid_server,maint
Fallback value	edge_encryption,mid_server,maint
Category	Session management
Security risk	• Severity score: 6.4 • CVSS score: Medium • Consider an active session timeout limit exception only for internal integration account roles. If a user is a victim of a session hijacking attempt, and has a role with an exception, attackers using that session can continue to authenticate to that session indefinitely. This may increase the impact of a security incident by enabling an attacker more time to make use of a hijacked account.
Dependencies and prerequisites	None

Enable UserCookie version 3.1 [Updated in Security Center 2.0]

Manage the version of UserCookie that is enabled on your instance to secure the storage of the secret key in the source code.

UserCookie v3 is generated only when property *glide.ui.secure.cookies.use_kmf is disabled*. UserCookie v3 is not secure due to storing secret key for HMAC in source code and identical for all customers. That can support malicious actors to use this one secret key for attempts to hijacking user sessions. By setting the property *glide.ui.secure.cookies.use_kmf* to true UserCookie v3.1 will be used and secret key will be stored in security storage such as KMF.

More information

Attribute	Description
Configuration name	*glide.ui.secure.cookies.use_kmf*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Session management

Attribute	Description
Security risk	• Severity score: 7.1 • CVSS score: High • Security risk details: Setting this to **false** is a security vulnerability due to the secret key for hash-based message authentication codes (HMAC) being stored in the source code.
Dependencies and prerequisites	None

Enforce password reset on api requests [Updated in Security Center 1.5]

Manage how the password reset functionality operates on your instance.

When a user is marked for **Password needs reset**, they must provide a new password at the next authentication attempt. This property controls whether the password reset is mandatory before making API calls. If this property is not set to the recommended value of **true**, user accounts marked as **Password needs reset** can still perform operations by querying the table API through basic authentication. This security vulnerability could enable information leakage if an inactive account is compromised.

More information

Attribute	Description
Configuration name	*glide.authenticate.api.user.reset_password.mandato*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Session management
Security risk	• Severity score: 8.1 • CVSS score: High • Security risk details: Setting this property to **false** could lead to information leakage if an inactive account is compromised.
Dependencies and prerequisites	None

Enable HTTP Only Cookie Flag [Updated in Security Center 1.3]

Use the *glide.cookies.http_only* property to enable the HTTPOnly attribute for sensitive cookies.

Use the HTTPOnly attribute to prevent attacks, such as cross-site scripting, because it doesn't allow access to the cookie using a client-side script, such as JavaScript. It does not eliminate cross site scripting risks but does eliminate some exploitation vectors.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.cookies.http_only*
Configuration type	System Properties (/sys_properties_list.do)
Category	Session management
Purpose	To mitigate the risk of client-side script accessing the protected cookie.
Recommended value	true
Default value	true
Security risk rating	8
Functional impact	This remediation adds an extra HTTPOnly flag in on session cookies, thus protecting them from being stolen. • If you have custom functionality that requires JavaScript to access the user's cookie, it breaks that functionality. It should not be the case under normal circumstances. • The ServiceNow AI Platform handles session management and there shouldn't be a reason for a custom script to access the user's cookies.
Security risk	(Moderate) Session cookies in the application authenticate an end user and provide implicit access permissions on the application. That means there is a need to secure them from being stolen or exported. HTTP Only flags protect the session cookies from JavaScript injections or cross site scripting vulnerabilities stealing them.
References	Available system properties

To learn more about adding or creating a system property, see Add a system property.

Invalidate Session After OAuth Token Expiration [New in Security Center 2.0]

Configure the *glide.authenticate.oauth.post.token.expiration.cookie_auth.disabled* property to the secure value to prevent users from continuing to use a session via cookies after the OAuth token used to create the session expires.

When the *glide.authenticate.oauth.post.token.expiration.cookie_auth.disabled* property is not set to the secure value of true, a user may continue to use a session via cookies after the OAuth token used to create the session expires. This increases the risk of cookies being leaked and the session being hijacked by a malicious user to access unauthorized resources. Ensure that the glide property

glide.authenticate.oauth.post.token.expiration.cookie_auth.disabled is set to true. If the record does not exist in the sys_properties table, the default value is false.

More information

Attribute	Description
Configuration name	*glide.authenticate.oauth.post.token.expiration.coo*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	For zboot instances, the property is true. For the update instances, the property is false by default.
Category	Session management
Security risk	• Severity score: 5.4 • CVSS score: Medium • Security risk details: When this property is not set to the secure value of true, a user may continue to use a session even after the OAuth token used to create the session has expired, increasing the likelihood of the session being hijacked by a malicious user.
Dependencies and prerequisites	None
Functional impact	True: Cookie authentication is only honored until the OAuth access token expires; after the expiration, authentication is not honored. False: Cookie authentication is honored even after the OAuth access token expires.

Minimize concurrent interactive session quantity [Updated in Security Center 1.3]

Use this property with the Limit Concurrent Sessions plugin to control the number of active sessions that can be opened by a user.

Use the **Glide Authenticate Max Concurrent Interactive Sessions** property with the **Limit Concurrent Session** (*com.glide.limit.concurrent.sessions*) plugin to control the number of active sessions opened for a user. The recommended value is **1**, which reduces the number of sessions opened (a greater number increases the likelihood that a session could be hijacked).

More information

Attribute	Description
Configuration name	*glide.authenticate.max.concurrent.interactive.sess*
Configuration type	System Properties (/sys_properties_list.do)

Attribute	Description
Data type	integer
Recommended value	1
Default value	1
Category	Session management
Security risk	• Severity score: 3.7 • CVSS score: Low • Security risk details: Setting the default value of the property greater than 1 increases the chance of session hijacking.
Dependencies and prerequisites	The **Concurrent Interactive Sessions** (*com.glide.limit.concurrent.sessions*) plugin should be active.
References	Limit concurrent sessions

Limit concurrent sessions across all nodes [Updated in Security Center 1.3]

Use the *glide.authenticate.limit.concurrent.sessions.across.all.nodes* property with the Limit Concurrent Sessions plugin to manage the number of sessions tracked across all nodes.

When the Limit concurrent sessions plugin is active, the number of open sessions can be limited per user. Ensure that when this plugin is active that the (**Glide authenticate limit concurrent sessions across all nodes**) property is set to **true** so that the number of open sessions are tracked across all nodes instead of a single application node. If this property is set to **false**, multiple sessions can be open across multiple nodes, which increases the chances of session hijacking.

More information

Attribute	Description
Configuration name	*glide.authenticate.limit.concurrent.sessions.acros*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Session management
Security risk	• Severity score: 3.7 • CVSS score: Low • Security risk details: When using the **Limit Concurrent Sessions** plugin, setting this property to **false** enables

Attribute	Description
	multiple sessions across multiple nodes to be open which increases the chance of a security vulnerability like a session hijacking.
Dependencies and prerequisites	None
References	Limit concurrent sessions

Limit concurrent sessions plugin

Configure the com.glide.limit.concurrent.sessions plugin to reduce the chance of session hijacking on your instance.

When the Limit Concurrent Sessions (com.glide.limit.concurrent.sessions) plugin is not active, the *glide.authenticate.limit.concurrent.interactive.sessions* property is not set to **true**, or the *glide.authenticate.max.concurrent.interactive.sessions* property is set beyond an organizationally-defined threshold, then ServiceNow instance user accounts are not limited to a defined number of concurrent interactive sessions.

1. Navigate to **All** > **System Definition** > **Plugins**.
2. Find and select the Limit Concurrent Sessions plugin. The plugin ID is *com.glide.limit.concurrent.sessions*
3. On the System Plugin form, review the plugin details and then select the **Activate/Upgrade** related link.
4. Select **Activate**.
5. After the plugin has successfully activate, navigate to **All** > **System Properties** > **All Properties**.
6. Open the *glide.authenticate.limit.concurrent.interactive.sessions* system property, and set the value to **true**.
7. Open the *glide.authenticate.max.concurrent.interactive.sessions* system property, and set the maximum concurrent sessions. This value depends on the needs of your organization.

More information

Attribute	Description
Configuration name	• *com.glide.limit.concurrent.sessions* (plugin) • *glide.authenticate.limit.concurrent.interactive.* property) • *glide.authenticate.max.concurrent.interactive.se* property)
Configuration type	System Properties (/sys_properties_list.do)

Attribute	Description
Data type	• plugin • system property (Boolean) • system property (Integer) plugin
Recommended value	• *com.glide.limit.concurrent.sessions* is enabled • *glide.authenticate.limit.concurrent.interactive.* system property set to **true** • *glide.authenticate.max.concurrent.interactive.se* set to a numeric value depending on the needs of your organization.
Default value	None
Category	Session management
Security risk	• Severity score: 3.7 • CVSS score: Low • Security risk details: A greater number of open sessions means there are more sessions that could potentially be hijacked. Limiting the number of allowed sessions per user is helpful in limiting risks related to denial-of-service (DoS) attacks.
Dependencies and prerequisites	None

Limit guest's active session life span [New in Security Center 1.3]

Use the *glide.guest.active.session.life_span* property to control the duration of an active guest's HTTP sessions.

The *glide.guest.active.session.life_span* system property enforces a maximum lifespan on active guest HTTP sessions, regardless of session inactivity. The configured value is in minutes. A value of 0 disables the lifespan limit entirely, allowing sessions to persist until the inactive timeout fires. Guest users are unauthenticated users who access the instance without logging in.

Set the *glide.guest.active.session.life_span* system property to 720.

More information

Attribute	Description
Configuration name	*glide.guest.active.session.life_span*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Integer
Recommended value	720

Attribute	Description
Default value	0
Category	Session management
Security risk	• Severity score: 4.2 • CVSS score: Medium • Security risk details: A larger maximum lifespan could allow an attacker to persist a stolen session for longer, increasing the scope of a security incident.
Functional impact	This configuration enforces max life-span on active guest HTTP sessions irrespective of inactive timeout. The configured value is in minutes. A value of zero disables the lifespan limit entirely. The max life-span should be more than the inactive timeout *glide.ui.session_timeout* (default 30 minutes).
Dependencies and prerequisites	None

Limit concurrent interactive sessions [Updated in Security Center 1.3]

Manage the number of interactive sessions on your instance.

This property is meant to be used with the Limit Concurrent Sessions (*com.glide.limit.concurrent.sessions*) plugin. When the plugin is active and the property is set to false, a user can have any number of concurrent interactive sessions on an instance. A greater number of open sessions means there is a great possibility for session hijacking to occur.

More information

Attribute	Description
Configuration name	*glide.authenticate.limit.concurrent.interactive.se*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Session management
Security risk	• Severity score: 3.7 • CVSS score: Low • Security risk details: When the plugin is active and the property is set to **false**, a user can have any number of concurrent interactive sessions which increases the possibility for a session hijacking.
Dependencies and prerequisites	None

Attribute	Description
References	Limit concurrent interactive sessions [Updated in Security Center 1.3]

Limit integrations' active session life span [New in Security Center 1.3]

The *glide.integrations.active.session.life_span* property enforces max lifespan on active guest HTTP sessions irrespective of inactive timeout. The configured value is in minutes. A value of zero will disable timing out the active sessions.

More information

Attribute	Description
Configuration name	*glide.integrations.active.session.life_span*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	0 to 720
Default value	0
Category	Session management
Security risk	• Severity score: 4.2 • CVSS score: Medium • Security risk details: A larger maximum lifespan could allow an attacker to persist in a stolen session for longer, increasing the scope of a security incident. This property is limited to integrations that have low-privilege access to an instance. Set the glide property, *glide.integrations.active.session.life_span* to a value of 0, and less than or equal to 720.
Dependencies and prerequisites	None

Limit policy based session access mobile refresh token interval [New in Security Center 1.5]

Use the *glide.authenticate.session_access.mobile.refresh_token_interval* property to govern the length of time that must elapse before a mobile device user will be forced to re-authenticate.

A user will be asked to re-authenticate only if the admin has configured the Identity Provider attributes in the session policy (attributes can vary each login), and the user authenticates using Single Sign On (SSO). The default value represents the time in seconds that a user has before being re-authenticated. A larger default value provides a bad actor more time for session access in the event of a session hijacking.

More information

Attribute	Description
Configuration name	*glide.authenticate.session_access.mobile.refresh_t*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	1800 (seconds)
Default value	1800 (seconds)
Category	Session management
Security risk	• Severity score: 4.3 • CVSS score: Medium • Security risk details: If the ZTA policy is enabled on the instance, then users who are using SSO during mobile login will be forced to logout and re-login after the default value of 1800 seconds (30 minutes) have eclipsed. If a higher value is used, then users will be forced to wait that elapsed time.
Dependencies and prerequisites	Zero Trust- Policy Based Session Access
Functional impact	This setting governs the time in seconds after login, that users will be forced to logout from mobile devices if they are using Single Sign On to authenticate, and admin has configured the Identify provider attributes in the session access policy.

Limit UI active session life span [New in Security Center 1.3]

The *glide.ui.active.session.life_span* property enforces max lifespan on active authenticated HTTP sessions irrespective of inactive timeout.

Reduce the scope of potential security incidents by decreasing the lifespan of active HTTP sessions. The **glide.ui.active.session.life_span** system property enforces a maximum lifespan on active HTTP sessions irrespective of inactive timeout. Longer maximum lifespans can allow an attacker to use a stolen session for a longer time, increasing the scope of a security incident. The default value of **0** disables timeout of active sessions

Set the **glide.ui.active.session.life_span** to a value between 1 and 720. This value represents the time in minutes that HTTP sessions can remain active.

Note: The **glide.ui.active.session.life_span** is limited to UI session timeout.

More information

Attribute	Description
Configuration name	*glide.ui.active.session.life_span*
Configuration type	System Properties (/sys_properties_list.do)

Attribute	Description
Data type	integer
Recommended value	1-720
Default value	0
Category	Session management
Security risk	• Severity score: 4.2 • CVSS score: Medium • Security risk details: A larger maximum lifespan could allow an attacker to remain in a stolen session longer, increasing the possibility of a security incident.
Dependencies and prerequisites	None
Functional impact	Enforces max life-span on active authenticated HTTP sessions irrespective of inactive timeout. The configured value is in minutes. A value of zero will disable timing out the active sessions. The max life-span should be more than inactive timeout glide.ui.session_timeout (default 30 minutes).

Limit session length for high assurance sessions

Reduce the risk of account takeover in high assurance sessions by limiting session length

Reduce the risk of account takeover in high assurance sessions by limiting the length of the session. After the specified length of time, end-users must re-authenticate.

Use the **glide.zta.high_assurance.session.timeout** system property to set a time, in minutes, after which users must re-authenticate. The value of this property must be between 1 and 480 minutes. Consider limiting this value to 30 or lower to reduce the risk of account takeover.

More information

Attribute	Description
Technical configuration name	glide.zta.high_assurance.session.timeout
Plugin applicability	Zero Trust - Continuous Authentication (com.snc.zero_trust_continuous_authentication)
Security risk	Shorter session lengths reduce the risk of account takeover by forcing users to re-authenticate.
Common Vulnerability Scoring System (CVSS) score	3.3
Common Vulnerability Scoring System (CVSS) rating	Low
Functional impact	Users in high assurance sessions must re-authenticate at the interval set by this property.
Dependencies and prerequisites	None

Attribute	Description
Data type	Integer
Base system value	30
Fallback value	30
Recommended value	30 or lower

To learn more about adding or creating a system property, see Add a system property.

Proactively invalidate inactive sessions [New in Security Center 1.3 and updated in 1.5 and 2.0]

The *glide.active.session.timeout.invalidate.session* property controls if a timeout session is proactively invalidated before the Tomcat server.

When *glide.active.session.timeout.invalidate.session* is not set to **true**, there can be a small interval of time where a timed out session is not invalidated (60 or more seconds depending on queue size). If a session is hijacked, an attacker may be able to use a session during this small period of time. To remediate this security risk, set *glide.active.session.timeout.invalidate.session* to **true**.

More information

Attribute	Description
Configuration name	*glide.active.session.timeout.invalidate.session*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Session management
Security risk	• Severity score: 4.6 • CVSS score: Medium • Security risk details: Not setting this property to the recommended value of true could cause a timed out session to not be validated. This increases the chances of a bad actor hijacking a session.
Dependencies and prerequisites	None

Rotate HTTP session identifiers

Use the *glide.ui.rotate_sessions* property to enable rotation of the HTTP session identifiers to reduce security vulnerabilities.

If an unauthenticated user's session ID doesn't change after authentication, a web application is vulnerable to a session fixation attack. A malicious user could start an unauthenticated session and give the associated session ID to the victim. Once the victim authenticates, the malicious user now shares that authenticated session.

More information

Attribute	Description
Property name	*glide.ui.rotate_sessions*
Configuration type	System Properties (/sys_properties_list.do)
Category	Session management
Purpose	To achieve more secure session authentication.
Recommended value	true
Default value	true
Security risk rating	8.8
Functional impact	This remediation modified the SessionID when user navigates from unauthenticated page to authenticated pages. • If you are using a proxy or hardcoding the SessionID when a user first logs in, or for any purpose, then there can be a potential functionality impact. • If you are using the SAML 2.0 plugin for Single Sign-on authentication, it might interfere with the session information sharing between the instance and the Identity Provider. In such case, you can set this property to false.
Security risk	(Moderate) SessionID is used to process and authenticate the instance user by maintaining the session state on the browser. Thus, SessionID is deemed as sensitive data and should be secure by default. Session Rotation is a security control that enforces the alteration of sessionID whenever the user navigates from unauthenticated pages to authenticate pages.
References	Authentication with SAML

To learn more about adding or creating a system property, see Add a system property.

Minimize concurrent interactive session quantity [Updated in Security Center 1.3]

Use this property with the Limit Concurrent Sessions plugin to control the number of active sessions that can be opened by a user.

Use the **Glide Authenticate Max Concurrent Interactive Sessions** property with the **Limit Concurrent Session** (*com.glide.limit.concurrent.sessions*) plugin to control the number of active sessions opened for a user. The recommended value is **1**, which reduces the number of sessions opened (a greater number increases the likelihood that a session could be hijacked).

More information

Attribute	Description
Configuration name	*glide.authenticate.max.concurrent.interactive.sess*
Configuration type	System Properties (/sys_properties_list.do)
Data type	integer
Recommended value	1
Default value	1
Category	Session management
Security risk	• Severity score: 3.7 • CVSS score: Low • Security risk details: Setting the default value of the property greater than 1 increases the chance of session hijacking.
Dependencies and prerequisites	The **Concurrent Interactive Sessions** (*com.glide.limit.concurrent.sessions*) plugin should be active.
References	Limit concurrent sessions

Minimize session activity timeout duration [Updated in Security Center 1.3]

Use the *glide.ui.session_timeout* property to designate, in minutes, activity timeout value.

There are several functional impacts from setting this property:

- The longer the specified session time-out, the greater the amount of memory is utilized during a processing session. The base system uses a default Apache Tomcat timeout duration of 30 minutes.
- The ServiceNow AI Platform still logs out users out with Remember Me. After 30 minutes of inactivity in the application, the platform logs the user out automatically, unless the **Remember Me** check box in the login page is selected. What 's different is that they don't log in again to continue.
- If there are gauges or content on users' home pages that refresh automatically, then this timeout may never be reached.

More information

Attribute	Description
Property name	*glide.ui.session_timeout*
Configuration type	System Properties (/sys_properties_list.do)
Category	Session management
Purpose	To enforce session timeout.

Attribute	Description
Recommended value	User specified timeout in minutes. 60 minutes is the recommended value, but this value may vary depending on functionality and security requirement. Do not set this value to more than one day.
Security risk rating	7.5
Functional impact	This remediation enforces timely expiration of user account. No functionality impact, however User experience is altered.
Security risk	(High) User sessions being active for indefinite amount of time is a security risk and should expire on a time-based configuration.
References	Manage user sessions

To learn more about adding or creating a system property, see Add a system property.

Minimize session window timeout duration [Updated in Security Center 1.3]

Use the *glide.ui.user_cookie.life_span_in_days* property to set the expiration time period for the Remember Me cookie. The default value is 15 days and the maximum cap is at 30 days.

More information

Attribute	Description
Property name	*glide.ui.user_cookie.life_span_in_days*
Configuration type	System Properties (/sys_properties_list.do)
Category	Session management
Purpose	To enable the default expiry of the **remember me** cookie.
Data type	Integer
Recommended value	15
Default value	15
Security risk rating	4.9
Functional impact	This property is enabled by the end user when the end user checks the **Remember me** check box from the login page and logs in to the ServiceNow AI Platform.
Security risk	(Moderate) The user cookies being active for an indefinite amount of time is a security risk and should expire on a time-based configuration.
References	Available system properties Remember me

To learn more about adding or creating a system property, see Add a system property.

Stored cryptography

This category focuses on the encryption of stored data. It encompasses several key aspects, such as employing established algorithms and cryptographic modules, ensuring the proper generation of pseudo-random values, implementing encryption based on data classification, and securely storing and isolating key material.

Enable glide KMF encrypter [Removed in Security Center 1.3.2]

Manage the encrypters used for Password2 fields on your instance.

Use the *glide.kmf.encrypter.enabled* property to set KMF encrypter as the default encrypter for Password2 fields. This property ensures that strong and compliant encryption standards are being used instead of a legacy encrypter. To ensure that KMF encrypter is used instead of the legacy encrypter, set this property to **true**.

More information

Attribute	Description
Configuration name	*glide.kmf.encrypter.enabled*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Stored cryptography
Security risk	• Severity score: 4.9 • CVSS score: Medium • Security risk details: Setting this property to **true** ensures that KMF encrypter is being used instead of legacy encrypter.
References	Password2 encryption with the Key Management Framework (KMF)

Disable use of TripleDES/3DES encryption algorithm

Avoid the security risks of outdated encryption methods by disabling the use of the TripleDES/3DES encryption algorithm.

Disable the use of legacy, deprecated Triple Data Encryption Standard (3DES or Triple-DES) on your instance, which may lead to loss and leakage of sensitive information.

The National Institute of Standards and Technology (NIST) has advised against using Triple DES (3DES) to encrypt data. For more details, refer to NIST 800-131A Rev 2.#3DES is now banned for encryption in federal systems. Going forward, TDES can only be used for historical purposes, such as decrypting old messages, key unwrapping, and MAC verification.

The **glide.security.3des.encryption.allow** system property controls whether 3DES encryption is enabled on your instance, however there are several steps which may need to be taken to verify that your instance is ready for deprecation. Review KB1704481 for further information on the deprecation process before setting this property.

More information

Attribute	Description
Technical configuration name	glide.security.3des.encryption.allow
Plugin applicability	None
Security risk	Usage of outdated and weak encryption like 3DES may lead to loss and leakage of sensitive information.
Common Vulnerability Scoring System (CVSS) score	4.2
Common Vulnerability Scoring System (CVSS) rating	Medium
Functional impact	Setting this property to **false** helps prevent your instance from use of outdated and weak encryption. This affects how Password2 data is stored. Review KB1704481 for further information on functional impact of 3DES deprecation.
Dependencies and prerequisites	Perform the tasks outlined in KB1704481 before changing the value of this property.
Data type	Boolean
Base system value	false
Fallback value	false
Recommended value	false

To learn more about adding or creating a system property, see Add a system property.

Prevent usage of 3DES keys [New in Security Center 7.0]

Disable the use of 3DES static keys on your instance with a system property.

Use the *glide.security.3des.static_keys_usable* system property to disable the usage of 3DES static keys on your instance. This property should be set automatically by a scheduled job, the status of which can be checked via the property *glide.security.3des.removal_job_status*. In cases where this property isn't configured to **false**, after the scheduled job runs and the status is KEYS_DEACTIVATED, the scheduled job waits until this property is configured to **false** before proceeding with deletion of the 3DES static keys.

Ensure the *glide.security.3des.static_keys_usable* exists on the System Properties [sys_properties] table and is set to a value of **false**. Review KB1704481 for further information on 3DES deprecation.

More information

Attribute	Description
Configuration name	*glide.security.3des.static_keys_usable*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean

Attribute	Description
Recommended value	false
Default value	true
Fallback value	true
Category	Stored cryptography
Security risk	• Severity score: 5 • CVSS score: Medium • 3DES is deprecated, and the use of 3DES and 3DES static keys outside of temporary backwards compatibility isn't recommended. Continued use may lead to situations of sensitive information disclosure if a user were to obtain access to encrypted data.
Functional impact	When set to **false**, there may be unlikely situations where code and data residing on the instance still relied on by 3DES static keys are now inaccessible.
Dependencies and prerequisites	None

To learn more about adding or creating a system property, see Add a system property.

Validation, sanitization, and encoding

Validation, sanitization, and encoding addresses input validation to prevent against vulnerabilities like Cross-Site Scripting (XSS), SQL injection and other attacks.

This control ensures input validation and output encoding are in place and correctly configured, such as encoding or escaping output data. This category also includes checks for items such as deserialization of objects and positive validation through allow lists.

Allow HTML Links to Trusted Domains in the Description Fields of the Impact Workspace Module [New in Security Center 7.0]

Use a system property to help sanitize the HTML allowed in the descriptions fields. This property limits the allowed links to only those from the trusted domains listed in the property.

Note: This hardening setting is not a part of the hardening baseline. It does not appear in Security Center hardening pages and affect your hardening score.

The Impact Workspace module allows HTML in a number of description-related fields. When configured, the *sn_impact_common.whitelisted.url_HTML_injection* system property contains a comma-separated list of domain names. Description fields for the Impact Workspace module are allowed to contain HREFs with URLs only from the domains listed in the property.

Ensure the *sn_impact_common.whitelisted.url_HTML_injection* system property is set to a comma-separated list of domain names that represent the domains allowed in HTTP reference URLs of description fields for the Impact Workspace module.

To disallow HREFs in these fields, set the property to an empty string. If the property doesn't exist on the System Properties [sys_properties] table, it defaults to this list: `servicenow.com, service-now.com, youtube.com, google.com, youtu.be, soti.net, dpdhl.sharepoint.com, documentation.avaya.com, www.juniper.net,`

`servicenow.sharepoint.com, servicenow-my.sharepoint.com, scaledagileframework.com`.

More information

Attribute	Description
Configuration name	*sn_impact_common.whitelisted.url_HTML_injection*
Configuration type	System Properties (/sys_properties_list.do)
Data type	String list
Recommended value	servicenow.com, service-now.com, youtube.com, google.com, youtu.be, soti.net, dpdhl.sharepoint.com, documentation.avaya.com, www.juniper.net, servicenow.sharepoint.com, servicenow-my.sharepoint.com, scaledagileframework.com
Default value	servicenow.com, service-now.com, youtube.com, google.com, youtu.be, soti.net, dpdhl.sharepoint.com, documentation.avaya.com, www.juniper.net, servicenow.sharepoint.com, servicenow-my.sharepoint.com, scaledagileframework.com
Fallback value	servicenow.com, service-now.com, youtube.com, google.com, youtu.be, soti.net, dpdhl.sharepoint.com, documentation.avaya.com, www.juniper.net, servicenow.sharepoint.com, servicenow-my.sharepoint.com, scaledagileframework.com
Category	Validation, sanitization, and encoding
Security risk	• Severity score: 4.4 • CVSS score: Medium • If an untrusted domain is added to the property, this opens these fields up to containing links to risky sources which can lead HTML injection attacks. The exact risk is dependent on the customer instance.
Functional impact	If the property is empty, no HREFs are allowed in the field text and all HREFs are removed. Any links using domains not listed in the property are removed. An improper value for this field could result in corrupted data for the affected fields.
Dependencies and prerequisites	If the *sn_impact_common.blacklist_tags_HTML_injection* system property contains HTML tags that surround HREF links, then all links within those tags will be removed.

To learn more about adding or creating a system property, see Add a system property.

Restrict access to GlideSystemUserSession scriptable API [Updated in Security Center 1.3 and 2.0]

The client callable GlideSystemUserSessionSandbox scriptable API exposes GlideSystemUserSession's addErrorMessageNoSanitization and addInfoMessageNoSanitization methods to the JavaScript sandbox. This allows all users to call this method via script.

`gs.addErrorMessageNoSanitizationMessaging()` and *`gs.addInfoMessageNoSanitization()`* are used within the scripting environment for logging and notifications. Both of these are available in the sandbox if this property is not set to the recommended value of false. The sandbox is a low privileged scripting environment available to unauthenticated and no role users. Both of these methods can be used to display unsanitized input to a user. Displaying unsanitized input to the user is dangerous, as unsanitized input may contain dangerous code that runs in the user's browser. This can be utilized for traditional reflected XSS attacks. Reflected XSS attacks can be used in multiple scenarios, including session hijacking.

Set **glide.sandbox.usersession.allow_unsanitized_messages** system property to **false**. If there is not record of this property in the System Properties [sys_properties] table, create one.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*`glide.sandbox.usersession.allow_unsanitized_messag`*
Configuration type	System Properties (/sys_properties_list.do)
Category	Access control
Purpose	This property will restrict unsanitized informational or error messages from being called in a sandboxed user session.
Type	Boolean
Recommended value	false
Default value	true
Security risk rating	8.1
Functional impact	Set the property with the value false will result in no message creation or logging should those functions get called.
Security risk	(High) Without appropriate sanitization, potentially dangerous content may be accessed and the unsanitized error function is available to script.
References	Access control

Disable JavaScript tags in embedded HTML [Updated in Security Center 1.3]

Use the *`glide.ui.security.codetag.allow_script`* property to disable support for embedding HTML JavaScript code created using of the [code] tag.

The ServiceNow AI Platform mitigates many injection and cross-site attacks by implementing escaping and encoding techniques. As a result, users can't write and submit HTML formatted inputs for journal fields. However, journal fields can render text enclosed within code tags as HTML. Ensure the *glide.ui.security.codetag.allow_script* property exists in the sys_properties table and is set to false.

- However, there is an associated security risk. If set to **true**, malicious users can write harmful HTML JavaScript code that may be executed on a different client browser after rendering of journal fields.
- Set this property to **false** so that administrators can prevent journal fields from rendering HTML JavaScript code by disabling support for the `[code]` tag.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.ui.security.codetag.allow_script*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	Protects against cross-site scripting and malicious script execution
Recommended value	false
Default value	false
Security risk rating	8.8
Functional impact	This remediation enforces JavaScript escaping to occur on the UI and renders the encoded results to the user. It can have a functionality impact based on the instance user interaction with the resulted data.
Security risk	(High) Input validation must occur in the application to defend against cross-site scripting attacks. These attacks enable foreign scripts to execute on the user session in the logged in browser's context. Attackers can use it to steal session information and sensitive data.
References	Restrict the CODE tag in journal fields Render journal field entries as HTML High Security Settings

To learn more about adding or creating a system property, see Add a system property.

Enable the hardened java security manager [New in Security Center 1.3]

The *glide.security.manager* property contains the Java classname of the current Java security manager.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Configuration name	*glide.security.manager*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	*com.glide.sys.security.ContextualSecurityManager*
Default value	*com.glide.sys.security.ContextualSecurityManager*
Category	Validation, sanitization, and encoding
Security risk	• Severity score: 7.2 • CVSS score: High • Security risk details: If *glide.security.manager* is not set to the recommended value of *com.glide.sys.security.ContextualSecurityManager* then the instance may be using an obsolete Java security manager which is missing expected hardening policies. Without this hardening, a malicious actor with script execution access could get remote code execution on the instance.
Dependencies and prerequisites	None

Enforce HTML Sanitization [Updated in Security Center 1.3]

Use the *com.glide.security.check_unsanitized_html* property to enforce sanitization behavior of translated_html fields on a global level for field assignments.

HTML is one of the types that can be assigned to the dictionary fields. Assigning HTML fields to any field type provides the functionality to format content using HTML tags (for example, `<p>`, `<a href>`, `<b>`, `<font>`, `<img>`). To prevent malicious activity, certain HTML tags can be disallowed using a block list. This property will prevent disallowed tags from being used in translated_html fields on your instance.

- Set this property to **enforce** to enforce sanitization behavior of translated_html fields.
- Set property to **disable** to turn off the html sanitization to allow blocked html tags on translated_html fields.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*com.glide.security.check_unsanitized_html*

Attribute	Description
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	Prevents the use of insecure HTML tags to protect against attacks such as cross-site scripting.
Type	String
Recommended value	enforce
Default value	enforce
Security risk rating	7.3
Functional impact	This remediation enforces HTML sanitization to occur on the UI and renders translated html fields to the user. It can have an impact on readability and formatting.
Security risk	(High) Input validation must occur on the application to defend against cross-site scripting attacks. These attacks enable foreign scripts to execute on user sessions in the logged in browser's context. Attackers can use it to steal session information and sensitive data.
References	HTML sanitizer

Ensure Contextual Search Do Not Contain An Unvalidated Redirect [New in Security Center 7.0]

Prevent Contextual Search results from containing referral links outside the current domain with a system property.

The Contextual Search plugin displays search results in a new window using the *cxs_new_window* UI page. This UI page contains a referral link which can be set by providing a value to *sysparm_url*. When the *com.snc.contextual_search.cxs_new_window.force_relative_link* system property is set to **true**, *sysparm_url* can only contain links that are relative to the current domain. This restriction prevents the UI page from being used as an unvalidated redirect to an attacker-controlled website. When the property is set to **false**, *sysparm_url* can link to any website.

Set the *com.snc.contextual_search.cxs_new_window.force_relative_link* property to **true**. If the property doesn't exist on the System Properties [sys_properties] table, the default value is **false**. If the property exists on the table, it defaults to **true**.

More information

Attribute	Description
Configuration name	*com.snc.contextual_search.cxs_new_window.force_rel*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true

Attribute	Description
Fallback value	false
Category	Validation, sanitization, and encoding
Security risk	• Severity score: 3.1 • CVSS score: Medium • When set to **false**, *sysparm_url* can link to any website, which allows the UI page to be used as an unvalidated redirect to an attacker-controlled website.
Functional impact	When set to **true**, *sysparm_url* is only allowed to contain links that are relative to the current domain. This restriction means that the UI page can only ever link to web pages on the current domain. However, the UI page is meant to display search results from the current domain and should only ever link to the current domain.
Dependencies and prerequisites	The Contextual Search (com.snc.contextual_search) plugin must be active.

To learn more about adding or creating a system property, see Add a system property.

Disable AJAXEvaluate

Use the *glide.script.allow.ajaxevaluate* to protect the system API from vulnerabilities of Client script execution through AJAX calls.

Elevation to the security_admin role is required to edit the property.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.script.allow.ajaxevaluate*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	To prevent a user from executing scripts as an admin privilege.
Recommended value	false
Default value	false
Configuration type	Boolean
Functional impact	This remediation forces the AJAXEvaluate processor to be turned off. It could impact functionality if you are explicitly using the AJAX evaluate processor as part of any customized scripts.

Attribute	Description
Security risk	(High) The AjaxEvaluator processor executes Client scripts in sandbox, however there are several additional properties which can allow the scope of activities in the sandbox to expand.
Security risk rating	7.3
References	This property belongs to the same family of properties that secure and restrict execution of scripts originating from the client, such as *glide.script.allow.ajaxevaluate*. For more information, see Enable AJAXEvaluate.

To learn more about adding or creating a system property, see Add a system property.

Disable Entity Expansion within the XMLDocument2 Streaming Parser [Updated in Security Center 1.5]

If customizations do not require entity expansion, use the *glide.stax.allow_entity_resolution* property to completely disable external entity expansion. The XML completes parsing but doesn't include any internal or external entities.

Disable entity expansion on your instance to secure your instance from attacks such as ability to read system files, and Denial of Service. Use the system property to disallow XML entities to be expanded during parsing by the streaming parser (XMLDocument2).

Set the **glide.stax.allow_entity_resolution** system property to **false** to disable entity expansion on your instance. If this property does not appear in the System Properties [sys_properties] table, the default value is **true**. Create the property record and set the value to **false** to change it's value.

Prerequisites

Before setting this property:

- Set the *glide.xml.entity.whitelist.enabled* and *glide.stax.whitelist_enabled* properties to true. To learn more, see Restrict XML external entities [Updated in Security Center 1.3 and 2.0] and Require XMLdoc2 entity validation with allowlistDisable entity expansion [Updated in Security Center 1.3].
- Define a listing of comma-delimited FQDN in the *glide.xml.entity.whitelist* property, which is the only URLs that can be reached using XML Entity processing. property. To learn more, see Restrict XML external entities [Updated in Security Center 1.3 and 2.0].

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.stax.allow_entity_resolution*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding

Attribute	Description
Purpose	This remediation control must be enabled to defend against an XML Entity Expansion/Billion Laugh attack.
Recommended value	false
Default value	true
Functional impact	If the customization is using entity expansion, then, the ServiceNow AI Platform might block further processing.
Security risk	(Critical) An attacker can use this vulnerability to expand data exponentially, quickly consuming all system resources.
Workaround	If the customization requires entity expansion, set this property to true and follow the steps documented in Require XMLdoc2 entity validation with allowlistDisable entity expansion [Updated in Security Center 1.3].

To learn more about adding or creating a system property, see Add a system property.

For more information about OWASp resources, see OWASp.

Disable external content url [Updated in Security Center 2.0]

Manage how external link metadata is used in your instance with Connect Chat.

Use the *glide.ui.url.external.content* property to manage external link metadata in your instance. If the property is set to the recommended value of **false**, then no external link metadata will be rendered. If set to **true** then Connect Chat will retrieve external link metadata from sources such as YouTube or news articles to render richer messages. This could lead to Server Side Request Forgery (SSRF) attacks.

Ensure the Glide Property *glide.ui.url.external.content* exists and is set to the value false. If the property does not appear in the sys_properties table, add a new record.

More information

Attribute	Description
Configuration name	*glide.ui.url.external.content*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	true
Category	Validation, sanitization, and encoding
Security risk	• Severity score: 7.2 • CVSS score: High • Security risk details: Setting this property to **true** could expose your instance to Server Side Request Forgery (SSRF) attacks.

Attribute	Description
Dependencies and prerequisites	None
References	Connect Chat

Restrict downloadable MIME types [Updated in Security Center 1.3 and 2.0]

The *glide.ui.attachment.download_mime_types* property will force the specified list of dangerous file types to be downloaded to the client and not viewed inline in the browser.

If the property *glide.ui.attachment.force_download_all_mime_types* is set to true, then the *glide.ui.attachment.download_mime_types* property will be overridden so that all MIME types will be downloaded rather than rendered by the browser. For example, downloading text/html forces an HTML file to be downloaded to the client as a file rather than viewed inline in the browser, preventing a XSS attack. XSS can lead to easily attained privilege escalation to higher roles such as admin where more lateral movement can be taken.

New remediation: Ensure the property *glide.ui.attachment.force_download_all_mime_types* is set to true. If the property does not exist in the sys_properties table, the default value is false.

Note: The security_admin role is required to edit the property.

More information

Attribute	Description
Property name	*glide.ui.attachment.download_mime_types*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	Maintaining the list properly of dangerous file types that cannot be viewed in the browser will prevent cross site scripting attacks (XSS).
Recommended value	List of applicable MIME types or the recommended value: `text/html,image/svg,image/svg+xml,application/xml`
Default value	List of applicable MIME types for the default value: `text/html,image/svg,image/svg+xml,application/xml`
Configuration type	String: any comma separated values of application mime types.
Functional impact	This remediation enforces performance of validation checks before performing an action when you click an attachment in a ServiceNow AI Platform application. There is no potential impact, but the user experience is altered.
Security risk	(Moderate) Attackers can abuse MIME types and place unintended script content in the attachment on the victim's side to capture sensitive information. The ability to have XSS can lead to easily attained privilege escalation to higher roles, such as admin, where more lateral movement can be taken.

Attribute	Description
	In the current context, populate the property with a list of comma-separated attachment MIME types that should not render inline in the browser.
Security risk rating	6.4
Related properties	• *glide.ui.attachment.force_download_all_mime_type* • *glide.ui.attachment.tables_ignore_force_download*
References	Define restricted downloadable MIME types [Updated in Security Center 1.3, 1.5, and 2.0].

Disable embedded HTML code [Updated in Security Center 1.3]

Use the *glide.ui.security.allow_codetag* property to disable support for embedding HTML code created using the [code] tag.

Disable support for displaying HTML code embedded using the `[code]` tag. This tag allows rendered HTML to display in journal fields and may lead to cross-site scripting (XSS) attacks. These attacks can enable foreign scripts to execute on a user session in the logged in browser's context. Attackers can use these scripts to steal session information and sensitive data. The HTML language was not designed to separate script from formatting, so allowing user-controlled HTML in any system has inherent risk.

If setting *glide.ui.security.allow_codetag* to `false` disrupts instance functionality, for example, if your instance uses a feature that relies on HTML rendering in journal fields, you can maintain a compliant security posture by keeping *glide.ui.security.allow_codetag* set to `true` and setting *glide.ui.security.codetag.allow_script* to `false`. This change disables script execution within `[code]` tags while preserving HTML rendering. Note that this approach carries some residual risk, as it relies on sanitizing all known script conventions within HTML rather than prohibiting HTML code tags entirely.

Set the *glide.ui.security.allow_codetag* system property to `false` to completely prohibit journal fields and forms from displaying rendered HTML.

The ServiceNow AI Platform mitigates many injection and cross-site attacks by implementing escaping and encoding techniques. As a result, users can't write/submit HTML formatted inputs for journal fields. But journal fields can render text enclosed within code tags as HTML.

- However, there is an associated security risk. If set to `true`, malicious users can write harmful HTML JS code that may be executed on a different client browser after rendering of journal fields.
- Set this property to `false` so that administrators can prevent journal fields from rendering HTML code by disabling support for the `[code]` tag.

More information

Attribute	Description
Property name	glide.ui.security.allow_codetag
Configuration type	System Properties (/sys_properties_list.do)

Attribute	Description
Category	Validation, sanitization, and encoding
Configure in Instance Security Center	Yes
Purpose	Protect against cross-site scripting and malicious script execution
Recommended value	false
Default value	true
Security risk rating	4.2
Functional impact	This remediation enforces HTML encoding to occur on the UI and renders the encoded results to the user. This property is set to `true` by default. In this state, your instance displays rendered HTML in journal fields and forms. If this property is set to `false`, HTML is not rendered properly and HTML tags may appear in journal fields on forms. It can have an adverse impact on functionality, and on user interactions with the resulting data. If this property negatively affects functionality, set *`glide.ui.security.codetag.allow_script`* to `false` to disable script execution within `[code]` tags while preserving HTML rendering.
Security risk	(Medium) Input validation must occur in the application to defend against cross-site scripting attacks. These attacks enable foreign scripts to execute on a user session in the logged in browser's context. Attackers can use it to steal session information and sensitive data.

Enable HTML Sanitizer within Virtual Agent [Updated in Security Center 1.3 and 1.5]

Use the *`com.glide.cs.html.sanitizer.enabled`* property to enable HTMLSanitizerService.

This property controls the whether the HtmlSanitizerService is enabled. If *`com.glide.cs.html.sanitizer.enabled`* is not set to true, then a Stored Cross-Site Scripting (XSS) attack is possible in the VA web client.

More information

⚠ **Warning:** This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

Attribute	Description
Property name	*`com.glide.cs.html.sanitizer.enabled`*
Configuration type	System Properties (/sys_properties_list.do)

Attribute	Description
Category	Validation, sanitization, and encoding
Purpose	Prevents application against cross site scripting and HTML injection attacks.
Recommended value	true
Default value	true
Security risk rating	8
Functional impact	This remediation enforces HTML-output encoding mechanism before the user data is rendered back to the user. If customer has any customization that involves rendering of the HTML attribute or content data, then there is a functionality impact.
Security risk	(High) User input should be securely treated when the data is being stored and processed on the application. This reduces client-side cross-site scripting attacks by output encoding the data.

To learn more about adding or creating a system property, see Add a system property.

Enable Jelly JS Interpolation Protection

Use the *glide.ui.jelly.js_interpolation.protect* property to ensure that any JavaScript about to be executed on a Jelly page is protected from injection with the help of Jelly interpolation.

When you set property to **true**, an application goes through a Jelly script tree (nested). It wraps potentially dangerous Jelly expressions with a filter that:

- Escapes their results to be safe, or
- If their safety can't be guaranteed, generates a SecurityException because the expression that was going to be evaluated represents a possible security issue.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.ui.jelly.js_interpolation.protect*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	To mitigate against malicious code execution attacks that can occur using Jelly Injection.
Recommended value	true
Default value	false
Security risk rating	9

Attribute	Description
Functional impact	This property makes a best guess at whether an expression is quoted. It may wrongly quote legitimate expression. In that case manually marking an expression as safe may be necessary.
Security risk	(Moderate) JEXL injection is a form of input injection unique to the ServiceNow AI Platform that can lead to both cross-site request forgery and code execution. Completely turning off the protection may potentially open many P1 security vulnerabilities.
Workaround	To manually mark an expression as safe add SAFE prefix to Jelly expression: `${SAFE:sysparm_input};` Blindly adding SAFE to each expression is the wrong way to approach the problem, because it may open a security vulnerability. • Only add SAFE to an expression if you can guarantee that the expression does not contain input from the client. • If it does, it's possible for a malicious client to cause evaluation of privileged JavaScript.
References	Jelly tags High Security Settings

To learn more about adding or creating a system property, see Add a system property.

Enable Jelly JS interpolation protection for nested expressions [Updated in Security Center 2.0]

Manage the interpolation protection on your instance.

Use the *glide.ui.jelly.js_interpolation.protect_nested_expressions* property to manage interpolation protection. Interpolation protection ensures that when Jelly expressions are used in JavaScript, that they must be deemed as safe by either falling under certain categories or being marked as SAFE in the expression itself. Without this mitigation enabled, a bad actor can send a GET parameter to a Jelly page and cause the contents of that parameter to be evaluated as server-side JavaScript with admin privileges. If this property is not set to the recommended value of **true**, malicious Jelly expressions interpolated in JavaScript are allowed and a user can execute code using a Jelly template.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Configuration name	*glide.ui.jelly.js_interpolation.protect_nested_exp*

Attribute	Description
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	false
Category	Validation, sanitization, and encoding
Security risk	• Severity score: 9 • CVSS score: Critical • Security risk details: If the property is set to **false**, then malicious Jelly expressions are allowed.
Dependencies and prerequisites	None

Enforce relative links [Updated in Security Center 1.3 and 1.5]

Use the *glide.cms.catalog_uri_relative* property to enforce relative links from the URI parameter on `/ess/catalog.do`.

The *glide.cms.catalog_uri_relative* property enforces relative links from the URI parameter on /ess/catalog.do. If *glide.cms.catalog_uri_relative* is not set to the recommended value of true, then the URL will not be sanitized with the enforceRelativeURL(url) function. Absolute URLs can pose a security risk when used as a part of parameter or a field value, thus redirecting the source page to an adversary-controlled website.

More information

Attribute	Description
Property name	*glide.cms.catalog_uri_relative*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	To restrict attempts to link external unauthorized content.
Recommended value	true
Default value	false
Security risk rating	2.6
Functional impact	This remediation enforces validation on Catalog page such that only Relative URLs are permitted. Existing links to external web applications become broken.
Security risk	(High) Absolute URLs can pose a security risk when used as a part of parameter or a field value, thus redirecting the source page to an adversary-controlled website.

To learn more about adding or creating a system property, see Add a system property.

Enforce URL allowlist check [Updated in Security Center 1.3, 1.5, and 2.0]

Use the *glide.security.url.whitelist* system property to add extra layer of validation to ensure whether any external URL introduced should be a part of inclusion listed URLs.

Protect your users from client-side open redirection, which enable attackers to redirect your users to untrusted and malicious pages.

If **glide.security.url.whitelist.strict_check** is not set to the recommended value of **true**, all external URLs are allowed for redirection as long as the **glide.security.url.whitelist** system property is empty. If **glide.security.url.whitelist** is not empty, then only external URLs listed in that property are allowed.

Set **glide.security.url.whitelist.strict_check** to **true** or ensure that **glide.security.url.whitelist** is configured with the allowed external URLs to help secure your instance from open redirection attacks.

This property is applicable in the following cases:

- `/logout.do?sysparm_goto_url={External URL}`
- `/cms_login_redirect.do?sysparm_goto_url={External URL}`

Users are directed to an external trusted site after they log out of the instance:

- `/logout_redirect.do?sysparm_url={External URL}`
- `/saml_redirector.do?sysparm_uri={External URL}`
- `/logout_success.do?RelayState={External URL}`

When SAML is enabled, it invokes an identity provider (IDP) logout URL.

Ensure the property *glide.security.url.whitelist.strict_check* is set to true or the property *glide.security.url.whitelist* is set to a value.

More information

Attribute	Description
Property name	*glide.security.url.whitelist*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	To implement safe URL redirect during login, logout, or other redirects. This property mitigates one of the OWASP top 10 attacks called Invalidated Redirects and forwards.
Type	String
Default value	true
Recommended value	true
Value	Your organization's approved URLs [Some Defined FQDN (Fully Qualified Domain Name) Ex. http://www.servicenow.com]
Security risk rating	6.3

Attribute	Description
Functional impact	This remediation enforces validation on logout page. It might have a functional impact on a user of an instance with an SSO/SAML configuration.
Security risk	(High) Client-side open redirection can enable attacker to redirect victims/users to attacker-controlled website and is viewed as a security risk.
References	Multi-SSO (SAML 2.0) errors and fixes

To learn more about adding or creating a system property, see Add a system property.

Escape Excel Formulas [Updated in Security Center 1.3]

Use the *glide.export.escape_formulas* property to prevent Excel Injection, also, known as formula injection.

Prevent potentially malicious formulas in programs such as Excel from being executed after exporting and opening the file by escaping formulas in these files. Excel injection occurs when websites embed untrusted entries inside Excel files. When you use a spreadsheet application such as Microsoft Excel, or LibreOffice Call, to open a file, any cells starting with +, -, =, or @ are interpreted as a formula unless properly escaped. Malicious formulas pose a risk even when the spreadsheet doesn't contain any sensitive information, as they can be used to compromise the viewer's computer through code execution.

Set the **glide.export.escape_formulas** system property to **true** to escape these formulas from executing.

More information

Attribute	Description
Property name	*glide.export.escape_formulas*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	To prevent application against the Excel or formula injection.
Recommended value	true
Default value	false
Security risk rating	6.4
Functional impact	Maliciously crafted formulas can be used for hijacking the user's computer by exploiting vulnerabilities in the spreadsheet software.
Security risk	(Moderate) Malicious formulae pose a risk even when the embedding spreadsheet doesn't contain any sensitive information, as they can be used to compromise the viewer's computer.

Attribute	Description
Workaround	As an alternative consider stripping all trailing white spaces where possible, and limiting all client-supplied data to alpha-numeric characters.
References	Available system properties

To learn more about adding or creating a system property, see Add a system property.

Escape HTML in list views [Updated in Security Center 1.3 and 1.5]

Use the *glide.ui.escape_html_list_field* property to force HTML escapes for HTML fields in a list view.

Set *glide.ui.escape_html_list_field* to **true** to prevent HTML from being rendered in HTML fields in list view. Leaving HTML sanitization inactive platform wide (via system property) or by field (via a schema attribute), may lead to XSS style attacks. XSS attacks may allow a low privileged user to hijack the session of a high privileged user or interfere in standard web application behaviors, including redirects or defacement.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It's non-revertible.

More information

Attribute	Description
Property name	*glide.ui.escape_html_list_field*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	To help prevent application against cross-site scripting attacks
Recommended value	true
Default value	true
Security risk rating	3.1
Functional impact	This remediation enforces HTML encoding to occur on the UI at the HTML parser level and thus renders back encoded results to the user. It can have a functionality impact based on the instance user interaction with the resulted data.
Security risk	(High) Input validation must occur on the application to defend against cross-site scripting attacks. These attacks enable foreign scripts to execute on user sessions in the logged in browser's context. Attackers can use it to steal session information and sensitive data.
References	High Security Settings

To learn more about adding or creating a system property, see Add a system property.

Escape JavaScript [Updated in Security Center 1.3]

Use the *glide.html.escape_script* property to force escape from JavaScript (`<script></script>`) tags in HTML fields during list views.

The glide property *glide.html.escape_script* helps sanitize HTML fields. If *glide.html.escape_script* is not set to the recommended value of true, then inputs will not be sanitized for HTML fields (output encoding) from a backend Java context by removing embedded JavaScript. Javascript in HTML fields can lead to stored and reflected XSS. The ability to have XSS can lead to easily attained privilege escalation to higher roles such as admin where more lateral movement can be taken.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.html.escape_script*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	To prevent cross-site scripting attacks against an application.
Recommended value	true
Default value	true
Security risk rating	8.8
Functional impact	This remediation enforces JavaScript escaping to occur on the UI and renders encoded results to the user. It can have an impact on functionality, based on the instance user interaction with the resulting data
Security risk	(High) Input validation must occur in the application to defend against cross-site scripting attacks. These attacks enable foreign scripts to execute on user session in the logged in browser's context. Attackers can use it to steal session information and sensitive data.
References	Available system properties High Security Settings

To learn more about adding or creating a system property, see Add a system property.

Escape jelly script [Updated in Security Center 1.3 and 1.5]

Use the *glide.ui.escape_all_script* property to force escape of all scripts injected into Jelly.

This property escapes all the JS and HTML strings included in <j:jelly> ... </j:jelly> before they are written to the output stream, preventing several XSS issues from occurring. If *glide.ui.escape_all_script* is not set to the recommended value of true, then escaping of scripts injected into Jelly is disabled. Without this mitigation, the platform becomes

widely open to a variety of script injection attacks. An attacker could execute arbitrary Rhino scripts on the instance.

Warning: Be careful when using these tags. If user input is displayed here it can open a security vulnerability.

More information

Attribute	Description
Property name	*glide.ui.escape_all_script*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Category	Validation, sanitization, and encoding
Purpose	If the property is not set to **true**, developers have to perform several steps on each custom Jelly script to prevent XSS issues. These steps include locating the Jelly variables being sent to output stream to render on web pages, and performing escaping on each of the following tags: `$â {JS:expression}` `$â {HTML:expression}` OR `$â {JS,HTML:expression}`
Recommended value	true
Default value	true
Security risk rating	7.3
Functional impact	This remediation enforces Jelly escaping at the parser level. It can have a functionality impact on user interaction with the resulting data.
Security risk	(High) Input validation has to occur on all the user input being entered on the application. By doing so, injection attacks against the platform can be defended and protected.
Workaround	The UI may be affected because some of the scripts and HTML tags designed for rendering on a web page may appear broken. This remediation sends the output encoded page to the browser to render. For example, instead of `'my string here'`, it might display `'<u>my string here</u>` as the `<u>` tag was properly escaped. In this case, to prevent escaping, add the NOESC prefix to Jelly expression to prevent JS escaping. For example:

Attribute	Description
	• Before: ($[jvar_context_menus]); • After: `($[NOESC:jvar_context_menus]);` • Before: $[jvar_ui_policy_scripts] • After: `$[NOESC:jvar_ui_policy_scripts]` **Warning:** Be careful when using these tags. If user input is displayed here it can open a security vulnerability.
References	High Security Settings Jelly tags

To learn more about adding or creating a system property, see Add a system property.

Escape scripts in scratchpad [Updated in Security Center 1.3]

Learn how scratchpad factors into the security posture of your instance and how to manage it so that malicious scripts can't be executed on it.

The scratchpad is an easy way to set information on the server that you can access in the browser. An admin can script anything to be on it, including arbitrary records. If this property is not set to the recommended value of **true**, then it is possible to execute malicious scripts like a cross-site scripting vulnerability.

More information

Attribute	Description
Configuration name	*glide.ui.escape_scratchpad*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Validation, sanitization, and encoding
Security risk	• Severity score: 6.5 • CVSS score: Medium • Security risk details: If the property is not set to the recommended value of **true**, then it is possible to execute malicious scripts like a cross-site scripting vulnerability.
Dependencies and prerequisites	None
References	Workflow administration

Escape XML markup [Updated in Security Center 1.3]

Use the *glide.ui.escape_text* property to force escape of XML values at the parser level before transmitting them to the client's browser.

Cross-site scripting occurs when an attacker injects malicious JavaScript into an entry point. The platform/application fails to escape the malicious JavaScript before transmitting it to the victim's browser for execution. Escaping in this context means the following:

- **&** --> `&`
- **<** --> `<`
- **>** --> `>`
- **"** --> `"`
- **'** --> `'`
- **/** --> `/`

Example: `<script>alert('XSS Attack');</script>`

Escaping: `<script>alert('XSS Attack');</script>`

Ensure the *glide.ui.escape_text* property exists in the sys_properties table and is set to true.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.ui.escape_text*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	Escaping XML ensures that browsers do not parse the malicious JavaScript embedded in untrusted data, and execute it as JavaScript. • A malicious user may try XSS attack to either hijack other users' session or redirect the user to a malicious website. • The NOW Platform contains code to secure cookies, but escaping it relies on this property being set to **true**.
Recommended value	true
Security risk rating	8.8
Functional impact	This remediation enforces XML encoding at the XML parser level on the UI. It renders the encoded results for the user, which can have a functionality impact based on the instance user interaction with the resulted data.

Attribute	Description
Security risk	(High) Input validation must occur on the application to defend against cross-site scripting attacks. These attacks enable foreign scripts to execute on user session in the logged in browser's context. Attackers can use it to steal session information and sensitive data.
Workaround	After you set this property to **true**, rendering stops on the HTML tags in the catalog item description or in the catalog item variable help text. You may not be able to use HTML formatting for some fields. However, if the *glide.ui.escape_text* property is turned of, all JEXL expressions would be prefixed with an output encoder: `${JS:expression}` `${HTML:expression}` or `${JS,HTML:expression}`

Escape xml response

Manage how XML escapes are handled on your instance.

Use this property to manage if XML responses are escaped. If the property is set to the recommended value of **false**, then XML responses will not be escaped which can lead to XML injection attack. The injection of unintended XML content into an XML message can alter the intended logic of an application.

More information

Attribute	Description
Configuration name	*glide.soaprequest.unescape_xml_response*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	false
Default value	false
Category	Validation, sanitization, and encoding
Security risk	• Severity score: 6.4 • CVSS score: Medium • Security risk details: Setting this property to **false** disables XML escaping, which could lead to XML injection attack.
Dependencies and prerequisites	None

Enable HTML Sanitizer [Updated in Security Center 1.3]

Use the *glide.html.sanitize_all_fields* property to enable the HTMLSanitizer script include, which sanitizes HTML input based on exclusion listed and inclusion listed attributes configured in a script.

The field types available with dictionary/fields include HTML and Translated HTML. These HTML input fields enable users to write HTML formatted input, for example:

`<h1>Test</h1>`, using the most basic HTML tags such as `<img>`, `<a href ...>`, and `<iframe>`.

It can open a door for a malicious attacker to inject malicious vector with HTML tags such as:

`[<IMG SRC=" &#14; JavaScript:alert('XSS');">][<IMG onmouseover="alert('xss')">],[a href="" onclick=alert(/xss/)]`.

More information

Attribute	Description
Property name	*glide.html.sanitize_all_fields*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	Prevents application against cross site scripting and HTML injection attacks
Recommended value	true
Default value	true
Security risk rating	8.8
Functional impact	This remediation enforces HTML-output encoding mechanism before the user data is rendered back to the user. If customer has any customization that involves rendering of the HTML attribute or content data, then there is a functionality impact.
Security risk	(High) User input should be securely treated when the data is being stored and processed on the application.This reduces client-side cross-site scripting attacks by output encoding the data.
Workaround	This property sanitizes all HTML fields in the system. If you must enable HTML sanitization on individual fields, see Enable sanitization on individual fields. You can also configure the inclusion list or exclusion list to sanitize HTML tags and attributes as per your organizations policy.
References	Enabling HTML sanitizer HTML sanitizer

To learn more about adding or creating a system property, see Add a system property.

Restrict allowed Java packages [Updated in Security Center 1.3]

Configuring these properties protect from dangerous APIs being exposed to the scripting engine.

Configure the system tables and install the recommended plugin accordingly.

If the *sys_whitelist_member* and *sys_whitelist_package* table are not empty values, then dangerous APIs may be exposed to the scripting engine. Entries correspond to the Java namespace that have not been approved by ServiceNow security teams.

Install the Packages call removal tool. See Packages call removal tool for details.

Contact Customer Service and Support to edit these tables.

More information

Attribute	Description
Table, plugin name	Tables: • *sys_whitelist_member* • *sys_whitelist_package* Plugin: com.glide.script.packages_call_removal
Configuration type	Tabular Configuration, Plugins
Category	Validation, sanitization, and encoding
Purpose	Protect from dangerous APIs being exposed to the scripting engine.
Recommended value	Empty
Default value	None. This is a table configuration and not a Glide Property, so there is no default value.
Configuration type	table list, plugin
Security risk	(High) Dangerous APIs may be exposed to the scripting engine. These supported APIs will likely introduce instability and insecurity within the instance.
Security risk rating	8.2

To learn more about adding or creating a system property, see Add a system property.

Packages call removal tool

Activate and run the Packages Call Removal Tool (*com.glide.script.packages_call_removal*) plugin, and then consider whether each of the proposed changes should be completed or rejected.

The Packages Call Removal Tool is a plugin that:

- Scans scripts for package calls to ServiceNow AI Platform Java classes.
- Proposes changes to replace them with preferred GlideScriptable names.
- Facilitates the script changes.

Note: If this record is a base system record, using the recommendation from the tool causes the item to be marked as customer_update. However, it may still be useful to use this tool because it flags Packages,xxx calls.

The Packages Call Removal Tool might report some package calls used in `sa_mapping_ext_commands` and `sa_custom_operation`. These package calls belong to the MID Server. As there are no classes, the code runs in MID Server. If you find any of the following listed package calls in the Errors section, mark them as Rejected (Ignored). The tool doesn't report that package call again.

- `Packages.com.snc.sw.util.JSONUtil.toJSONPlain(file_content);`
- `Packages.com.snc.sw.util.JSONUtil.toJSONPlain(file_name);`
- `Packages.com.snc.sw.commands.HttpCallHandler;`
- `Packages.com.snc.sw.dto.ProviderType.SSH`

More information

Attribute	Description
Plugin Name	com.glide.script.packages_call_removal
Configuration type	System Definition > Plugins
Purpose	To remove/replace unauthorized package/member calls with Glide Acceptable (GlideScriptable) names that only allow authorized access to data.
Recommended value	Active
Functional impact	This remediation would replace the package calls with *`GlideScriptable`* APIs, and can affect the customizations that include package calls. The tool doesn't actually replace package calls automatically. Instead, it provides suggestions that are stored into the packages_call_item table. Your administrator can then decide whether to accept or reject the proposed change.
Security risk	(Medium) Client-side API calls that result in data retrieval or object access on server are deemed to be dangerous from a security standpoint. They should be validated for authorization and restriction for sensitive object access.

Steps to configure

1. Navigate to **System Definition** > **Plugins**

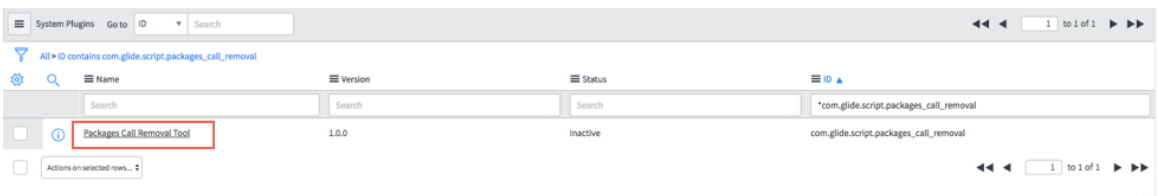

2. Search for the plugin ID = *`com.glide.script.packages_call_removal`*.

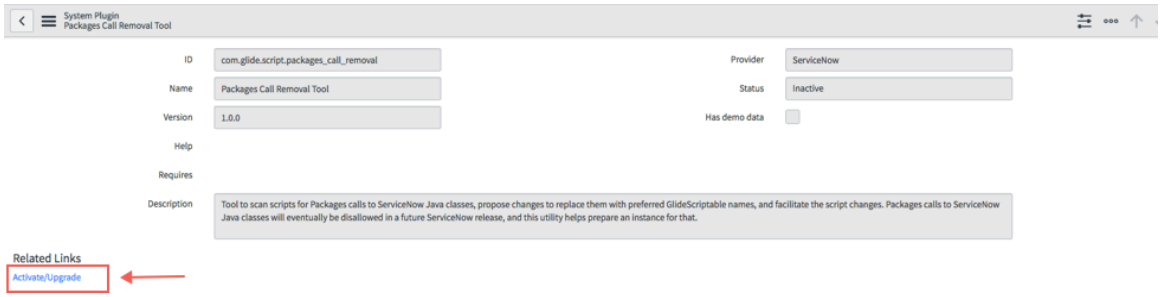

3. Click **Activate/Upgrade** to activate the plugin.

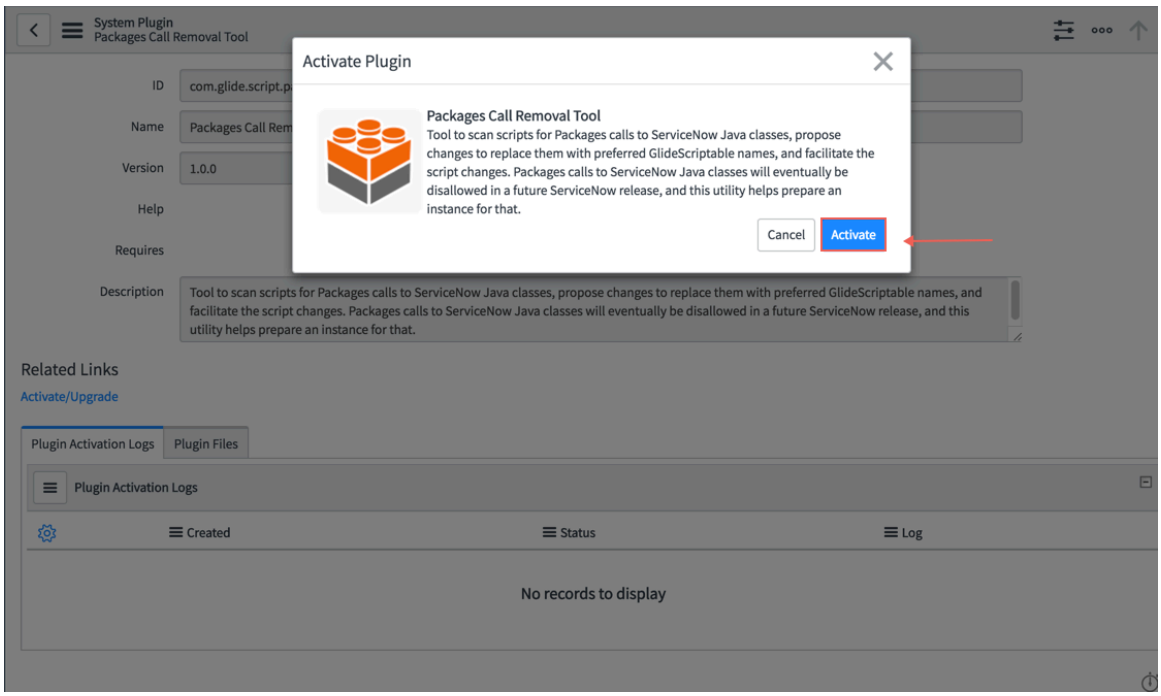

4. To check inclusion list package calls and inclusion list member calls, complete the actions outlined in the Steps to Configure sections in Restrict allowed Java packages [Updated in Security Center 1.3].

Unset LDAP Initial distinguished name [Updated in Security Center 1.3 and removed in 2.0]

Use this property to manage the distinguished name of a LDAP Server record.

This property controls the distinguished name of a LDAP Server record which is inserted when running an out-of-the-box (OOB) fix script. If it is set to the recommended value of "" or blank, then LDAP server data can be enumerated by a lower privilege user.

More information

Attribute	Description
Configuration name	*glide.ldap.initial.dn*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string

Attribute	Description
Recommended value	blank
Default value	blank
Category	Validation, sanitization, and encoding
Security risk	• Severity score: 2.7 • CVSS score: Low • Security risk details: Setting the property value to "" or blank could make LDAP server data accessible to a lower privilege user.
Dependencies and prerequisites	None

Enforce strict security of session cookies [Updated in Security Center 1.3]

Use the *glide.ui.secure_cookies* property to require properly formatted cookies

When you set the property is to true, your instance will reject a session if the associated cookie is not in the expected format.

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.ui.secure_cookies*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	To achieve more secure session authentication.
Recommended value	true
Default value	true
Security risk rating	8.8
Functional impact	When the property is set to true, improperly formatted cookies are rejected. When such a cookie is rejected, the user must login again.

To learn more about adding or creating a system property, see Add a system property.

Minimize Entity Expansion Threshold for GlideXMLUtil Scriptable [Updated in Security Center 1.3, 1.5, and 2.0]

Use the *glide.xmlutil.max_entity_expansion* property to change the maximum entity expansion limit to a smaller number.

This property controls the maximum amount of entity expansion within an XML Parser. If *glide.xmlutil.max_entity_expansion* is not set to the recommended value of 3000 or less, then the GlideXMLUtil parsing scriptable may be vulnerable to denial of service attacks.

Ensure the property *glide.xmlutil.max_entity_expansion* is set to 3000 or less. If the instance is on Washington or later, the default implied value is 3000 if the sys_properties record does not exist. If the instance is not on Washington or later, the recommendaiton is for the instance admin to create a sys_properties record with name *glide.xmlutil.max_entity_expansion* and the value 3000.

Note: 500 is the default minimum imposed by the ServiceNow AI Platform, which is considered to be a safe threshold.

More information

Attribute	Description
Property name	*glide.xmlutil.max_entity_expansion*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	This remediation control must be enabled to defend against XML Entity Expansion/Billion Laugh attack.
Recommended value	3000
Default value	100000
Security risk rating	5.3
Functional impact	If the customization is using large entity expansion, then, the ServiceNow AI Platform might block further processing.
Security risk	(Moderate) An attacker can use this vulnerability to expand data exponentially, quickly consuming all system resources.

To learn more about adding or creating a system property, see Add a system property.

Prevent Empty ACL Creation [New in Security Center 2.0]

Set the *glide.security.empty_acl.popup_window.enabled* property to the secure value of true to block attempts to create, update, or save an invalid ACL. This setting will also provide a client-side model to configure a role or security attribute for the ACL.

The *glide.security.empty_acl.popup_window.enabled* property determines whether users making form-based edits to access control lists (ACLs), specifically sys_security_acl, can create, update, or save an invalid ACL that has an invalid data condition, script, security attribute, or roles list. Otherwise, it remains unconfigured (an empty ACL). As of the Xanadu release, any empty ACL will deny access. In ServiceNow versions prior to Xanadu, an empty ACL will permit unconditional access.

When the *glide.security.empty_acl.popup_window.enabled* property is set to the secure value of true, it blocks attempts to create, update, or save an invalid or empty ACL, and provides a client-side model to configure a role or security attribute for the ACL. If the property is set to the unsecure value of false, then such attempts will be permitted, and no client-side model will be displayed.

Note: This property is case-sensitive. For example, a value of True (capital "T") will be evaluated as false. Moreover, this property only functions when the High Security (com.glide.high_security) plugin is installed and active.

More information

Attribute	Description
Configuration name	*glide.security.empty_acl.popup_window.enabled*
Configuration type	System Properties (/sys_properties_list.do)
Data type	string
Recommended value	true
Default value	true
Category	Validation, sanitization, and encoding
Security risk	• Severity score: 6.5 • CVSS score: Medium • Security risk details: If this property is set to true, the empty ACL warning popup will prevent the user from submitting an empty ACL on the client side. If it is set to false, the popup will no longer show.
Dependencies and prerequisites	None
Functional impact	This property allows the user to toggle the empty ACL warning popup on and off.
References	Prevent Empty ACL Creation [New in Security Center 2.0]

Define restricted downloadable MIME types [Updated in Security Center 1.3, 1.5, and 2.0]

Use the *glide.ui.attachment.force_download_all_mime_types* property to download MIME types and not to render inline in the browser.

If *glide.ui.attachment.download_mime_types* does include dangerous MIME types such as text/html, image/svg ,image/svg+xml,application/xml, then dangerous files could be rendered inline in the browser, which could lead to Cross Site Scripting attacks (XSS). This property is the list of comma-separated attachment mime types, which won't render inline in the browser. For example, including text/html forces HTML files to be downloaded to the client as attachments rather than viewed inline in the browser. Maintaining this list properly prevents cross-site scripting attacks.

If the *glide.ui.attachment.download_mime_types* system property doesn't include dangerous MIME types such as "text/html, image/svg,image/svg +xml,application/xml", then dangerous files could be rendered inline in the browser. This can lead to Cross Site Scripting (XSS) attacks. This check is only relevant when *glide.ui.attachment.force_download_all_mime_types* is set to **false**.

This property is a list of comma-separated attachment MIME types, which don't render inline in the browser. For example, including `text/html` forces HTML files to be downloaded to the client as attachments rather than viewed inline in the browser.

If *glide.ui.attachment.force_download_all_mime_types* is set to **false**, verify that the *glide.ui.attachment.download_mime_types* system property includes the dangerous MIME types `text/html`, `image/svg`, `image/svg+xml`, `application/xml`.

More information

Attribute	Description
Configuration name	*glide.ui.attachment.force_download_all_mime_types*
Configuration type	System Properties (/sys_properties_list.do)
Data type	String (Comma-separated list of MIME types)
Recommended value	text/html,image/svg,image/svg+xml,application/xml
Default value	text/html,image/svg,image/svg+xml,application/xml
Fallback value	text/html,image/svg,image/svg+xml,application/xml
Category	Validation, sanitization, and encoding
Security risk	• Severity score: 6.3 • CVSS score: Medium • Security Risk: Maintaining this list properly can prevent cross site scripting attacks.
Dependencies and prerequisites	This check is only relevant when *glide.ui.attachment.download_mime_types* is set to **false** or doesn't exist in the System Properties [sys_properties] table.

Restrict uploaded MIME types [Updated in Security Center 1.3 and 2.0]

Use the *glide.security.file.mime_type.validation* property to activate MIME type checking for uploads. You can enable (set the property to **true**) or disable (set it to **false**) MIME type validation for file attachments.

Prerequisites

Before setting this property, set the *glide.attachment.extensions* property. Only those extensions specified in *glide.attachment.extensions* are checked for MIME type during upload. To learn more, see Restrict file extensions.

Ensure the Glide Property *glide.ui.jelly.js_interpolation.protect_nested_expressions* exists and is set to the value true. If the property does not appear in the sys_properties table, add a new record.

More information

Attribute	Description
Property name	*glide.security.file.mime_type.validation*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	To enforce checking of MIME type / magic bytes during file uploads.

Attribute	Description
Recommended value	true
Default value	true
Security risk rating	5.4
Functional impact	This remediation enables MIME type verification on the attachments to the application. No functionality impact, unless there is a malicious intent in uploading the files as this validation is merely checking for mis-sync between the MIME type and the data.
Security risk	(Medium) To reduce vulnerabilities such as file inclusion and malicious file uploads, MIME type verification should be enabled.
References	Administering attachments

See Hardening settings for details on configuring properties for hardening.

To learn more about adding or creating a system property, see Add a system property.

Restrict XML external entities [Updated in Security Center 1.3 and 2.0]

Configure system properties to ensure that your instance only processes XML from trusted sources to help prevent XML external entity (XXE) attacks.

Use the *glide.xml.entity.whitelist* and *glide.xml.entity.whitelist* system properties to prevent your instance from processing XML from untrusted sources.

XML external entity (XXE) attacks occur when a malicious actor modifies incoming XML (such as adding HTTP requests) to access data or intact with otherwise restricted systems. To help prevent these attacks, the *glide.xml.entity.whitelist.enabled* system property limits the sources from which your instance executes XML. Use the *glide.xml.entity.whitelist* property to define a set of trusted sources.

Ensure that the *glide.xml.entity.whitelist* system property exists in the System Properties [sys_properties] table, and is set to `http://java.sun.com/j2ee/dtds/`. Ensure that the *glide.xml.entity.whitelist.enabled* system property exists in the System Properties [sys_properties] table and is set to the value `true`.

Tip:

Values other than `http://java.sun.com/j2ee/dtds/` can be included in the *glide.xml.entity.whitelist* property, but are unnecessary for the out of the box platform state. Review any additional values to determine if they are safe.

Warning: s a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Configuration name	• *glide.xml.entity.whitelist* • *glide.xml.entity.whitelist.enabled*
Configuration type	System Properties (/sys_properties_list.do)
Data type	• String • Boolean String
Recommended value	• http://java.sun.com/j2ee/dtds/ • true
Default value	• http://java.sun.com/j2ee/dtds/ • true
Fallback value	• http://java.sun.com/j2ee/dtds/ • true
Category	Validation, sanitization, and encoding
Security risk	• Severity score: 9.8 • CVSS score: Critical • Security risk details: An XML Eternal Entity (XEE) attack can allow attackers to access data or perform unauthorized actions via crafted XML payloads.
Functional impact	If the customization is using external entity, not inclusion listed in the *glide.xml.entity.whitelist* property, the NOW Platform might block further processing.
Dependencies and prerequisites	None

Require XMLdoc2 entity validation with allowlistDisable entity expansion [Updated in Security Center 1.3]

If customizations do not require entity expansion, use the *glide.xmlutil.max_entity_expansion* property to completely disable external entity expansion. The XML completes parsing but doesn't include any internal or external entities.

If the glide property *glide.stax.whitelist_enabled* doesn't exist in the System Properties [sys_properties] table, or is not set to the recommended value of **true**, then all external entities are allowed when the glide property *glide.stax.allow_entity_resolution* is set to the value of **true**.

If customizations don't require entity expansion, use the *glide.stax.allow_entity_resolution* property to completely disable external entity expansion. The XML completes parsing but doesn't include any internal or external entities.

- If you set *glide.stax.allow_entity_resolution* to **true**, all external entities attempt to resolve or expand subject entities, subject to the setting of the *glide.stax.whitelist_enabled* property.
- If you set *glide.stax.allow_entity_resolution* to **false**, all entity resolution and expansion is blocked. To learn more about this property, see Disable Entity Expansion within the XMLDocument2 Streaming Parser [Updated in Security Center 1.5].

When **glide.stax.whitelist_enabled** is set to **true**, define a listing of comma-delimited FQDN in the **glide.xml.entity.whitelist** property, which are the only URLs that can be reached using the XML entity processing property. To learn more, see Restrict XML external entities [Updated in Security Center 1.3 and 2.0]. Attackers can use this vulnerability to expand data exponentially in an External Entities Expansion (XXE) attack, quickly consuming all system resources.

Prerequisites

Before setting this property:

- Set the *glide.xml.entity.whitelist.enabled* and *glide.stax.whitelist_enabled* properties to **true**. To learn more, see Restrict XML external entities [Updated in Security Center 1.3 and 2.0].
- Define a listing of comma-delimited FQDN in the *glide.xml.entity.whitelist* property, which is the only URLs that can be reached using XML Entity processing property. To learn more, see Restrict XML external entities [Updated in Security Center 1.3 and 2.0].

Warning: This is a safe harbor property, meaning the value can't be altered once it's changed. It is non-revertible.

More information

Attribute	Description
Property name	*glide.stax.whitelist_enabled*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	This remediation control must be enabled to defend against an XML Entity Expansion/Billion Laugh attack.
Recommended value	true
Default value	false
Security risk rating	9.8
Functional impact	If the customization is using entity expansion, then, the ServiceNow AI Platform might block further processing.
Security risk	An attacker can use this vulnerability to expand data exponentially in an External Entities Expansion (XXE) attack, quickly consuming all system resources.
Workaround	If the customization requires entity expansion, set this property to true and follow the steps documented in

Attribute	Description
	Restrict XML external entities [Updated in Security Center 1.3 and 2.0].

To learn more about adding or creating a system property, see Add a system property.

For more information about OWASp resources, see OWASp.

Sanitize All Translated HTML Fields [New in Security Center 2.0]

Learn how to configure the *glide.translated_html.sanitize_all_fields* property to the secure value to ensure that all *translated_html* elements are sanitized with an HTML sanitizer.

When the *glide.translated_html.sanitize_all_fields* property is set to true, all *translated_html* elements are sanitized with an HTML sanitizer. If the property is set to false, only elements with the dictionary attribute *html_sanitize* set to true will be sanitized. This sanitization helps prevent attackers from embedding malicious content that could lead to cross-site scripting (XSS) attacks.

More information

Attribute	Description
Configuration name	*glide.translated_html.sanitize_all_fields*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Validation, sanitization, and encoding
Security risk	• Severity score: 4.6 • CVSS score: Medium • Security risk details: Not configuring this property to the secure value of true, doesn't sanitize all HTML elements, which increases the likelihood of a bad actor embedding malicious content in a field.
Dependencies and prerequisites	None
Functional impact	Allows customers to access any table information if the widget is set to public and included in the property's value.

Sanitize HTML in the Description Fields of the Impact Workspace Module [New in Security Center 7.0]

Sanitize the HTML in the description fields by removing HTML tags that are sources of HTML injection attacks with the *sn_impact_common.blacklist_tags_HTML_injection* property.

The Impact Workspace module allows HTML in the following description fields:

- The customer_notes field of the sn_impact_common_capabilities_map and sn_impact_common_par_version_phase_app_mapping tables.
- The manual_description field of the sn_impact_common_manual_capability_description table.

When this system property contains a comma-separated list of HTML tags (for example, scripts), those tags and their contents are removed from the HTML portions of the listed fields. Removing these tags helps sanitize the HTML in the description fields by removing HTML tags that are sources of HTML injection attacks. If this property isn't set in the System Properties [sys_properties] table, the value defaults to a default list of denied HTML tags. If the property is empty, all HTML tags are allowed.

Use the *sn_impact_common.blacklist_tags_HTML_injection* provide a comma-separated list of HTML tags which are removed from the description fields for the Impact Workspace module. This removal helps to prevent HTML injection attacks. At minimum, this list should contain the contents of the default list. If the property isn't set in the System Properties [sys_properties] table, it defaults to the list `script,iframe,object,embed,form,onerror,onload,style,img,video,audio,source,b`

More information

Attribute	Description
Property name	*sn_impact_common.blacklist_tags_HTML_injection*
Configuration type	System Properties (/sys_properties_list.do)
Category	Validation, sanitization, and encoding
Purpose	Sanitize the HTML in the description fields by removing HTML tags that are sources of HTML injection attacks.
Recommended value	At minimum, the default value of `script,iframe,object,embed,form,onerror,onload,sty`
Default value	`script,iframe,object,embed,form,onerror,onload,sty`
Security risk rating	4.4
Functional impact	If an HTML tag is added to default list, it may limit the required HTML functionality of the description fields. The exact impact is dependent on the customer instance.
Security risk	(Medium)
References	High Security Settings

To learn more about adding or creating a system property, see Add a system property.

Set safe content security policy for svg files [New in Security Center 1.3]

The *com.glide.csp.self_script_src_svg* property adds the **script-src none** directive to the HTTP Content-Security-Policy header when Scalable Vector Graphics (SVGs) are accessed through the Translation Memory Index (IIX) file extension.

The *com.glide.csp.self_script_src_svg* property prevents malicious file attachments that stores cross site scripting (XSS) attacks from running in an instance. Without

this policy, a bad actor could cause a user to run arbitrary JavaScript code in their web browser which could lead to security vulnerabilities such as data exfiltration and session takeover.

More information

Attribute	Description
Configuration name	*com.glide.csp.self_script_src_svg*
Configuration type	System Properties (/sys_properties_list.do)
Data type	Boolean
Recommended value	true
Default value	true
Category	Validation, sanitization, and encoding
Security risk	• Severity score: 7.1 • CVSS score: High • Security risk details: Not setting this property to the recommended value of true could cause a user to run arbitrary JavaScript code from a bad actor.
Dependencies and prerequisites	None
Functional impact	This property prevents scalable vector graphics (SVG) files from accessing external scripts.

Log Export Service (LES)

Log Export Service (LES) lets you seamlessly export your instance system and application logs into your enterprise security analytic tools.

Get started

Explore	Configure	Administer
		
Learn about Log Export Service (LES)	Configure LES for Kafka and MID server consumers	Administering LES

Use	Reference	
 Review log reports using LES	 Learn about miscellaneous information about the app	

Troubleshoot and get help

- Ask or answer questions about Log Export Service (LES) in the ServiceNow Community
- Search the Known Error Portal for known error articles
- For cloning recommendations for LES, including tables to preserve and exclude, see Tables to be preserved to keep the LES configuration.

Exploring Log Export Service (LES)

The LES service provides a highly scalable and near real-time integration with your analytic tools that is easy to set up and maintain. If you're new to LES, read this overview section to learn what the tool can do.

Check your entitlements to determine whether you have access to Log Export Service.

Log Export Service overview

The integration tool allows you to leverage your analytic solutions to perform the following:

- Detect ServiceNow security threats and analyze security incidents
- Troubleshoot and optimize ServiceNow app performance
- Monitor and optimize ServiceNow user experience

LES leverages a ServiceNow AI Platform capability called the Hermes Messaging Service, which is a multi-tenant, multi-cluster, data transport, and queuing service built on Apache Kafka that enables your instance to produce and consume large volumes of Kafka events. Apache Kafka is an open-source data streaming platform that provides a single integration point for exchanging data across business systems in your organization.

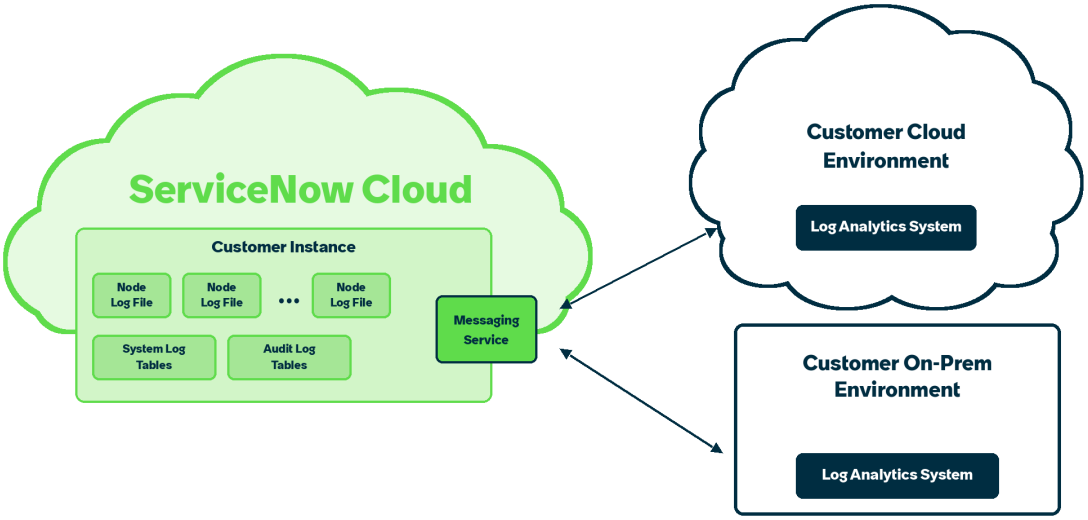

LES forwards a copy of the log events as they're generated to the Hermes Messaging Service.

The Hermes Messaging Service is a multi-tenant, multi-cluster, data transport, and queuing service built on Apache Kafka that enables your instance to produce and consume large volumes of Kafka events. The Hermes Messaging Service is a ServiceNow AI Platform capability that is available as part of Stream Connect, Log Export Service (LES), and Instance Data Replication (IDR).

The external log analytic systems, either in the cloud or on-prem, can use and consume the log events from the Hermes Messaging Service. LES provides three connectivity options to consume the logs:

- Dedicated MID Server: A dedicated MID Server is installed on-prem or in the cloud that automatically connects to Hermes Messaging Service, pulls log events from it continuously and then pushes them to log analytic tools via a REST connection.
- Leverage Kafka connector from your log analytic solution (for example, Splunk): A Kafka connector from your log analytics product of choice is installed on-prem or in the cloud that automatically connects to Hermes Messaging Service, pulls log events from it continuously and then pushes them to log analytics tools.
- Directly from your Kafka system: Your Kafka system connect directly with the Hermes Messaging Service and use its native Kafka protocol commands and connectivity to pull logs events from it.

Note: If your Kafka message exceeds the configured memory buffer, Hermes may return an error indicating that it's larger than the total memory buffer you've configured.

To configure and manage LES you need to install it from ServiceNow Store. The LES application provides Guided Setups to help you install the service, pages to configure the service (log sources, consumers and destinations) and reports to understand log creation and consumption.

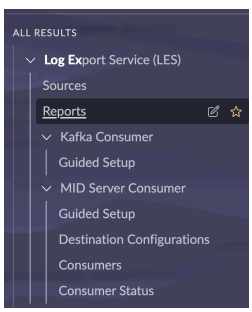

Note: You can also create a new source configuration. See Create a log source configuration for more information.

Log Export Service users

Log Export Service has the following users.

Users	Description
Application admin [sn_logstoanalytics.admin]	This role is installed along with the LES application and allows a non-admin to use the application.
System administrator [admin]	Admin role is required for the setup of the LES store application.

Log Export Service benefits

Benefit	Feature	Users
Create log source configuration to set filters on the logs	Create a log source configuration	Application admin
Experience guided setup for Kafka consumers	Guided setup for Kafka consumers	System administrator
Experience guided setup for MID server consumers	Guided setup for MID Server consumers	System administrator
Examine the log report dashboard to analyze the size of each data log	Review log report	System administrator or Application admin

What to explore next

To learn more about using Log Export Service, see:

- Administering Log Export Service (LES)
- Configuring Log Export Service (LES)
- Using Log Export Service (LES)
- Log Export Service (LES) references

Log sources

Log Export Service (LES) can export log sources from some System Log Tables, the Audit Table, and Application Node Log Files.

The following are the log sources that can be exported by LES.

- System Log Tables
  - syslog table: View warnings and errors for instance processes, records, and non-critical events, such as memory usage on the server machine
  - syslog_transaction table: view all browser activity for an instance
  - sys_outbound_http_log table: view all requests and responses for outbound web services such as REST and SOAP

- Audit Table: Use the sys_audit table view record changes made to tables chosen to be audited
- Application Node Log Files: Use the localhost log files to view application node errors. Your instance will have multiple nodes and each node will have multiple log files.

See System logs to learn more about the schema and purpose for the above log sources.

Note: LES forwards logs exclusively when they are generated through the logging framework. The following are the two sources for which logs are not getting forwarded:

Source name	Reason
datacollector	Data is inserted directly into the syslog child table, which causes it to appear in the syslog table as well. However, this bypasses the logging framework, so the records reach the syslog table but are not forwarded to the topic.
com.glide.ui.ServletErrorListener	Logs are inserted into the syslog table using the NoBroadcast logging variant, where the broadcast flag is set to false. In this mode, logs are written to syslog but are not forwarded to external consumers because they do not pass through the logging framework.

Administering Log Export Service (LES)

Use LES to create log source configuration and multi-topics for each source type.

Create a log source configuration

Regulate and set filters on the logs to be forwarded by creating a log source configuration.

Before you begin

Role required: admin or sn_logstoanalytics.admin

Procedure

1. Navigate to **All** > **Log Export Service**.
   A list of source configurations shows up.
2. Select **New** if you want to create a new source configuration.
   You can also select an existing source configuration if you want to modify it.
   The Source form shows up.
3. On the form, fill up the fields.

Source form

Fields	Description
Source Type	Types of log sources ○ Node Log ○ Table See Log sources for more information.

Fields	Description
Log Level	A set of standard logging levels that can be used to control logging output. Following convention, each level will forward logs of equal or greater severity. **Note:** This field is visible only when one of the following conditions is met. ○ When you select Node Log as the Source Type ○ When you select Table as the Source Type and the table is syslog
Accepts	Specifies the format in which the logs are forwarded to Hermes. They can either be sent as JSON or as plain text.
Table	Selection of table for table type logs. **Note:** This field is visible only when you select Table as the Source Type.
Filter Type	Conditions to forward logs selectively. **Note:** This field is visible only if you select either syslog or sys_audit as the table.

4. Review the Source Topics related list details.
   You can review each of the log tables and can create its own topics.

   Note: The related list is visible only if you select Log Table in the Filter Type field.

   The topic name is not auto-populated, and you can either select or create its own topics.

5. Select a topic for the log table.

   You can either select an existing topic or can create a new topic for a log table.

   a. Select **New** to create a new source topic for a log table. The Source Topics form shows up.
   b. Select the required table in the Log Table field.
   c. Select the lookup icon in the Topic field.

      Note: You can select an existing topic from the list. You can also create a new topic by selecting **New** in the Kafka Topics list. See Create source type and multi topics in the LES source table to create a new Kafka topic.

   d. Select **Submit** on the Source Topics form.

6. View the recently created log table and its corresponding topic in the Source Topics related list.
7. Select **Submit** to create a new source configuration.

Create source type and multi topics in the LES source table

Consume logs for each source type by creating multiple topics per source type. You can now leverage the option of customized selection of specific topics for different log sources during the debugging process, without impacting the other log tables.

Before you begin

Note: If you delete a Log Source, the corresponding topic in the sys_kafka_topic table is not deleted. If you create the Log Source again, the existing topic can be reused, ensuring continuity and avoiding unnecessary topic recreation.

Role required: admin or sn_logstoanalytics.admin

Procedure

1. Navigate to **All** > **Log Export Service (LES)** > **Sources**.
2. Click **New** to create a new source.
   The Source form shows up. Previously, the **Topic** field was auto-filled once the **Source Type** is selected. Starting Yokohama, each source type can create their own topic. It doesn't populate by default on selecting the **Source Type** or the **Table** name. They can be created either directly from the Kafka topics or from the reference field.

   Note: Since sys_audit table has multiple log tables, no specific topic is shown in the Topic column of the Sources list. Previously, all the source types had the same topic. Now you can select different topics for different sources.

3. Select the lookup icon to go to the list of Kafka topics.
   The list of Kafka topics shows up.
4. **Optional:** Create the new Kafka topic with the following steps.

   Note: This step is applicable only if you want to create a new topic for a selected Source Type.

   a. Select **New** to create a new Kafka topic for the selected Source Type. The Kafka topic form shows up.
   b. On the form, fill up the fields.

Source form

Fields	Description
Name	Name of the topic you are creating
Application ID	Enter sn_logstoanalytics
Namespace	Enter Default Namespace
Partition	The partition field of a topic in Hermes refers to the partitions into which the topic's data is divided. It plays a key role in scalability and parallelism

   c. Select **Submit** to create the new Kafka topic.
5. Select the **Source Type** you want to create for a particular topic.
6. Select the required topic from the list for the selected **Source Type**.

7. Select **Submit** to create the source type with the particular selections.
   The source type shows up on the Sources list with the selected topic name and other information.

Update system property

Update the glide.les.disable_logs_forwarding system property within the Log Export Service application to control log forwarding during migration or database reseeding operations.

Before you begin
Role required: admin

About this task
This property controls whether log forwarding is active in Log Export Service. By default, the value is set to **false**, meaning logs are forwarded normally. During database reseeding or migration, historical log records can be replayed or reprocessed, which may result in duplicate log exports to downstream systems. To prevent duplicate logs from being sent, set this property to **true** to temporarily pause LES forwarding. After the migration is complete, revert the value to **false** to resume normal operation.

Procedure

1. Navigate to **All** > **Log Export Service** > **Sources**.
2. In the workspace filter navigator, enter *`sys_properties.list`* to access the system properties list.
   This opens the full list of system properties.
3. In the Name column, search for *`glide.les.disable_logs_forwarding`*.
   In case you do not find the property *`glide.les.disable_logs_forwarding`*, you can create a new system property by following these steps:

   a. Select **New** on the top-right corner of your screen to create new system property.

   b. In the **Name** field enter *`glide.les.disable_logs_forwarding`*.

   c. In the **Type** field, select **true | false**.
      Only the **Name** and **Type** fields are required. All other fields are optional and can be left empty.

4. Open the record.
5. The default value is **false**, which means log forwarding is active.
6. In the **Value** field, set the value based on your requirement:

Property	Value	When to use
`glide.les.disable_logs_forwarding`	**true**	During migration or DB reseeding, this disables log forwarding.
`glide.les.disable_logs_forwarding`	**false**	Default. After migration is complete, this re-enables log forwarding.

7. Select **Update** to save the change.

Result
When the property is set to **true**, Log Export Service stops forwarding logs, ensuring no data is sent during a migration or DB reseeding event. Once the operation is complete and the value is reverted to **false**, log forwarding resumes automatically. This configuration allows administrators to safely pause log forwarding during maintenance windows without affecting the underlying LES configuration.

⚠ **Warning:** After migration or database reseeding is complete, revert the property value to **false**. If you do not revert this setting, log forwarding stops permanently and may result in missing log data in downstream systems.

Configuring Log Export Service (LES)

Use guided setup to step through the initial configuration of LES. Guided setup assists you with planning the roll-out of the product and performing the basic configuration to go live.

Guided setup organizes configuration activities into categories. Each category provides information, such as planning guidance, pre-setup steps, and links to useful help content. Categories also provide a set of links to the pages in your instance where you perform the configuration. The guided setup process keeps track of what you have completed, so you can stop and start again where you left off.

Kafka consumer

Use guided setup to step through the initial configuration of LES for Kafka consumers.

Guided setup home page

The home page for guided setup contains an overview of configuration types for your guided setup. You can select your guided setup type and select **Continue** to open the guided setup steps and begin configuration.

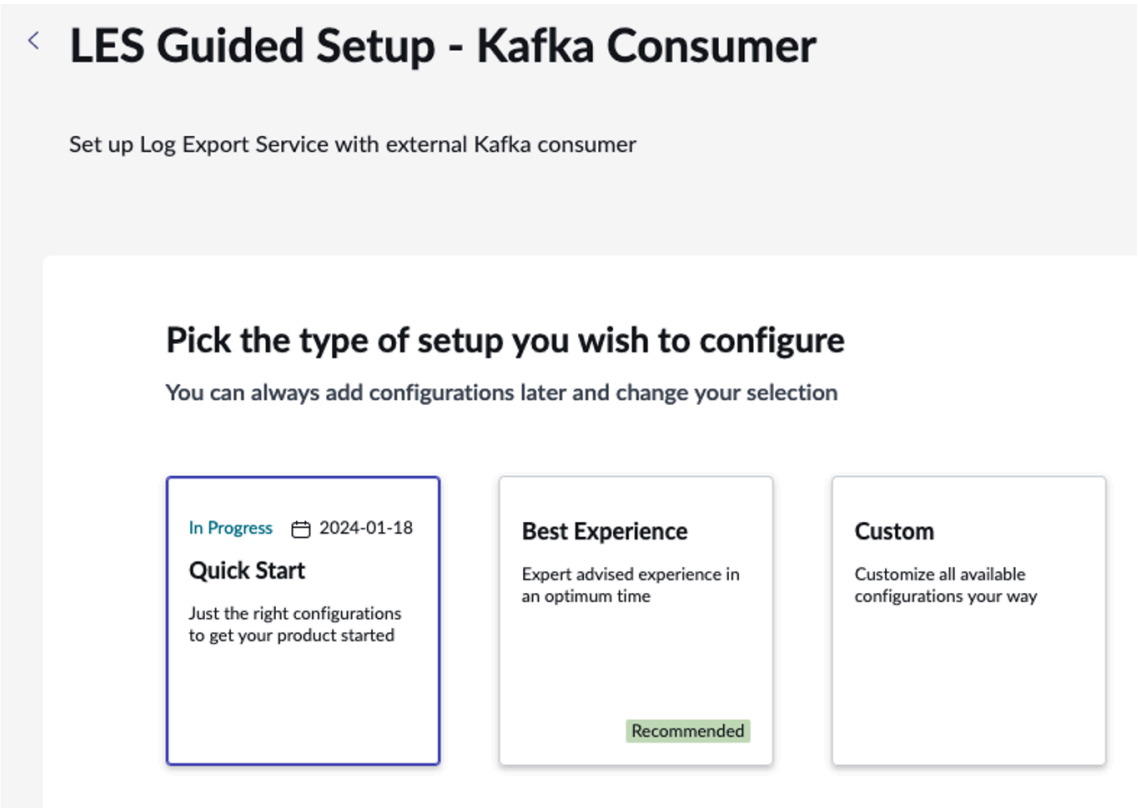

All three setups on this page, Quick Start/Best Experience/Custom provide the same tasks and functionality. It is required that to coordinate the integration between the ServiceNow instance and the destination log analytic tool with the respective administrator. The log analytic administrator must configure their tool to connect to the ServiceNow instance securely. It is recommended to share the Set up a secure connection to the Hermes Messaging Service document with the log analytic administrator ahead of time. Even if you have marked a task as completed, you can go back and uncheck it back to in progress. To do so, first click on the **Edit** box in the upper right corner for the category. Then click on the **Edit** box for the task you want to uncheck. The **Mark as complete** box will no longer be marked.

Guided setup categories page

The categories page contains an overview and descriptions of the categories and associated tasks. You can click either the drop down arrow to view information about the category or the Start button to open the guided setup steps and begin configuration.

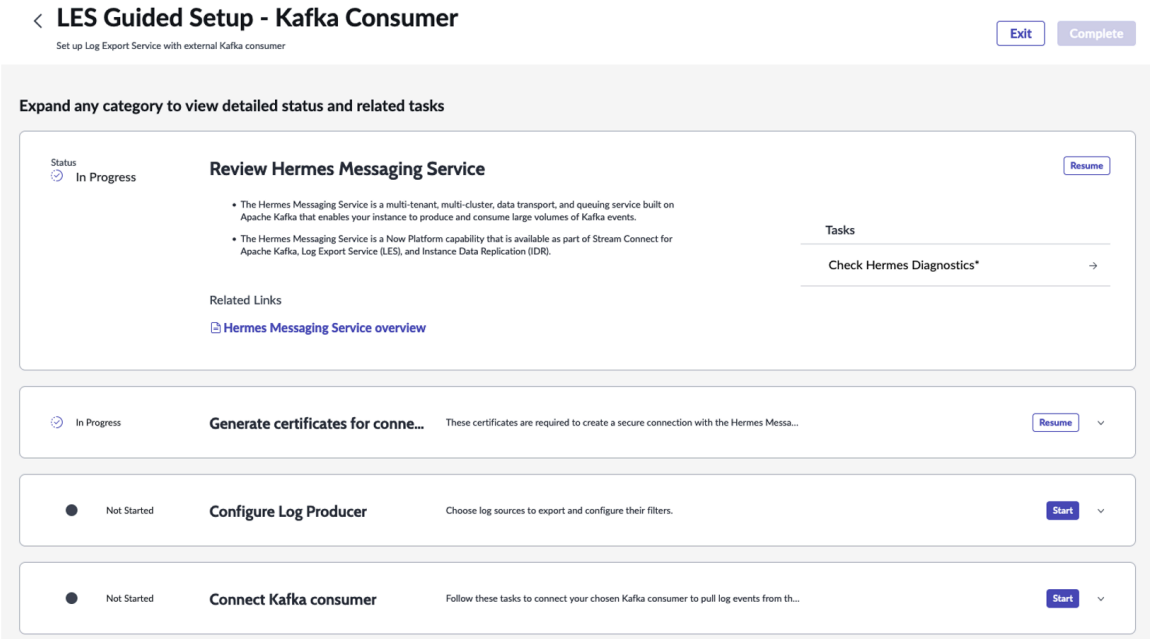

Complete the tasks under each category by following the setup instructions.

Guided setup for Kafka consumers

Implement the following steps for a complete guided setup for Kafka consumers.

Before you begin

Navigate to **Log Export Service (LES)** > **Kafka Consumer** > **Guided Setup**. Select the type of setup you wish to configure and select **Continue**.

Note: During the Log Export Service application installation, ServiceNow will provision the underlying Hermes Messaging Service infrastructure. Be aware that this process can take up to a couple of hours to complete from the time you request the Log Export Service application installation.

Role required: admin

Procedure

1. Review Hermes Messaging Service Diagnostics.

   It is recommended that you verify that the Hermes Messaging Service is up and running with the Hermes diagnostic tool, which displays on screen during this step. If you see a "Page not found" error on this page, Hermes is not installed and you should contact your system admin.

   - Setup information: The following bootstrap information is used to connect to the Hermes Messaging Service. The "Producer Bootstrap" is the connection used to send messages into Hermes and "Consumer Bootstrap 1 & 2" are used to retrieve messages from Hermes.
     - Producer Bootstrap
     - Consumer Bootstrap 1
     - Consumer Bootstrap 2
   - Instance PKI: The Instance Public Key Infrastructure (PKI) component allows a ServiceNow instance to act as an issuer in a X.509 trust hierarchy.
   - Bootstrap Connectivity: Select **Run Test** to confirm external client is able to connect to the defined instance ports (producer and consumer).

     Note: Work with your network administrator to ensure that the following port ranges are open before you begin:
     - Consumer1: 4100-4150
     - Consumer2: 4200-4250
   - Instance Connectivity: Select **Run Test** to confirm the instance is able to send and receive messages.
   - View Topics: Select the listed topic to retrieve the timestamp of the last known message.

   Note: You can access Hermes Diagnostics in the future to troubleshoot potential connectivity issues by returning to this step of the guided setup or by navigating to **All** > *Hermes Messaging Service* > **Diagnostics**.

2. Generate certificates for a secure connection to Hermes Messaging Service and pull log events from it.

   You are going to use these certificates when connecting your external system.

   Setup secure connection to Hermes Messaging Service. See Set up a secure connection to the Hermes Messaging Service for LES for more information. You will need these certificates for authentication and authorization in the client which will pull the logs from Hermes.

   Note: admin or Hermes_admin roles are required for this step.

3. Configure Log Producer: Choose log sources to export and configure their filters. Complete the following tasks to configure the Log Producer.
   - Configure log sources to export: Create one Source record for each of the log sources that you want to export.

     Note: admin or sn_logstoanalytics.admin roles are required to complete this step

     a. Select **New** at the top right corner

     b. Select a Source Type

c. Select a Table

d. Select a Log Level that can be used to control logging output

e. Select or create a topic to which that log source will be exported. If you are creating a new topic, fill out the following fields:

- **Name**: Name of the topic you're creating
- **Application ID**: Enter `sn_logstoanalytics`
- **Namespace**: Enter Default Namespace
- **Partition**: The partition field of a topic in Hermes refers to the partitions into which the topic's data is divided. It plays a key role in scalability and parallelism.

f. Select a **Filter Type** condition to forward logs selectively.

> **Note:** Filters are different depending on the selected source type

g. Select Update
When successfully created, it will display the name of the Hermes topic to which that log source will be exported to. Write down the topic name, you will need it later when configuring your log consumer system.

The Active field controls whether or not that log source is going to be exported or not. If you see errors, go back to the Check Hermes Diagnostics task and verify Hermes status.

- Validate Log Producer: Once you have created a Source to produce logs from, you can see live log records in the topic using **Hermes Messaging Service** > **Hermes Topic Inspector**.

a. Select External Topics

b. Select List Topics

c. Select row with your topic from previous step (listed in Sources)

d. Adjust message start date if necessary

e. Select **View** to see a log message that was exported to the topic

4. Connect Kafka consumer: Follow these tasks to connect your chosen Kafka consumer to pull log events from the Hermes.
   - Identify Kafka consumer: You have two options based on your log analytics architecture.
     - If you have your own Kafka system and choose that for log aggregation, you can connect directly to the Hermes Messaging Service via the native Kafka protocol.
     - If you choose to have your log analytics tool connect directly to the Hermes Messaging Service then you need to deploy a Kafka connector supported by your log analytics system (i.e. Splunk Connect for Kafka).

     > **Note:** In either case, you will need to work with the Administrator for those systems to coordinate the connection with the Hermes Messaging Service.

   - Import Hermes certificates to Kafka consumer system: Log into your Kafka consumer system and make sure you have appropriate admin entitlements to configure it and connect it to an external system. Import the certificates generated in "Set up secure connection to Hermes Messaging Service" task into your Kafka connector or Kafka server. Follow the instructions in the documentation for your chosen Kafka consumer.
   - Configure Kafka processes: The Hermes Messaging Service is designed for high availability. Two processes are required to consume messages from Hermes. Two processes are

required because Hermes uses a pair of Kafka clusters for failover purposes. If one cluster goes down, data is produced to the other Hermes Kafka cluster.

In your Kafka consumer system, you will need to create two separate consumer processes to connect to both Hermes Kafka clusters. For both processes you will specify the same Hermes Kafka topic but you will need to configure two separate bootstrap addresses:

- <instance_name>.service-now.com:4100,<instance_name>.service-now.com:4101,<instance_name>.service-now.com:4102,<instance_name>.service-now.com:4103
- <instance_name>.service-now.com:4200,<instance_name>.service-now.com:4201,<instance_name>.service-now.com:4202,<instance_name>.service-now.com:4203

Important notes:

- When accessing the Kafka topic from external systems, you must prepend “snc.<instance name>.” to the topic that the logs are being forwarded to.
- Configure each consumer with the same Kafka Consumer Group ID.
- Install your keystore and truststore files in a location where your consumers can access them.
- If your consumers require it, specify the Kafka JSON Converters properties to disable schemas: “key.converter.schemas.enable=false”, “value.converter.schemas.enable=false”

- Verify Kafka consumer pulling logs from Hermes: Verify in your chosen Kakfa consumer that you can pull log events from the Hermes Messaging Service.

MID server consumer

Use guided setup to step through the initial configuration of LES for MID server consumer.

Guided setup home page

The home page for guided setup contains an overview of configuration types for your guided setup. You can select your guided setup type and select **Continue** to open the guided setup steps and begin configuration.

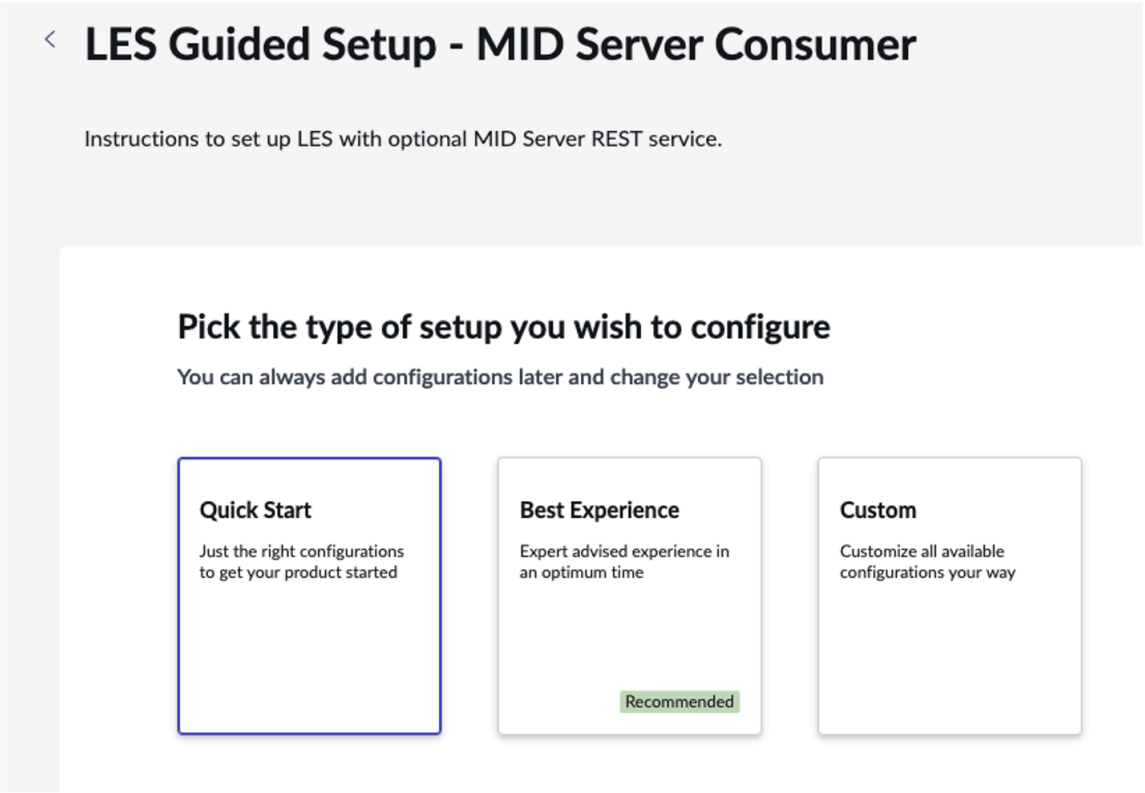

All three setups on this page, Quick Start/Best Experience/Custom provide the same tasks and functionality. A dedicated MID Server is required to continuously stream logs from your instance to your log analytic system. The MID Server needs a one-time setup to establish a secure connection to the Hermes Messaging Service. Even if you have marked a task as completed, you can go back and uncheck it back to in-progress. To do so, first click on the **Edit** box on the upper right corner for the category. Then click on the **Edit** box for the task you want to uncheck. The **Mark as complete** box will be no longer marked.

Guided setup categories page

The categories page contains an overview and descriptions of the categories and associated tasks. You can click either the drop down arrow to view information about the category or click **Start** to open the guided setup steps and begin configuration.

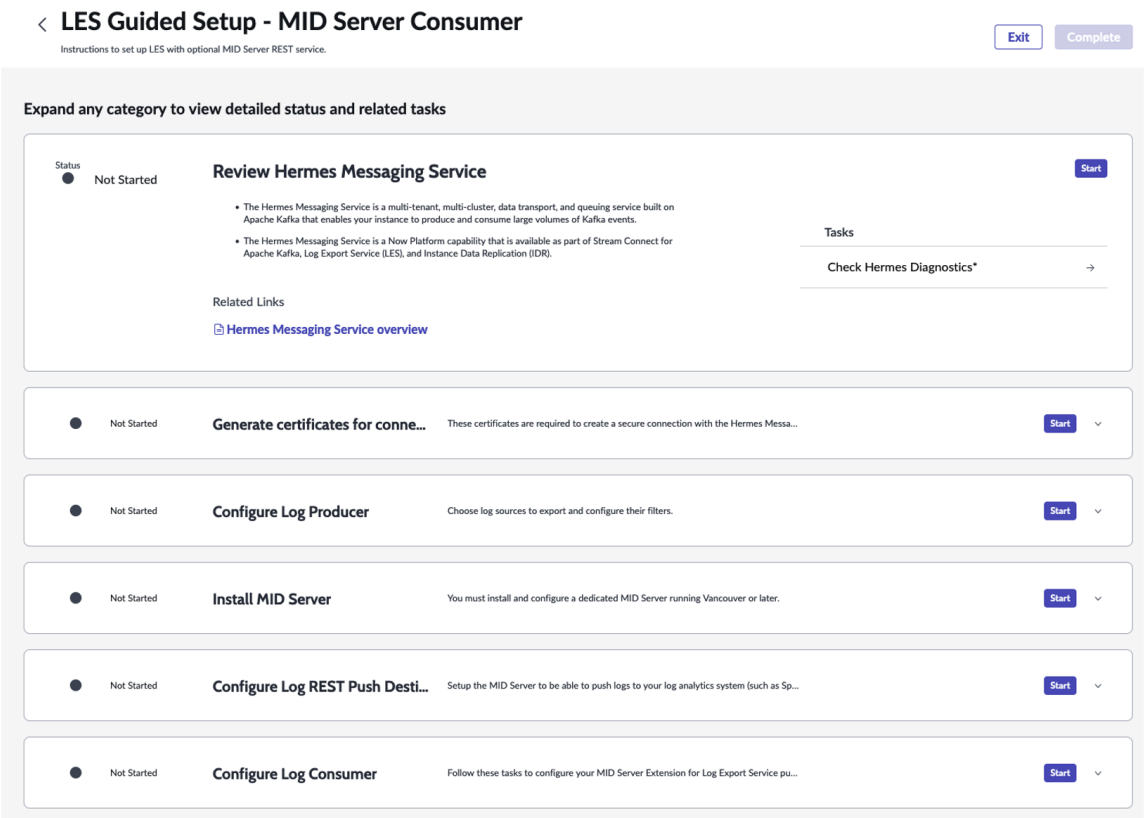

Complete the tasks under each category by following the setup instructions.

Guided setup for MID Server consumers

Implement the following steps for a complete guided setup for MID Server consumers.

Before you begin

Navigate to **Log Export Service (LES)** > **MID Server Consumer** > **Guided Setup**. Select the type of setup you wish to configure and select **Continue**.

Note: During the Log Export Service application installation, ServiceNow will provision the underlying Hermes Messaging Service infrastructure. Be aware that this process can take up to a couple of hours to complete from the time you request the Log Export Service application installation.

Role required: admin

Procedure

1. Review Hermes Messaging Service Diagnostics.

   It is recommended that you verify that the Hermes Messaging Service is up and running with the Hermes diagnostic tool, which displays on screen during this step. If you see a "Page not found" error on this page, Hermes is not installed and you should contact your system admin.

   - Setup information: The following bootstrap information is used to connect to the Hermes Messaging Service. The "Producer Bootstrap" is the connection used to send messages into Hermes and "Consumer Bootstrap 1 & 2" are used to retrieve messages from Hermes.
     - Producer Bootstrap
     - Consumer Bootstrap 1
     - Consumer Bootstrap 2

- Instance PKI: The Instance Public Key Infrastructure (PKI) component allows a ServiceNow instance to act as an issuer in a X.509 trust hierarchy.
- Bootstrap Connectivity: Select **Run Test** to confirm external client is able to connect to the defined instance ports (producer and consumer).
- Instance Connectivity: Select **Run Test** to confirm the instance is able to send and receive messages.
- View Topics: Select the listed topic to retrieve the timestamp of the last known message.

Note: You can access Hermes Diagnostics in the future to troubleshoot potential connectivity issues by returning to this step of the guided setup or by navigating to **All** > *Hermes Messaging Service* > **Diagnostics**.

2. Generate certificates for a secure connection to Hermes Messaging Service and pull log events from it.

   Setup secure connection to Hermes Messaging Service. See Set up a secure connection to the Hermes Messaging Service for LES for more information. You will need these certificates for authentication and authorization in the client which will pull the logs from Hermes.

   Note: admin or Hermes_admin roles are required for this step.

3. Configure Log Producer: Choose log sources to export and configure their filters.
   Complete the following tasks to configure the Log Producer.
   - Configure log sources to export: Create one Source record for each of the log sources that you want to export.

     Note: admin or sn_logstoanalytics.admin roles are required to complete this step

     To create a new Source, navigate to **Log Export Service** > **Sources**

     a. Select **New** at the top right corner

     b. Select a Source Type

     c. Select a Table

     d. Select a Log Level that can be used to control logging output

     e. Select or create a topic to which that log source will be exported. If you are creating a new topic, fill out the following fields:
        - **Name**: Name of the topic you're creating
        - **Application ID**: Enter `sn_logstoanalytics`
        - **Namespace**: Enter Default Namespace
        - **Partition**: The partition field of a topic in Hermes refers to the partitions into which the topic's data is divided. It plays a key role in scalability and parallelism.

     f. Select a **Filter Type** condition to forward logs selectively.

     Note: Filter Type options differ depending on the selected source type.

     g. Select Update
     When successfully created, it will display the name of the Hermes topic to which that log source will be exported to. Write down the topic name, you will need it later when configuring your log consumer system.

The Active field controls whether or not that log source is going to be exported or not. If you see errors, go back to the Check Hermes Diagnostics task and verify Hermes status.

- Validate Log Producer: Once you have created a Source to produce logs from, you can see live log records in the topic using **Hermes Messaging Service** > **Hermes Topic Inspector**.

  a. Select External Topics

  b. Select List Topics

  c. Select row with your topic from previous step (listed in Sources)

  d. Adjust message start date if necessary

  e. Select **View** to see a log message that was exported to the topic

4. Install MID Server: You must install and configure a dedicated MID Server running Vancouver or later.
   Complete the following tasks to install the MID Server.

   - Install dedicated MID Server: The MID Server that Log Export Service uses should be dedicated for only this purpose and should not be expected to run other processes. This is important to ensure timely delivery of exported log messages to your REST endpoint. You can install the new MID Server either by using the Use MID Server guided setup or by installing it manually. For the manual installation, follow the Configure MID Server network connectivity documentation first and then the Installing the MID Server documentation after that.
   - Validate MID Server: You must manually validate the MID Server after it is installed to enable it to execute automation tasks. To validate the MID Server you are dedicating for LES, see Validate the MID Server

   Note: If you configure a MID Server for this application, MID Server proxy settings do not apply to Hermes connectivity. MID Server proxy configuration applies only to HTTP-based communication. Because Hermes uses Kafka-native TCP connections, traffic to the Hermes cluster bypasses proxy settings regardless of MID Server configuration.

5. Configure Log REST Push Destination: Setup the MID Server to be able to push logs to your log analytics system (such as Splunk).
   Complete the following tasks to configure log REST push destination.

   - Add MID Properties: You must add MID Server properties so that it's able to connect to Hermes. Navigate to **MID Server** >**Properties** and set the appropriate values for each of the properties listed below.

Name	Value
mid.les.kafka.ssl.truststore.password	<password>
mid.les.kafka.ssl.keystore.password	<password>
mid.les.kafka.ssl.key.password	<password>
mid.les.kafka.ssl.truststore.location	<your_path>/<truststore>.p12
mid.les.kafka.ssl.keystore.location	<your_path>/<keystore>.p12
mid.les.kafka.ssl.truststore.type	PKCS12
mid.les.kafka.ssl.keystore.type	PKCS12
mid.les.kafka.client.id	<instance_name>

Name	Value
mid.les.kafka.group.id	snc.<instance_name>.group1
mid.les.kafka.bootstrap.servers	<instance_name>.servicenow.com:4100,<instance_name>.se
mid.les.kafka.set2.bootstrap.servers	<instance_name>.servicenow.com:4200,<instance_name>.se

Follow these notes on how to obtain some of the above values

- <password> is the password you set for the keystore and truststore
- <your_path> is the file path to the directory where you keep the keystore and truststore files you downloaded. The certificates need to be on the server where you are running MID
- <instance_name> is the name of your ServiceNow instance. If you are not sure, you can find it in the sys_properties table
- You can obtain the values for both mid.les.kafka.bootstrap.servers and mid.les.kafka.set2.bootstrap.servers from the Hermes Diagnostics page. Navigate to **Hermes Messaging Service** > **Diagnostics** and copy the strings under Consumer Bootstrap 1 and Consumer Bootstrap 2 respectively.

- Configure Destination: Create a new Destination Configuration record, which defines the REST endpoint that this Extension will forward logs to.

Note: admin or sn_logstoanalytics.admin roles are required to complete this step.

a. Navigate to **Log Export Service (LES)** > **Destination Configurations**

b. Create a new configuration record

c. Specify the URL for your desired endpoint for the exported log sources

d. Search for or create new credentials to connect to your endpoint. When creating credentials for your endpoint, note that only the following credential types are valid with LES: Basic Auth and OAuth

e. Search for or create a new transform script. We ship with the prewritten script, **SplunkTransform** for Splunk

6. Configure LOG Consumer: Follow these tasks to configure your MID Server Extension for Log Export Service purposes.

- Configure LES Consumer Context: In this step, you will update the LES Consumer record to execute on the dedicated MID server you just installed for the Log Export Service. Navigate to **MID Server** > **Extensions** > **LES Consumer Context** and update the LES Consumer record by setting the following fields:
  - Select on 'LES Consumer' to open the MID Server context record
  - Select Specific MID Server for the "Execute on" field
  - Enter the name of the MID you validated in the previous step for the “MID Server” field

Note:

- Each Consumer is designed to process data from a single Hermes topic using its own Consumer Context.
- Each Consumer Context runs on a dedicated MID Server to ensure optimal performance.
- We recommend configuring a separate MID Server for each Consumer Context you create.

- Configure Consumer: Create a new Consumer record, which represents the process that's part of the Log Export Service MID Server Extension. Navigate to **Log Export Service (LES)** > **Consumers** and create a new configuration record specifying the Hermes topic to retrieve log messages from and the Destination to relay them to. When you select the **Consumers** module, it shows information about the consumer name and the destination configuration.

  a. Create new Consumer record

  b. Choose a source topic from the dropdown

  c. Choose destination configuration

- Test MID Connection: Verify connectivity from the MID Server environment to the Hermes cluster before starting the Consumer.

  Roles required: admin or sn_logstoanalytics.admin

  Navigate to **Log Export Service (LES)** > **Consumers** and test the connection:

  a. Open the existing Consumer record.

  b. Click the **Test MID Connection** UI Action

  i. This validates connectivity between the MID Server environment and the Hermes cluster

  ii. The Status field will change to "Checking Kafka Network Connectivity"

  c. Wait for the test to complete, then refresh the page

  d. Review the Status and Status Detail fields

  e. If successful: Network connectivity is confirmed. Refresh the page and verify the Status and Status Detail fields reflect success.

  Start the consumer.

  f. If unsuccessful: Please pause before proceeding. Review the MID Server agent logs to understand the cause. Common areas to verify include:

  i. MID Server properties are correctly configured

  ii. Network connectivity is available between the MID Server and the Hermes cluster

  Note: The Test MID Connection must succeed before starting the Consumer.

- Verify MID Server integration: Navigate to#**Log Export Service (LES)** > **Consumer Status** and view the#Status and Status Details fields of the record defined. The information in these fields reports the current state of the process running on the MID server, including any errors that may have been encountered while relaying messages to the REST endpoint. This is a status view page only. Navigate to **Log Export Service (LES)** > **Consumers** if you want to create a new consumer record. #If the Consumer status indicates that the process has started, you should be able to inspect your endpoint to view the logs that have been relayed to it.#Additionally, you can view the logs on the MID Server to see if there are additional details about any errors that might be encountered. You can also enable Debug logging on the MID Server to get additional information if needed.

  Note: If you make any change in one of the Consumer records, it shows up on the Consumer Status view page. If you select a consumer record name on the Consumer Status list, the Consumer form for the selected record opens. You can then update the Name and Destination Configuration of the selected record.

Multi-consumer support using unique mid servers

You can now precisely manage log consumption with a new multi-consumer system, enabling dedicated consumers and MID servers for each specific log stream.

The system now supports multi-consumer log consumption which means each log source can be consumed separately have its own dedicated topic.

Previously, all logs were consumed from the same topic. However, with the new multi-topic concept, you can create multiple consumers for different topics. To create a consumer, you can select the following:

- A specific topic from the dropdown menu
- Respective destination configuration
- Respective consumer context

> **Note:** Starting Zurich release, the **Consumer context** is a new field added in the Consumer form. Each consumer context is associated with one unique mid server.

To create multiple consumers, you are required to select distinct topics for each consumer, which in turn would require to create separate consumer context records and its corresponding unique mid servers.

Parallel processing with multiple MID servers

A major benefit of this multi-consumer architecture is the ability to run multiple consumers in parallel. This parallel processing significantly enhances the overall throughput and efficiency of log consumption, allowing the system to handle larger volumes of diverse log data more effectively than before.

Set up a secure connection to the Hermes Messaging Service for LES

Secure your Kafka topics by generating a ServiceNow® instance-signed certificate.

Before you begin
Setting up the Hermes Messaging Service requires coordination with your network administrator and with your Kafka administrator. Work with your network administrator to obtain required security certificates and open the required ports. Work with your Kafka administrator to ensure that your Kafka environment is configured correctly and that your applications can connect to the Hermes Messaging Service using the standard Kafka protocol.

Make sure the following setup is in place:

- The Hermes Messaging Service is activated. See Activating the Hermes Messaging Service.
- The Key Management Framework plugin (com.glide.kmf.global) is activated.
- The Certificates [sys_kmf_certificate] table contains a ServiceNow instance root CA certificate.
- The instance isn't configured with a Custom URL. Custom URLs are not supported with the Instance PKI Certificate Generator.

Role required: hermes_admin, sn_kmf.cryptographic_manager, or admin

For details on assigning KMF roles, see Roles installed with Key Management Framework.

Procedure

1. Navigate to **All** > **Certificate Generator** > **Instance PKI Certificate Generator**.
2. **Optional:** Control access to topics by configuring Access Control Lists (ACLs) at the namespace or topic-level.

Option	Description
Apply ACLs to namespaces	**a.** Select **Configure ACLs**. **b.** In the Topic ACLs dialog box, select **Name spaces**. **c.** Enter a namespace that you want to config ure. **d.** Set the permission level by selecting either **Read Only** or **Read/Write**. **e.** Select **Add**.
Apply ACLs to defined topics	**a.** Select **Configure ACLs**. **b.** In the Topic ACLs dialog box, select **De fined topics**. **c.** Enter an existing topic that you want to con figure. **d.** Set the permission level by selecting either **Read Only** or **Read/Write**. **e.** Select **Add**.

   The bearer of the certificate is granted read or read/write access to the topics in the namespace or the existing topic that you selected.
3. Set up security for the Hermes Messaging Service.
   a. Navigate back to the Instance PKI Certificate Generator page.
   b. Enter a keystore password in the **Certificate Password** field.
   c. Select **Generate**.
   The system generates an instance-signed certificate in the Certificates [sys_kmf_certificate] table, creates a keystore, and creates a truststore.

   If Restricted Caller Access isn't allowed for the IPKI Certificate Generator, a cross scope access error appears. Contact Customer Service and Support for assistance with allowing Restricted Caller Access. To resolve this issue, Customer Service and Support can reference source_scope=76f9d51369115083f4ea77aab1677cc0 in the Restricted Caller Access Privileges [sys_restricted_caller_access] table.
4. Save a copy of the keystore by selecting **Download Keystore**.
5. Save a copy of the truststore by selecting **Download Truststore**.
6. Copy the keystore and truststore files to each producer and consumer client that will connect to the Hermes Messaging Service.

Result

You can now create a secure connection to the Hermes Messaging Service.

Note: You must use the keystore that you generated using the Instance PKI Certificate Generator to connect to Hermes. Custom-generated keystores that aren't created according to the ServiceNow documentation aren't supported.

What to do next

.

Using Log Export Service (LES)

Use LES to review the log report dashboard.

Review log report

Analyze the size of each data log by reviewing the log report dashboard.

Before you begin

Role required: admin or sn_logstoanalytics.admin

Procedure

1. Navigate to **All** > **Log Export Service** > **Reports**.
2. Review log data in the following widgets.
   - Logs Exported (Monthly): It shows the megabytes exported per month in the initial view. You can drilldown the graph to view the data for each week or daily data. You can view data from last 395 days.

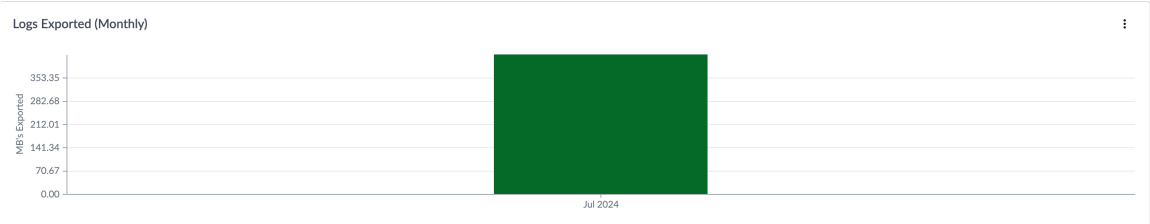

Note: Click on the arrows on the up right corner of the widget to drill down the data. You can also move back up to the parent data.

Weekly:

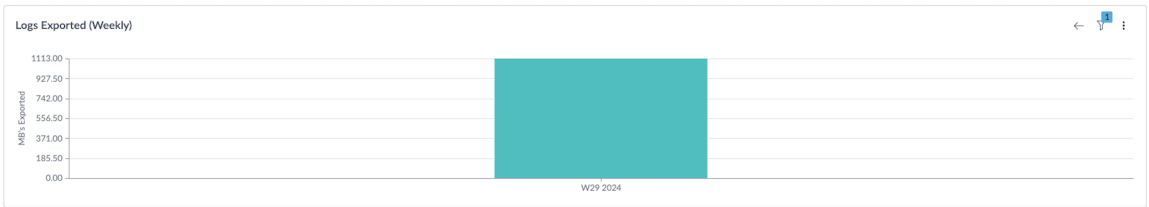

Daily:

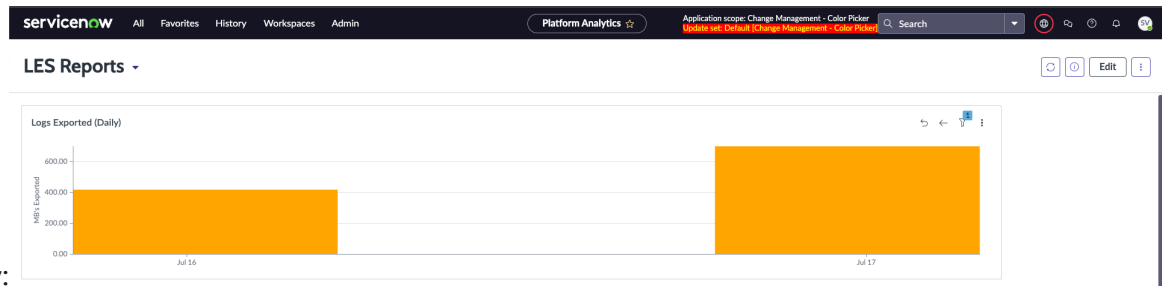

The data in this report reflects collection taken 4 hours prior.

- Logs exported per source (Monthly): This widget shows the kind of logs were exported on a monthly basis. You can also drilldown to see weekly and daily log data. You can view data from last 395 days.

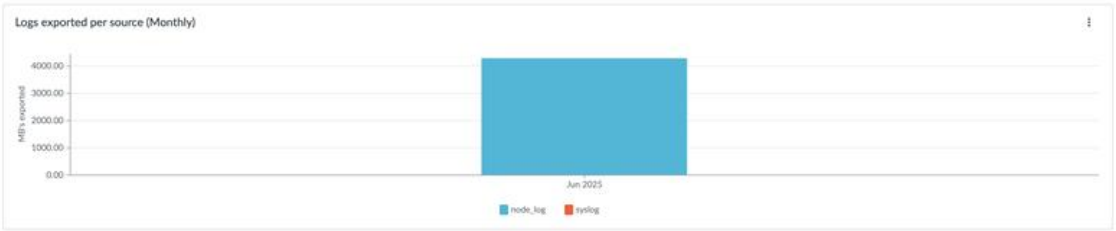

Note: You can select or de-select the source you want to view the data from. The sources from which you can view the logs are: node_log, syslog, syslog_transaction, and sys_outbound_http_log.

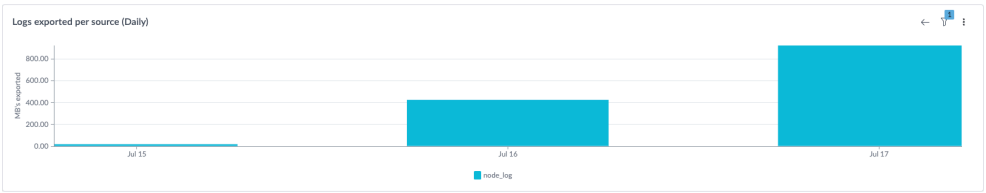

- Log Exported Per Source (Last 24 hrs): It shows the data from last 24 hrs. The graph progresses chronologically from left to right on a 24 hr scale, with the oldest data (24th hour) positioned on the right and the latest data on the left side of the report. You can select or de-select the source you want to view the data from. The sources from which you can view the logs are: node_log, syslog, syslog_transaction, and sys_outbound_http_log.

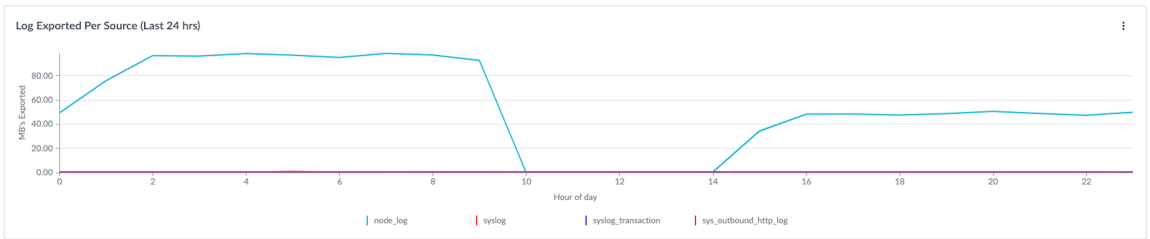

Note: Depending on the hour of the day you check this data, the widget dynamically updates to reflect exactly the past 24 hours.

- Daily Log Events by Table: It shows the log events generated by each log tables per day. You can select or de-select the tables you want to view the data from.

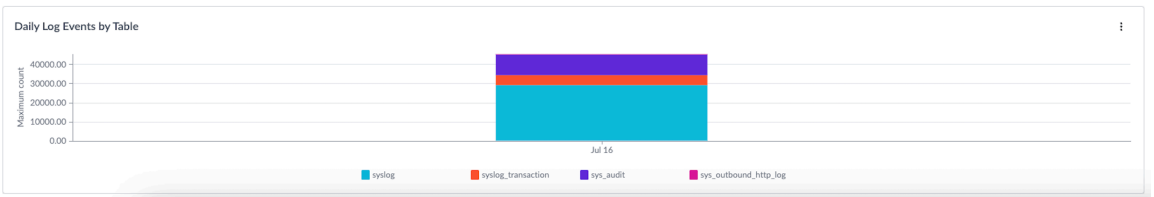

Note: The tables from which you can view the log events are: syslog, syslog_transaction, sys_audit, and sys_outbound_http_log.

Log Export Service (LES) references

Find all the miscellaneous information about LES in the reference section.

Log Export Service roles

Log Export Service is installed with these roles.

Application admin [sn_logstoanalytics.admin]

To learn more about managing per-user subscriptions, see Managing per-user subscriptions in Subscription Management and contact your account representative.

This role is installed along with the LES application and allows a non-admin to use the application.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, seeElevate to a privileged role.

No.

Special considerations

None.

System administrator [admin]

To learn more about managing per-user subscriptions, see Managing per-user subscriptions in Subscription Management and contact your account representative.

Admin role is required for the setup of the LES store application.

Contains Roles

List of roles contained within the role.

- sn_templated_snip.template_snippet_admin
- sn_employee.admin
- taxonomy_admin
- sn_ace.ace_user
- sn_hr_sp.esc_admin

Groups

List of groups this role is assigned to by default.

None.

Elevated

Whether the role is an elevated role. Elevated roles aren't assigned to users or groups, and must be used by elevation. For details, see Elevate to a privileged role.

No.

Special considerations

None.

Logs

Logs module provides a variety of logs that you can use to troubleshoot and debug transactions and events that take place within the instance.

Get started

System Logs

The System Logs module provides a variety of logs that you can use to troubleshoot and debug transactions and events that take place within the instance.

Logging, auditing, and errors

Apply a logging and auditing strategy so that you can identify and act on suspicious activity in a timely manner.

System logs

The System Logs module provides a variety of logs that you can use to troubleshoot and debug transactions and events that take place within the instance.

Access the following logs from the System Logs module:

Log	Description
Transactions	All application activity for an instance.
Email and Push	All email notifications and Push messages sent from all instances within the system.
Event Logs	All system events that occur within the system.
Import	Data import activity within the platform.
Table Changes	Changes made to all tables in the system.
Outbound web services logging	All outbound web services requests such as REST and SOAP requests.
Signature Images	Electronic signatures for the HR signature pad.

Log	Description
System	Warnings and errors for instance processes, records, and non-critical events, such as memory usage on the server machine.

Use the Log File Browser to search and download logs. You can also search archived logs in the log history.

Other logs

Your instance offers other logs in addition to those in the System Logs module. For example, the System Diagnostics module provides upgrade history and slow query logs, which you can use to gain insight into how queries are affecting platform performance. The Customer Updates table records every change that is made in the system.

System log

View warnings and errors for instance processes, records, and non-critical events, such as memory usage on the server machine.

The following information is tracked in the system log:

- Workflows
- Configuration
- Chats sessions
- Transactions for each view of each page in the system, including load times for network, server, and browser
- Inbound and outbound email
- Events triggered in the system
- Imports and integrations
- System warnings, errors, and script logs
- Upgrade information for any plugin activations, update sets, or system upgrades

Log entries appear for the current day only. To view other log files, use the log file browser.

Syslog messages serve diagnostic and maintenance purposes. Log messages can have different owners and may not require action from all readers. Check the message source to determine relevance and ownership. The following fields describe the attributes of each message.

System log

Field	Description
Created	Date and time of the logging activity for the locale of the machine running the instance.
Level	The severity classification of the message. Levels indicate the nature of the event, not necessarily whether action is required. Messages are generated by the ServiceNow platform, installed applications, and custom scripts. The levels are Debug, Error, Warning, and Information. • Debug: Detailed diagnostic information, typically used during development or active troubleshooting. • Information: Confirmation that an operation completed as expected.

System log (continued)

Field	Description
	• Warning: An unexpected condition was encountered, but the operation completed successfully. May require investigation depending on context and source. • Error: An operation failed or could not complete as intended. Instances can generate a high volume of Error messages, primarily from the platform itself. These do not necessarily indicate a platform issue and do not always require further investigation.
Message	System-generated message regarding the nature of the occurrence.
Source	Name of the process or area affected by the occurrence. For example, the source of the occurrence might be EMAIL or Memory.
Source Package	Name of the application package associated with the occurrence. Select the name to view the Store Application [sys_store_app] record for this package.

Workflow logging

- Each activity executed, including:
  - Date and time started
  - Date and time ended
  - State, for example, Finished, Cancelled, Timed Out, Error
  - Result
  - Fault description, if there was an error
- Transition history, including:
  - Time of transition
  - Activity transitioned from
  - Activity transitioned to
  - Which transition was triggered
- Workflow log, including any log statements added to the workflow

Configuration information

- Action taken, including insert, update, and delete
- Category of change
- Comments recorded with the change
- Name of the change
- XML difference of the change
- Update set associated with the change
- Date and time of the change
- User who made the change
- Table where the change was made

- Name of the object being changed
- Type of object being changed
- View in which the change was made, for form or list changes

Transaction logs

The transaction log records browser activity for an instance. To aid in debugging of system issues, you can filter transaction logs by application scope, limiting transactions that appear to only those transactions originating in specific scopes.

Note: Background and scheduler transactions are logged only when their execution time exceeds 1000ms. For a complete list of logged transactions, see KB0778299.

Access the transaction logs by navigating to **All** > **System Logs** > **Transactions**. The transaction log provides the following information for all activities.

Field	Description
Created	Date and time of the application action for the locale of the machine running the instance.
Type	Type of recorded transaction.
Created by	User who created this activity.
Origin application	Application scope the transaction originated in. Global appears if the transaction originated in the global scope.
Response time	Round trip response time for the application request, in milliseconds.
Network time	Latency time of the network response after the application request is made, in milliseconds.
Output length	Size of the output string sent by the instance to the application, in bytes.
SQL count	Number of SQL server commands executed for this activity.
Business rule count	Number of business rules executed for this activity.
Business rule time	Elapsed time for the execution of the business rules for this activity.
URL	Application or module connected to by the client application.
System ID	System generated identifier of the client instance making the request. This ID is used for cluster environments in which several instances (nodes) communicate with the database.
IP address	IP address of the client making the request.
GZipped	Indication of whether a compressed Web page was requested by the application.
Protocol	HTTP protocol used by the application for this instance.

Client transaction timings

The Client Transaction Timings plugin enhances the system logs by providing additional information about the durations of transactions taking place between the client and the server.

You can track down performance issues to their source by viewing where the time is being consumed, and how time was spent during a transaction.

This plugin requires the Response Time Indicator Response Time Indicator to be enabled, and collects information from the following browsers:

- Firefox
- Internet Explorer
- Chrome

Client Transactions Information

Installing the plugin adds the *Client Transactions* module to the *System Logs* application. It provides a list of every logged transaction between client and server within the last day:

Client Transactions Information

Field	Description
Created	Moment the transaction was recorded.
Response Time	Number of ms spent by the server in fulfilling the transaction.
Business Rule Time	Number of ms spent by business rules triggered by the transaction.
SQL Time	Number of ms spent by the SQL database.
Client Response Time	(Load_completion_time) - (start_time). It is inclusive of server time.
Client Network Time	Number of ms spent by the network the client is connecting through.
Browser Time	Number of ms spent by the browser during the transaction.
Client Script Time	The number of ms spent executing client scripts.
UI Policy Time	Number of ms spent executing ui policy.
Type	Type of transaction: • Form • List • Other
Table	Name of table that appeared. For example, incident, change_request.
View	View for this form/list.

Client Detailed Information

A more detailed breakdown of the client timings for all Form rendering (but not list rendering) is also tracked. To see details, drill into a particular client transaction record and observe the related list at the base of the screen.

Client Detailed Information

Field	Description
Order	Order during the load in which this operation occurred.
Type	Type of operation.

Client Detailed Information (continued)

Field	Description
Name	Descriptive name of this particular operation
Duration	Number of ms this operation took to complete.

Push logs

Consult the push log to track the status of push notifications that are queued to send from your system.

To view the push log, navigate to **System Logs** > **Push Notifications**. Users are required to have the push_admin or admin role to view the push log.

Push log fields

Field	Description
Claim	Identification number that is generated by the scheduled job that sends the push notification. This number is applied to the **Claim** field to ensure consistency across multiple push scheduled jobs.
Payload	Content of the push notification.
Queue count	Number of times that the system tried to send the push notification. The status of the push notification relates to its queue count. • If the queue count is 0 for longer than expected, then there are no scheduled jobs that are trying to send the push notification. • If the queue count is greater than 0 and the **Type** is **success**, then you can infer that this is how many times the system tried to send the push notification before finally sending. • If the queue count reaches 10, the system stops trying to send the push notification. The **Type** changes to **failure**.
Request ID	Unique identification number for the push notification. Similar to the message ID for an email, the request ID is used as a correlation token for the push notification.
Type	Status to indicate whether the push notification has been sent. The Type column can have these values: • **failure**: The message could not be sent. • **pending**: The message is queued for processing. • **success**: The message was successfully sent, although not necessarily received by the mobile device.

Related topics

Scheduled jobs

System email log and mailboxes

The system email log records all emails that the instance creates or receives. System mailboxes are filtered views of this log.

Every notification email that the instance creates or receives is recorded in an Email [sys_email] record. You can navigate to a log of these records at **System Logs** > **Emails**.

The System Mailboxes are filtered views of the Emails [sys_email] table. The instance assigns an email record to a system mailbox depending on the values of the **Type** and **State** fields. For more information, see System mailboxes.

The following fields can be included in the layout of the system log and any of the system mailboxes:

Email log

Field	Description
Mailbox	The system mailbox that lists this email record. The instance sets the value of this field according to the values of the **Type** and **State** fields.
State	The current state of the email (Error, Ignored, Processed, or Ready).
Receive type	The type of inbound email (None, Forward, New, or Reply).
Type	The status of the email. Choices are: • received: The server received this email. • received - ignored: The server received this email, but it was ignored by the instance for inbound email action purposes. Typically, these emails are either spam or auto-replies. See the Error String field for details. • send - failed: The server has attempted to send the email and failed. See the Error String field for details. • send - ignored: The server skipped sending this email. Typically, this is for an email which was generated but lacked a recipient email address or is a duplicate email. See the Error String field for details. • send - ready: The email is ready to be sent, but has not been sent out by the mail server. Typically, an email remains in this state for only a short time. • send - translation - ready: The email is generated during email translation and sent out. Typically, an email remains in this state for only a short time. • sent: The email was sent by the instance without any errors or issues.
Target	A Document ID reference to the record if the email is generated by an insert, update, or delete of a particular record.
User	The name of the user, from the user record, of the instance from which the email notification was sent. **Note:** This is a string field.
Notification Type	The type of notification. Choices are: • None • SMS • SMTP
UID	The unique ID of the email stored on the server.

Email log (continued)

Field	Description
Created	The date and time of the email activity for the locale of the machine running the instance.
Deleted	For inbound email, indicates whether the email was deleted from the email server.
Weight	The weight of the email, which determines the sending priority relative to other notifications on the same table.
Importance	An indication that the email was sent with a changed level of importance, such as Urgent.
Originating Event and Notification	For emails generated by notifications, an embedded list that stores the event and notification that created the email.
Subject	The email subject. For notifications, you create the subject text in **System Notification** > **Email** > **Notifications**.
Error String	The error string captured from the email server to determine why the email was not sent. This is logged only if the email is send-failed.
Recipients	The email addresses of the recipients.
Body	The body of the email, displayed in raw HTML markup. Use the related link Preview HTML Body to see the body text as rendered HTML.
Content type	The email content type.
Headers	Any headers embedded in the email.

Event logs

The event log records all system events that occur within the ServiceNow AI Platform.

This log provides the following information for all events that occur:

Event log

Field	Description
Created	Date and time of the event for the locale of the machine running the instance.
Name	Name of the event as listed in the Event Registry.
URI	HTTP query that generated the event.
Parm1	Event-specific value that depends on the event and the recipient.
Parm2	Event-specific value that depends on the event and the recipient.
Table	Database table acted on for this event.
Processed	Date and time the event started processing. This time reflects the locale of the machine running the instance.
Processing time	Time taken to process this event, in milliseconds.

Event log (continued)

Field	Description
Queue	Processor queue name.

Import logs

The import log displays information in a verbose format about any data import activity within the platform.

For a more detailed view of the import sets that produced a particular log, see **Import Sets** > **Transform History**.

This log provides the following information for all imports:

Import log

Field	Description
Created	Date and time of the import for the locale of the machine running the instance.
Level	Type of message displayed. For import files, the level is Information.
Message	System-generated message regarding the status of the import.
Source	Name of the external source of the import, such as an integration.

System Diagnostics module

The System Diagnostics application provides logs that relate to the platform.

These logs are available:

Upgrade History

Tracks every upgrade to an instance.

Slow Queries

Provides insight into how queries affect platform performance. See Use a slow query log.

Customer Updates table

Changes made in the system are recorded on the Customer Updates [`sys_update_xml`] table chronologically. There are a few exceptions, as noted below.

To navigate to this table, enter `sys_update_xml.list` into the navigation filter. For information about update sets, see System update sets.

The following information is stored about each update:

Customer Updates table

Field	Description
Name	A name that identifies the updated record.
Created	The date and time the Customer Update record was created.
Created By	The user who performed the change.

Customer Updates table (continued)

Field	Description
Type	The type of the update.
Updated	The date and time the Customer Update record was updated.
Updated By	The user who performed the update.
Updates	The number of times the record has been updated.
Target Name	The name of the element that was altered.
View	The view of the form that was altered if it was a form layout change.
Payload	The XML contents of the record after the change.
Remote Update Set	A reference to that update set if the change was performed by a remote update set.
Local Update Set	The update set the change is associated with.

Note: Some application changes are not represented by Customer Update records (sys_update_xml) Examples:

- Types of metadata where updateSynch = false
- Cascading changes to tables and fields such as display name changes
- Unresolved metadata references from other applications (resulting in no “display value” on the element)
- sys_id changes for coalescing files
- Changes to Flow/Flow Actions that may generate sys_documentation
- ua_table_license_config records generated on table creation
- Jobs running in the background such as natural language processing
- Cases where sys_update_xml is manually modified or removed

To learn more, see Commit changes.

Log history

The system uses table rotation and table extension to archive older logs.

By default, the system uses the following schedule to archive common logs:

Common log archive schedule

Table	Archive schedule	Rotations	Type
Event [ecc_event]	Every day	7	Rotation
Queue [ecc_queue]	Every day	7	Rotation
Event [sysevent]	Every day	7	Rotation
Log [syslog]	Every week	8	Rotation
Transaction Log [syslog_transaction]	Every week	8	Rotation

Common log archive schedule (continued)

Table	Archive schedule	Rotations	Type
Email [sys_email]	Every 30 days	8	Extension

Use the log file browser

The instance provides the utilities log file browser and log file download.

Use **System Logs** > **Utilities** > **Node Log File Browser** to view any system log entry. You can search for log files by using the following filters:

Log file browser

Field	Description
Start time	Start date and time of the range you want to search, for the locale of the machine running the instance.
Session ID	System-generated hexadecimal string that identifies the session that generated the log entry.
End time	End date and time of the range you want to search, for the locale of the machine running the instance.
Message	System-generated description of the occurrence.
Level	Type of message displayed. The levels are Debug, Error, Warning, and Information. A warning is an error that has been handled and recovered. An error is something that must be fixed.
Thread name	System-generated identifier of the thread that created the log file.
Max rows	Maximum number of records returned for a particular filter.

The instance creates compressed archives of system logs every 2 days and purges log archives after 21 days.

Important: When a node is retired, the node's log files are purged immediately, which means they aren't archived for an additional 21 days.

You can download log file archives and view them with **System Logs** > **Utilities** > **Node Log File Download**. Select a log archive from the list, and then click **Download log** under *Related Links* to open or save the archive.

Note: Log files are only available for the node you are currently logged into. To see the currently logged into node, navigate to **System Diagnostics** > **Stats**.

Use the new **Show Syslog Records** button on the Transaction and Active Transaction forms to view any System Log entries that were generated during the execution of the transaction. A transaction can have any number of syslog entries. The multiple syslog entries for all the transactions make it difficult to co-relate a transaction with their respective syslog entries. The **Show Syslog Records** UI action helps in co-relating the active and completed transactions to their respective syslog entries by building a URL to query the syslog table. Identifying the correct syslog entries for a particular transaction helps in debugging and addressing security concerns.

Enhanced logging security

Explore the **Attribution** field in the node log lines to identify the script or component that generated the log message. Transaction start lines include the new field to identify the type of request made.

Achieve the following using the new enhancements:

- Trace the source originator of each of the log lines
- In case the originator info is not available, print the Java class name and an attribution
- At the start of each transaction line, it contains the transaction ID and transaction type

Use the transaction ID of each log line to understand the information given at each log line. Once you identify the transaction type, you get the originator information of each log line. Both the transaction type and originator information of each log line together gives you the required source info of each node log line.

Note: SYS_UI_MACRO and SERVICE_PORTAL_WIDGET script types in attribution are not reported.

Transaction types

The following is the list of transaction types:

- List
- Form
- XMLHttp
- Report
- SOAP
- Export
- Scheduler
- TextSearch
- Other
- REST
- JSON
- AMB
- Archive
- Batch REST
- Instance Scan

System properties

The following are the system properties required for the feature:

- Glide.log.append.attribution: This property is enabled by default. It turns on/off the attribution information of each node line
- Glide.db.log.append.classname.attribution: This property is enabled by default. It turns on/off logging java class name attribution

Avoid log tampering

Configure system log table protection rules to limit the scope of modification and deletion of application log records. The rules enable you to determine the logging of changes or attempts to changes in these tables.

If you are a security_admin, activate the Protected Tables plugin (com.glide.protected_tables) that allows the platform to restrict update, insert and delete operations on the following system log tables:

- syslog
- syslog_transaction
- sys_outbound_http_log
- sysevent
- sys_audit
- sys_push_notification
- syslog_app_scope
- protected_table_configuration (config not modifiable)

Note: The com.glide.protected_tables plugin gives protection only to the system log tables mentioned above. Any attempt to update, insert or delete a record logs a message in the protected_table_log table.

See Installing and configuring the log protection plugin for more details.

You can specify one of the following log protection levels for each of the system log table.

- Block and log the attempt: Blocks any modification and logs the attempt
- Only block the attempt: Blocks any modification and doesn't log the attempt
- Only log the attempt: Doesn't block the modification but logs the attempt
- Don't block and don't log the attempt: Doesn't block the modification and doesn't log the attempt

Platform uses the log protection levels specified for each of the system log tables to block and/or log any attempts to modify a record after being initially created.

Note: If you are a security_admin, you have the ability to override the default log protection levels in each of the system log tables to adapt to the customizations on your instance.

If there have been any attempts to modify the system log tables, they are logged into the protected_table_log table.

Note: If the protection level is not specified for a table, any attempts of modification are not logged into the protected_table_log table.

In order to disable the plugin operations on tables in the Admin Panel, set the com.glide.security.protected_table.enabled property to false. See Create log protection property for more information.

Configuring the log protection plugin

Configure the protection rules for each table and operation to complete the configuration of the log protection plugin.

Before you begin
Role required: security_admin

Procedure

1. Navigate to **All** > **Protected Tables** > **Log Protection**.
   The Admin Panel page is displayed.

   Note: Starting in the Utah release, the Protected Tables plugin is installed by default, but disabled.

2. Elevate your role to security_admin in order to continue configuring the plugin.

   a. Select **System Administrator**.

   b. Select **Elevate role**.
   The Elevate role modal shows up.

   c. Select the security_admin option to elevate your role and select **Update**.

3. Configure the protection rules for each table and operation. The protection rules apply to update, insert and delete. The protection levels can't be changed for some tables. The syslog and syslog_app_scope tables have fixed values for update and delete protections. The protected_table_configuration table has fixed protection values for all three operations.

   Note: For sysevent table, insert protections can't be set to block.

   If you had **Apply to Child Tables** turned on for syslog in the previous release, the child tables are added to Log Protection with the same protection rules as syslog upon upgrade to the Utah release. This only happens for syslog, not for any other tables.

4. Enable the feature by selecting the **Enable Log Protection** toggle.

   Note: You can disable this plugin only by changing the com.glide.security.protected_table.enabled property in the sys_properties table.

Create log protection property

Create a log protection property to avoid the risk of log tampering.

Before you begin
Role required: admin

Procedure

1. If the *com.glide.security.protected_table.enabled* property doesn't exist in the System Properties list, select **New**.
   The new system property form shows up.
2. On the form, fill in the details

Fields	Description
Name	Name of the property as com.glide.security.protected_table.enabled
Application	Application that has the property.

Fields	Description
Description	Description of the property
Choices	
Type	Value type - true or false
Value	Actual value of the property
Ignore cache	Option to ignore the cache content
Private	Option to make the property private ○ Read roles ○ Write roles

3. Select **Submit** to create the property.

Logging, auditing, and errors (instance security hardening)

Apply a logging and auditing strategy so that you can identify and act on suspicious activity in a timely manner.

To learn more about what can be logged in the instance, see System logs. Ensure that there is a schedule for monitoring system events such as logins and failed logins by using **System Logs** > **Events**.

Disabling SQL error messages (instance security hardening)

Use the *glide.db.loguser* property to disable SQL error messages from rendering in a browser.

More information

Attribute	Description
Property name	glide.db.loguser
Configuration type	System Properties (/sys_properties_list.do)
Configurable in Instance Security Center	No
Purpose	To disable SQL error messages from displaying within the browser.
Type	true|false
Recommended value	false
Functional impact	(Low) This remediation disables rendering of SQL error messages. There is no impact to any functionality.
Security risk	(Medium) No sensitive SQL information that could help an attacker should appear as a part of error message on a web page.

To learn more about adding or creating a system property, see Add a system property.

Secrets Management

Secrets Management lets you control which applications and users can access sensitive credentials stored on your instance.

Important:

Secrets Management has begun its End of Life process, and has reached the End Of Sale and End of Renewal milestones as of the Yokohama release. For support storing passwords with two-way encryption, see Password2 encryption with the Key Management Framework (KMF).

Use Secrets Management to restrict access to passwords and other sensitive credentials beyond standard role-based controls, such as when different teams or applications share an instance but shouldn't have access to each other's credentials, or when your organization's security policies require fine-grained control over who can decrypt specific secrets.

Secrets Management has two available versions. Secrets Management Core is included on the ServiceNow platform at no additional cost and Secrets Management Enterprise is a premium, subscription version of the product. For more information on the difference between these versions, see Exploring Secrets Management.

Explore  Learn the key features and business value of Secrets Management.	Analyze  Learn more about Secrets Management dashboard.
Configure  Plan your core configurations.	

Exploring Secrets Management

Use ServiceNow Secrets Management for granular management of access to your passwords to fit your business needs.

Important: Admins must have the role to see modules and records related to Secrets Management. For secrets management role information, see Secrets management roles.

Select from Core and Enterprise versions of Secrets Management

Choose from Secrets Management Core and Secrets Management Enterprise depending on your business needs.

The Secrets Management Core plugin (com.glide.sm.core) is available by default. No installation is required on the instance to use this plugin. The Secrets Management Enterprise plugin is only available with a ServiceNow Vault v1, PROD18537 license. Contact Customer Support for assistance with the Secrets Management Enterprise plugin.

Secrets Management Core	Secrets Management Enterprise
Secrets Management Core is available by default to install on your instance at no additional cost. The plugin provides the ability to use secrets groups with criteria in non-custom tables provided in the ServiceNow platform that have been created by ServiceNow application engineering teams.	Secrets Management Enterprise includes additional functions to help admins create and manage secrets groups. Enterprise provides the following features in addition to the features listed in Core. • Use granular access controls to create secrets groups based on any of these criteria: ○ Scope ○ Package ○ Table ○ Column ○ Record • Create client-accessible secrets that are encrypted using your own key which ServiceNow can't access. • Use the Secrets Management Dashboard to review the secret groups configured on your instance and learn about potential security issues. **Note:** Secrets Management Enterprise is a paid plugin that ServiceNow personnel must activate on your production instance.

Use secret groups to organize your secrets

Use Secrets Management to organize your secrets into groups. Then, apply access policies to those secrets at a group level.

Basic secret group

These groups apply to all secrets in a scope. These secrets are decrypted by a common cryptographic module and module access policies (MAPs).

Secret group with criteria

Secret groups with criteria function the same as a basic secret group, but further refine what is included using criteria. These criteria include:

- Application scope
- Package
- Table
- Secret column
- Filter record

Secret groups of either type can be made instance accessible or client accessible.

Instance-side secret groups

Instance-side secret groups contain secrets that can be decrypted by your instance.

Client-side secret groups

Client-side secrets groups use a public/private key pair so that secrets can only be decrypted by the client. When you create a client-accessible secrets group, you upload the public key to the instance and retain the private key on your MID Server. The instance uses the public key to encrypt your secrets, but they can only be decrypted using the private key.

Note: For more information on these group types, see About client-side Secrets Management.

Use secrets groups for more granular control

While Password2 is available on the ServiceNow platform, Secrets Management provides these additional features.

Granular access controls	**Password2** With Password2, admins can control access to an application scope but can't restrict access to elements within the scope. **Secrets Management** With Secrets Management, admins can restrict access based on criteria they define. Criteria types can be based on criteria such as package, table, or column.
Secure storage	For client-side secret groups, Secrets Management uses a new encryption scheme. In this encryption scheme, ServiceNow doesn't save the encryption key. For this reason, the security of your data doesn't depend on ServiceNow's security.

Apply module access policies to your groups

After you've grouped your secrets into a secret group, you can apply policies that determine how you can access them at a group level. Module access policies are the access control mechanisms that you apply to cryptographic modules to define instance-level controls, such as a validity time frame for the cryptographic key. For more information on module access policies, see Module access policy overview.

Tables installed with Secrets Management

Secrets Management adds or modifies these tables.

New tables	
[sn_sm_secret_group]	Stores secret groups
[sn_sm_secret_group_criteria]	Stores criteria secret groups
[sn_sm_secret]	Stores wrapped secrets
[sn_sm_identity_group]	Defines the identity group for mapping a group of identities to the public key
[sys_kmf_wrapped_module_key]	Stores the wrapped symmetric cryptographic keys
Modified Tables	
[sys_kmf_crypto_module]	Added cryptographic module type (identity cryptographic module or secret group cryptographic module)
[sys_kmf_module_key]	• Stores conceptual secret encryption key (with no key material) • Stores the identity public key
[sys_kmf_crypto_caller_policy]	Added new module access policy type

Secrets Management use case examples

Help ensure secure ITOM Discovery

This infographic shows a simplified reference architecture of how ServiceNow IT Operations Management (ITOM) Discovery can be deployed by your organization. As shown in the infographic, multiple Windows and Linux servers connect to the Management, Instrumentation, and Discovery (MID) Server and several MID Server agents enable the discovery process to update the Configuration Management Database (CMDB). Every MID Server transaction requires a secure authentication, so managing the authentication credentials is critical from a security perspective.

Accelerating workflow connectivity with Integration Hub securely

Use ServiceNow's Integration Hub to connect to different systems using automated application programming interface (APIs). Each time Integration Hub connects to a system using an API, an authentication credential is required to establish connectivity. Management of a multitude of applications and APIs for connectivity is made easier by using a secrets management solution.

Secrets Management is a key part of ensuring your organization’s cybersecurity. It covers all processes and tools related to the creation, storage, transmission, and management of digital credentials such as encryption keys, API tokens, and passwords. To manage secrets both securely and effectively, you can build a core secrets management policy that establishes standard rules and procedures for all phases of a secret’s lifecycle.

About client-side Secrets Management

Learn how use Secrets Management to manage access to secrets and groups.

Terminology

Client-side secrets management is designed to provide a method for managing secrets without the use of proxies, and without giving ServiceNow access to your decrypted data. To understand this process, begin with the following encryption terms.

Term	Description
Symmetric encryption	Symmetric encryption uses a single same key both to encrypt and decrypt data. If data is encrypted with a symmetric key, this key is all that is needed to decrypt it.
Symmetric key	The symmetric key encrypts a secret, turning your clear text password into unreadable cyphertext.
Asymmetric encryption	Asymmetric Encryption uses two keys, one to encrypt and the other to decrypt.
Public key	The public key is one half of the asymmetric key pair. This key is stored on your instance, which uses the key to encrypt a symmetric key. This encrypted symmetric key can only be decrypted when paired with the private key.
Private key	The private key is one half of the asymmetric key pair. This key is stored in a keystore on your MID Server. ServiceNow has no access to this key. Combined with the public key, the asymmetric key pair is used to decrypt your secrets.

Client-side encryption process

A symmetric key encrypts a credential (in this case, an admin password), changing it from readable cleartext into encrypted cyphertext.	
The symmetric key (represented in green) can be applied to the credential to encrypt or decrypt it.	
At this point, asymmetric encryption begins using public (green) and private key (blue) keys.	

The public key encrypts the credential along with the symmetric key. The symmetric key is now protected, so it can't be used to decrypt the credential. Although the public key can perform this encryption, it can't be used alone to decrypt.	
After being encrypted with the public key, the private key is needed to decrypt the credential. Since the customer alone has this key, they're the only ones who may access the encrypted credential.	

Configuring client accessible secrets

Learn how to configure your instance to use client accessible secrets.

Use this example implementation to configure Secrets Management without using proxies, or giving ServiceNow access to your decrypted data.

For more detail on using client-side Secrets Management to manage access to passwords and groups, see About client-side Secrets Management.

These instructions assume you have a MID Server configured on your local network. For information on this process see MID Server.

Process overview

1. **Create encryption keys and certificate**

   Create encryption keys and a certificate using terminal commands on your local environment.

2. **Add your certificate to the ServiceNow Trusted Key Store**

   Upload your key and certificate to the ServiceNow Trusted Key Store.

3. **Create a secret group with criteria**

   Create a group for your secrets. Secret groups are used to organize your secrets into groups. Using these groups enables you to apply access policies to those secrets at a group level. Then associate your secrets group to an identity group and add your MID Server to that identity group.

4. **Upload the public/private keypair to the MID Server**

   Upload your public/private keypair to your MID Server. This keypair enables the MID Server to handle authentication requests from your instance.

5. **Create credentials and test credential encryption**

   Create a credential to authenticate into a third-party system and test that ServiceNow can't access the credential.

6. **Configure Flow Designer to manage the integration**

   On your instance, use Workflow Studio to manage an integration between your local network and your instance.

7. **Test the end-to-end client-side encrypted secrets integration**

   Test your integration, and review the execution details to confirm your configuration is working.

Create encryption keys and certificate

Create encryption keys and a certificate using terminal commands on your local environment.

Before you begin

Role required: none

Procedure

1. In your local environment, open Terminal (on Mac or Linux), or Command Line (in Windows).
2. Using the terminal, use `cd` to move into the folder where you want to store your encryption keys.
3. Using the terminal, enter the following:

```
openssl req -newkey rsa:4096 -nodes -keyout sm_private_key.pem
 -x509 -days 365 -out sm_public_cert.pem
```

> **Note:** This example uses OpenSSL to generate keys and certificates. You may substitute other comparable tools based on your requirements.

   The command generates a private key and a public certificate (with the matching public key). A series of prompts for required information follows, starting with "Country Name."

4. Fill in the prompts with the requested information.
   The following prompts appear.
   - Country Name
   - State of Province Name
   - Locality Name (for example, city)
   - Organization Name (for example, company)
   - Organizational Unit name (for example, section)
   - Common Name (for example, fully qualified host name)
   - Email address

   Work with your security team to verify that you enter the correct certificate information.

```
Country Name (2 letter code) []:US
State or Province Name (full name) []:CO
Locality Name (eg, city) []:Boulder
Organization Name (eg, company) []:ServiceNow
Organizational Unit Name (eg, section) []:Product Management
Common Name (eg, fully qualified host name) []:fake@servicenow.com
Email Address []:fake@servicenow.com
```

5. Check the folder you that chose in step 2 to verify that the private key and public certificate have been created.
   If you used the same file names as in the step 3 example, you should see the following files:
   - sm_private_key.pem
   - sm_public_cert.pem
6. In the same folder, use the following command:

> **Important:** The specific command to use depends on your operating system.

For Linux:	`cat sm_private_key.pem sm_public_cert.pem > sm_keypair_bundle.pem`
For Windows:	`sm_private_key.pem sm_public_cert.pem > sm_keypair_bundle.pem`

This command bundles the private key and public certificate into a single file to load into your MID Server in later steps.

7. Check the folder again to verify that the new file containing your private key (sm_keypair_bundle.pem) and public certificate has been created.

Add your certificate to the ServiceNow Trusted Key Store

Upload your key and certificate to the ServiceNow Trusted Key Store.

Before you begin

Role required: admin

The public certificate that you created in this example is considered a “self-signed” certificate (meaning it didn't come from a trusted root authority). You must add the certificate to the ServiceNow Trusted Key Store for it to be used. When using a certificate that came from a Certificate Authority, you don't need to complete this step.

Procedure

1. In your local environment, open Terminal (on Mac or Linux), or Command Line (in Windows).
2. Using the terminal, use `cd` to move into the folder where you have created your encryption keys.
3. In the terminal, enter the following command:

   ```
   cat sm_public_cert.pem
   ```

   You must view the contents of your public certificate to copy it into the trusted key store. This `cat` command displays the certificate.

4. Copy the certificate information to your clipboard.
   Start with and include the `-----BEGIN CERTIFICATE-----` line and end with and include the `-----END CERTIFICATE-----` line.
5. On your ServiceNow instance, navigate to the **X.509 certificate** list by entering `sys_certificate.LIST` in the navigation filter.

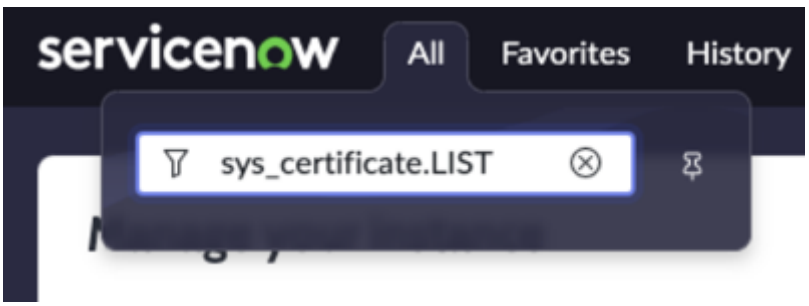

6. Select **New** to create a **X.509 Certificate** record.
7. In the form, fill in the following fields.

X.509 Certificate fields

Field	Value
Name	Name of the certificate. This name can be any name you choose.
Format	Select **PEM** **Note:** Privacy Enhanced Mail (PEM) files are a type of Public Key Infrastructure (PKI) file used for keys and certificates. The records you created in the previous steps are of this file type.
Type	Select **Trusted Store Cert**
Short Description	Description of the certificate. Enter a value that lets you know what this certificate is being used for.
PEM Certificate	Paste the certificate information that you copied in step 4.

8. Select **Submit** to save the record.

Create a secret group with criteria

Secret groups organize secrets and apply access policies at the group level. Associate a secret group with an identity group to control MID Server access.

Before you begin

Role required: admin, sn_kmf.admin, sn_secrets.admin

Procedure

1. Navigate to **All** > **Secrets Management** > **Secret Groups with Criteria**.
2. Select **New** to create a **Secret Group with Criteria** record.
3. In the form, fill in the following fields.

Secret Group with Criteria fields

Field	Value
Group name	Name of the secret group. This name can be any name you choose.
Secret type	Select **Client accessible**
Autogen module	Check the box
Short description	Description of the secret group. Enter a value that indicates what this group is being used for.
Criterion type	Select **Target table**
Target table	Select **SSH Credentials [ssh_credentials]**

4. Select and hold (or right-click) the form header and select **Save** from the context menu to save the record.
5. Make sure that the **Active** check box isn't checked.
6. Select the **Manage client side access** button to create an identity group.
   An **Identity group alias** window displays.
7. Select the **New** button.
8. Select the reference icon ( ) next to the **Identity group alias** field.
9. In the **Group Alias Name** field, enter a value.
   This name can be any value you want.
10. If you chose a unique name in Step 9, select **Submit**.
11. Select the **Upload identity key** button.
    The **Import identity public key certificate** window displays.
12. In the **Identity key alias** field, enter a value.

    Important: This value can be anything you want, but it must be an exact match for what you insert into the MID Server in later steps.

13. Select **Import**
    The **Attachment** window displays.
14. Select **Choose file**.
15. Select the public certificate that you created in the earlier steps.
    This certificate should be the `sm_public_cert.pem` file.
16. Select the X icon to close the window.
17. Select **OK** to close the **Import identity public key certificate** window.
    A blue **Keys and certificates are successfully imported to the instance** banner displays confirming a successful import.
18. Select **Submit**.
    The **Identity groups** list displays.
19. Select the check box to the left of your identity group record in the list.
20. Select the **Associate secret group** button.
    You're returned to your **Secret Group with Criteria** record. The **Secret Group – Identity Group Alias**" and **Secret Group – Identity Group** related lists are visible. These related lists display the records that you created in the previous steps.
21. In the **Secret Group – Identity Group** related list, select the **Created** field for the record on that list.

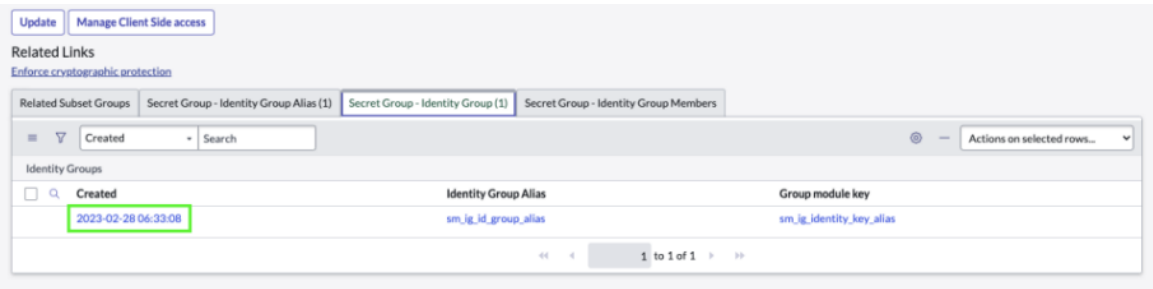

    An **Identity group** record displays.
22. In the **Identity group members** related list, select the **New** button.

A **Identity group member** record displays.

23. In the **Member table** field, select **MID Server [ecc_agent]**.
24. Select the reference icon ( ) next to the **Identity group member record** field, and select your MID Server.

    Note: If you enable the **Include all records** check box, all MID Servers connected to your instance are added to the identity group.

25. Select **OK** to close the **Select the document** window.
26. Select **Submit**.
27. Navigate back to **All** > **Secrets Management** > **Secret Groups with Criteria** and open the record you created in step 2.
28. Enable the **Active** field.
29. Select **Update** to save the record.

Upload the public/private keypair to the MID Server

Upload your public/private keypair to your MID Server. This keypair enables the MID Server to handle authentication requests from your instance.

Before you begin

Role required: none

Since ServiceNow lacks access to the private key, it can't pair it with the public key to decrypt the symmetric key and then decrypt the credential. If the MID Server tries to use this encrypted credential, it's unable to decrypt the credential for authentication without access to the private key.

In these steps, you upload the private key to the MID Server to complete the Public/Private keychain. This upload grants the MID Server access without giving ServiceNow access.

To grant the MID Server access to the private key, you must construct a command to be run as administrator in Powershell. In this example, the command is for the Azure Windows virtual machine.

Important:

Ensure that the system you perform these steps on has access to both the MID Server and the keypair file.

Procedure

1. In your local environment, locate the folder where you created your key pair in the Create encryption keys and certificate steps.
2. Find and copy the full path to the `manage-certificates.bat` file.

   Note: This file is located on your MID Server folder. Depending on where you've stored your MID Server folder, your path may look like this example:

   ```
   C:
   \Users\<your_user_account>\Documents\SM_Implementation\mid.
   utah-07-08-2022__patch4b01-31-2023_02-07-2023_1702.windows.
   x86-64\sm_ig_MIDS\bin\scripts\manage-certificates.bat
   ```

3. Create a text file and paste the path into the file.
4. In the text file, add the following after the path:
   `-a your_identity_key_alias`

   Replace `your_identity_key_alias` with the name of the identity key alias that you created when you uploaded your public certificate.
5. Find and copy the full path to your key pair file.

   Note: If you used the names in these steps that file is named `sm_keypair_bundle.pem`.
6. In your text file, add this path to the end of the line, adding a space between this path and the previous information.
   The text within your text file should look similar to this example:

```
C:
\Users\<your_user_account>\Documents\SM_Implementation\mid.utah
-07-08-2022__patch4b01-31-2023_02-07-2023_1702.windows.x86-64\s
m_ig_MIDS\bin\scripts\manage-certificates.bat -a
 your_identity_key_alias
 C:\Users\<your_user_account>\Desktop\sm_keypair_bundle.pem
```

   Note: In this example, the `sm_keypair_bundle.pem` file is on the desktop to make the path shorter.
7. Copy the entire text of your text file to the clipboard.
8. Find Powershell on your system, and choose the **Run as Administrator** option.
9. Paste the text from your text file into Powershell, and press Enter.
   If successful, you can see the following message:

```
Installed certificate with alias: <your_identity_key_alias>
 into the MID keystore.
```

   Tip: If you don't see this message, ensure that your command has no errors, spaces, or unnecessary quotes. Ensure that the full path is typed correctly.
10. Restart your MID Server by navigating to your MID Server record, right-clicking the record, and selecting **Restart MID**.

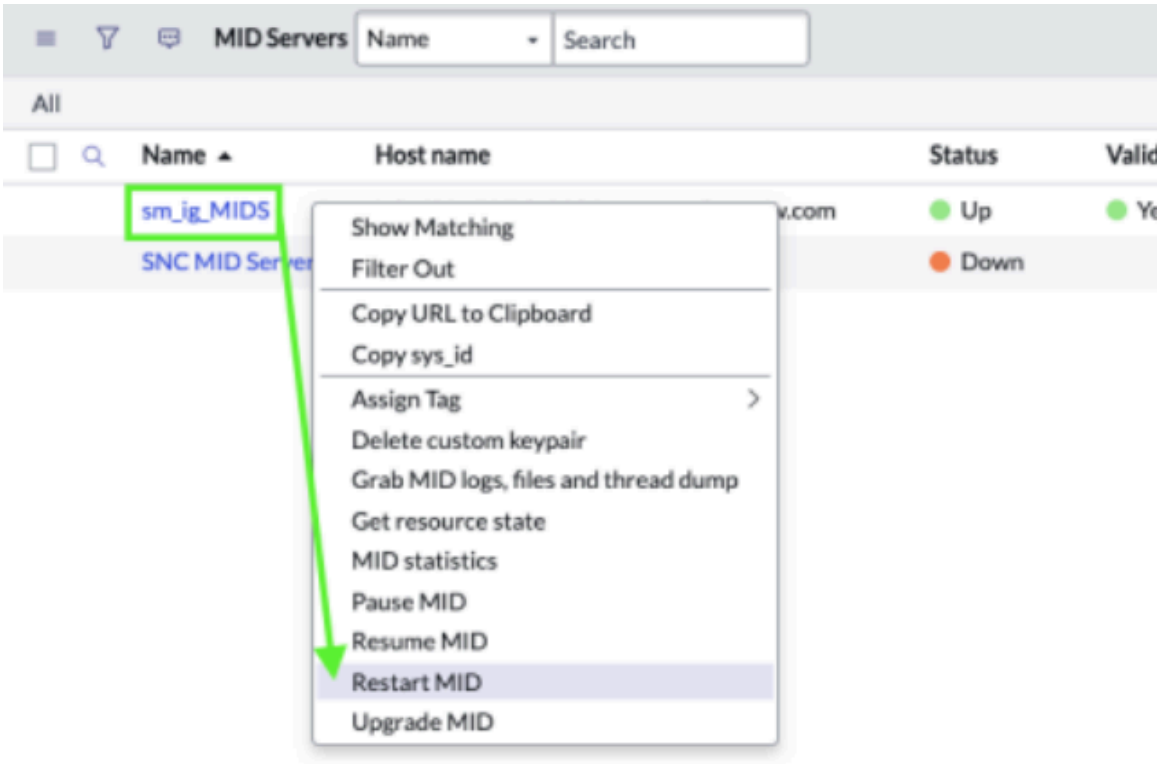

Restarting the MID Server synchronizes the uploaded key pair in the MID Keystore for use with operations. Wait for the MID Server to restart, with a status of **Up** and a validated value of **Yes** before continuing.

Create credentials and test credential encryption

Create a credential to authenticate into a third-party system.

Before you begin

Role required: admin, KMF_admin, sn_secrets.secret_manager, and sn_kmf.cryptographic_manager

Important: The logged in user of the MID Server also is required to have the following roles: admin, KMF_admin, sn_secrets.secret_manager, and sn_kmf.cryptographic_manager.

Procedure

1. Navigate to the **SSH credentials** list by entering `ssh_credentials.list` in the navigation filter.
2. Select **New** to create an **SSH credential** record.
3. In the form, fill in the following fields.

SSH credential form

Field	Value
Name	Enter a name for your credentials record. This name can be any value you want.

Field	Value
Applies to	Select **Specific MID Servers**
MID Servers	Select your MID Server.
User name	Enter the user name.
Password	Enter the password related to the user in the **User name** field.

4. Select the lock icon ( ) next to the **Credential alias** field.
5. Select the reference icon ( ) to open the **Connection & credential aliases** list.
6. Select **New** to create a **Connection & credential alias** record.
7. Enter a name in the **Name** field.
8. Select **Credential** in the **Type** field.
9. Select **Submit**.
   You're returned to the **SSH Credentials** record. In the next steps, you test that the credential is encrypted.
10. Select and hold (or right-click) the form header and select **Show XML**.
11. Find the `<password>` XML tag within the XML.
12. Copy the sys_id within this `<password>` tag to your clipboard.
    The sys_id is a 32 character code representing the symmetric key that is being used to encrypt this credential. The encrypted password that you entered into the SSH Credentials table is to the right of the two sets of boxes on this same line.

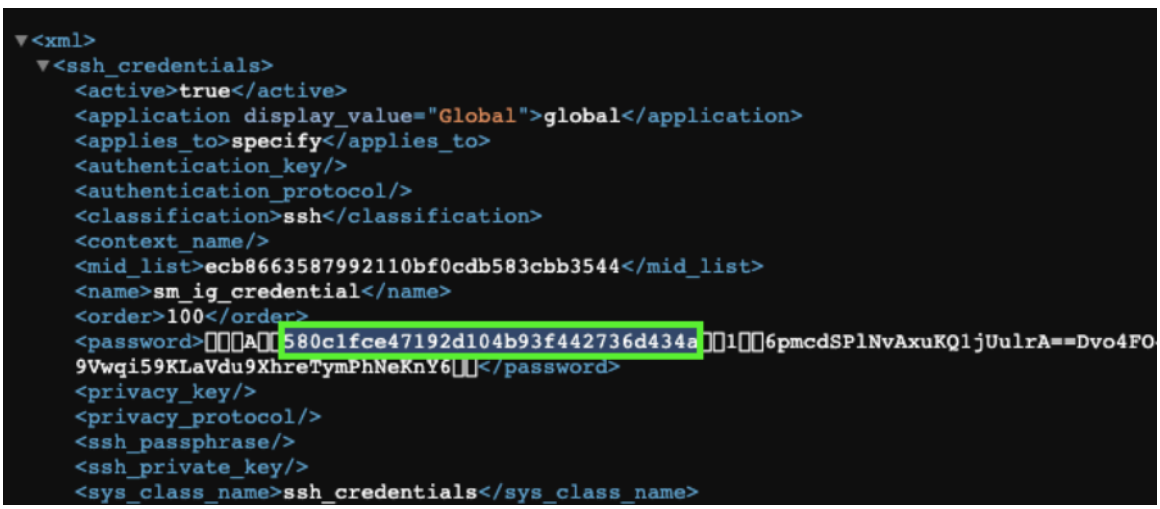

13. Navigate to the **Module keys** list by entering `sys_kmf_module_key.list` in the navigation filter.
14. Filter the list for records where the **Sys ID** field matches the sys_id you copied in step 12, and select **Run**.

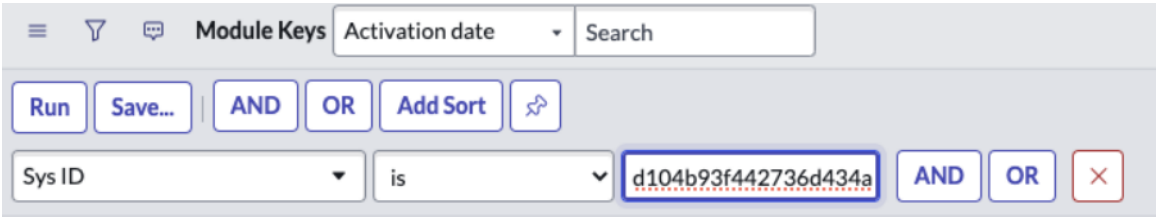

Your search should return a single **Module key** record. This record shows you that you have successfully created, and are using, a symmetric key.

15. Navigate to the **Wrapped module keys** list by entering `sys_kmf_wrapped_module_key.list` in the navigation filter.
16. Filter the list for records where the **Crypto module** field matches the cryptographic module you created in the earlier steps, and select **Run**.
    Your search should return a single **Wrapped module key** record. From this list you can verify the following:
    - The **Wrapped key material** column shows that the symmetric key within the Crypto Module (that is being used to encrypt the SSH credential) is encrypted by the public key that you uploaded to the Identity Group.
    - The **Wrapped key sys id** field shows that it's the key (the Crypto Module symmetric key) that is being encrypted by the **Wrapping key sys id** (the Public Key uploaded to the Identity Group).

    If the preceding fields aren't on your list by default, you can add them to the list by selecting the **Personalize list** icon ( ).

    Note: Your instance executes a clean-up job every 10 minutes, which removes orphaned keys and helps prevent unassociated key proliferation after you update your credentials.

Configure Flow Designer to manage the integration

On your instance, use Workflow Studio to manage an integration between your local network and your instance.

Before you begin

Role required: admin, KMF_admin, sn_secrets.secret_manager, and sn_kmf.cryptographic_manager

In these steps, you create a Workflow Studio workflow to create a text file on your local system.

Procedure

1. On your instance, navigate to **All** > **Process Automation** > **Flow Designer**.
2. Create an action in ServiceNow by selecting **New** and selecting **Action**.

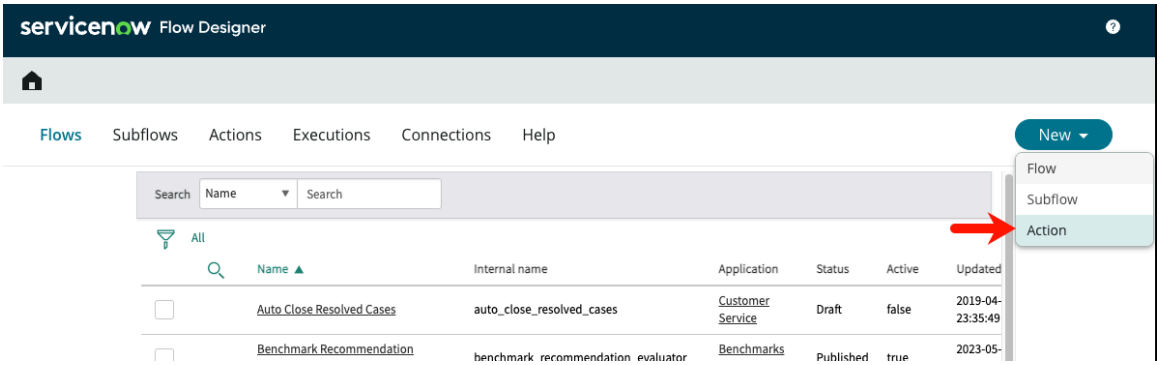

3. Enter a name in the **Action name** field, and select **Submit**.
4. Create a step by selecting the plus sign between **Inputs** and **Error Evaluation** in the **Action Outline**.

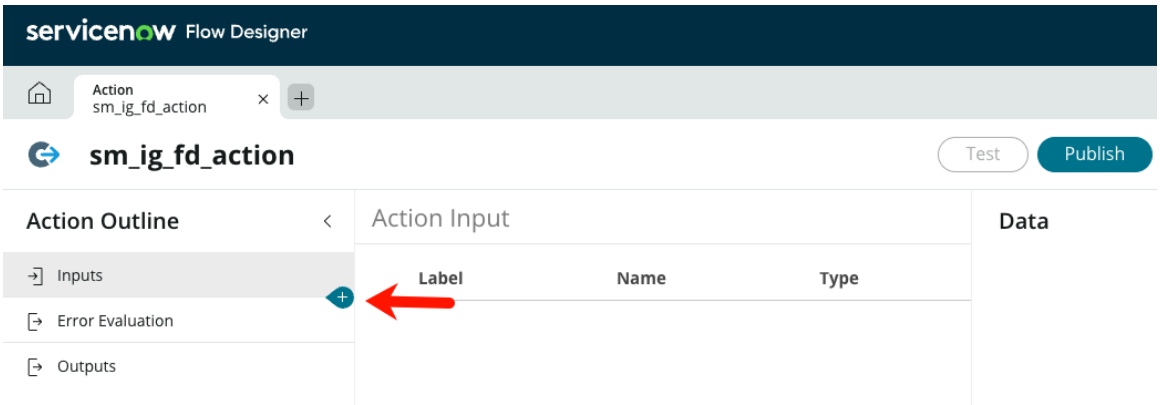

5. In the **Choose a step to add to your action** window, select **Look Up Record**.

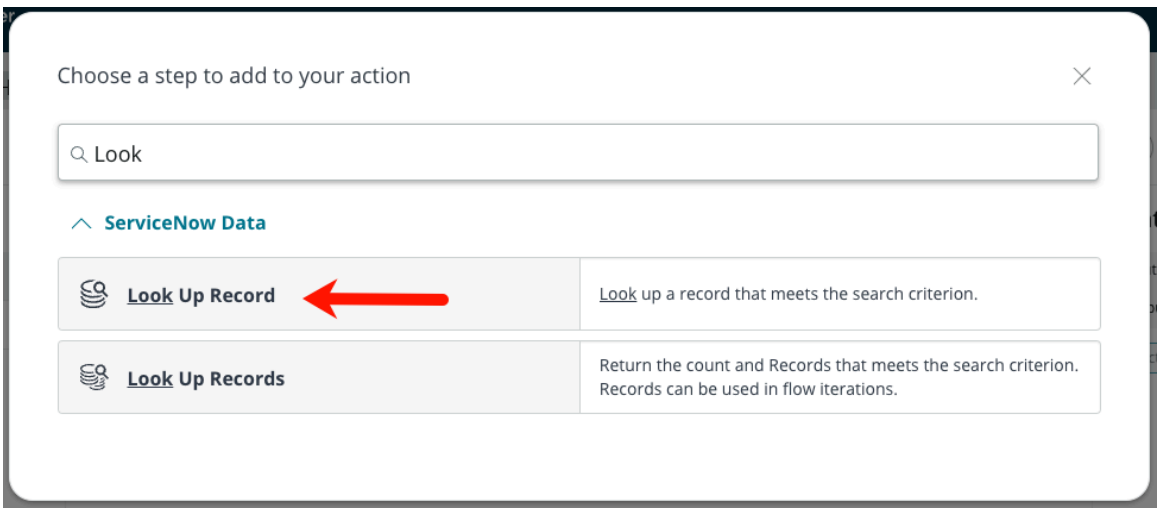

6. In the **Look Up Record step** section, select **MID Server [ecc_agent]** in the **Table** field.
7. Create another step by selecting the plus sign under your **Look Up Record step**.
8. In the **Choose a step to add to your action** window, select **SSH**.

Note: If you don't see the **SSH** option, you must activate the required plugin.

9. In the **SSH Step** section, enter the following information:

Field	Value
Connection	Select **Define Connection Inline**
Credential Alias	Select the credential alias for the SSH credential that was created in earlier steps.
Host	Enter the IP address for the host that you're connecting to via SSH.
Port	Enter `22`.
MID Selection	Select **Specific MID Server**.

10. To fill the **MID Server** field, drag the **MID Server Record** pill from the **Data** section into the field.

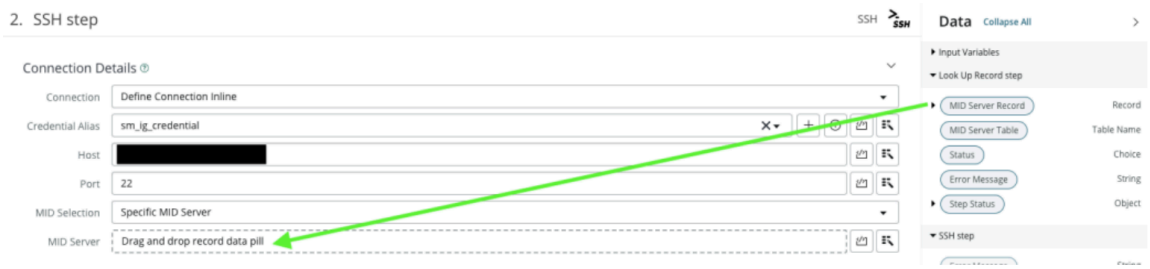

> **Warning:** When dragging the pill into the field, select the pill and not the black arrow next to the pill.

11. In the **SSH Configuration** section, enter the following value into the **Command** field.

```
/bin/date > sm_ig_text_file.txt
```

This command creates a text file in your local system using the decrypted secrets from the MID Server. The MID Server grants access to the ServiceNow instance (through Workflow Studio) without ever giving the ServiceNow instance access to the decrypted secret.

> **Tip:** The `/bin/date` command is inserting the current date/time into the created text file. This command demonstrates that the integration is happening in real time based on the current date/time versus when the text file was created.

12. Select the **Save** button to save the workflow.

Test the end-to-end client-side encrypted secrets integration

Test your integration, and review the execution details to confirm your configuration is working.

Before you begin
Role required: admin, KMF_admin, sn_secrets.secret_manager, and sn_kmf.cryptographic_manager

Procedure

1. In Workflow Studio select the **Test** button in the top-right corner of the screen.

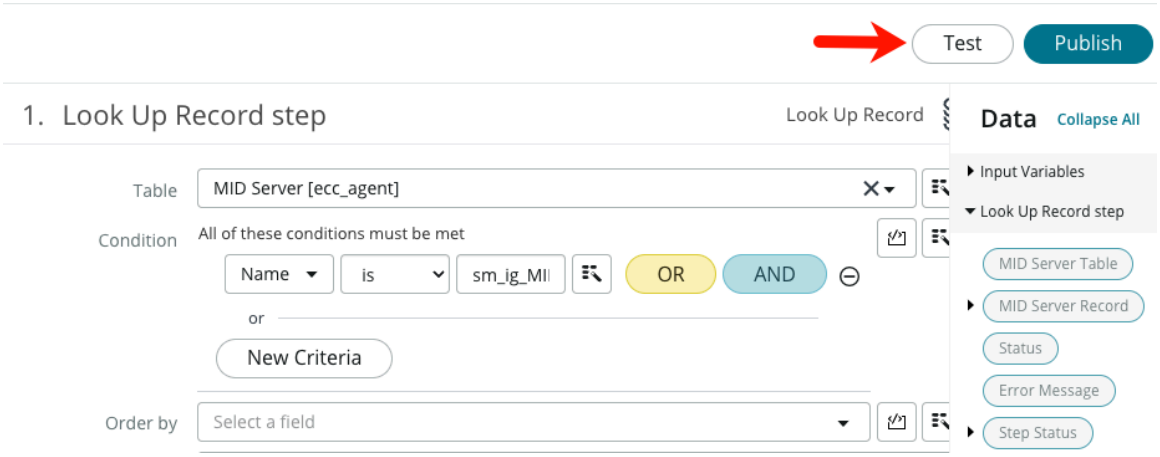

2. In the **Test Action** window, select **Run Test**.
3. Select **Your test has finished running. View the Action execution details.**
4. Refresh your screen until you see **Test Run – Completed** in the top-right side of the screen.
5. Select the **Steps** arrow at the bottom left of the screen.
6. Scroll down until you see a **Step Output Data** heading with the following success message:

```
{"Step Status";{"code";0,"message";"Success"}}
```

7. After seeing this message in Workflow Studio, verify that your text file has been created in your local system.

Test a Windows Management Instrumentation credential encrypted with Secrets Management

Verify that your Windows Management Instrumentation (WMI) credential is encrypted with Secrets Management and use an Integration Hub workflow to complete end-to-end testing.

Before you begin

Role required: admin, KMF_admin, sn_secrets.secret_manager, and sn_kmf.cryptographic_manager

- Integration Hub Enterprise must be active on your instance. For details, see Request an Integration Hub plugin.
- You must have a configured secrets group on your instance. See steps 1 through 3 listed in Configuring client accessible secrets.
- You must have a configured credential encrypted with Secrets Management. See steps 1 through 9 listed in Create credentials and test credential encryption.

Procedure

1. Confirm that the desired MID Server is associated with the Secret Group.
   Open the Secret Group with Criteria [sn_sm_criteria_secret_group] record, and look for your MID Server in the **Secret Group – Identity Group Members** list. If your MID Server isn't associated with a group see Create a secret group with criteria.
2. Confirm that the credential is encrypted with Secrets Management.

a. Navigate to **All** > **IntegrationHub** > **Connections & Credentials** > **Credentials**.

b. Open your credential record.
Verify that you have a credential alias listed in the **Credential alias** field. If you have not created one already, see Create a Connection & Credential alias.

c. Select and hold (or right-click) the header of your credentials record, and select **Show XML** from the list.

d. In the XML view of the record, find the password and confirm that the value starts with □□□A□□.

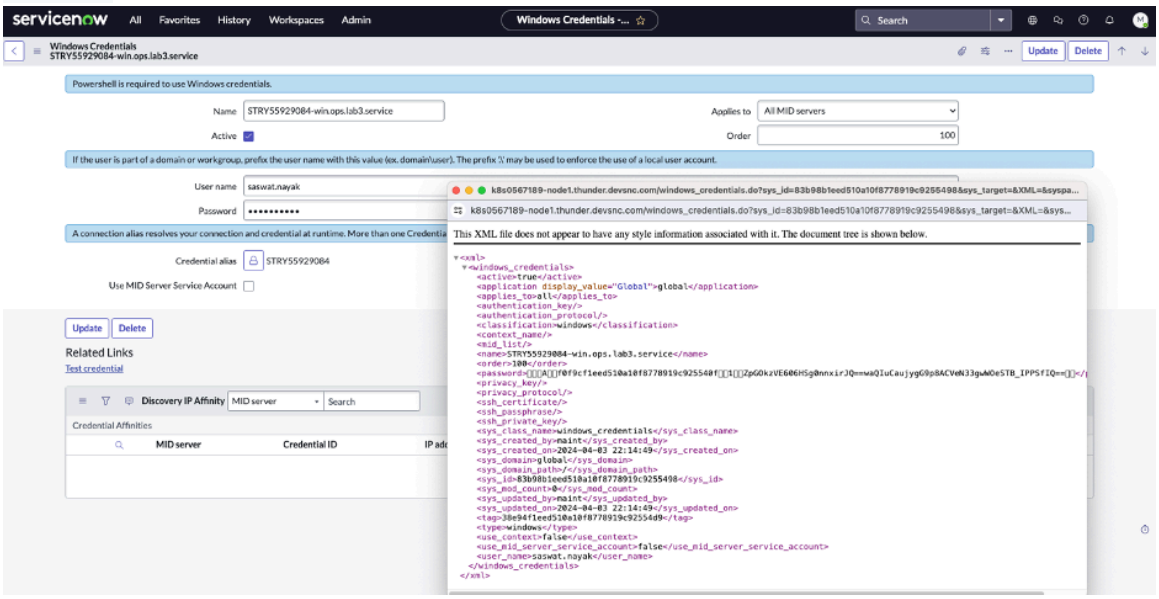

3. Create a testing workflow in Workflow Studio.

a. Navigate to **All** > **Process Automation** > **Workflow Studio**.

b. Select **New**, then select **Action** from the list.

c. Enter a name in the **Action name** field.
Leave the **Application** field set to Global.

d. Select the **Build action** button.

4. Configure a record lookup step in Workflow Studio.

a. In the **Action Outline**, select the plus button to add a new step.
Find and select the **Look Up Record** step type.

b. In the **Table** field, select **MID Server [ecc_agent]**.

c. For the condition, select **Name is**, followed by the name of your MID Server.

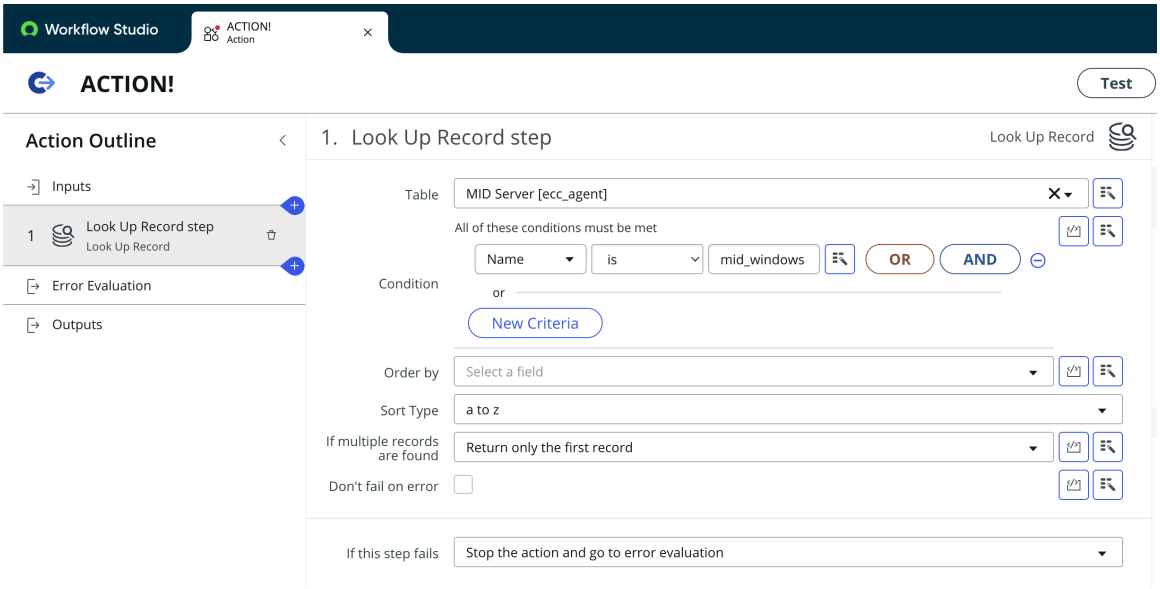

Important: Make sure that the selected MID Server is the same as the one associated with the secret group.

5. Configure a record lookup step in Workflow Studio.

   a. In the **Action Outline**, select the plus button to add a new step below your Look Up Record step.
   Find and select the **PowerShell** step type.

   b. In the **PowerShell Step** form, fill in the fields as needed.

PowerShell step fields

Field	Value
Connection	Select **Define Connection Inline**
Credential Alias	Select your credential alias
Host	Enter the IP address of the target Windows server
Port	Enter 22 if it isn't already the default.
MID Selection	Select **Specific MID Server**
MID Server	Drag in the **MID Server Record** pill from the **Data** panel This data pill is visible in the **Data** panel on the right edge of the screen, under **Look Up Record step**.
Remoting Type	Select **Run on a MID Server or have your script establish a remote session**

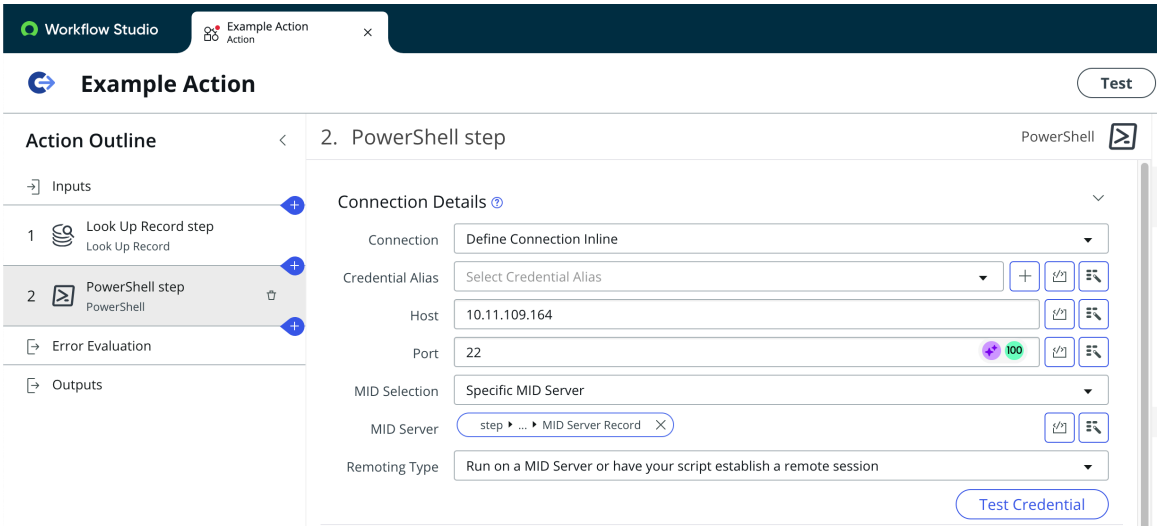

6. Create a script for your test action.

   a. In the **Script Type** field, select **Inline Script**.

   b. Use the following script, replacing the example path with the path from your test host.

```
Get-Date | Out-File -FilePath
 C:\Users\<path>\sm_ig_text_file.txt
```

   Note: You may also run your own script that suits your remote host configuration. It's important to confirm that a connection was established using the SM-encrypted value.

   c. Select the **Save** button to save your action.

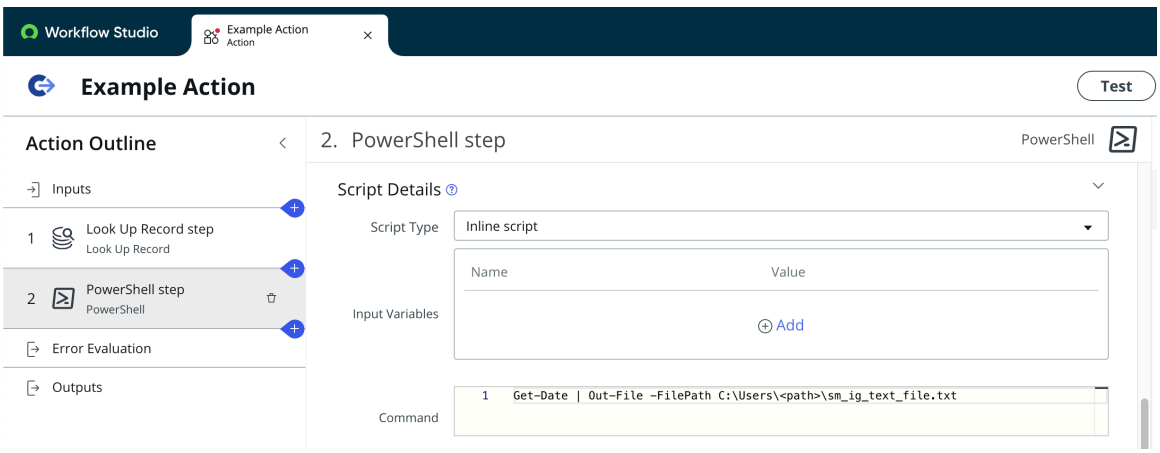

7. Select the **Test** button to test your action.
8. Review the output logs for any error messages.

9. Verify that a test file has been created on the host server in the file path you provided in the script.

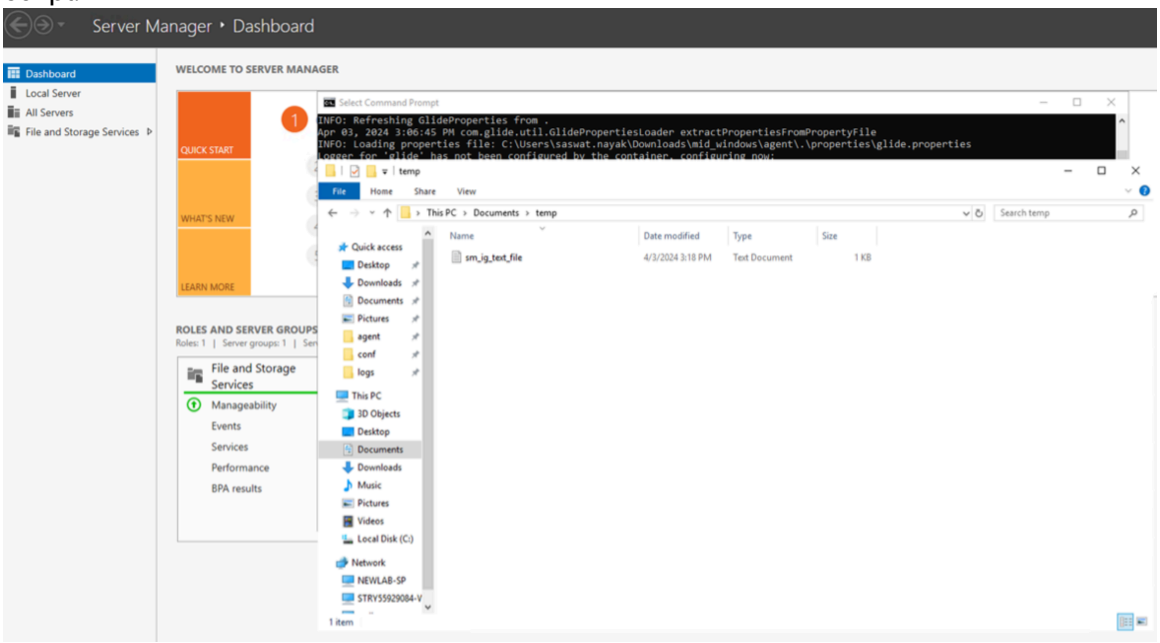

Cloning and Secrets Management

Learn how to reconfigure secrets groups and client secrets groups after a clone.

After you clone an instance, your secrets groups and client secrets groups may require reconfiguration to operate as expected.

Secrets group origin	Behavior after a clone
Original instance secrets groups from the target instance	Works as expected after manually importing the missing cryptographic module. See Exporting and importing data via XML.
Original client secrets groups from target instance	Works as expected after manually importing the missing cryptographic module. See Exporting and importing data via XML.
Cloned instance secrets groups from source instance	Doesn't work unless manually set up
Cloned client secrets groups from source instance	Work if manually config sg/identity group/ alias/mid

Secrets Management dashboard

Use the Secrets Management dashboard to review the secret groups configured on your instance and learn about any security issues.

Secret Group Overview

The **Secret Group Overview** tab displays information about your configured secret groups. Use this tab to see information about the secrets groups configured on your instance.

Secret Group by Type

Displays a pie chart showing secret groups installed on your instance, grouped by secret type (instance side of client side).

Secret Group by Criterion Type

Displays a pie chart showing secret groups with criteria configured on your instance by the type of criteria used.

Number of Inactive Secrets Groups

Displays a count of the inactive secret groups configured on your instance.

Number of Dedicated Secrets

Displays a count of secrets within basic secret groups.

Secret Group Warnings

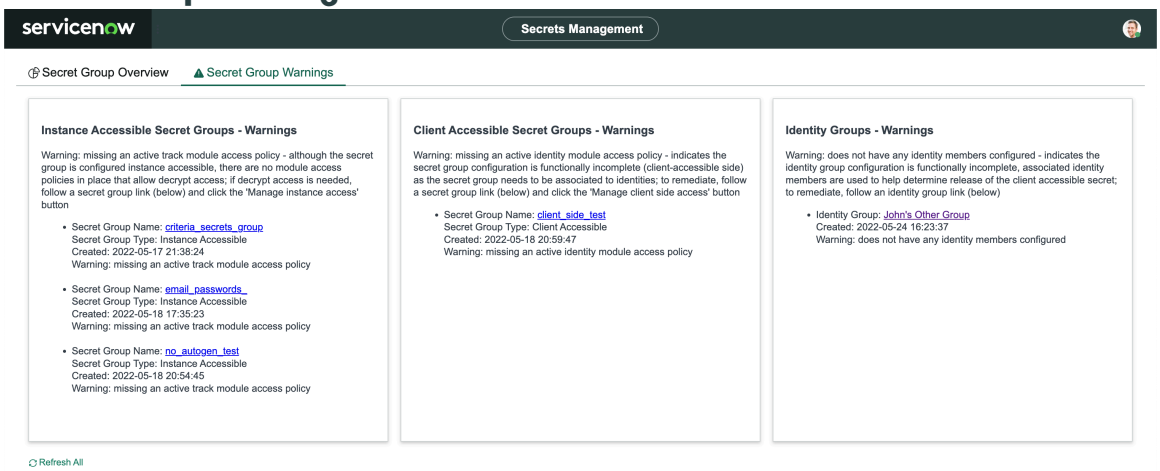

The **Secret Group Warnings** tab displays warnings related to your secret groups and identity groups.

Instance Accessible Secret Groups- Warnings

This card displays warnings if there are secret groups with no active access policies in place. Select a secret group name to view that record.

Client Accessible Secret Groups- Warnings

This card displays warnings if there are client accessible secret groups that don't have an active identity module access policy (MAP). Select a secret group name to view that record.

Identity Groups - Warnings

This card displays warnings if there are identity groups that don't have group members configured. Select the identity group name to view the record.

Note: The Secrets Management Dashboard is a part of Secrets Management Enterprise. Secrets Management Enterprise is a Secrets Management Enterprise. Secrets Management our production instance.

Secrets management roles

Secrets management adds these roles.

Secrets administrator [sn_secrets.admin]

Assign non-admin secrets roles to other users. Secrets administrators have the same privileges as secrets manager and viewer.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Important: Avoid granting an admin role when more specialized roles are available.

- A user must have both the admin and security_admin roles to be assigned the sn_secret.admin role.
- Avoid granting an admin role when more targeted roles are available.

Secrets manager [sn_secrets.secret_manager]

Grant this role to users who must be able to perform any of the following functions.

- View secret and secret group records
- View access history, and other usage activity info
- Create secret groups and filters
- Create secret providers
- Move secrets across secret groups
- Change secret group and secret provider settings

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Important: Secrets managers can't see secrets in plain text.

Secrets viewer [sn_secrets.viewer]

Grant this role to users who must be able to view secret and secret group records.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

None.

Create a secret group cryptographic module

Create a secret group cryptographic module to perform encryption and decryption.

Before you begin

Role required:

- admin
- sn_secrets.admin
- sn_secrets.secret_manager
- sn_kmf.cryptographic_manager
- sn_kmf.admin

Procedure

1. Navigate to **All** > **Secrets Management** > **Create Secret Group Crypto Module**.
2. Select the type of group crypto module you would like to create.

Cryptographic module type	Description
Create an instance accessible secret crypto module	Create instance accessible secrets that can be decrypted by your instance.
Create a client accessible secret crypto module	Create client accessible secrets that are encrypted using your own key, which ServiceNow can't access.

3. In the **Cryptographic Module** form, fill in the fields.

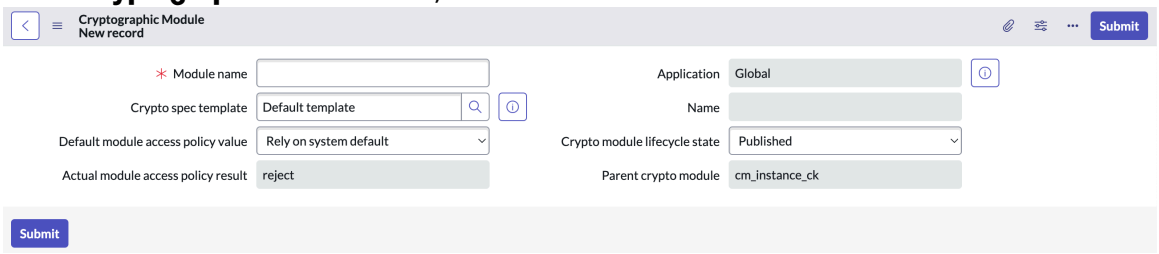

Field	Description
Module name	Descriptive name for your module
Application	The application scope that contains your module. This field is automatically populated with the current module.
Crypto spec template	**Default Template** is selected by default.
Name	Name of the module. This name is automatically generated by based on the application and module name.
Default module access policy value	○ Rely on system default ○ Reject ○ Track
Crypto module lifecycle state	Whether the module is in a **Draft** or **Published** state.
Actual module access policy result	This field is for information purposes, and is read-only.
Parent crypto module	The parent cryptographic module, which is determined by the cryptographic module type you selected in the previous step. This field is read-only.

4. Select **Submit**.

Create a basic secret group

Create a basic secret group to group any secrets, regardless of their criteria.

Before you begin

Role required: admin

Basic secret groups can contain any secret you add to them, regardless of their table, scope, or application. After creating it, you manually add secrets to the group. To create a group specifically for all secrets that share a common attribute such as those, create a secret group with criteria using the instructions in Create a secret group with criteria.

Procedure

1. Navigate to **All** > **Secrets Management** > **Secret Groups**.
2. Select **New**.

3. For the **What type of Secret Group would you like to create?** option, select **Basic Secret Group**.
4. In the **Secret Group** form, fill in the fields.

Secret group fields

Field	Description
Group Name	Name for the group **Note:** Secret group names can only contain lowercase characters, numbers, and underscores(_)
Secret Type	Whether the group is **Instance accessible** or **Client accessible**.
Autogen Module	Generates a new cryptographic module for this secret group. This module encrypts and decrypts your data. This field is enabled by default.
Application	Scoped application for this record. This read-only field is automatically populated with the current scope.
Short Description	Description of the group
Crypto Module	Select the cryptographic module to use with this group. This module encrypts and decrypts your data. This field is only visible when **Autogen Module** isn't selected. For details on module access policies, see Module access policy overview **Note:** You can review the module access policies related to your secret group using the **Manage instance access** button.

5. Click **Submit**.

Note: When created, a secret group is inactive by default. Return to the group record and select **Active** to activate the group.

Create a secret group with criteria

Create a secret group with criteria to organize secrets entered in Password2 fields automatically when they share a common criteria, such as table, scope, or application.

Before you begin

Role required: admin, KMF_admin, sn_secrets.secret_manager, and sn_kmf.cryptographic_manager

Secrets within this type of secret group must all share common criteria. For groups without this restriction, consider creating a basic secret group. Learn about creating a basic secret group in Create a basic secret group.

Procedure

1. Navigate to **All** > **Secrets Management** > **Secret Groups**.
2. Select **New**.
3. At the **What type of Secret Group would you like to create?** prompt, select **Secret Group with Criteria**.
4. In the **Secret Group** form, fill in the fields.

Secret group fields

Field	Description
Group Name	Name for the group **Note:** Secret group names can only contain lowercase characters, numbers, and underscores(_)
Secret Type	Whether the group is **Instance accessible** or **Client accessible.**
Autogen Module	Generates a new cryptographic module for this secret group. This module encrypts and decrypts your data. This field is enabled by default.
Application	Scoped application for this record. This read-only field is automatically populated with the current scope.
Short Description	Description of the group
Criterion Type	The criteria the secrets in this group shares. ○ Scope ○ Package ○ Target table ○ Secret column ○ Filter record
Crypto Module	Select the cryptographic module to use with this group. This module encrypts and decrypts your data. This field is only visible when **Autogen Module** isn't selected. For details on module access policies, see Module access policy overview **Note:** You can review the module access policies related to your secret group using the **Manage instance access** button.

Note: Depending on your configuration, the **Crypto Module** might use an automatically selected value.

When the **Criterion Type** field is set to `Package`, and the **Autogen Module** field is selected:	The **Crypto Module** field is empty and read-only. An existing Password2 submodule is used. If a Password2 submodule isn't found, the instance level Glide Encrypter module is used.
When the **Criterion Type** field is set to `Package`, and the **Autogen Module** field is deselected: (The **Autogen Module** field can only be deselected by Enterprise users)	The **Crypto Module** field is editable, and admins can select a crypto module to use.

5. Select and hold (or right-click) the form header and select **Save**.

Note: When created, a secret group is inactive by default.

6. After saving the record, additional fields might appear based on how you've configured your group.

Additional secret group fields

Field	Description
Target Scope	Scope shared by the secrets to be assigned to this group. This field is only available when you select `Scope` in the **Criteria Type** field.
Target Package	Package shared by the secrets to be assigned to this group. This field is only available when you select `Package` in the **Criteria Type** field.
Target Table	Table shared by the secrets to be assigned to this group. This field is only available when you select `Table` or `Secret Column` in the **Criteria Type** field.
Application Scope	Application scope of the table selected in the **Target Table** field. This field is only visible when you select `Table`, `Filter Column`, or `Secret Column` in the **Criteria Type** field.
Secret Column	Table column that contains the Password2 secrets you include in this group. The fields available in this list are determined by the table selected in the **Target Table** field. **Note:** If there are no columns on the select table that contain secrets, this field only displays `– None --` as a selection.

Field	Description
Filter Column	The column on the table selected in **Target Table** you want to use as a filter. This field can't be a Password2 field.
Filter value	Value to filter by, based on the column selected in the **Filter Column** field.

Example: An instance accessible group containing all email account passwords for an email server

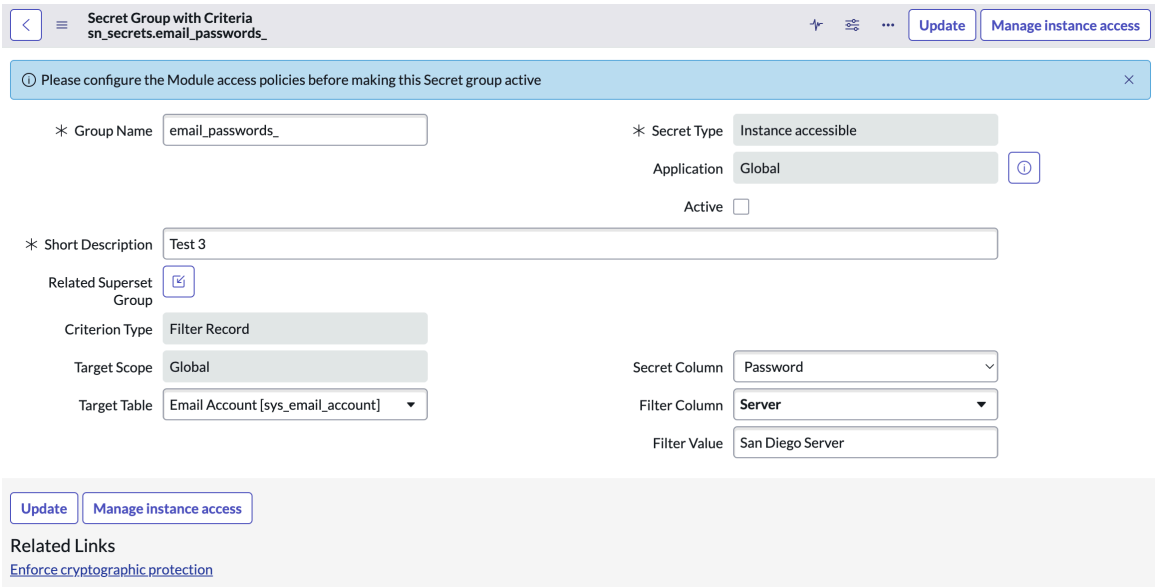

What to do next

After creating your group, any new records matching the criteria will be encrypted. To encrypt existing records using this group's cryptographic module, you must run a security job. For details, see Run Secrets Management security jobs.

Client-accessible groups need a customer-provided public key to encrypt your secrets. For steps on uploading this key, see Upload a public key for Secrets Management.

Upload a public key for Secrets Management

Upload a public key to encrypt your secrets in Secrets Management.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Secrets Management** > **Secret Groups** and open the secret group record.
2. In the record, select the **Manage Client Side Access** button.
   A new identity group record opens.
3. Select the **Upload Identity Key** button.
   A **Import identity public key certificate** window displays.
4. Enter an alias for your key in the **Identity Key alias** field.
5. Select the **Import** button, to upload the key from your local environment.
6. Select the **OK** button.

The **Group module key** field uses the identity key alias name.

7. Select **Submit** to save the identity group record.

Run Secrets Management security jobs

Schedule a Secrets Management job to perform encryption tasks on secrets fields on your instance.

Before you begin

Role required: sn_kmf.admin, security_admin, and sn_secrets.admin

To perform these steps, you must elevate to the security_admin role. For details on this process see, Elevate to a privileged role

Procedure

1. Navigate to **All** > **System Security** > **Security Jobs** > **Create New**.
2. At the **What type of Security Job would you like to create?** prompt, select **Secrets Management Job**.
3. On the form, fill in the fields.

Secrets Management Job form

Field	Description
Name	Name of the security job
State	The initial job state is New. After the job has been executed as scheduled, the state will update accordingly
Time window start	Start time for the job in 24-hour format. The job begins execution at the chosen time.
Time window end	End time for the job in 24-hour format. If the job isn't finished by this time, it continues during the next specified processing window until the job is complete.
Enforcement Level	Whether the job affects all tables, or a selection of specific tables or fields. Select from ○ All Tables ○ Specific Tables ○ Specific Fields ○ Specific Packages ⚠ **Warning:** Selecting the **All Tables** option may affect instance performance. Consider scheduling at non-peak hours.
Packages	The packages to include in this job. Encryption is applied to selected packages. This option displays only when the

Field	Description
	Enforcement Level field is set to **Specific Packages**
Tables	The tables to include in this job. Encryption is applied to all applicable fields within the selected tables. This option displays only when **Enforcement Level** is set to **Specific Tables**
Fields	The tables to include in this job. Encryption is applied to all selected fields. This option displays only when **Enforcement Level** is set to **Specific Fields**
Job Mode	Select from **Password2 to Secrets Management** Encrypt all Password2 fields within your secrets groups using the cryptographic modules defined in each group's module access policy. **Secrets Management to Password2** Re-Encrypt data in your secrets groups using password2 encryption. For details on this encryption type see Password2 encryption with KMF. **Secret Group Enforcement** Queries all data that should match the group selected in the **Secret Group** field. If all the data found by the query is already in the group, the job makes no changes. If the query finds data that is not yet in the group, the job re-encrypts this data within the Secret Group. **Note:** If the data found in this query is already encrypted and your instances can't decrypt that data, it isn't encrypted and added to the secrets group.
Secret Group	Secret group containing the secrets to encrypt. This field is only available when **Secret Group** is selected in the **Job Mode** field.
Force rekeying data	Creates a new encryption key to replace an existing key. The data decrypts using the older key and re-encrypts using the new key.

Field	Description
Summary	Displays information about the job progress. Summary also displays records that couldn't be encrypted by the job.

4. Select **Submit**.

What to do next

The job queries all data that should match the selected secret group. If all the data found by the query is already in the group, the job makes no changes. If the query finds data that is not yet in the group, the job re-encrypts this data within the Secret Group. (If the instance can decrypt it, which may not be the case for client side-encrypted secrets).

Code Signing

Use Code Signing to create digital signatures that prevent unauthorized or tampered External Communication Channel (ECC) queue records from being processed by MID Servers. This cryptographic verification helps maintain the integrity of integrations between ServiceNow and external systems.

Code signing and Circle of Trust

The Circle of Trust (COT) is a prerequisite for Code Signing that creates secure communication between your trusted and protected instances to ensure that only authorized users can access the Code Signing feature.

Multiple security measures help to prevent malicious actors from disabling or misusing code signing in the case a protected instance is compromised. As part of the defense-in-depth strategy, the COT uses the following components:

- Controls that restrict even the most powerful administrator accounts are established in the protected instance to help protect Code Signing processes and configuration.
- Trusted instances are required to work together with protected instances to establish the Circle of Trust relationship. At least one trusted instance is required, but multiple trusted instances may be configured to collaborate with the protected instance.

Circle of Trust overview

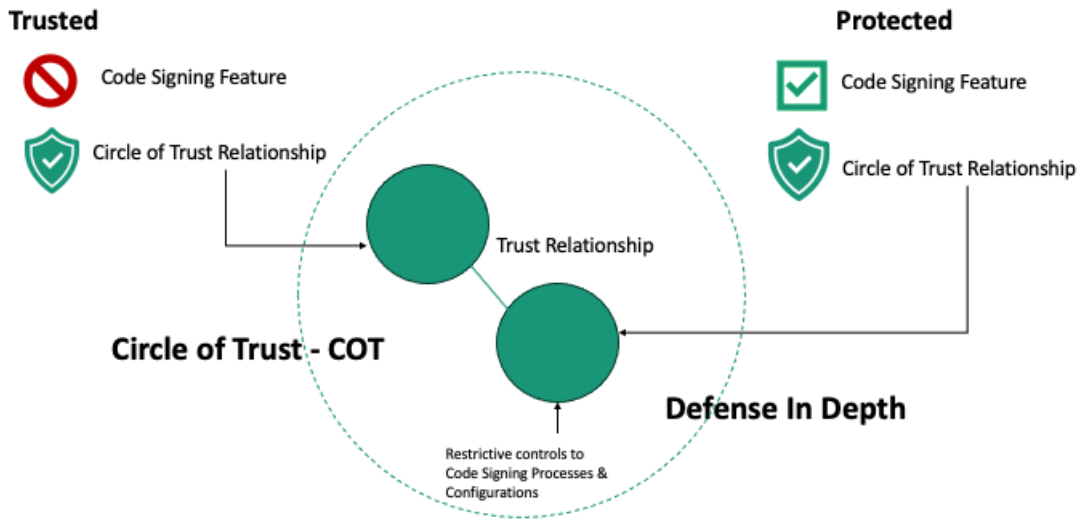

The Circle of Trust uses jobs, scripts, and business rules along with a key pair to generate signatures to sign update sets to the protected instance. When the job is called, the signature is verified along with the trusted certificate to execute protected instance updates.

Trusted update sets process

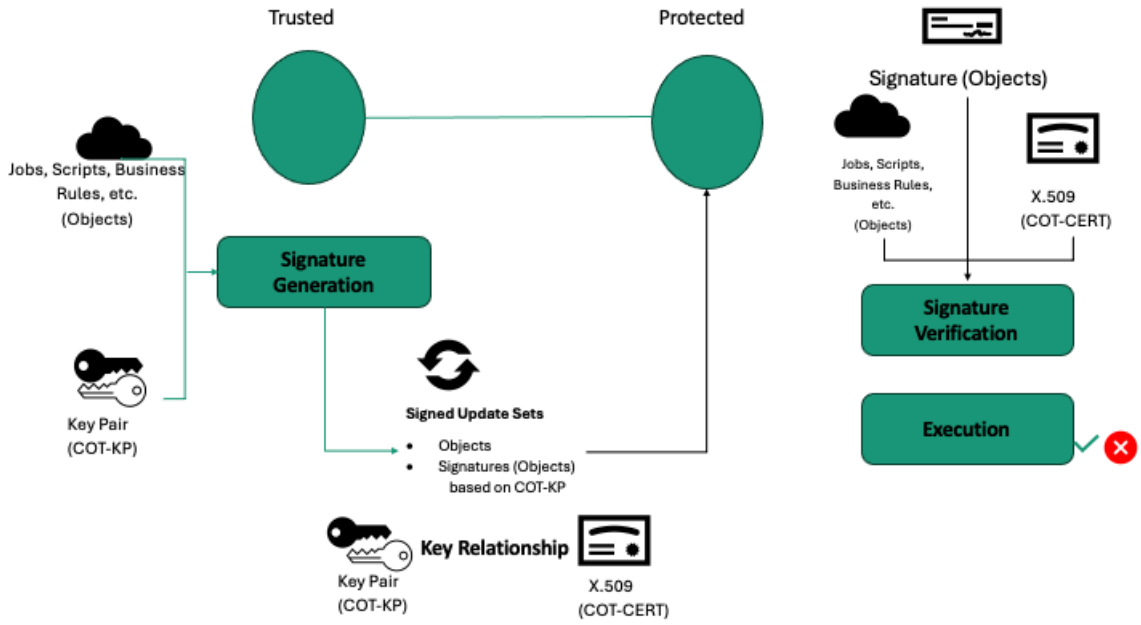

Code Signing flow

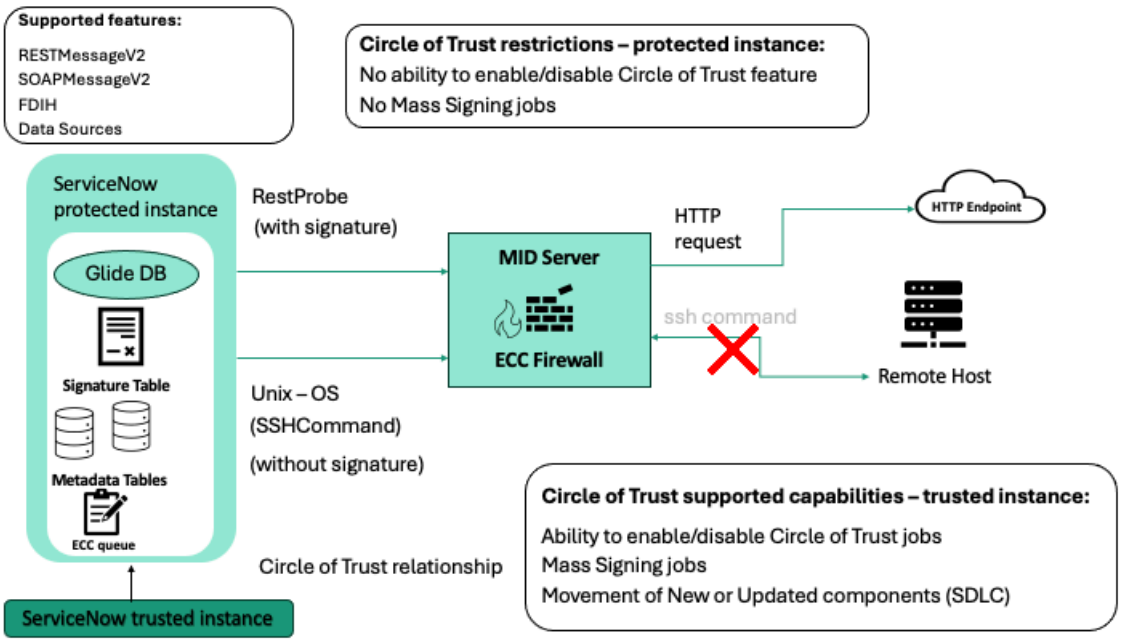

The Circle of Trust requires an initial trust relationship between trusted and protected instances that prevents any unauthorized user with any authorization level from accessing unapproved activities.

Get started

Explore	Configure	Reference
 Learn the key features and business value of Code Signing.	 Activate and configure Code Signing.	 Get details about properties and troubleshooting
	Use  Learn how to use Code Signing to help verify the authenticity and integrity of your data.	

Troubleshoot and get help

- https://www.servicenow.com/community/secops/ct-p/security-operations
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Code Signing

Code Signing provides cryptographic verification to ensure that only authorized scripts can execute on MID Servers. Code Signing prevents unauthorized or tampered ECC queue records from being processed by MID Servers, maintaining the integrity of integrations between ServiceNow and external systems.

Code Signing creates digital signatures for your data, which are later checked to confirm the authenticity and integrity of the data. Code Signing is a module licensed as a component of ServiceNow Vault.

Note: The Customer Service and Support team must grant access to Code Signing.

Code Signing declares the intent behind the operation being performed and validates whether the resource or record may be used for the intended purpose. To facilitate Code Signing, the Key Management Framework (KMF) uses digital certificates and industry standard asymmetric encryption for digital signatures.

Use Code Signing internally on the platform and infrastructure side. Code signing provides a way to sign the content of specific tables or of a subset of records in a given metadata table.

Code Signing uses a secure Circle of Trust (COT) between your trusted and protected instances to help ensure that only authorized, secure trusted instances can access the Code Signing feature.

Note: Code Signing is enabled on the protected instance and not on the trusted instance.

How Code Signing protects your environment

Without Code Signing, an attacker who gains access to ServiceNow records can modify SQL statements in a protected instance. When the MID Server processes this data source request, it would execute the malicious SQL commands, potentially compromising system integrity and security.

When you implement a Circle of Trust architecture with Code Signing, transfer of data to the MID Server follows the following verification process. This process helps ensure that only authorized code originating from the trusted instance can execute on the MID Server. The processes reduces potential attack vectors that could otherwise compromise your systems.

1. Digital signatures are applied to data sources created or updated within the trusted instance.
2. Use the Code signing process to transfer the signed data from the trusted instance to the protected instance
3. The MID Server verifies the digital signature on all incoming requests, automatically rejecting any requests lacking valid signatures.
4. If the MID Server rejects a request, it logs this rejection and sends a notification to the protected instance.

Benefits of implementing Code Signing

Code Signing provides several key advantages:

Execution Control

Only cryptographically verified scripts can run on MID Servers

Tamper Detection

Any modifications to signed records are immediately identified and blocked.

Automated Protection

The system handles security enforcement without requiring manual intervention.

Comprehensive Logging

All signature verification failures generate detailed audit records.

Code Signing validation and jobs

All the metadata tables with valid configurations are signed at build time using the Code Signing metadata plugin (*com.glide.code_signing*). Installing this plugin automatically installs the Code Signing OOB App Signatures plugin (*com.glide.code_signing.oob_apps_signatures*) which contains build time signatures for all relevant records in the trued-up ServiceNow® Store application versions. If you choose to sign tables, admin users with the Security administrator role have access to Code Signing jobs:

- Sign update sets.
- Mass sign records.
- Mass sign attachments.

Sign update set

This job signs records that match a signature configuration in the update set. The job also adds all the new signature records and the verification certificates to the update set.

KMF signature record for update set

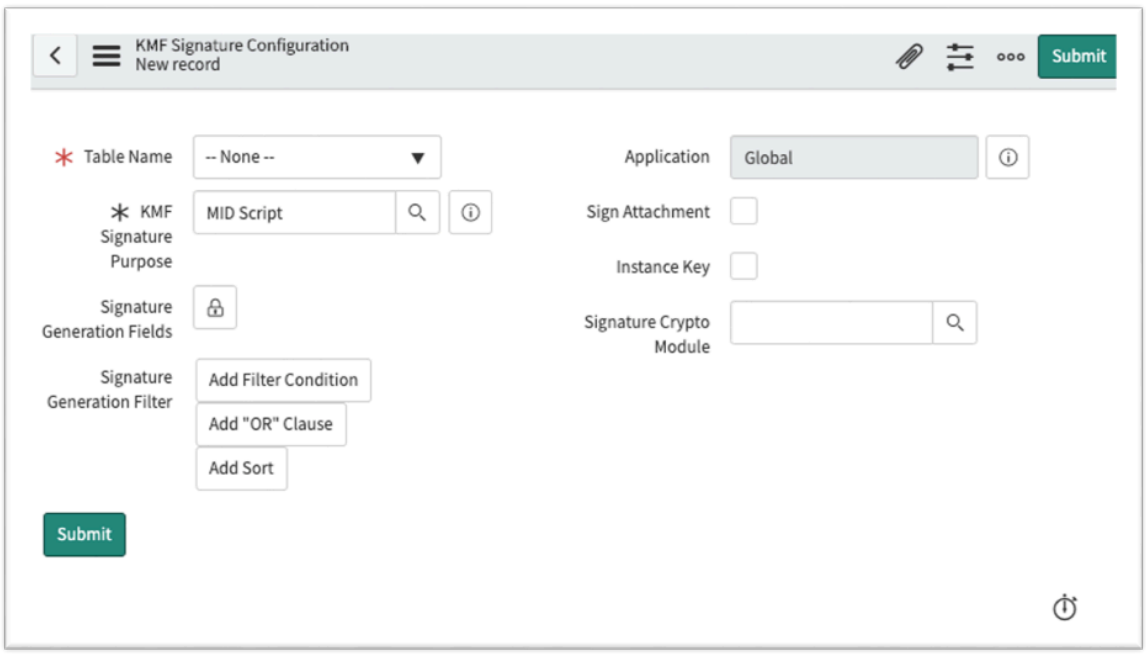

Mass sign records

This job signs all the records that match the signature configuration applied on a specific metadata table.

Mass sign attachments

This job signs all the attachment records that are attached to a table that matches a specified signature configuration.

Encryption job to mass sign records

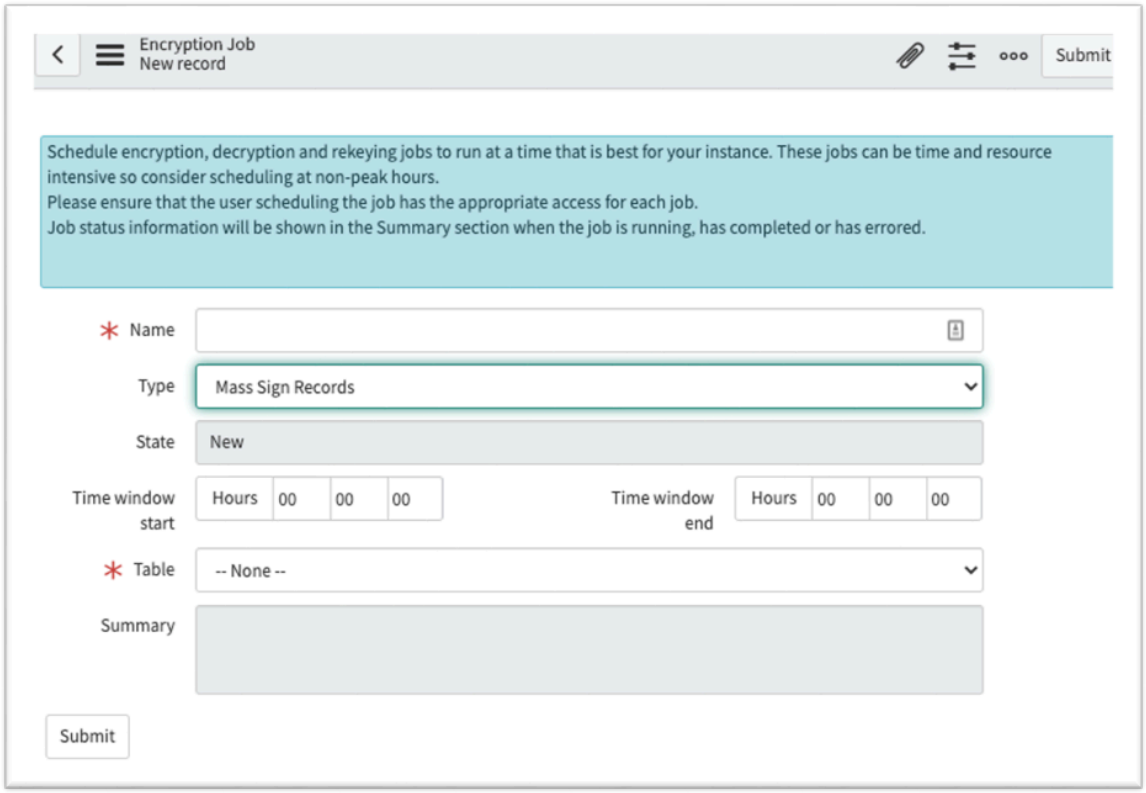

Configuring Code Signing

Activate and configure Code Signing to verify the authenticity and integrity of your data.

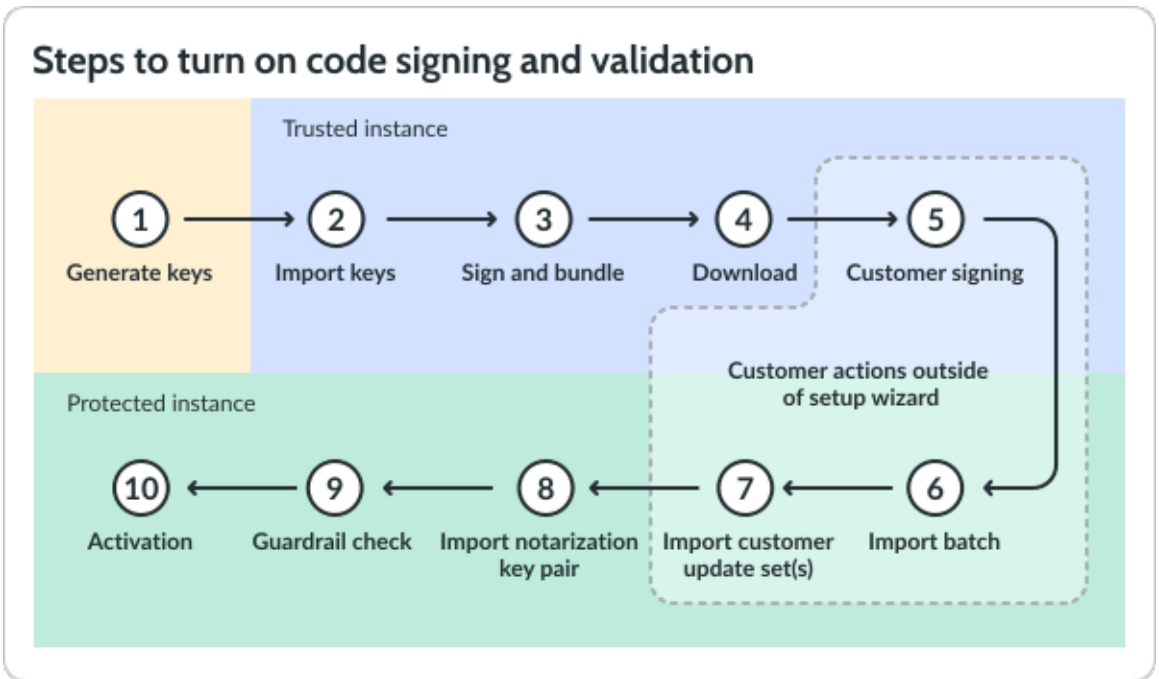

Code Signing Enterprise requires an initial trust relationship between trusted and protected instances that helps to prevent any unauthorized user with any authorization level from accessing unapproved activities.

Refer to each topic to complete the configuration steps to establish the Circle of Trust with Code Signing Enterprise:

Note: This optional step isn't part of the initial configuration for Code Signing

Assign the Code Signing Administrator Role

Assign the Code Signing Administrator role to a user to access the Code Signing configuration experience.

Before you begin

Roles required:

- admin
- security_admin

Important: The user that is assigned the Code Signing Administrator role must be in the Global scope.

Procedure

1. Select a trusted instance to use.
   The trusted instance is used to create jobs, sign records, and perform other necessary tasks before enabling signature validation on your protected Instance.
2. Navigate to **All** > **Code Signing** > **Administration** > **Role Administration**.
3. Grant a user the Code Signing Administrator role by moving a user from the **Available Users** list to the **Selected User(s)** list.
4. Select **Save**.

Important: The user assigned to the role must log out and lot in again to perform Code Signing tasks.

Result

The selected user has the Code Signing Admin role. A user with this role can access the Code Signing configuration experience, and assign other Code Signing roles to users.

To learn more about Code Signing roles, see Roles installed with Code Signing.

Configure Code Signing Enterprise on your trusted instance

Turn on and configure Code Signing on your trusted instance.

Before you begin

Roles required:

- admin
- security_admin
- codesigning_admin

> **Note:** The codesigning_admin role can be assigned using the process detailed in Assign the Code Signing Administrator Role.

- sn_kmf.cryptographic_manager

You must have at least one cryptographic key pair and certificate (p12 file extension) for customer signing and Circle of Trust (COT) administration. For improved security, consider using separate cryptographic key pairs for customer signing and COT administration.

> **Warning:** After completing this process, you will have downloaded a configuration file that must be installed on your protected instance within an hour. Verify that you have time available after this process to upload the configuration file to your protected instance. For details on that process, see Upload your Code Signing configuration file to your protected instance.

Procedure

1. On your trusted instance, navigate to **All** > **Code Signing** > **Configuration** > **Guided Setup** to open the Code Signing configuration page.
2. In the **Instance type** field, select **trusted instance**.
3. Select the **Next** button.
4. In the **Action** field under **Select the action you'd like to accomplish**, select **Turn on Code Signing**.
5. Next to **Attachments** in the **Customer signing key pair and certificate** section select **+Add File** to upload a cryptographic key pair (p12 file extension) to use for customer signing.

   > **Tip:** If the **+Add File** option is not available, verify that you are in the Global scope, and that you have the sn_kmf.cryptographic_manager role.

6. In the **Password** field, enter the password for the uploaded key pair.
7. Select **Import**.
8. Select **Continue** to move to the next section.
9. Next to **Attachments** in the **COT administration key pair and certificate** section select **+Add File** to upload a cryptographic key pair (p12 file extension) to use for customer signing.
10. In the **Password** field, enter the password for the uploaded key pair.
11. Select **Import**.
12. Select **Continue** to move to the next section.
13. On the **Export Configuration** file page, select **Export** to create and download a configuration file used to turn on Code Signing on your protected instance.
    The export process downloads an XML file to your local machine for use in the steps details in Configure Code Signing Enterprise on your protected instance.

    > **Note:** Code Signing enforces limits on large update sets to improve the user experience. The maximum size for an update set is 10,000 records.

Upload your Code Signing configuration file to your protected instance

Upload the configuration file generated on your trusted instance.

Before you begin

Roles required:

- admin
- security_admin
- codesigning_admin

> **Note:** The codesigning_admin role can be assigned using the process detailed in Assign the Code Signing Administrator Role.

- sn_kmf.cryptographic_manager

You also need the configuration file generated in the steps detailed in Configure Code Signing Enterprise on your trusted instance.

Procedure

1. Navigate to **All** > **System Update Sets** > **Retrieved Update Sets**.
2. Select **Import Update Set from XML** at the bottom of the Retrieved **Update Sets** list.
3. Select **Choose File**, and select your configuration file (xml file extension).
4. Select **Upload**.
   You should see the **Codesigning configuration** update set in the **Retrieved Update Sets** list.
5. Select and open the **Codesigning configuration** update set.
6. Select **Preview Update Set Batch**.
   If any conflicts occur during the preview, resolve them by selecting **Accept remote update** or **Resolve known CS conflicts**.
7. After resolving any conflicts, select **Commit Update Set Batch**.

Configure Code Signing Enterprise on your protected instance

Turn on and configure Code Signing on your protected instance.

Before you begin

Roles required:

- admin
- security_admin
- codesigning_admin

> **Note:** The codesigning_admin role can be assigned using the process detailed in Assign the Code Signing Administrator Role.

- sn_kmf.cryptographic_manager

Procedure

1. On your PPI, navigate to **All** > **Code Signing** > **Configuration** > **Guided Setup** to open the Code Signing configuration page.
2. In the **Instance type** field, select **protected instance**.

3. Select the **Next** button.
4. Select the **Please confirm the Code Signing configuration update set was imported and committed** check box.
5. Select the **Next** button.
6. Wait for the certificate installation to complete.
   A blue alert message displays indicating that items from the configuration file are installing. The alert disappears in a few seconds when the installation finishes.
7. Next to **Attachments** in the **Runtime/notarization key pair and certificate** section select **+Add File** to upload a cryptographic key pair (p12 file extension) to use for runtime/notarization purposes.
   You can re-use a key pair from the trusted instance or upload a new one specifically for this use.

   Tip: If the **+Add File** option is not available, verify that you are in the Global scope, and that you have the sn_kmf.cryptographic_manager role.

8. In the **Password** field, enter the password for the uploaded key pair.
9. Select **Import**.
10. Select **Continue** to move to the next section.
11. Wait while your instance performs a guardrail check.
    This check scans your instance for any invalid signatures, and might take some time to complete. Don't exit or refresh the page until the scan completes.

    This workflow identifies signatures that were generated with expired or inactive certificates and re-signs the affected records. To improve performance, the workflow now divides the process into multiple events and runs them in parallel.

12. If the scan finds invalid signatures, select **Download Report**.

    Selecting **Download Report** downloads a report (`scan_report.txt`) that you can use to investigate and fix the invalid signatures.

    After resolving errors, refresh the page to re-run the check.

13. If there are no errors, select **Complete setup**.
14. Wait for the configuration jobs to complete.
    Your instance runs one or more jobs to complete the configuration process. Don’t exit or refresh the page until the jobs have finished. When finished, you should see a **Code Signing configuration completed successfully** message.

Result

Signature validation is now being enforced on your production instance. You can verify successful completion by looking at your system properties. Look for the **com.snc.kmf.signature.validation.flag** property on the System Properties [sys_property] table, and verify that it has a value of **true**. Verify that the **com.snc.kmf.signature.validation.certificate_trust** property doesn’t have an empty value.

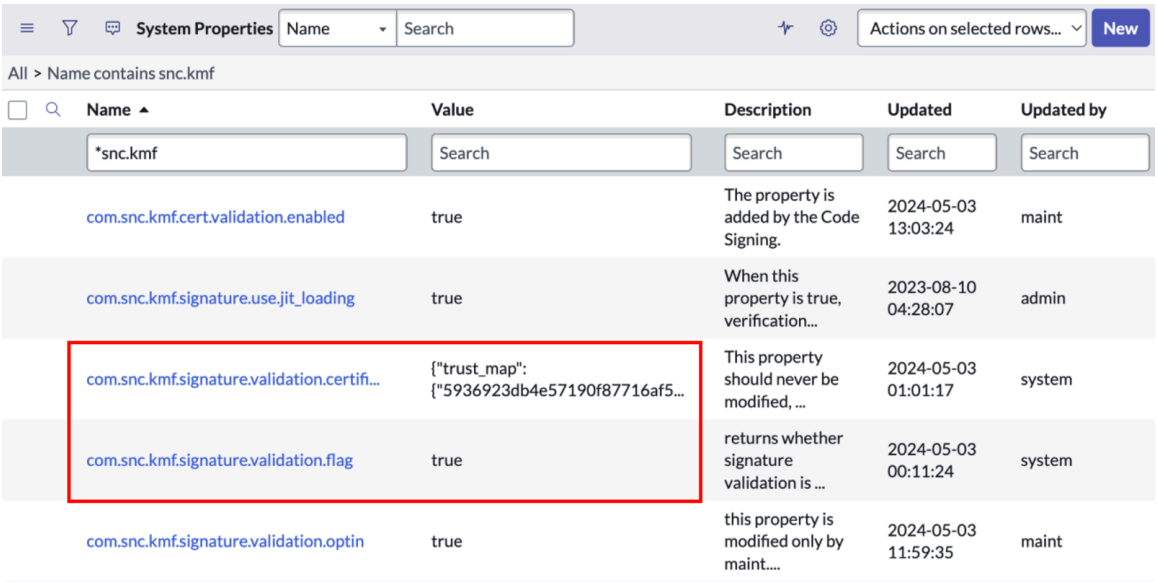

Turn on certificate validation

Protect your instance with certificate based validation.

Before you begin

Role required:

- codesigning_admin
- security_admin
- sn_kmf.cryptographic_manager

Procedure

1. On your trusted instance, navigate to **All** > **Code Signing** > **Configuration** > **Guided Setup** to open the Code Signing configuration page.
2. In the **Instance type** field, select **trusted instance**.
3. Select the **Next** button.
4. In the **Action** field under **Select the action you'd like to accomplish**, select **Turn on Cert Validation**.
5. Select the **Next** button.
6. Next to **Attachments** in the **Customer signing key pair and certificate** section select **+Add File** to upload a cryptographic key pair (p12 file extension) to use for customer signing.

   Tip: If the **+Add File** option is not available, verify that you are in the Global scope, and that you have the sn_kmf.cryptographic_manager role.

7. In the **Password** field, enter the password for the uploaded key pair.
8. Select **Import**.
9. Select **Continue** to move to the next section.
10. Next to **Attachments** in the **COT administration key pair and certificate** section select **+Add File** to upload a cryptographic key pair (p12 file extension) to use for customer signing.
11. In the **Password** field, enter the password for the uploaded key pair.
12. Select **Import**.

13. Select **Continue** to move to the next section.
14. In the **Perform trusted instance tasks** section, wait for all tasks to be completed.

    Your instance generates and executes these tasks automatically. If you used Code Signing prior to the Vancouver release, tasks are created and executed to update your signatures.

    In some cases, no tasks are needed. **No tasks needed** displays on this page.
15. Select **Continue** to move to the next section.
16. On the **Export Configuration** file page, select **Export** to create and download a configuration file used to turn on Code Signing on your protected instance.
    The export process downloads an XML file to your local machine for use in the steps detailed in Configure Code Signing Enterprise on your protected instance.
17. On your protected instance, navigate to **All** > **System Update Sets** > **Retrieved Update Sets**.
18. Select **Import Update Set from XML** at the bottom of the Retrieved **Update Sets** list.
19. Select **Choose File**, and select your configuration file (xml file extension).
20. Select **Upload**.
21. Return to the code signing configuration page at **All** > **Code Signing** > **Code Signing Configuration**.
22. Use the wizard to complete your configuration, selecting the options for completing certificate validation activation.

Quorum Controlled Certificate Revocation

The quorum-controlled certificate revocation for Code Signing certificates provides a secure mechanism for a Code Signing admin to revoke Code Signing certificates. The revocation process involves submitting a request that requires approval from multiple stakeholders. This workflow helps to prevent accidental or unauthorized revocations.

This topic provides the high-level process overview for securely revoking certificates between a trusted instance and a protected instance using a quorum-based approval workflow.

- Request creation and export (Trusted instance): Initiate a quorum certificate revocation by creating a request with the required configuration properties. Export the update set, which includes the revocation request and related data. See Export Revocation Request Configuration
- Import and approval (Protected instance): Import the update set into the protected instance. Approvers receive notifications and must complete the approval workflow before the certificate is deleted. See Import the revocation request configuration
- Optional MID Server restart: If enabled in the configuration, MID Servers on the protected instance are restarted after the certificate revocation. This forces immediate certificate resynchronization but can result in data loss for unprocessed or cached events. See Approve certificate revocation

Note:

Your *update set* expires after the time window defined by the `com.snc.kmf.signature.validity_window` property. If it expires, you can export a new signed update set from the trusted instance. This validity window applies to all update set operations, including export, import, and enabling Code Signing. This validity window is not related to the request time window that you specify when creating the request.

Export Revocation Request Configuration

Start the certificate revocation process by selecting the certificate that you want to revoke. Provide the required configuration properties. Export this transaction as part of an update set, which is imported into the protected instance for approval and execution.

Before you begin

Role required: sn_cse.codesigning_admin, sn_cse.quorum_requester, security_admin

Procedure

1. On your trusted instance, navigate to **All** > **Code Signing** > **Configuration** > **Guided Setup** and open the Code Signing configuration page.
2. In the **Instance type** field, select **trusted instance**.
3. Select the **Next** button.
4. In the **Action** field under **Select the action you'd like to accomplish**, select **Certificate Revocation**.
   The Certification Revocation Request page appears.
5. In the Select Certificate to Revoke page, select the certificate that you want to revoke and select **Initiate Revocation**.
6. Configure the approval requirements for certification revocation.

   a. In the Quorum Requirements menu, enter the appropriate information in the text fields.

   Quorum Requirements Properties

Property	Description
Minimum approvals	Minimum number of approvers required to approve the certificate revocation request.
Time window	Expiration time for the revocation request.
Approvers	Email address of the users who are authorized to approve the revocation request.

   b. In the **Request description** field, enter the reason for initiating the certificate revocation.

   c. Select **Save**.
   The Export configuration file menu appears.

   d. On the Export Configuration file page, select **Export** to create and download a configuration file used to run certificate revocation workflow on your protected instance.

Result

The export process downloads an XML file to your local machine for use in the steps detailed in Configure Code Signing Enterprise on your protected instance.

Import the revocation request configuration

Import the update set into the protected instance to initiate the certificate revocation process. Approvers receive email notifications and they should complete the approval workflow before

the certificate is revoked. The approval means that the revocations are confirmed, authorized, and traceable for security and compliance purposes.

Before you begin

Role required: sn_cse.codesigning_admin, sn_cse.quorum_requester, security_admin

Procedure

1. Log in to your protected instance and navigate to **All** > **System Update Sets** > **Retrieved Update Sets** > **Import XML**.
   The Import XML page appears.
2. Select **Choose file** and select the configuration XML file from your local system.
3. Select **Upload**.
4. Return to the code signing configuration page at **All** > **Code Signing** > **Code Signing Configuration**.

   a. Review the following configuration files in the **Customer Updates in Batch** tab.
      - Code Signing Configuration Property (`time_window`)
      - Code Signing Configuration Property (`approver_email_ids`)
      - Code Signing Configuration Property (`minimum_approvals`)
      - Code Signing Configuration Property(`restart_mid_servers`)
      - Code Signing Quorum Request (CSEQC`xxxxxxx`)
      - Code Signing Transaction (CSETRANS`xxxxxxx`)
      - KMF Signature Records

   b. Select **Batch Update Set Preview** > **Commit Update Set Batch**.

5. Return to the code signing configuration page at **All** > **Code Signing** > **Configuration** > **Guided Setup**.
6. In the **Instance type** field, select **Protected instance** and select **Next**.
   The **Start or continue your configuration** page appears.
7. Select the **Confirm that the Code Signing configuration update set was imported and committed.** check box and select **Next**.
   In the Quorum Control Configuration page, review the information that you entered and select **Trigger Quorum Approval**.
8. In the Protected instance, select **Certification Revocation Status** in the Trigger Quorum Approval page to review updates on the quorum control request approvals.
9. Select **Approver details** and review the status of the quorum requests.

Approve certificate revocation

Review and approve certificate revocation requests from the email approval notifications that are sent to your registered email address. Review the approval notification and select the **Click here to approve** or **Click here to reject** link to access the protected instance and act. You can also access the approval request and the code-signing quorum request directly from the email.

Before you begin

Role required: sn_cse.codesigning_admin, approver_user

Procedure

1. Access the protected instance by following the links provided in the notification email.
2. Review the transaction requests under the Approvals menu.
3. Approve the task by selecting **Approve**.
4. In the approval confirmation page, select **Yes** to access the specific code signing quorum request.
5. Review the Code Signing quorum request and add any additional notes in the **Comments** menu, if necessary.
6. Select **Update**.

Result

The Code Signing quorum request is approved. If you have enabled the MID Server restart while exporting the revocation request, the MID Server restarts for the request to be activated.

Turn off Code Signing

Disable code signing on your protected instance.

Before you begin

Role required: admin, codesigning_admin

Procedure

1. On your trusted instance, navigate to **All** > **Code Signing** > **Configuration** > **Guided Setup** to open the Code Signing configuration page.
2. In the **Instance type** field, select **trusted instance**.
3. Select the **Next** button.
4. In the **Action** field under **Select the action you'd like to accomplish**, select **Turn off Code Signing**.
5. In the **Export configuration file** panel, select **Export** to download the update set.
   The export process downloads an XML file to your local machine for use in the next steps.
6. Select **Done**.
7. Log in to your protected instance.
8. Upload your configuration files using the steps in Upload your Code Signing configuration file to your protected instance.
9. Navigate to **All** > **Code Signing** > **Configuration** > **Guided Setup** to open the Code Signing configuration page.
10. In the **Instance type** field, select **protected instance**.
11. Select the **Next** button.
12. Select the **Please confirm the Code Signing configuration update set was imported and committed** check box.
13. Select the **Next** button.
14. Wait for the certificate installation to complete.
    After completion, a **Code Signing configuration completed successfully** message displays.

Result
You can verify successful completion by looking at your system properties. Look for the **com.snc.kmf.signature.validation.flag** property on the System Properties [sys_property] table, and verify that it has a value of **false**.

Create Code Signing key pairs and certificates

Create two key pairs to signed certificates to establish trust between your protected and trusted instances.

To establish trust between your instances, you must create a key pair and certificate for each of the `cm_code_attest` and `cm_code_signing` cryptographic modules.

Creating key pairs and certificates is done using a cryptographic tool installed on your local device, such as the OpenSSL tool. For more information on this tool, see https://www.openssl.org. If your organization uses other cryptographic tools, such as LibreSSL or GnuTLS, refer to the documentation for those products for similar steps.

Key pair specifications

The key pairs you create must meet these requirements.

Type	RSA
Key length	4096
Signing algorithm	RSASSA_PKCS1_V1_5_SHA_512

Certificate specifications

Certificates must be signed by a public certificate authority.

Specify custom rules in ECC firewall

Configure the External Communication Channel (ECC) firewall in your MID Server by specifying the custom rules to selectively allow or reject the incoming message and override the Code Signing configuration.

Security administrators can use the ECC firewall tags to override the Code Signing configuration and allow or reject specific operations on MID Server. These custom rules must be specified in the YAML file of the located at: `agent/boot-config.yaml`.

These tags are specific to a protocol. The configuration specified for the parent tag is applicable to the child tag. For example, if `Http` protocol is allowed, `rest` and `soap` protocols are also allowed. This table outlines the available parents and child tags.

Parent tag	Child tag
DNS	
HTTP	• REST • SOAP
DIRECTORY_SERVICES	LDAP
SNMP	

Parent tag	Child tag
SSH	• SCP • SFTP
SYSLOG	
WINDOWS	• CIM • POWERSHELL • WMI • WINRM
JAVASCRIPT	
GROOVY	
VCS	GIT
DATABASES	JDBC
DATA_SOURCES	
INTEGRATION_HUB	
ITOM	• CLOUD_PROVISIONING_GOVERNANCE • DISCOVERY • EVENT_MANAGEMENT • HEALTH_LOG_ANALYTICS • SERVICE_MAPPING
ORCHESTRATION	

To configure the custom rules:

1. In the MID Server, identify the file `boot-config-sample.yaml`.
2. Rename the YAML file to `boot-config.yaml` and move the file to the location: `agent/boot-config.yaml`.
3. In the YAML file, specify the custom rules and save the changes. An example of the YAML file:

```
security:
  eccFirewall:
    mode: enforcing
    rules:
      - tags: [rest]
        action: accept
      - tags: [soap]
        action: accept
      - tags: [jdbc]
        action: reject
```

4. Restart the MID Server.

Change your Root of Trust configuration

Trust and use your own certificates instead of relying on ServiceNow build certificates (default) by changing to use your Root of Trust (ROT). ServiceNow components like script includes, business rules, etc., are signed at build time using a ServiceNow build time key (verification certificate is the ServiceNow build certificate).

Changing the root of trust

To change the Root of Trust for these records signatures, you must follow the change of Root of Trust process.

- Generate and migrate a new set of signatures for all provided components, using your provided certificate.
- Disable the Root of Trust property using a scheduled job.

Learn more about these steps in Migrate signatures to use a customer certificate and Disable ServiceNow Root of Trust.

Impact on the signature generation and verification process

By default, Code Signing build certificates are trusted during the signature verification process. After making this change, your instance accepts signatures only from your own Code Signing certificate.

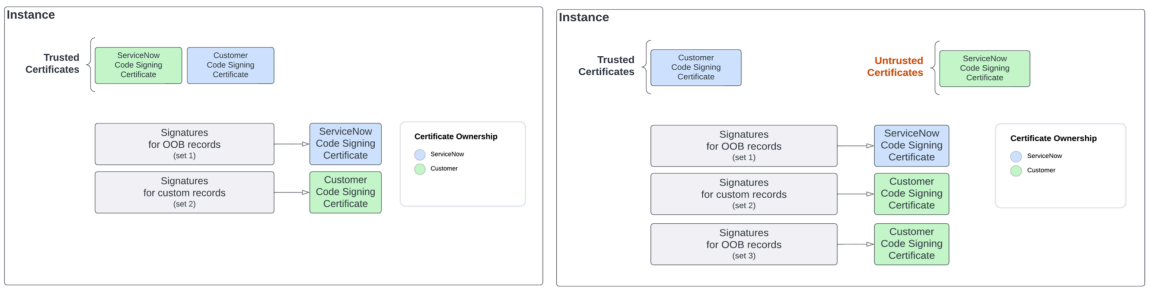

ROT property set to false (default)	ROT property set to true
• When verifying, signatures with build certificates are trusted. • When signing, if you don't provide keys, the instance signing key is used as a backup key. • The signature REST end point *api/sn_kmf/signature/certificates* returns ServiceNow Code Signing build certificates along with other certificates present on the instance.	• When verifying, signatures with build certificates aren't trusted. • When signing, if you don't provide keys, signing isn't performed. • The signature REST end point *api/sn_kmf/signature/certificates* excludes ServiceNow build certificates (San Diego, Vancouver PKI, W PKI).

Impact on your MID Server

When the ROT property is set to false

If you choose to leave your ROT property at its default value (false), there's no impact on your MID Server.

When Code Signing is enabled and the ROT property is set to true

- The `isTrusted()` API returns `false` for signatures with a build certificate.
- The `isTrusted()` API returns `true` for signatures with your certificate.
- The REST API call for certificates excludes build certificates.
- You may see MID Server issues, such as `signature validation failed` messages in the logs.

Migrate signatures to use a customer certificate

Run a signing job to migrate your signatures to a customer Root of Trust (ROT).

Before you begin

Role required: admin, security_admin, and sn_kmf.cryptographic_manager

Code Signing must be enabled on your protected and trusted instances. You can verify by checking that the **com.snc.kmf.signature.validation.flag** system property is set to `true`.

This procedure is part of a series of procedures to change a customer Root of Trust (ROT) on your instances. For an overview of this process, see Change your Root of Trust configuration.

Procedure

1. Log in to your protected instance.
2. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
3. Find and open the **ROT - Generate Updateset of records to migrate signatures using customer certificate** scheduled job.
4. At the bottom of the form, select **Execute Now**.
5. Navigate to **All** > **System Security** > **Security Jobs** > **Create New**.
6. At the prompt, select **Signing Job**.
7. In the **Signing Job** form, fill in the fields as needed.

Field	Value
Name	Create a unique name for your job.
Type	Select **Sign Update Set**.
Table	Select the update set created in the previous steps. The update set has a name beginning with `change_root_of_trust_updateSet`.

8. Right-click the form header and select **Save** to save this record.
9. Right-click the form header and select **Export** > **XML (This Record)** to export this record as an XML file.
10. Log in to your trusted instance.
11. Navigate to **All** > **System Security** > **Security Jobs** > **All**.
12. Right-click the list header and select **Import XML**.
13. In the **Import XML** form, select **Choose File** and select the XML file you downloaded in the previous steps.
14. Select **Upload**.
15. From the list, open the imported security job.

16. Select **Export Code Signing job to production**.
    This action signs the job and places it in a new update set you can import into your protected instance.

    Important: After signing the job, you must perform the next steps within 10 minutes. If you exceed this time period, you can re-sign the job using these steps, which creates another signed update set.

17. Navigate to **All** > **System Update Sets** > **Local Update Sets**.
18. Find the update set created in the previous steps.
    The name starts with `SIGN_UPDATE_SET_updateSet`.
19. Select **Export XML** to export your update set as an XML file.
20. Log in to your protected instance.
21. Navigate to **All** > **System Update Sets** > **Retrieved Update Sets**.
22. At the bottom of the list, select **Import Update Set from XML**.
23. In the **Import XML** form, select **Choose File** and select the XML file you downloaded in the previous steps.
24. Select **Upload**.
25. Navigate to **All** > **System Update Sets** > **Retrieved Update Sets**, and open the update set starting with `SIGN_UPDATE_SET_updateSet`.
26. Select **Preview Update Set**.
27. After the preview is completed, select **Commit Update Set**.
28. Navigate to **All** > **System Security** > **Security Jobs** > **All**.
29. Open the imported security job.
30. Select **Start** to run the security job.

    After the security job is completed, information regarding the status of the job appears in the **Summary** field.

    When the job is in Done state, all signatures of update set records must use the customer provided certificate as verification certificate. You can verify this on the KMF Signature Records [sn_kmf_record_signature] table.

What to do next

To continue the Root of Trust configuration process, see Disable ServiceNow Root of Trust.

Disable ServiceNow Root of Trust

Run a scheduled job on your trusted instance to disable Root of Trust.

Before you begin

Role required: admin, security_admin, and sn_kmf.cryptographic_manager

Code Signing must be enabled on your protected and trusted instances. You can verify by checking that the **com.snc.kmf.signature.validation.flag** system property is set to `true`.

This procedure is part of a series of procedures to change to a customer Root of Trust (ROT) on your instances. For an overview of this process, see Change your Root of Trust configuration.

Procedure

1. Log into your trusted instance.
2. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
3. Open the **Disable ServiceNow Root of Trust** scheduled job.
4. Select the **Export signed job to production**.
5. Navigate to **All** > **System Update Sets** > **Local Update Sets**.
6. Find and open the **Disable ServiceNow Root of Trust** update set.
7. Select **Export XML** to export your update set as an XML file.
8. Log in to your protected instance.
9. Navigate to **All** > **System Update Sets** > **Retrieved Update Sets**.
10. At the bottom of the list, select **Import Update Set from XML**.
11. In the **Import XML** form, select **Choose File** and select the XML file you downloaded in the previous steps.
12. Select **Upload**.
13. Navigate to **All** > **System Update Sets** > **Retrieved Update Sets**, and open the **Disable ServiceNow Root of Trust** update set.
14. Select **Preview Update Set**.
15. After the preview is completed, select **Commit Update Set**.
16. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
17. Open the scheduled job imported in the update set.
18. Select the **Execute Now** button to run the job.

Result

Executing the scheduled job sets the ROT property to true. Your instance is configured to use the customer root of trust.

Using Code Signing

Learn how to sign records, messages, and attachments to help verify the authenticity and integrity of your data.

Sign the JDBC data source records in the protected instance

Use update sets to sign and validate the JDBC data sources by enabling the code signing in protected and trusted instances.

Sign the REST and SOAP messages in the production instance

Use update sets to sign and validate the REST and SOAP messages by enabling the code signing in protected and trusted instances.

Sign the flows, subflows, and actions in the protected instance

Use update sets to sign and validate the flows, subflows, and actions by enabling the Code Signing in protected and trusted instances.

Sign specific records or attachments

Create a security job to sign specific records or attachments rather than all records or attachments on a table.

Standalone signing tool

Use the standalone Signing Tool to sign supported records in ServiceNow applications using your own private key.

Standalone signing tool

Use the standalone Signing Tool to sign supported records in ServiceNow applications using your own private key.

Use the Signing Tool to sign records in ServiceNow apps. The tool generates signatures for records in your local environment using your own private key.

Signing tool workflow

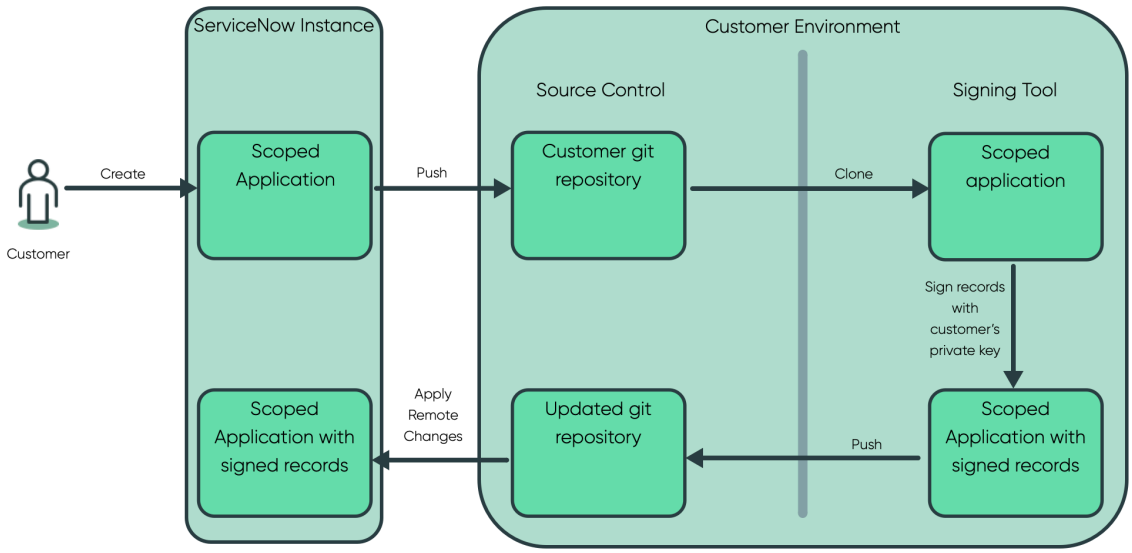

1. Create or select an existing ServiceNow application with records to be signed, such as business rules or script includes.
2. Push the application into your Git repository, which resides in your environment.

   Note: Applications can be synced between a Git repository and your instance using the Source Control Integration. For details on configuring and using this integration, see .

3. Clone the application in your local environment.
4. Use the Signing Tool (also in your local environment) to sign the supported records from the cloned ServiceNow application using your private key. The Signing Tool creates signature records and X.509 Certificate [sys_certificate] records. For details on using the Signing Tool, see Using the Signing Tool.
5. Push the updated application to your Git repository.
6. In your instance, import the updated application by applying remote changes.

Using the Signing Tool

Learn how to use the Signing Tool to sign supported records in ServiceNow applications.

Before you begin
Role required: admin

To perform these steps, you must have the following:

- A ServiceNow application that has records to be signed.
- A private key for signing records.
- The *signRecords.sh* script in your local environment, with execute permission.

Important: The *signRecords.sh* script is included as part of the signing-tool jar file, which you must request from ServiceNow customer support.

Procedure

1. In your local environment, navigate to the directory containing the *signRecords.sh* script.
2. Use the following command format to sign your records:

```
./signRecords.sh -d [Path to the root directory of the
 ServiceNow Application to Sign] -f [Path to the Keystore file]
```

   For example:

```
./signRecords.sh -d /users/abc/ServiceNow-App-1
 -f /users/abc/codesigning.p12
```

3. If prompted, enter the password for the keystore.
   Press Enter if there's no password.
4. Review the output to confirm a successful signing.

```
Sep 26, 2022 2:41:09 PM com.snc.java.commands.ACommand start
INFO: CODESIGN: executing codesigning...
Sep 26, 2022 2:41:09 PM
 com.snc.core.codesigning.CodeSignerSupplier get
INFO: CODESIGN: signing record for documentId:
 65e81132770211057416efe7c5a994f
Sep 26, 2022 2:41:11 PM com.snc.java.commands.ACommand start
INFO: CODESIGN: codesigning successfully completed!
```

   In the preceding output example, the Signing Tool used the provided keystore file to sign the record. You can also see that:

   - The script found a supported record `65e81132770211057416efe7c5a994f` and signed it.
   - In the ServiceNow-App-1 directory, two records created: `sys_certificate.xml` and `sn_kmf_record_signature.xml`.
5. Import the updated application back into your instance by applying remote changes in Studio. For more information, see .

Signing Tool arguments

Learn about the available arguments for the Signing Tool.

Command line arguments

Argument	Required	Description
-d	Yes	Root directory of project to sign. Should contain project directory (random 32 alphanumeric), sn_source_control.properties file, and a <project_name>.iml file
-f	Yes	The file path of the keystore.

Command line arguments (continued)

Argument	Required	Description
-a	No	Alias used to access a specific entry within the keystore.
-c	No	Concatenate record signatures into one file.
-k	No	Password to access the key stored within the keystore. You can also enter this password when prompted instead of within the argument.
-o	No	Sign with a new certificate in place of any existing sys_cert files.
-p	No	Password to access the keystore if it has a password. You can also enter this password when prompted instead of within the argument.
-w	No	Wipe all existing signature record files.
-h	No	Show this help message and exit.

Sign the JDBC data source records in the protected instance

Use update sets to sign and validate the JDBC data sources by enabling the code signing in protected and trusted instances.

- Establish Circle of Trust between the protected and trusted instances.
- Role required: security_admin

Note:

- MID Server doesn't handle the file data sources and hence, these data sources are not code signed.
- LDAP data sources can't be code signed.

Sign existing data sources of the JDBC type

Use update sets to bring mass signing jobs to the protected instance.

Before you begin

Role required: security_admin

Procedure

1. In the trusted instance, configure the KMF signing job to sign the data sources.

   a. Navigate to **System Security** > **Security Jobs** > **All**.

   b. Click **New**.

   c. On the form, fill these values.

Field	Description
Name	Name to identify the record.
Type	Type of the encryption job. Select **Mass Sign Records**.

Field	Description
Table	Table from which the records should be signed. Select **Data Source**.

d. Click **Export Code Signing job to production**.
A confirmation message is displayed that the update set is signed.

e. Export the generated update set to an XML file.

2. In the protected instance, import and commit the update set to retrieve the mass signed jobs from trusted instance.

a. Navigate to **System Security** > **Security Jobs** > **All**.

b. Open the update set exported from the trusted instance.

c. Click **Start**.
A confirmation message is displayed that the records are signed.

Sign new data sources of the JDBC type

Use update sets to bring the signed update set to the protected instance.

Before you begin

Role required: sn_kmf.cryptographic_manager

Procedure

1. In the trusted instance, start an update set.

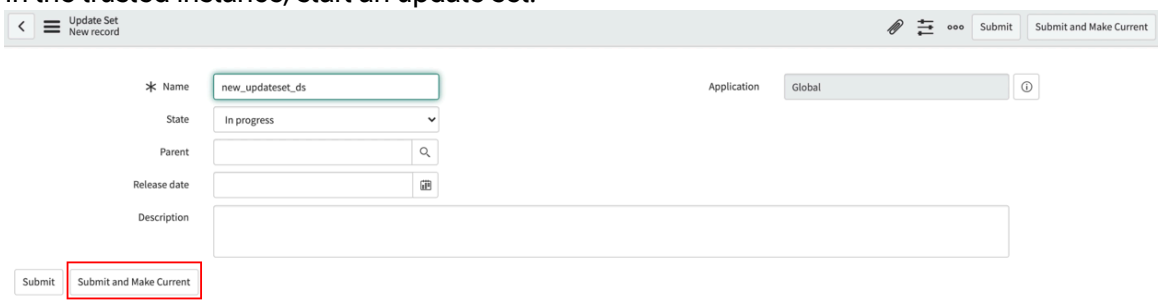

2. In the trusted instance, create the required data sources.

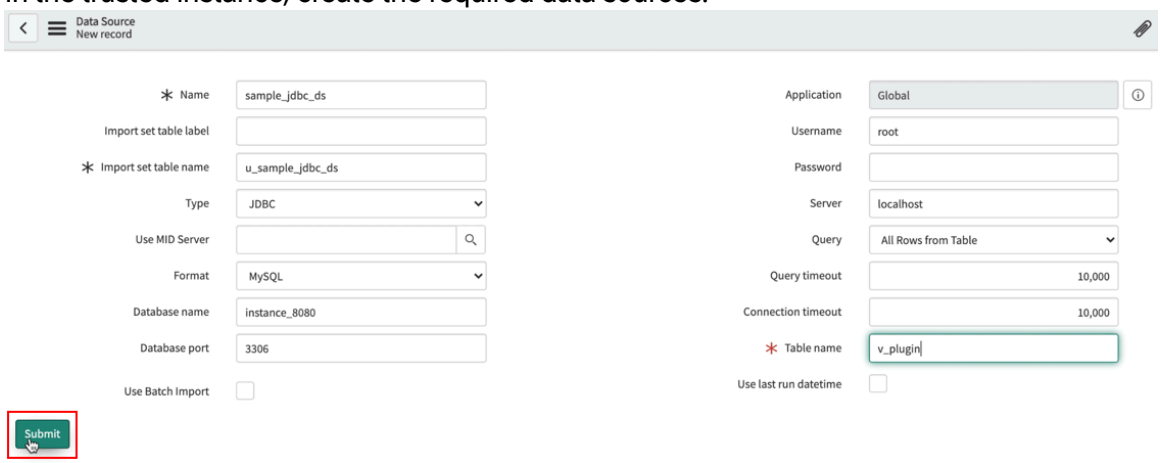

The data sources are added to the update set.

3. In the trusted instance, change the state of the update set to **Complete** and click **Update**.

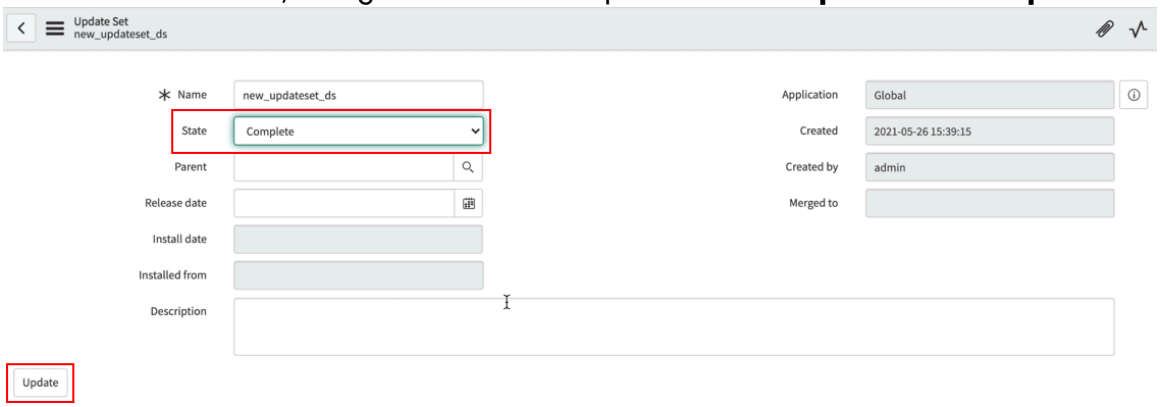

4. In the trusted instance, sign the update set by creating an encryption job.

   a. Navigate to **System Security** > **Security Jobs** > **All**.

   b. Click **New**.

   c. On the form, fill these values.

Field	Description
Name	Name to identify the record.
Type	Type of the encryption job. Select **Sign Update Set**.
Table	Update set from which the records should be signed.

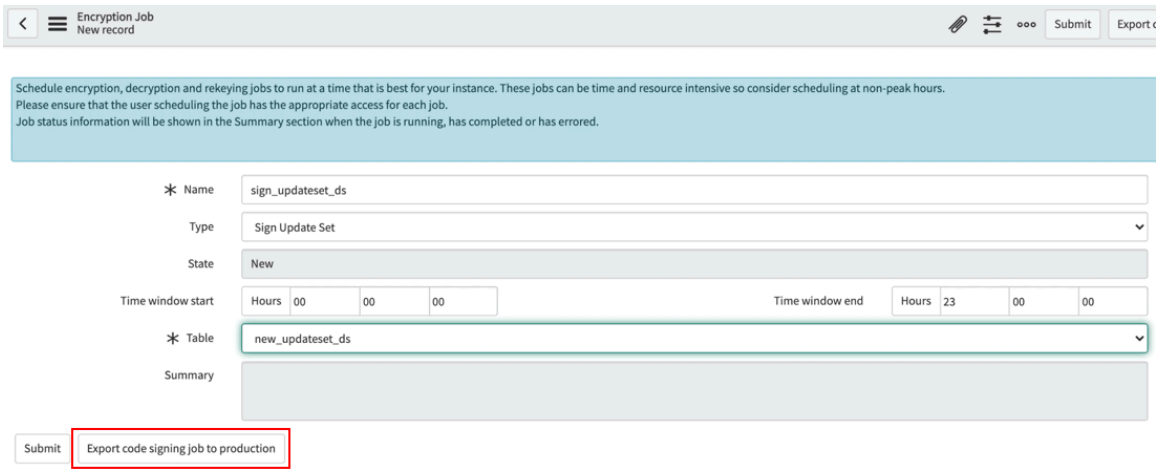

   d. Click **Submit**.

   e. Click **Start** to sign the update set.

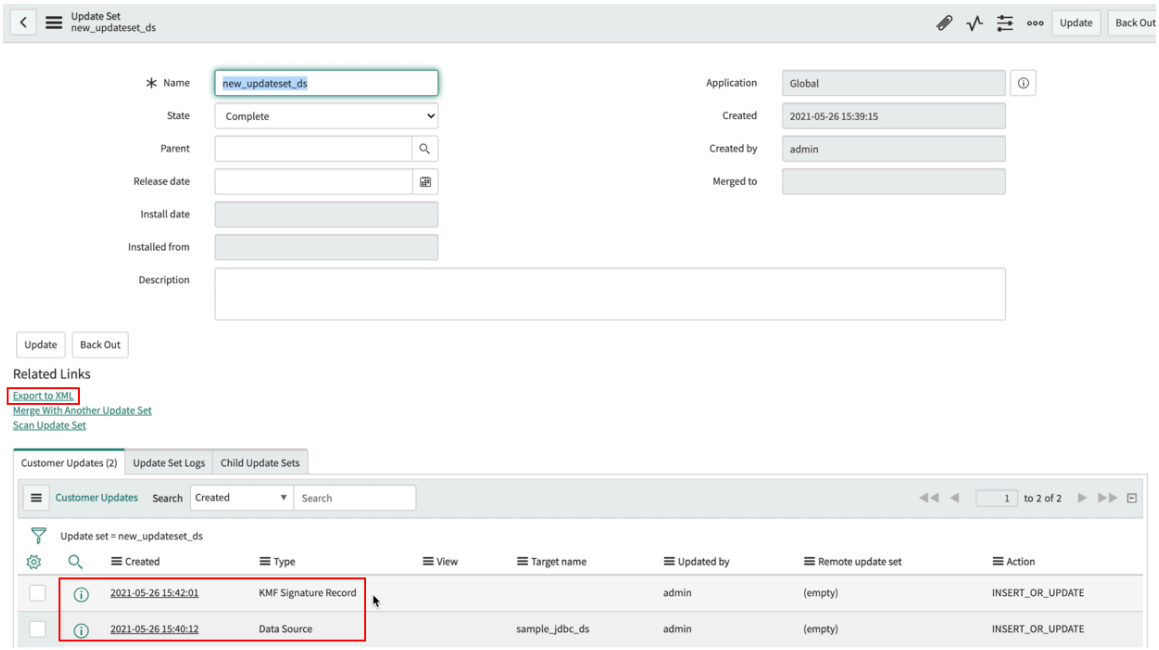

- **Summary** is updated that the records are signed.
- The update set is updated and includes the signature.

5. In the trusted instance, open the signed update set record and export it to an XML.
6. In the protected instance, import the signed update set.

   a. Navigate to **System Update Sets** > **Retrieved Update Sets**.

   b. Click the **Import Update Set from XML** related link to import the update set that is exported from the trusted instance.
   For more information, see Import and commit the quick-start update set.
   The update set is committed successfully.

Sign the REST and SOAP messages in the production instance

Use update sets to sign and validate the REST and SOAP messages by enabling the code signing in protected and trusted instances.

Before you begin

- Establish Circle of Trust between the protected and trusted instances.
- Role required: security_admin

Sign the existing REST and SOAP messages

Sign and validate the existing REST and SOAP messages by enabling the Code Signing in protected and trusted instances.

Before you begin

Role required: sn_kmf.cryptographic_manager

Procedure

1. In the trusted instance, configure the KMF signing job to sign the UI actions.

   a. Navigate to KMF Signature Configuration.

   b. On the form, fill these values.

   KMF Signature Configuration form

Field	Description
Table Name	Glide table name. For example, select **UI Actions [sys_ui_action]**.
KMF Signature Purpose	Purpose of signing the records. Select **ECC Queue**.
Signature Generation Fields	Fields in the data source that you want to sign. If any changes are made to the values in one or more of these fields, the previously generated signature becomes invalid. Select **Name** and **Script**.
Signature Generation Filter	Filter criteria that must be met to sign the records.
Sign Attachment	Option to sign the attachment in the glide record.
Instance Key	Option to use the instance key.

   c. Right-click the form header and click **Save**.

2. In the trusted instance, sign the required records.

   a. Navigate to **System Security** > **Security Jobs** > **All**.

   b. Click **New**.

   c. On the form, fill these values.

Field	Description
Name	Name to identify the record.
Type	Type of the encryption job. Select **Mass Sign Records**.
Table	Table from which the records should be signed. Select **UI Action**.

   d. Click **Export Code Signing job to production**.
   Two locally signed update sets are created.
   - One update set for the UI action configuration.
   - Another update set from the encryption job to export the code signing job.

3. In the trusted instance, export the local update set to an XML file.
   a. Navigate to **System Update Sets** > **Local Update Sets**.
   b. Open the update set you had created for mass signing the records.
   c. Click the **Export to XML** related link and save the XML file.
4. In the protected instance, import the update sets.
   a. Navigate to **System Update Sets** > **Retrieved Update Sets**.
   b. Click the **Import Update Set from XML** related link to import the update set that is exported from the trusted instance.
   For more information, see Import and commit the quick-start update set.
   The update set is committed successfully.
5. In the protected instance, run the encryption job you had earlier created in the trusted instance by selecting **Start**.
   A confirmation message is displayed mentioning that the records are signed.

Sign new REST and SOAP messages

Sign and validate the new REST and SOAP messages from the trusted instance by enabling the Code Signing in protected and trusted instances.

Before you begin

Role required: security_admin

Procedure

1. In the trusted instance, start an update set.
2. In the trusted instance, create the required REST or SOAP messages.
   The messages are added to the update set.
3. In the trusted instance, change the state of the update set to **Complete** and click **Update**.
4. In the trusted instance, sign the update set by creating an encryption job.
   a. Navigate to **System Security** > **Security Jobs** > **All**.
   b. Click **New**.
   c. On the form, fill these values.

Field	Description
Name	Name to identify the record.
Type	Type of the encryption job. Select **Sign Update Set**.
Table	Update set from which the records should be signed. Select **Sign new Rest V2 update set - 1**.

   d. Click **Submit**.

e. Click **Start** to sign the update set.
   - **Summary** is updated that the records are signed.
   - The update set is updated and includes the signature.

5. In the trusted instance, open the signed update set record and export it to an XML.
6. In the protected instance, import the update set.

   a. Navigate to **System Update Sets** > **Retrieved Update Sets**.

   b. Select the **Import Update Set from XML** related link to import the update set that is exported from the trusted instance.
   For more information, see Import and commit the quick-start update set.
   The update set is committed successfully.

Sign specific records or attachments

Create a security job to sign specific records or attachments rather than all records or attachments on a table.

Beginning in the Vancouver release, security administrators can use security jobs to sign specific records on a table rather than all records on a table. Sign specific records to avoid accidentally signing any unreviewed code.

These jobs include an attached log file, which is generated at the end of the signing job. This log file contains information about signed records and the signature configuration used.

Signing process

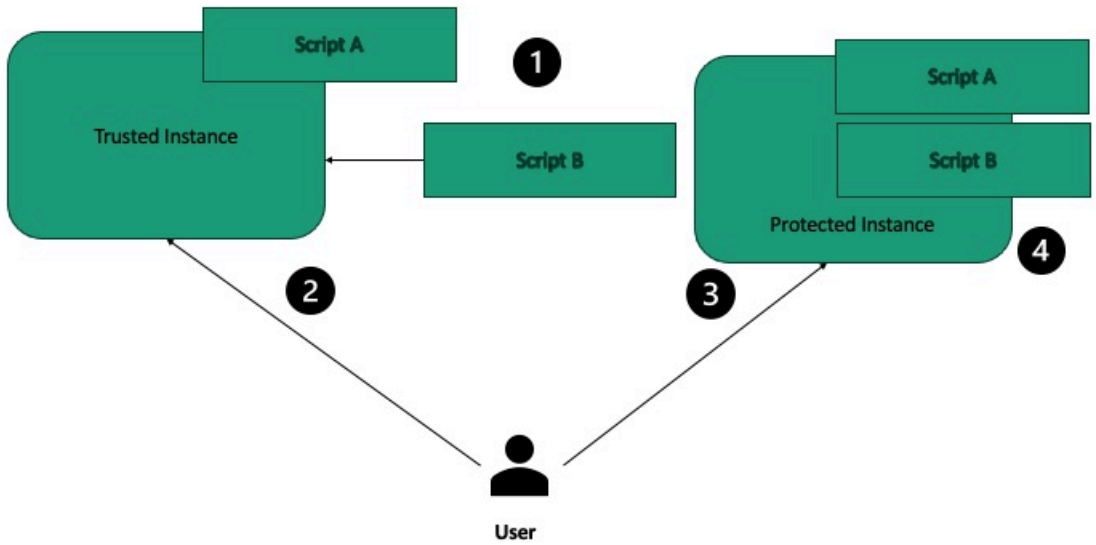

This diagram shows an example of how you can use the signing process. In this example, a script, labeled `Script B` exists only on the protected instance, and must be imported to a trusted instance to be signed. `Script A` exists on both instances already, doesn't need to be signed. It has either already been reviewed and signed, or hasn't been reviewed yet and shouldn't be signed.

1. Import the record to your trusted environment.
2. In the trusted instance, create a signing job to sign the record. This process is detailed in Create a job to sign specific records or attachments on a trusted instance.
3. Import the signed signing job to your protected instance using an update set.
4. In the protected instance, run the imported signing job.

Create a job to sign specific records or attachments on a trusted instance

Sign a specific record, or group of records you define on a trusted instance.

Before you begin

Role required: security_admin or sn_kmf.cryptographic_manager

Procedure

1. Navigate to **All** > **System Security** > **Security Jobs** > **Create New**.
2. At the **What type of Security Job would you like to create?** prompt, select **Signing Job**. A new **Signing Job** record displays.
3. In the form, fill in the fields as needed:

Signing job fields

Field	Description
Name	Descriptive name for this job.
Type	Type of security job. To sign specific records, select the **Sign Specific Records** option. To sign specific attachments, select **Sign Specific Attachments**.
State	State of this job. This field begins with the value of **New**. This field is read only.
Table	Table containing the records or attachments you want to sign. If signing attachments, select the table with records that the attachments are associated with, not the Attachment [sys_attachment] table. **Tip:** Check the Key Management Framework (KMF) Signature Record [sn_kmf_record_signature] table to ensure that there aren't already signatures for the table you've selected.
Filter Records	Filter conditions used to limit which records appear in the **Select records for signing** table.
Select records for signing	List of records on the table selected in the **Table** field, limited by filter created in the **Filter Records** field. Move records from the **Available** window to the **Selected** window to include them in the signing job.
Time window start	The start of the time window to run this job. The job will run after the time entered in this field.

Field	Description
	A valid time value is in Coordinated Universal Time based on a 24-hour time notation.
Time window end	The end of the time window to run this job. The job runs before the time entered in this field. If the job hasn't yet completed, the job will pause and resume at the next time window start. The end time must be after the start time. A valid time value is in Coordinated Universal Time based on a 24-hour time notation.
Summary	Summary of the execution of this job. This field is read only.

4. Right-click the form header, and select **Save**.
5. Open the job that you saved, and click **Start**.
   This action signs the signing job, and it's ready for export.
6. Navigate to **All** > **System Update Sets** > **Local Update Sets**
7. Find and open the update set for your signing job.
   In the customer updates tab, you can see that this update set includes the signing job and the signature record.
8. Select **Export to XML**.
   This action creates an XML file containing your update set on your local device.
9. On your protected instance, navigate to **All** > **System Update Sets** > **Retrieved Update Sets** > **.**
10. At the bottom of the page, select **Import Update Set from XML**.
11. Select the **Choose File** button, and select the XML file created in previous steps.
12. Select **Upload**.
    Your update set is loaded, and appears in the **Retrieved Update Sets** list.
13. In the **Retrieved Update Sets** list, open the record for your imported update set.
14. Select the **Preview Update Set** button.
    After the preview completes successfully, you'll see the **Commit Update Set** button appear.
15. Select **Commit Update Set**.

Sign the flows, subflows, and actions in the protected instance

Use update sets to sign and validate the flows, subflows, and actions by enabling the Code Signing in protected and trusted instances.

- Establish Circle of Trust between the protected and trusted instances.
- Role required: security_admin

Sign the existing flow, subflows, and actions

Sign and validate the existing flow, subflows, and actions by enabling the code signing in production and trusted instances.

Before you begin

Role required: sn_kmf.cryptographic_manager

Procedure

1. In the trusted instance, sign the records in the Step Instance table.

   a. Navigate to **System Security** > **Security Jobs** > **All**.

   b. Click **New**.

   c. On the form, fill these values.

Field	Description
Name	Name to identify the record.
Type	Type of the encryption job. Select **Mass Sign Records**.
Table	Table from which the records should be signed. Select **Step Instance**.

   d. Click **Export Code Signing job to production**.
   Two locally signed update sets are created.
   - One update set for the KMF signature.
   - Another update set from the encryption job to export the code signing job.

2. In the trusted instance, export the local update set to an XML file.
   a. Navigate to **System Update Sets** > **Local Update Sets**.
   b. Open the update set you had created for mass signing the records.
   c. Click the **Export to XML** related link and save the XML file.
3. In the protected instance, import the XML file.

   a. Navigate to **System Update Sets** > **Retrieved Update Sets**.

   b. Click the **Import Update Set from XML** related link to import the update set that is exported from the trusted instance.
   For more information, see .
   The update set is committed successfully.

4. In the protected instance, run the encryption job you had earlier created in the trusted instance.
   a. Navigate to **System Security** > **Security Jobs** > **All**.
   b. Open the encryption job you had earlier created in the trusted instance.
   c. Click **Start** to start the job.
   A confirmation message is displayed mentioning that the records are signed.

Sign new flow, subflows, and actions

Sign and validate new flow, subflows, and actions by enabling the Code Signing in protected and trusted instances.

Before you begin

Role required: sn_kmf.cryptographic_manager

Procedure

1. In the trusted instance, start an update set.
2. In the trusted instance, create the required flows, subflows, or actions, and publish them.
   The flows, subflows, or actions are added to the update set.
3. In the trusted instance, change the state of the update set to **Complete** and click **Update**.
4. In the trusted instance, sign the update set by creating an encryption job.
   a. Navigate to **System Security** > **Security Jobs** > **All**.
   b. Click **New**.
   c. On the form, fill these values.

Field	Description
Name	Name to identify the record.
Type	Type of the encryption job. Select **Sign Update Set**.
Table	Update set from which the records should be signed.

   d. Click **Submit**.
   e. Click **Start** to sign the update set.
      - **Summary** is updated that the records are signed.
      - The update set is updated and includes the signature.
5. In the trusted instance, open the signed update set record and export it to an XML.
6. In the protected instance, import the signed update set.
   a. Navigate to **System Update Sets** > **Retrieved Update Sets**.
   b. Select the **Import Update Set from XML** related link to import the update set that is exported from the trusted instance.
   For more information, see Import and commit the quick-start update set.
   The update set is committed successfully.

Code Signing Health and Status Dashboard

The Code Signing Health and Status dashboard provides a centralized, user-friendly view of your Code Signing environment's health and configuration. Use it to identify issues, verify configuration accuracy, and support secure, uninterrupted code-signing operations.

The Code Signing Health and Status dashboard highlights configuration issues and includes intuitive guidance to help you resolve them quickly, reducing the need for escalation. You can use the dashboard to verify that Code Signing configurations are set up correctly and performing

securely. It helps you identify potential issues that could lead to failures and enables you to provide accurate guidance to the stakeholders involved.

The dashboard lets you search, filter, and rescan the following parameters:

- Configurations
- Certificate status
- Plugin details
- Signature verification certificate
- MID Server setup

To access the Code Signing Health and Status Dashboard, go to **All** > **Code Signing** > **System Health** > **Dashboard**.

Overview Dashboard

The Overview dashboard provides a centralized view of your Code Signing environment, offering real-time insights into key components and their status.

The Overview dashboard displays the following reports:

Overview Dashboard

Title	Type	Description
Enterprise plugin	Text field	Status of essential plug-ins and system properties required for Code Signing functionality. **Enterprise plugin** displays the core plugin needed for signing operations.
Opt-in property	Text field	Configuration flag that enables Code Signing capabilities. **Note:** This setting must be **Active** to confirm that code signing functions according to the specific requirement.
Signature validation status	Text field	Indicates whether signature validation is enforced during code execution to verify that only the trusted scripts are executed.
Signature states	Pie chart	Comprehensive overview of all the script signatures, including their current validation status. The signatures are categorized into:

Overview Dashboard (continued)

Title	Type	Description
		• **Trusted** signatures are valid and associated with trusted code signing certificate. • **Untrusted** signatures fail validation or are linked to untrusted verification certificate. • **Orphan** signatures are present but no longer linked to a known or active certificate. **Note:** Use this information to assess script integrity and take corrective action as needed.
MID Server configuration state	Pie chart	Status and trust relationship of all the MID Servers in your instance. Use this section to view the total number of MID Servers and check their status. • **Active MID Servers**: The number of servers that are currently running and successfully connected. • **Inactive MID Servers**: The number of servers that aren't running or are disconnected.
Signed records by type	Pie chart	Distribution of signed records by record type, showing the percentage coverage across the following categories:

Overview Dashboard (continued)

Title	Type	Description
		• **Business Rules**: The percentage of signed business rule records. • **Script Includes**: The percentage of signed script include records. • **Others**: The percentage of signed records in other supported categories.
Signed records by application	Pie chart	Distribution of signed records across different application modules, showing the percentage of code signing coverage. • **User Management**: Percentage of signed records in the User Management module. • **Form Validation**: Percentage of signed records related to form validation. • **Others**: Percentage of signed records in all other application modules.

Signature Verification Status

View the status of valid, invalid, and missing signatures across different applications to assess code signing coverage. Use this information to identify areas that may require additional attention or action.

The Signature Verification Status dashboard displays the following reports:

Signature Verification Status Dashboard

Title	Type	Description
Script name	Text field	Identifier or title assigned to a specific script within the system.
Type	Text field	Category of the script within the system. It helps define the role and functionality of the script.

Signature Verification Status Dashboard (continued)

Title	Type	Description
		Example: • Business rule • Script include • Client script • Flow action
Application	Text field	Module or area within the system where the script is applied. Example: • User management • Reporting • Form validation • Workflow • Notifications
Status	Text field	Indicates the current verification status of the script signature. • **Valid**: The script signature is verified and is trusted. • **Invalid**: The script signature isn't verified or is deemed untrusted. • **Missing**: The source document is eligible for signing but doesn't have a signature yet. Use this status to identify eligible records that remain unsigned and require code signing.
Last scanned	Text field	Date and time when the script was last scanned for signature verification in the format: `DD/MM/YY/H:S (Day/Month/Year/ Hour:Minute)`

Code Signing MID Server Configuration

Manage and configure the trust relationships and certificate settings for MID Servers.

The Code Signing MID Server Configuration dashboard displays the following reports:

Code Signing MID Server Configuration

Title	Type	Description
MID Server	Text field	Name of the specific MID Server instance.
Status	Text field	Current operational state of the MID Server. • **Active** • **Inactive**
Version	Text field	Software version of the MID Server currently in use. Use this information to: • Ensure compatibility with your instance. • Ensure that the server is running the latest supported features and security updates.
Last check-in	Text field	Most recent date and time that the MID Server successfully communicated with the instance in the following format: `DD/MM/YY/H:S (Day/Month/Year/ Hour:Minute)`

Key Pair and Certificates

The Key Pair and Certificates dashboard displays details about the cryptographic keys and digital certificates used for Code Signing. It includes information such as key type, certificate issuer, expiration date, and validity status. Use this dashboard to manage code signing certificate credentials, verify their validity, and help ensure secure and trusted Code Signing operations.

The Key Pair and Certificates Configuration dashboard displays the following reports:

Key Pair and Certificates Configuration

Title	Type	Description
Certificate Name	Text field	The unique identifier (name) assigned to a digital certificate used for Code Signing.
Crypto Module	Text field	The cryptographic component used to generate, store, and manage private keys for Code Signing. It ensures secure key operations and helps protect sensitive cryptographic

Key Pair and Certificates Configuration (continued)

Title	Type	Description
		material from unauthorized access.
Type	Text field	The classification of the cryptographic key or certificate used for Code Signing. Example: • ServiceNow • Third party
Expires	Text field	Date when the digital certificate expires and is no longer valid for use.
Status	Text field	Current validity of the digital certificate. • **Valid**: The certificate is active and can be used for Code Signing. • **Expired**: The certificate has passed its expiration date and is no longer valid. • **Expiring Soon**: The certificate is nearing its expiration date. Use this status to take proactive steps for renewal to avoid disruptions in code signing operations.
Chain	Text field	The certificate chain, which includes the digital certificate along with the intermediate and root certificates that establish a trust path. Select **View Chain** to view the certificate chain.

Code Signing Configuration

The Code Signing Configuration dashboard displays the system properties and key settings that control Code Signing in your environment, including flags and enforcement policies. These settings enable features, enforce signature validation, and define trusted sources.

The Code Signing Configuration dashboard displays the following reports:

Code Signing Configuration

Title	Type	Description
Setting	Text field	Name of a specific configuration property that controls a feature or behavior related to code signing. Each setting defines how the system should handle aspects like signature validation, enforcement, or trusted sources.
Value	Text field	Current state or input assigned to a specific setting. It determines how the setting behaves. For example, whether Code Signing is enabled (true) or disabled (false). The value directly affects the system's Code Signing operations and enforcement.
Last Updated	Text field	Most recent date and time when the setting was modified in the following format: `DD/MM/YY/H:S (Day/Month/Year/ Hour:Minute)`

Code Signing reference

Reference topics provide additional information to administer and troubleshoot Code Signing.

Properties installed with Code Signing

Code Signing adds the following properties.

Roles installed with Code Signing

Code Signing includes the following roles.

Troubleshooting and accessing logs

Access various logs to troubleshoot and identify the failure reasons.

Properties installed with Code Signing

Code Signing adds the following properties.

Property	Type	Description
com.glide.codesigning.expanded.tracking.enabled	true|false	When true, the validation length for the meta stack is increased for ecc_queue topics listed in the **com.glide.codesigning.expanded_tracking.topic.list** property.

Property	Type	Description
		Important: Elevated Security is needed to modify this property.
com.glide.codesigning.expanded.tracking.depth	Integer	Levels of Code Signing validation to occur when Code Signing is enabled. The default value is 3. **Important:** Elevated Security is needed to modify this property.
com.glide.codesigning.expanded_tracking.topic.list	String	Comma-separated list of topics to be subject to increased meta-stack tracking. **Important:** Elevated Security is needed to modify this property.
com.glide.codesigning.tracking.agent.validation.exclusion.list	String	Comma-separated list of ecc_queue agents for which Code Signing should be skipped
com.glide.codesigning.tracking.debug	true|false	When true, debug logging for the Code Signing tracker is enabled.
com.glide.codesigning.tracking.enabled	true|false	When true, enables Code Signing caller tracking. **Important:** Elevated Security is needed to modify this property.
com.glide.codesigning.tracking.logging.enabled	true|false	When true, enables logging for Code Signing tracking.
com.glide.codesigning.tracking.unsupported_script_tracking.enabled	true|false	When true, ecc_queue records inserted via unsupported scripts (if detected) aren't notarized. **Important:** Elevated Security is needed to modify this property.
com.glide.codesigning.tracking.validation.fail_fast	true|false	When true, Code Signing verification fails at the first script validation failure instead of verifying all scripts. **Important:** Elevated Security is needed to modify this property.
com.glide.event_handler.code_signing.tracking	String	Defines the Event Handler that helps ensure that customers newly enabling Code Signing are configured to be as secure as possible.

Property	Type	Description
		Important: Elevated Security is needed to modify this property.
com.glide.web_service_outbound.impl.codesigning.tracking	true|false	When true, enables SOAPMessageV2 Code Signing tracking **Important:** Elevated Security is needed to modify this property.
com.snc.csf.maximum_update_size	Integer	The maximum number of records allowed in a Code Signing update set. This value should be limited to a value between 6000 to 10000. If this value is exceeded, multiple update sets are generated and are linked to the same parent update set to enable batch processing. This limitation prevents UI issues described in KB0557104.
com.snc.csf.servicenow_root_of_trust.disabled	true|false	Whether the Root of Trust feature is active. The default value is `false`, meaning that ServiceNow build certificates are trusted. **Important:** This property can only be changed using a signed scheduled job from a user with the admin, security admin and KMF manager roles. For details on changing your Root of Trust, see Change your Root of Trust configuration.
com.snc.kmf.signature.validation.optin	true|false	When true, enables Code Signing on your instance. **Important:** This property can only be changed via a request to Customer Service and Support.
glide.jdbcprobeloader.tracking	true|false	Toggles Code Signing on or off for JDBC data sources. **Important:** Elevated Security is needed to modify this property.
glide.rest.codesigning.tracking	true|false	When true, enables RESTMessageV2 Code Signing tracking. **Important:** Elevated Security is needed to modify this property.

Roles installed with Code Signing

Code Signing includes the following roles.

Code signing admin [codesigning_admin]

Use the code signing admin role to assign codesigning_manager and codesigning_auditor roles to other users.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Avoid granting an admin role when more specialized roles are available.

Code signing manager [codesigning_manager]

Use the code signing manager role to create and update signature configuration, and create and run code signing jobs.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

None.

Code signing auditor [codesigning_auditor]

Use the code signing auditor role to view signature configurations and signing jobs. The auditor role does not have create or write access to code signing assets.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

None.

Troubleshooting and accessing logs

Access various logs to troubleshoot and identify the failure reasons.

Code Signing logs

If any of the ECC queue records is not signed by the Code Signing Tracker API, the unsigned messages and the required details are displayed in the Code Signing module. Navigate to **System Logs** > **System Log** > **Code Signing** to access the list of records that are not trusted.

For additional debug node logs, enable *com.glide.codesigning.tracking.debug* and set its value to `true`.

REST message signature validation failure on MID Server

Access the error message regarding to the signature validation failure, by navigating to **System Web Services** > **Outbound** > **REST Message** and opening the required REST message record.

Signature validation failure on MID Server

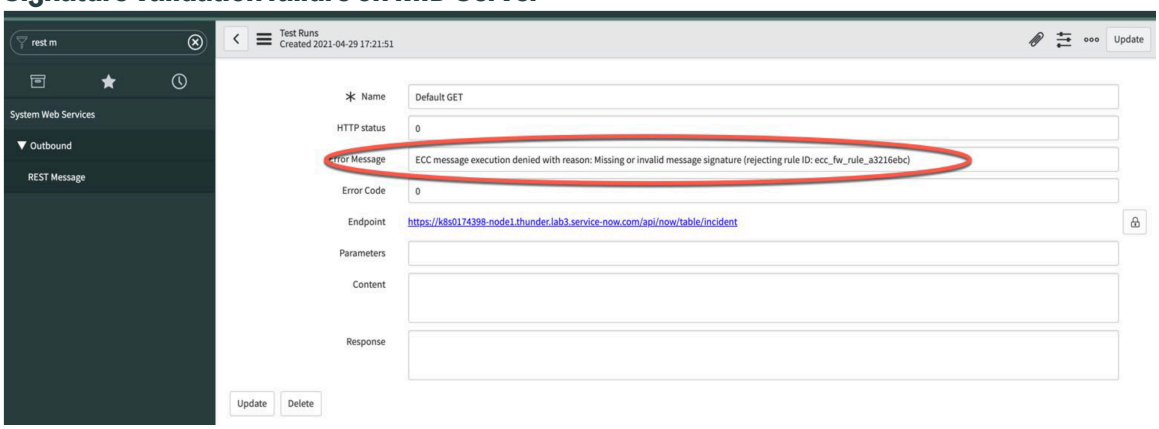

Note: Error messages related to the ECC firewall rejections start with `ECC message execution denied`.

ECC queue when signature validation fails on MID Server

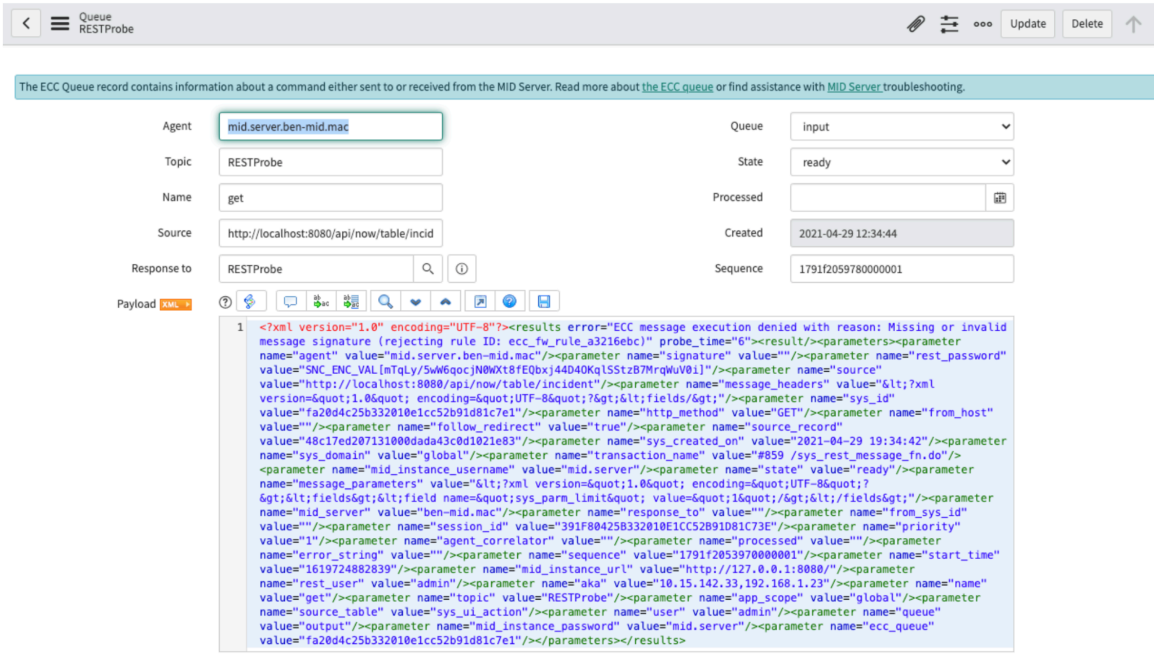

Error message when ECC message is blocked by the user rule

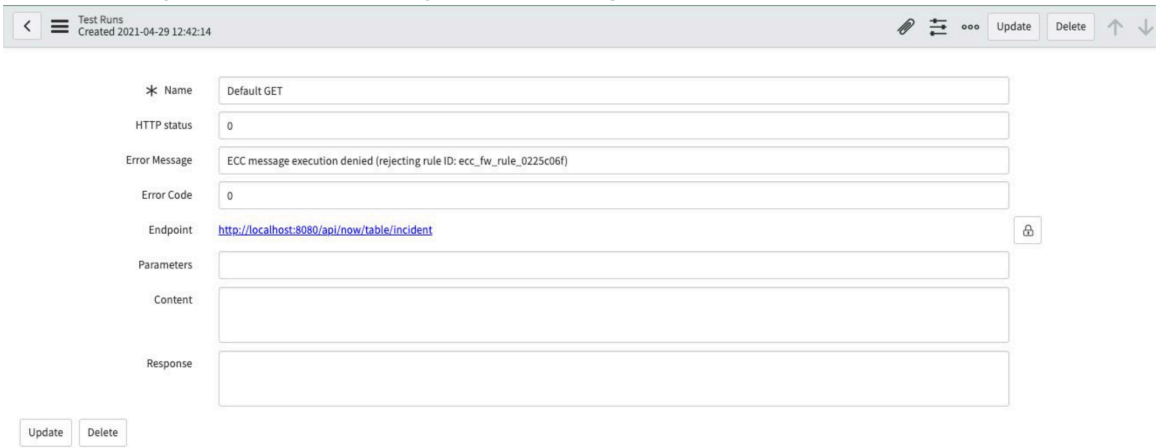

JDBC probe

When a JDBC data source with an invalid or missing signature is executed on a MID Server, an error message with the required details is displayed.

Source also displays the details of the error message.

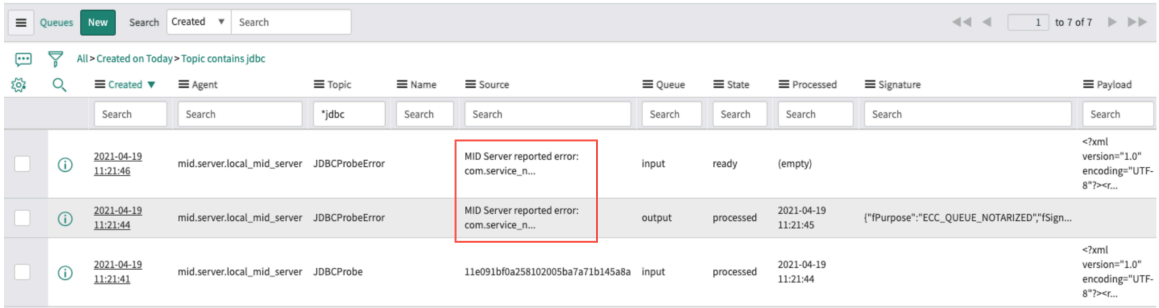

MID Server logs

To enable the detailed ECC firewall logging, increase the log level by setting the value of the MID Server configuration parameter, *`mid.log.level`*, to TRACE. The detailed logs provide information about:

- Rules that the MID Server loaded from the boot configuration file.
- Granular execution trace of rules.
- Specific rule due to which an ECC message is to be accepted or rejected.

Note: If `boot-config.xml` is invalid, the MID Server fails to start and the failure details are logged in the MID agent logs.

Antivirus Scanning

Use Antivirus Scanning to help protect your instance against virus infections that can be introduced by file attachments to your system records, such as incidents, problems, and stories.

Explore Antivirus scanning

Learn the value of Antivirus Scanning.

Configure Antivirus protection

Understand how to configure Antivirus protection.

Resolve Infected Files

Learn what to do with infected files.

Reference for Antivirus scanning

Know about the Dictionary Attributes for Antivirus Scanning.

Exploring Antivirus Scanning

Use Antivirus Scanning to help protect your instance against virus infections that can be introduced by file attachments to your system records, such as incidents, problems, and stories.

Antivirus Scanning scans file attachments stored in your attachment [sys_attachment] table to help protect users from uploading and downloading infected files. All the document types supported by the Platform are scanned by Antivirus Scanning.

If Antivirus Scanning is enabled, all file attachments in the Attachments table [sys_attachment.do] are scanned by default.

The Antivirus Protection plugin (com.glide.snap) is activated and enabled by default on your instance. As an administrator, you can deactivate and reactivate the Antivirus Scanning feature across your instance at the switch of a toggle, set configuration options, and review antivirus activity on the instance.

Note:

- Antivirus Scanning is also available for customers in the Government Community Cloud (GCC) and commercial environment.

  GCC Users must set the (com.glide.snap.fed_enable_scan) property to `true` to start using the feature.

  Commercial users must set *`com.glide.snap.enable_scan`* to `true`.

- HTTP and HTTPS communication protocols are supported.
- Edge-encrypted files are excluded from this scan.
- Antivirus definitions are updated everyday.
- Any file above 100-MB file size isn't scanned.

Email scanning

Inbound emails are scanned for viruses by the system email filters, not by Antivirus Scanning.

File Attachment field in a table

The addition of a **File Attachment** field in a table generates zz_yy tables. These tables are dynamic and virtual. They are automatically generated when the column type **file_attachment** is added to parent tables.

Consider adding a user photo to the Users table and incorporating it into the form view. When a photo is uploaded to a record, it automatically uploads the attachment to the sys_attachment table. The sys_attachment table maps the photo to the zz_yyUsers table.

By default only attachments attached to zz_yylive_profile tables are scanned. To scan other tables that have column type **file_attachment** create the system property *`com.glide.snap.scan.zz_yytables`* and insert the table name.

Example

The "zz_yyincident" and "zz_yycase" tables are dynamic tables created when the column is added to the parent tables: Incident and Case, then the property value should be *`zz_yyincident,zz_yycase`*.

After this property is set, attachments for the zz_yyincident and zz_yycase tables are scanned.

Scanning scenarios

Review these upload and download scenarios to understand how the system identifies potential security threats from files attached to your records.

Scenario 1 - Upload a file

1. The user unknowingly uploads an infected file to a record.
2. The system scans the file and moves it to quarantine.
3. The file appears in the Attachments window, where it's marked as unavailable.
4. The user selects the file and this error message appears: `The file Infected_testing.txt did not pass the security scan. Please remove the file from record INC0000059 and try again.`

5. The system sends an email notification to the user and the antivirus administrator.
6. The user closes the Attachments window and is returned to the record. The infected file is displays in the header as unavailable. Example, `infected_testing123.txtZ [unavailable]`.

Scenario 2 - Download a file

1. The user opens a record to download a file that is attached to it.
2. Unaware that the file is infected, the user selects it for download.
3. The system scans the file, moves it to quarantine, and displays a message similar to `The file infected_testing123.txt did not pass security scan and cannot be downloaded.`
4. The user closes the message and the screen refreshes showing that the file is unavailable.
5. The system sends an email notification to the user and the antivirus administrator.

Scenario 3 - Download a ZIP file

1. A user opens a record and downloads a ZIP file that is attached to it.
2. The system scans the ZIP files individually.
3. One file doesn't pass the security scan and is marked as unavailable. The remaining files are zipped and downloaded successfully.
4. The user opens the ZIP file and sees an "error.txt" file in addition to the successfully downloaded file. This file contains an error message specifying which file didn't pass scanning and was therefore not included in the ZIP.
5. The user opens the record again, and sees that the unavailable file has been moved into the **Potential security risks** section in the Attachments window and can't be downloaded.

Configuring Antivirus Scanning

Configure Antivirus Scanning across your instance and at the table level.

Before you begin

Role required: antivirus_admin or admin

About this task

Antivirus Scanning is active by default in your instance, where it automatically scans attachments to identify any files that are infected by viruses. Configure the feature by ensuring the scan is enabled across your instance, and by identifying any tables that you want to exclude from the scan.

Procedure

1. Navigate to **All** > **Antivirus** > **Configuration**.
2. As you configure the feature, consider the following.

Option	Description
Enable Antivirus scanning	Antivirus scanning is active and enabled on the instance by default, its toggle is set to the on position and appears green. **Note:** To set the property to be `false` contact customer support.
Allow attachments to be downloaded when Antivirus scanner is unavailable	If this option is set to the on position, antivirus scanning is bypassed if the scanner times out, and a response can't be obtained. In this situation, the file download proceeds without completing the scan. If the option is set to off, the file download is prohibited until the scan can complete successfully.
List of Tables Excluded	Any file attachments associated with a table in this list are excluded from antivirus scan ning. Proceed to Step 4 if you want to define the tables the system excludes from scan ning.

3. Select **Save**.
4. Exclude tables from the Antivirus scan by adding them to the **List of Tables Excluded**.
   a. Navigate to **System Definition** → **Dictionary**
   b. Search for the table you want to exclude from the scan and select the table with Type set to collection.
   c. In the **Attributes** tab, select **New**.
   d. Add `Exclude_from_antivirus_scan` in the Attributes field and enter `True` in the **Value** field.
   e. Select **Submit**.

Result
Antivirus Scanning is enabled in your instance, and the **List of Tables Excluded** on the Antivirus Configuration page is populated with all the tables that you excluded from the scan.

Reviewing quarantined files

Review quarantined file attachments and take further action as needed.

Before you begin
Role required: antivirus_admin or admin

About this task
Monitor your quarantined file entries in the Antivirus Quarantine page at regular intervals to perform any of the following actions.

Procedure

1. Navigate to **All** > **Antivirus** > **Quarantine**.
2. Select the check box next to each quarantined entry on which you want to perform an available action.
3. Select the **Actions on selected rows** drop-down in the banner to choose the action to be performed on the selected row quarantined entry.

Action	Description
Delete	Select this action to delete the quarantined file if it is not required any more by users.
Restore	Select this action to restore the file based on your assessment of its status or on the user request. However, scan the file with your third-party antivirus product before you re store it to mitigate any virus threats. Once re stored, the file is available for download to the user.
Download	Select this action to download the file to your local system for further scanning and analy sis.

Result
The system asks for confirmation and performs the selected action per your input.

Related topics
Instance Security Center

Antivirus metrics

Review antivirus activity

Review the Antivirus Activities log that tracks all activities that occur on potentially-infected files from the point that they are discovered and placed into quarantine.

Before you begin
Role required: antivirus_admin or admin

About this task
This log functions as a report that captures antivirus activity, such as discovery, deletion, and other possible quarantined file events.

Procedure
1. Navigate to **All** > **Antivirus** > **Activity**.
2. Review the log of quarantined files.
   You can view the actions performed in the Event column for each quarantined file.

What to do next
Determine which records you want to delete, restore, download, or keep in the log. See Review quarantined files

Understanding Dictionary attributes for Antivirus Scanning

Dictionary attributes alter the behavior of the table or element that the dictionary record describes. As an administrator, you can set the values of dictionary attributes to modify the behavior of the default Antivirus Scanning configuration.

Dictionary attributes for Antivirus Scanning

Name	Value	Target element	Description
Exclude_from_antivirus_scan	true/false	any table	If true, file attachments on the table are excluded from the antivirus scan. See Configuring Antivirus Scanning
Supress_antivirus_email_notification	true/false	any table	If true, stops sending Platform-generated email notifications when a potentiality-infected file is identified.
Suppress_antivirus_ui_notification	true/false	any table	If true, stops Platform-generated UI notifications when a potentially-infected file is identified.

Related topics

Altering tables and fields using dictionary attributes

HTML sanitizer

Remove unwanted code and protect against security concerns such as cross-site scripting attacks by sanitizing HTML markup in HTML fields and translated HTML fields.

Explore the HTML Sanitizer

Learn how the HTML sanitizer works.

Configure HTML Sanitizer

Configure the HTML sanitizer.

Activate HTML Sanitizer

Learn how to enable the HTML sanitizer.

Exploring HTML sanitizer

Remove unwanted code and protect against security concerns such as cross-site scripting attacks by sanitizing HTML markup in HTML fields and translated HTML fields.

Use HTML sanitization to ensure HTML content within your instance doesn't contain potentially harmful content. HTML sanitization works by removing HTML tags that could be used to compromise your instance, such as `<script>`, `<link>`, or `<embed>` tags that can be used to run unwanted scripts on your instance or direct your users to unwanted content. Safe tags that control the formatting of your content are preserved. As an administrator, you're able to customize what content is removed or preserved. You're also able to control whether sanitization applies to all content, or just fields you specify.

The HTML sanitizer works by checking the built-in inclusion list for markup that you always want to preserve. The sanitizer provides the *HTMLSanitizerConfig* script include that administrators can use to modify the built-in inclusion list. Items can also be added to the exclusion list to remove HTML markup. Contents of the exclusion list override the inclusion list.

The following types of items can be added to the inclusion and exclusion lists:

- Global attributes
- Any HTML elements

Note: By default, URL attributes like `href` and `src` support only these protocols:

- `http`
- `https`
- `mailto`
- `data`

For example:

```
<a href="https://community.servicenow.com/community">ServiceNow
Community</a>
```

Note: To learn more about the *`glide.html.sanitize_all_fields`* property that controls use of the HTML sanitizer, see Enable HTML Sanitizer [Updated in Security Center 1.3] in Instance Security Hardening Settings.

Configure urlAttributes and the protocols

You can configure urlAttributes and their protocols in the **HTMLSanitizerConfig** script include. In the script include, `HTML_WHITELIST` configures the inclusion list and `HTML_BLACKLIST` configures the exclusion list. For example:

```
HTML_WHITELIST : {
urlAttributes: { "protocols" : [ "file", "notes" ] },
                 -   -
                 -   -
        }
```

Because `notes` is in the inclusion list in this example, this URL isn't sanitized:

```
<a title="Lotus"
 href="Notes://
ABC/X575C90019DE33/ABC594DCB76D86EB4925653E0011C4C1/ZZ90B7E2D339
64749257EEA003456FD">Lotus</a></p>
```

Default Inclusion List

Note: The Default Inclusion List is a system list and is not accessible by users in the Instance.

```
BUILTIN_HTML_WHITELIST :{

 globalAttributes:{ attribute:["id","class","lang","title","styl
e"],

                                attributeValuePattern:{}},

    label:{ attribute:["for"]},

    font:{ attribute:["color","face","size"]},

    a:{ attribute:["href","nohref","name","shape"]},
```

```
 img:{ attribute:["src","name","alt","border","hspace","vspace",
"align","height","width"]},

 table:{ attribute:["border","cellpadding","cellspacing","bgcolo
r","background","align","no
 resize","height","width","summary","frame","rules"]},

 th:
{ attribute:["background","bgcolor","abbr","axis","headers","sco
pe","nowrap","height","width","align","valign","char
 off","char","colspan","rowspan"]},

 td:
{ attribute:["background","bgcolor","abbr","axis","headers","sco
pe","nowrap","height","width","align","valign","char
 off","char","colspan","rowspan"]},

 tr:
{ attribute:["background","height","width","align","valign","c
har off","char"]},

    thead:{attribute:["align","valign","char off","char"]},

    tbody:{attribute:["align","valign","char off","char"]},

    tfoot:{attribute:["align","valign","char off","char"]},

    colgroup:{attribute:["align","valign","char
 off","char","span","width"]},

    col:{attribute:["align","valign","char
 off","char","span","width"]},

    p:{attribute:["align"]},

    style:{attributeValuePattern:{"type":"text/css"}},

    canvas:{ attribute:["height","width"]},

    details:{ attribute:["open"]},

    summary:{ attribute:["open","valign","char off","char"]},

    button:{ attribute:["disabled","accesskey","type"]},

    form:{},

 input:{ attribute:["size","maxlength","checked","alt","src","ty
pe","disabled","readonly","accesskey","border","usemap"]},

    select:{ attribute:["disabled","multiple","size"]},
```

```
 textarea:{ attribute:["rows","cols","disabled","readonly","acce
sskey"]},

    option:{ attribute:["disabled","label","selected"]},

    div:{ attribute:["align"]},

    ol:{ attribute:["start","type","square"]},


 ul:
{ attribute:["type","square","itemscope","itemtype","itemref"]},

    li:{ attribute:["value","fb__id","itemprop"]},


 span:{ attribute:["color","size","data-mce-bogus","itemprop","f
ace"]},

    br:{ attribute:["clear"]},

    h3:{ attribute:["itemprop"]},

    html:{ attribute:["xmlns","lang","xml:lang"]},

    link:{ attribute:["rel","type","href","charset"]},


 meta:{ attribute:["name","content","scheme","charset","http-equ
iv"]},

    pre:{ attribute:["xml:space"]},

    noscript:{},    h1:{},    h2:{},      h4:{},    h5:{},
 h6:{},

    i:{},    b:{},    u:{},    strong:{},    em:{},    small:{},
    big:{},

    pre:{},    code:{},    cite:{},    samp:{},    sub:{},
 sup:{},

    strike:{},   center:{},  blockquote:{},    hr:{},
 map:{},

    dd:{},    dt:{},    dl:{},  fieldset:{},    legend:{},
 figure:{},  tt:{},

    body:{},   caption:{},   head:{},   title:{},  shape:{},},
```

Using variables and templates in HTML fields

HTML/Translated HTML fields undergo HTML sanitization by default. This process sanitizes the input HTML to protect it from cross-site scripting (XSS) and related security attacks. Storing

templates or variables such as ${description} or {{description}} or similar and replacing them with a true description post sanitization will reduce the effectiveness of the sanitization process. This is due to the sanitization being called solely on the placeholder template and not the HTML content. Storing only HTML content in the HTML/Translated HTML fields helps ensure an effective sanitization process.

Configuring HTML sanitizer

You must modify a script include to make configuration changes to the HTML sanitizer.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Script Includes**.
2. Open **HTMLSanitizerConfig**.
3. To add items to the exclusion list, use the HTML_BLACKLIST class.

   To add items to the inclusion list, use the HTML_WHITELIST class.

   Use this format:

```
HTML_XXXXLIST :{
        globalAttributes :{

            attribute:[attribute-name1,...],

 attributeValuePattern:{ attribute-name2:attribute-value-reg
ex-pattern,...}

        },<html-element-name>:{// Same as Above},----}
```

   - *globalAttributes* contains attribute or attributeValuePattern items that are applicable globally for all the HTML elements.
   - *attribute* is a comma-separated list of attributes.
   - *attributeValuePattern* is a dictionary of attribute to attribute-value-regex-pattern pairs. The attribute-value-regex-pattern is a regular expression which has to match the attribute value.

Example:
Consider the following example:

```
HTML_WHITELIST:{
        globalAttributes:{
            attribute:["id","name"],},
        img:{
            attribute:["style","align"],
            attributeValuePattern:{src:".*jpeg"}},
        iframe:{},}
```

It adds the following items to the inclusion list:

- The global attributes id and name. This is a list of strings that can be applied globally to all the elements.
- The img element where the attributes are style and align.
- The img element where the source attribute of the image is a file with the .jpeg extension. This is an example of a regular expression pattern that matches an attribute value.
- The iframe element.

Enabling HTML sanitizer

The HTML sanitizer provides a property to enable or disable the sanitizer for all HTML fields in the system.

Before you begin

Role required: admin

About this task

By default, the property is set to true for new instances.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
2. Set the properties *`glide.html.sanitize_all_fields`* and *`glide.translated_html.sanitize_all_fields`* to **true**.

   Note: To learn more about this property, see Enable HTML Sanitizer [Updated in Security Center 1.3] in Instance Security Hardening Settings.

Trouble?

If the properties do not exist in the System Properties table, you can add them.

Enable sanitization on individual fields

You can use field attributes to enable or disable the sanitizer on individual fields.

Before you begin

Role required: admin

About this task

You need to first set the sanitizer property to false, and then enable the sanitizer on a per-field basis for any form.

Procedure

1. Navigate to the sys_properties table and set the *`glide.html.sanitize_all_fields`* to **false**.
   This disables the sanitizer for all HTML fields in the system.
2. Navigate to the form that contains the HTML field.
3. Right-click the HTML field label, and select **Configure Dictionary**.
   The Dictionary Entry form opens for the HTML field.
4. Enter one of the following in the Attributes field:

- To disable sanitization enter `html_sanitize=false`
- To enable sanitization enter `html_sanitize=true`

5. Click **Update**.
6. To enable the HTML sanitizer for translated HTML fields, set the *`glide.translated_html.sanitize_all_fields`* property is **true**.

Enable HTML Sanitizer logging

When the HTML sanitizer removes elements or attributes, they are added to the system log.

Before you begin

Role required: admin

About this task

You can review these sanitized elements by adding /syslog_list.do?sysparm_query=source %3DHTMLSanitizer to your instance URL.

Procedure

1. To review these sanitized elements add `/syslog_list.do?sysparm_query=source %3DHTMLSanitizer` to your instance URL.
2. To enable or disable logging, add the `glide.html_sanitize.discarded_log.enable` property to the system properties and set the value to **true** (enabled) or **false** (disabled).
   This property is **true** by default.

Auditing

Track record changes on auditing-enabled tables. By default, the system tracks changes to the incident, change, and problem tables, among others.

Explore

Learn the features and business values of Auditing.

Configure

Understand how to configure Auditing.

View

Review the Sys Audit and Audit Relationship Change tables.

Reference

Know about the History sets.

Exploring Auditing

Track record changes on auditing-enabled tables. By default, the system tracks changes to the incident, change, and problem tables, among others.

Auditing overview

Enabling auditing tracks the creation, update, and deletion of all records in the table. If you want to audit individual fields in a table, you can hide fields you don't want to track using a dictionary attribute.

Auditing information is kept in the following tables:

- Audit
- Knowing about History sets

Warning: Auditing system tables that receive a large amount of traffic, such as workflow Contexts [wf_context] or Event Management Alerts [em_alert], can impact performance. For this reason, you can't audit the em_alert table as a whole. Instead, audit selected fields of interest. Set `audit=true` on both the em_alert table and the selected fields. Try to audit as few fields as possible.

Auditing users

Auditing has the following users.

- admin
- security_admin

Auditing benefits

Benefit	Feature	Users
Enable table auditing to track changes to all or some of the table's fields	Configuring auditing for a table	admin
Experience a more enhanced way of defining and configuring the audit capability	Configure auditing using Audit Management Console	admin
Automate and simplify the deletion of audit data	Setup your audit retention	security_admin

What to explore next

To learn more about using Auditing, see:

- Configuring auditing for a table
- Configure auditing using Audit Management Console
- Viewing Sys Audit and Audit Relationship Change tables
- Knowing about History sets

Auditing components

Explore the following auditing components for a better understanding of auditing tables, deletions, and exemptions.

Auditing parent and child tables

Tables don't derive the audit flags from parent or child audited tables.

- For example, if you enable auditing for the Configuration Items [cmdb_ci] table, only CIs stored in that base table are audited.
- Likewise, if you enable auditing for the Computers [cmdb_ci_computer] table, only the computer CI records are audited, including any fields on the Computers [cmdb_ci_computer] table that is derived from the Configuration Items [cmdb_ci] table.

Auditing system tables

By default, the system doesn't audit the deletion of a record from system tables. To audit a system table, add it to the list of tables in the `glide.ui.audit_deleted_tables` property list.

Auditing deletions from a form or list

By default, the system audits deletions of individual records from a form. To prevent auditing, set the table's dictionary attribute *no_audit_delete*.

The system audits deletions from a list when it **audit** is selected on the table dictionary, and the table isn't listed in the *glide.db.audit.ignore.delete* property.

Note: By default, the *glide.db.audit.ignore.delete* property isn't present in the System Property [sys_properties] table. To change the property, and its associated values, you must first manually add it. However, when manually added, it overwrites the following default values:

```
glide.db.audit.ignore.delete =
sys_mutex,sys_db_cache,sys_lucene_block,sys_lucene_file,sys_lucene_directo
cmdb_ci_windows_service, cmdb_sam_sw_install,
cmdb_software_instance, cmdb_sam_sw_usage,
sam_sw_counter_detail
```

To learn more about adding system properties, see Add a system property

It is to be noted that by default, the audit deletes are enabled whether the record is deleted from the form view, list view, or through a script/scheduled job.

Information audited

Auditing tracks the following record changes:

- Unique Record Identifier (sys_id) of the record that changed
- Field that changed
- New field value
- Old field value
- Number of times this record and field have been updated
- Date and time when the change occurred
- User who made the change
- Reason for the change (if any reason is associated with the change)
- Internal checkpoint ID for the record, if the record has multiple versions.

Information exempted from auditing

Some updates aren't audited despite enabling auditing on a table. For example, you may see 132 updates in a record's history, but only seven audited ones.

Auditing excludes the following information:

- Updates made by an upgrade.
- Updates are made through import sets.
- Records in parent or child tables.
- Fields with the no_audit dictionary attribute.
- System tables are not listed in the *glide.ui.audit_deleted_tables* property list.
- Fields that begin with the sys_ prefix (system fields), except the sys_class_name and sys_domain_id columns.
- UI Pages can sometimes trigger updates to a record without creating an audit log.

- Anytime an inactivity monitor touches a record. It prevents you seeing possibly hundreds of updates listed against an incident, with the noise drowning out the useful data.
- Manual changes to Performance Analytics scores.
- Applying Update Set
- Import XML

Auditing a table

For instructions on how to audit a table, see Configuring auditing for a table.

By default, the system tracks all fields in an audited table. You can audit a subset of fields in a table in one of two ways:

- You can enable auditing for the entire table, then exclude those fields you don't want to include. It's appropriate when you want to audit most, but not all, fields, and is referred to as an exclusion listing. For more information, see Exclude a field from being audited (exclusion listing).
- You can enable auditing for the table, but only for specified fields. It's appropriate when you want to audit only a small number of the table's fields and is referred to as an inclusion listing. For information on how to include a field using an inclusion listing, see Include a table field in auditing (inclusion listing).

Non-cancellable audit records

Reduce the chances of audit records not being recorded when a transaction is canceled with the new default setting.

Audits have been set to create a record immediately in the same transaction with the target record write operations. If the target record gets deleted, the audit still gets created and retained under **NCA Test Audit Delete** module.

Note: The enhanced audit process is enabled by default. If the `glide.db.audit.lazy` property is set to `True`, the enhanced audit process is disabled.

Prior to Zurich release, if a transaction is cancelled, certain auditable operations were missed being recorded. This is because the platform executes some operations between the record change and is cancelled before audit creation. But, now audits are created immediately after the record is changed, reducing the chances of a cancelled transaction aborting the operation before the audit is recorded.

Audits are now recorded in the same thread as the transaction. Earlier audits were created in a background thread. This change redefines the default value of the `glide.db.audit.lazy` property from True to False. This property is not usually defined in the Properties table because the majority of instances start using the new default value and behavior. On some instances, this property might be already present and set as True, which means that these instances won't be able to use this change to audit behavior.

Note: It is recommended to delete this property to leverage the update.

Configuring auditing for a table

You can enable table auditing to track changes to all or some of the table's fields.

Before you begin

Note: Encrypted fields aren't audited by design. This behavior isn't configurable.

Role required: admin.

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
   The system displays the list of dictionary entries. The list includes a row for each table as well as a row for each column (field) in the table.
2. In the list of dictionary entries, find the row corresponding to the table you want to audit, for example `cmdb_ci_computer`.
   You can distinguish the row for the table itself – versus a row for a column in the table – by finding the row with the correct table name, an empty entry for the **Column** name, and a **Type** of **Collection**.
3. Select the dictionary entry for the table.
   The system displays the dictionary entry form.
4. Check the **Audit** check box.
5. Select **Update**.

What to do next

If you want to audit only a few fields in the table Enable inclusion list auditing for a table. If you want to audit most – but exclude some – fields, see Exclude a field from being audited (exclusion listing).

Enable inclusion list auditing for a table

Enable a table to audit only those fields you explicitly designate. This is useful when you want to audit only a small number of fields in an audited table.

Before you begin

The table must be enabled for auditing.

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
   The system displays the list of dictionary entries. The list includes a row for each table as well as a row for each column (field) in the table.
2. If necessary, customize the list view to show the **Attributes** column.
3. In the list of dictionary entries, find the row corresponding to the table you want to audit, for example `cmdb_ci_computer`.
   You can distinguish the row for the table itself – versus a row for a column in the table – by finding the row with the correct table name, an empty entry for Column name, and a type of *`collection`*.
4. In the **Attributes** field for that row, enter *`audit_type=whitelist`*.

What to do next

Designate which fields you want to audit in this table.

Exclude a field from being audited (exclusion listing)

Prevent the ServiceNow AI Platform from tracking a subset of fields in an audited table by excluding those fields from an audit.

Before you begin

To exclude a field in a table from being audited, you must have first enable auditing for that table.

Role required: admin

About this task

Add a set of fields to an exclusion list when you want to audit most of the fields in an auditable table. If you want to audit only a few fields, follow the inclusion listing procedure instead.

Note: Disabling auditing on journal-based fields can impact the functionality of features, such as the Activity Formatter.

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. If necessary, customize the list view to show the **Attributes** column.
3. Navigate to the row corresponding to the table and field (column) you want to exclude from auditing.
4. In the **Attributes** column for that row, enter `no_audit`.

Include a table field in auditing (inclusion listing)

Track a subset of fields in an audited table by add those fields to an inclusion listing.

Before you begin

To add fields in a table to an inclusion list, you must have first enabled auditing for that table and enabled inclusion list auditing for that table.

Role required: admin

About this task

Add a set of fields to an inclusion list when you want to audit only a small number of an audited table's fields. If you need to audit most fields, and exclude only a few, follow the exclusion list procedure instead.

Procedure

1. Navigate to **All** > **System Definition** > **Dictionary**.
2. If necessary, customize the list view to include showing the **Attributes** column.
3. Navigate to the table and field (column) you want to the inclusion list.
4. In the **Attributes** field, enter *`audit=true`*.

Enable auditing for a system table

Deletions from tables with a sys_ prefix are not audited by default. To track deletions from these tables, add the table name to the *`glide.ui.audit_deleted_tables`* property. Enabling the Restore Deleted Records plugin adds several default values to this property.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Properties** > **UI Properties**.
2. Locate the **List of system tables (beginning with "sys_", comma separated) that will have the delete audited** property.
3. Add or remove table names.
   Table names should be separated by commas, without any spaces.
4. Select **Save**.

> **Note:** For more information about auditing, see Understanding the sys audit Table.

Configure auditing using Audit Management Console

Use Audit Management Console module to experience a more enhanced way of defining and configuring the audit capability within your instance.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Audit Management Console**.
   A list of all tables in your instance shows up.

   > **Note:** By default, the list of tables that has set Audit enabled shows up.

   You can select **Disabled** tab to see the list of tables that has set Audit disabled. You can also select the All tab to see all the tables in your instance with either enabled or disabled Audit option.

2. Select the table from the list you want to update the auditing configurations.
   The table with its respective columns show up. You can also see the total number of columns in the selected table.

   > **Note:** **Columns** and **Retention** are the tabs that show up on the table.

3. Modify the Audit toggle depending on if you want to enable or disable audit for the selected table.

   > **Note:** When Audit is enabled in a table, all the columns and fields within the table are enabled by default.

4. **Optional:** Select **Edit Audit Status** if you don't want all the columns to be enabled within an audit enabled table.
   The Select Columns to be Edited modal shows up.
5. **Optional:** Unselect the columns that you don't want to be enabled.
   You can also add any column by checking the column from the Available columns list.
6. Select **Save** to save the latest modifications.
   Select **Clear All** to remove all the columns being enabled. See Setup your audit retention for more information about audit data retention.

Setup your audit retention

Use the Retention option to automate and simplify the deletion of audit data as per your requirement.

Before you begin
Role required: security_admin

Procedure
1. Navigate to **All** > **Audit Management Console**.
2. Select the table from the list you want to update the retention policy.

   Note: By default, you land on the Columns tab at the end of this step.

3. Select the Retention tab to update the retention policy for the selected table audit data.
   A modal showing the retention option shows up.
4. Enable the Automatically Purge Audit Records toggle.
5. Select the duration from the Duration dropdown menu.
   Select **Yes** if you want to proceed with the selected duration. You can also select **Cancel** if you want to select a different duration.

   Note: Logs older than your set duration will be purged and can't be restored.

   The Retention Policy modal shows up confirming the selected duration.
6. Select **Save** to update the retention policy for the selected table.
7. **Optional:** To disable deletion of Audit records of a table, disable the Automatically Purge Audit Records toggle and select **Save**.

   Note: The audit records that were previously deleted due to the selected retention duration are permanently unavailable.

Viewing Sys Audit and Audit Relationship Change tables

The ServiceNow AI Platform tracks inserts and updates to audited records in the Sys Audit (sys_audit) and Audit Relationship Change (sys_audit_relation) tables.

The ServiceNow AI Platform tracks audit tables. To track tables, select the **Audit** check box in the dictionary record to set the value to true. By default, it does not audit records from system tables, such as update sets tables.

Note: To prevent performance issues and infinite loops, the system skips any business rule or workflow triggered by inserts to the Sys Audit table.

Sys audit table columns

Access the Sys Audit table on your instance by entering `sys_audit.list` in the navigation filter.

The following columns appear in the sys_audit table records:

Field	Description
Table name	Table that the audit record is for (for example, "`incident`")
Field name	Column in the table that the audit record is for (for example, "`assigned_to`")

Field	Description
Document key	Sys_id (Unique Record Identifier) for the originating record associated with the audit record.
User	Name of the user who created the change. **Note:** Some automated processes use the system or guest user to apply and track changes to records. For more information, see System and guest users.
Old value	Old value of the field change represented by this sys_audit record. • Reference fields: Unique sys_id value of the changed record. • Date and time fields: Value in Coordinated Universal Time (UTC) as stored in the database.
New value	New value of the field change represented by this sys_audit record. • Reference fields: Unique sys_id value of the changed record. • Date and time fields: Value in Coordinated Universal Time (UTC) as stored in the database.

How the Audit Relationship Change (sys_audit_relation) table works

The Sys Audit [sys_audit] table tracks changes to reference fields in tables that are flagged for auditing. This activity includes journal field entries and history sets. The Audit Relationship Change [sys_audit_relation] table tracks relationship changes between sys_audit table records and the source tables that the audited records originate from. It also tracks when a record might have been deleted.

-
- Whenever you audit any record in a table, a relationship is created between the various originating tables to the store that records data. This relationship information is saved in the sys_history_set, sys_history_line, and sys_journal tables.
- If you delete a field that is related to an audited table record, the sys_audit_relation table records the deletion. In other words, anytime you change an audited record, it first deletes past elements, and then creates a relationship in the sys_audit_relation table with new document IDs.

Knowing about History sets

The system automatically generates History Set records as needed from the Audit table when a user either creates a record or views its history.

If a record is in an audited table, its history set is generated when the record is inserted or when a user views the record.

Note: Don't use history sets to generate reports.

Several fields of information are captured in the History Set record, displayed in the list view.

List View Record Fields

Field	Input Value
ID	Document ID for the record whose history is being recorded.
Table	Audited table for the record whose history is being recorded.
Load Time	Amount of time it took to generate the history set.

Audit History Record Fields

Field	Input Value
Label	The label of the field that was changed.
Old	Value before the change.
New	Value after the change.
Type	Indicates if the entry is for a normal field, an email record, or a relationship change record.
Update Number	The number of times this field has been changed. A value of -1 indicates when the record was created or deleted.
Update Time	Date and time of the change. **Note:** The Update time for auto-generated history lines doesn't match the Created or the Updated time for a record in a specific processing situation. When you view a history set of a record for the first time, an initial set of history line records is auto-generated. Since file changes in an upgrade aren't audited, this date mismatch occurs when: • You view a history set after a change is made to a record, but • Before another change is made to it in a future upgrade.
User Name	Name of the user who created the change.

History Sets in a Calendar View

After History Sets are active, the History context menu choice populates using information from the history set, rather than from the `sys_audit` table. From the user's perspective, the same historical data is available in the same user interface, but how the information is stored is different.

Since the History view includes a calendar view, but doesn't use the normal list interface to filter and interact with the history records, it enables:

- Searching and filtering historic data.
- Exporting historic data.

Viewing history sets

There are two ways of viewing the history set, accessible through the Context Menu action **History**.

Differences Between Audit and History Sets

The Audit [sys_audit], History Sets [sys_history_set], and History [sys_history_line] tables store the same data, but they serve different purposes and manage data differently.

Audit [sys_audit] table

The Audit [sys_audit] table is where the system stores historical information for all records. These records are intended to be kept forever so that administrators can always track the history of audited records. As the number of auditing records grows over time, it becomes more inefficient to directly query the Audit table for historical information. It is much more efficient to run queries only on the smaller subset records you actually want to view historical information for.

History Set [sys_history_set] table

The History Set [sys_history_set] table identifies which particular records from an audited table have historical information. The History [sys_history_line] table stores the actual changes to field values that occurred.

- The system automatically generates History Set and History records as needed from the Audit table when a user either creates a record or requests its history.
- Rather than containing a complete history of all changes in the system, History Set and History records only contain a recent subset of historical information for records where users have created or requested such information.
- In addition to audit data, history sets also include the information that is set during record insert, including journal field entries. Journal field entries you create before creating a record are handled in the same manner as journal entries created at the time of record creation. These journal entries appear in history sets with the same creation time and created by data as the associated record itself.

The system limits History Set and History records by:

- Having the table cleaner delete History Set records that have not been updated in 30 days.
- Using table rotation to rotate between four History tables every seven days. Because the platform truncates the leading table when rotation occurs, the maximum available retention period is 21 days, not 28. One of the four tables is always in the process of being cleared in preparation for the next rotation.

Should someone need historical information again at a later date, the system can regenerate it from auditing source records.

After the system generates History Set records, the context menu choice **History** uses the History Set rather than Audit records. From the user's perspective, the same historical data is available in the same user interface, but the way the information is stored is different.

Control access to history

You can give a role access to view audit history by setting a system property.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Properties** > **System**.
2. Select the *glide.history.role* property from the table.

3. In the *List of roles (comma-separated) that can access the history of a record*, select the user roles you want to access history.
4. Select **Save**.

Result
Any changes to a field are omitted if a user without read-access views the history of a record.

Change the number of history entries

By default, the history displays a maximum of 250 history entries, but you can change this value.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Properties** > **System**.
2. Select the *glide.history.max_entries* property.
3. Set the *Maximum number of field entries displayed in record history, default is 250* value with a new maximum number of entries to be displayed.

History List

The history list displays each change as its own row in the change list.

View History List

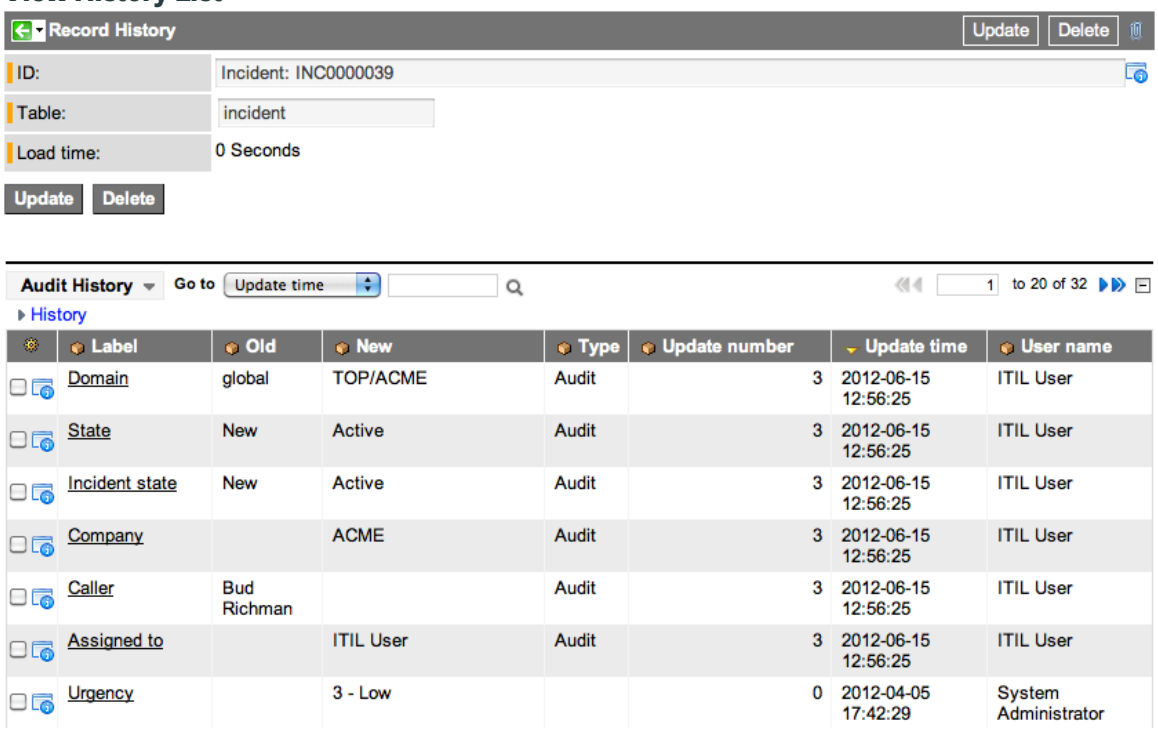

	Label	Old	New	Type	Update number	Update time	User name
	Domain	global	TOP/ACME	Audit	3	2012-06-15 12:56:25	ITIL User
	State	New	Active	Audit	3	2012-06-15 12:56:25	ITIL User
	Incident state	New	Active	Audit	3	2012-06-15 12:56:25	ITIL User
	Company		ACME	Audit	3	2012-06-15 12:56:25	ITIL User
	Caller	Bud Richman		Audit	3	2012-06-15 12:56:25	ITIL User
	Assigned to		ITIL User	Audit	3	2012-06-15 12:56:25	ITIL User
	Urgency		3 - Low		0	2012-04-05 17:42:29	System Administrator

Click on a row item to view additional details about the change.

View List Change

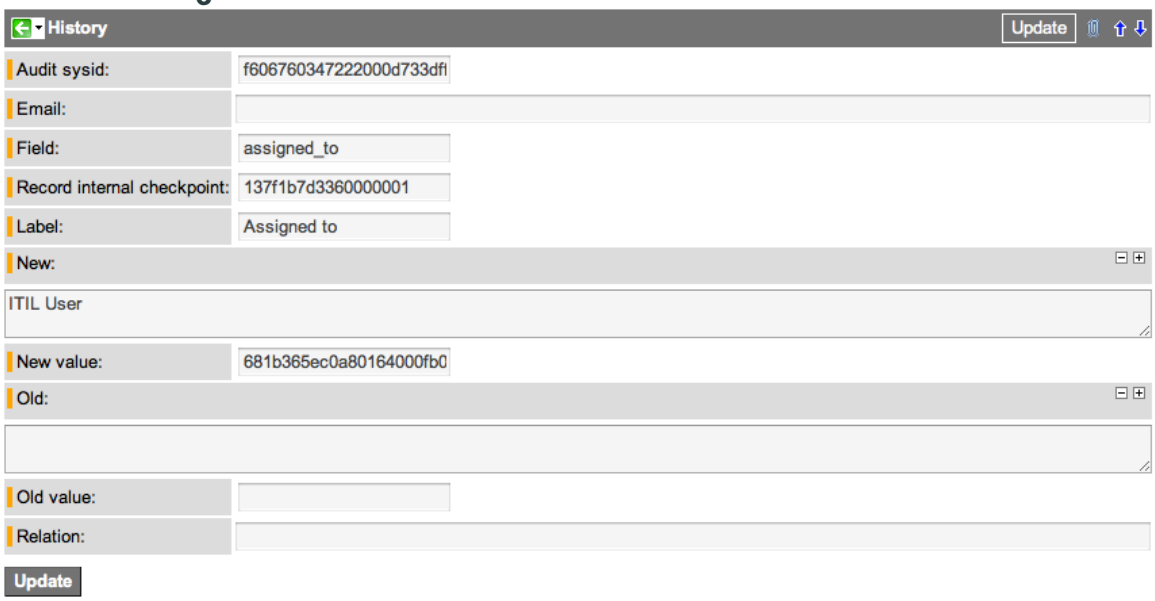

Requirements

To view a history list, the following requirements must be met.

Auditing

Auditing for the table must be enabled to view a history list.

ACLs

By default, the *List* history option is only available to users with the admin user role. To enable this option to non-admins, create a custom ACL rule granting read access to the Record History [sys_history_set] table.

Roles

At least one of the roles that the user has must be included in the *`glide.history.role`* property, which includes the itil role by default.

History Calendar

The history calendar shows you the days where the record was changed, who made the change, and when.

The History Calendar is sorted by update number. Each user is assigned a color so you can tell at a glance how many times a record was changed by a specific user. For example:

Changes made by system administrator

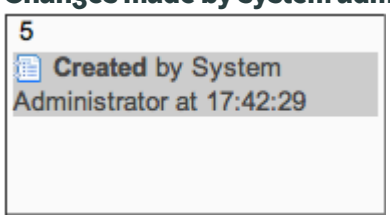

Changes made by ITIL user

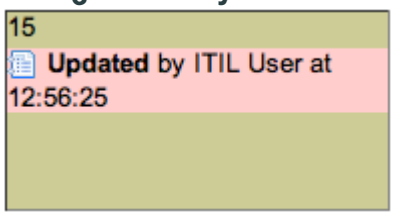

To highlight changes to a particular field, select the field from the **Highlight changes to field** selection box. Picking a field from this selection box changes the calendar to highlight the times that field was changed. Hover over the text of one of a highlighted change to see the change in value.

View highlighted changes

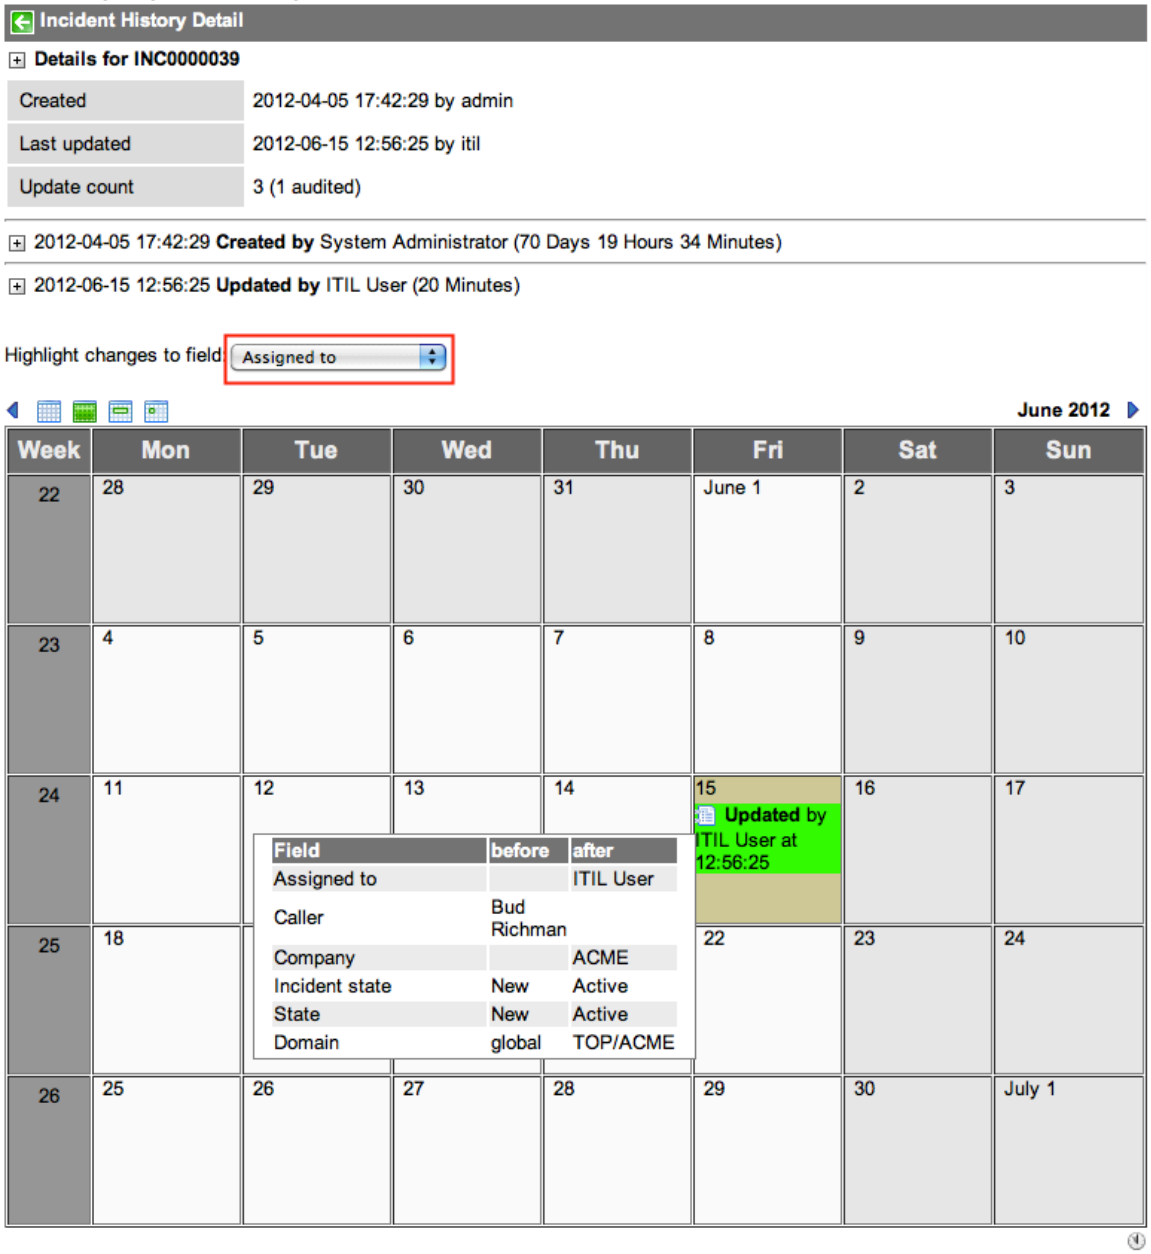

If you hover over the icon within an entry, a popup displays all the value changes. This is the same information that is displayed in the top part of the form.

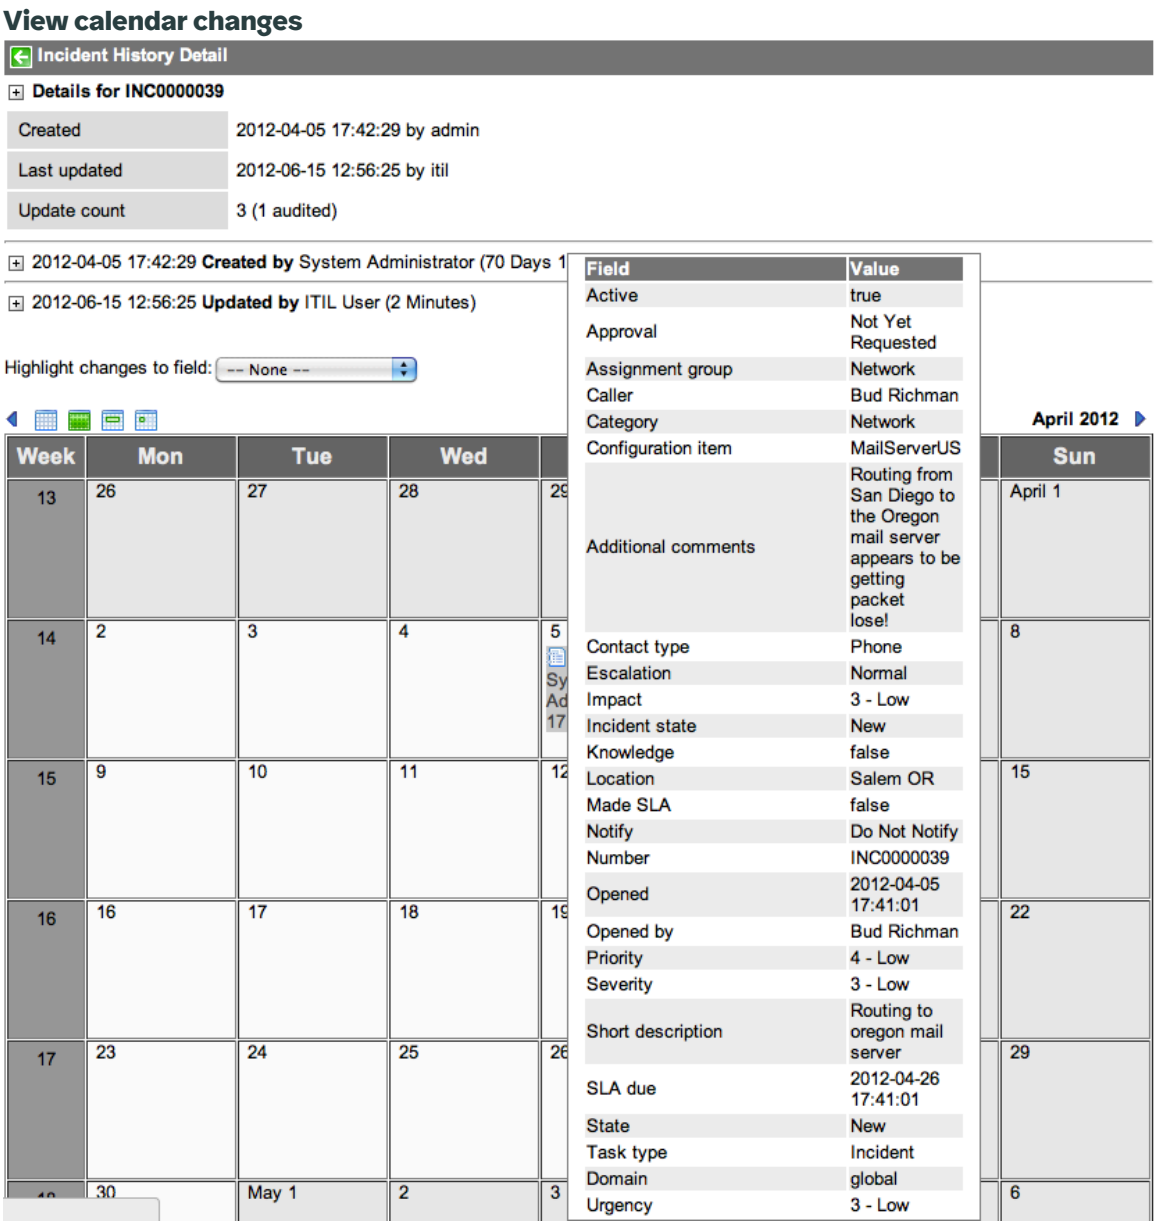

You can click on the day number to get a view of the changes for that day. You can also click on the week number to the left to get a the week view. You can scroll to and from month to month to see changes.

History Timeline

You can view a timeline of changes for a CI and for its related records, relationships, baselines, and proposed changes for the CI. Timelines are available for CIs in the Configuration Item [cmdb_ci] table or a descendant of this table, if auditing is enabled for the tables.

Role required: The ACL for this view is based on the roles defined in the *`glide.history.role`* system property, which by default is set to `itil`. Also, the user must have read access to the History Set [sys_history_set] table, which by default is granted to admin.

You can open a timeline when you view the history of a CI. You can specify the time period, time range, and properties that are displayed in the timeline. You can view either what has changed in a particular change set, or view the entire CI to better troubleshoot any issues. You can also display a timeline of changes to the CI's related records, and export and compare snapshots of the CI at any point in time.

CI changes are represented by bubbles in different shapes and colors along the timeline. The shape of each bubble represents a different type of change and the color of each bubble specifies whether the change is valid or invalid. CI baselines are represented by black circles that you can hover over to display more details. Click the **?** icon to display bubble shape and color definitions, and point to a bubble to display details about the change set.

A change to a relationship is considered valid only if it was applied through change management. If the change was applied via the Proposed Changes framework, it is valid. For additional validation steps, see Create or edit a planned validation script .

History Timeline view

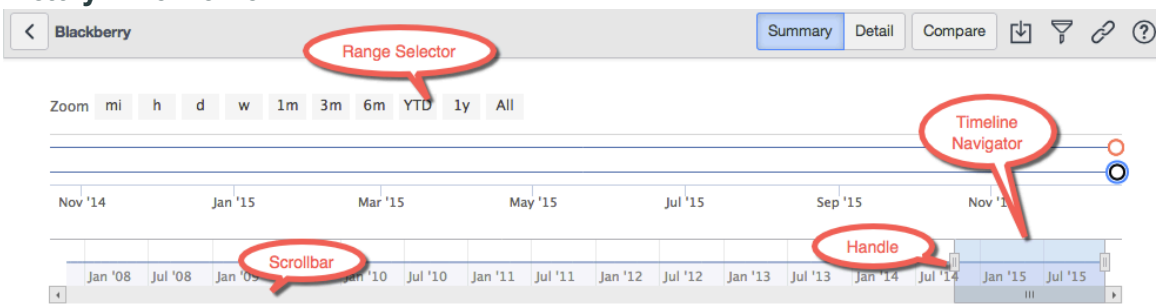

Timeline bubbles

Note: Proposed changes that do not have a planned start date are placed at future points of time.

Timeline navigator

Use the handles on both ends of the timeline navigator to extend or to shorten the time period that is shown.

You can scroll to a different period of time by clicking on the bottom part of the timeline navigator and then dragging the navigator to the left or right.

Zoom

By default, the timeline for the last month is shown. Next to the **Zoom** label above the timeline, you can select another time interval. You can select intervals from a minute to the entire period of data.

If there are many changes of the CI during the time period, the bubbles displayed might get too crowded. You can zoom in or out to spread the bubbles in either method:

- Change the time interval on the timeline. As you shorten the time interval, you zoom in, and as you lengthen the time interval, you zoom out.
- Select the section of the timeline that you want to zoom into.

Property filter

You can filter the bubbles that are displayed. By default, all bubbles are displayed, representing changes to all of the CI's properties. You can limit the view to display only the bubbles in which selected properties have changed and exclude bubbles in which only unselected properties changed.

The **Detail** and **Summary** views highlight properties within your filter scope that have changed. The changed properties are highlighted in light blue.

In the **Summary** view, you can choose to include all the properties of the CI, or only properties that have changed. If you choose to display all properties in the summary view, then changed properties are listed before unchanged properties.

Summary view

The **Summary** view displays snapshots of the CI's represented by each bubble. Each snapshot displays the changes to the CI's fields and relationships according to the change set. It displays old and new values before and after the change, and any relationships that were added or deleted.

Use the **>** and **<** buttons on both sides of the snapshot display to scroll through the next and previous change set records in a chronological order.

Detail view

The **Detail** view displays snapshots of the CI that correspond with the bubbles. Each snapshot includes the fields that are within the property filter scope, displaying the properties that have changed with a light blue background. Click on a bubble to display its corresponding snapshot of the CI. The data that is displayed is read-only.

Use the **>** and **<** buttons on both sides to scroll through the next and previous change set records in a chronological order.

View timeline of changes to related records

On the time line of changes for a CI record, you can also view a timeline of changes for the CI's related records.

Before you begin

Role required: admin

- Target table: the CI record must be in the Configuration Item [cmdb_ci] table or a descendant of this table.
- Auditing: must be enabled for the table containing the CI.

Procedure

1. Open the timeline for the CI.
2. Select the **Related Records** icon and select related records from the **List of Related Records** to view.
   Click the Related Records icon again to display the related records timeline.

Result

The timeline of changes to the CI's related records is displayed right above the CI's timeline. If you uncheck all related records, the related records timeline is hidden.

What to do next
Hover over a change bubble on the related records timeline to display details about the change, such as date and number of changed properties. As you change the time interval in focus, or zoom in or out, it affects both the CI timeline and the related records timeline simultaneously.

Export a snapshot of a configuration item

You can export a snapshot of a configuration item from its timeline.

Before you begin
The configuration item must be in the Configuration Item [cmdb_ci] table or a descendant of this table. Auditing must be enabled for the table containing the CI.
Role required: admin.

About this task
You can export a snapshot of the CI to an XML, PDF (Portal), or PDF (Landscape) format.

Procedure
1. Navigate to **All** > **Configuration** > **Base Items** > **All** to open the configuration item list.
2. Open a configuration item record.
3. Open the timeline for the record.
4. Select the bubble representing the time for which you want to export a snapshot of the CI.
5. Click the export icon ( ).
6. Select the file format to use for the export.
   You can download the file to your system for viewing.

Compare CI snapshots

You can compare the properties and relationships of a CI at two different points in its timeline.

Before you begin
The CI must be in the Configuration Item [cmdb_ci] table or a descendant of this table. Auditing must be enabled for the table containing the CI.
Role required: admin.

Procedure
1. Open the timeline for the CI.
2. Click **Compare**.
3. Select a **Start** date and an **End** date.
4. Click **Compare**.

Tracking changes to reference fields

Administrators can track changes to reference field display values.

Since reference fields only store an ID value, the system can normally only audit changes when the ID value changes. By default, the system does not audit changes when a reference field display value changes.

Consider the following situation. A user changes their name from Jane Smith to Jane Miller. Since the user name is the display value for the User table, any previous reference to Jane Smith instead refers to Jane Miller. If the administrator just updates the name of the existing

user record, audit and history records will only display the new name Jane Miller. By default, the system does not provide a way to distinguish between changes made under the original user name versus those made with the new user name.

If your auditing policy requires tracking user name changes, you can:

- Create a new user record for the new name and deactivate the previous user record. The system preserves audit records for the old user name and creates future audit records with the new user name.
- Create custom fields and a business rule to save the previous name and the date of the name change. The system can use this information to construct the proper names in audit and history records.

Tracking inserts

By default, the system does not create Audit records for inserts because in a typical instance, inserts can account for over 80% of the size of the Audit table.

Not tracking inserts allows for better performance and a much smaller Audit table. Administrators can enable auditing of inserts by setting the *`glide.sys.audit_inserts`* property to *true*.

Tracking CI Relationships

Changes to a CI relationship (CI Relations, CI/User Relations, or CI/Group Relations) appear in the history of the items on both sides of the changed relationship regardless of whether the change was manual or a result of Discovery.

For example, if the computer alpha has a used by CI Relation with the computer beta, then the history for alpha has a record of when the relationship with beta was established, and likewise, the history for beta has a record of when the relationship with alpha was established. This example illustrates the history displayed when some CI Relations are established, and then one of the relations is removed:

CI Relationship History

⊟ 2008-12-03 10:49:37 **Updated by** Guest (19 Hours 54 minutes ago) - CI Relationship Change

- created 2 Days 2 Hours 51 minutes earlier

Relationship	Before	After
Runs	(relationship added)	Tomcat@tomdmac
Runs	(relationship added)	MySQL Server@tomdmac

⊞ 2008-12-03 10:49:38 **Updated by** Guest (19 Hours 54 minutes ago) - Mac OS X - Disks

⊞ 2008-12-03 10:49:43 **Updated by** Guest (19 Hours 54 minutes ago) - Mac OS X - Active Connections

⊞ 2008-12-03 10:50:02 **Updated by** Guest (19 Hours 53 minutes ago)

⊟ 2008-12-04 06:43:39 **Updated by** Glide Maintenance (just now) - CI Relationship Change

- last activity was 19 Hours 53 minutes earlier
- created 2 Days 22 Hours 45 minutes earlier

Relationship	Before	After
Runs	MySQL Server@tomdmac	(relationship removed)

The created bullet indicates the date that the CI, user, or group was created. The last activity bullet refers to when the relationships were last changed. If you don't want to show CI

relationship history for any or all CI relationship types, you can turn it off by disabling auditing on the CI relationship tables (CI Relationship [`cmdb_rel_ci`], CI/User Relationship Type [`cmdb_rel_user_type`], or Group Relationship [`cmdb_rel_group`]).

High Security Settings

High Security Settings refer to several security options available in your instance.

Explore High Security setting	Activate High Security Settings
	
Learn the features and business values of High Security Settings.	Activate the High Security Settings.

Exploring High Security Settings

High Security Settings refer to several security options available in your instance.

The High Security Settings module is activated with the High Security Settings plugin, which is active by default on new instances. If High Security Settings are not active on your instance, see Requesting High Security Settings activation. To learn more about this plugin, see Enable High Security Plugin [Updated in Security Center 1.3] in Instance Security Hardening Settings. Properties for these types of high security settings are available:

- Default property values: To harden security on your platform by centralizing all critical security settings to one location for management and auditing.
- Default deny property: Provides a security manager property to control the default security behavior for table access.
- Security Administrator role: Provides a role to prevent modification of key security settings and resources. The Security Administrator role is not inherited by the admin role and must be explicitly assigned.
- Elevated privileges: Allows users with the security admin role to operate in the context of a normal user and elevate to higher security role when needed.
- Property access controls: Allows security administrators to set the roles required to read and write properties.
- System logs: Are read only.
- Access control rules: Control what data users can access and how they can access it.

Note:

- High Security Settings also automatically activates the Contextual Security plugin, if it is not already active. In addition, Platform Security Settings - High delivers settings and features in the context of increasing the security of your instance.
- The Instance Security Hardening Settings content contains detailed descriptions, and compliance values, for the security-related system properties and plugins in the ServiceNow AI Platform.
- To learn more about each of these properties, see Hardening settings.

There are two ways to set or change High Security Settings properties.

- Navigate to **System Security** > **High Security Settings**.

  Options on the High Security Properties page are **Yes** or **No**.

- Navigate to the **sys_properties.list** and search for the property you want to set or change.

  Options in the System Properties table [sys_properties.list] are **true** or **false**.

Property access control

Two additional columns are created in the Properties [sys_properties] table when High Security Settings are active:

- **read_roles**: A comma-separated list of role names that are allowed to read all fields of this property.
- **write_roles**: A comma-separated list of role names that are allowed to write/modify all fields of this property.

Properties listed in the Properties table have **read_roles** of admin, and **write_roles** of security_admin. Users with the admin role can view and read the property values, but must elevate to the security_admin role to modify them.

Notifications

Activation of high security settings also activates security warning messages. The following is an example of a message that appears after an approval.

Security Warning notification

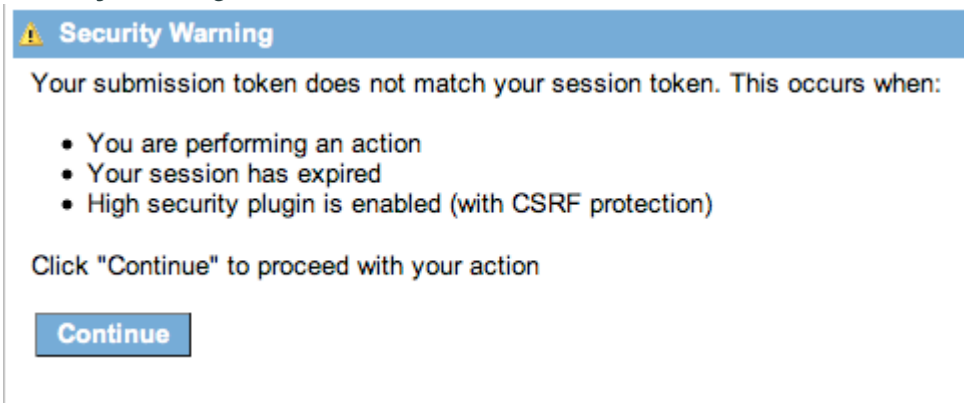

High Security Settings properties

Property	Description	Default Value	Instance Security Hardening Settings
glide.ui.escape_text	Escape XML values at the parser level for the user interface. Prevents reflected and stored cross-site scripting attacks. This property is not	Yes	Escape XML markup [Updated in Security Center 1.3]

Property	Description	Default Value	Instance Security Hardening Settings
	applicable in Service Portal. **Note:** This property is set to **true** by default in Vancouver and later releases, and can't be changed by administrators. For a use case where the property has to be changed, contact customer support.		
glide.ui.escape_all_script	Forces all expressions within Jelly JavaScript `<script type="text/javascript">` tags to be escaped by default. Enforces escaping only if the type attribute in the `<script>` tag is empty, or if the value is `text/javascript`, `text/ecmascript`, `application/javascript`, `application/ecmascript`, or `application/x-javascript`.	Yes in new instances	Escape jelly script [Updated in Security Center 1.3 and 1.5]
glide.ui.rotate_sessions	Rotate HTTP session identifiers	Yes	Rotate HTTP session identifiers

Property	Description	Default Value	Instance Security Hardening Settings
	to reduce security vulnerabilities. See: http://www.owasp.org/index.php/Session_Management#	**Note:** If you are using the SAML 2.0 plugin for Single Sign-on authentication, set this property to No. Otherwise, it interferes with the session information sharing that takes place between the instance and the Identity Provider.	fiers .
glide.ui.secure_cookies	Enable secure session cookies: Enable additional cookie security. If **Yes**, strict session cookie validation is enforced.	Yes	Enforce strict security of session cookies [Updated in Security Center 1.3]
glide.security.password_reset.uri	For mobile Password Reset, URL that the user is taken to when the user clicks the **Forgot password?** button.		None
glide.security.strict.updates	Double-check security on inbound transactions during form submission (rights are always checked on form generation).	Yes	Double check inbound transactions [Updated in Security Center 1.3]

Property	Description	Default Value	Instance Security Hardening Settings
	Note: This property is set to **true** by default in Vancouver and later releases, and can't be changed by administrators. For a use case where the property has to be changed, contact customer support.		
glide.security.strict.actions	Check conditions on UI actions before execution. Normally conditions are checked only during form rendering.	Yes	Check UI action conditions before execution
glide.security.use_csrf_token	Enable usage of a secure token to identify and validate incoming requests. This token is used to prevent cross-site request forgery attacks.	Yes	Enable Anti-CSRF token [New in Security Center 1.3, updated in 1.5, and removed in 2.0]
glide.ui.escape_html_list_field	Escape HTML for HTML fields in a list view.	Yes	Escape HTML in list views [Updated in Security Center 1.3 and 1.5]
glide.html.escape_script	Escape JavaScript tags in HTML fields.	Yes	Escape JavaScript [Updated in Security Center 1.3]
glide.ui.forgetme	Remove the **Remember me** check box from the login page.	Yes	Remove remember me
glide.smtp.auth	Authenticate with the SMTP server by the user name and password properties.	Yes	

Property	Description	Default Value	Instance Security Hardening Settings
	Note: This property is deprecated.		
glide.soap.strict_security	Enforce strict security on incoming SOAP requests. Requires incoming SOAP requests to go through the security manager for table and field access and checks SOAP users for the correct roles for using the web service.	Yes	Enforce SOAP request strict security [Updated in Security Center 1.3]
glide.basicauth.required.wsdl	Require authorization for incoming WSDL requests. **Note:** If you choose not to require authorization for incoming WSDL requests, you must modify the Access Control (ACL) rules to allow guest users to access the WSDL content.	Yes	Require authorization for WSDL request [Updated in Security Center 1.3 and 1.5]
glide.basicauth.required.csv	Require basic authorization for incoming CSV requests .	Yes	Require authorization for csv requests [Updated in Security Center 1.3]
glide.basicauth.required.excel	Require basic authorization for incoming Excel requests.	Yes	Require authorization for excel requests [Updated in Security Center 1.3]
glide.basicauth.required.importprocessor	Require basic authorization for	Yes	Require authorization for import requests

Property	Description	Default Value	Instance Security Hardening Settings
	incoming import requests.		[Updated in Security Center 1.3]
glide.basicauth.required.pdf	Require basic authorization for incoming PDF requests.	Yes	Require authorization for pdf requests [Updated in Security Center 1.3]
glide.basicauth.required.rss	Require basic authorization for incoming RSS requests.	Yes	Require authorization for RSS requests [Updated in Security Center 1.3]
glide.basicauth.required.scriptedprocessor	Require basic authorization for incoming script requests.	Yes	Require authorization for script requests [Updated in Security Center 1.3]
glide.basicauth.required.soap	Require basic authorization for incoming SOAP requests.	Yes	Require authorization for SOAP requests [Updated in Security Center 1.3, 1.5, and 2.0]
glide.basicauth.required.unl	Require basic authorization for incoming unload requests.	Yes	Require authorization for unload requests [Updated in Security Center 1.3]
glide.basicauth.required.xml	Require basic authorization for incoming XML requests.	Yes	Require authorization for XML requests [Updated in Security Center 1.3]
glide.basicauth.required.xsd	Require basic authorization for incoming XSD requests.	Yes	Require Authorization for XSD Requests [Updated in Security Center 1.3]
glide.cms.catalog_uri_relative	Enforce relative links from the URI parameter on /ess/catalog.do. If **Yes**, only relative URLs are permitted through the /ess/catalog.do page using the *`uri`* parameter. If **No**, all URLs are permitted, which may permit	Yes	Enforce relative links [Updated in Security Center 1.3 and 1.5]

Property	Description	Default Value	Instance Security Hardening Settings
	linking to external unauthorized content.		
glide.set_x_frame_options	Enable this property to set the X-Frame-Options response header to SAMEORIGIN for all UI pages. The X-Frame-Options HTTP response header can be used to indicate whether a browser should be allowed to render a page in a <frame> or <iframe>. Sites can use this property to avoid clickjacking attacks by ensuring that their content is not embedded into other sites. https://developer.mozilla.org/en/the_x-frame-options_response_header	Yes	Implement the x-frame-options: SAMEORIGIN security header [Updated in Security Center 1.3]
glide.ui.attachment.download_mime_types	A list of comma-separated attachment mime types that do not render inline in the browser. Prevents cross-site scripting attacks. For example, `text/html` forces HTML files to be downloaded to the client as attachments rather than viewed inline in the browser.	`text/html,image/svg,image/svg+xml`	Restrict downloadable MIME types [Updated in Security Center 1.3 and 2.0]
glide.security.groupby_acl_check	When this property is enabled, ACL checks for GroupBy operations are performed for the group names based on the actual data from the groups.	Yes	None
glide.security.diag_txns_admin	If **Yes**, only the admin user or	No	Restrict performance monitoring access

Property	Description	Default Value	Instance Security Hardening Settings
	user from allowed IP address can access `stats.do`, `threads.do`, and `replication.do`.		[Updated in Security Center 1.3]
glide.ui.security.codetag.allow_script	Allow embedded HTML (using [code] tags) to contain JavaScript tags. **Note:** This property is set to **true** by default in Vancouver and later releases, and can't be changed by administrators. For a use case where the property has to be changed, contact customer support.	No	Disable embedded HTML code [Updated in Security Center 1.3]
glide.script.allow.ajaxevaluate	Enable the AJAXEvaluate processor. The *`AJAXEvaluate`* API call allows the client to send and execute arbitrary scripts on the server.	No	Disable AJAXEvaluate

The following properties are defined in the sys_properties table, but are not visible on the High Security Settings page.

Property	Description	Default value	Instance Security Hardening Settings
com.glide.communications.httpclient.verify_hostname	Verify the hostname and certificate chain presented by remote SSL hosts. Protect against Man-In-The-Middle (MITM) attacks.	true	None

Property	Description	Default value	Instance Security Hardening Settings
	For more detail, see Set up Kubernetes spoke **Note:** This property overrides the com.glide.communications.trustmanager_trust_all property.		
glide.basicauth.required.schema	Require basic authentication for inbound table schema requests.	true	None
glide.security.csrf_previous.allow	Allow usage of an expired secure token to identify and validate incoming requests. This token is used to prevent cross-site request forgery attacks.	false	None
glide.security.csrf_previous.time_limit	Time in seconds for a secure token to expire. Allows control over the length of time that the previous CSRF token is valid. When the user session expires, the secure token expires with it unless the *glide.security.csrf_previous.allow* property is enabled and it is within the timeframe described by this property. This token is used to prevent cross-site request forgery attacks.	86400 **Note:** Value in seconds. Equivalent to 1 day.	None
glide.security.csrf.strict.validation.mode	Enforces strict validation on CSRF tokens so that users cannot resubmit a request if the CSRF token does not match.	false	Prevent Users From Accepting Warning To Bypass CSRF Validation [Updated in Security Center 1.3 and 1.5]

Property	Description	Default value	Instance Security Hardening Settings
com.glide.security.check_unsanitized_html	Enforces sanitization behavior of translated_html fields on a global level for field assignments.	enforce	None

Activating High Security Settings

The High Security Settings plugin is active by default on all new instances. If it is not active on your instance, you can request the plugin.

Before you begin

Role required: None

Before activating High Security Settings on an existing instance:

1. Review the following information to understand the new behavior:
   - Access Control List Rules
   - High Security Settings
   - Default deny property
2. Enable the plugin on a non-production instance. A recent clone of production is preferable.
3. Test the revised functionality, especially the added ACLs and default-deny functionality. Continue testing until the system performs as expected. If users cannot access expected resources, ensure they have appropriate roles and ACL rules to grant them the access.
4. Create update sets of any needed changes so you can apply them to production.

Note: To learn more about this plugin, see Enable High Security Plugin [Updated in Security Center 1.3] in Instance Security Hardening Settings.

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Select **Request plugin** to open the **Activate Plugin** form on Now Support.
3. On the **Activate Plugin** form, provide the following information.

Activate Plugin form

Field	Description
What is your target instance	Select the instance that you want to activate the plugin on.
Which plugin would you like to activate	Select the name of the plugin to activate.

Field	Description
	Note: If the system doesn't list the plugin you want or if you're activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	Select the date and time to activate the plugin.

Example

For example, see the following form to activate the Event Management plugin on an instance named SNC Instance.

4. Select **Submit**.
   After the maintenance window, the system installs the plugin on your instance. To confirm the installation, go to the Installed tab in the Application Manager.

Virtual Private Network (VPN)

Use a virtual private network (VPN) to integrate your instance with external data sources over the Internet.

Explore

Learn the features and business values of virtual private network.

Activate

Active the virtual private network.

Configure

Understand how to configure virtual private network.

Exploring Virtual Private Network (VPN)

Use a virtual private network (VPN) to integrate your instance with external data sources over the Internet.

When configuring an integration that uses an encrypted protocol, such as Lightweight Directory Access Protocol (LDAP) or HTTPS, it is good practice to use the Internet as a transport mechanism.

However, there may be security or network architecture requirements that dictate the use of a site-to-site Internet Protocol Security (IPSEC) Virtual Private Network (VPN) connection between the datacenters and your business networks. The VPN supports the necessary encrypted communication between the instance and your network.

⚠ Warning:

When a VPN tunnel is initiated, it operates as site-to-site connection. This means that the endpoint on your infrastructure receives an IP address, referred to as an encryption domain. This public IP can be accessed by any instance within the same data center.

For example, if you have an internal web service and establishes a VPN tunnel, your instance is able to reach the internal endpoint as well as all your other instances in the same data center.

VPN connections

The ServiceNow VPN infrastructure uses pairs of Cisco adaptive security appliance (ASA) devices that serve as VPN termination points.

The VPN between the instance and your network utilizes your existing networking hardware to support communications. It is not necessary to install a piece of hardware. Because each customer has a unique configuration, the instance has a flexible VPN solution. the instance has built tunnels to Checkpoint, Juniper, Nortel, and other IPSEC VPN-capable devices.

The VPN connections between the instance and your network are created to support the encrypted flow of traffic into your network. Frequently, integrations that use the VPN do not have encryption as part of the underlying protocol. For example, LDAP over the VPN versus LDAPS over the Internet and HTTP over the VPN versus HTTPS over the Internet.

The network does not allow any inbound-to-ServiceNow integration or end-user-to-ServiceNow traffic to traverse a VPN connection. This restricted communication includes end-user access to the platform, administration of the platform, web services integrations, and other integrations that are configured to use a MID Server. All such inbound communication to the instance must be performed over the Internet using HTTPS. This configuration provides an encrypted communication channel. The encryption channel, along with IP access control, meets the security requirements for this traffic flow.

This restriction applies to inbound traffic only. Responses to outbound requests that the instance initiates do traverse the VPN tunnel.

Addresses for VPN communication

To prevent conflict or overlap with internal ServiceNow networks or with another internal IP address schemes in your network, all tunneled traffic in the encryption domain must use non-RFC-1918 addresses on both sides of the tunnel.

ServiceNow provides a single IP address for the source of queries into your network. You must provide Network Address Translation (NAT), non-RFC-1918 addresses for each host that is integrating with your instance. These public addresses need to be owned by your organization. Third-party addresses cannot be used inside tunnels. Additionally, the encryption domain must not contain the IP address of the VPN peer.

Redundant tunnels

There are two ways to build redundancy for your tunnels:

- Using the same encryption domain behind both of your peers. This is the preferred method.
- Using a different encryption domain behind each peer.

With the first method, you need to provide the same NAT address behind each of your peers to create a connection path using that address to your server. The path to your server could be the same physical machine or a mirror which provides identical services. With this method, your instance would use the same IP address to connect to your servers regardless of whether your primary or secondary tunnel is active. If you have more than one server, follow this same scheme for your additional servers. This method provides the most transparency to your users and is recommended.

The second method requires configuration in your instance to provide the redundancy. When the tunnel is used for LDAP, for example, you could provide redundant LDAP servers in your instance. Note that this method requires the connection to the first configured LDAP server to timeout before the instance attempts to connect to the secondary server. Because of this additional time delay, this solution should only be implemented if the first option is unattainable. Also note that

not all services can be configured for redundantly in your instance. If you are using a VPN tunnel for something other than LDAP and redundancy is required, check that your configuration can support multiple addresses, or see the first option above.

Alternatives to using a VPN

These alternatives provide a simpler way to connect your instance to the resources in the ServiceNow data centers and provide better encryption. Additionally, you can avoid any issues that VPN downtime might cause, such as making your instance unavailable to users if there is an issue with the VPN tunnel.

Single sign-on and MID server

Consider using a combination of Single Sign-On (SSO) for authentication and the MID Server for user data synchronization, rather than using a VPN to connect your LDAP server to your instance. For integrations other than LDAP, consider using certificate-based encryption.

You can use the LDAP listener on a MID server to synchronize your user table in near real time.

The advantage of this approach is that there are no firewall holes, routes, VPN tunnels, or other special network settings to configure and maintain. The SSO/MID-Server solution is the most flexible, secure, and cost-effective method to achieve the complete LDAP integration.

LDAP over SSL

Another alternative to using a VPN tunnel is to configure LDAP Over SSL (LDAPS) directly over the Internet. You can configure a read-only domain controller and lock the instance down in your DMZ using only the instance's source addresses and the destination ports of your choice. Since the ports for LDAP are configurable in your instance, you can perform a port address translation (PAT) if desired. With LDAPS, you control the certificate that is uploaded over an encrypted channel to the instance, (see Uploading a certificate to an instance). The packets cannot be encrypted or decrypted without the certificate.

The advantage of this approach is that it provides a stronger encryption and decryption mechanism. A VPN can only encrypt and decrypt the traffic between the two peers sitting on the Internet with a coordinated pre-shared key, similar to a password. LDAPS provides a longer encrypted path, end-to-end, at the application layer and with a certificate that is far more complicated than a pre-shared key that the IPSec tunnel uses.

VPN setup

From the time that a VPN request is submitted, it typically takes one week or less to complete the VPN build. To support the redundancy requirements of your instance and your organization, a minimum of two and a maximum of four VPNs are provisioned (from the active site to your active site or the active site to your DR site, and so on).

It is good practice for the encryption domain to be as specific as possible. Ideally, the encryption domain would include only the specific hosts that are required for the integrations. A large encryption domain can create opportunities for routing discrepancies (VPN versus Internet).

To create the VPN, the instance does the following:

1. Provides the VPN peer and host addresses from each data center.
2. Builds the necessary VPN connectivity from two data centers into your network. To support redundancy and disaster recovery (DR) requirements, the VPNs can be provisioned from two data centers into two networks.

The instance does not support building multiple VPN tunnels into a customer network for the purpose of connecting to multiple geographic regions or subsidiaries. You should perform any inter-site routing, traffic distribution, or traffic shaping within your own internal network, rather than having multiple VPN tunnels.

Activating a VPN service

For all VPN requests, including provisioning, modifications, or general questions, use the Service Catalog VPN Request form.

Before you begin

Role required: admin

Procedure

1. Navigate to .
2. Select the **Automation Store** tab.
3. Use tree on the left to navigate to **All automations** > **Service catalog** > **Cloud Infrastructure**.
4. Select **VPN Requests**
5. Select the appropriate VPN request type.
6. Answer the questions.
   Questions vary depending on the request type selected.
7. Click **Submit**.

Result

Once your request is submitted, ServiceNow will work with your network engineer(s) to test and validate that the VPN is successfully passing traffic. To ensure that your questions are answered in a timely manner, please address VPN-related questions during this process.

Configuring an address for VPN communication

To prevent conflict or overlap with internal ServiceNow networks or with another customer's internal IP address schemes, the instance requires that all tunneled traffic in the encryption domain use non-RFC-1918 addresses on both sides of the tunnel.

Before you begin

Role required: admin

About this task

The instance provides a single IP address for the source of queries into your network.

Procedure

Provide Network Address Translation (NAT), non-RFC-1918 addresses for each host that is integrating with the instance.

Platform Privacy

Privacy enables you to mask sensitive data on your instance.

Store Applications

Data Privacy	Data Discovery
	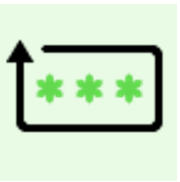
Use Data Privacy to classify sensitive data and to remove personally identifiable information (PII) from user data in a production instance and anonymize data in non-production instances.	Use Data Discovery to identify sensitive data within an instance to classify, protect, or report.

Plugins

Data Privacy (Classic)	Data Anonymization	Data Classification
		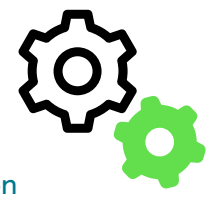
Use the legacy Data Privacy(Classic) plugin.	Data anonymization provides a way to easily transform data so that it is unidentifiable and more compliant with data privacy regulations.	Group data by type, using pre-defined or user-defined data classifications. If you have an assigned data classification administrator or auditor role, you can administer different data classes or visually analyze the current state of different types of data within the instance.

Exploring Data Privacy

Use Data Privacy to classify sensitive data and to remove personally identifiable information (PII) from user data in a production instance and anonymize data in non-production instances. Once anonymized, the user data is no longer considered regulated private information.

Developers must work with data on non-production instances to ensure that their implementations are working as expected. While importing data from your production instance is a useful way to simulate production, it presents a security risk. Administrators can use data privacy to provide developers with data that does not contain private information to work safely in a non-production environment.

Data classification

Identify and classify your sensitive data according to pre-defined criteria determined by the level of sensitivity of the data types in your instance. Data sensitivity levels help determine how each type of classified data should be handled. There are several pre-defined classes provided with base level data privacy. Use the classification section of Data Privacy to label and group data within your instance. Add classes, view data class structure and classify data. Group data by type, using pre-defined or user-defined data classifications.

User data anonymization

As an administrator, you define whether to anonymize all information for all users or for a subset of users. When anonymized, data for the selected user records is replaced with randomized values or values you define. When replacing values, the data structure can be preserved using various techniques.

Data Privacy

Use Data privacy to classify sensitive data and to remove personally identifiable information (PII) from user data in a production instance and anonymize data in non-production instances.

Note: You can also use the legacy plugin Data privacy (Classic).

Installation details

You must have the following applications installed on your instance:

- Data Privacy (Classic) [com.glide.data_privacy]
- Data Privacy [sn_dp_store_app]
- Data Discovery [sn_data_discovery]
- Data Discovery APIs[com.glide.data_discovery]

Here is the how the installation works:

- Installing the Data Privacy Store App will auto install the Data Discovery Store App, Data Privacy (Classic) plug-in, and the Data Classification plugin.
- Installing the Data Discovery Store App will auto install the Data Discovery APIs plugin

Considerations

- Only classified data can be anonymized. For information on data classes and classification, see Data classification (Classic) or Data classification Store App.
- PII in logs and other auditing data are not anonymized.
- Only structured data can be anonymized. Unstructured data, such as Journal fields, comments, attachments, and other fields where partial text may represent PII is not anonymized. See Supported field types for anonymization for more information.

- Integrations with single sign-on (SSO) systems may resynchronize user information from their source of truth systems. There is no mechanism in place to ensure the permanency of the de-identification of sys_user data. For information on user administration and sys_users see User Administration.

Data Privacy

Use Data Privacy to classify sensitive data and to remove personally identifiable information (PII) from user data in a production instance and anonymize data in non-production instances.

Explore Data Privacy

Learn about Data Privacy.

Configure Data Privacy

Learn how to configure Data Privacy.

Roles in Data Privacy

Get details about roles in Data Privacy.

Data Privacy Advanced Features

Learn about the advanced features of Data Privacy.

Data Privacy for Now Assist

Set up and configure how to discover and anonymize sensitive data from generative AI prompts.

Get started

Explore

Learn more about Data Privacy for Now Assist

Configure

Configure Data Privacy for Now Assist

Note: Data Privacy for Now Assist detects and masks sensitive data based on Regex patterns and does not support contextual(model-type) data patterns.

Important:

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

Exploring Data Privacy for Now Assist

Learn more how Data Privacy for Now Assist enhances your ability to protect sensitive data.

Data Privacy for Now Assist overview

Sensitive data such as age, phone number, and other personally identifiable information(PII) can be masked so that it does not get processed from generative AI prompts. Placeholder text and anonymized data are sent with the prompt instead, and these values are replaced with the original text after the large language learning module(LLM) response has been received. This two-way masking ensures that end users receive accurate responses, but sensitive data is not exposed to the LLM.

There are some considerations for configuring Data Privacy for Now Assist. See Configuring Data Privacy for Now Assist for configuration details.

Important:

Data Privacy for Now Assist detects and masks sensitive data based on Regex patterns and does not support contextual(model-type) data patterns

Configuring Data Privacy for Now Assist

Configure a data privacy advanced configuration to de-identify personally identifiable information (PII) in generative AI applications.

Before you begin

You must have the following applications installed on your instance:

- Data Privacy (Classic) [com.glide.data_privacy]
- Data Privacy [sn_dp_store_app]
- Data Discovery [sn_data_discovery]
- Data Discovery APIs[com.glide.data_discovery]

Installing the latest version of the Generative AI controller will auto install the Data Privacy Store App(sn_dp_store_app). The Data Privacy Store App will auto install the Data Discovery store app[sn_data_discovery], Data Privacy (Classic)(com.glide.data_privacy) plug-in, and the Data Discovery APIs[com.glide.data_discovery].

Role required: now_assist_data_privacy_admin

Important: You do not need a full active license to configure Data Privacy for Now Assist.

Procedure

1. Navigate to **All** > **Data Privacy (Classic)** > **Privacy Policy Advanced Configuration**.
   If you previously used the Sensitive Data Handler to help de-identify data for generative AI, you may already see a privacy policy configured. Your previously configured regular expressions have been migrated as part of your upgrade. If you already have a data policy for Now Assist, skip to step 6.
2. Select **New**.
3. Enter a name for the privacy policy.
4. In the **Data Channel** field, select the data channel to be used.

Channel	Description
Data Kit	Data, which AI models are using for evalua tion, is sanitized by discovering and de-identi fying sensitive data
Data Extraction	Data is sanitized before being sent for model training
Now Assist	Data is sanitized before being sent to GenAI Controller

5. Set **Active** to `true`, then select **Submit** to create the policy advanced configuration.
   Only one policy configuration for each data channel can be active at a time. To activate a new policy advanced configuration, you must set **Active** to `false` for all other policy configurations on that data channel.
6. After you're redirected to the list of policy advanced configurations, open the record you created.
   Open the existing record with the Now Assist data channel if one is already present.
7. To add a data pattern to de-identify, select **Select Data Patterns**.
8. **Optional:** To create your own data pattern, see Configure Data Discovery patterns.
9. Select your data patterns, then select **Save**.

Result

Data caught by the regular expressions selected in the data patterns is de-identified for generative AI applications. In this example policy, it is configured to catch a series of active data patterns and then use the data extraction data channel.

Data privacy

The data privacy store app is a Next Experience refresh for data classification and data privacy with modern look, feel, and usability. Data privacy store app is supported in Utah and above.

Data privacy is comprised of several components, Overview, Classification, and Anonymization.

Overview

Use the Overview section as a starting point to manage your data and data privacy compliance. See Data privacy overview for details.

Classification

Data classification is the process of organizing data into categories that make it easy to retrieve, sort, and store for future use. Using a classification system helps to focus on confidentiality and security policy requirements. Classify data to be used for anonymization. For information on data classes and classification, see Data classification.

Data privacy overview

The Overview homepage is a starting point to manage your data and data privacy compliance.

The Overview dashboard reports the current state of data classifications and anonymization jobs within your instance. You can view your data with the Overview dashboard to see how much data can be classified, the data classes available, the amount of data in each class, and an overall status of the anonymization jobs. The classified data is broken into different categories. Selecting a sub-category in either chart will add or remove the category from the overall chart and adjust the counts.

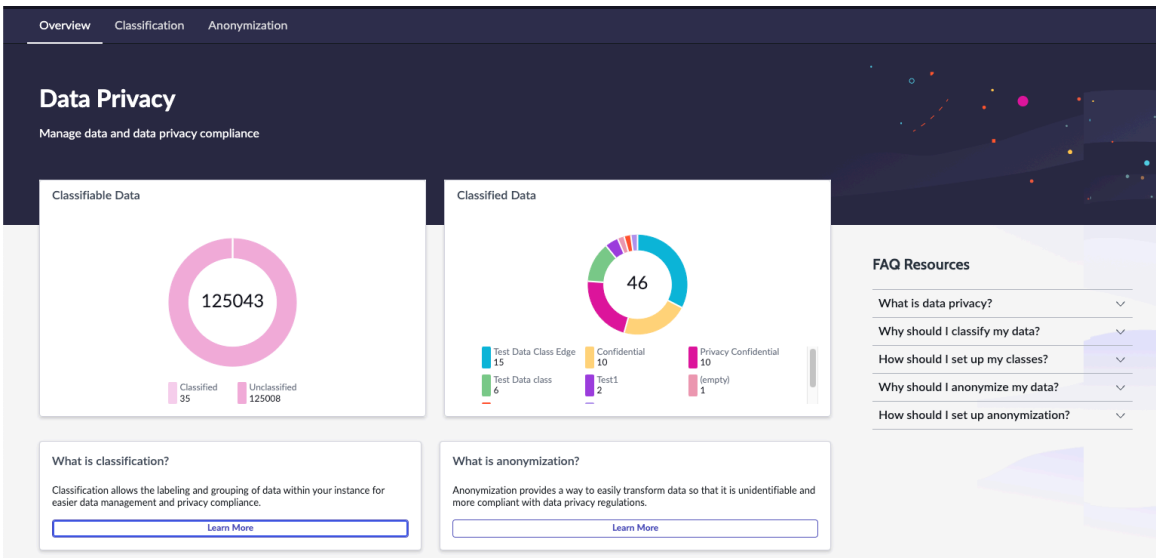

Classifiable Data

Displays the total number of data records in the instance that can be classified. The number is broken down into the total number of classified data and the total of unclassified data. This gives administrators a quick view of potential classification opportunities.

Classified Data

Displays the total number of data records in the instance that have been classified. The number is broken down into the total number of classified data for each assigned data class. This provides a quick understanding how much data has been classified in each area.

Anonymization Jobs

Displays data privacy jobs that are currently in process of jobs that have completed on your non-production instance for user and data class jobs.

Learn More

The **Learn More** sections provide a quick view into understanding classification and anonymization, as well as an easy way to get started in the process.

FAQ Resources

Access in product learning resources about classification or anonymization for additional information on how to get started securing your instance with data privacy.

Data classification

Group data by type, using pre-defined or user-defined data classifications. If you have an assigned data classification administrator or auditor role, you can administer different data classes or visually analyze the current state of different types of data within the instance.

Data Classification enables support for:

- Visibility into the types of data hosted on a ServiceNow AI Platform instances.
- Compliance with privacy laws, and meeting regulation requirements for industries such as financial services and medical device manufacturing.

Data classifications

Data classification is a standalone process in which you manually apply data classifications to existing dictionary entries in any table. See Data dictionary tables for additional information.

- You classify data as you find appropriate for your business and you can alter the available data classes as necessary.
- When you classify data, you can use the pre-defined data classifications, or create your own. Although use of pre-defined data classifications is optional, it is advisable do so as a starting point. These pre-defined data classifications are included in demo data that you can install in your instance.
- If you create your own data classifications, you can also design a tiered hierarchical system with parent and child data classifications.

Note: Data Classification supports domain separation, and the data_classification table itself is process separated. See Domain separation and Data Classification for additional information.

Use cases

General Data Protection Regulation (GDPR) is a European Union regulation whose purpose is to provide individuals with control over their own personal data. You can use data classifications, such as Personally Identifiable Information, to identify where personal data is being stored in your instance. By applying the appropriate security mechanisms to protect that personal data from leaking out, your organization satisfies GDPR requirements.

If you store customer information in the ServiceNow AI Platform, use the Personally Identifiable Information (PII) classification code where needed to track data subject to regulation by local privacy laws.

You can apply a Restricted data classification to Employee table columns that store sensitive employee information such as Social Security Numbers (SSN). Administrators and auditors can then use the Overview dashboard to confirm that you have assigned data classifications to the correct columns. They can also view the classification details for restricted types of information.

Create data classifications

Create your own user-defined data classifications in the [data_classification] table that you can then assign to specific columns in specific tables. Create new data classes to start the classification process.

Before you begin

Role required: data_classification_admin, admin

Procedure

1. Navigate to **All** > **System Security** > **Data privacy** > **Classification**.
2. Select **+Add data class**.

   Note: There are several data classes included with the base system.

3. Fill in the fields on the form.

Field	Description
Class Name	Name of the data classification.
Parent Class	Name of the parent data classification that this data classification is subordinate to. Leave the field empty if this data classification isn't a child to a parent data classification.
Description	Description of the data classification.

4. Click **Submit**.
   The new data class is added. If the data class is a child, it will be listed under the parent in the left-hand navigation bar.

Classify data

Group data by type, using pre-defined or user-defined data classifications. Assign data classifications to specific table columns in the Dictionary [sys_dictionary] table. When you assign data classifications, it creates entries in the Dictionary-Data Class [m2m_dictionary_dataclass] table, which you can then review in the Overview dashboard.

Before you begin

Role required: data_classification_admin, admin

Procedure

1. Select **New** to assign data from a table to be classified.
2. Select the data class from the drop-down.
3. Select the records reflecting the tables and columns to be classified.
   Choose additional rows to display per page in the table to simplify locating a particular table.
4. Select **Classify data**.

The data is classified under the Classification Name in the Classified Data table.

5. View or export the data to Excel, CSV, JSON, or PDF.
   Choose to download the data in the chosen format or receive it via email.

Data anonymization

Anonymization provides a way to easily transform data so that it is unidentifiable and more compliant with data privacy regulations.

Administrators can perform the anonymization process in both production and non-production instances. Data privacy can be used in production instances to de-identify users. Anonymization of data classes should only be used with non-production instances. This preservation ensures that data like email addresses or physical addresses are replaced with similarly formatted, but anonymized versions.

Administrators can also use anonymization as part of their General Data Protection Regulation (GDPR) Right to be forgotten (RTBF) processes to anonymize user information. See https://gdpr-info.eu/art-17-gdpr/ for more information.

Use the anonymization section of data privacy to create and view privacy policies and techniques,and to perform privacy bulk assignment. View all of the jobs, with the description, private policy used, and status. In order to access the anonymization section, an admin must first elevate to the data_privacy_admin and data_privacy_processor roles.

Anonymization techniques

Anonymization techniques are options you select to determine how your data is anonymized. You must create an anonymization technique to reference in the anonymization job. See Create anonymization techniques to associate a privacy technique to an associated **Anonymization technique configuration**.

Anonymization policies

Configure an anonymization policy to specify which data privacy techniques are used when anonymizing your data. See for Create anonymization policies details.

Anonymization jobs

Anonymization jobs use all of these components to anonymize your data. For more information on these jobs, see Create anonymization job.

Create anonymization techniques

Create a data privacy technique configuration to customize how data privacy anonymizes your data.

Before you begin

Role required: data_privacy_admin and admin

Procedure

1. Elevate to the **data_privacy_admin** role.
   For details on role elevation, see Elevate to a privileged role.
2. Navigate to **System Security** > **Data Privacy** > **Anonymization**.
3. Select **View techniques**.
   There are several pre-defined techniques available for selection.

Technique	Description
Selective Replace	This technique does a selective replace of String data. All characters between the input's start and end indices are replaced with the character you choose. You can specify characters to exclude from masking: ○ start_index: Technique masks data starting at the specified character. If left blank, masking begins with the first character. ○ end_index: Technique masks data from the start of the string to the specified character. If left blank, masking ends with the last character. ○ exclude_char: Define a character to exclude from masking. ○ replacement_char: Define a character used for masking. If none is provided, asterisks(*) are used by default.
Static Replace	This technique swaps values with static values. String, Number, and Date data can use this technique: ○ date_time_value: Replace Date values with this date. Use the yyyy-MM-dd HH:mm:ss format. ○ date_value: Replace Date values with this date. Use the yyyy-MM-dd format. ○ number_value: Replace Number values with this number. ○ string_value: Replace String values with this text. ○ number_type: Accepts an integer number
Random Replace	This technique swaps values with randomly generated values. String and Number data can use this technique.
Remove	This technique removes values, replacing them with empty (null) values.
No Action	This technique is a placeholder. It does not modify fields when selected.
Selective Replace with X	Transforms String data and selectively replaces sensitive characters with the letter X. **Note:** Default technique for data patterns in Exploring Data Discovery (Classic).
Data pattern anonymization	Only anonymizes discovered data patterns within unstructured data fields while keeping underlying context intact. **Note:** Settings for this technique reference data pattern anonymization technique settings in Exploring Data Discovery (Classic).

4. Select **Add custom technique**, if not using a pre-defined technique.
5. Fill in the fields in the **Customize technique** form.

Field1	Description
Base technique	Select a pre-defined technique, as custom techniques are based on the pre-defined techniques.

Field1	Description
Technique name	Enter a name for the technique.
Technique description	Enter a description for the technique.

6. Select **Next**.
7. Enter the technique parameters.
   The available parameterized values depend on which privacy technique you have selected. There are no parameterized values for the **No Action** and **Remove** techniques.

Privacy parameterized values for the Base techniques

Base Technique	Privacy Technique Parameter value	Description	Default value
Selective Replace	end_index	Technique masks data from the start of the string to the specified character. If left empty, masking ends with the last character.	(Empty)
Selective Replace	exclude_char	Character to skip masking. Only a single character can be used in this value. If more than one is entered, the first character is used.	(Empty)
Selective Replace	replacement_char	Character to use when replacing values using a selective replace.	An asterisk (*) is used if no other value is entered.
Selective Replace	start_index	Technique masks data starting at the specified character.	If left blank, masking starts at the first character.
Static Replace	date_time_value	Replace date and time values with this date. Use the yyyy-MM-dd HH:mm:ss format.	1988-11-11 10:10:10
Static Replace	date_value	Replace date values with this date. Use the yyyy-MM-dd format.	1988-11-11
Static Replace	number_value	Replace Number values with this number.	1234567
Static Replace	string_value	Replace String values with this text.	TEXT123

Base Technique	Privacy Technique Parameter value	Description	Default value
Random Replace	preserve_data_length	Set to **true** to preserve data length. De-identified data will have the same length as the original data.	True

8. Select **Create Custom Technique**.
   Your custom technique is added to the Anonymization techniques.

What to do next
See Create anonymization policies to configure an anonymization policy to specify which techniques are used when anonymizing your data.

Create anonymization policies

Configure an anonymization policy to specify which techniques are used when anonymizing your data.

Before you begin
The data privacy configuration defines tables, sys_user and other, and columns to the de-identified, depending on the use case and specifies parameterized types of the techniques to be used while de-identifying data.

Note: To complete a privacy configuration, you must first configure a data privacy technique configuration. See Create anonymization techniquesfor more information.

Role required: data_privacy_admin and admin

Procedure
1. Elevate to the **data_privacy_admin** role.
   For details on role elevation, see Elevate to a privileged role.
2. Navigate to **System Security** > **Data Privacy** > **Anonymization**.
   All anonymization policies display. Published policies are available to schedule the anonymization job.
3. Select **Create new policy**.
4. Select to either anonymize **Data tables or columns**, **User specific data**, or **Real time data**.

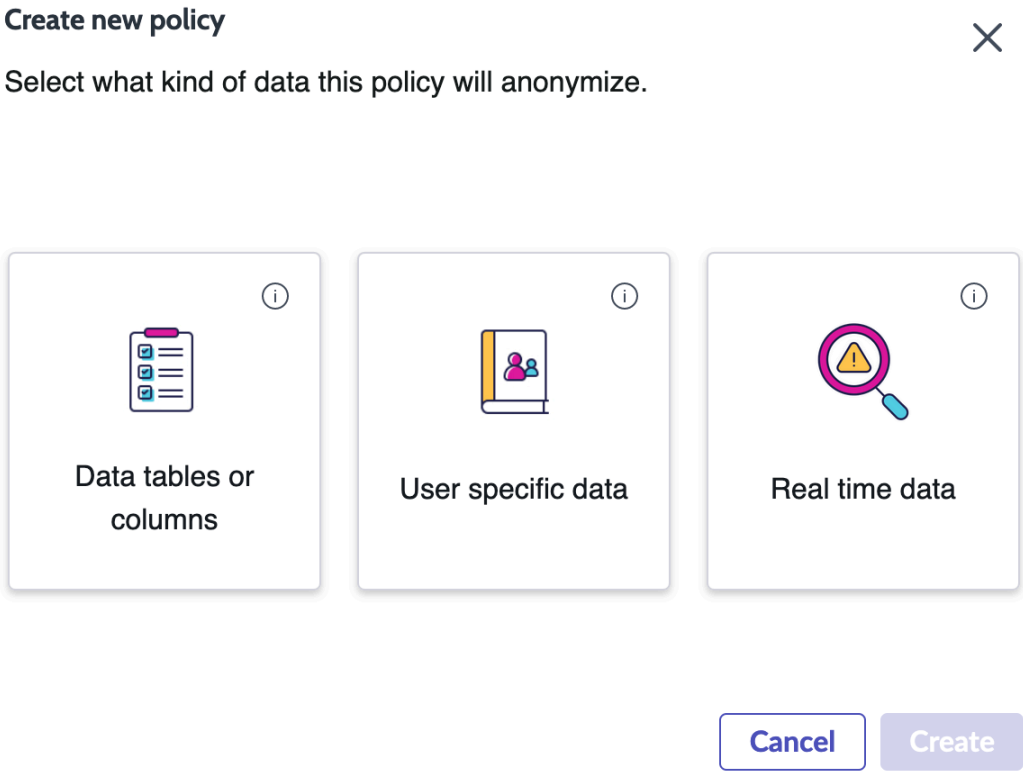

Data Type	Description
Data tables or columns	Records that match the data policy will be anonymized.
User specific data	Select a set of users or user groups to be anonymized.
Real time data	Anonymize real time entries for a set of columns.

Data privacy policies can only apply to classified data, for more information on data classification, see Data classification.

5. Select **Create**.
   There are sequential steps required to complete the policy, **Define details**and **Assign techniques**. **Select user reference** is also required when defining the policy for user specific data.
6. Define the details for the new anonymization policy.
   - Enter the policy name in the **Name** field, and the policy description in the **Description** field.
   - Define what channels automatically activate the policy and the channel priority in **Activation Channels**
   - In the **Data Class** field, select the data class to use with this policy.
   - Turn on or off real time data anonymization. See Step 8 if real time data anonymization is on

> **Note:** If you are not anonymizing an entry, select the **DoNothing** technique rather than leaving the entry empty. Policies with empty values in the Privacy Technique Configuration field cannot execute when used in data privacy jobs.

After selecting a data class, the Assign techniques form displays for each record returned for the defined data class.

7. Assign anonymization techniques for the selected data class.

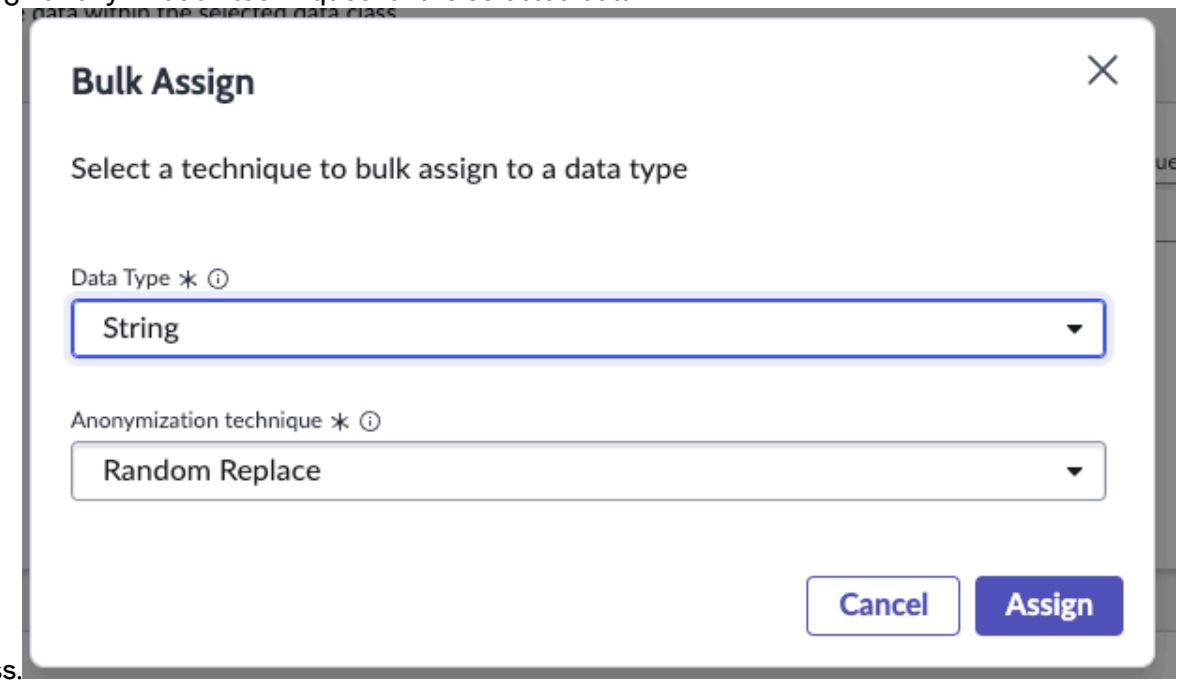

Option	Description
Select Bulk Assign Techniques	Applies anonymization to all data records in the chosen data class. Select the data type and the anonymization technique to apply to all entries with the selected data type. Repeat this step for additional bulk assignments of different data types. See Supported field types for anonymization for a list of data types.
Select an anonymization technique for each data column record	Your data privacy processor users can choose which records to anonymize when creating data privacy jobs. Individually apply anonymization to each data record in the cho sen data class.

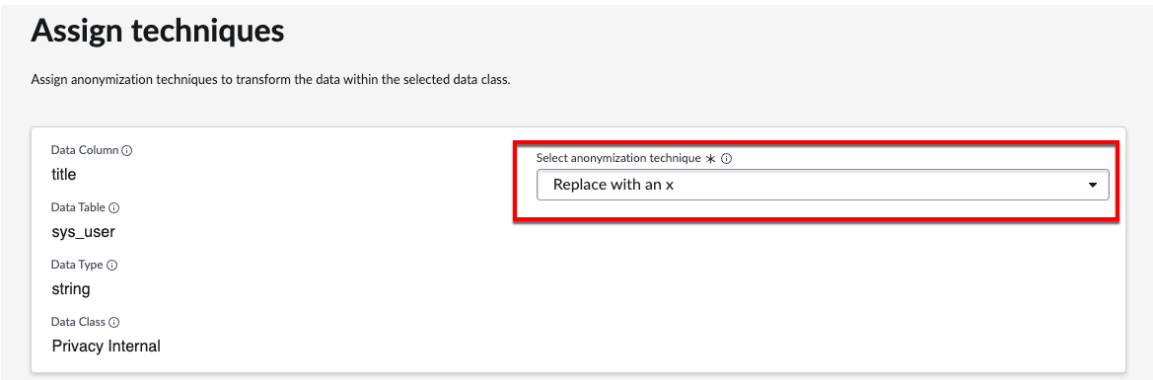

8. **Optional:** Enter child tables to be scanned.

Child tables of the parent will be anonymized, if a table has no children this option will not be available.

Warning: A parent job will fail if a child job fails.

9. **Optional:** If Data Pattern Anonymization is selected, select the anonymization technique to be used.
10. **Optional:** Set the ordering for data patterns.
11. **Optional:**

    Tip: Use the **Test** feature to test sample inputs. You can review metrics from the result like scan time, result, and discovered patterns.

    Select the **Test** button to test the policy.

12. **Important:** All tables must have a correct sys_dictionary entry.

    Select **Save**.
13. Select **Publish** to update the anonymization policy for scheduling and be returned to Anonymization policies.

    Note: Only published policies can be used for anonymization job scheduling.

What to do next

Create anonymization job.

Configure data anonymization clone request

Data privacy clone integration is configured using a PostClone script to create and execute data privacy jobs for configured policies on the target. After running the script, users will see de-identified data and will not have access to the original data.

Before you begin

The data privacy PostClone script is installed with the activation of the data privacy plugin (sn_dp_store_app). See Activate data privacy for details.

Role required: data_privacy_clone_processor, data_privacy_admin, and admin

Procedure

1. Activate the data privacy plugin (sn_dp_store_app) on the source instance.
   The data privacy PostClone script is installed.
2. Elevate to the **data_privacy_admin** role.
   For details on role elevation, see Elevate to a privileged role.
3. Navigate to **System Security** > **Data Privacy** > **Anonymization**.
4. Select **Create new policy**.
   See Create anonymization policies for details.
5. Select **Data tables or columns**.
6. Select **Create**.
7. Enter a name and select a data class.
8. Select to **Activate the policy during cloning**.
9. Select the policy order to run if there are multiple clone policies.

A data privacy job for Postclone configuration with a higher Application Order might start before another job of lower order, if the job with the higher order does not involve any table related to other lower order job.

10. Select **Continue**.
11. Complete the policy configuration and publish the policy.
12. Back up data privacy configurations.
13. Schedule the anonymization job.
    See Create anonymization jobfor details.
14. As the data privacy admin, submit a clone request.

Result

The data privacy PostClone script executes on the target Instance and The PostClone script creates a data privacy federated job record on the target instance. The federated job creates and executes a data privacy job for each post-clone policy, in Application Order, on the target instance.The backup source is cloned to the target Instance. Data privacy PostClone script creates and executes data privacy jobs for configured policies on the target instance.

The elevated data privacy clone processor can log on to the target instance and monitor the post-clone federated job state on the dp_federated_job.list and dp_job.list.

Create anonymization job

Configure a data privacy job on your production instance to use anonymized data on your non-production instance for user and data class jobs.

Before you begin

The data privacy job supports two de-identification use cases:

- Sensitive data of specific sys_users
- Sensitive data of particular data class.

Role required: data_privacy_processor and admin

Procedure

1. Elevate to the **data_privacy_processor** role.
   For details on role elevation, see Elevate to a privileged role.
2. Navigate to **System Security** > **Data Privacy** > **Anonymization**.
3. In an Anonymization policy, select **Schedule job** for the policy to be used in the job.
   A policy must be in a published state in order to schedule an anonymization job.

> ⚠ **Warning:** Anonymization jobs are highly destructive and can only be reversed on rollback. Double check all information, such as records and table processed, before scheduling a job.

4. In the form, fill in the fields.

Data privacy job fields

Field	Description
Policy used	Read only name of the selected privacy policy configuration to use for this job. Edit the policy to view additional information about the policy.

Field	Description
	For detail on privacy policy configurations, see Create anonymization policies.
Job description	Description of the job.
Start time	The start of the time to run this job in **HH:MM:SS**.
End time	The end of the time window to run this job in **HH:MM:SS**. The end time must be after the start time. The job runs before the time entered in this field. If the job has not yet completed, the job will pause and resume at the next time window start.
Dry Run	Run the job as a test. No records are affected when running this job. Results are displayed in the **Jobs** list, as though the job had executed. **Note:** **Dry Run** must be turned off when configuring a data privacy job with rollback. See Data privacy job rollback for details.
Type of user selection	Select either users or groups to be anonymized.
Select users/groups	Select a the specific set of users or groups to be anonymized in this job. Supports up to a maximum of 1000 users. **Note:** This required field displays only when the selected privacy policy condition requires a selection of user records.

5. **Important:** All tables must have a correct sys_dictionary entry before scheduling and during its job.

Select **Schedule job** on the form to place the anonymization in the job queue.

The job runs between the times selected in the **Start time** and **End time** fields. If the job has not completed during the start and end time window, the job will continue at the next time window start.

A job can only be executed once, even if **Dry Run** is selected. To run a job again based on the same policy, select **Schedule job** and complete the form using the same field values.

Warning: Anonymization jobs on encrypted columns will you to decrypt and re-encrypt any encrypted columns targeted by the job. To prevent this, select **No Action** as the policy technique.

The job is listed in the **Jobs**

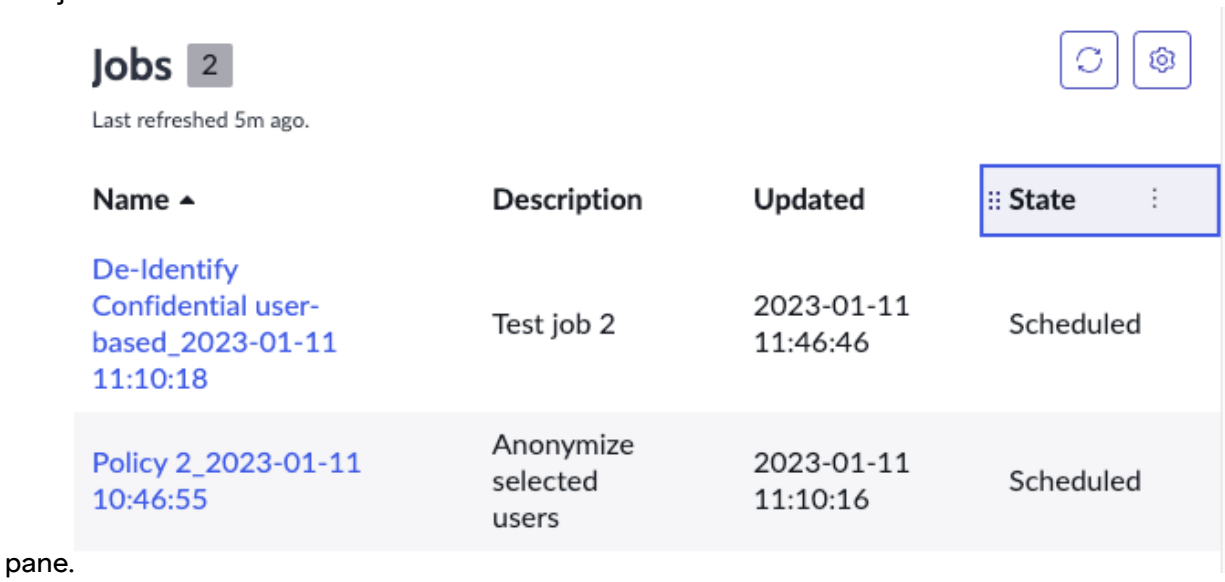

pane.

Field	Description
Name	Name of the anonymization job.
Description	Description of the anonymization job.
Updated	The date and time that the job was last updated.
State	State of the data privacy job: ○ **Scheduled**: Default state for new jobs. ○ **Completed**: The job has successfully anonymized the selected data. ○ **Cancelled**: The job was manually cancelled. ○ **Error**: There was a problem saving the job. Re-schedule the job or create a new job. There may be issues with the configuration if the error continues. ○ **Rollback in progress**: The job has been set to roll back anonymization. ○ **Rollback complete**: The anonymization job rollback has successfully completed. A read-only field.

6. Select a job from the **Jobs** pane to open the Job summary.
After a job is scheduled, the **Cancel Job** and **Pause** buttons appear in the Job summary.

Additional fields for scheduled jobs

Field	Description
Estimated record count	The estimated number of records this dry-run job affects before being executed. A read-only field
Total data records processed	The total number of individual data records affected by this job. A read-only field.

Field	Description
Total data tables processed	The total number of data tables processed by this job. A read-only field.
Time remaining to rollback	The remaining time that a completed data anonymization job can be rolled back and de-anonymize the data. A read-only field.
Total users processed	The total number of individual user records this job affects. A read-only field.
Cancel job	Select to cancel the anonymization job. This must be selected prior to the job start time. When selected, the job status updates to *Cancelled*.
Pause	Select to pause the job and rollback recording, if rollback was selected. A warning message will display after a three-day expiry period for rollback contexts. This must be selected after the job start time and before the job end time. When selected, the job status updates to *Paused*.
Resume	Restarts a paused job. Rollback is not supported for resumed jobs if paused. Cancel the job and create a data privacy job. The recording uses an unexpired rollback context. When selected, the job status updates to *Scheduled*.
Export	Downloads a .PDF file of the data privacy job details.

Activate parallel jobs for data anonymization

Use parallel jobs to reduce your anonymization job execution time.

Before you begin

Role required: admin

Procedure

By default data class and user based anonymization jobs run using a single thread. Multiple threads may be activated for data class and user based jobs. In addition, federated jobs used during cloning uses parallel jobs with 3 workers by default.

1. In the navigation filter enter `sys_properties.list`.
2. Create the property `com.glide.data_privacy.max_parallel_workers`
3. **Optional:** To adjust the number of federated job parallel workers, edit the `dp.max_concurrent_clone_item_workers` property.
4. Verify the **Type** is `integer`.
5. **Important:** It is recommended to start with a small number of workers per node(2 or 3).

   Enter the number of parallel job workers in the **Value** field up to 5.

Real time anonymization

Use the real time anonymization(RTA) policy to anonymize data entries in real time.

Real time anonymization Overview

Users create an RTA policy by selecting real time anonymization in the Anonymization Policies page, and then selecting the appropriate data channel. For example you can use Virtual Agent with real time anonymization: create an anonymization policy with the Virtual Agent selected as its **Data Channel**.

Columns from the target tables may be selected for RTA, whereupon active data patterns are used and their policies applied to any valid record entries to the columns targeted for RTA. If an entry matches an active data pattern its associated anonymization technique will be used for anonymization.

Tip: If you need to change the anonymization technique see Configure Data Discovery patterns.

Real time anonymization failures

If an RTA policy fails, you can review its status with the Real time anonymization failures table.

Real time anonymization failures

Review real time anonymization(RTA) failures.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Security** > **Data Privacy(Classic)**.
2. Review the RTA failures in the table.

Field title	Description
Table Column	The table column the failure occurred in.
Failure Reason	Category of the failures cause.
Failure Reason Details	Short description summarizing the cause of the failure.
Privacy Configuration	The RTA privacy configuration that failed.
Record	The specific record that failed.
Table Name	Table the failure occurred in.
Timestamp	Time of the failure.

Activate data privacy

Data Privacy includes data classification and anonymization and is installed from the ServiceNow Store.

Before you begin

To use data anonymization, Data Privacy (Classic) must first be activated with the ServiceNow Vault entitlement. See Activate data privacy (Classic) for additional information.

Note: Installing the Data Privacy Store App will auto install the Data Discovery Store App, Data Privacy (Classic) plug-in, and the Data Classification plugin.

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All** > **Data Privacy**.
2. Find the application using the filter criteria and search bar.

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

   Visit the ServiceNow Store website to view all the available apps and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the ServiceNow Store version history release notes.
3. Select a version from the list and select **Install**.

   In the Install dialog that is displayed, any dependencies that are installed along with your application are listed.
4. If you are prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
5. **Optional:** If demo data is available and you want to install it, select the **Load demo data** check box.

   (Optional) Demo data comprises sample records that describe application features for common use cases. Load demo data when you first install the application on a development or test instance.

   Important: If you don't load the demo data during installation, it's unavailable to load later.
6. Select **Install**.

Data Privacy for Virtual Agent

You can use Data Privacy to detect and mask the sensitive data and PII during a Virtual Agent conversation.

Before you begin

Role required: virtual_agent_data_privacy_admin

About this task

Use Data Privacy to detect and mask the sensitive data and PII during a Virtual Agent conversation. Data Privacy is replacing the depreciated Sensitive Data Handler, see Deprecation Process [KB0867184] for more information. The setup process may vary depending on if Sensitive Data Handler was previously installed and configured.

Note: If you have previously configured Sensitive Data Handler your settings will be migrated to Data Privacy as a policy. See Create anonymization policies for more information.

Procedure

1. Navigate to **All** > **Conversational Interfaces** > **Settings**.
2. Under **Sensitive data detection**, select **View all**.
   If the button reads **Get data privacy**, you will need to install the Data Privacy plugin to continue. See Activate data privacy (Classic) for more information
3. Set the slider to **Active**.
4. **Optional:** If Data Privacy is not yet installed you will be prompted to install the Data Privacy plugin to continue, see Activate data privacy (Classic) for more information.
5. Select the appropriate conversation flow:
   a. Requester to agent - detect and mask sensitive data entered by the requester in an Agent Chat conversation.
   b. Agent to requester - detect and mask sensitive data entered by the agent in an Agent Chat conversation.
   c. Requester to Virtual Agent - detect and mask sensitive data entered by the requester during a Virtual Agent conversation.

   Note: You must select at least one conversation flow.

6. Select **Manage in Data Privacy**.
   If you have configured Sensitive Data Handler before, your settings will be migrated as a policy to Data Privacy
7. Create a new Data Privacy policy for Virtual Agent, see Create anonymization policies for more information on Data Privacy policies.

Domain separation and data privacy

Domain separation is unsupported for data privacy. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: No support

- The domain field may exist on data tables but there is no business logic to manage the data.
- This level is not considered domain-separated.

For more information on support levels, see Application support for domain separation.

Related topics

Domain separation for service providers

Supported field types for anonymization

Check which field types are supported when anonymizing data.

Note: Not all field types that have been classified are available for anonymization.

Some high risk field types are turned off by default, as detailed in the table. For more information about fields, see Field Types.

Supported field types for anonymization

Field type	Available by default
audio	No
condition	No
condition_string	No
currency	Yes
decimal	Yes
due_date	Yes
float	Yes
glide_date	Yes
glide_date_time	Yes
glide_duration	No
glide_time	Yes
html	No
icon	No
integer	Yes
ip_addr	No
ip_address	No
journal	Yes
journal_input	Yes
journal_list	Yes
longint	Yes
name_values	No
percent_complete	No
phone_number_e164	Yes
price	Yes
string	Yes
string_full_utf8	Yes
translated_html	No
translated_text	No
url	No
user_image	No
video	No
wiki_text	No

Data privacy roles

Data privacy adds these roles.

Data privacy administrator

The data privacy administrator roles is an admin role used to create data privacy techniques and policies.

Data privacy administrator [data_privacy_admin]

Contains roles

None

Assigned to groups

None

Subscription

No

Elevated

Yes

Considerations

Avoid assigning this role to your users when more targeted roles are available.

In order to assign "Data privacy administrator" role elevate the user to security admin role in order to add the role.

Data privacy auditor

Data privacy auditor is a read-only role used to view data privacy records.

Data privacy auditor [data_privacy_auditor]

Contains roles

None

Assigned to groups

None

Subscription

No

Elevated

No

Considerations

None

Data privacy clone processor

Users with the Data privacy clone processor role can create and execute data-class data privacy jobs.

Data privacy clone processor [data_privacy_clone_processor]

Contains roles

None

Assigned to groups

None

Subscription

No

Elevated

Yes

Considerations

None

Data privacy processor

Users with the Data privacy processor role create and execute data privacy jobs on the user [sys_user] table.

Data privacy processor [data_privacy_processor]

Contains roles

None

Assigned to groups

None

Subscription

No

Elevated

Yes

Considerations

None

Data privacy (Classic)

Data Privacy (Classic)is a legacy application that provides data classification, techniques, and jobs to anonymize PII.

Data Privacy Classic

Note: This section is for Data Privacy(Classic). See Data privacy for the latest up to date store version.

Activate data privacy (Classic)

You can activate the data privacy plugin (com.glide.data_privacy) for Platform Security if you have the admin role. The application includes demo data and installs related ServiceNow® Store applications and plugins if they aren't already installed.

Before you begin

Important: This section is for Data Privacy(Classic). See Data privacy for the latest up to date store version.

Data privacy requires a separate subscription from the rest of the ServiceNow AI Platform®.

To purchase a subscription, contact your ServiceNow account manager. When you purchase a subscription, certain plugins are activated automatically. If a paid plugin isn't activated automatically, you can manually activate it from the All Applications list in your instance.

Note: Before purchasing a subscription, you can evaluate Data Privacy(Classic) and Data Discovery with a 30 day trial period. After your trial expires you will no longer be able to discover data or run anonymization jobs in without a license.

Role required: admin

About this task

The following items are installed with data privacy:

- Plugins
- Roles
- Tables

For more information, see Installed with data privacy (Classic).

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Data Privacy plugin (com.glide.data_privacy) using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Installed with data privacy (Classic)

Learn about the components installed with the data privacy plugin (com.glide.data_privacy).

Tables installed

Table	Description
Data Privacy Federated Job [dp_federated_job]	Federated jobs for data privacy
Data Privacy Job [dp_job]	Data privacy jobs
Data Privacy Job	Data privacy job summaries

Table	Description
[dp_job_summary]	
Data Privacy Technique [dp_technique]	Data privacy techniques
Primary Reference Link [dp_primary_reference]	Primary reference links
Privacy Classified Field Technique [dp_field_technique]	Data privacy classified field techniques
Privacy Configuration [dp_configuration]	Data privacy configurations
Privacy Technique Configuration [dp_technique_with_params]	Data privacy techniques
Privacy Technique Parameter [dp_technique_with_parameter]	Data privacy techniques with parameters
Privacy Technique Parameter Value [dp_technique_with_parameter_value]	Parameter values used in Data privacy techniques with parameters.

Data privacy (Classic) configuration

Learn how to create data privacy techniques and policies, and how to create and execute data privacy jobs.

Important: This section is for Data Privacy(Classic). See Data privacy for the latest up to date store version.

Data privacy techniques

Data privacy techniques are options you select to determine how your data is anonymized. You must create a data privacy technique to reference in the data privacy job. See Create a data privacy technique configuration to associate a privacy technique to an associated **Privacy technique configuration**.

Data privacy policies

Configure a data privacy policy to specify which data privacy techniques are used when anonymizing your data. See Create a data privacy policy for details.

Data privacy jobs

Data privacy jobs use all of these components to anonymize your data. For more information on these jobs, see Configure a data privacy job.

Create a data privacy technique configuration

Create a data privacy technique configuration to customize how data privacy anonymizes your data.

Before you begin
Role required: data_privacy_admin and admin

Procedure

1. Elevate to the **data_privacy_admin** role.
   For details on role elevation, see Elevate to a privileged role.
2. Navigate to **System Security** > **Data Privacy** > **Privacy Technique Configuration**.
3. Click **New**.
4. In the **Name** field, enter a name for your privacy technique configuration.
5. In the **Privacy Technique** field, select a privacy technique.

Data privacy techniques

Technique	Description
No Action	This technique is a placeholder. It does not modify fields when selected.
Random Replace	This technique swaps values with randomly generated values. String and Number data can use this technique.
Selective Replace	This technique does a selective replace of String data. All characters between the input's start and end indices are replaced with the character you choose. You can specify characters to exclude from masking: ○ start_index: Technique masks data starting at the specified character. If left blank, masking begins with the first character. ○ end_index: Technique masks data from the start of the string to the specified character. If left blank, masking ends with the last character. ○ exclude_char: Define a character to exclude from masking. ○ replacement_char: Define a character used for masking. If none is provided, asterisks(*) are used by default.
Static Replace	This technique swaps values with static values. String, Number, and Date data can use this technique:

Technique	Description
	○ date_time_value: Replace Date values with this date. Use the yyyy-MM-dd HH:mm:ss format. ○ date_value: Replace Date values with this date. Use the yyyy-MM-dd format. ○ number_value: Replace Number values with this number. ○ string_value: Replace String values with this text.
Remove	This technique removes values, replacing them with empty (null) values.

Note: Previously supported value **Replace** is deprecated and renamed **Replace-Deprecated** and should not be used.

6. Right-click the header and click **Save** in the context menu.
   After the record is saved, the **Privacy Parameterized Values** list appears.
7. Use the records in the **Privacy Parameterized Values** list to customize your data privacy technique configuration.
   The available parameterized values depend on which privacy technique you have selected. There are no parameterized values for the **No Action** and **Remove** techniques.

Privacy parameterized values for Selective Replace

Privacy Technique Parameter value	Description	Default value
char_to_replace	Character to use when replacing values using a selective replace.	*
end_index	Technique masks data from the start of the string to the specified character. If left empty, masking ends with the last character.	(Empty)
exclude_char	Character to skip masking. Only a single character can be used in this value. If more than one is entered, the first character is used.	(Empty)
start_index	Technique masks data starting at the specified character.	1

Privacy parameterized values for Replace

Privacy Technique Parameter value	Description	Default value
date_time_value	Replace date and time values with this date. Use the yyyy-MM-dd HH:mm:ss format.	1988-11-11 10:10:10
date_value	Replace date values with this date. Use the yyyy-MM-dd format.	1988-11-11
number_value	Replace Number values with this number.	1234567
preserve_data_length	Set to **true** to preserve data length. De-identified data will have the same length as the original data.	true
string_value	Replace String values with this text.	TEXT123
use_random_generated_value	Set to **true** to replace data with randomly generated values. Only String and Number data can be replaced with random values. This option overrides static values.	false

8. Click **Save**.

Create a data privacy policy

Configure a data privacy policy to specify which data privacy techniques are used when anonymizing your data.

Before you begin

The data privacy configuration defines tables, sys_user and other, and columns to the de-identified, depending on the use case and specifies parameterized types of the techniques to be used while de-identifying data.

Note: To complete a privacy configuration, you must first configure a data privacy technique configuration. See Create a data privacy technique configuration for more information.

Role required: data_privacy_admin and admin

Procedure

1. Elevate to the **data_privacy_admin** role.
   For details on role elevation, see Elevate to a privileged role.
2. Navigate to **System Security** > **Data Privacy (Classic)** > **Privacy Policy Configuration**.
3. Select **New**.
4. In the **Name** field, enter a name for your privacy policy configuration.
5. In the **Data Class** field, select the data class to use with this policy.

Data privacy policies can only apply to classified data, for more information on data classification, see Data classification.
After selecting a data class, the **Privacy Classified Field Techniques** and **Privacy Primary Reference Links** lists display on the form.

6. **Optional:** Select **Apply to All Data in Class** to apply anonymization to all data in the chosen data class.
   If you don't select this field, your data privacy processor users can choose which users to anonymize when creating data privacy jobs. If you select this field, that option is not available.
   - **Apply when Cloning**: This option becomes available. When selected, the privacy configuration executes during data privacy clone.
   - **Application Order**: A data privacy job for Postclone configuration with a higher Application Order might start before another job with lower order.

   > **Important:** Avoid creating multiple data privacy policies with the same Application Order, as the resultant processing order for those with the same order will be inconsistent.

7. **Optional:** Select **Supports Rollback** to enable the ability to de-anonymize the data from a data privacy job.
   See Roll back a data privacy job for more information.
   After selecting **Supports Rollback** when creating a data privacy job, the option to roll back the job becomes available.
8. Select the **Privacy Classified Field Techniques** tab to display the **Privacy Classified Field Techniques** list.
9. Select an entry in the **Table** field to open the **Privacy Technique Configuration** field for each list entry.
   The **Privacy Classified Field Techniques** list displays all the data to be anonymized in your selected data class. For each of these entries, you must select a privacy technique to apply.
10. Select a **Privacy Technique Configuration** to apply.

    > **Important:** If you aren't anonymizing an entry, select the **DoNothing** technique rather than leaving the entry empty. Policies with empty values in the **Privacy Technique Configuration** field can't execute when used in data privacy jobs.

11. Select **Submit** or **Save** to save the record.

What to do next
Configure a data privacy job.

Configure a data privacy job

Configure a data privacy job on your production instance to use anonymized data on your non-production instance for user and data class jobs.

Before you begin
The data privacy job supports two de-identification use cases:

- Sensitive data of specific sys_users
- Sensitive data of particular data class.

Role required: data_privacy_processor and admin

Procedure

1. Elevate to the **data_privacy_processor** role.
   For details on role elevation, see Elevate to a privileged role.
2. Navigate to **System Security** > **Data Privacy** > **Data Privacy Job**.
3. In the Data Privacy Jobs list, click **New**.
4. In the form, fill in the fields.

Data privacy job fields

Field	Description
Name	Name of the job.
Description	Description of the job.
Privacy Configuration	The privacy policy configuration to use for this job. For detail on Privacy policy configurations, see Create a data privacy policy.
Frequency	○ **Run Once**: The job will run once in the specified time window. ○ **Weekly**: The job will run each day of the **Day of the week** field in the specified time window. ○ **Monthly**: The job will run each day of month as entered in the **Day of the month** field in the specified time window.
Users	Select the users or group of users to be anonymized in this job. Up to 1000 users of a group can be processed. **Note:** This field displays only when the selected privacy policy condition requires a selection of user records.
Dry Run	Run the job as a test. No records are affected when running this job. Results are displayed in the **Summary** field as though the job had executed. **Note:** **Dry Run** must be turned off when configuring a data privacy job with rollback. See Roll back a data privacy job for details.
State	State of the data privacy job: ○ **Completed**: Job completed successfully. ○ **Ready to Schedule**: Default state for new jobs. ○ **Rollback in progress**: The job has been set to roll back anonymization.

Field	Description
	○ **Rollback complete**: The anonymization job rollback has successfully completed. ○ **Completed with Errors**: The job completed but has errors. See Data Privacy Job Logs for more information. ○ **Error**: Job did not complete and had an error. See Data Privacy Job Logs for more information A read-only field.
Estimated record count	The estimated number of records this job affects. A read-only field.
Summary	A read-only field that displays the results of the job when you execute it.
Time window start	The start of the time window to run this job. The job will run after the time entered in this field. A valid time value is in Coordinated Universal Time based on a 24 hour time notation.
Time window end	The end of the time window to run this job. The job runs before the time entered in this field. If the job has not yet completed, the job will pause and resume at the next time window start. The end time must be after the start time. A valid time value is in Coordinated Universal Time based on a 24 hour time notation.

5. Right-click the form header and select **Save** from the context menu.
   After saving the record, the **Schedule Job** and **Delete Job** buttons appear.
6. Click **Schedule Job** to run your job.
   The job runs between the times selected in the **Time window start** and **Time window end** fields. If the job has not completed during the start and end time window, the job will continue at the next time window start.

   Note: A job can only be executed once, even if **Dry Run** is selected. To run the same job again, create a data privacy job using the same field values.

7. **Optional:** Choose one of the following functions:
   - **Cancel Job**: Cancels the data privacy job.
   - **Pause**: Pauses job and rollback recording, if rollback has been selected. A warning message will display after a three-day expiry period for rollback contexts. See Roll back a data privacy job for details.
   - **Resume**: Restarts a paused job. Rollback is not supported for resumed jobs if paused. Cancel the job and create a data privacy job. The recording uses an unexpired rollback context.

Data privacy job rollback

Database changes are captured for actions like jobs and scripts so that the changes can be rolled back. Roll back a data privacy job for when human error inadvertently anonymizes incorrect user information. The rollback de-anonymizes the data from the data privacy job.

Overview

- Rollback is limited to a few days, per the configured expiry duration of the RollbackContext of the new RollbackType *REDACT*. After expiration of the RollbackContext associated with a data privacy job, the rollback function is no longer available for that job.
  - A rollback context from de-anonymization is saved for three days by default.
  - The default expiry time can be set to a value greater than one by the data privacy administrator in the **RollbackContext** of the new **RollbackType** *REDACT*. Set the value in the Glide system property *glide.rollback.expiration_days_redact*. See Roll back and delete recovery.

    To learn more about adding or creating a system property, see Add a system property.

    for additional information.
- Rollback is available for data privacy jobs that are in a Completed, Canceled, or Error state.
- For every successful sys_user de-anonymization job which is configured with a data privacy policy with rollback support turned on, a rollback context is created. There can be at most one rollback context per data privacy job.

Roll back a data privacy job

Roll back a data privacy job on your non-production instance that uses anonymized data from your production instance to a state prior to de-identification of a data class or user job.

Before you begin

Role required: data_privacy_admin or data_privacy_processor, and admin

Procedure

1. Elevate to the **data_privacy_processor** or **data_privacy_admin** role.
   For details on role elevation, see Elevate to a privileged role.
2. Navigate to **All** > **System Security** > **Data Privacy** > **Data Privacy Job**.

   Note: In advance, a data privacy policy configuration that supports rollback must be created. Create a data privacy policy.

3. Create a data privacy job and select a Privacy Configuration that supports Rollback.
   See Configure a data privacy job.
4. Schedule the job for data anonymization.
   - After the job has run, the data will be anonymized for the selected configuration.
   - A message displays in the job informing of the expiration time.
     Within the expiry period you have the ability to roll back the job.

     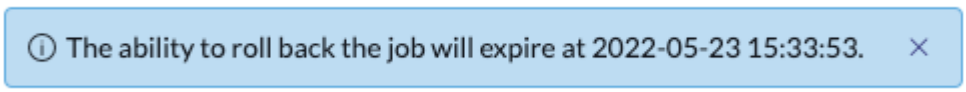

5. Elevate to the **data_privacy_processor** role.

6. Open the data privacy job to be rolled back.
7. Select **Rollback** to de-anonymize the data.

Data privacy clone

As customer data are cloned from a source to a target instance, typically from production to non-production, sensitive data are de-identified on the target instance.

Important: This section is for Data Privacy(Classic). See Data privacy for the latest up to date store version.

A data privacy administrator configures post-clone policies. After the post-clone script completes on the target instance, the users will see de-identified data and will not have access to the original data. Data privacy administrators can configure de-identification policies to apply on the target instance when cloning to ensure that the target instance will not have original sensitive data. An order is specified for the policy relative to other policies to be executed.

Note: Data privacy cloning is not available on self-hosted instances.

Data privacy clone has the following additional attributes:

- The data privacy plugin creates the post clone script to be executed on the target instance.
- Data privacy jobs created for PostClone configuration can run in parallel if they don't involve the same tables.
- A data privacy job for Postclone configuration with a higher Application Order might start before another job of lower order, if the job with the higher order does not involve any table related to other lower order job.
- With the data privacy plugin, data privacy tables are by default in the Clone Data Preservers table set.
- An attempt to add a data privacy table (dp_[table]) to Clone Exclude Tables will get a warning, that table should not be excluded.

Configure data privacy clone request

Data privacy clone integration is configured using a PostClone script to create and execute data privacy jobs for configured policies on the target. After running the script, users will see de-identified data and will not have access to the original data.

Before you begin

Role required: data_privacy_clone_processor, data_privacy_admin, and admin

Procedure

1. Activate the data privacy plugin (sn_dp_store_app) on the source instance.
   The plugin can only be installed by Customer Service and Support.
   The data privacy PostClone script is installed.
2. Elevate to the **data_privacy_admin** role.
   For details on role elevation, see Elevate to a privileged role.
3. Navigate to **System Security** > **Data Privacy** > **Privacy Policy Configuration**.
4. Create a Privacy Policy configuration.
   Select **Apply to All in Data Class** and **Apply when Cloning**. See Create a data privacy policy for more information.
5. Back up data privacy configurations.
6. As the data privacy admin, submit a clone request.

Result

The data privacy PostClone script executes on the target Instance and The PostClone script creates a data privacy federated job record on the target instance. The federated job creates and executes a data privacy job for each post-clone policy, in Application Order, on the target instance.The backup source is cloned to the target Instance. Data privacy PostClone script creates and executes data privacy jobs for configured policies on the target instance.

The elevated data privacy clone processor can log on to the target instance and monitor the post-clone federated job state on the dp_federated_job.list and dp_job.list.

Data Privacy Job Logs

Review errors from Data Privacy jobs.

Before you begin

Important: This section is for Data Privacy(Classic). See Data privacy for the latest up to date store version.

Role required: data_privacy_admin, data_privacy_clone_processor, data_privacy_processor, data_privacy_auditor

About this task

Use the **Data Privacy Job Logs** to review failed discovery, classification, and anonymization jobs.

Procedure

1. Navigate to **System Security** > **Data Privacy(Classic)** > **Data Privacy Job Logs**.
2. Review the table

Data Privacy Job Logs table

Label	Description
Created	Who or what created the log entry
Level	The severity of the error
Message	Description of the cause of the error
Source	Plugin source of the error
Type	The type of error that occurred
Code	Error code that occurred
Job Table Name	The table the job is located in
Job Id	The job ID
Target Table Name	Table the job executed on
Target Table Column	Column the job executed on
Target Table Record	Record the job executed on

Data Discovery

Use Data Discovery to identify sensitive data within an instance, such as credit card information, emails, or social security numbers.

Explore Data Discovery  Learn about Data Discovery.	Configure Data Discovery  Get help configuring Data Discovery.
Roles in Data Discovery  Learn about roles in Data Discovery.	Data Discovery Findings  Review Data Discovery Findings.

Exploring Data Discovery (Classic)

Use Data Discovery to identify sensitive data within an instance, such as credit card information, emails, or social security numbers.

Data Discovery runs a user-defined set of jobs on a set of tables. The jobs search for and report sensitive information for review on the Data Discovery dashboard. A scheduled job automatically uses all active data patterns and target tables when it runs.

Data Discovery also includes pre-allocated roles with varying levels of access to data.

Access the Data Discovery dashboard

Navigate to **All** > **System Security** > **Data Discovery (Classic)** > **Dashboard** to view the Data Discovery Dashboard and to review your current job findings.

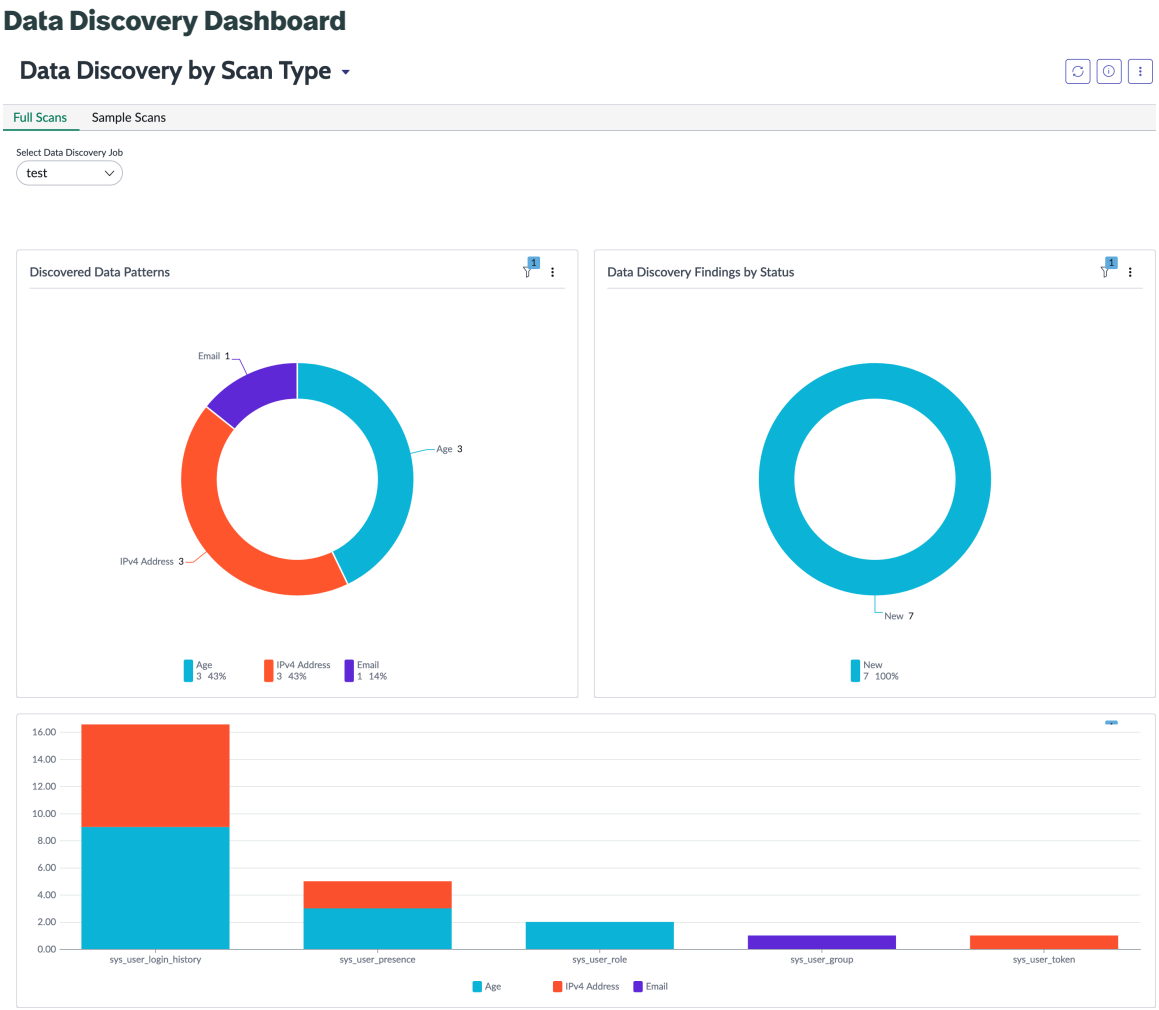

Discovered Data Patterns

Tip: Select a record to view more information and take further action with the findings. See Data Discovery job results for more.

The Discovered Data Patterns chart displays the number and percentage of currently discovered data patterns. Select a section to take further action.

Data Discovery Findings by Status

The Data Discovery Findings chart reviews the number and percentage of currently discovered data patterns by their status.

Discovered Tables

The Discovered tables highlight which tables currently discovered data is in, and their data pattern.

Columns with Discovered Data

The Columns with Discovered Data section reviews currently discovered data by their column.

Activating Data Discovery

The application installs Data Discovery and related ServiceNow® Store applications and plugins if they aren't already installed.

Before you begin

Data discovery requires a separate subscription from the rest of the ServiceNow AI Platform.

To purchase a subscription, contact your ServiceNow account manager. When you purchase a subscription, certain plugins are activated automatically. If a paid plugin isn't activated automatically, you can manually activate it from the All Applications list in your instance.

Note:

Before purchasing a subscription, you can evaluate this feature on a non-production instance without charge by requesting it from the Now Support Service Catalog.

Installing the Data Discovery Store App will auto install the Data Discovery APIs plugin

Role required: admin

About this task

The following items are installed with data discovery:

- Roles
- Tables

For more information on the roles and tables installed, see Data Discovery roles and Default data patterns.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Data Discovery plugin (sn_data_discovery) using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Classify data in Data Discovery Findings page

Classify sensitive data found by all successful jobs through Data Discovery Findings.

Before you begin

Roles required: data_discovery_admin

A job must complete a successful run before any result data can appear to be classified.

About this task

The Data Discovery Findings page shows all data entries that have been found by a job. The Findings page enables you to review the entries and associate them with a classification.

Procedure

1. Navigate to **System Security** > **Data Discovery(Classic)** > **Data Discovery Findings**.
2. Select the entry that you want to classify.
3. In the Data Discovery Findings list, select **Classify**.
4. Select the data classes that you want to associate with the table entries.
5. Select **Classify**.

Data Discovery jobs

Data Discovery reviews your targeted information using user-defined data patterns and target tables.

Data Discovery works by first scheduling a job in the **Data Discovery Job** section. When a scheduled job runs, all active data patterns are searched for on the current target tables. For more information on creating and configuring a Data Discovery job, see Configure a Data Discovery job.

Note: Sample Scan jobs scan up to 10,000 records are scanned per table per pattern.

Warning: If a job is executed twice in a small amount of time, the number of scanned rows may return 0 for second run.

Data Discovery data patterns

The **Data Patterns** section shows all current data patterns, both active and inactive. A data pattern is a regular expression used to match sensitive data. For more information on how to access and configure data patterns, see Configure Data Discovery patterns. The **Active Data Patterns** section shows only active patterns. Inactive patterns don't appear.

Target tables

The **Target Tables** section lists all tables that can be processed for a job. To access the target tables, select **Target Tables**.

The following tables aren't supported:

- sr
- sysx
- v
- sh$
- syslog
- ua
- usageanalytics
- ecc
- clone
- jrobin
- pa

- sla_repair_log
- scan
- gcf
- fm_log
- log
- np$
- sn_data_discovery
- dp_configuration
- dp_federated_job
- dp_field_technique
- dp_job
- dp_job_summary
- dp_primary_reference
- dp_technique
- data_classification
- m2m_dictionary_dataclass

Granular Configuration

The Granular Configuration tool enables more specific and fine tuned controls for discovering sensitive data in a table. However there are key differences in its operation compared to a traditional data discovery job. A granular configuration scans only specified columns of a table for discovery, reporting its findings by the record. Findings from the scan of Granular Configurations, called Granular Findings have their actions specified per record, and require the data_privacy_admin role to be anonymized.

Configure a Data Discovery job

Configure a Data Discovery job and review the status of ongoing jobs. A Data Discovery job defines when a pattern is executed on a target table.

Before you begin

Role required: data_discovery_admin

Procedure

1. Navigate to **System Security** > **Data Discovery(Classic)** > **Data Discovery Job**.
2. In the Data Discovery Jobs list, select **New**.
3. On the Data Discovery job fields form, fill in the fields.

Data Discovery job fields

Field	Description
Name	Name of the job.
Description	Description of the job.
Scan Type	Number of entries to be scanned. Possible states are as follows:

Field	Description
	○ **Sample**: Scans 10,000 entries. ○ **Incremental**:Scans for sensitive data patterns in new or modified records. This scan type is recurring. ○ **Full**: Scans all entries.
Policy	The Data Discovery policy to be used in this job. See Data Discovery policies for more information.
Scan Attachments	Scans attachments for any sensitive data patterns. Images embedded under .pdf or .doc files will not be scanned for sensitive data. **Warning:** This feature requires sending your data to a controlled ServiceNow environment. Before enabling, please confirm you have agreed to the additional terms and conditions required to use this feature.
Context	Job details about patterns scanned, target tables entries hit, and time elapsed.
State	State of the Data Discovery job. The possible states are as follows: ○ **Ready to Schedule**: Default state for new jobs. ○ **Scheduled**: The job is scheduled to run. ○ **In Progress**: Job is actively running. ○ **Completed**: Job has finished running successfully. ○ **Error**: The job has stopped running due to an error. ○ **Canceled**: The job has been canceled. ○ **Paused**: The job is paused.
Summary	Displays the results of the job after you execute it.
Start Date	Sets the start date for the job.
Time window start	The start of the time window to run this job. The job will run after the time entered in this field. The time entered in the **Time window start** field must happen before the time entered in the **Time window end** field. **Note:** A valid time value is in Coordinated Universal Time based on a 24-hour time notation.

Field	Description
Time window end	The end of the time window to run this job. The job runs until the time entered in this field. If the job hasn't complete this time, the job pauses and resumes at the next time window start. The time entered in the **Time window end** field must happen after the time entered in the **Time window start** field. **Note:** A valid time value is in Coordinated Universal Time based on a 24-hour time notation.

4. Select **Submit**.
   The **Schedule Job** and **Update** buttons appear.
5. Select **Schedule Job** to run your job.
   The job runs between the times selected in the **Time window start** and **Time window end** fields. If the job hasn't completed during the start and end time window, the job will continue at the next time window start.
6. **Optional:** Choose one of the following functions:
   - **Cancel Job**: Cancels the job.
   - **Pause**: Pauses the job.
   - **Resume**: Restarts a paused job.

Attachment scanning in Data Discovery jobs

Attachment scanning in Data Discovery enables you to scan, discover, and report on sensitive data in file attachments.

Sensitive data such as social security Numbers (SSN), credit card numbers, and other personally identifiable information (PII) are often contained in file attachments. Data Discovery jobs attachment scanning, enables you to discover and report sensitive data in certain attachments.The attachment scanning feature is only available for scheduled scans.

Warning: This feature requires sending your data to a controlled ServiceNow environment. Contact your account and support teams for enabling the attachment scanning feature

Limitations

- Embedded images and attachments in PDF and DOC files will not be scanned for sensitive data.
- Files must be less than 10MB.
- Support is available for the following file types:
  - PDF
  - DOC(X)
  - TXT
  - XLS
  - CSV

Configure Data Discovery patterns

Configure a Data Discovery pattern and review current patterns. A Data Discovery pattern defines the regular expression used to match data against a target table.

Before you begin

Role required: data_discovery_admin

About this task

Custom Data Discovery patterns can be used with Now Assist anonymization in addition to the base system patterns provided with the platform. A pattern applies to Now Assist prompts when it is associated with the **Generative AI** data channel. Base system patterns that don't include "(Generative AI)" in their name can also apply to Now Assist, provided they are associated with the **Generative AI** data channel. Configured patterns apply consistently across Now Assist skills, Now Assist Virtual Agent, AI Agents, and custom skills built with the Now Assist Skill Kit.

Note: Data Privacy for Now Assist is available in Yokohama and later releases. On Xanadu instances, use the Sensitive Data Handler to mask sensitive data for generative AI.

Important: Now Assist anonymization uses two-way masking. When a Now Assist skill such as incident summarization processes a record, sensitive data matching an active pattern is replaced with placeholder tokens in the prompt sent to the large language model (LLM). The original values are then restored in the response returned to the end user. End users therefore see unmasked data in Now Assist responses; this is by design. The purpose of masking is to prevent sensitive data from being transmitted to the LLM, not to hide it from the end user in the final response.

Procedure

1. Navigate to **System Security** > **Data Discovery** > **All Data Patterns**.
2. In the Data Discovery Pattern list, select **New**.
3. In the Data Discovery job fields form, fill in the fields.

Data Discovery job fields

Field	Description
Internal Scope	Scope of the pattern.
Description	Description of the job.
Name	Name of the data pattern.
Application	Application scope of pattern.
Expression	Regular expression used to discover the data pattern. **Note:** Expression length must be less than 1000 characters.
Keyword(Optional)	A specific word(or words separated by comma) to be searched for around a expression. Must be used with **Keyword Proximity**

Field	Description
	Note: A keyword can be used to search for additional context for a pattern. For example, using keyword can help differentiate between a date of birth or a date of hire given they have the same MM/DD/YY formatting.
Keyword Proximity(Optional)	How far from the expression to search for keywords. Must be used with **Keyword** **Note:** Default is 30, upper bound of 64
Privacy technique configuration	The masking technique applied to matched data before it is sent to the LLM. Common techniques include: **Synthetic replacement** Replaces the matched value with a realistic but fictitious substitute (for example, substituting a different email address). Use when the LLM needs plausible values to maintain response quality. **Static replacement** Replaces the matched value with a fixed non-inferable placeholder (for example, replacing any SSN with "999-99-9999"). **Selective replacement with x** Obscures part or all of the value using wildcard characters (for example, masking most digits of a card number while retaining the last four). Use when partial visibility is acceptable and helps the LLM understand context. **Remove** Deletes the matched value from the prompt entirely.
Synthetic Value	List of values substituted for the patterns
Type	Type of pattern ○ **Local**: The pattern is regex-based ○ **Model**: Uses AI/ML service

4. Select **Submit**.
   - The **Test** button enables you to test your regular expression before submitting the data pattern list.

The data pattern must be set as active to be used with scheduled jobs.

5. Navigate to **System Security** > **Data Discovery** > **Active Data Patterns**.
6. In the Data Discovery Active Pattern list, select **Edit**.
7. Select the pattern list from **Available Lists** and move it to **Selected Lists**.

Default data patterns

Review the default data pattern regular expressions included in Data Discovery. These default data patterns can be used to filter table entries for further classification.

The following are available default patterns for data discovery.

Name	Description	Regular Expression	Keywords	Examples
Age	A person's age between 0-99	\b([0-9]|[1-9][0-9]|1[012][0-9])\b	age	**Matching** 24 **Non matching** • 103 • -2
Date of birth	Date of birth using the DD/MM/YYYY format	\b[0-3]?[0-9]/[0-3]?[0-9]/(?:[0-9]{2})?[0-9]{2}\b	dob, birthday, date of birth	**Matching** • 06/18/2012 • 1/1/19
Email	Standard email address	\b[\w!#$%&'*+/=?`{\|}~^-]+(?:\.[\w!#$%&'*+/=?`{|}~^-]+)*@(?:[a-zA-Z0-9-]+\.)+[a-zA-Z]{2,6}\b		**Matching** • johndoe@emailserver.com • historyprofessor@college **Non matching** • notanemail.com • bademail@.org
Vehicle identification number	A vehicle identification number (VIN)	\b[A-HJ-NPR-Z0-9]{17}\b		**Matching** AHUYA31581L000000
IP Address	Standard IP address	**4 digit IP** \b(?:(?:25[0-5]|2[0-4][0-9]|[01]?[0-9]		**Matching 4 digit IP** 102.28.46.103 **Matching 6 digit IP** 914b:d45a:61ea:6346:59bc:

Name	Description	Regular Expression	Keywords	Examples
		[0-9]?)\.){3}(?:25[0-5]|2[0-4][0-9]|[01]?[0-9][0-9]?)\b **6 digit IP** \b(((?:[0-9A-Fa-f]{1,4}(?::[0-9A-Fa-f]{1,4})*)?)::((?:[0-9A-Fa-f]{1,4}(?::[0-9A-Fa-f]{1,4})*)?))|((?:[0-9a-fA-F]{1,4}:){7}[0-9a-fA-F]{1,4})\b		
Credit Card- Visa	Visa credit card number	\b4[0-9]{12}(?:[0-9]{3})?\b		**Matching** 4444434342424242
Credit Card- American Express	American Express credit card number	\b3[47][0-9]{13}\b		**Matching** 378225246366005
Credit Card- Mastercard	Mastercard credit card number	\b(?:5[1-5][0-9]{2}|222[1-9]|22[3-9][0-9]|2[3-6][0-9]{2}|27[01][0-9]|2720)[0-9]{12}\b		**Matching** 5555444455554444
Credit Card- Diners Club	Diners Club credit card number	\b3(?:0[0-5]|[68][0-9])[0-9]{11}\b		**Matching** 3056930009020004
Credit Card- Discover	Discovery credit card number	\b6(?:011|5[0-9]{2})[0-9]{12}\b		**Matching** 6011025690875424

Name	Description	Regular Expression	Keywords	Examples
Credit Card- CCV	Credit card security number	\b[0-9]{3,4}\b	cvv,verification code,security code	**Matching** 124
Credit Card- Expire Date	Credit card expiration in MM/YYYY format	\b(([1-9])|(0[1-9]|1[0-2]))\/?([0-9]{4}|[0-9]{2})\b	expire,exp	**Matching** • 02/2027 • 04/23 **Non matching** 03/9
USA- Social security number	USA citizen social security number	\b(?!666|000|9\d{2})\d{3}-(?!00)\d{2}-(?!0{4})\d{4}\b		**Matching** 001-22-111
USA- Phone Number	USA phone number **Warning:** Does not use the USA calling code.	\b\(?([0-9]{3})\)?[-.]?([0-9]{3})[-.]?([0-9]{4})\b		**Matching** 2065550199 **Non matching** 1 555 238 0199
USA- Passport Number	9 digit USA passport number	\b[a-zA-Z0-9]\d{8}\b		**Matching** 770022122
USA- Taxpayer ID	USA taxpayer ID number	\b(9\d{2})([\-]?)([7]\d|8[0-8])([\-]?)(\d{4})\b		**Matching** 927 70 5828
USA- California State Driver License number	State of California, USA driver license number	\b[a-zA-Z]\d{7}\b		**Matching** A0002144
USA - Bank Routing number	US Bank Routing (ABA) number	\b((0[0-9])|(1[0-2])|(2[1-9])|(3[0-2])|(6[1-9])|(7[0-2])|80)([0-9]{7})\b		**Matching** 125210305

Configure Data Discovery target table

Add a target tables to be used in Data Discovery jobs. Only target tables will be scanned for data patterns.

Before you begin

Role required: data_discovery_admin

About this task
When a Data Discovery job is run it will run against all target tables with all active patterns.

Procedure
1. Select **System Security** > **Data Discovery (Classic)** > **Target Tables**.
2. Select **New**.
3. Select your target table in the **Table Name** field.
4. Select **Submit**.

Activate parallel jobs for Data Discovery

Use parallel jobs to reduce your Data Discovery job execution time.

Before you begin
Role required: admin

Procedure
By default Data Discovery runs jobs using a single thread. Multiple threads may be activated for Data Discovery full scans.

1. In the navigation filter enter `sys_properties.list`.
2. Create the property `com.glide.data_discovery.max_concurrent_item_workers`
3. Verify the **Type** is `integer`.
4. **Important:** It is recommended to start with a small number of workers per node(2 or 3).

   Enter the number of parallel job workers in the **Value** field up to 5.

Data Discovery roles

You can assign Data Discovery roles to limit user access to certain data types.

Data discovery administrator [sn_data_discovery.data_discovery_admin]

View, create, and modify, data patterns and related jobs.

- Data patterns:
  - Create
  - Read
  - Update
  - Delete
- Active data patterns:
  - Delete
  - Read
- Jobs:
  - Create
  - Read
  - Update

- Delete
- Schedule
- Pause
- Resume
- Cancel

- Target tables:
  - Create
  - Read
  - Write
- Granular configuration:
  - Create
  - Read
  - Update
  - Delete
- Granular findings:
  - Create
  - Read
  - Actions

> **Note:** Users with this role may only take the ignore action

Contains Roles

List of roles contained within the role.

- data_classification_auditor
- data_classification_admin
- sn_data_discovery.data_discovery_api_processor

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Admins are automatically assigned this role on product installation.

> **Note:** Avoid granting an admin role when more specialized roles are available.

Data discovery auditor [sn_data_discovery.data_discovery_auditor]

Read data patterns and target tables.

- Read data patterns
- Read active data patterns
- Read Data Discovery jobs

- Read target tables
- Review granular configuration
- Review granular findings

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

None.

Data classification administrator [data_classification_admin]

Read data patterns, discovery jobs, and the data_classification table when an enabled specific pattern finding is classified.

Contains Roles

The Data classification role contains the List of roles contained within the data_classification_auditor role.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Note: For more information about the Data Classification admin role, see Installing Data Classification plugin demo data.

Data Discovery job results

The Data Discovery Findings page shows details on the data found by a job. You can use the Findings page to review the results of a job and begin classifying data.

The Data Discovery Findings page displays the following details after a completed job.

Column	Description
Dictionary Entry	Column of the target table where the data was
Table	Target table where the data was found
Data Pattern	Data pattern used to find the data
Data Pattern Match Count	Number of data entries that match the data pattern

Column	Description
Total Row Scan Count	Number of rows scanned during the job
Percentage of Matching Rows	Percentage of rows in the target table that match the data pattern
Data Discovery Job	Job used on the target table
Status	Status of the entry

After you run a Data Discovery job, the results have a status of New. If no action is necessary, you can leave the data alone, which automatically sets the status to Ignored. Otherwise, you can classify the data, such as by creating user-defined data classifications, to prepare for data anonymization using the Data Classification tool.

Granular Findings

Users may select the **Track Granular Findings** action to use the Granular Findings page to take action on specific discovered records. See Granular Findings for more information.

Available Protections

Select the **Available Protections** button to review an entries control service, status, notes, and last check in. As of Xanadu this feature supports only Field Encryption.

Protecton services can also be accessed from the Dashboard. For more information on the dashboard see, Data Discovery.

Consolidate findings

Job results can be consolidated using the **Consolidate Findings** button. Job results shared between two separate jobs will then be consolidated under the most recently run job.

Note: Only finished full scan jobs may be consolidated

Classify data in the Data Discovery job results page

Classify data in Data Discovery directly from a job's results page.

Before you begin
Role required: data_discovery_admin and admin

A job must complete a successful run before any result data can appear to be classified.

Procedure

1. Go to **System Security** > **Data Discovery** > **Data Discovery Job**.
2. Select the entry that you want to classify.
3. In the Data Discovery Findings list, select **Classify**.
4. Select the data classifications that you want to associate with the table entries.
5. Select **Classify**.

Data Discovery supported data types

Check which field types are supported when using Data Discovery.

Note: Not all field types that have been classified are available for Data Discovery.

Some high risk field types are turned off by default, as detailed in the table. For more information about fields, see Field Types.

Supported field types for data discovery

Field type	Available by default
audio	No
condition	No
condition_string	No
currency	Yes
decimal	Yes
due_date	Yes
float	Yes
glide_date	Yes
glide_date_time	Yes
glide_duration	No
glide_time	Yes
html	No
icon	No
integer	Yes
ip_addr	No
ip_address	No
journal	Yes
journal_input	Yes
journal_list	Yes
longint	Yes
name_values	No
percent_complete	No
phone_number_e164	Yes
price	Yes
string	Yes
string_full_utf8	Yes
translated_html	No
translated_text	No
url	No
user_image	No
video	No
wiki_text	No

Scanning with Granular Configuration

Granular scan can be used to scan specific table columns for discovery. Traditional Data Discovery jobs scan the entire table to discover data, whereas granular scan targets specific columns of the table thereby offers more control over the discovery process

Before you begin

Role required: admin, sn_data_discovery_admin, data_privacy_admin

Procedure

1. Navigate to **All** > **System Security** > **Data Discovery(Classic)** > **Granular configuration**.
2. In the **Granular Configurations** list, select **New**.
3. In the **Granular Configuration** fields form, fill in the fields.

Granular Configuration fields

Field	Description
Dictionary Entry	Column to be scanned.
Table	Table to be scanned.
Column label	Label of column to be scanned
Scan start point	Sensitive data will only be discovered for the day of and after the scan start point. **Note:** If the scan start point is left empty, all entries in the column will be scanned. If scan start point is changed, the scanning will be reset to start from the newly configured timestamp.
Active	Check to activate the granular configuration.

4. Select **Submit**.
   Children of extensible tables may be selected for configuration.
5. **Optional:** Select the Granular Configuration for the parent table, select **Select Child Tables** and select child tables for scanning.

Result

The specific column of the target table will be scanned for discovery. Granular Scan makes use of Data patterns from **Active Data Patterns**. You can review the findings and take action at **All** > **System Security** > **Data Discovery(Classic)** > **Granular findings**. See Granular Findings for more information.

Granular Findings

Granular findings may be reviewed using the Granular Findings tool.

Reviewing Granular Findings

Granular Findings is used to review up to 500 findings from Granular Configurations. When there are 500 findings, scanning of Granular Configurations is paused until action is taken.

Column label	Description
Record	Discovered record.
Table	Parent table of record.
Column	Column of the record where sensitive information was discovered.
Data Patterns	Pattern used to discover record.
Action	Action to take on findings. Select to change **Important:** The data_privacy_admin role is required to take the Anonymize action on a record. **Review** Record is pending review. This is assigned to new granular discoveries. **Ignore** No action will be taken on the record **Anonymize** Record will be anonymized.
Status	Status of record. **New** Status assigned to finding when it is first reported **Processed** When user chosen action has been successfully applied on the finding **Note:** Processed findings are store for 3 days in the Granular Findings table before deletion **Manual Review** When applying user chosen action has failed. **Warning:** Findings in Manual Review should be deleted by users after taking appropriate actions.

Warning: Rollback is not supported for anonymization triggered from Granular Findings.

Contextual based discovery

Use contextual based discovery to help discover sensitive data that does not follow fixed patterns.

Data Discovery supports using a Named Entity Recognition (NER) model to discover data such as names, organizations, nationalities, and political affiliations. Data patterns with the type **Model** use this feature, see pattern type in Configure Data Discovery patterns for more details.

Warning: This feature requires a license check from the customer before it is enabled.

AI model-based discovery supports 5 patterns by default.

Contextual discovery patterns

Pattern	Description	Example text	Matching text
Person	The name of individual(s)	John Doe came by for a coffee.	John Doe
Nationality, religious or political groups (NRPs)	National, religious and political affiliations.	He was an American.	American
Location	Addresses or location information	He was in New York	New York
Date_Time	Dates and time information	He came at 9:30 today.	9:30 today
Organization	Organization names	He works at ServiceNow.	ServiceNow

Note: AI Model-based discovery is only supports English for use with Real time anonymization(RTA) for tables and the Data Kit channel.

Data Discovery policies

Use Data Discovery policies to scan specific tables and enable column based jobs.

Data Discovery policies allow you to have fine-grained control over defining specific data patterns, tables and columns that should be included in a job. Additionally, using policies enables you to create multiple discovery job settings at the same time, instead of one global configuration. You can review your Data Discovery policies by navigating to **All** > **System Security** > **Data Discovery (Classic)** > **Data Discovery Policies**.

Note: If no columns are selected for a policy, all columns of the target table will be scanned

See Create a Data Discovery policy to learn how to create a Data Discovery policy.

Create a Data Discovery policy

Create a Data Discovery Policy for granular control over your Data Discovery jobs.

Before you begin
Role required: admin, sn_data_discovery_admin, data_privacy_admin

Procedure

1. Navigate to **All** > **System Security** > **Data Discovery (Classic)** > **Data Discovery Policies**.
2. Select the **New** button.
3. Fill in the form.

Data Discovery Policy

Label	Description
Name	Name of the policy
Target Column(s)	Select the target columns to use in the policy **Note:** If no columns are selected, all columns of the target table will be scanned
Target Table	Select the target table to use in the policy
Active	Determines if the policy is active or inactive **Note:** Defaults to active
Data Pattern	The data pattern the policy will use.

4. Select **Submit**.

Result

Select the policy when creating a Data Discovery jobs.

Data Discovery Store

Discover and track sensitive data within a ServiceNow instance

Data Discovery Store

Use Data Discovery to discover sensitive data within your instance. Scan for structured sensitive data such as, credit card numbers, phone numbers, social security numbers (SSN), etc. Get started quickly by using out-of-the- box common data patterns or create their own for a more granular experience. For unstructured and often context-based sensitive data such as names, addresses, etc. you can use AI and machine learning to aid with identification

Data Discovery Store overview

Get an overview of your Data Discovery metrics, and a starting point for using its tools.

Data Discovery policy

Review current Data Discovery policies and configure them.

Data Discovery sources

See the data patterns and target tables currently used by your discovery jobs

Data Discovery scheduled discovery

Schedule and review discovery jobs, including granular column discovery jobs.

Data Discovery Store overview

Get visibility into sensitive data on your instance and work towards implementing security measures to prevent loss or exposure of sensitive data.

The **Overview** tab provides an easy starting point for using Data Discovery, alongside showing important discovery metrics and findings. Navigate to **All** > **Data Discovery** > **Overview** to begin using the tab.

Create discovery policy and schedule discovery job

You can begin using Data Discovery with the recommended action items. Select **Create Policy** to define what sensitive data patterns should be identified and how they should be handled. See Data Discovery policy for more information on creating a new policy. Select **Schedule Job** to schedule a data discovery jobs that will scan your instance for sensitive information. See Data Discovery scheduled discovery for more information on scheduling a new discovery job.

Data discovery summary

You can review at a glance important Data Discovery metrics. Additionally, you can select the charts to drill-down further. For more information on reviewing discovery findings see Review discovery findings.

Discovered data

A bar chart overlaying the classifications and amount of discovered data. It summarizes where and how many instances of PII were discovered across different tables in your instance, categorized by data pattern types.

Discovery status

A donut chart displaying the proportions of the current state of all discovered patterns- whether they are new findings pending review, have been classified, or marked as ignored

Discovered data and recent jobs

You can review newly discovered data and recent discovery jobs in this section. For a more detailed review of these tables and their contents see Data Discovery scheduled discovery

.

Data Discovery policy

Use Data Discovery policies to define what sensitive data patterns should be identified and how they should be handled.

The **Policy** tab of Data Discovery reviews current discovery policies. View additional details of a policy by selecting the **View details** button. Edit, delete, and switch the state of current policies by selecting the three dots ⋮. To create a new policy see Create new policy.

Create new policy

Create a Data Discovery policy to begin scanning tables for data patterns.

Before you begin
Role required: discovery.admin

Procedure

1. Navigate to **All** > **Data Discovery** > **Policy**.
2. Select the **Create new policy** button.
   Edit or delete an existing policy by selecting the three dots icon.
3. Fill in the fields of the pop-up window.

New Discovery Policy form

Field	Description
Name	Enter a unique name for the policy
Data patterns	Select which sensitive data patterns to scan for. You can select multiple data patterns **Note:** AI/ML model based data patterns require an additional license and are subject to additional terms and conditions.
Select user tables and columns	Check the table(s) to be scanned **Note:** All selected table columns will be scanned by default. See the step below to narrow your selection to specific columns
Selected tables and columns	Displays a list of currently selected tables and columns for the policy

4. **Optional:** Select the arrow icon to view and select specific columns of the table.
5. Select **Save**.

Result

By default the new policy will be set to **Active**. Select the three dots icon and select **Deactivate policy**/**Activate policy** to switch its state.

Data Discovery sources

Create, and select the data patterns to be used in Data Discovery, and what tables to scan.

The **Sources** tab of Data Discovery helps you review your all your data patterns, set which pattern to use for scans, and select the tables to scan.

All Patterns

The **All Patterns** table lists all current patterns. Note that inactive patterns will show up in this list.See Create new data pattern to create a new data pattern.

Note: Patterns musts be active to be used in discovery jobs.

All Patterns table

Label	Description
Name	Name of the data pattern.
Expression	Regular expression used to discover the data pattern.
Keyword	A specific word(or words separated by comma) to be searched for around a expression.
Keyword Proximity	How far from the expression to search for keywords.
Privacy Technique Configuration	Privacy technique used for the pattern

Active Patterns

The **Active Patterns** table lists all of the current active data patterns. See Select active data patterns to learn how set a pattern to active and used by discovery jobs.

Active Patterns table

Label	Description
Name	Name of the data pattern.
Expression	Regular expression used to discover the data pattern.
Keyword	A specific word(or words separated by comma) to be searched for around a expression.
Keyword Proximity	How far from the expression to search for keywords.
Privacy Technique Configuration	Privacy technique used for the pattern

Target Tables

The **Target Tables** table lists all tables currently selected for discovery jobs. See Select target tables to select tables for use in discovery jobs.

Target Tables table

Label	Description
Table Name	The name of the table
Display name	The display name of the table (Used in reporting)
Application	The application scope of the table.

Create new data pattern

Create a new Data Discovery store pattern.

Before you begin
Role required: discovery.admin

Procedure
1. Navigate to **All** > **Data Discovery** > **Sources**.
2. Select **All Patterns** in the navigation pane.
3. Select the **Create new** button.
4. Fill in the form.

New data pattern form

Field	Description
Description	Description of the pattern.
Name	Name of the data pattern.
Expression	Regular expression used to discover the data pattern. **Note:** Expression length must be less than 1000 characters.
Keyword	A specific word(or words separated by comma) to be searched for around a expression. Must be used with **Keyword Proximity** **Note:** A keyword can be used to search for additional context for a pattern. For example, using keyword can help differentiate between a date of birth or a date of hire given they have the same MM/DD/YY formatting.
Keyword Proximity	How far from the expression to search for keywords. Must be used with **Keyword** **Note:** Default is 30, upper bound of 64
Privacy Technique Configuration	Privacy technique used for the pattern
Synthetic Value	List of values substituted for the patterns
Type	Type of pattern
Application	Application scope of pattern.
Scope	Scope of the pattern.

5. **Optional:** Select the **Test** button to test your regular expressions if necessary.
6. Select **Submit**.

Result
The new data pattern must be set as active to be used in discovery jobs. See Select active data patterns for more information.

Select active data patterns

Select the active data patterns to be used for Data Discovery jobs.

Before you begin
Role required: discovery.admin

Procedure
1. Navigate to **All** > **Data Discovery** > **Sources**.
2. Select **Active Patterns** in the navigation pane.
3. Select the **Edit** button.
4. Check the patterns to activate, they will show in the right side of the popup.
5. **Optional:** Select and drag patterns to re-order them.

   Note:

   The order of the data patterns will also determine the order of anonymization techniques applied for data pattern anonymization.

6. Select the **Save** button.

Select target tables

Select target tables to be used in discovery jobs.

Before you begin
Role required: discovery.admin

Procedure
1. Navigate to **All** > **Data Discovery** > **Sources**.
2. Select **Target Tables** in the navigation pane.
3. Select the **Edit** button.
4. Check the tables to target, they will show in the right side of the pop-up.

   Note: Target tables will scan all columns, unless specified otherwise in policy.

5. Select the **Save** button.

Result
Selected tables will now be targeted by scheduled discovery jobs.

Data Discovery scheduled discovery

Set up and schedule data discovery jobs to scan your instance for sensitive information.

Discovery Jobs

The **Discovery Jobs** table lists all discovery jobs. See Create discovery job to create a new discovery job.

Discovery jobs table

Name	Name of the job
Description	The job description
Percent Complete	The percentage completion of the job
Scan Type	Scan type of the job. Possible states are as follows: • **Sample**: Scans 10,000 entries. • **Full**: Scans all entries.
Start Date	The start date of the job
Updated	When the job was last updated.
State	State of the Data Discovery job. The possible states are as follows: • **Ready to Schedule**: Default state for new jobs. • **Scheduled**: The job is scheduled to run. • **In Progress**: Job is actively running. • **Completed**: Job has finished running successfully. • **Error**: The job has stopped running due to an error. • **Canceled**: The job has been canceled. • **Paused**: The job is paused.
Run	How often the job is scheduled to run.

Discovery Findings

The **Discovery Findings** table lists all the discovered findings from discovery jobs. See Review discovery findings to review and classify findings from discovery jobs.

Discovery Findings table

Label	Description
Dictionary Entry	Column of the target table where the data was in
Table	Target table where the data was found
Data Pattern	Data pattern used to find the data
Data Pattern Match Count	Number of data entries that match the data pattern
Total Row Scan Count	Number of rows scanned during the job

Discovery Findings table (continued)

Label	Description
Percentage of Matching Rows	Percentage of rows in the target table that match the data pattern
Data Discovery Job	Job used on the target table
Status	Status of the entry

Granular Configuration

The **Granular Configuration** table lists all granular discovery jobs. See Create granular job to learn how to create a granular discovery job.

Granular Configuration table

Label	Description
Table	Table of the granular discovery job
Column label	Column label of the granular discovery job
Scan Start Point	Sensitive data will only be discovered for the day of and after the scan start point. **Note:** If the scan start point is left empty, all entries in the column are scanned.
Active	The state of the granular discovery job.

Granular Findings

The **Granular Findings** table lists all the granular findings from granular discovery jobs. See Review granular findings to review the findings from granular discovery jobs.

Granular Findings table

Label	Description
Record	Discovered record
Table	Parent table of the record
Data Pattern	Pattern used to discover the record
Action	Action to be taken on the record **Review** Record is pending review. This is assigned to new granular discoveries **Ignore** No action will be taken on the record **Anonymize**

Granular Findings table (continued)

Label	Description
	Record will be anonymized
Status	The status of the record **New** Status assigned to finding when it is first reported **Processed** When user chosen action has been successfully applied on the finding **Note:** Processed findings are store for 3 days in the Granular Findings table before deletion **Manual Review** When applying user chosen action has failed **Warning:** Findings in Manual Review should be deleted by users after taking appropriate actions.

Create discovery job

Create and schedule a new Data Discovery Store job.

Before you begin

Role required: discovery.admin

Procedure

1. Navigate to **All** > **Data Discovery** > **Scheduled Discovery**.
2. Select **Discovery Jobs** in the right side navigation pane.
3. Fill in the form.

Schedule new discovery form

Field	Description
Name	Name of the job.
Description	Description of the job.
Scan Type	Number of entries to be scanned. Possible states are as follows: ○ **Sample**: Scans 10,000 entries. ○ **Full**: Scans all entries.
Select policy	The policy to use for the scheduled job.

Field	Description
Start Date	Sets the start date for the job.
Time window start	The start of the time window to run this job. The job will run after the time entered in this field. The time entered in the **Time window start** field must happen before the time entered in the **Time window end** field. **Note:** A valid time value is in Coordinated Universal Time based on a 24-hour time notation.
Time window end	The end of the time window to run this job. The job runs until the time entered in this field. If the job hasn't complete this time, the job pauses and resumes at the next time window start. The time entered in the **Time window end** field must happen after the time entered in the **Time window start** field. **Note:** A valid time value is in Coordinated Universal Time based on a 24-hour time notation.

4. Select the **Schedule** button.

Review discovery findings

Classify data in Data Discovery Store

Before you begin
Role required: discovery.admin

Procedure

1. Go to **All** > **Data Discovery** > **Scheduled Discovery**.
2. Select **Discovery Findings** in the right-side navigation pane.
3. Check the entries to be classified.
   Select the **Edit** button to edit the classification of selected entries.
4. Select the **Classify** button.

   Note: Select the **Available Protections** button to review protection options for classified data.

5. Select the data classifications that you want to associate with the table entries.
6. Select **Classify**.

Result
The selected data is now classified.

Create granular job

Scan specific table columns in Data Discovery Store.

Before you begin
Role required: discovery.admin

Procedure
1. Navigate to **All** > **Data Discovery** > **Scheduled Discovery**.
2. Select **Granular Configuration** in the right side navigation pane.
3. Select **Create new**.
4. Fill in the form.

Create new granular configuration form

Field	Description
Table	The table to be scanned
Column label	The column to be scanned
Scan start point	Sensitive data will only be discovered for the day of and after the scan start point. **Note:** If the scan start point is left empty, all entries in the column will be scanned. If scan start point is changed, the scanning will be reset to start from the newly configured timestamp.

5. Set **Active** slider.
6. Select **Save**.

Result
A granular scan is scheduled to run on the target table and column, after it executes you can review the scan findings.

Review granular findings

Review granular findings in Data Discovery Store

Before you begin
Role required: discovery.admin

Procedure
1. Navigate to **All** > **Data Discovery** > **Scheduled Discovery**.
2. Select **Granular Findings** in the right side navigation pane.
3. Review the table entries.

Column label	Description
Record	Discovered record.
Table	Parent table of record.
Data Patterns	Pattern used to discover record.
Action	Action to take on findings. Select to change

Column label	Description
	Important: The data_privacy_admin role is required to take the Anonymize action on a record. **Review** Record is pending review. This is assigned to new granular discoveries. **Ignore** No action will be taken on the record **Anonymize** Record will be anonymized.
Status	Status of record. **New** Status assigned to finding when it is first reported **Processed** When user chosen action has been successfully applied on the finding **Note:** Processed findings are store for 3 days in the Granular Findings table before deletion **Manual Review** When applying user chosen action has failed. **Warning:** Findings in Manual Review should be deleted by users after taking appropriate actions.

Warning:

Rollback is not supported for anonymization triggered from Granular Findings.

Data Classification

Group data by type, using pre-defined or user-defined data classifications. If you have an assigned data classification administrator or auditor role, you can administer different data classes or visually analyze the current state of different types of data within the instance.

Explore Data Classification  Learn about Data Classification.	Configure Data Classification  Create and configure your own data classes.
Reference for Data Classification  Learn how Data Classification works with demo data.	Analyze Data Classifications  Learn how to analyze Data Classifications.

Exploring Data Classification

Explore about data classification.

Data Classification enables support for:

- Visibility into the types of data hosted on a ServiceNow AI Platform instances.
- Compliance with privacy laws, and meeting regulation requirements for industries such as financial services and medical device manufacturing.

Data classifications

Data classification is a standalone process in which you manually apply data classifications to existing dictionary entries in any table. See Data dictionary tables for additional information.

- You classify data as you find appropriate for your business and you can alter the available data classes as necessary.
- When you classify data, you can use the pre-defined data classifications, or create your own. Although use of pre-defined data classifications is optional, it is advisable do so as a starting point. These pre-defined data classifications are included in demo data that you can install in your instance. To learn more, see Installing Data Classification plugin demo data and Components installed with Data Classification demo data.
- If you create your own data classifications, you can also design a tiered hierarchical system with parent and child data classifications.

- When creating manual data classifications on an extended or child table, base fields inherited from the parent table are not available for selection.

  Classification is supported only for dictionary entries. You can't assign different classifications to inherited columns because they share the same dictionary entry. For example, you can't classify task.description as PII while classifying incident.description as Public.

Overview dashboard

Use the Overview dashboard to understand how your current data tables map to different data classifications. You can also analyze how your global, regional, international users may require different approaches to data classification, regarding the use or access to data. You can also customize the Overview dashboard content and layout to fit your needs.

To learn how to use the available scripted and REST APIs to apply the classification metadata within existing processes, workflows, and applications, see the following:

- Data Classification - REST API
- DCManager - Global
- ScopedDCManager - Scoped

Note: Data Classification supports domain separation, and the data_classification table itself is process separated.

Use cases

General Data Protection Regulation (GDPR) is a European Union regulation whose purpose is to provide individuals with control over their own personal data. You can use data classifications, such as Personally Identifiable Information, to identify where personal data is being stored in your instance. By applying the appropriate security mechanisms to protect that personal data from leaking out, your organization satisfies GDPR requirements.

If you store customer information in the ServiceNow AI Platform, use the Personally Identifiable Information (PII) classification code where needed to track data subject to regulation by local privacy laws. When you install demo data, it automatically applies this classification code to certain security-sensitive fields in the User [sys_user] table. To learn more, see:

- Components installed with Data Classification demo data
- Assigning data classifications to dictionary entries

You can apply a Restricted data classification to Employee table columns that store sensitive employee information such as Social Security Numbers (SSN). Administrators and auditors can then use the Overview dashboard to confirm that you have assigned data classifications to the correct columns. They can also view the classification details for restricted types of information.

Installing Data Classification plugin demo data

When you upgrade to or install Zurich (and above), the Data Classification (com.glide.data_classification) plugin is automatically activated. However, you should manually install the demo data that comes with the plugin. It includes several important pre-defined data classifications, and it also assigns one of them to specific User [sys_user] table columns in your instance.

Before you begin
Role required: admin

About this task

Regardless of whether you install demo data, the activated Data Classification plugin adds the following user roles in your instance:

- data_classification_admin: Can administer all aspects of the Data Classification application, including data classification setup and assignment.
- data_classification_auditor: Can audit Data Classification code assignments made to user tables and columns.

Important: Notice regarding use by Customers:

All decisions in connection with the implementation of this application are at the sole decision of the Customer. Customers acknowledge and agree that use of the application is not a representation by ServiceNow of compliance with any law or regulation and any suggested language, field or classification provided out of the box with the application does not constitute legal advice by ServiceNow.

Customers remain solely responsible for complying with their legal obligations under applicable law, including (but not limited to) data protection, security requirements, and privacy laws, and are responsible for configuring and making any necessary modifications to this application, including (but not limited to) templates, to meet the Customers' requirements.

Procedure

1. Navigate to **System Applications** > **All Available Applications** > **All**.
2. Find the Data Classification plugin using the filter criteria and search bar.
   An `Installed` message appears after you locate the plugin.
3. Click the icon with three vertical dots, then select **Repair** to access the Activate Plugin dialog.
4. Select **Load demo data**, and then click **Repair**.

Components installed with Data Classification demo data

When you install the demo data included in the Data Classification (com.glide.data_classification) plugin, several types of components are installed in your instance. These components include pre-defined data classifications and code assignments for specific User [sys_user} columns.

Data Classifications installed

data classification	Description
Confidential	Sensitive data that if compromised could negatively affect operations.
Internal	Internal data not meant for public disclosure.
Personally Identifiable Information	Also known as PII. Data that could potentially be used to identify a particular person.
Public	Data that may be freely disclosed to the public.
Restricted	Highly sensitive corporate data that if compromised could put the organization at financial or legal risk.

Data Classification assignments

Table	Assigned Column	Assigned data classification
sys_user	zip	Personally identifiable information
sys_user	first_name	Personally identifiable information
sys_user	email	Personally identifiable information
sys_user	city	Personally identifiable information
sys_user	middle_name	Personally identifiable information
sys_user	street	Personally identifiable information
sys_user	mobile_phone	Personally identifiable information
sys_user	last_name	Personally identifiable information
sys_user	country	Personally identifiable information
sys_user	gender	Personally identifiable information
sys_user	name	Personally identifiable information
sys_user	photo	Personally identifiable information
sys_user	state	Personally identifiable information
sys_user	home_phone	Personally identifiable information

Creating data classifications

Create your own user-defined data classifications in the Data Classification [data_classification] table that you can then assign to specific columns in specific tables.

Before you begin

Role required: data_classification_admin, admin

Procedure

1. Navigate to **All** > **System Security** > **Data classification** > **Data classes**.
2. Select **New**.
3. Fill in the fields on the form.

Field	Description
Classification Name	Name of the data classification.
Description	Description of the data classification.
Parent	Name of the parent data classification that this data classification is subordinate to. Leave the field empty if this data classification is not a parent to child data classifications.
Application	Application scope for this data classification.

4. If this data classification should be a parent to child data classifications, click **New**.

If you do not want to create child data classifications, skip this step.

5. Fill in the fields on the form.

Title

Field	Description
Classification Name	Name of the child data classification.
Description	Description of the child data classification.
Parent	Name of the parent data classification that this data classification is subordinate to. Leave the field empty if this data classification is not a parent to child data classifications.
Application	Application scope for this child data classification.

6. Click **Submit**.

Assigning data classifications to dictionary entries

Assign data classifications to specific table columns in the Dictionary [sys_dictionary] table. When you assign data classifications, it creates entries in the Dictionary-Data Class [m2m_dictionary_dataclass] table, which you can then review in the Overview dashboard.

Before you begin

Role required: data_classification_admin and admin

Procedure

1. In the Navigator pane, type `sys_dictionary.list`.
2. In Dictionary Entries, select each of the elements you want to assign specific data classifications to.
3. After selecting the elements, click **Actions on selected rows**, and then select **Classify**.

   Note: To clear previously assigned data classifications for selected dictionary elements, you can select **Clear classification**.

4. When the Assign to data class dialog appears, select the data classifications you want to assign to your selected dictionary elements, then click **Classify**.

   Warning: This will overwrite any existing classifications for the selected dictionary items.

   You can select multiple data classifications as needed.

   See Data dictionary tables for additional information.

Analyzing data classifications using the Overview dashboard

The Overview dashboard reports the current state of data classifications within your instance and how your users are distributed by location.

If you have a data classification administrator or auditor role, you can visualize the current sensitivity of instance data, which helps increase security and compliance with privacy laws. By using the Location field in user records, administrators map users across different regions that have differing privacy regulations.

Tables and columns tab

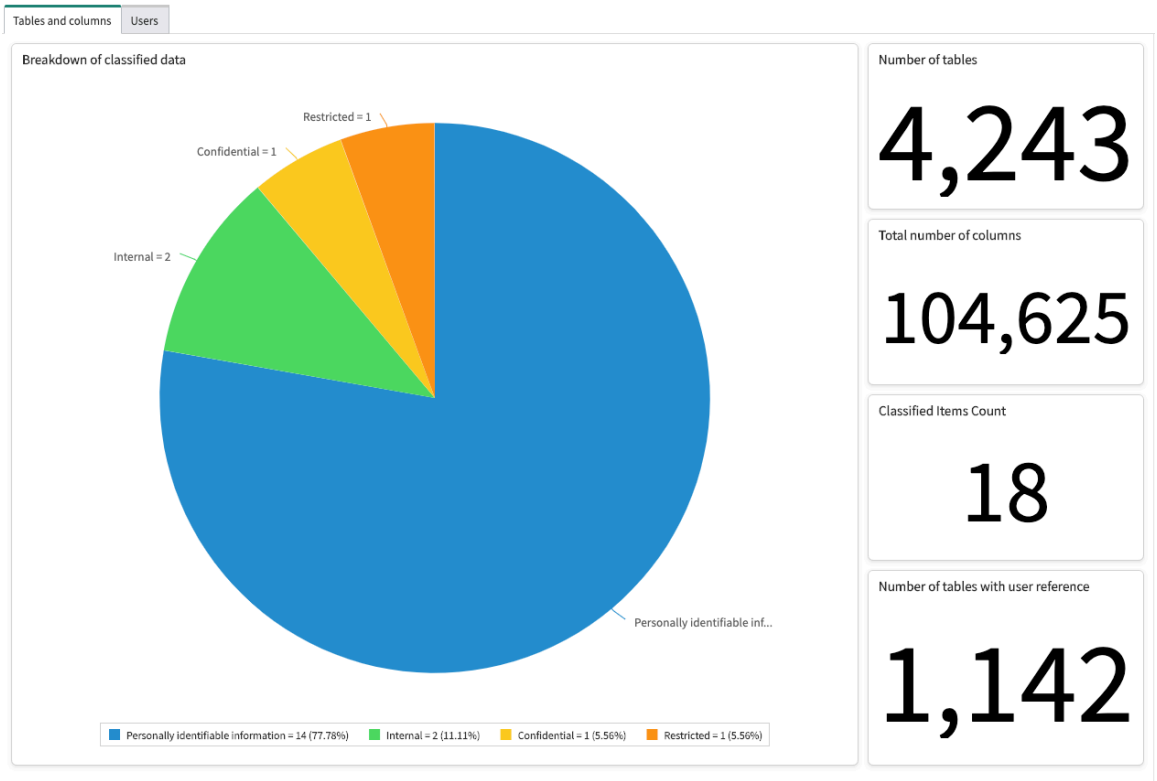

Users tab

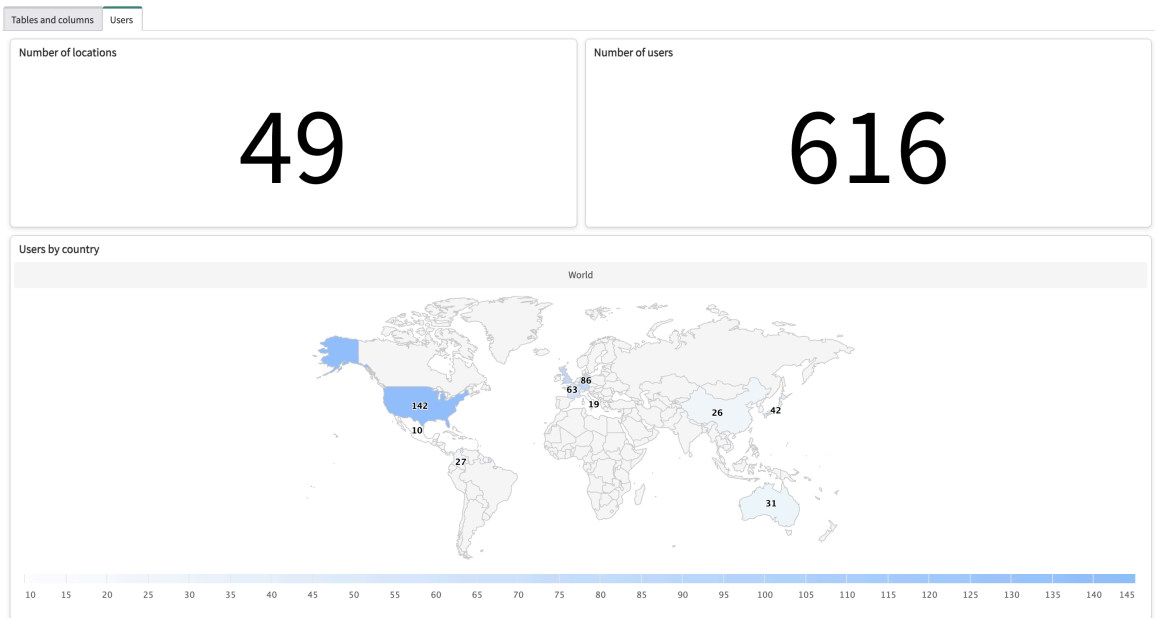

Required ServiceNow AI Platform roles

- data_classification_admin: Administers all aspects of the Data Classification application, including data classification setup, and assignment.
- data_classification_auditor: Audits Data Classification code assignments.

Use cases

For additional examples of how different people in your organization would use this dashboard, see Use Cases in Data Classification.

User	Dashboard use
data_classification_admin	Confirm that you have assigned data classifications to the correct fields in your instance.
data_classification_auditor	Audit the senility and security of data in your instance.

Reports

Title	Type	Source table	Description
Breakdown of classified data	Donut chart	m2m_dictionary_dataclass	Provides a breakdown of instance data classifications, by data class, total displayed in center.
Total tables	Single Score	sys_dictionary	Total number of data tables in the entire dictionary.
Total columns	Single Score	sys_dictionary	Total number of classified data columns in the entire dictionary.
User reference columns	Single Score	sys_dictionary	Total number of user reference columns in the dictionary.
User location count	Single Score	sys_user	Total number of distinct locations found for users
User count	Single Score	sys_user	Total number of user records in your instance.
Users by country	Map	sys_user	Breakdown of user records by country.

Domain separation and Data Classification

Domain separation is supported for Data Classification . Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: Enhanced

- Includes all aspects of **Basic** and **Standard** levels of support.
- Data-driven process enables service provider customers to modify business logic that is based on defined use cases. These configurations are UI-based and fail-safe so that configurations by one customer cannot affect another.
- Tenants of the instance must be able to configure minimum viable product (MVP) business logic and data parameters for themselves. This logic and parameters would be expected for the application's normal function.

Sample use case: Tenant-customers of a shared environment must be able to modify the impact, urgency, or priority matrix to set priority within their domain.

For more information on support levels, see Application support for domain separation.

How domain separation works in Data Classification

For domain separation, the application uses process separation for the Data Classification [sys_data_classification] table. For the Dictionary-Data Class [m2m_dictionary_dataclass] table, it uses data separation. For learn more about data and process separation, see Domain separation explained.

Related topics

Domain separation for service providers

servicenow

Encryption

Protect your sensitive data and stay compliant with regulatory#equirements and standards.

Key Management Framework

Use the Key Management Framework (KMF) to fully customize and manage how cryptographic operations are performed on your ServiceNow instance.

Field Encryption

Use

Field Encryption

to permit and deny access to encrypted data based on user role.

Field Encryption

includes basic key management using encryption modules.

Field Encryption Enterprise

Use

Field Encryption

to permit and deny access to encrypted data based on user role.

Field Encryption

includes basic key management using encryption modules.

Cloud Encryption

Encrypt your instance database using block encryption, along with enhanced key management.

Platform Encryption entitlement bundle

Upgrade to unlimited-use

Field Encryption Enterprise

, Cloud Encryption, and Database Encryption.

Full Disk Encryption (FDE)

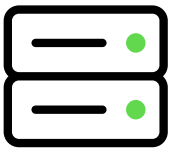

Full disk encryption applies encryption to the entire storage system within the database server only. Because this is the only customer data-storing component.

Edge Encryption

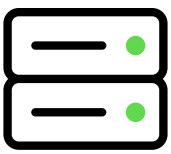

Certificates

Database Encryption

Encrypts sensitive data on your company premises before sending data over the internet to your ServiceNow instance. Data remains encrypted at rest on the instance.	Use certificates to establish secure connections and validate signatures.	Encrypt all stored data in real-time, providing protection for data online and offline with no loss of functionality.

Key Management Framework

Use the Key Management Framework (KMF) to generate, exchange, store, use, and replace the cryptographic keys used to encrypt and decrypt sensitive data on your ServiceNow instance.

Key Management refers to the activities involved in handling your cryptographic keys and related security parameters during the key's life cycle. Key Management Framework is based on National Institute of Standards and Technology (NIST) 800-57 guidelines. In accordance with these guidelines, you can use KMF to:

- Assign dedicated roles for cryptographic management and operations, auditing, and integration.
- Create cryptographic modules to configure of cryptographic specifications for unique cryptographic purposes and key types.
  - Symmetric key: encryption and decryption, key wrapping and unwrapping, and authentication
  - Asymmetric key: digital signature generation and verification, encryption and decryption, key wrapping and unwrapping
- Manage your key life cycle to generate, rotate, revoke, and suspend keys, including support of several key life cycle states
- Create module access policies (MAPs) to enforce access controls, to grant access only to users and scripts that you choose.
- Protect your cryptographic keys with the Federal Information Processing Standard (FIPS) 140-2-L3 hardware Root of Trust (RoT), Public Key Infrastructure (PKI), key hierarchy, and envelope encryption.
- Assign the auditing role to users to can then view auditing information such as key usage statistics.

Get started

Exploring the Key Management Framework	**Configuring the Key Management Framework**	**Key Management Framework Reference**
		
		Review additional Key Management reference materials

Learn about the components of the Key Management Framework, and how to use them to manage how cryptographic operations are performed on your instance.	Create and maintain Key Management components to customize and manage how cryptographic operations are performed on your ServiceNow instance.	
	Key Management Framework actions  One of the core features of KMF is to provide the capability#to manage #keys, such as revoking or rotating keys.#KMF properly secures sensitive data with the most up-to-date encryption materials and life cycle operations.	

Activation information

The ServiceNow Platform Encryption subscription bundle is a group commercial entitlement that includes Key Management Framework, Field Encryption Enterprise, Cloud Encryption, and Database Encryption.

Field Encryption Enterprise is the unlimited license of Field Encryption. The Field Encryption Enterprise plugin is available with the activation of the com.glide.now.platform.encryption plugin. For details, see Encryption and Key Management subscription bundle.

Note: KMF doesn't support domain separation, but can be used with on-premise instances.

Exploring the Key Management Framework

Learn about the components of the Key Management Framework (KMF), and how to use them to manage how cryptographic operations are performed on your instance.

Components of the Key Management Framework

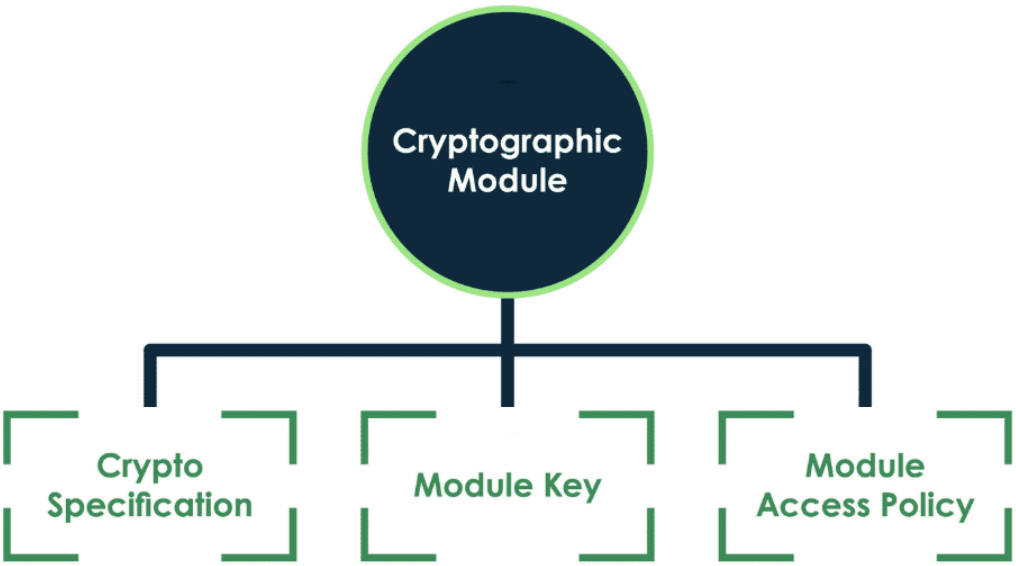

Key Management Framework consists of the following components.

Cryptographic modules

KMF is centered around managing cryptographic modules. These modules act as the parent record for the other components. They define what data on your instance is encrypted, and what method of encryption to use. Using multiple modules, you can encrypt different areas of your instance with different specifications.

For example, you can create a module to secure the data in your Human Resources application to users with a specific role. You could then create another module to encrypt Incident descriptions which are visible to certain users based on a script you create.

Module access policies are found by navigating to **All** > **Key Management** > **Cryptographic Modules** > **All**. For more information on these modules, see Cryptographic module overview.

Module keys

Cryptographic keys are strings of characters used in cryptography. When used together with a cryptographic algorithm, they can encode or decode your data. These keys are used by the cryptographic specifications assigned to your modules. You can choose to use a key generated by ServiceNow, or upload your own key.

You can access the module keys for a cryptographic module in the **Module Keys** related list in cryptographic module records. For more information on module keys, see Instance level keys in the Key Management Framework.

Cryptographic specifications

A cryptographic specification defines algorithms used to encrypt your data. These algorithms use a cryptographic key to encode or decode your data. Assigning a cryptographic specification to the module determines how the data assigned to that module is encrypted.

You can access the module keys for a cryptographic module in the **Crypto Specifications** related list in cryptographic module records. For more information on module keys, see Cryptographic specification overview.

Module access policies

Module access policies (MAPs) are the access controls you apply to your cryptographic modules. Use these policies to determine which users and scripts can access data encrypted by a cryptographic module.

Find module access policies by selecting the **View access policies** link in cryptographic module records. For more information, see Module access policy overview.

Key Management Framework workflow

1. Assign KMF roles

Administrators must begin by assigning themselves the sn_kmf.admin role. This role enables you to use KMF features and assign KMF roles to other users.

2. Configure KMF settings

Configure your field encryption settings to select either supplied keys or your own customer-supplied keys (CSK) for encryption.

3.Create cryptographic modules

Use cryptographic modules to select a set of data on your instance to be encrypted. In later steps, you assign a cryptographic specification to determine how to encrypt this data, and a module access policy to determine who can decrypt the data.

4. Create a cryptographic specification

The cryptographic specification defines a method of encryption. Once assigned to a module, it defines how the data assigned to that module is encrypted.

5. Create module access policies

After creating modules to secure your data, create module access policies to control which users and scripts are able to access the encrypted data.

6. Create a cryptographic module life-cycle policy

These policies place limits on cryptographic modules, such as how long a cryptographic key is valid. These policies can safeguard your cryptographic modules by limiting their exposure.

Key Management Framework benefits

Benefit	Feature	Users
Protect your sensitive and proprietary data.	Encryption and key Management	All
Maintain compliance with NIST 800-57 guidelines. These guidelines are provided by the National Institute of Standards and Technology to reduce cybersecurity risk to your networks and data.	Encryption and key Management	Security administrators
Use the Key Management Framework to generate, upload, view, and manage your cryptographic keys. Use key rotation for manual or scheduled rotation of your keys for increased security.	Key Management Framework	Security administrators

Cryptographic module overview

The Key Management Framework (KMF) is centered around managing Cryptographic modules. Use these modules to select a cryptographic mechanism and define where they're applied on your instance.

Cryptographic modules are the centerpiece of KMF. They define the specific cryptographic mechanisms used for cryptographic operations for a given use case.

For example, you want to secure the data in your Human Resources application with an AES-CBC with a 256-bit symmetric key. You can create a module for that purpose.

Cryptographic modules also support key life-cycle management. You can create and rotate your cryptographic keys, and define your encryption method. Cryptographic modules are composed of the following components:

Cryptographic specification

Defines which algorithm to use for encryption, and where the key will come from. All keys use the Advanced Encryption Standard with Cipher Block Chaining (AES CBC), but you can select either 128 or 256 bit. This specification covers both asymmetric and symmetric key-based cryptographic operations.

Note: Symmetric encryption uses a single key for both encryption and decryption. Asymmetric encryption uses a pair of keys, a public key for encryption and a private key for decryption.

Cryptographic keys

The key your module uses to encode or decode cryptographic data. This key can be generated by your instance, or a customer-supplied key you create and upload.

Module access policies

Module access policies are the access control mechanisms that place limits on whether data can be encrypted or decrypted.

Module policy exceptions

A control mechanism to define exceptions to a module access policy.

The following screen shows these high-level components in a cryptographic module:

Cryptographic module components

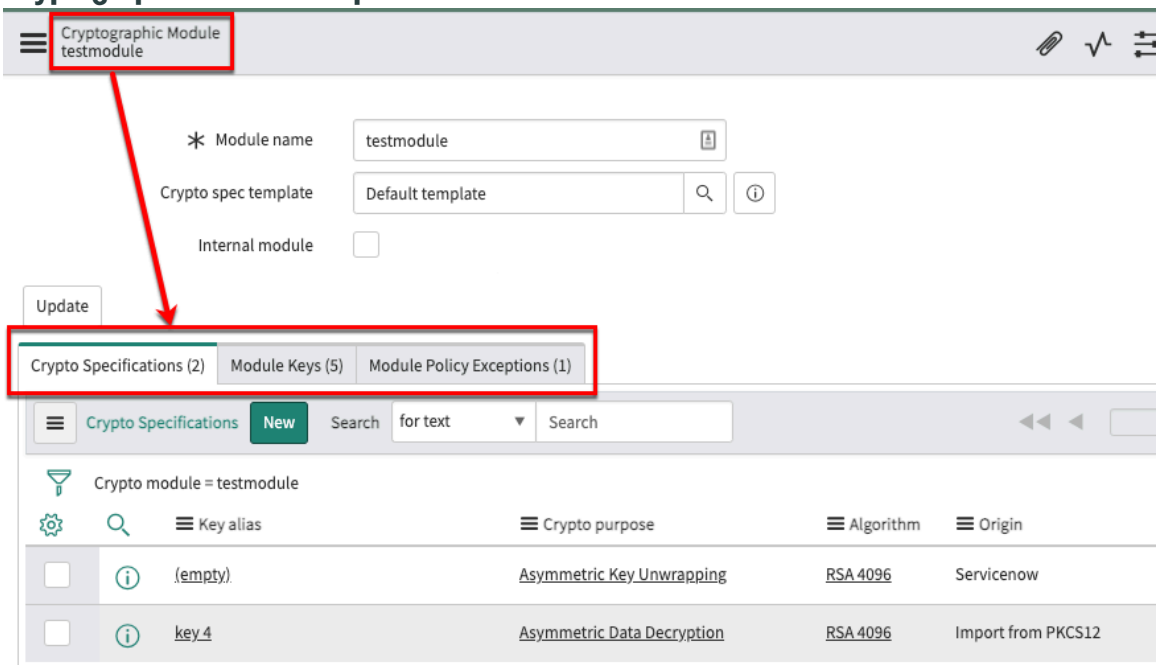

For details on creating cryptographic modules, see Create a cryptographic module.

Cryptographic specification overview

The Cryptographic specification is the component that defines aspects of your cryptographic module, including its cryptographic purpose and which encryption algorithm to use.

Cryptographic specifications can be tailored to a specified cryptographic purpose, covering both asymmetric and symmetric key-based cryptographic operations. Selection of a cryptographic purpose offers further choices, namely to a set of supported algorithms and key length configurations.

Cryptographic purposes, algorithms, and key information

Cryptographic Purpose	Algorithm	Key Information
Asymmetric Data Decryption	RSA	Asymmetric – 2048-bit, 3072-bit, and 4096-bit key
Asymmetric Data Encryption	RSA	Asymmetric – 2048-bit, 3072-bit, and 4096-bit key
Asymmetric Key Unwrapping	RSA	Asymmetric – 2048-bit, 3072-bit, and 4096-bit key
Asymmetric Key Wrapping	RSA	Asymmetric – 2048-bit, 3072-bit, and 4096-bit key
Signature Generation	RSA	Asymmetric – 2048-bit, 3072-bit, and 4096-bit key
Signature Verification	RSA	Asymmetric – 2048-bit, 3072-bit, and 4096-bit key
Symmetric Authenticity	HMAC	Symmetric – 256-bit, 384-bit, and 512-bit key
Symmetric Data Encryption/ Decryption*	AES-CBC *	Symmetric – 128-bit, 192-bit, 256-bit key
	AES-CFB	

Cryptographic purposes, algorithms, and key information (continued)

Cryptographic Purpose	Algorithm	Key Information
	AES-OFB	
	AES-CTR	
	AES-GCM **	
Symmetric Key Wrapping/ Unwrapping*	AES-CBC *	Symmetric – 128-bit, 192-bit, 256-bit key
	AES-CFB	
	AES-OFB	
	AES-CTR	
	AES-GCM **	

* AES-CBC supports equality-preserving options.#Field Encryption Enterprise utilizes AES-CBC.

** AES-GCM has built-in data integrity.

The configuration of these parameters is covered in Create a cryptographic module.

Module access policy overview

Module access policies (MAPs) are access controls that you apply to your cryptographic modules. Use these access policies to decide which users and scripts can access data encrypted by a cryptographic module.

Module access policies

Note: A subscription is required to utilize the Field Encryption Enterprise functionality. See Activate Field Encryption for more information on Field Encryption Enterprise.

Module access policies are introduced with the Key Management Framework (KMF) in the base system.

Module access policies expand on the role-based designations that were provided with the encryption modules. Module access policies can be based on the following:

- Basic (scope)
- Role
- System user
- Script
- Resource Exchange

Note: See Key Management Framework Resource Exchange for details.

In a cryptographic module, you must configure the correct module access policies to permit access to encrypted data. Without a module access policy associated with a cryptographic module, encrypted data isn't visible to users and associated fields and columns in lists display as empty.

In this example, the absence of a module access policy on the encrypted Short Description field hides the content from all users accessing the Incident table. With a module access policy in place, only users with a specific role are able to see the encrypted data.

Encrypted short descriptions with and without module access policies

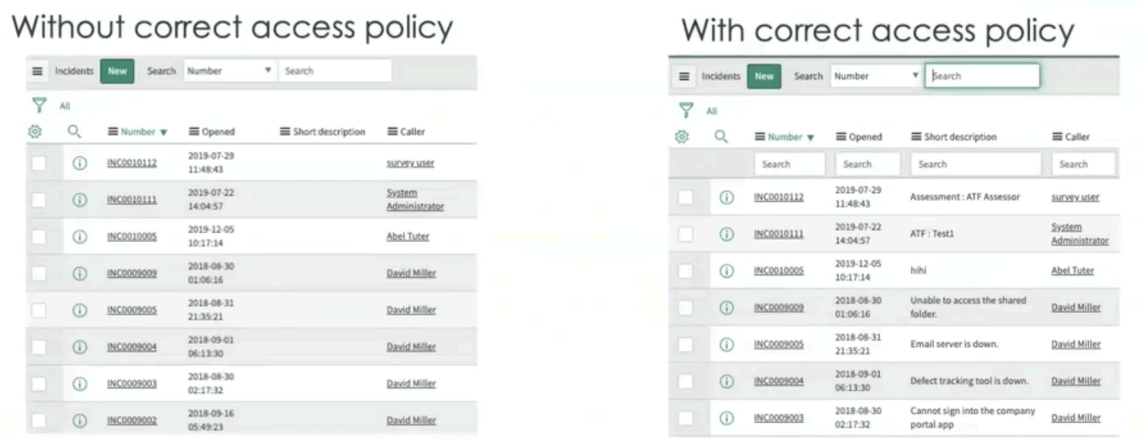

Note: The data in the column also appears empty to users without the correct role specified in the module access policy.

Refer to Create a module access policy for setup.

Autogen policies

Autogen policies are automatically system generated based on the default module access policy defined for the given cryptographic module. If there are no granular level policies defined when the system or a script tries to access the given cryptographic module, these global policies are generated and applied.

Important:

Autogen policy rules aren't applied for scheduled jobs types, or field encryption modules (modules where the parent module is Field Encryption).

Instance level keys in the Key Management Framework

Learn about the Key Management Framework (KMF) key structure, which uses envelope encryption to ensure that all platform keys under KMF management are protected through a chain of keys. Customer Data Encryption Keys (CDEKs) created by KMF are also included in this structure

.

KMF key storage architecture

The KMF key structure uses the SafeNet KeySecure Hardware Security Module (HSM). The HSM is designed to be physically and electronically tamper-proofed to meet the FIPS 140-2-L3 security standard. KMF uses envelope encryption to ensure that all platform keys under KMF management are protected through a chain of keys, including the module keys that can be generated by KMF.

Envelope encryption

Envelope encryption is the practice of encrypting a key with another key, also referred to as wrapping. Module keys are envelope encrypted by the Instance Key Encryption Key (IKEK),

which in turn is envelope encrypted by the Instance Root Key (IRK), which is finally envelope encrypted by the Root Key (RK). Since the IRK can only be accessed by the HSM, the IKEK must be uploaded for decryption.

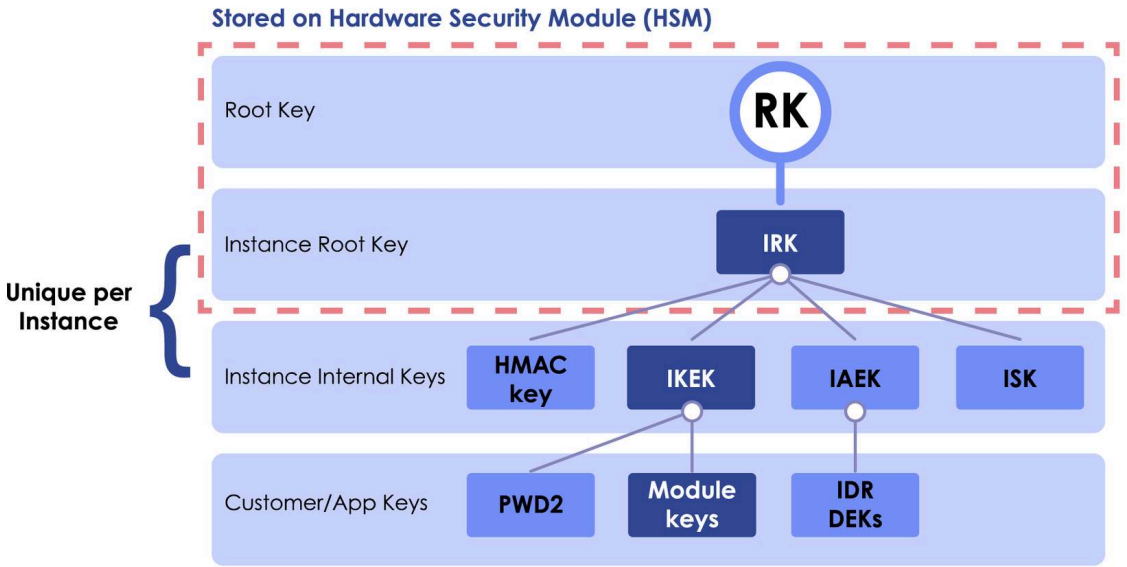

At the instance level, KMF defines several keys that are used internally for varying cryptographic purposes throughout the ServiceNow AI Platform.

This table provides examples of a subset of available keys that are managed and protected by KMF.

Key	Location	Description
Root Key (RK)	Hardware Security Model (HSM)	Root key used to decrypt the IRK.
Instance Root Key (IRK)	HSM	A key unique to your instance that is used to envelope-encrypt several instance internal keys.
Instance HMAC Key (IHK)	Instance	Unique per instance, the IHK is used internally for Hash-Based Message Authentication Code (HMAC) purposes. The IHK helps to verify the authenticity and integrity of module keys and is wrapped on either KeySecure or the File Key Store.
Instance Key Encryption Key (IKEK)	Instance	The IKEK wraps the module keys and is wrapped on either KeySecure or the File Key Store.
Instance Asymmetric Encryption Key (IAEK)	Instance	A key unique to your instance that is used internally for asymmetric encryption purposes. The IAEK is used to transmit confidential messages between an instance during Key Exchange or Instance Data Replication consumer approval.

Key	Location	Description
Instance Signature Key (ISK)	Instance	A key unique to your instance that is used internally for signing purposes.
Password2 (PW2)	Instance	With KMF, the key for PW2 fields is fully managed by KMF.
Customer Data Encryption Key (CDEK)	Instance	Encryption keys created through KMF are envelope-encrypted by the IKEK.
Instance Data Replication (IDR) Data Encryption Key (DEK)	Instance	Specific encryption keys used for the IDR process.

Configuring the Key Management Framework

Create and maintain Key Management components to customize and manage how cryptographic operations are performed on your ServiceNow instance.

Assign Key Management Framework roles

Administrators with the security_admin role can assign Key Management Framework (KMF) admins, who in turn can assign other Key Management Framework roles.

Before you begin

Role required: admin and security_admin

You must elevate to the security_admin role before assigning the KMF admin role. For instructions, see Elevate to a privileged role.

Important: KMF roles are required to use the Key Management Framework. Users without KMF roles aren't be able to access lists, tables, and modules used to configure key management.

Procedure

1. Elevate to the security admin role.
2. Navigate to **User Administration** > **Users** and select the user you want to be the KMF admin.
3. If the user doesn't already have the admin and security_admin roles, select **Edit** under the Roles related list and add **admin** and **security _admin**.
4. Navigate to **System Security** > **Key Management Administration**.
5. Select the user that you want to be KMF admin in the **Available Users** column and move them to the **Selected User(s)** column.

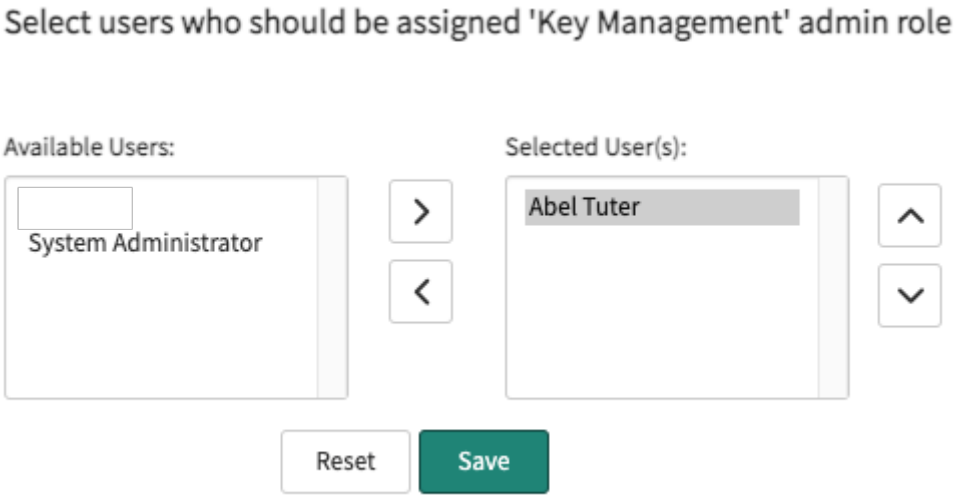

6. Select **Save**.
7. Navigate to **User Administration** > **Users** and select the user you just gave the KMF admin role to.
   That user now has the sn_kmf.admin role in the Roles related list. That user can now assign other KMF roles.

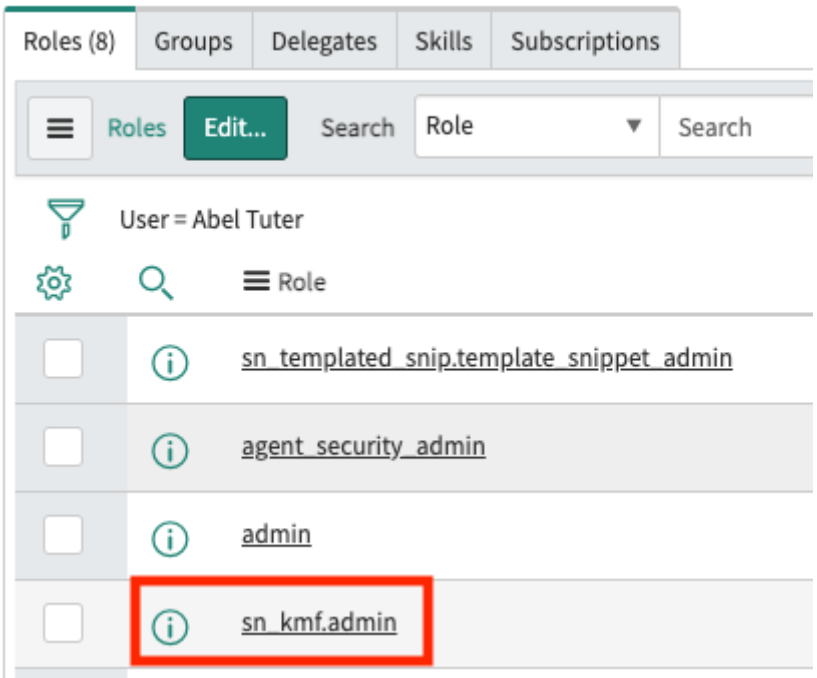

What to do next

If you have the KMF admin role, follow these steps for assigning other KMF roles:

1. Navigate to **User Administration** > **Users** and select the user you want to have another KMF role, such as KMF Cryptographic Manager.
2. In the Roles related list, click **Edit** and select the KMF roles you want to assign the users. All KMF roles start with `sn_kmf`.

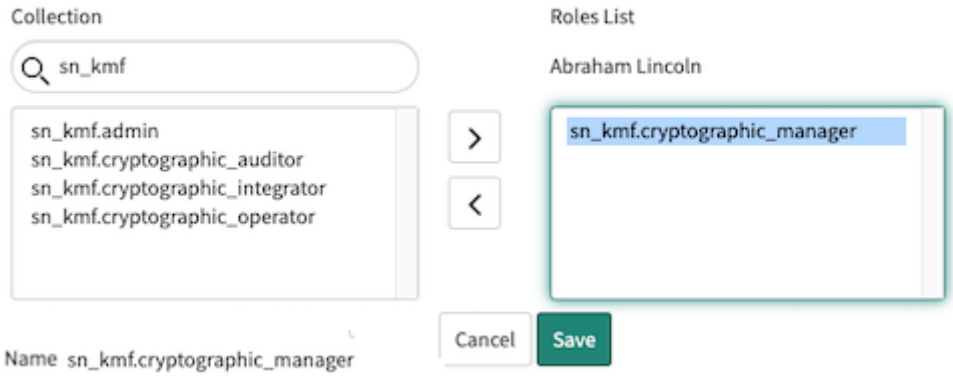

What to do next

To learn more about the available KMF roles, see Roles installed with Key Management Framework.

Configure field encryption settings to select key type

Configure your field encryption settings to use ServiceNow supplied keys or your own customer-supplied keys (CSK) for encryption on the ServiceNow AI Platform.

Before you begin

Customer-supplied keys are only supported with Field Encryption Enterprise.

Role required: sn_kmf.cryptographic_manager and security_admin

Procedure

1. Navigate to **All** > **System Security** > **Field Encryption Settings**.
2. From the Field Encryption Settings, select either **ServiceNow Generated Keys** or **Customer Supplied Keys** from the **Key Source** list.

Key source selection

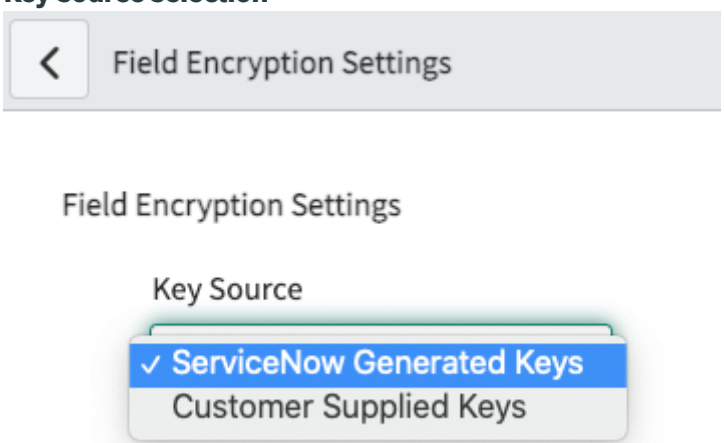

This option changes the `com.glide.encryption.cle_kmf.key_source` property to either **ServiceNow Generated Keys** or **Customer Supplied Keys**.

3. Select **Save**.

What to do next

- If you're using your own customer-supplied keys, see Using customer supplied keys with Field Encryption Enterprise.
- If you're using ServiceNow supplied keys, start creating your cryptographic module. See Create a cryptographic module.

Create a cryptographic module

Create a cryptographic module to define the mechanisms used for cryptographic operations. After you create the module, you create a cryptographic specification, where you define an algorithm for encryption and generates a key.

Before you begin

If you're supplying your own keys, go to Configure and upload your customer supplied key.

Role required: sn_kmf.cryptographic_manager

About this task

This procedure describes options that are available with KMF in the ServiceNow platform base system. Field Encryption Enterprise functionality is available only when the *com.glide.now.platform.encryption* plugin is active. See Activate Field Encryption for more information on obtaining Field Encryption Enterprise. See Create cryptographic module for Field Encryption.

Note: Cryptographic module [sys_kmf_crypto_module] records can't be deleted.

Procedure

1. Navigate to **All** > **Key Management** > **Cryptographic Modules** > **Create New**.
2. On the form, fill in the fields:

Cryptographic Module fields

Field	Description
Module name	Alphanumeric string to be referenced when running scripts.
Crypto spec template	Select the **Default** template to use to create the cryptographic module, as it contains mappings of supported algorithms for crypto specifications.
Default module access policy value	○ Rely on system default: ○ Reject ○ Track
Actual module access policy result	Reject or track, based on the default policy value or the value selected during the creation of the module access policy.
Name	Crypto module name prepended with application scope name.

Field	Description
Crypto module life-cycle state	Life cycle refers to the creation, use, and deactivation of a cryptographic module. Set to **Draft** initially during configuration. When using the module, set this field to **Published**. The Default template is automatically set to Published.

3. Select **Submit**.

> **Warning:**
>
> **For legacy encryption support users:**
>
> If you're using the non-enterprise version of Field Encryption, you're limited to five modules. If you've exceeded this limit, you receive the following warning:
>
> This insertion exceeds the number of published modules limit for Field Encryption entitled with the Subscription Product. The Enterprise subscription for Field Encryption is required for additional modules. Please reach out to your Account team.

After submitting successfully, your cryptographic module is listed in the Cryptographic Modules table. The system prepends the name with the scope to avoid conflict with other scoped applications. For example, if you created a module with the name `my_crypto_module` in the global application scope, the name is saved as `global.my_crypto_module`.

What to do next

Create a cryptographic specification

.

Create a cryptographic specification

After you create a cryptographic module, create a cryptographic specification to define the module algorithms.

Before you begin

Role required: sn_kmf.cryptographic_manager

Procedure

1. Navigate to **Key Management** > **Cryptographic Modules** > **All**.
2. Select the cryptographic module for definition to open the configuration options.
3. On the **Crypto Specifications** tab, select **New**.

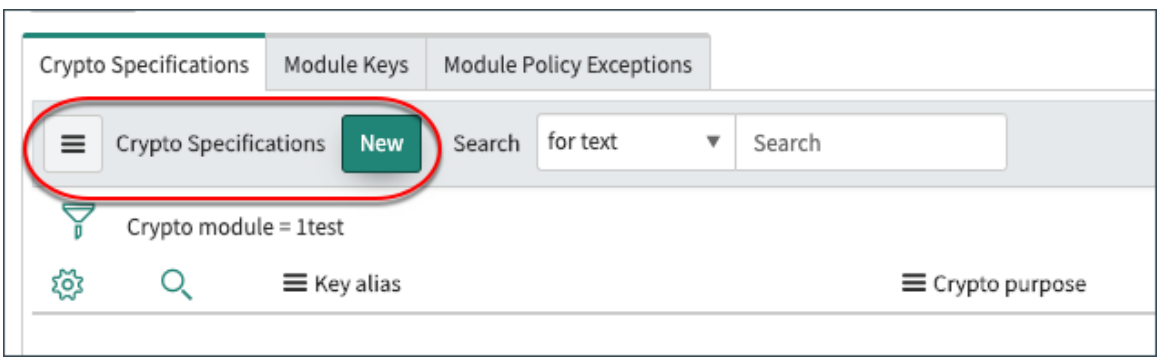

4. Complete the Algorithm Definition form.
   See Cryptographic specification overview for details.

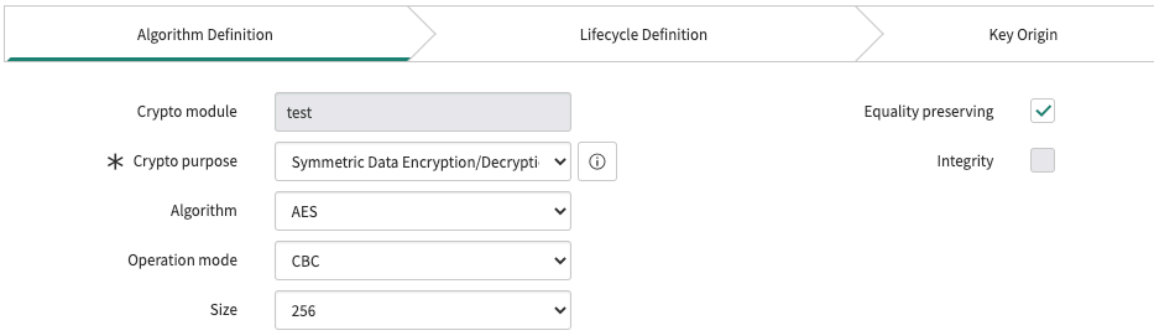

The algorithm definition screen opens. Select options for the key generation. Repeat this step to generate multiple keys for the selected crypto module.

Algorithm Definition fields

Field	
Crypto module	Read only. Name of the selected cryptographic module displays.
Crypto purpose	Select the purpose of this module. For example, you might use it for data encryption, signature generation, or key wrapping. The available algorithms adjust based on the selected crypto purpose. See Cryptographic specification overview for details.
Algorithm	Type of algorithm used to accomplish the crypto purpose. The algorithm also controls the key origin. Adjusts automatically based on the selected crypto purpose. Cryptographic specification overview for details.
Operation mode	This field may display based on the selected crypto purpose.
Size	Select the bit size.
Hash	This field becomes available based on the algorithm selected.
Equality preserving	Enables non-deterministic encryption. This option appears when you select **Symmetric Data Encryption/Decryption** with **AES** and in **Cipher Block Chaining (CBC)** mode. Selecting this option means that if the same data is encrypted again the encoded data is the same each time. Non-deterministic encryption doesn't support filtering a list of

Field	
	encrypted data using equality comparison operators.
Integrity	GCM operation mode provides Integrity.

5. Select **Next**.
   The crypto specification is listed on the Key Lifecycle table based on the algorithms selected.

What to do next
Perform one of the following operations:

- Select an entry in the Key Lifecycle table to define key lifecycle behavior. See Configure key lifecycle states for details to complete the lifecycle definition for the key.
- Select **Next** to create a cryptographic key. See one of the following tasks for key generation:
  - Generate a ServiceNow cryptographic key.
  - Configure properties for customer-supplied keys.
  - Import the wrapping / unwrapping key pair.

Configure key lifecycle states

After you have created a cryptographic specification, you can configure the lifecycle actions for the keys in your instance.

Before you begin
Role required: sn_kmf.admin

Procedure
1. Navigate to **Key Management** > **Cryptographic Modules** > **All**.
2. Select the cryptographic module to configure the lifecycle of a key.
3. Select a key alias on the Crypto Specifications tab.

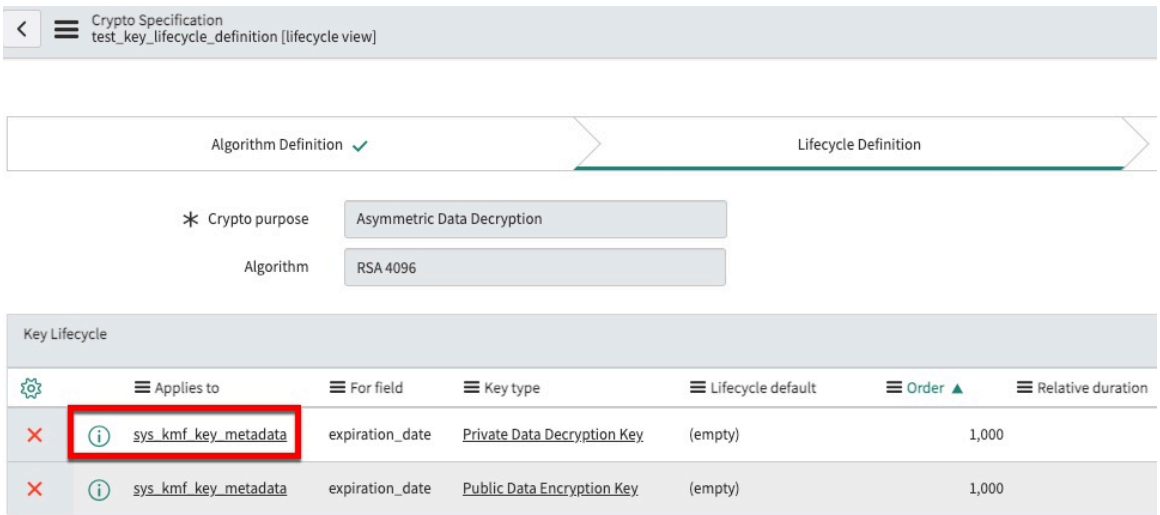

4. Select **Next**.
   The Field Lifecycle Template loads. Default Key Lifecycle values are created based on the selected algorithms for the defined cryptographic specification.

5. Select a Key Lifecycle from the **Applies to** column on the Lifecycle Definition step for the crypto specification.

Key Lifecycle fields

Field	Description
Applies to	Selected key that the lifecycle applies to.
For field	Select the type of control for the key that the lifecycle applies to. **Key lifecycle management "For field" values** 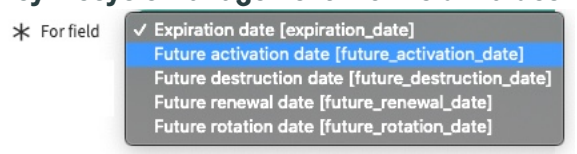
Type	Select if the valuation for the key lifecycle is a relative value or an absolute value: ○ **Relative**: Enter a value that depends on other data entries in the system, such as key generation, activation, and deactivation. ○ **Absolute**: Enter an exact value, such as a date.
Lifecycle default	Read only. Displays a value if set.
Order	Enter the sequence in which to process the key lifecycle state for the crypto specification.
Relative duration type	Duration of the lifecycle: **Years**, **Months**, or **Days**.
Relative duration	Number of years, months, or days the key is valid.
Relative operation	**Before**or **After**.
Relative to	Field the duration is relative to. Displays if a relative duration or operation is selected. 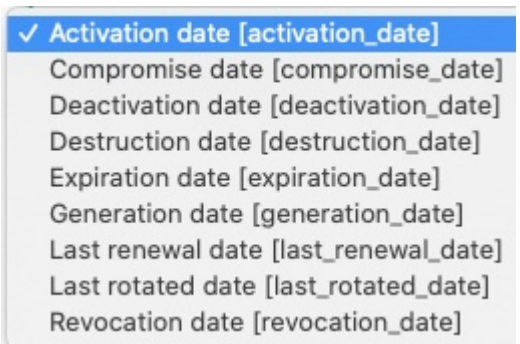

Generate a ServiceNow cryptographic key

Follow this procedure to upload and configure a ServiceNow cryptographic key to encrypt sensitive data.

Before you begin

Role required: sn_kmf.cryptographic_manager

About this task

Cryptographic managers have the choice to use ServiceNow supplied keys or their own customer-supplied keys (CSK) for encryption on the ServiceNow AI Platform with Field Encryption Enterprise. For information on CSK, see Configure properties for customer-supplied keys.

Procedure

1. Set field encryption settings to use ServiceNow Generated Keys.
   See Configure field encryption settings to select key type for details.
2. Navigate to **Key Management** > **Cryptographic Modules** > **All**.
3. Select the corresponding cryptographic module to open the Cryptographic Module details page.
4. Select the row for the key alias entry on the Crypto Specifications tab.
   If a key hasn't yet been generated, the key alias field is empty.
5. Select **Next** to navigate to the Key Origin tab of the Crypto Specification components.
   The Lifecycle Definition tab displays along with the Key Lifecycle table and can be reviewed or edited. See Configure key lifecycle states for details.
6. Select **Servicenow** in the Origin

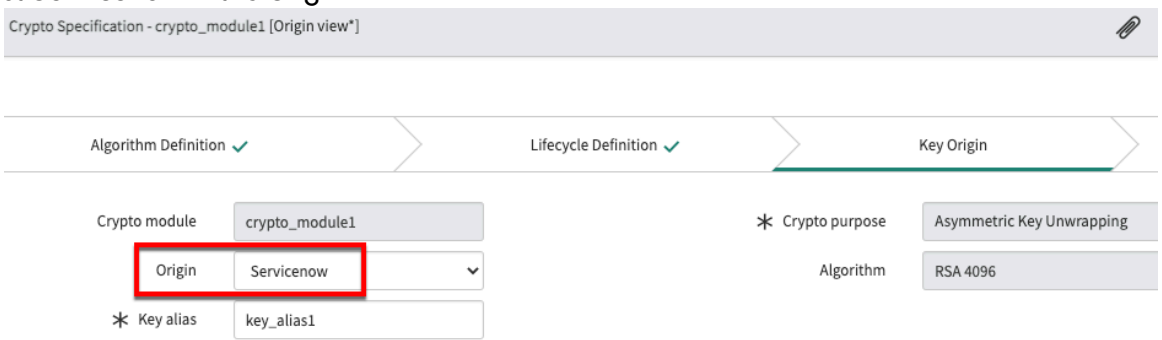

   field.
   This field varies based on the field encryption settings from Step 1 and the algorithm selected. To use an imported key, see Import the wrapping / unwrapping key pair. See Configure properties for customer-supplied keys if you're using your own key.
7. Enter a friendly name for the Key alias.
8. Select **Next** to move to the Key Creation tab.
9. Select **Generate Key**.
   After you generate the key, the Cryptographic Module form reloads displaying the cryptographic specification.
10. Select the **Module Keys** tab to view the keys.
    Secure information for the key is stored on the Module Keys tab along with the number of keys that exist for the cryptographic specification.

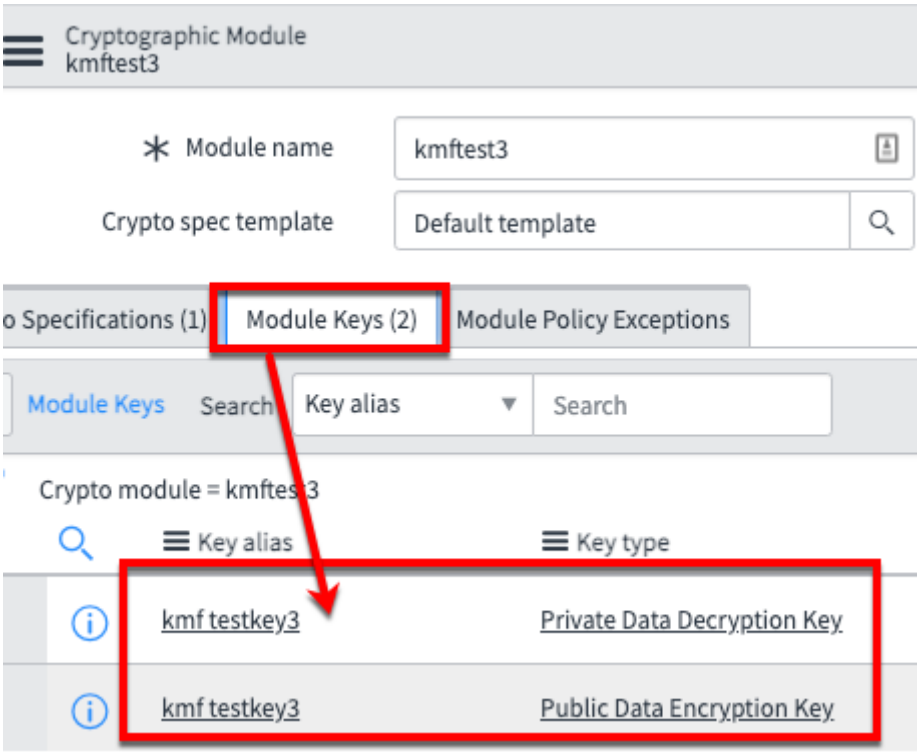

11. Select a key to perform key management actions.
See Key management actions for details.

Create a module access policy

Create module access policies to decide which users and scripts can access data encrypted by a cryptographic module.

Before you begin

Role required: sn_kmf.cryptographic_manager or sn_kmf.admin

About this task

Field Encryption supports role-based module access policies and additional configuration options become available with (CLE_Ent) functionality.

- Configure the specific cryptographic operation in module access policies for cryptographic modules that support symmetric operations. For instance, a user can be enabled to encrypt data but not decrypt data.
- Set a default module access policy value or according to a cryptographic module.
- Associate script versions where changes to the script are tracked and invalidate the script policy providing better security for script-type module access policies.

CLE_Ent functionality is available with a paid subscription. Refer to for supported features and options available with each offering. For more information, see Field Encryption Enterprise.

Note: The default behavior of the module access policies (MAPs) is Reject to help prevent any unauthorized access, unless explicitly declared in MAP records.

Procedure

1. Navigate to **All** > **Key Management** > **Module Access Policies** > **All**.
   If you don't create a cryptographic module configured for Symmetric Data Encryption/ Decryption, an auto-generated module access policy is created and listed in the table.
2. Select **New**.
   - Select **Specify purpose** to choose a **Crypto spec** and set the **Granular operation**.

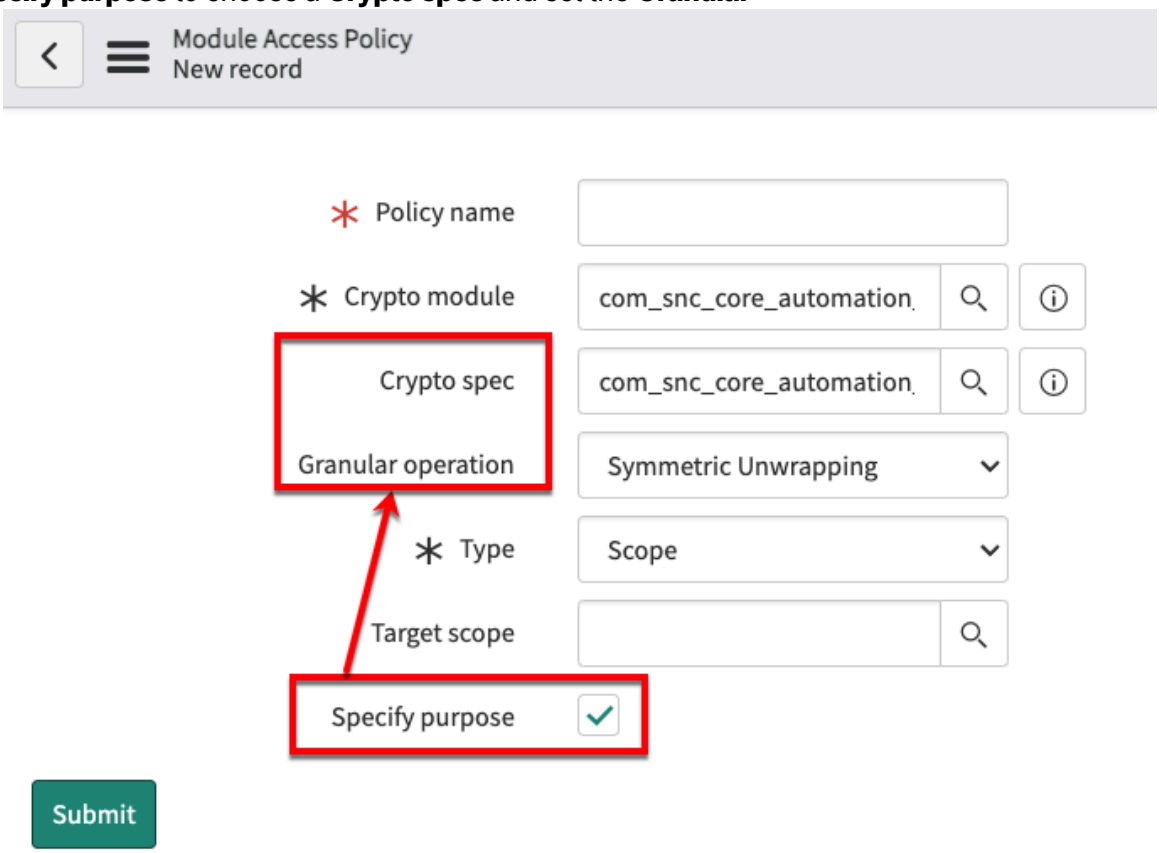

- With cryptographic specifications for symmetric data encryption/decryption and symmetric wrapping/unwrapping, the **Granular operation** field is available if you select the **Specify purpose** check box.

3. Complete the form.

   Module Access Policies fields

Field	Description
Policy name	Enter a name for the policy.
Crypto module	Select the search icon ( ) to select a module.
Crypto spec	Select or create cryptographic specification while generating the module access policy. This field becomes available when the **Specify purpose** check box is selected.
Granular operation	Select the cryptographic purpose for the cryptographic specification. The values available depend upon the type of cryptographic specification that is selected. See for details on crypto purposes.
Type	○ **Scope**: Controls access by the application scope. ○ **System user**: Allows access for system users to crypto modules. ○ **Script**: Control access by script. See for more information

Field	Description
	○ **Role**: Controls access by user role. ○ **Resource Exchange**: Control access using the Resource Exchange. See for more information. **Note:** Only Role type is supported with Field Encryption. All other types are available with Field Encryption Enterprise.
Target Scope	Field is visible as an identifier for the **Scope** type. Refers to the functionality for the policy. Select the applications from the search menu. **Note:** Target scope isn't supported and can only be set with Field Encryption Enterprise
Target Role	Field is visible as an identifier for the **Role** type. Role to which this policy applies.
Script Table Target Script	These fields appear when you select **Script** as the type. Field is visible as an identifier for the **Script** Type. Select a table to which this policy applies. Document to which this policy applies. Select the **Table name** and then the related document for the policy. The first time a script calls a cryptographic module, access to the module is denied, and the developer receives an error. This error gives the module owner the ability to grant or refuse access to the module.
Resource Exchange: ○ Crypto spec ○ Approval type ○ Target instance host	These options appear when you select as the type. Resource Exchange is supported by both KMF and by when the parent module is **column_level_encryption**. Select the crypto specification, One-time or Recurring, and the URL of the target instance. See for more information.
Impersonation	In role-based module access policies, users can access encrypted data using an impersonation session. When users, such as admins, impersonate other users, such impersonation-enabled module access policies are applied.
Specify purpose	Select to toggle the Cryptographic specification field as an available field for the policy.
Active	Select to activate the policy.
Result	Select one of the following:

Field	Description
	○ **StrictReject** rejects access under all circumstances. ○ **Reject** rejects users with the **Target Role** or **Target Scope** from accessing this cryptographic module unless another policy grants them access. ○ **Track** to permit access and monitor use of the module.

4. Select **Submit**.

Warning:

For legacy encryption support users:

If you're using the non-enterprise version of Field Encryption, you're limited to five modules. If you have exceeded this limit, you receive the following warning:

This insertion exceeds the number of published modules allowed for Field Encryption entitled with the subscription product. The Enterprise subscription for Field Encryption is required for additional modules. Please reach out to your Account team.

5. Select the policy name associated with the cryptographic module that you want to examine. Using Script type module access policy:

   A module access policy is auto-generated based on the default access setting when the script is run. The module name is preceded with `AutoGen-`. For example, the `Module-TestPolicy` module is listed as `AutoGen-Module-TestPolicy` in the Policy name column.

   The Cryptographic Caller Policy form lists the caller policy that you selected. The Target Scope field specifies the scope of the script attempting to use the module. See for additional information.

   Note: A maximum of five module access policies are permitted with Field Encryption. See for configuration options.

Create a granular role module access policy for symmetric encryption

Create a granular role module access policy (MAP) to secure data while permitting users not assigned to a specific MAP to submit forms on a public record.

Before you begin

Role required: security_admin

Procedure

1. Navigate to **All** > **Key Management** > **Module Access Policies** > **All**.
2. Select **New**.
3. Enter a name for the policy.
4. Select the **Crypto module** that you'll be granting access to.

5. Select the appropriate **Target role**.
6. Select the Specify purpose box and enter the purpose details in the newly displayed fields as follows:

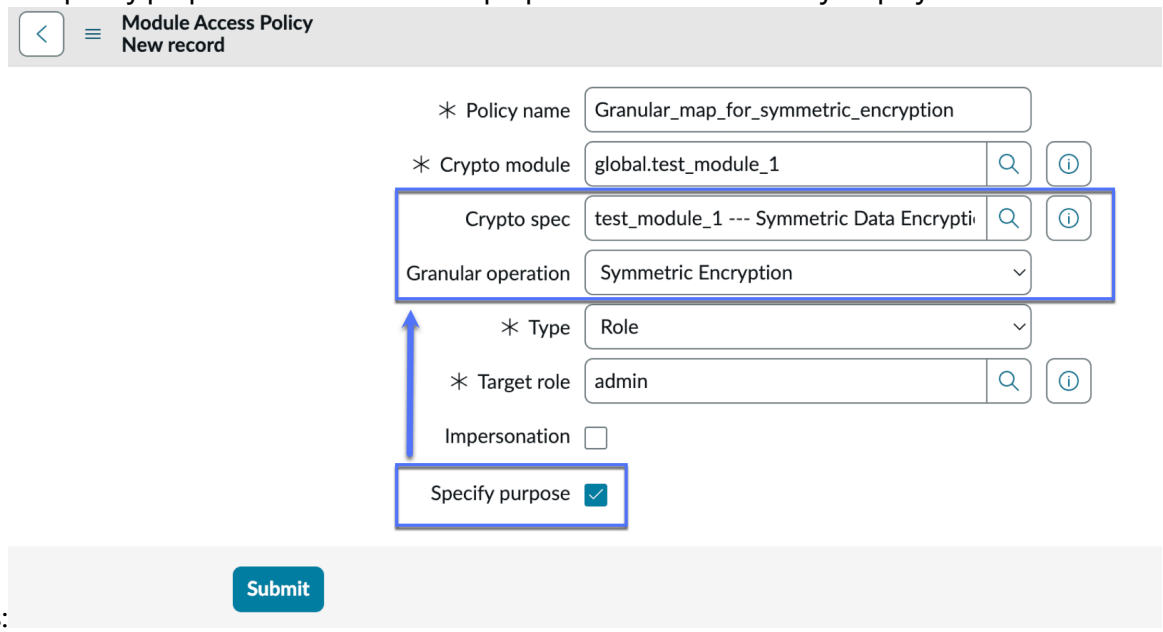

- **Crypto spec**: Select the crypto spec to be used.
- **Granular operation**: Select Symmetric Encryption from the drop-down list.

7. Select the **Active** box.
8. For **Result** select **Track**.

Result

The granular role MAP enables a user that is not assigned to a MAP to submit public record forms while securing data.

Create a cryptographic module life-cycle policy

Create a cryptographic module life-cycle policy to place limits on cryptographic modules, such as how long the key is good for. Create policies to safeguard cryptographic modules by limiting their exposure.

Before you begin

Role required: sn_kmf.cryptographic_manager

About this task

A cryptographic module life-cycle policy is an instance-level policy. The more exposure that a cryptographic key has, the more likely it can be compromised. Safeguard keys by limiting how long the keys can be used and who can use them.

The following features govern cryptographic modules:

- Instance policies set boundaries for the instance. For example, if you specify in an instance policy that the expiration date should never be more than two years after the activation date, you can't use the life-cycle rules to set an expiration date five years after the activation date.

- Instance life-cycle templates enable you to set different policies for different keys. Templates offer default life-cycle rules for cryptographic modules so that they don't have to be re-created for every module. For example, you can set different expiration dates for symmetric data encryption keys than for public key wrapping keys.

- Life-cycle rules affect the keys directly. For example, if you specify in the life-cycle rules that the expiration date should be two years after the activation date, keys will expire two years after the activation date.

Procedure

1. Navigate to **All** > **Key Management** > **Lifecycle Policies** > **Instance Policies**.
2. Select **New**.
3. Complete the form.

   Cryptographic Life-cycle Policies fields

Field	Description
Applies to	Read only. The key that the life cycle applies to.
Active	Select to activate the policy.
Policy condition	Conditional statements that specify when to activate, renew, deactivate, and destroy the cryptographic module.
Result	**Reject** to revoke access to the cryptographic module, or **Track** to permit access and monitor use of the cryptographic module.

What to do next

If you want to add exceptions to this life-cycle policy at the module level, see Create module life-cycle policy exceptions.

Create module life-cycle policy exceptions

Create a module policy exception to change the life-cycle policy of a key only for a specific on one instance.

Before you begin

Role required: sn_kmf.cryptographic_manager and sn_kmf.admin

Exceptions apply only to that module and not to the entire instance. For example, an administrator configured symmetric keys to be limited to one year at the instance level. An exception can be made at the module level to be two years.

Procedure

1. Navigate to **All** > **Key Management** > **Cryptographic Modules All**.
2. Select the cryptographic module that will use the policy exceptions.
3. In the Cryptographic Module table, select the **Module Policy Exceptions** tab.
4. Select **New**.
5. Complete the form.

Module Policy Exceptions fields

Field	Description
Crypto Module	Name of the module selected. This field is read only.
Applies to	Specified key is auto-populated.
Key Type	Key type that the exception policies are related to. **Note:** You may only select a single key type, but multiple exception policies can be created per cryptographic module.
Policy condition	Customizable condition which determines when the policy exception applies.
Result	The result that occurs when the condition in the **Policy Condition** field is met. ○ **Reject** rejects usage of the key. ○ **Track** allows the key to be used.

6. Select **Submit** to be returned to the Cryptographic Module table.

Key Management Framework Reference

The Key Management Framework (KMF) API/UX lets you fully customize and manage how cryptographic operations are performed on your ServiceNow instance. The ServiceNow Key Management Framework provides a secure and comprehensive interface for instance-side cryptographic key management services.

Key Management Framework key life-cycle states

KMF supports several cryptographic key life-cycle states through the enforcement of specific allowable actions. For example, only keys that are in the active state can be used fully for their intended cryptographic purpose. The following table provides further detail on the varying key life-cycle states.

Roles installed with Key Management Framework

The Key Management Framework (KMF) introduces specific roles for cryptographic module and key management-related configurations.

Module access policy visualization

Use module access policy visualization to view all relevant cryptographic module information on a single UI page.

Module access policy debugger

Use the module access policy debugger to review logging information and understand why your users are or aren't granted access to an encryption context.

Encryption and Key Management subscription bundle

With key management, Field Encryption is upgraded at no additional charge to the highly configurable encryption modules. You also have the option to upgrade to the unlimited-use license. Subscribe to the new encryption entitlement bundle, Platform Encryption, which includes Field Encryption Enterprise and Cloud Encryption.

Key Management Framework key life-cycle states

KMF supports several cryptographic key life-cycle states through the enforcement of specific allowable actions. For example, only keys that are in the active state can be used fully for their intended cryptographic purpose. The following table provides further detail on the varying key life-cycle states.

Key life-cycle state or action	Description
Active	The active key is used to generate new content, such as encrypting or signing. There can be only one active key for a given cryptographic specification in a cryptographic module.
Compromised	Compromised keys can't be used to generate new content, such as encrypting or signing, but may still be used to identify the purpose of existing content, such as decryption or verification. Several keys can exist in the compromised state for revocation in a given cryptographic specification in a cryptographic module. Any active or suspended key can be moved to a compromised state.
Deactivated	Any active key can be deactivated. There could be several keys in a deactivated state for a given cryptographic specification in a cryptographic module. For example, when the key is rotated, the current active key is deactivated. Deactivated keys cannot be used to generate new content, such as encrypting and signing, but may still be used to identify purposes of existing content, such as decryption or verification. **Note:** Compromised and revoked keys are treated as deactivated keys.
Destroyed	When a key is destroyed key material is permanently removed and can no longer be used for any cryptographic purpose. Any deactivated key can be destroyed using lifecycle automation when it hasn't been used in the configured designated time frame. There could be several keys in a destroyed state for a given cryptographic specification in a cryptographic module. **Warning:** Data associated with a destroyed key can no longer be accessed, therefore extreme caution should be exercised when performing a destroy key action.
Generated	Multiple keys can exist in the generated state for a given cryptographic specification in a cryptographic module.

Key life-cycle state or action	Description
	A generated key can be moved to an active state when no active key exists for the given cryptographic specification. The first key generated is automatically set to active. **Note:** If the choice is to generate a new key, then a new key is generated and made active even though there are keys in a generated state for the given cryptographic specification.
Renewed	An active key that has an expiration date can be renewed any number of times to extend the life-cycle period of the key. **Note:** The difference between the activation date and expiration date is calculated and the expiration date is postponed by that duration from the current day.
Resume	The UI action is available on suspended keys to move them back to an active state when no other active key exists for the given cryptographic specification.
Revoked	Any active or suspended key can be moved to the revoked state. Revoked keys can't be used to generate new content, such as encrypting or signing, but may still be used to identify the purpose of existing content, such as for decryption or verification. Several keys in a revoked state may exist for a given cryptographic specification in a cryptographic module.
Rotated	Key rotation results in deactivating the current active key and making another key active. Select the new active key from the following: • Generation of a new key. • Point to an existing imported key. Any active key can be rotated.
Suspended	There could be several keys in the suspended state for a given cryptographic specification in a cryptographic module. When the key is suspended, the key can be resumed and reassigned to an active state when no other active key exists for that cryptographic specification.

Roles installed with Key Management Framework

The Key Management Framework (KMF) introduces specific roles for cryptographic module and key management-related configurations.

Important:

KMF roles are required to use the Key Management Framework. Users without KMF roles are not able to access lists, tables, and modules used to configure key management.

To assign the KMF admin role, you must have the *admin*, *security_admin*, and *sn_kmf.admin* roles. Use the KMF admin role to assign other KMF roles. For details on assigning KMF roles, see Assign Key Management Framework roles.

The *sn_kmf.admin* role is also required to modify any group record that includes the *sn_kmf.cryptographic_manager* role. This requirement applies to all updates to the group record, not only to role assignment operations.

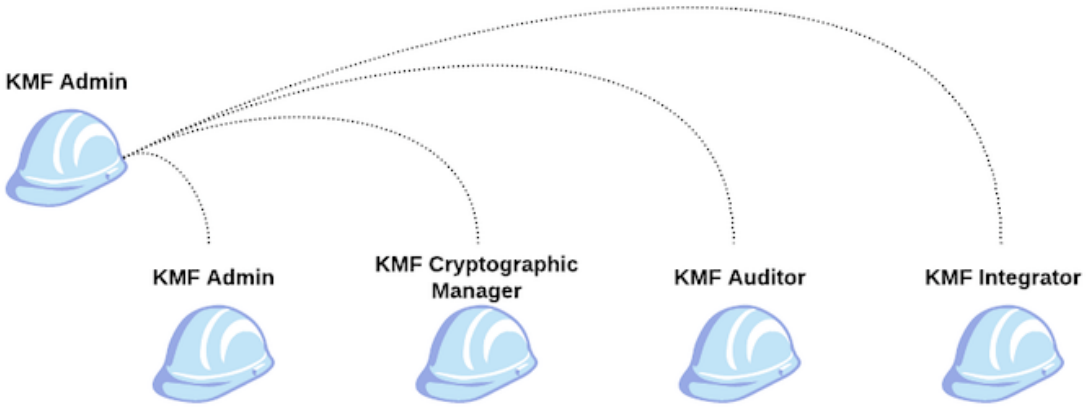

KMF admin [sn_kmf.admin]

Assigns roles to other users to perform operations around the ServiceNow Key Management Framework.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Important: Avoid granting an admin role when more specialized roles are available.

- This role is assigned via the process shown in Assign Key Management Framework roles.
- Users with this role must also have the admin and security_admin
- You must have this role to assign#KMF roles, and in addition can perform all the capabilities of the KMF cryptographic manager.

KMF cryptographic manager [sn_kmf.cryptographic_manager]

Create, read, and update operations on cryptographic modules (association of keys to cryptographic usage and algorithm configurations) and module access policies. Also, KMF

cryptographic managers can perform key management (generate, rotate, revoke) and life cycle operations.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

This role can only be assigned to a user by a KMF admin.

KMF cryptographic auditor [sn_kmf.cryptographic_auditor]

View cryptographic module information, key metadata, and life cycle-related details, as well as module access policy (MAP) information.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

This role can only be assigned to a user by a KMF admin.

KMF cryptographic integrator [sn_kmf.cryptographic_integrator]

Integrate Key Management Framework with external keystores or systems.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

This role can only be assigned to a user by a KMF admin.

KMF cryptographic operator [sn_kmf.cryptographic_operator]

Access part of the ServiceNow Key Management Framework key lifecycle: renewal, rotation, revocation.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

None.

Assign KMF roles

Assign KMF roles to admins, who in turn can assign other KMF roles.

Before you begin
Role required: admin and security_admin

You must elevate to the security_admin role before assigning the KMF admin role. For instructions, see Elevate to a privileged role

Procedure

1. Elevate to the security admin role.
2. Navigate to **User Administration** > **Users** and select the user you want to be the KMF admin.
3. Verify that the user already has the admin and security_admin roles.
   If not, select **Edit** under the Roles related list and add **admin** and **security _admin**.
4. Navigate to **System Security** > **Key Management Administration**.
5. Select the user that you want to be KMF admin in the **Available Users** column and move them to the **Selected User(s)** column.

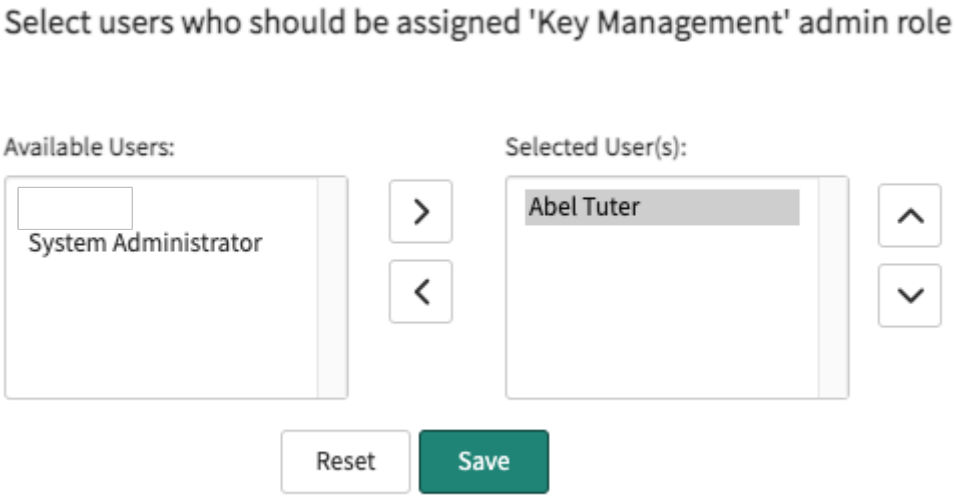

6. Select **Save**.
7. Navigate to **User Administration** > **Users** and select the user you just gave the sn_kmf.admin role to.
   The user has the sn_kmf.admin role in the **Roles** related list, and can assign other KMF roles.

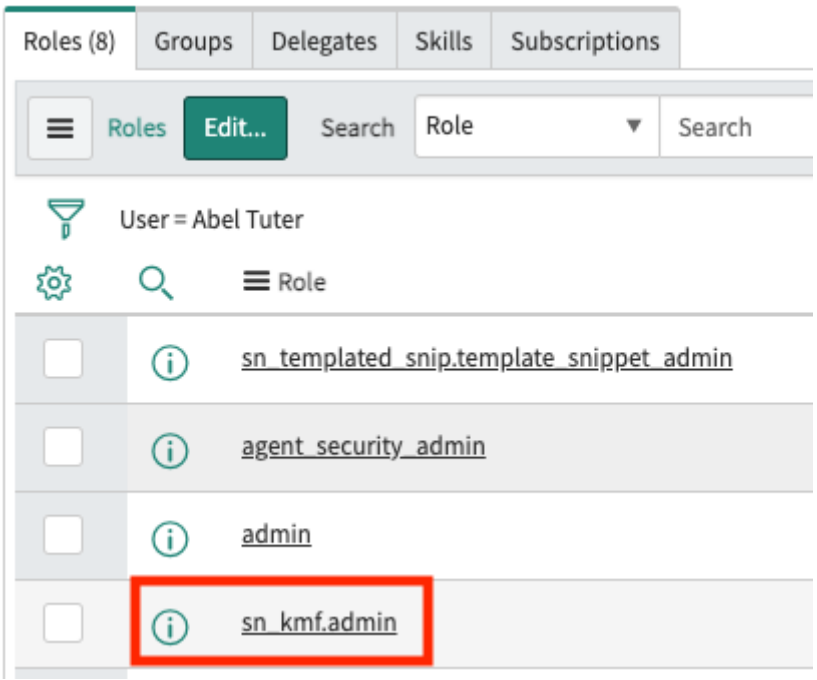

What to do next

If you have the KMF admin role, follow these steps for assigning other KMF roles:

1. Navigate to **User Administration** > **Users** and select the user you want to have another KMF role, such as KMF Cryptographic Manager.
2. In the Roles related list, select **Edit** and select the KMF roles you want to assign the users. All KMF roles start with `sn_kmf`.

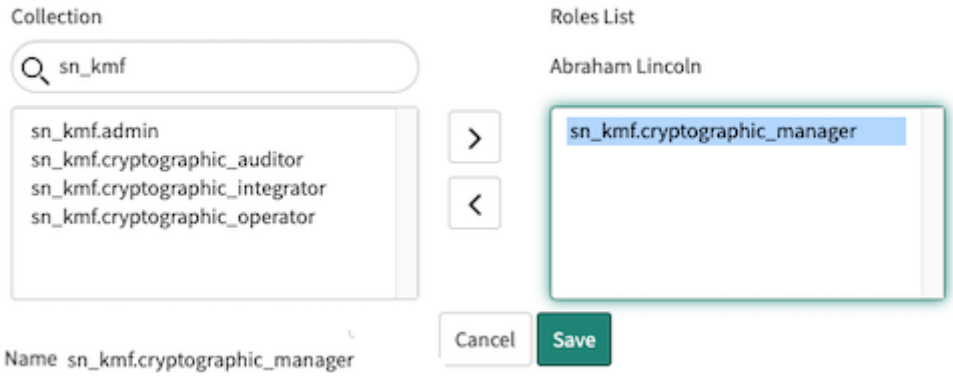

Module access policy visualization

Use module access policy visualization to view all relevant cryptographic module information on a single UI page.

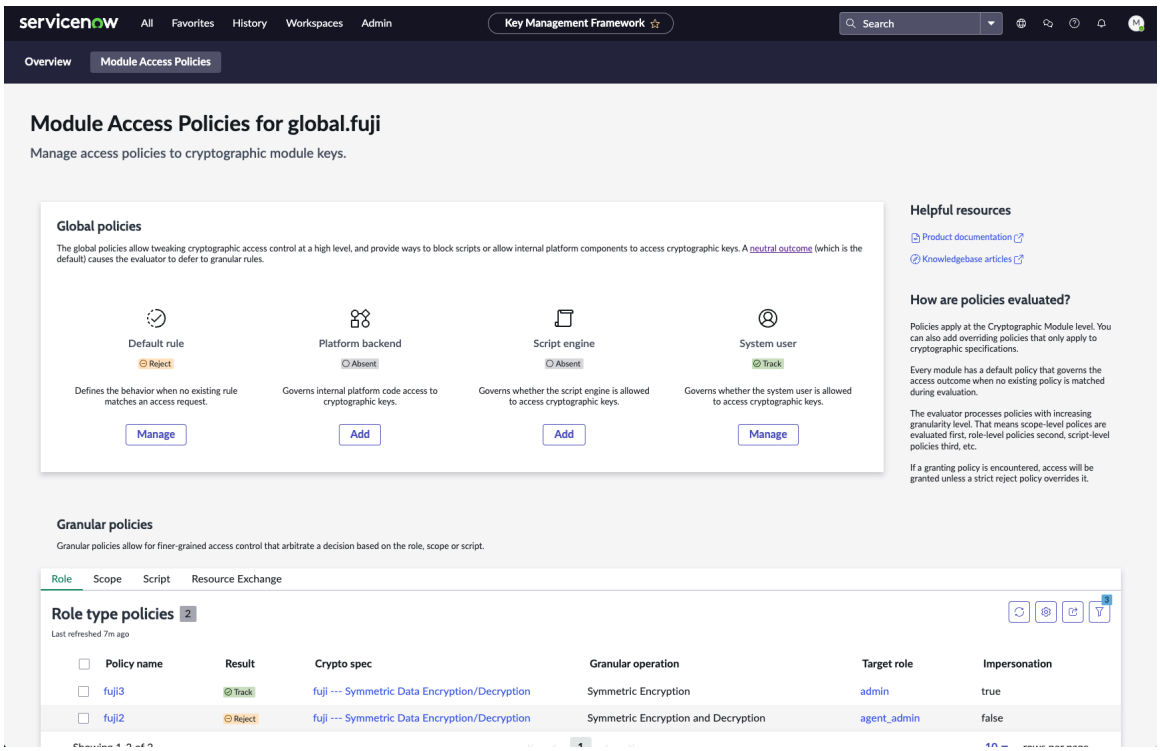

Key Management Framework admins and cryptographic managers can use the module access policy UI page to view all access control mechanisms related to a single cryptographic module. Use the information collected on this UI page to determine who has access to encrypted information on your instance.

Users with the *`sn_kmf.admin`* or *`sn_kmf.cryptographic_manager`* roles can access the module access policy visualization UI page by navigating to **All** > **Key Management** > **Cryptographic Modules** > **All**.

Results Labels

Module access policies contain a **Result** field, which determines whether to grant access to the selected cryptographic module. The UI page displays a label on elements on the UI page based on the value of that field.

UI label	Result field value	Definition
	*Track* or *Allow*	Access is granted to all users, including scripts.
	*Reject*	Access is denied unless a track module access policy is found.
	*StrictReject*	Access is denied.
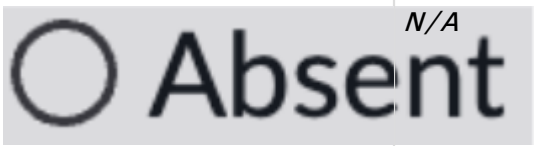	*N/A*	The module access policy doesn't exist on the instance. Access is denied to all.

Global policies

Use the **Global policies** section to review the module access policies that control platform-level access.

Select the **Manage** button below any of the policies to navigate to that policy record. If the policy doesn't exist, an **Add** button appears below that entry. Select the **Add** button to navigate to a new policy record where you can define the policy.

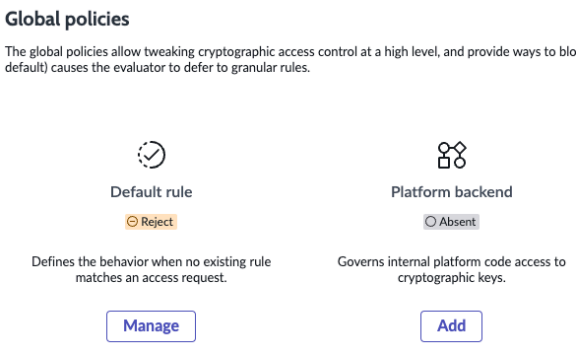

Policy	Definition
Default rule	The default rule policy defines the behavior when no existing rule matches an access request.
Platform backend	The platform backend policy governs internal platform code access to cryptographic keys.
Script engine	The script engine policy governs whether the script engine is permitted to access cryptographic keys.

Policy	Definition
System user	The system user policy governs whether the system user is permitted to access cryptographic keys.

Helpful resources

Use the **Helpful resources** section to find links to product documentation, relevant knowledge articles, and a brief description on how module access policies are evaluated on the platform. For a deeper look into how module access policies are evaluated, see Module access policy debugger.

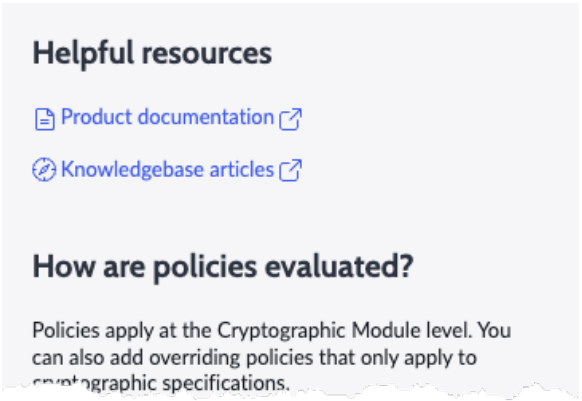

Granular policies

Use the **Granular policies** section to view lists of module access policies, separated by policy type. Use the tabs above the list to select a policy category to display.

- Role
- Scope
- Scope and Domain (if Domain Separation is active)
- Script
- Resource exchange (if the cryptographic module is a Password2 or Field Encryption submodule)
- Identity (if Secrets Management Enterprise is active)

By default, the each list displays only active policies. Use the filter icon to change the default filter for the list.

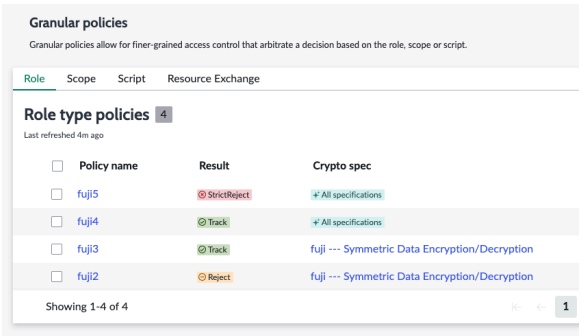

Users with access

Use the **Users with access** section to see a list of all users that have access to the selected cryptographic module. The list is grouped by user, as single users can posses multiple roles that grant access to a cryptographic module.

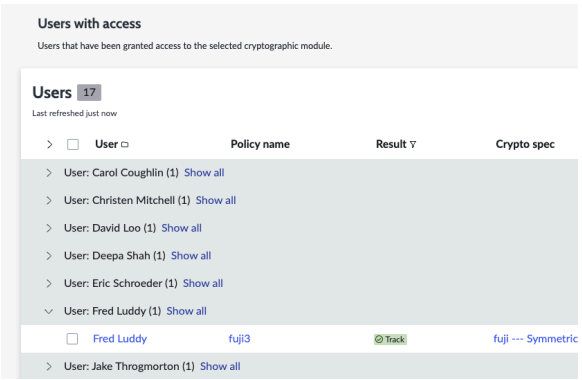

Module access policy debugger

Use the module access policy debugger to review logging information and understand why your users are or aren't granted access to an encryption context.

Module access policies (MAPs) define instance-level controls for access to cryptographic modules. Callers (for example, a user or script) require explicit access to use a cryptographic module for encryption and decryption. Use the debugger to see which policies are evaluated when a caller attempts to access a cryptographic module. You can also use the debugger and learn why access is or isn't being granted.

This flowchart shows how your instance evaluates requests for access to a cryptographic module.

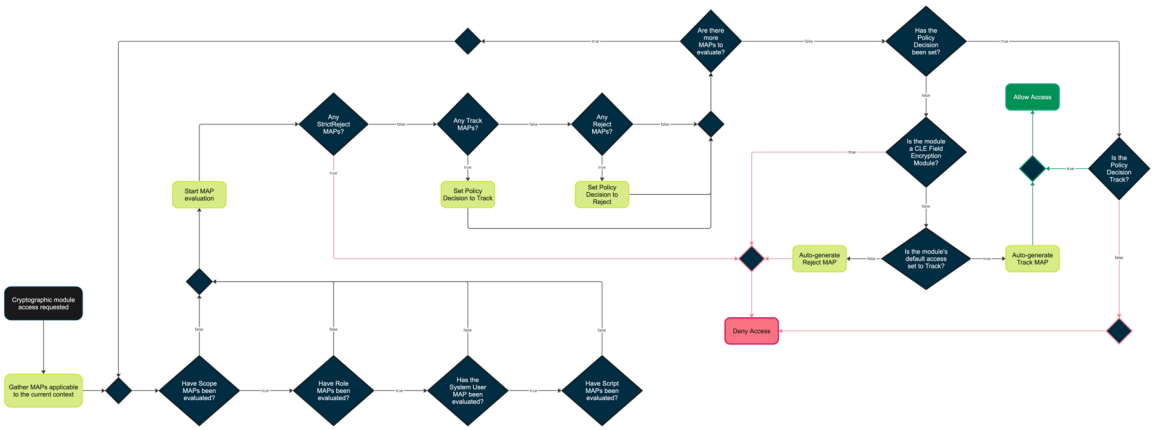

Control access to the debug logs

Access to the module access debug logs is determined by role. Users with the *sn_kmf.admin* and *sn_kmf.cryptographic_manager* roles have access to the debugger. Grant access to other roles using the *glide.kmf.module_access_policies.debugger.authorized.roles* system property. The value of this property is a comma-separated list of roles that access the debug logs.

Enable or disable the debugger

To enable debug logging messages for module access policies, navigate to **All** > **Diagnostics** > **Session Debug** > **Debug Module Access Policies** > **.**

When you're finished debugging, you can disable the logging messages by navigating to **All** > **Diagnostics** > **Session Debug** > **Disable All** > **.**

Access the logs

After enabling debugging, navigate to a page that triggers a MAP evaluation to view the MAP debug logs. Debug messages appear at the bottom of the page.

Tip: You can use impersonation to troubleshoot access for other users. For details on impersonation, see Impersonating users. To view the debug logs from the perspective of another user, make sure that your module access policies with the *role* type have the **Impersonation** field set as **true**.

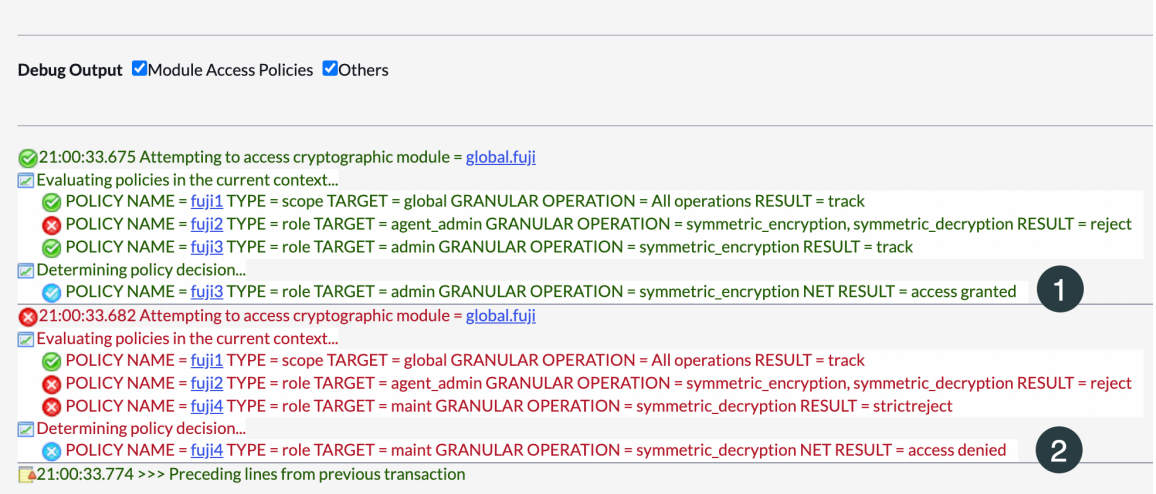

In this example, a caller invokes two access requests to the `global.fuji` cryptographic module. A symmetric encryption, which is granted, and a symmetric decryption, which was denied.

Understanding log entries

Debugging information is structured using this format.

1. This first line displays the cryptographic module receiving the access request.
2. The lines between the first and last line displays the evaluated MAPs in the order that they were evaluated, and includes their name, type, target, granular operation, and result.
3. The last line displays the Policy Decision (if applicable) and the net access result for the caller (whether the caller is granted access).

Each line starts with an icon that indicates its message type.

Message icons

Icon	Message type
	Informational message
	Module access policy grants access
	Module access policy denies access
	Caller is granted access

Message icons (continued)

Icon	Message type
	Caller is denied access
	No module access policy to evaluate

Debug log examples

Access granted message

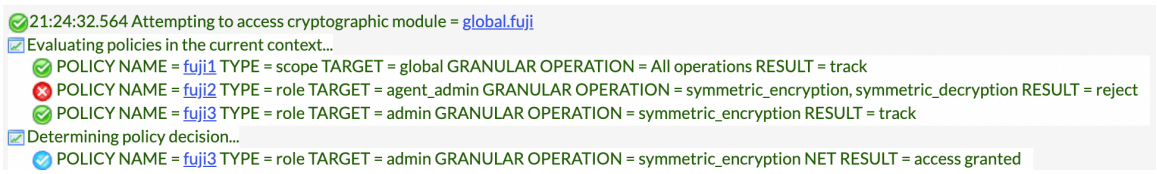

Access denied message

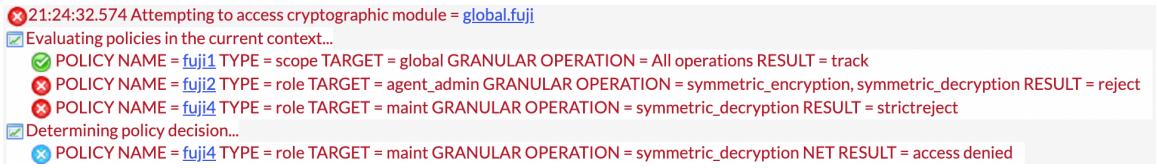

Access denied (No module access policies to evaluate

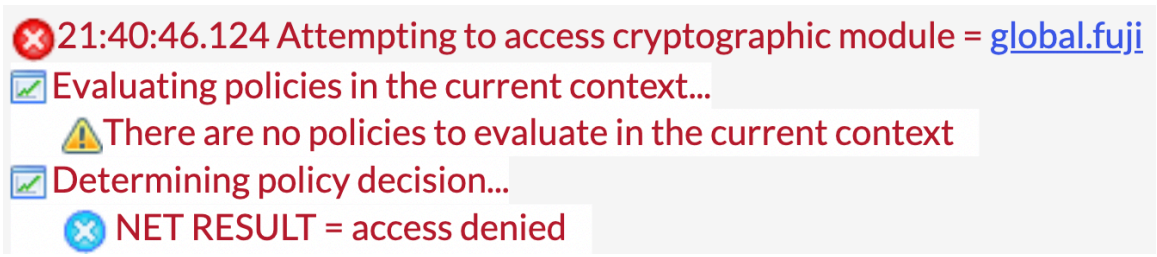

Access denied (insufficient privileges)

Encryption and Key Management subscription bundle

With Key Management, Field Encryption is upgraded at no additional charge to include highly configurable encryption modules. You can also optionally upgrade to the unlimited-use license. Subscribe to the new encryption entitlement bundle, Platform Encryption, which includes Field Encryption Enterprise and Cloud Encryption.

Field Encryption features

Field Encryption with encryption modules is included for free in your instance and includes NIST 800-57 key management.

ServiceNow Platform Encryption group features

The Platform Encryption group adds the following features and offerings:

- Encryption.
- Cloud Encryption with Key Management.

Additional information

To learn more about key management, see Exploring the Key Management FrameworkExploring the Key Management Framework.

Key management actions

One of the core features of KMF is to provide the capability to manage keys, such as revoking or rotating keys. KMF properly secures sensitive data with the most up-to-date encryption materials and life cycle operations.

The following table provides a summary of the key life cycle operations and management actions. The cryptographic module purpose is applied to the data with the cryptographic module configuration and has no impact on data.

Key management action	Description
Generate key	Generates a new key for the given cryptographic module. A first generated key is set to *active*.
Rotate key	Deactivates the current key and generates a new one. The new module key is set to current (active).
Revoke key	Marks the current key and life cycle state as revoked. The cryptographic module auto-generates a new key on new data and sets the key status to active. Revoked means that the key is no longer used for encryption. However, it can still be used for decryption. You can't destroy a key.
Suspend key	Marks the current key as suspended. Manually resume the suspended key or revoke the suspended key to generate a new module key before using the cryptographic module again.
Resume key	Marks a suspended key as the active key.
Renew key	Extends the life of the current key. The **Renew** button becomes available under the following circumstances: • You're assigned the cryptographic manager role. • The life-cycle state is marked to either **Active** or **Renewed**. • An expiration date is set in the module life cycle definition.

View and manage keys

Review the status of any key to determine further key action, such as when to renew, rotate, suspend, deactivate, or destroy a current key.

Before you begin

Role required: sn_kmf.cryptographic_manager

Procedure

1. Navigate to **All** > **Key Management** > **Cryptographic Modules** > **All**.
2. Select a cryptographic module.

The Cryptographic Module <module-name> form appears.

3. On the Module Keys tab, select the key alias to review the key status on the life cycle <key name> form.
4. Review the form, as all fields are read-only.

Cryptographic life cycle Key fields

Field	Description
Generation date	Displays the date when the key was generated.
Activation date	Displays the date when the key was activated.
Last renewal date	Displays the date when the key was last renewed.
Last rotated date	Displays the date when the key was last rotated.
Deactivation date	Displays the date when the key was deactivated.
Destruction date	Displays the date when the key was destroyed.
Key life-cycle state	Displays the key life-cycle state.
Future activation date	Displays the future key activation date.
Future renewal date	Displays the future key renewal date.
Future rotation date	Displays the future date for key rotation.
Future destruction date	Displays the future date when the key is destroyed.
Expiration date	Displays the date when the key expires.

5. To perform an action on the key, select one of the following to take effect immediately:

Note: These actions can only be performed for custom Cryptographic Modules, not for base system Cryptographic Modules.

- **Revoke Key**: Select to deactivate the key and generate a new key. Enter the reason why you're revoking the key.
- **Rotate Key**: Select to deactivate the current key and to generate a new key in its place. The new key is listed in the Module Key table and the Key Version number increments by 1. See for details.
- **Suspend Key**: Select to deactivate the current key.
- **Resume Key**: Select to mark a suspended key as the active key. This option is only available after the active key has been suspended.

Rotate keys

For increased security, you can rotate your cryptographic keys on a pre-determined schedule. Key rotation is when you retire an encryption key and replace that old key by generating a new cryptographic key.

Before you begin

Role required: sn_kmf.cryptographic_manager

About this task

Encryption modules, unlike encryption contexts, support a rekey of records for re-encryption with a new key. The following demonstrates how to perform a key rotation operation manually on a cryptographic module.

Procedure

1. Navigate to **Key Management** > **Cryptographic Modules** > **All**.
2. Select the cryptographic module for key rotation.
3. On the **Module Keys** tab, select the Active key.

Select the active key

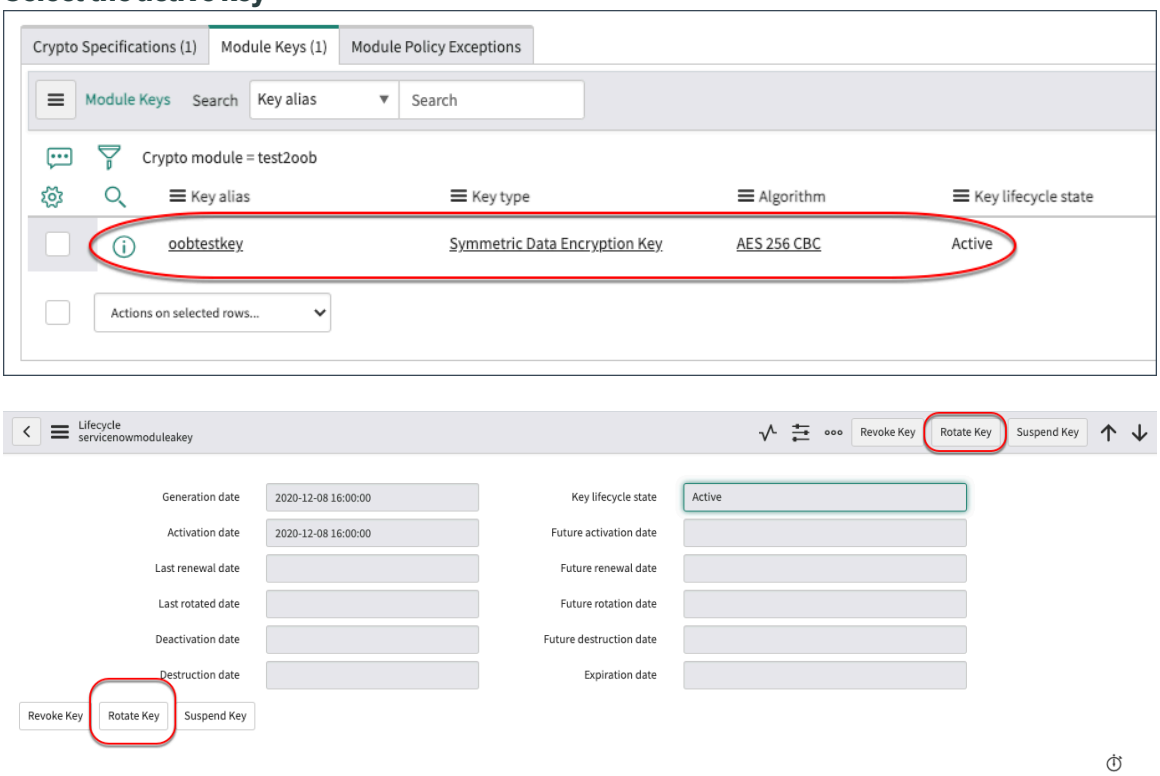

4. Select **Rotate Key**.
   The key life-cycle state changes to "Deactivated." The **Last rotated date**, **Deactivation date**, and **Key version** fields update.
5. Return to **Cryptographic Module** > **Module Keys**.

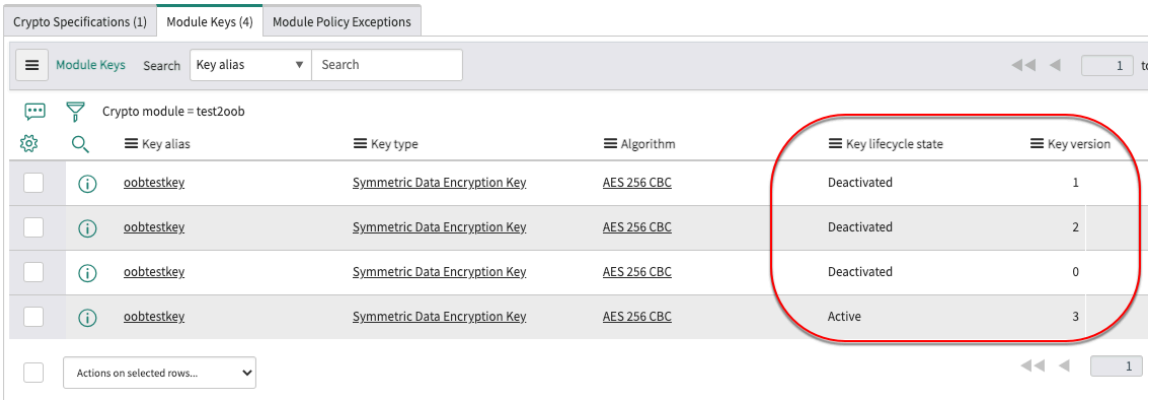

There's an extra module key listed in the table. The newly rotated key becomes "Active" and the last key is "Deactivated."

Import a key from a web service

Securely upload an external customer key onto your instance using import a key from a web service (for example the key REST API). Both symmetric and asymmetric public keys can be imported into a targeted KMF cryptographic module.

The key to be imported (the target key) must be encrypted with a wrapping key before it's uploaded into the instance's target cryptographic module. This wrapping key is the public component of a public/private key pair, which must be present on the instance. The key is a pre-requisite before the wrapped target key can be uploaded via Import From Web Services.

These two separate procedures (importing the wrapping key pair and importing the wrapped target key from a web service) are detailed in the following documentation. This key pair must be generated and uploaded to be available in the instance's internal Key Import cryptographic module.

Note: This example uses OpenSSL for key and certificate generation and the Postman API test tool to show REST API use. Substitute other comparable tools based on your company requirements.

Import the wrapping / unwrapping key pair

Configure Key Management Framework import settings before importing a key.

Before you begin

Role required: sn_kmf.cryptographic_manager

About this task

This example uses OpenSSL for key and certificate generation. Substitute other comparable tools based on your company requirements.

Procedure

1. In your local environment, use the terminal to create a certificate.
   For example: `openssl req -x509 -sha256 -nodes -days 365 -newkey rsa:4096 -keyout wrapping_private.key -out wrapping_public.crt`

   This certificate is a public component that contains a key. The certificate is used to wrap an AES symmetric key.

2. In your local environment, use the terminal to create a keystore containing public cert (with the wrapping key), and private unwrapping key.
   For example:`openssl pkcs12 -export -in wrapping_public.crt -inkey wrapping_private.key -name "wrapping_key_alias" -out wrapping_keystore.p12`
3. On your instance, navigate to **All** > **Key Management** > **Import Settings** > **Key Import Settings**.
4. In the Algorithm Definition section, verify the **Crypto Purpose** is set to *`Asymmetric Key Unwrapping`*.

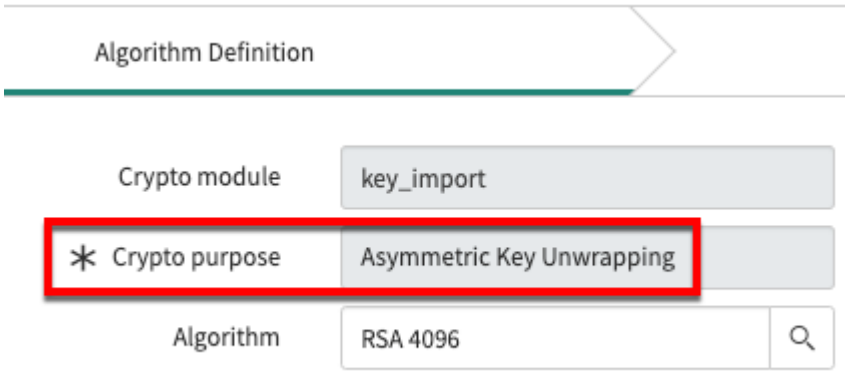

5. Select an appropriate algorithm that aligns with asymmetric key material for the imported keystore.
   See Cryptographic specification overview for additional information.
6. Select **Next**.
7. In the **Lifecycle Definition** section, select **Next** to continue.
8. In the **Key Origin** section, select either **Import from PKCS12** or **Import from BCFKS** in the **Origin** field.

   Note: If using the example keystore from step 1, select **Import from PKCS12**.

9. Enter a **Key Alias** to identify the key.
   This alias should match the key alias (or "friendly name") that was specified when generating the certificate or keystore to be uploaded. Continuing the example above, this would be `wrapping_key_alias`.
10. Select **Next**.
    The **Key Creation** section includes an **Import Key** link, which displays a dialog to upload the keystore. Continuing the example, this would be `wrapping_keystore.p12`.

Import a wrapped key from a web service

Upload your wrapped key into a cryptographic module using the import key from web service functionality. The example uses a symmetric key. Similar steps can be used to import an asymmetric key.

Before you begin

Role required: sn_kmf.cryptographic_manager (module configuration), sn_kmf.cryptographic_operator (REST operation basic authentication)

About this task

KMF Import key endpoint access is required to complete the key import process.

This example uses OpenSSL to generate keys and certificates. You may substitute other comparable tools based on your requirements.

Procedure

1. Using the terminal on your local device, wrap your symmetric key using the Key Import module public key wrapping key.
   For example: `openssl pkeyutl -encrypt -pubin -inkey public_wrapping_key.pem -in symmetric_key.bin -pkeyopt rsa_padding_mode:oaep -pkeyopt rsa_oaep_md:sha256 -out wrapped_symmetric_key.txt`
   This example creates a wrapped key file named `wrapped_symmetric_key.txt`.
2. Create a cryptographic module to be tied to the API.
   See or for additional information.
3. Add a cryptographic specification with the following selections.
   - **Crypto Purpose**: *Symmetric Data Encryption/Decryption*.
   - **Key Origin**: *Import from web service*

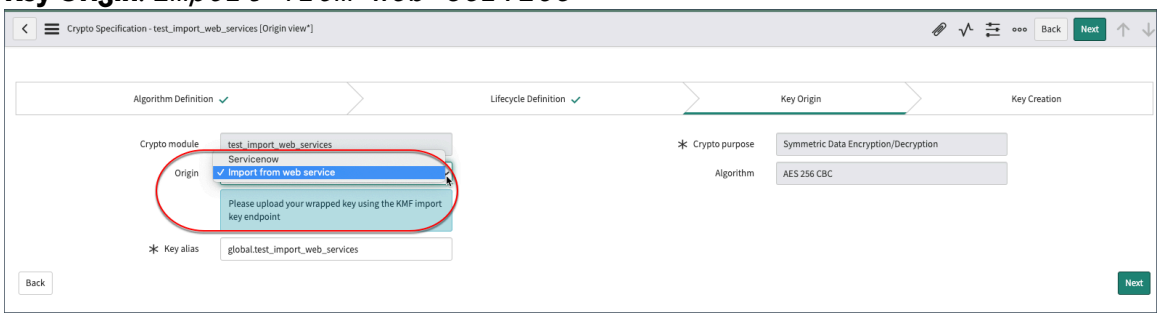

   See or for more information.

4. Execute an *HTTP POST request* to the import from a web service REST endpoint.

Option	Value/Format
URL of the endpoint	*https://<instance>/api/sn_kmf/key/import?cryptoSpecSysID=<sys_id_of_crypto_spec>*.
CryptoSpecSysID parameter	The *sys_id* of the newly created crypto specification. **Tip:** Right-click the header of the cryp to specification to copy the *sys_id*.
Header-Content-Type	Application/octet-stream.
Body	Must contain a file attachment-binary and the public key to import (wrapped_symmetric_ key.txt).
Import from web service REST endpoint	Uses basic authentication of *<username/password>*. **Note:** Ensure that the designated user has the sn_kmf_cryptographic_operator role.

Successful import of the public key results in an HTTP response message with *Status 200*.

5. Verify that the key successfully imported to the targeted cryptographic module.

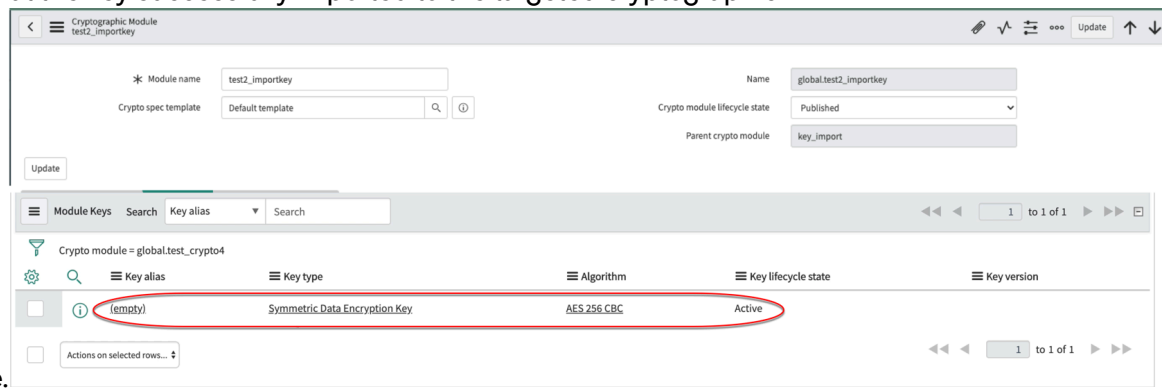

Key Management Framework Health

Access on-demand health status information for the Key Management Framework. Warning and malfunction errors contain a detailed message.

Before you begin

Role required: sn_kmf_admin and either sn_kmf.cryptographic_auditor or sn_kmf.cryptographic_manager

About this task

Each component of the Key Management Framework is outlined and reports the following statuses and colors:

- Green/Operational: The component is operational, no errors to report.
- Gray/Disabled: The component is inactive, therefore no health check is performed.
- Yellow/Degraded: Warning, the component is working, but delays/transient issues are susceptible to occur.
- Red/Malfunction: A fatal error is preventing the component from operating, which is likely to cause partial outages.

Components can include subcomponents with individual reports and their own health status impacts the parent as follows:

- If all subcomponents are inactive, the parent shows as inactive. Inactive subcomponents don't impact the health of their parents.
- If one or more subcomponents is degraded or malfunctioned, the parent health shows as degraded.
- If all subcomponents report as malfunctioned, then the parent also reports as malfunctioned.

For additional information on subcomponents, see Instance level keys in the Key Management Framework.

Note: Health checks run every 15 seconds. Refresh the health page to rerun the report.

Procedure

1. Navigate to **All** > **Key Management** > > **Diagnostics**.
2. Review the following health status information:

Diagnostic Information

Category	Details
Key Secure	Checks if encryption is being attempted.
File Key Store	Checks if an Instance Root Key (IRK) fetch attempt is occurring. **Note:** The File Key Store is an offline alternative to Key Secure used for on-premise instances and developer instances.
GlideEncrypter	Checks if a GlideEncrypter instance-level cryptographic module, specification, and key are present. **Note:** GlideEncrypter is a scriptable component that enables transparent encryption of Password2 fields and other legacy encryption usages through the Key Management Framework.
Instance Key Encryption Key (IKEK)	Checks if the key can be fetched from the File Key Store or KeySecure.
Instance HMAC Key	Checks if the key can be fetched from the File Key Store or KeySecure.
Vault PKI	Checks Vault connectivity to verify if the Instance Asymmetric Encryption Key (IAEK) and Instance Signature Key (ISK) are usable and can be fetched from Vault.
EJBCA PKI	Checks LDAP connectivity to verify if IAEK and ISK are usable and can be fetched from cache and LDAP.
Instance PKI	Checks the File Key Store and KeySecure for a key and whether the certificate is present and matches the symmetric key. **Note:** Instance PKI is only available on instances within a ServiceNow datacenter.

For assistance in troubleshooting, contact Customer Service and Support.

Prepare your instance for GlideEncrypter deprecation

Use an instance scan script to find and remove GlideEncrypter API calls on your instance. Removing these calls is a necessary step in deprecating 3DES encryption on your instance.

Before you begin

Role required: admin

The GlideEncrypter API is planned for deprecation as of the Zurich release of ServiceNow. Removing GlideEncrypter calls from your script is also a necessary step before deprecating 3DES encryption on your instance.

Procedure

1. Navigate to **All** > **Instance Scan** > **Suites**.
2. In the **Suites** list, select **GlideEncrypter** to identify the records with GlideEncrypter usages.

3. In the **Suites** record, select **Execute Suite Scan**.
4. In the **Scan Suites Now** window, select **Full Instance**, then select **Execute Scan**.
   The suite scan executes. During the scan execution a **Execute Test Scan**, a window displays the progress of the scan.

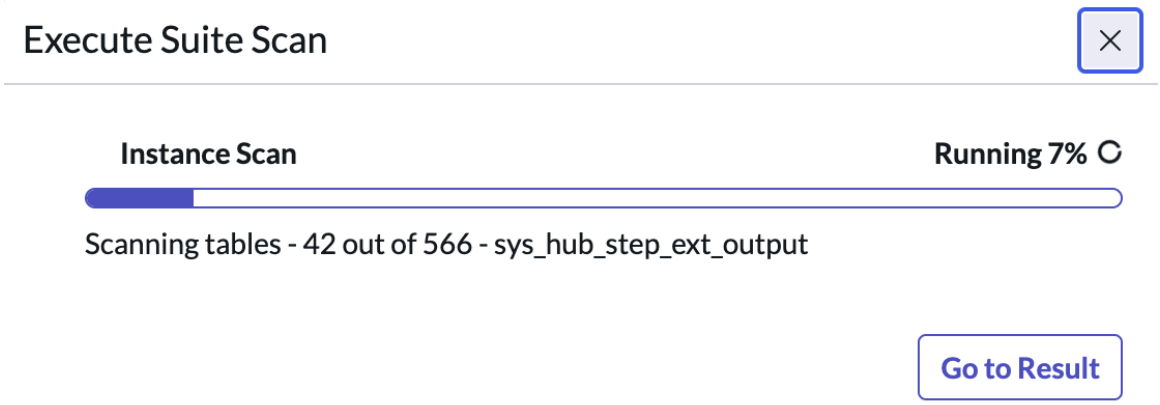

Note: This scan checks only records that have been created or modified by the customer.

5. When the scan is complete, select **Go to Result** to view the **Scan Result** record.

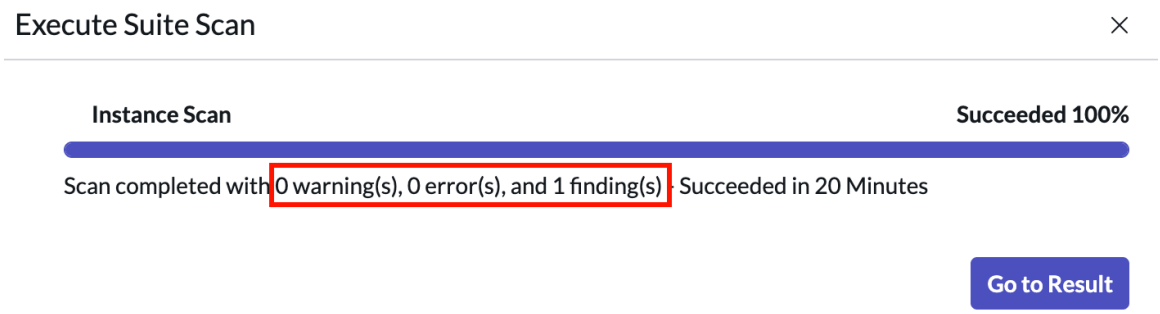

6. In the **Scan Result** record, select the **Count** field on the **Scan Findings** list to navigate to a record.

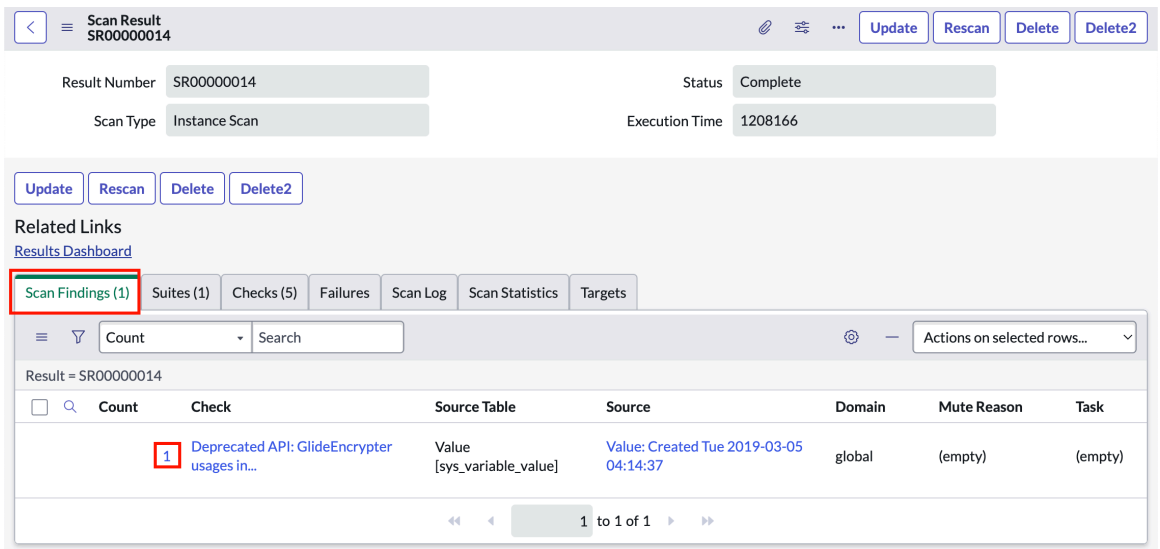

7. Modify any scripts in the record that uses the **GlideEncrypter** API.
   For details on alternatives to **GlideEncrypter**, see Alternatives to deprecated GlideEncrypter APIs.

   Note: You do not need to modify base system (out-of-box) records if you have not made any changes to them. These records are handled by ServiceNow.

8. After removing **GlideEncrypter** calls from your scripts, run the scan again to ensure that there are no remaining calls to the API.

What to do next

Deprecate GlideEncrypter usage of 3DES for password2 fields

GlideEncrypter deprecation

Learn how to remove the use legacy GlideEncrypter calls from the scripts on your instance.

GlideEncrypter availability in the Zurich release

The availability of GlideEncrypter in the Zurich release depends on whether your instance was created in or upgraded from a previous release.

New instances

In new instances created in Zurich or later, this API turned off by default. All base system scripts have been updated and no longer use calls to this API.

Upgraded instances

Zurich instances that have been upgraded from a previous release are still able to use the legacy GlideEncypter API, but the API has been updated to use AES256-GCM encryption via the Key Management Framework. This change replaces the use of the legacy 3DES encryption standard to comply with NIST guidelines, while allowing any scripts still using GlideEncrypter to operate.

Enable or disable GlideEncrypter

The availability of GlideEncrypter is controlled by the *`glide.security.glideencrypter.allow`* system property. This system property is set to **false** by default in new Zurich instances and can't be updated. In instances upgraded to Zurich, this property can be changed by administrators with the security_admin role.

The glide.security.glideencrypter.allow system property

Value	Behavior
true	When true, GlideEncypter can still be called in scripts, but uses the AES256-GCM encryption via the Key Management Framework.
false	When false, GlideEncrypter calls return `null`, and administrators see this error: `Unsupported call to GlideEncrypter. Details: GlideEncrypter is deprecated and now returns null, please refer`

Deprecate GlideEncrypter usage of 3DES for password2 fields

Deprecate GlideEncrypter usage of 3DES encryption standard on your instance ensure that your instance uses the more secure Advanced Encryption Standard (AES) exclusively for the encryption and decryption of your Password2 data.

Beginning in Rome, password2 data is protected using the Key Management Framework, which uses the more modern Advanced Encryption Standard (AES) algorithm. However, some configurations and fallbacks in password2 logic can still use the 3DES algorithm for encryption and decryption.

In the Vancouver release, administrators can choose to deprecate the 3DES algorithm entirely. After completing this change, your instance uses AES encryption exclusively for all encryption and decryption tasks relating to password2 data. This change provides better instance security than compared with 3DES encryption, and is necessary to remain NIST compliant.

Considerations before deprecation

Transferring password2 data between instances

When transferring password2 encrypted texts to other instances, you must ensure that KMF Key Exchange is enabled between source and target instances. This configuration ensures that the keys used to encrypt password2 texts are available on both instances to decrypt the password2 encrypted texts. Before deprecating 3DES, Consider the following use cases that can impact password2 data between instances.

- If you have applications on your instance that use password2 data, ensure that KMF Resource Exchange is installed on that instance. KMF Resource Exchange ensures that instance level keys used to encrypt the password2 data on the source instance are available on the target instances for decryption. For more information, see Key Management Framework Resource Exchange.
- If you plan on exporting password2 data through XML or Data Sources, ensure that the target instance has KMF Key Exchange enabled. This configuration ensures that the instance level keys used to encrypt the password2 data on the source instance are available on the target instances for decryption. For details on this configuration, see Key Management Framework Key Exchange.

> **Important:** The examples above are more common scenarios, but if you're using any other means of transferring password2 encrypted text between instances, you must configure KMF Resource Exchange to ensure the target instance can decrypt password2 data.

Downgrading an instance after the 3DES deprecation

The following only applies for instances that have password2 fields have input lengths greater than 125 characters and you have already deprecated 3DES encryption.

To downgrade an instance to release earlier than Vancouver via Instance Cloning, take the following steps before initiating the clone.

1. Check if data preservation is configured to preserve password2 field data.
2. If yes, then before requesting a clone, contact ServiceNow support to disable 3DES deprecation. In the **Reason** field, use "Clone downgrade pre-requisite for password2 support."

Legacy password2 fields

Your instance uses 3DES encryption to convert password2 data to legacy (pre-Rome) password2 data. After deprecating 3DES encryption, this option is no longer available. If you still need this feature, request partial deprecation (see details in the next section).

How to deprecate 3DES

After you've reviewed the preceding use cases, use knowledge base article KB1704481, for a step by step process to safely deprecate the usage of DES or Triple DES algorithm in instance. For details see KB1704481.

Important: You must elevate to security admin to see the **Security Compliance** module and perform these steps. For details on that process, see Elevate to a privileged role.

After GlideEncrypter deprecation

After the deprecation process is complete, the following information applies to your instance.

- password2 fields still support decryption (but not encryption) of 3DES encrypted data.
- Existing 3DES encrypted data in password2 fields remain as is until the field value is updated by a user or workflow.
- Any update to the value of a password2 field removes 3DES encrypted text and replaces it with the text encrypted by KMF using AES.
- In some situations, your instance may display an error when saving password data:

```
Action Aborted: Password value cannot be saved due to technical
issue. Please see KB1296997 for help.
```

If you see this error refer to support information in knowledge base article KB1296997.

Key Management Framework Resource Exchange

ServiceNow® Resource Exchange is a KMF feature that gives you the capability to exchange resources between instances in a secure manner.

Terminology

When using the Resource Exchange, reference the following terminology:

Resource Exchange Terminology

Name	Description
Resource Exchange	The process to exchange resources across instances.
Key Exchange (KE)	The process to exchange keys over instances.
Key Source instance (Key Source)	The instance that owns the keys.
Key Target instance (Key Target)	The instance that requests the keys.

Overview

Resource Exchange utilizes the KMF cryptographic APIs to provide confidentiality, integrity, authentication, and non-repudiation. Currently, Resource Exchange supports the Key Exchange functionality. See Key Management Framework Key Exchange for additional information.

Key Management Framework Key Exchange

KMF Key Exchange is a subset function of KMF Resource Exchange. Key Exchange securely transfers encrypted data across multiple instances.

Overview of Key Exchange

Key Exchange securely transfers keys across instances.

KMF Key Exchange provides a secure way for customers to exchange KMF keys between instances. One application use case is the data cloning process. With Key Exchange, crypto module keys are copied over during the data cloning of KMF components. Cryptographic modules, module key specifications, and module access policies are included in the cloning process. Transfer of keys isn't included.

This functionality is included with the Key Management Framework, which is included in the ServiceNow AI Platform Encryption subscription bundle. For details on this product, see Key Management Framework.

Using Key Exchange

Administrators who use KMF for Field Encryption can use Key Exchange to clone the keys between production instances when performing data cloning. In data cloning, the administrator/KMF cryptographic manager can perform the following:

- Exchange all keys to the other instances.
- Exchange particular keys one time or periodically to the other instance.
- Send on-demand requests from the target instance to the key source instance.
- Exchange keys from source to target for rekeying ciphertext.
  - Manage the expiration time of the request with the ability to delete keys or reject the key exchange request if the request has expired.
  - After the request is completed and the key is imported, the used key will be set as expired and timestamped.
  - Rekey ciphertext on the target instance that was encrypted with keys from the source.

Supported modes

Key Exchange supports several modes on the encryption module crypto specification level:

Mode	Description
Automatic (no configuration, default behavior)	All keys are sent over automatically during the data cloning process without additional configuration.
Configurable (one-time configuration setup)	The administrator configures the keys to be sent over during the data cloning process.
Manual (person in the loop)	The administrator sends an on-demand request on the target instance to the source. The request must be approved by an administrator on the key source instance.
Rekey (automated request)	The administrator selects the option of rekey during the cloning setup process.

Configure Key Exchange

Key Management Framework (KMF) generates automatic key exchange requests for supported cryptographic modules during the fresh installation or upgrade of the instance, and manages the data encryption key locally for the instance.

Before you begin

A cryptographic module with a key must be created in both the target and source instances before using Key Exchange.

Role required: sn_kmf.cryptographic_manager

About this task

Key Exchange requests are initiated from the target instance.

Automatic Key Exchange is active by default when cloning an instance, where the property is cloned to the target instance. Along with KMF, configure system properties to manage how keys are handled during an instance clone:

- **Turn off automatic key exchange:** Set the *glide_encryption.auto_key_exchange.enabled* property to **false** for recurring clone requests.
- **Send auto key exchange requests**: Set this property to **true**.

> **Important:** The base system property is set to **true** by default, meaning that automatic key exchange is activated when cloning an instance. This value must be set to **false** if you're using the Rekey ciphertext with Key Exchange or the recurring Key Exchange functionality. See Recurring Key Exchange walkthrough for additional details.

Procedure

1. Navigate to **All** > **Key Management** > **Resource Exchange Requests** > **New**.
2. On the form, fill in the fields.

Resource Exchange Request form fields

Name	Description
Exchange Frequency	○ **Adhoc**: Sends requests from the key target instance to the source instance. Enter the instance sys_id and the Host information for the Source. Not supported with Rekey of Key Exchange. ○ **One Time Clone**: One-time exchange of the keys from the source crypto specifications to the target instance. ○ **Recurring Clone**: Exchange keys from the selected source crypto specifications to the target instance on a defined recurring clone.
<Source or Target> Instance sys_id	○ **Adhoc**: Enter the sys_id for the source instance to request the keys from. ○ **One Time Clone**, **Recurring Clone**: Enter the sys_id for the target instance that sends the requests. **Tip:** Enter `stats.do` in the application navigator to locate the instance ID.

Name	Description
<Source or Target> Instance Host	Enter the host location or name of the source or target instance. **Tip:** For example `instanceA.service-now.com`
Crypto Specifications	The keys from the crypto specification in a crypto module define the keys to clone. For both one-time and recurring clone requests, your instance automatically creates a **Resource Exchange** module access policy. You don't need to configure a policy manually. **Note:** Select the lookup using list icon (🔍) to browse the available cryptographic specifications.
Enable Rekeying after Key Imported	Option to enable auto rekeying.

3. Select **Submit**.
   If successful, a confirmation displays at the top of the form. The Requests table is updated with an entry of **Request Pending** in both the source instance and in the target instance. Open the Request Record to view the status of the request, the Imported Key Count, and the Total Key Count on the target or source host.

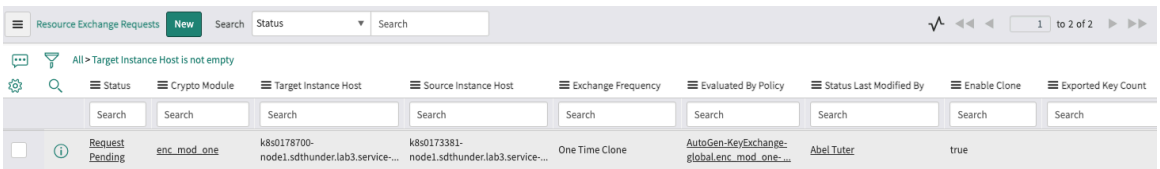

4. In the source instance, approve the pending request to complete the exchange.

   The approval process depends on the exchange frequency you selected.

   One Time Clone or Recurring Clone

   The module access policy on the source instance auto-approves the request at clone time. No manual action is required.

   Adhoc

   Manually approve each pending request in the source instance:

   a. Log in to the source instance and navigate to **All** > **Key Management** > **Resource Exchange Requests**.
   b. Open a request with a status of **Request Pending**.
   c. Change the **Status** field to **Request Approved**.
   d. Select **Update Request**.

   Important: Select **Update Request**, not the standard **Update** button. Using the standard Update button does not complete the approval.

   e. Repeat for each remaining **Request Pending** entry.

Note: Not all crypto specifications requested may exist in the source instance. Requests for specifications not found in the source instance do not require approval.

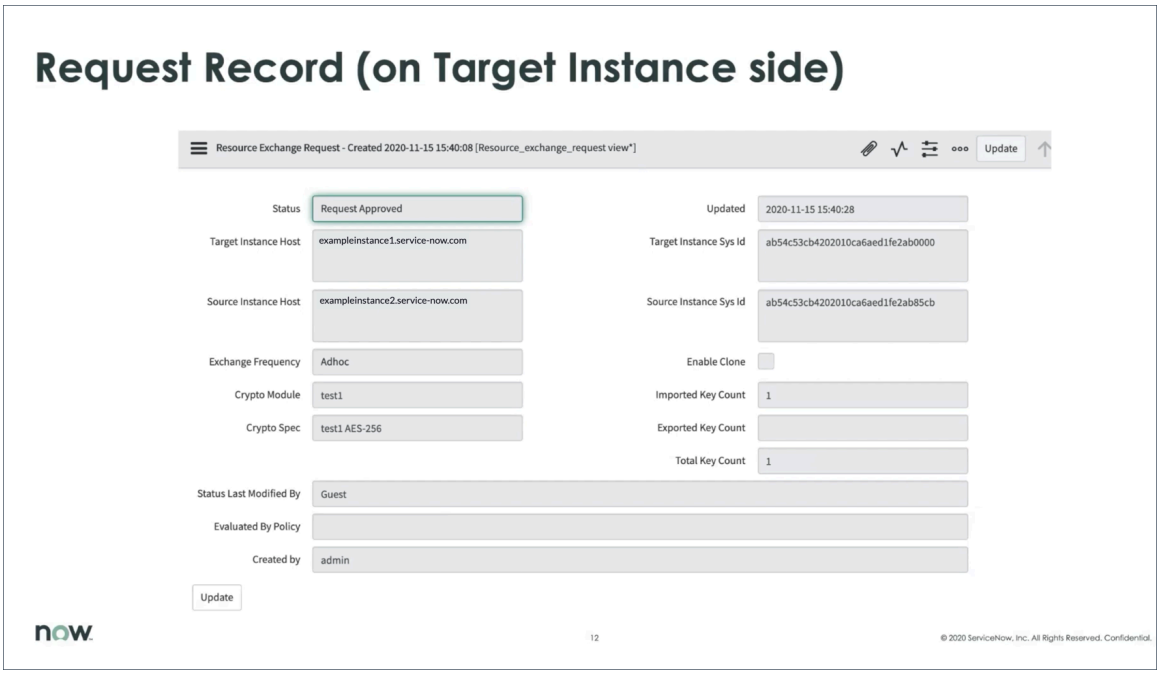

Result

After a key exchange is attempted, your non-production instance updates the *protected.script.values.kmf.rekeyed* system property. This property is visible in the System Properties [sys_properties] table after a key exchange is attempted. If the encryption using the exchanged key is successful, this property has a value of **true**. Otherwise, the property has a value of **false**. If the value is false, the instance will attempt to encrypt again the next day.

Rekey ciphertext with Key Exchange

Resource Exchange supports rekeying of ciphertext on the target instance that was encrypted with keys from the source. Rekey activity is tracked in the key life-cycle.

Overview

Administrators who use KMF for Field Encryption can use Key Exchange to rekey cryptographic keys between production instances when performing data cloning. An active key must first be available on the target instance for rekey, as rekey requires an active key. An encryption job is automatically created and run by the system to rotate and rekey the source key and re-encrypt the ciphertext.

Use Key Exchange to do the following:

- Set an expiration time frame for rekey.

  If the request has expired, then the request is rejected and the key is deleted.

- Automate rekeying ciphertext that was encrypted with keys from source instances.

  A new cloned crypto key is used to re-encrypt ciphertext on the target instance.

- The Rekey purpose is set up during the cloning process and is automated as part of the clone.
- Rekey activity is tracked on the **Modules Key** tab of the cryptographic module. Access the Key life cycle state and Key version for key activity. See Rotate keys for additional information.

Configure a Key Exchange and select the **Enable Rekeying After Key Imported** check box for activation. See Configure Key Exchange for details.

Recurring Key Exchange walkthrough

Use this walkthrough to set up a recurring key exchange in your instance using and Resource Exchange.

Before you begin

Roles required: sn_kmf.cryptographic_manager

About this task

This example shows you how a target instance requests keys from a host instance.

- Before you can perform this procedure, you must clone an instance. See System clone for more information.
- **Automatic Key Exchange**: The base system property *glide_encryption.auto_key_exchange.enabled* is **true** by default, meaning that automatic is activated when cloning an instance. The property is cloned over to the target instance.
- Turn off automatic by setting the property to **false**.

Procedure

1. On the source instance, create a crypto module or access an existing crypto module using column_level_encryption and set up the encrypted field configurations for the ciphertext encryption for Key Exchange.
   See and for details.
   a. Ensure that keys have been generated in the crypto module.

   Note: Your instance automatically creates a module access policy on the execution of the clone request.

2. From the cloned instance, navigate to **Key Management** > **Resource Exchange Requests** > **New**.
3. Complete the form and select **Recurring Clone** as the Exchange Frequency.
4. From the target instance of the clone, navigate to **Key Management** > **Resource Exchange Requests**.
   The request from the host instance is displayed in the table.

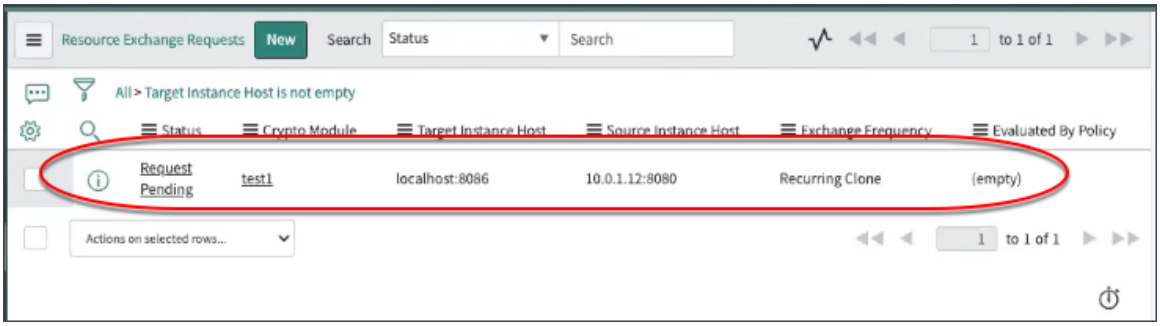

Important: For both one-time and recurring clone requests, your instance creates a module access policy automatically. You don't need to configure a policy manually. At clone time, this policy on the source instance is invoked to auto-approve the request and send keys to the newly cloned target.

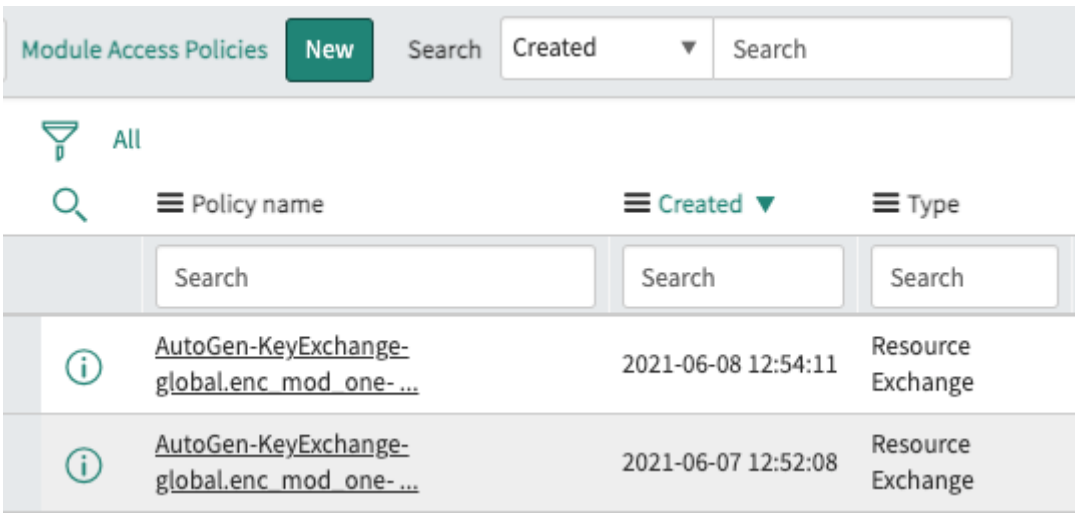

In the Requests form, the status updates to **Request Approved** and the **Imported Key Count** field appears on the record.

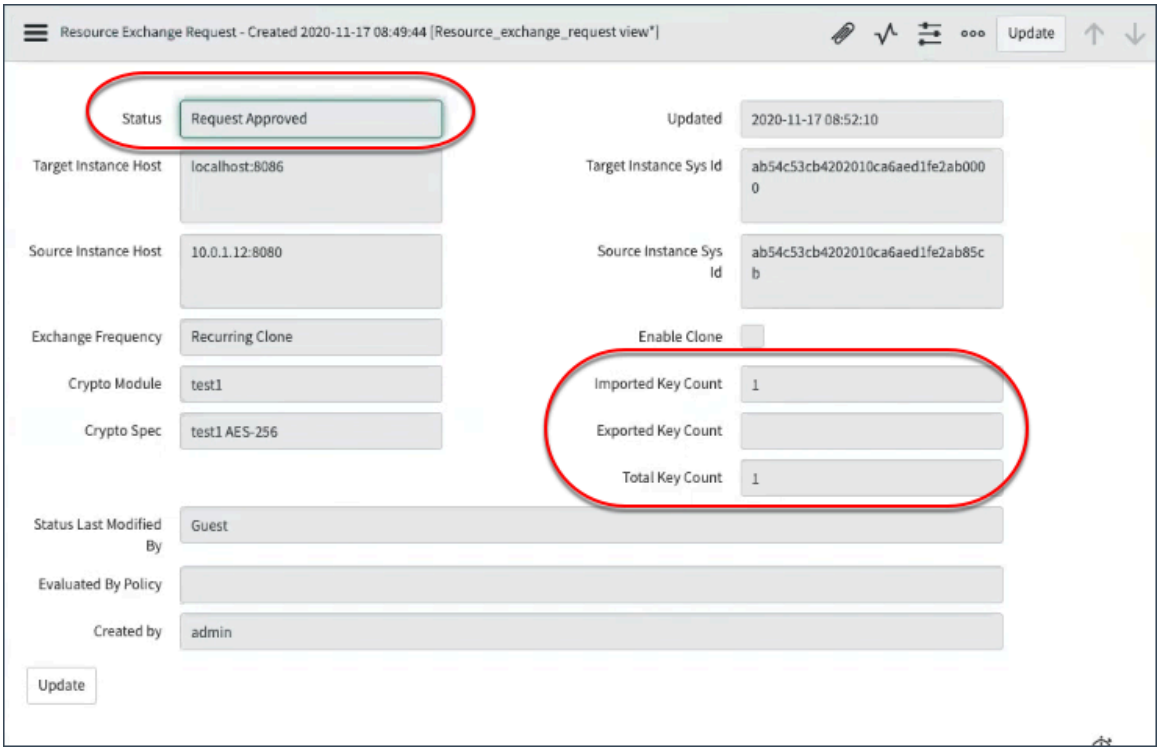

5. Return to the host instance.
6. View the Request record to see the number of exported keys.
7. View the Module Access Policy record to see that the **Type** is Resource Exchange.

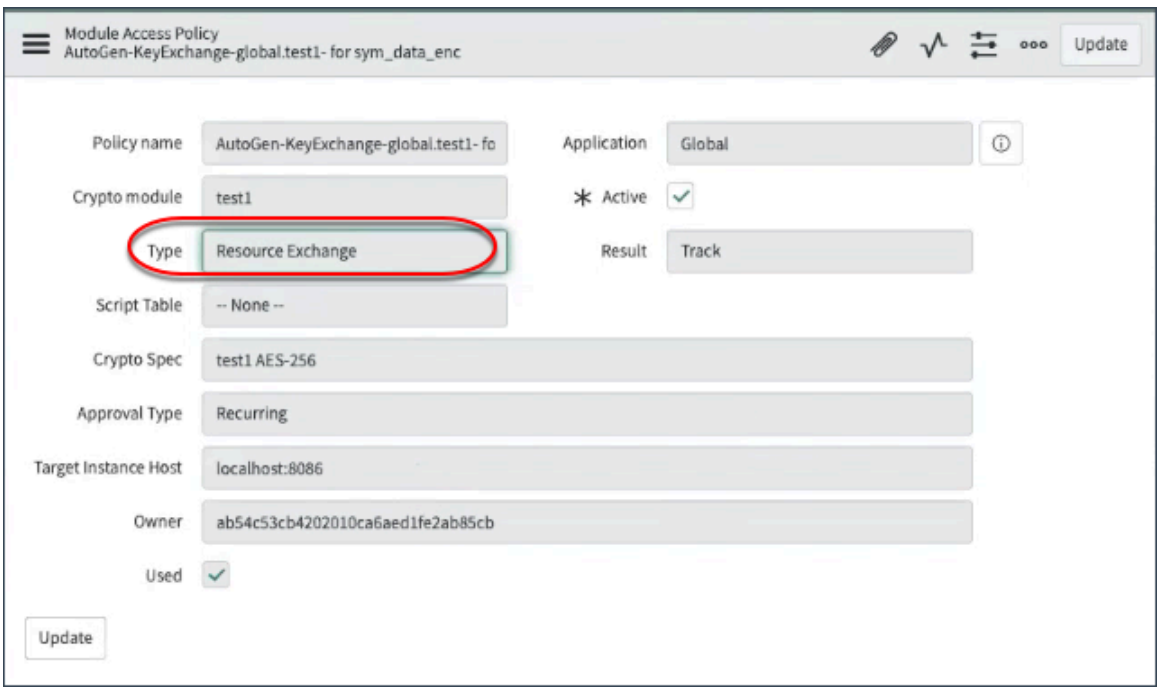

Note: Resource Exchange also supports rekeying of ciphertext on the target instance. See for details.

Result

After a key exchange is attempted, your non-production instance updates the *protected.script.values.kmf.rekeyed* system property. This property is visible in the System Properties [sys_properties] table. If the encryption using the exchanged key is successful, this property has a value of **true**. Otherwise, the property has a value of **false**. If the value is false, your instance will attempt to encrypt again the next day.

Infrastructure Security

Use Infrastructure security tools to create, upload, and manage certificates your instance uses to encrypt traffic from client to server.

The Infrastructure Security plugin provides the tools that you can use to manage the Transport Layer Security (TLS) ciphers and certificates. Your instance uses TLS to encrypt traffic from the client to your server.

Select the ciphers used on your instance

Navigate to **All** > **Infrastructure Security Settings** > **TLS Settings** to configure which TLS 1.2 ciphers your instance uses and the order in which they are tried. TLS 1.3 ciphers are fixed and cannot be modified. Custom ciphers can be configured through Customer Support.

Generate and upload your own certificates

Use the infrastructure security tools to generate your own certificate signing requests, which can be signed by the certificate authority of your choice. Navigate to **All** > **Infrastructure Security Settings** > **Upload Certificate** to upload the signed certificate to your instance's load balancer. See Generate a Certificate Signing Request.

Monitor the status of your ciphers and certificates

Use the **All** > **Infrastructure Security Settings** > **TLS Settings History** and **All** > **Infrastructure Security Settings** > **View SYOC Settings** pages to view the status of changes you have made to your ciphers and certificates.

Install the Infrastructure Security plugin

Install the ServiceNow Infrastructure Security Settings (com.glide.infrastructure_security) plugin to get started using these features. For details on plugin activation, see Activate a plugin.

After installing the plugin, enable Sign Your Own Security (SYOC) functionality by setting the *sn_infra_sec.syoc.enabled* system property to `true`.

Note: If the *sn_infra_sec.syoc.enabled* property isn't available on your instance, you must create it. For details on this process see Add a system property.

Generate a Certificate Signing Request

Use the Generate Certificate Signing (CSR) page to create a certificate signing request to support customer-signed certificates for your instance load balancer.

Before you begin

Role required: admin

The Infrastructure Security (com.glide.infrastructure_security) plugin must be installed to perform these steps. SeeInfrastructure Security for details on this plugin.

For information on using custom URLs with ServiceNow, see Set a custom URL as the instance URL.

Procedure

1. Navigate to **All** > **Infrastructure Security Settings** > **Generate CSR**.
2. Add one or more domains to your request.
   a. Select the **Add** button under the **Domains** heading.
   b. In the pop-up window, enter the domain and select **OK**.
   c. Repeat these steps as needed to add more domains.

      Note: Domains can be removed by selecting the X button to the left of each domain entry.
3. Enter any information that you want included in your request in the **Optional Certificate Fields**.
4. Select **Submit**.

Warning: You can't submit a request while another request is being generated. If this issue occurs, you see a "Resource in Conflict" error. To process, cancel the current request or wait for the current request to process before submitting another.

After selecting **Submit**, your instance will generate the certificate signing request. The request appears in the **Generated CSR** field.

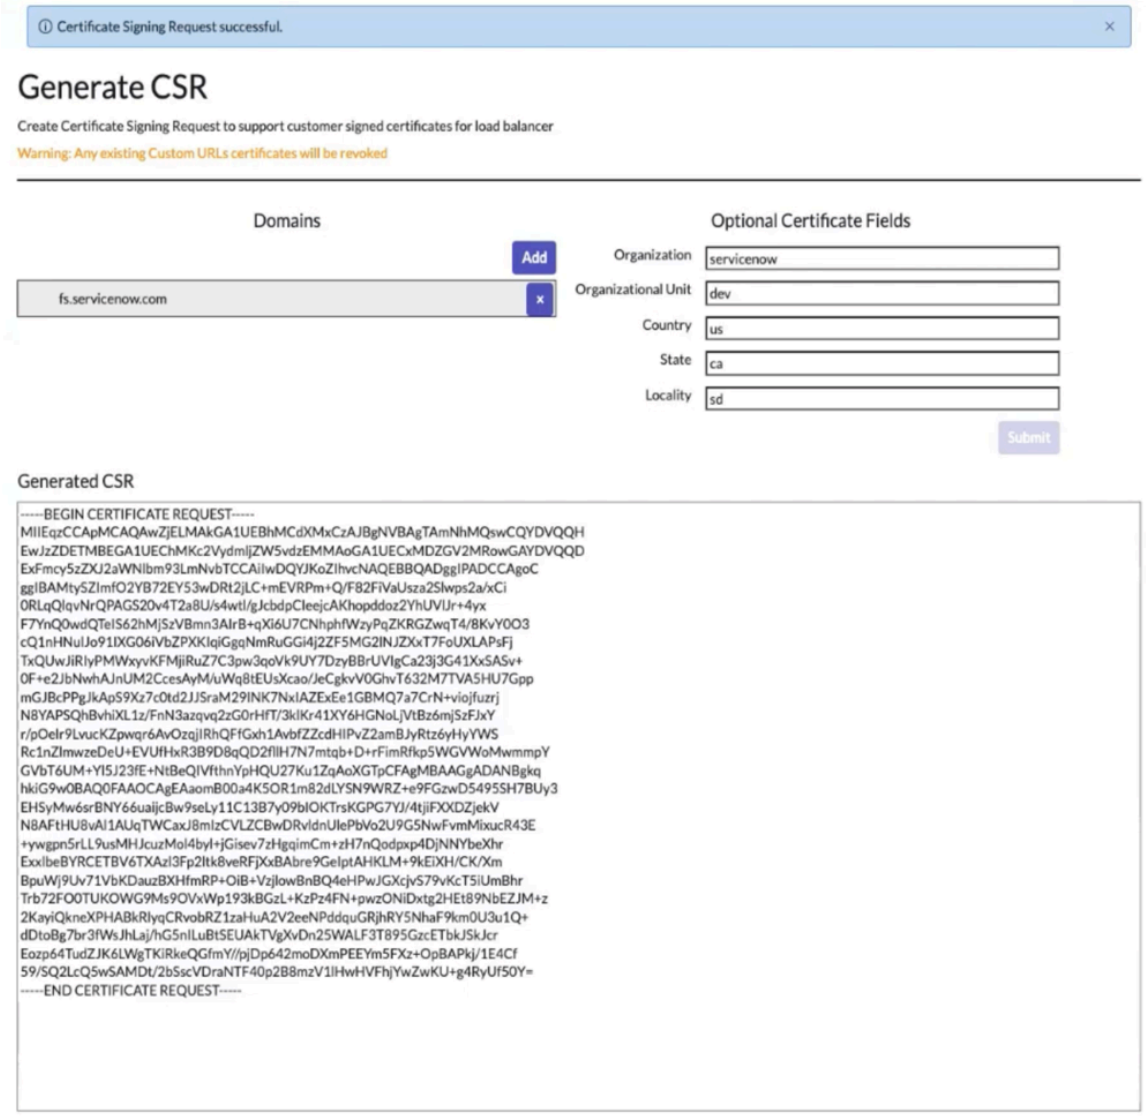

5. Copy the contents of the **Generated CSR** field.

What to do next

After generating the CSR, complete the following steps to install the signed certificate on your instance:

1. Submit the CSR to your Certificate Authority (CA). The CA signs the request and returns a signed certificate file.
2. Upload the signed certificate to your instance. See Uploading a certificate to an instance.

Password2 encryption with the Key Management Framework (KMF)

Supported by the Key Management Framework, use the Password2 (2-way encrypted) field type to encrypt and decrypt custom fields with segregation of duties, key protection, and life-cycle management. It works in accordance with NIST 800-57 guidelines and provides FIPS 140-2-L3 protection.

Password2 is a text field that stores passwords with two-way encryption. The two-way encryption stores passwords as a secure encrypted value that can be decrypted within the instance.

Tip: Beginning in the Vancouver release, administrators can deprecate 3DES encryption on password2 fields in favor of the newer Advanced Encryption Standard (AES). For details, see Deprecate GlideEncrypter usage of 3DES for password2 fields.

Activation

Password2 functionality is active by default. It's controlled by the *glide.kmf.encrypter.enabled* property, which is set to **true** for all new instances and upgrades. You don't need to enable Field Encryption Enterprise to use Password2.

How Password2 works

The Key Management Framework provides a base system parent cryptographic module **cm_glide_encrypter**. This module provides a cryptographic specification and a key that can decrypt legacy Password2 fields.

Cryptographic Module for Password2

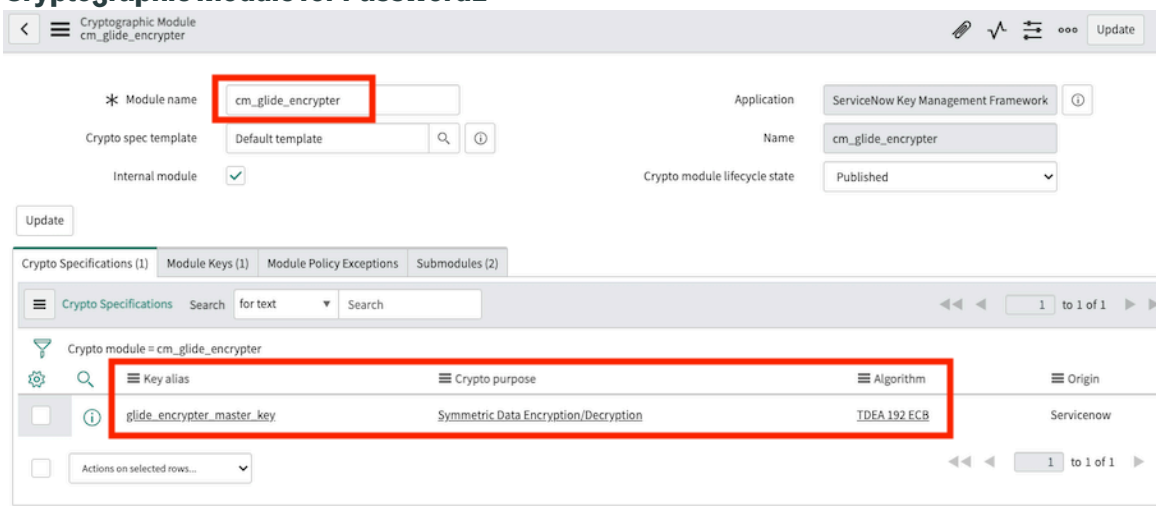

This cm_glide_encrypter module can have submodules, each with their own module key and specification. If a submodule is present with the same application scope as the application where the Password2 field is, the system uses the submodule. For example, if a table in the ServiceNow® Customer Service application has a submodule, and you write information to a Password2 field on a table in the Customer Service application scope, the cryptographic process calls the Customer Service submodule. The process also uses that submodule's key for encryption and decryption with a unique AES 256 GCM encryption key. One submodule per application scope is allowed. Parent module isn't always used for global scope. Generally, new fields use instance_level_glide_encrypter.

Note: You can't create your own submodules in Zurich. Submodules are provided in various application plugins on the ServiceNow AI Platform. You can rotate keys on submodules, but not the parent cm_glide_encrypter module.

Domain separation and on-premise customers

KMF Password2 doesn't support domain separation. You can use Password2 with on-premise instances.

Legacy Password2 and the current Password2

In Zurich, the existing Password2 field has been upgraded.

The current implementation of Password2:

- Uses the Key Management Framework in accordance with NIST 800-57 key wrapping guidelines and provides FIPS 140-2-L3 protection for the entire key hierarchy.
- Includes capabilities to create dedicated and unique KMF Password2 submodules for specific applications, providing control through application scope. Each submodule has its own unique AES 256 GCM encryption key.

Password2 fields in scripts

When accessing Password2 fields with a script, run the script under the same scope as the table scope. Use `setDisplayValue()` to encrypt Password2 values and `getDecryptedValue()` to decrypt and read the value.

Note: Don't use the `GlideEncrypter()` API on Password2 fields.

This example script shows you how to encrypt `my@Password` in the password2 column of table 'table_xyz'.

```
var  gr =  new GlideRecord('table_xyz');
gr.pwd2column_name.setDisplayValue('my@Password');

gr.insert();
```

Important: You can't use the `setValue()` API for the Password2 field.

This example script shows you how to decrypt the same field to retrieve the value:

```
var  gr =  new GlideRecord('table_xyz');
gr.query();
gr.next();
var ge=gr.getElement('pwd2column_name');
var ged1 = ge.getDecryptedValue();
```

Important: The `getDecryptedValue()` API isn't scoped. It's available globally.

1. When you encrypt data in a Password2 field, the system determines the scope of the application where the Password2 field resides.
2. The system then looks for a submodule of the **cm_glide_encrypter** parent module with the same scope as the application if the property is set to *true*.

   Note: If a submodule with the same scope is present, it uses the submodule specification and key to perform the encryption.

This illustration explains how your instance decrypts data in Password2 fields:

Password2 decryption flow

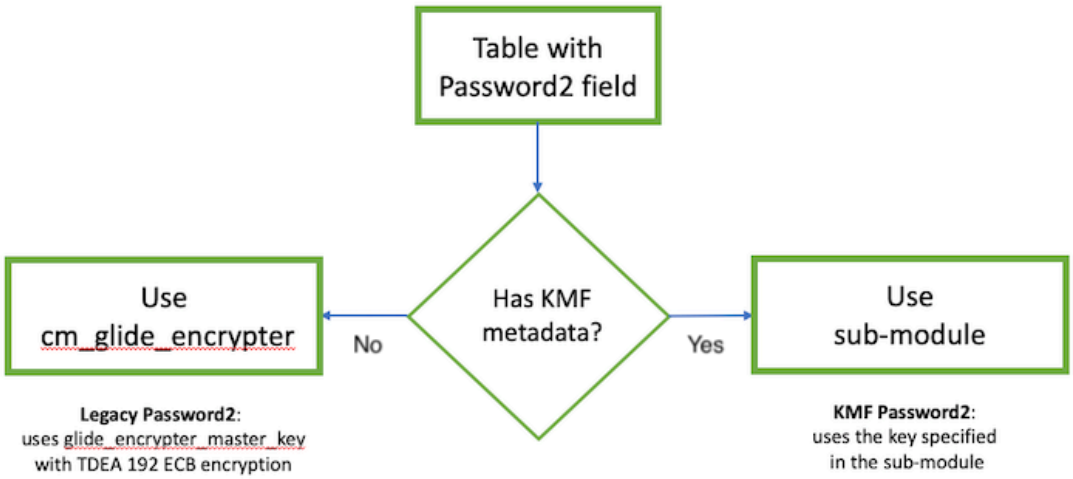

KMF Password2 migration job

A migration job is provided for customers upgrading from previous releases. It takes data encrypted with a legacy Password2 encryption and re-encrypts it with the key in a KMF Password2 submodule key. The re-encryption only applies to tables with Password2 fields in application scopes that also have submodules created for that scope. For example, a legacy Password2 field in **XYZ_example** application (with **XYZ_example** application scope) is re-encrypted only if a submodule for the **XYZ_example** application scope exists under the cm_glide_encrypter parent module.

The KMF Password2 encryption keys in the submodule are protected (envelope encrypted) in the KMF key hierarchy.

Certificates

Your instance requires certificates to establish secure connections and validate signatures.

Explore	Configure
Learn the key features and business value of Certificates.	Plan your core configurations.

Upload  Plan and set-up the application.	

Exploring Certificates

Your instance requires certificates to establish secure connections and validate signatures.

Certificates are used for features such as:

- LDAPS
- Outbound web service mutual authentication
- Web service security
- MID Server

In order to use a certificate, you must generate or purchase a certificate for the secured server, or client and upload it to an instance.

LDAP certificates

An SSL certificate is required for the instance to establish an LDAP over SSL (LDAPS protocol) connection with an LDAP server.

The instance accepts two types of LDAP certificates:

Certificate	Type	Required for
LDAP server certificate	Any supported type	All LDAP configurations
LDAP client certificate	Java keystore type	Mutual authentication

If there are multiple server certificates, the instance tries each server certificate in turn until the LDAP server allows the connection. If you use multiple LDAP servers, be sure to include the SSL certificate for each LDAP server.

Mutual authentication requires the client to present a certificate in addition to the server. If your LDAP server requires mutual authentication, you must also provide your LDAP server's client certificate in a Java keystore type certificate.

Certificate criteria

A valid certificate must meet these criteria:

- The certificate can have a key size up to 2048 bits.
- The certificate must have one of these file extensions:

Extension	Description
DER	The Distinguished Encoding Rules format is a binary message transfer syntax. This format also supports the .CER and .CRT file extensions.
CER	Certificate file extensions for certificates using the Distinguished Encoding Rules format.
CRT	Certificate file extensions for certificates using the Distinguished Encoding Rules format.
PEM	The Privacy Enhanced Mail format is a base-64 encoded DER certificate enclosed between "-----BEGIN CERTIFICATE-----" and "-----END CERTIFICATE-----" text strings.

Certificate trust

By default, your instance trusts only certificates from a Certificate Authority (CA) recognized in the Java Virtual Machine (JVM). Self-signed and enterprise-signed certificates aren't trusted.

Note: To learn more about the properties that affect the use of certificates, see in Instance Security Hardening Settings.

Generating an LDAP client certificate

Generate an LDAP client certificate for mutual authentication using OpenSSL. The final output is a PKCS#12 certificate stored within a Java keystore.

Before you begin

Role required: admin

About this task

See the OpenSSL documentation for more information about generating certificates. These steps assume you have access to OpenSSL.

Enter these commands in a command line interface.

Procedure

1. Generate a self-signed client certificate.

   Example

   For example, this command creates a client certificate test1-cert.crt based on the test1-key.key private key.

   ```
   openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout
    test1-key.key -out test1-cert.crt
   ```

2. Convert both the certificate file and private key to PKCS#12 (a file with a .pfx or .p12 extension).

   Example

   For example, this command converts the client certificate and private key to a PKCS#12 certificate called test1-certificate.pfx.

```
openssl pkcs12 -export -out test1-certificate.pfx -inkey
 test1-key.key -in test1-cert.crt
```

3. Generate the Java Key Store and import the pkcs12 file into it.

   Example
   For example, this command imports the certificate to the test1.jks Java keystore.

```
keytool -importkeystore -srckeystore test1-certificate.pfx
 -srcstoretype PKCS12 -destkeystore test1.jks
```

4. Upload the certificate in the keystore file (`test1.jks`) to the instance.

   Note:

   If you are uploading to an on-premise instance using a certificate with the .jks extension and you receive an error saying "No valid certificate found to process the application upload", use a certificate with the .pfx extension instead.

What to do next
Uploading a certificate to an instance

Generating a server certificate

You can use keytool to generate a new Java keystore file, create a certificate signing request (CSR), and import the private key, public certificate pair, and signed certificates into the keystore.

Before you begin
Role required: admin

About this task
See the Java keytool documentation for more information on generating keys and CSRs.

Enter these commands in a command line interface:

Procedure

1. Generate a Java keystore and key pair.

   Example
   For example, this command creates a keystore called my.keystore and generates a private key called mydomain within the keystore.

```
keytool -genkey -alias mydomain -keyalg RSA -keystore
 my.keystore
```

2. Generate a CSR for an existing Java keystore.

   Example
   For example, this command generates a CSR called mydomain.csr or the mydomain key.

```
keytool -certreq -alias mydomain -keystore my.keystore -file
 mydomain.csr
```

3. Import a root or intermediate certificate authority, or CA, certificate to the Java keystore.

Example
For example, this command imports the CA certificate for Thawte. This command assumes that Thwate was the CA that signed the CSR.

```
keytool -import -trustcacerts -alias root -file Thawte.crt
 -keystore my.keystore
```

4. Import a signed primary certificate to the Java keystore.

Example
For example, this command imports the signed certificate mydomain.crt into the keystore.

```
keytool -import -trustcacerts -alias mydomain -file
 mydomain.crt -keystore my.keystore
```

5. Upload the certificate in the keystore file (my.keystore) to the instance.

What to do next
Uploading a certificate to an instance

Uploading a certificate to an instance

Add a certificate to the instance from the Certificates module.

Upload a certificate to an instance

Before you begin
Role required: admin

About this task
When a certificate is updated on the ADFS server, you must also upload an updated certificate to the instance.

Procedure
1. Navigate to **All** > **System Definition** > **Certificates**.
2. Select **New**.
3. On the form, fill in the fields.

Field	Description
Name	Specify a unique name for the certificate.
Expiration notification	[Optional] Select whether you want to send a notification when the certificate is about to ex pire.
Active	Select whether the instance should use this certificate for secure communications and signing requests.
Short Description	[Optional] Enter a text description of the cer tificate such as the requester or server name.
Format	Select the certificate format. The instance sup ports the PEM and DER formats.

Field	Description
Type	Select the certificate container. The instance recognizes certificates from trust stores, Java keystore, and PKCS#12 keystores.
PEM Certificate	Enter the base-64 encoded PEM-formatted text containing the DER certificate. The in stance decodes the certificate to populate the **Valid from**, **Expires**, **Expires in days**, **Is suer**, and **Subject** fields.

4. Select **Submit**.

   During the upload, the module extracts and displays the certificate's read-only properties in these fields:

   - Valid from date
   - Expiration date
   - Issuer
   - Subject of the certificate

5. Select **Validate Stores/Certificates** to check if the certificate is correct.
   If the instance encounters any errors with the certificate or keystore, it displays an error message.

Uploading a trusted server certificate

By uploading the service provider's trusted server certificate, the instance ensures it is connecting to a valid and secure service.

Before you begin
Role required: admin

About this task
The instance validates outbound Web Service calls by using the certificate provided by the service provider.

Procedure

1. Create a new Certificate record with the type **Trust Store Cert**.
2. Do one of the following actions:
   - Attach the service provider's DER formatted certificate.
   - Copy and paste the service provider's PEM format certificate into the **PEM Certificate** field.

Field Encryption

Protect encrypted data on your instance from unauthorized users, scripts, or system processes using Field Encryption.

Field Encryption is an encryption product based on the ServiceNow Key Management Framework. Field Encryption allows for the encryption of specific fields or attachments within an instance. Use Field Encryption in combination with Cloud Encryption and Access Control Lists (ACLs) to help protect sensitive information from logged in user who are not authorized to view it.

Field Encryption has two available versions.

Field Encryption Starter

Field Encryption Starter is included on the ServiceNow platform at no cost and supports encryption for a limited number of fields.

Field Encryption Enterprise

Field Encryption Enterprise is a premium, subscription version of the product available when subscribing to ServiceNow Vault or the Platform Encryption bundle.

For more information on the difference between these versions, see Exploring Field Encryption.

Field Encryption replaces the Column Level Encryption product available in Xanadu and earlier releases.

Get started

Explore	Configure	Use
		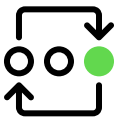
Learn the benefits of the Starter and Enterprise versions of Field Encryption.	Learn how to activate and configure Field Encryption Enterprise, and manage migration from Encryption Support or Column Level Encryption.	Use Field Encryption to manage access to encrypted data on your instances

Exploring Field Encryption

Learn the details of Field Encryption Starter and Field Encryption Enterprise

Encryption-backed access control

By default, Field Encryption blocks all users, scripts, and system processes from accessing encrypted data. However, Field Encryption has an access control feature that is used in combination with, but also separate from, Access Control Lists (ACLs) to ensure only the correct users, scripts, or system processes can access encrypted data.

You can configure the access control feature of Field Encryption through a combination of Field Encryption Modules, Encrypted Field Configurations, and Module Access Policies. The next image shows how these three components work together.

By default, encrypted data is locked down from all access. A MAP defines which accessor (users, scripts, and system processes) can be authorized to access the data.

You can configure multiple MAPs to apply different access rules to different encrypted fields. In this diagram, Module Access Policy A covers columns A, B, C, and D, and Module Access Policy B covers column E — each with its own rules per accessor.

Access rules can differ between two policies for each accessor type. The following table reflects the access rules defined for Module Access Policy A, applied to columns A, B, C, and D, and Module Access Policy B, applied to column E.

Accessor	MAP A Columns A, B, C, D	MAP B Column E
Role A	Allow	Block
Role B	Allow	Block
Role C	Block	Allow
Script A	Allow	Block
Script B	Block	Block
Script C	Block	Allow
System Context Processes	Block	Allow

Differences between Field Encryption Starter and Field Encryption Enterprise

The feature-set is different between Field Encryption Starter and Field Encryption Enterprise.

Feature	Field Encryption Starter	Field Encryption Enterprise
Number of encrypted fields	Up to 5 encrypted fields	No restriction on number of encrypted fields

Feature	Field Encryption Starter	Field Encryption Enterprise
	Note: Field Encryption Starter limits the number of encrypted fields, not encryption modules or contexts. Field Encryption replaces the deprecated Column Level Encryption product, which used a module and context-based limit.	
Attachment encryption	No	Yes
Key management	None (Contact ServiceNow Support for key rotation)	Manage keys from your instance with no involvement from ServiceNow Support
Supported data types	All supported data types	All supported data types
Number of Field Encryption Modules	No restriction	No restriction
Number of Module Access Policies	No restriction	No restriction

Field Encryption users

Users

User	Description
Key Management Framework (KMF)Admin or KMF Cryptographic Manager	These roles are used to configure elements of Field Encryption. • Field Encryption modules and module keys • Cryptographic Specifications • Module life-cycle policies • Encrypted field configurations for fields and attachments • Module Access Policies (MAPs) • Configures, wraps, and uploads customer supplied keys (for Field Encryption Enterprise) • Configures Access Observer and review Access Observer logs. • Schedule mass encryption, decryption, or re-keying
KMF Cryptographic Operator	Configures properties for customer supplied keys

Field Encryption and record history

Changes to fields encrypted with Field Encryption are not tracked in the activity stream for the record or in the record history [sys_history_set] table.

Encryption on system tables

Field Encryption currently doesn't support the encryption of fields and attachments of system tables (tables that begin with sys_).

What to explore next

To learn more about configuring and using Field Encryption, see:

- Configuring Field Encryption
- Using Field Encryption

Configuring Field Encryption

Learn how to activate and configure Field Encryption and manage migration from Encryption Support.

Activate Field Encryption

Learn how to active either Field Encryption or Field Encryption Enterprise.

Roles Required for Configuring Field Encryption

Learn about the roles required to configure Field Encryption.

Migration from encryption support

Use Scheduled jobs to migrate your keys and encrypted data from legacy Encryption Support to Field Encryption Enterprise. See details for this process at Migrating to Field Encryption

Change attachment encryption settings

Improve security by preventing users from attaching unencrypted files. For details, see Prevent users from attaching unencrypted files.

Activate Field Encryption

Activate either Field Encryption Starter or Field Encryption Enterprise.

Before you begin
Role required: admin

ServiceNow has replaced Column Level Encryption with Field Encryption beginning in the Yokohama release.

In the Yokohama release, customers that have not previously used Column Level Encryption may start using either Field Encryption Starter or Field Encryption Enterprise under the new entitlement structure.

Customers that were using Column Level Encryption in previous releases, and who want to begin using Field Encryption, have the following options:

Field Encryption Starter

Column Level Encryption Starter customers can install Field Encryption Starter with no need for re-implementation. Field Encryption Starter takes over the existing configuration and adds new features seamlessly.

> **Warning:** There are differences in entitlement between Column Level Encryption Starter and Field Encryption Starter. Before installing Field Encryption Starter, ensure that your configuration complies with the entitlements. For information on entitlements for Field Encryption Starter, see Exploring Field Encryption.

Field Encryption Enterprise

Column Level Encryption Enterprise customers should work with their account teams to ensure they have the correct entitlement for Field Encryption Enterprise. Once that happens, the Field Encryption Enterprise plugin will be available to install in their instances.

> **Important:** Column Level Encryption Enterprise does not automatically grant an entitlement to Field Encryption Enterprise, as it is a different, but replacement, product.

Procedure

1. Navigate to **All** > **System Definition** > **Plugins**
2. Under **Search your licensed applications and plugins**, search for `Field Encryption`. The search should reveal the plugin. If you purchase a subscription for Field Encryption Enterprise, you can also see this plugin available.

> **Important:** To activate Field Encryption Enterprise, you must first purchase a subscription. Your account manager can arrange to have the plugin activated on your organization's production and non-production instances, generally within a few days.

3. Select or Field Encryption Enterprise, then select **Install**.

> **Note:** When domain separation and delegated admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

Role requirements for Field Encryption

Learn about the roles required to configure Field Encryption.

Managing requires the following roles. Since is based on the Key Management Framework, there are roles common to both.

- Admin
- security_admin
- sn_kmf.admin
- sn_kmf.cryptographic_manager

For complete details on details on roles, see Roles installed with Key Management Framework.

Admin and Security Admin

Users must have the admin role to elevate to the Security Admin role. You need the Security Admin role to perform high Security Tasks, such as configuring encrypted field configurations and configuring Access Observer.

Admins can elevate to Security admin using this procedure.

1. Select on your profile picture at the top right of the screen.
2. In the drop-down menu, select **Elevate Role**.
3. Select **Security Admin**.
4. Select **Save**.

KMF Admin

Users with the Admin and sn_kmf.admin roles can assign users to the KMF admin role using this process.

1. Navigate to **All** > **System Security** > **Key Management Administration**.
2. From the **Available Users** list, move a user who needs the KMF Admin role over to the **Selected User(s)** list.
3. Select **Save**.

Important: For help prevent security issues, avoid granting the KMF Admin role to more than one users. Avoid assigning this role to users when more specialized roles are available.

KMF Cryptographic Manager

Users with the KMF Cryptographic Manager role can create and update operations on cryptographic modules and module access policies. KMF cryptographic managers can also perform key management and life cycle operations.

Use the following process to assign this role to a user.

1. Navigate to **All** > **System Security** > **Users**.
2. Select a user that needs to configure Field Encryption.
3. In the **Roles** related list, select **Edit**.
4. Search for `sn_kmf.cryptographic_manager` and add the role the selected user.
5. Select **Save**.

Configure Field Encryption modules

Learn how to configure Field Encryption modules.

Before you begin

Role required: KMF Admin or KMF Cryptographic Manager

Procedure

1. Navigate to **All** > **System Security** > **Field Encryption** > **Field Encryption Modules**.
2. Select **New**
3. In the Module form fill out the fields as shown here.

Field	Value
Module name	Choose a name for the module. This name is referenced when running scripts.

Field	Value
Crypto Spec Template	Automatically populated with **Default template**. This template is used to create the cryptographic module that contains mappings of many cryptographic purposes to cryptographic specifications and recommended algorithms.
Application	The application scope for this module. This field is automatically populated with the current application.
Name	This name automatically generated. It is the module name prepended with the application scope name to avoid conflict with other scoped applications. For example, if you create a module with the name `my_crypto_module` in the global application scope, the name is saved as `global.my_crypto_module`.
Crypto Module Lifecycle State	The term "lifecycle" refers to the creation, use, and deactivation of a cryptographic module. Set this value to **Draft** initially during configuration. Set it to **Published** for active use. **Note:** The Default template is automatically set to **Published**.
Parent Crypto Module	For Field Encryption, ensure this value is set to **column_level_encryption**.

4. Select **Submit**.

What to do next
Configure the purpose, algorithm, key length, mode, and origin of your encryption key in Cryptographic specifications for Field Encryption.

Cryptographic specifications for Field Encryption

Use cryptographic specifications to define the purpose, algorithm, key length, mode, and origin of your encryption key.

Before you begin
Role required: KMF Admin or KMF Cryptographic Manager

About this task
This procedure shows how to configure generated keys. A data encryption key (known as a Module Key) is automatically populated once you have configured the Crypto Specifications.

For customer-supplied key configuration, see Configure Customer-supplied keys for Field Encryption Enterprise.

Procedure
1. Navigate to **All** > **System Security** > **Field Encryption** > **Field Encryption Modules**.
2. In the **Cryptographic Modules** list, open the module record you want to configure.

   Note: The names displayed in the list are appended with the scope, for example, `global.[your Field Encryption Module name]`.

3. In the **Crypto Specifications** related list, open the cryptographic specific record by selecting the name under **Key alias**.

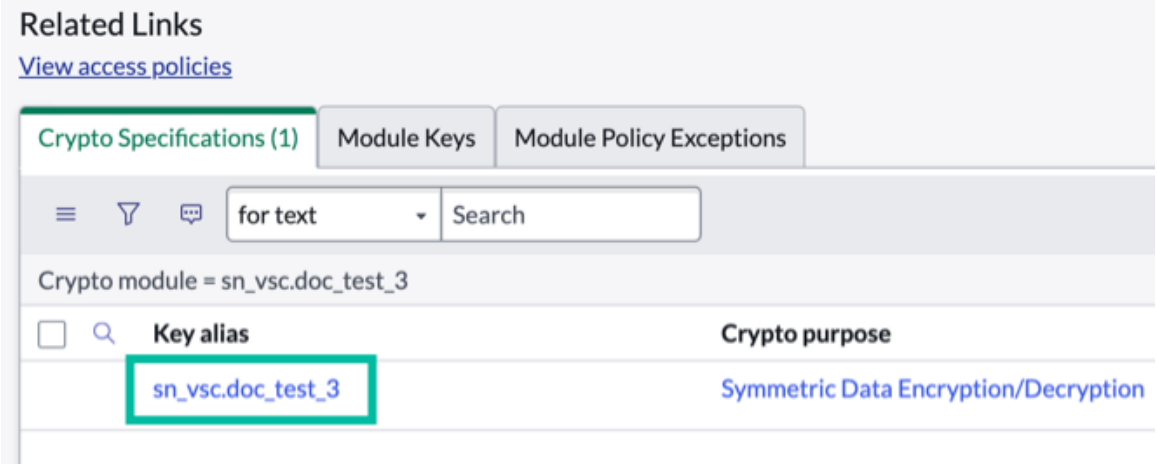

4. In the Cryptographic Specifications form, fill out the fields as needed.

Note: The fields are divided into sections. Select the **Next** or **Back** buttons to navigate between sections.

Section	Field	Description
Algorithm Definition	Crypto module	Displays the name of the selected cryptographic module.
	Crypto purpose	The purpose of the selected algorithm, key, length, and mode. For Field Encryption, this field is read only and has a value of **Symmetric Data Encryption/Decryption**.
	Algorithm	Select a type of algorithm used to accomplish the cryptographic purpose. The options available are filtered to align to the selected cryptographic purpose.
Lifecycle Definition	Applies to	Displays the selected key that the lifecycle applies to.
	For field	Select the type of control for the key that you want to apply for the lifecycle. ○ Expiration date ○ Future activation date ○ Future destruction date ○ Future renewal date ○ Future rotation date
	Key type	
	Lifecycle default	
	Order	Order in which to process the key lifecycle state for the crypto specification. Lower values execute before higher values.
	Relative duration	Number of years, months, or days the key is valid.
	Relative duration type	Duration type of the lifecycle. Select from **Years**, **Months**, or **Days**.

Section	Field	Description
	Relative operation	Choose **Before** or **Before**.
	Relative to	Select a field that the duration is relative to. ○ Activation date ○ Compromise date ○ Deactivation date ○ Destruction date ○ Expiration date ○ Generation date ○ Last renewal date ○ Last rotated date ○ Revocation date **Note:** This field works together with the value selected in the **Relative operation** field.
	Type	Select if the value for the key lifecycle is relative or absolute. **Relative** Enter a value that depends on other data entries in the system, such as key generation, activation, and deactivation. **Absolute** Enter an exact value, such as a date.
Key Origin	Crypto module	Displays the Name of the selected cryptographic module.
	Origin	Whether the key originated from or is supplied by the customer. ○ For starter, select **Servicenow** ○ For Enterprise, select **Servicenow** or **Customer Supplied**
	Key alias	Name of the cryptographic module with the scope appended to the front of the name.
	Crypto purpose	Displays the purpose of the selected algorithm, key, length, and mode. For Field Encryption, this field is read only and has a value of **Symmetric Data Encryption/Decryption**.
	Algorithm	Displays the algorithm used to accomplish the crypto purpose.
Key Creation	Crypto module	Displays the name of the selected cryptographic module.
	Key alias	Displays the name of the cryptographic module with the scope appended to the front of the name.
	Generate Key	Select this link to generate your data encryption key if you're using a generated key, and not a customer supplied key.

Section	Field	Description
	Auto generate key	If you don't select the **Generate key** link, a data encryption key is automatically generated during the first time data must be encrypted using the cryptographic module.
	Crypto purpose	Displays the purpose of the selected algorithm, key, length, and mode. For Field Encryption, this field is read only and has a value of **Symmetric Data Encryption/Decryption**.
	Origin	Displays the value that was selected during the **Key Origin** section.
	Algorithm	Displays the algorithm used to accomplish the crypto purpose.

5. Select **Go To Crypto Module** to return to your Module record.
   In the Module record, there's now an entry in the **Module Keys** related list. The **Key alias** field in the **Crypto Specifications** related list is now empty, since the key alias has moved to the new module key.

What to do next

For information on using a customer supplied key, see Configure Customer-supplied keys for Field Encryption Enterprise.

Module keys for Field Encryption

The Module Keys tab shows you summary level information about your Field Encryption Data Encryption Key(s). You can view the Key alias, Key type, Algorithm, Key lifecycle state, and Key version.

Accessing module key information

Access your module keys by from records in the Cryptographic Module [sys_kmf_crypto_module] table. From a Cryptographic Module record, you can select the **Key alias** field in records in the **Module Keys** related list.

The information available depends on which version of Field Encryption you have active on your instance.

Field Encryption Starter

Using Field Encryption Starter, you can view key usage audit data.

Field Encryption Enterprise

You to perform manual key operations such as Renew, Revoke, Rotate, or Suspend, in addition to viewing key usage audit data.

Module key information

Information on your module keys is displayed on the module key record in the fields listed here. These fields are read-only and used for information purposes only.

Field	Description
Generation date	When the Module key was generated
Activation date	When the Module key was activated.

Field	Description
Last renewal date	When the Module key was last renewed.
Last rotated date	When the Module Key was last rotated.
Deactivation date	When the Module Key was deactivated.
Destruction date	When the Module Key was destroyed.
Future activation date	When the Module Key will be activated.
Future renewal date	When the Module Key will be renewed.
Future rotation date	When the Module Key will be rotated.
Future destruction date	When the Module Key will be destroyed.
Key lifecycle state	Current the state of the Module Key.

Module lifecycle policy exceptions for Field Encryption

Use module lifecycle policy exceptions to customize the lifecycle of your module keys.

Before you begin

Role required: KMF Admin or KMF Cryptographic Manager

About this task

Module lifecycle policy exceptions change the lifecycle policy of Field Encryption modules from the standard Instance-level lifecycle policy. For example, if you've configured symmetric keys to be limited to one year at the instance level, you can create a module lifecycle policy exception for a specific Field Encryption module to allow its key to remain active for two years.

Procedure

1. Navigate to **All** > **System Security** > **Field Encryption** > **Field Encryption Modules**.
2. Select the field encryption module record that requires a module lifecycle policy exception.
3. In the field encryption module record, select **New** in the **Module Policy Exceptions** related list.
4. In the key lifecycle policy form, fill in the fields as needed.

Field	Description
Crypto Module	Displays the name of the field encryption module that will use this policy exception.
Applies To	The specified key is auto populated.
Key Type	Select the key type. Exception policies are related to a specific key. Multiple exception policies can be created per Field Encryption Module.
Policy Condition	Create qualifying conditions from the drop-down menu and complete the additional constraint criteria.
Result	Select **Reject** to reject use of the key or **Track** to allow use of it when the criteria are met.

5. Select **Submit**.

Configure Customer-supplied keys for Field Encryption Enterprise

Bring your own data encryption key to the platform instead of using the one that ServiceNow generates.

Before you begin

Role required: KMF Admin or KMF Cryptographic Manager

About this task

If you're using Field Encryption Enterprise, you can use your own data encryption key to the platform rather than one generated by ServiceNow.

You must have a symmetric key that has been generated outside of ServiceNow. The examples in this document rely on OpenSSL. For more information on OpenSSL, see details at https://www.openssl.org. If you are using other cryptographic tools, such as LibreSSL or GnuTLS, refer to the documentation for those products for similar steps.

Procedure

1. In a command line on your machine (example: Terminal), run the following command: `openssl rand -out mykey.bin -hex 32`.
   Save the `mykey.bin` file, which will be used in following steps.
2. On your instance, navigate to **All** > **System Security** > **Field Encryption** > **Field Encryption Settings**.
3. Change the **Key Source** field from **ServiceNow Generated Keys** to **Customer Supplied Keys**.
4. Select **Submit**.

What to do next

Use the symmetric key you've created on your instance by following these steps:

1. Configure properties for customer-supplied key
2. Wrap your customer-supplied key
3. Upload your customer-supplied key

Configure properties for customer-supplied key

Review the system properties for the ephemeral public wrapping key that your instance uses to unwrap customer-supplied keys.

Before you begin

Role required: KMF Admin or KMF Cryptographic Operator

About this task

You must wrap your symmetric data encryption key with a ServiceNow ephemeral public wrapping key before you can upload it to your instance.

When your key is uploaded to your instance, the instance unwraps it using the private side of the public key.

You can use system properties on your instance to define key padding, ephemeral key pair size, and a key validity period for this ephemeral public key.

System Property	Description	Default value
glide.kmf.ephemeral.key_padding	Key padding scheme for the ephemeral key.	`OAEPWithSHA256AndMGF1Padding` OAEP SHA256, but SHA1 is supported.
glide.kmf.ephemeral.key_size	Key size of the ephemeral key pair.	`4096` 4096 bits, but 2048 bits are also supported.
Glide.kmf.ephemeral.key_validity_period	Period for which the ephemeral key pair is valid.	`02 : 00 : 00` 2 Hours

Procedure

Contact ServiceNow Support if you need to change any of these properties.

Note: These system properties are not visible to admins, and do not appear in the System properties [sys_properties] list. Use the table above to see their default values.

What to do next

Once your properties are configured to your needs, proceed to Wrap your customer-supplied key.

Wrap your customer-supplied key

Wrap your symmetric data encryption key with an ephemeral public wrapping key before you can upload it to your instance.

Before you begin

Role required: KMF Admin or KMF Cryptographic Operator

You must have a symmetric data encryption key in a `.bin` to use these steps. For instructions on this process, see Configure Customer-supplied keys for Field Encryption Enterprise.

Important: Your symmetric data encryption key must be in a binary format (.BIN). If another format is used, the following error message:

```
Token failed validation. Please reattach the unmodified token.
```

About this task

To modify optional properties that control the size, padding algorithm, and validity period of the key, see Configure properties for customer-supplied key.

You must have a cryptographic tool to wrap your key. The example in this document uses OpenSSL 1.1. For more information on OpenSSL, see details at https://www.openssl.org. If you're using other cryptographic tools, such as LibreSSL or GnuTLS, refer to the documentation for those products for similar steps.

Procedure

1. Navigate to **All** > **System Security** > **Field Encryption** > **Field Encryption Modules**.
2. Open a field encryption module that you've previously created.

Note: If you haven't created a field encryption module yet, you can create one using the steps in Configure Field Encryption modules.

3. In the **Module** related list, open the cryptographic specific record by selecting the name under **Key alias**.
4. Select the **Next** button until you reach the **Key Origin** section.
5. Verify that the **Origin** field has a value of **Upload customer supplied key**.
   If that value can't be selected, refer to steps 3–5 in Configure Customer-supplied keys for Field Encryption Enterprise.
6. In the **Key Alias** field, create an alias.
   Your key uses this alias once it's uploaded.
7. Select **Next**.
8. Select the link in the **Download wrapping key** field.

   A `token_publickey` file downloads to your computer. Don't rename this file.
9. On your local machine, unzip and open the `token_publickey` folder.
   You should see an import token file (.txt) and a public key file (.PEM) in this folder.
10. Move your symmetric data encryption key that you generated into this folder.
11. Copy the name of the `token_publickey` file to your clipboard.
12. Open a terminal session and navigate to the `token_publickey` folder.
13. Enter the following command:

Important: Replace any bracketed text (<>) with your specific file names and information. Use the following key wrapping command examples table as a guide.

```
openssl pkeyutl -encrypt -pubin -inkey publickey_<keyname>.
PEM -in <keyname.bin> -out wrapped_key_material -pkeyopt
rsa_padding_mode:oaep -pkeyopt rsa_oaep_md:sha256
```

Key wrapping command examples

Directions	Command	Example
Input the publickey_<keyname>.PEM	`openssl pkeyutl -encrypt -pubin -inkey publickey_<keyname>.PE`	`openssl pkeyutl -encrypt -pubin -inkey publickey_567898643ffff.PEM`
Input the name of your symmetric data encryption key	`-in <keyname.bin>`	`-in mykey.bin`
Enter the <-out> command and specify whether the wrapped key material should use 256-bit encryption	`-out wrapped_key_material -pkeyopt rsa_padding_mode:oaep -pkeyopt rsa_oaep_md:sha256`	NA

What to do next

Now that your key is wrapped, you can upload it to your instance using the procedure in Upload your customer-supplied key.

Upload your customer-supplied key

Upload your wrapped symmetric data encryption key to your instance to begin using it work encryption.

Before you begin

Role required: KMF Admin or KMF Cryptographic Manager

About this task

Note: If you don't want to supply your own key, you can use the steps in Configure Field Encryption modules to use a ServiceNow. You cannot revoke a customer supplied key.

Procedure

1. Navigate to **All** > **System Security** > **Field Encryption** > **Field Encryption Modules**.
2. Open a field encryption module where you want to use your key.
3. In the **Module** related list, open the cryptographic specific record by selecting the name under **Key alias**.
4. Select the **Next** button until you reach the **Key Origin** section.
5. Verify that the **Origin** field has a value of **Upload customer supplied key**.
   If it doesn't, and you don't can choose that value, please refer to steps 3–5 in Configure Customer-supplied keys for Field Encryption Enterprise.
6. Confirm that you have a value in the **Key Alias** field.
7. Select **Next**.
8. Select the **Upload customer supplied key** link.
   This link should appear underneath the **Download wrapping key** link that you selected as part of wrapping your key.
9. Select **Browse**, and select two files:
   a. The wrapped_key_material file
   b. The "import token" file
10. Select **OK**.

Result

A confirmation message displays a successful upload of the customer-supplied key. The key is also listed in the **Module Keys** related list with an **Origin** of `customer-supplied key`.

Now that your encryption key is configured, you can begin to specify which fields and attachments are encrypted. For details, see Configure encrypted field configurations for fields or attachments.

Configure encrypted field configurations for fields or attachments

Create an encrypted field configuration to specify which fields are encrypted on a table, and whether that tables attachments are encrypted.

Before you begin

Role required: KMF Admin or KMF Cryptographic Manager, Security Admin

You must have a configured field encryption module with a ServiceNow or customer-supplied key. If you have not yet configured a module, see Configure Field Encryption modules.

Procedure

1. Ensure that you are in the same application scope as the table you want to encrypt.
2. Navigate to **All** > **System Security** > **Field Encryption** > **Encrypted Field Configurations**.
3. Select **New**.
4. In the **Encrypted Field Configuration** form, fields in the fields as needed.

Field	Value
Type	Select either **Column** or **Attachment** **Note:** Attachment encryption is only available with Field Encryption Enterprise.
Table	Select the table which will have it's fields or attachments encrypted.
Column	If you have chosen **Column** in the **Type** field, select the fields to be encrypted. **Note:** If the field you want to encrypt is not available, it may not be a supported type. The supported field types are: ○ String (including Full UTF-8) ○ Date ○ Date/Time ○ URL ○ HTML ○ Journal ○ Translated ○ Email ○ Phone
Active	Whether the configuration is active. **Important:** When active, your instance is actively encrypting new data in the selected fields or attachments. Users will not have access to this data unless they have permission via an associated Module Access Policy. Do not check if the field is not yet ready to begin encrypting and enforcing Module Access Policies. To ensure historical data is encrypted after an Encrypted Field Configuration is active, you'll need to run a Mass Encryption Job on the column. For details, see "Schedule Mass Encryption, Decryption, or Rekeying" .
Crypto Module	The field encryption module use by this encrypted field configuration.
Method	Select **Single Module** to ensure all fields or attachments are encrypted by a single field encryption module.

Field	Value
	Select **Multi Module** to allow for different field encryption modules to be used for different rows within a column or different attachments. For details on multi-module configuration, see .
Algorithm Equality Preserving	Displays whether **Equality Preserving** is enabled in the field encryption module selected in the **Crypto Module** field.

5. Select **Submit**.

Configure multi-module encrypted field configurations

Create an encrypted field configuration that uses more than one encryption module.

Before you begin

Role required: KMF Admin or KMF Cryptographic Manager, Security Admin

You must have a configured field encryption module with a ServiceNow or customer-supplied key. If you have not yet configured a module, see Configure Field Encryption modules.

About this task

Use multiple encryption modules for a single encrypted field configuration to encrypt different rows within a column (or different attachments on the same table) using different module keys. For example, users with different roles can encrypt data on the same table, but still be prevented from decrypting each others encrypted data.

⚠ Warning:

Note these limitations on multi-module encrypted field configuration before proceeding:

- Mass encryption isn't supported for multi-module encrypted field configurations.
- You can't change a field configuration from a multi-module to a single module. Instead, you must deactivate the multi-module field configuration and create a new single module one.
- Which module key a multi-module field configuration uses is determined by the first user to enter data into a field. Because the field encryption module is set on a per-record basis, fields in a list can be encrypted by different field encryption modules. However, within a single record, the field can be encrypted by only one field encryption module.

Procedure

1. Verify that you are in the same application scope as the table you want to encrypt.
2. Confirm that you have the field encryption modules you want to use created.
   If you have not done so, see Configure Field Encryption modules.
3. Confirm that each of your modules has a module access policy.
   If you have not done so, see Configure module access policies for field encryption.
4. Navigate to **All** > **System Security** > **Field Encryption** > **Encrypted Field Configurations**.
5. Open of create an encrypted field configuration record.
6. In the **Method** field, select **Multiple Modules**.
7. Select **Column** or **Attachment** in the **Type** field, depending on your need.

8. Select a table in the **Table** field, and column in the **Column** field, if applicable.
9. Select **Submit**.

Result

After the record is saved, and the **Active** field is enabled, new data created for the specified field is encrypted with the module key of the relevant field encryption module. When a user with the role in module access policy "A" writes to the specified table, the data is encrypted with field encryption modules "A"'s module key. In this case, only users with the same role are able to decrypt that data.

Example: Encrypting the Short Description column on the Incident table using a multi-module encrypted field configuration

1. Create two field encryption modules, referred to A and B in this example.
2. For each field encryption module, create a module access policy (MAPs A and B).
   a. For field encryption module A, give users with an HR role access to Module Access Policy A.
   b. For field encryption module B, give users with a Sales role access to Module Access Policy B.
3. Create an encrypted field configuration record specifying the Short Description column on the Incident table, and make sure that you select **Multiple Modules** in the **Method** field.
4. Have two users:
   - One with the HR Role related to MAP A and field encryption module A (User A)
   - One with the Sales role related to MAP B and field encryption module B (User B)

   create an incident record with a short description value. Have both users look at the list of incidents.
   a. The short description for the incident record created by User A is encrypted with the key for field encryption module A.
   b. The short description for the incident record created by User B is encrypted with the key for field encryption module B.
5. Users with the HR and Sales roles have access to incidents. Only a user with the HR role can decrypt and view the short description for those incidents created by User A (who had the HR Role). Only users with the Sales role can decrypt and view the short descriptions for those incidents created by User B (who has the Sales role).

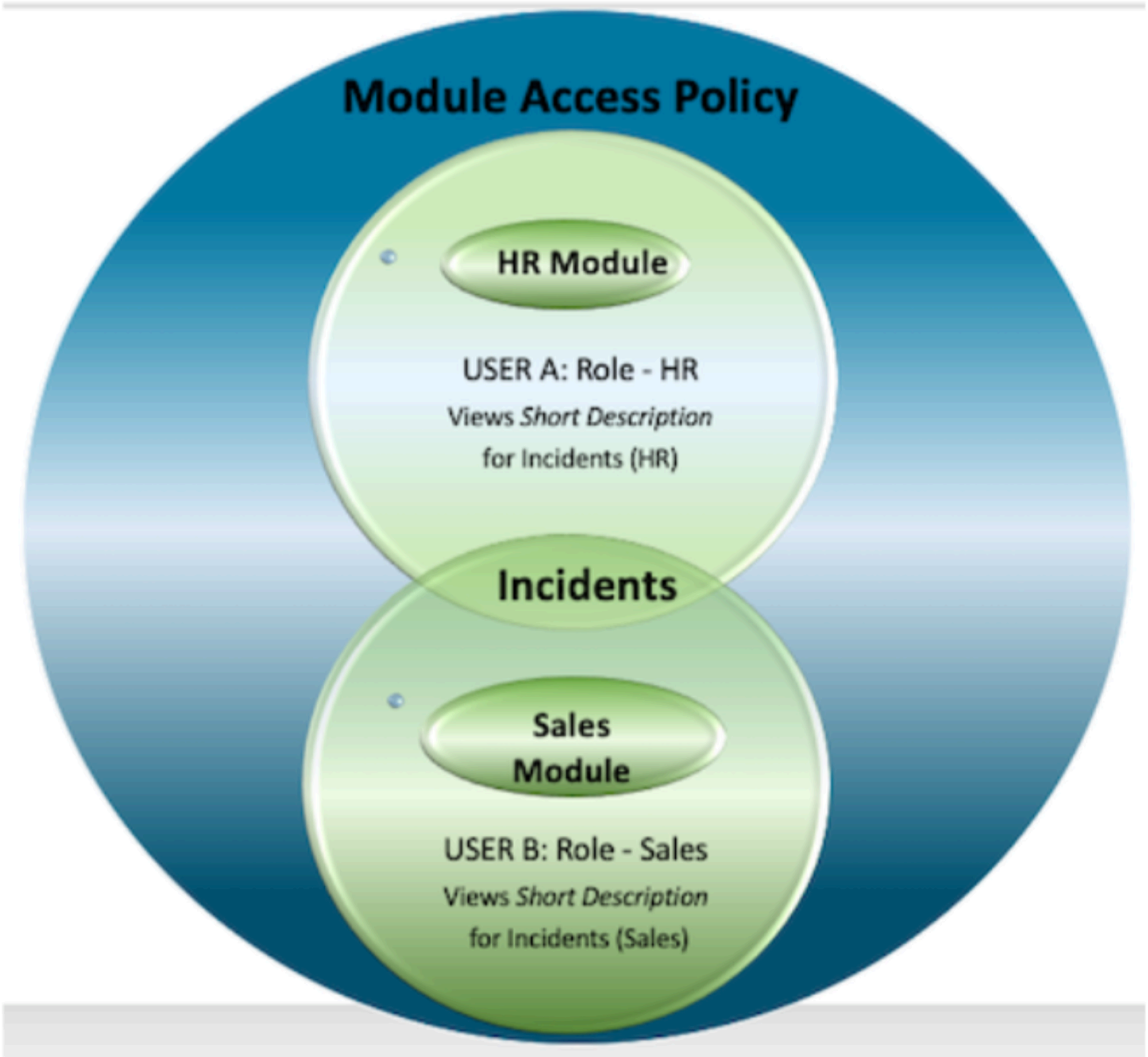

Configure module access policies for field encryption

Create a module access policy to control which users, scripts, or system processes can encrypt or decrypt data encrypted by a field encryption module.

Before you begin

Role required: KMF Admin or KMF Cryptographic Manager, Security Admin

You must have a published field encryption module to use this process. If you have not done so, see Configure Field Encryption modules.

About this task

Module access policies (MAPs) are the access controls you apply to your field encryption modules to define which users, scripts, or system processes can encrypt or decrypt data. Configure MAPs for users (via roles), scripts, or processes running in the “system” context. Without a MAP, users, scripts, or system processes aren’t able to encrypt or decrypt data, which can result in end-to-end workflow processes not working correctly.

MAPs are separate from access control lists (ACL), but can be used in combination with them. See Exploring Field Encryption for more information about the purpose behind MAPs.

For Field Encryption Enterprise, review to plan for which users, scripts, or system processes need a MAP.

Procedure

1. Navigate to **All** > **Key Management** > **Module Access Policies** > **All**.
2. Select **New**.
3. In the Module Access Policy form, fill in the fields as needed.

Field	Description
Policy Name	Name of your MAP
Crypto Module	Select the field encryption module to be governed by this MAP.
Crypto Spec	Optional. Select or create a new Cryptographic Specification for this MAP. This field appears only when the **Specify Purpose** field is enabled.
Type	Decide who or what should have access to this MAP to encrypt or decrypt data. **Scope** Anything within the specified Application Scope has access to this MAP. **Role** Only users with the specific role can access this MAP. **Script** Ensure a specified script can access this MAP. **System Access** Allows processes running in "System Context" access to this MAP. **Resource Exchange** Allows for the Resource Exchange feature access to this MAP. For more information on how these different types of MAP work, see Exploring Field Encryption.
Target Scope	Select a scope to that this MAP applies to. This field appears only if the **Type** field is set to **Scope**.
Specify Purpose	Optional. Enable to display the **Crypto Spec** field on the form. Enable this option to configure granular operations, such as some users being able to encrypt, but not decrypt.
Granular Operation	Optional. Select the cryptographic purpose for the Crypto Spec. The values available depend upon the type of Crypto Spec that is selected. For example, you can specify that this MAP only allows users to encrypt, but not decrypt, or the opposite, or both. This field appears only if there's a value in the **Crypto Spec** field.

Field	Description
	○ If a user has encrypt access, but not decrypt access, the field displays in edit mode and the data entered displays as asterisks. ○ If a user has decrypt access, but not encrypt access, the field displays the decrypted data in read-only mode. ○ If a user has encrypt and decrypt access, both read and write functionality are available for the encrypted field.
Target Role	Select which role should have access to this MAP. This field appears only when the **Type** field is set to **Role**
Script Table	Select which type of script applies to this MAP: ○ Access Control ○ Activity Designer ○ Business Rule ○ Inbound Email Action ○ Record Producer ○ Scheduled Script Execution ○ Script Include ○ UI Action ○ Widget ○ Workflow Activity This field appears only if the **Type** field is set to **Script**.
Target Script	Choose the specific script of the type table selected in the **Script Table** field that should have access to this MAP. This field appears only if the **Type** field is set to **Script**.
Check Script Version	When selected, the system checks the version of the script that is run with the version specified in the **Target Script** field. If the versions are different, the admin is notified. This field appears only if the **Type** field is set to **Script**.
Approval Type	Select either **One Time** or **Recurring**: **One Time** Allows for the symmetric data encryption key in the associated field encryption module to be securely shared to the target instance one time. **Recurring** Allows for the symmetric data encryption key in the associated field encryption module to be securely shared to the target instance on a recurring basis. This field appears only if the **Type** field is set to **Resource Exchange**.

Field	Description
Target Instance Host	Enter the URL for the target instance that the symmetric data encryption key in the associated field encryption module is being sent to. This field appears only if the **Type** field is set to **Resource Exchange**.
Impersonation	When enabled, a user impersonating another user gains any MAP permissions from both users. If disabled, a user impersonating another user only has any MAP permissions that were granted to them from before the impersonation.
Active	Enable to activate this MAP.
Result	Select one of the following: **Track** Permits access and monitors use of the MAP. **Reject** Rejects access unless a different MAP grants access. **StrictReject** Rejects access under all circumstances, even if a different MAP grants access.

4. Select **Submit**.

Migrating to Field Encryption

Scheduled jobs migrate your keys and encrypted data from Encryption Support to Field Encryption.

When you install Field Encryption, the migration process triggers automatically. Once the key migration job completes, encryption context keys are restricted to decryption only. Use cryptographic module keys for all new encryption operations going forward.

You can review the scheduled jobs by navigating to **System Security** > **High Security Settings** > **Security Jobs**:

- **autoKeyMigration**: Migrates encryption context keys to Key Management Framework (KMF) cryptographic module keys.
- **autoDataMigration**: Migrates data that you already encrypted to use the KMF cryptographic module key.

You can modify when these scheduled jobs run, and can pause or restart them at any time.

Verify that the encrypted field configurations are using your newly migrated module keys by navigating to **System Security** > **Field Encryption** > **Encrypted Field Configurations**. Look for the following items:

- The **Method** field is **Single Module**.
- The **Crypto module** field is populated with the name of the cryptographic module that the system automatically creates. You can review that module and the module access policy, both of which are active and published.

Field Encryption migration status page

Use the migration status page to track the migration of encryption contexts to encryption modules.

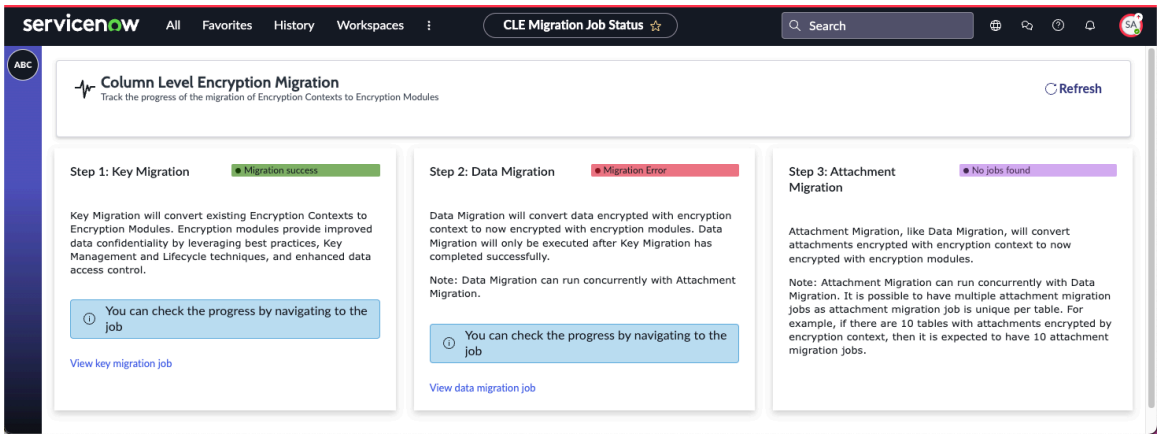

The Field Encryption Migration page displays the status of the steps involved in migrating encryption contexts to encryption modules. Each of the three sections displays the status of a specific step in the process.

Page section cards

The page contains three cards representing the steps in the migration progress. These cards display:

- (1) The status of the current step. This status will display whether the step has completed successfully, or if there are no jobs to process.
- (2) A description of the listed step.
- (3) A link to the relevant encryption job [sys_mass_encryption_job] record.

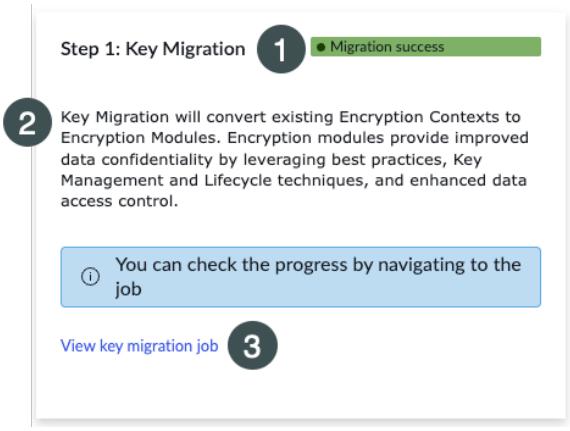

Migrate from Edge Encryption to Field Encryption

Migrate from Edge Encryption to Field Encryption to take advantage of the latest security features.

Migration process

This topic covers, which comprise the following steps:

1. Migrate columns and attachment from Edge Encryption to Field Encryption.
2. Stop and disabling Edge Encryption proxy servers.
3. De-tokenize your data
4. Decrypt Service Catalog item variables.

Limitations

- Field Encryption doesn't currently support tokenizing data in the same way that Edge Encryption does. Data that is tokenized through Edge Encryption must be included in an encrypted field configuration for Field Encryption.
- Field Encryption doesn't currently support encrypting Service Catalog item variables.

Configure Field Encryption for your Edge Encrypted fields

Before migrating your existing edge encrypted fields to Field Encryption, you must configure field encryption for these fields.

Before you begin

Role required: admin

This process assumes you have existing edge encrypted fields that you want to migrate to Field Encryption.

Procedure

1. Navigate to **All** > **Key Management** > **Cryptographic Modules** > **All** > **.**
2. In the Cryptographic Modules list, select **New**.
3. In the form, fill in these fields.

Field	Value
Name	Select descriptive name
Crypto spec template	Default template
Crypto module lifecycle state	Published
Parent crypto module	column_level_encryption

4. Right-click the form header and select **Save**.
5. In the **Crypto Specifications** related list, open the record on the list.
6. In the **Crypto Specification** record, select **Next** three times to complete the algorithm definition, lifecycle definition, and key origin sections.
   You don't need to modify any fields in these sections.
7. In the **Generate Key** field, select the **Generate Key** link to generate a key.
8. Navigate to **All** > **System Security** > **Field Encryption** > **Encrypted Field Configuration**.
9. In the encrypted field configuration list, select **New**.
10. In the form, fill in these fields.

Field	Value
Type	Select Column or Attachment, depending on what you want to encrypt.
Crypto module	Select the cryptographic module that you created in earlier steps.
Table	Select the table containing the data that you want encrypted.

Field	Value
Method	Select `Single Module`
Column	Select the column (field) on the table that you want to be encrypted.

Note: The data to encrypt is determined by the Table and Column fields. These fields should be the table and column (field) where you're currently using Edge Encryption.

11. Select **Submit**.
12. Navigate to **All** > **Key Management** > **Module Access Policies** > **All**.
13. Select **New**.
14. In the form, fill in these fields.

Field	Value
Policy name	Select descriptive name
Crypto module	Select the cryptographic module that you created in earlier steps.
Type	Select `Role`
Target role	Select a role to be used. This role must be able to encrypt and decrypt data in the column.
Result	Select the desired action.

15. Select **Submit**.
16. To verify your configuration, navigate to the table you want encrypted with Field Encryption, and open a record.

    For example, to access the User [sys_user] table, enter `sys_user.list` into the navigation filter.

    The field you have selected for encryption in the previous steps now has a lock icon next to the field label.

Result

You edge encrypted field is ready to be migrated to field encryption. To configure more fields, repeat the preceding steps for each of those fields.

Field Encryption and system clones

Cloning an instance with Field Encryption installed automatically generates new field encryption module encryption keys on the target clone instance.

If Field Encryption is installed on your instance, a new field encryption module encryption key is automatically generated on the target clone instance as a part of clone process. These keys are generated for all modules to which the user has access, and that does not have a key already.

Because of this, field encryption modules on the target clone instance may have two module encryption keys present:

- An active module encryption key. This is the new key generated after clone, as long as the module is accessible to the user and has no prior keys.
- A deactivated encryption module key (from the automated key exchange transfer

  The active module encryption key is used to encrypt inserted data as needed on the target clone instance. The deactivated module is used to decrypt existing data that was cloned over as part of the system clone.

  To use a single key to decrypt and encrypt all data, you can run a module rekeying job. For more information about module rekeying jobs, see Schedule mass encryption, decryption, and rekeying jobs.

Prevent users from attaching unencrypted files

Modify the com.glide.encryption.enable_attachment_key_ui property to prevent your users with access to an encryption module key from attaching unencrypted attachments.

Before you begin

Role required: security_admin

You must elevate to the security_admin role performing these steps. For instructions, see Elevate to a privileged role

By default, users who have access to an encryption module key are able to upload unencrypted attachments. Use the *com.glide.encryption.enable_attachment_key_ui* system property to change this behavior.

When attaching, your users see a UI picker on records that have a multi-module encrypted field configuration. When this property is set to false, users no longer see an option not to encrypt an attachment.

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. In the system properties list, find and open the system property.
3. Set the **value** of the property to `false`.

Using Field Encryption

Use Field Encryption to manage access to encrypted data on your instances.

There are two methods of encrypting field data:

- Single module - Permits data encryption using a single encryption module in a deterministic method.
- Multi module - Permits data encryption using multiple encryption modules in a non-deterministic method. Row Conditions is the new and preferred method of applying multiple modules to a field. Row Conditions applies the multi module capability in a deterministic way.

Use the related links to find information on common Field Encryption tasks.

Related topics

Create cryptographic module for Field Encryption

Create a cryptographic specification for Field Encryption

Configure advanced algorithms for Field Encryption Enterprise

Configure properties for customer-supplied keys

Encrypting fields and attachments

Field Encryption Enterprise examples

Create cryptographic module for Field Encryption

Create a Field Encryption cryptographic module to define the mechanisms used for cryptographic operations.

Before you begin

Role required: sn_kmf.cryptographic_manager or sn_kmf_admin, security_admin, admin

About this task

This procedure describes options that are available with Field Encryption with the base system and additional configuration options that become available with Field Encryption Enterprise functionality. Field Encryption Enterprise is available with a paid subscription. Refer to Encryption and Key Management subscription bundle for supported features and options available with each offering. See Activate Field Encryption for more information on obtaining Field Encryption Enterprise.

Procedure

1. Navigate to **All** > **System Security** > **Field Encryption Modules** > **New**.

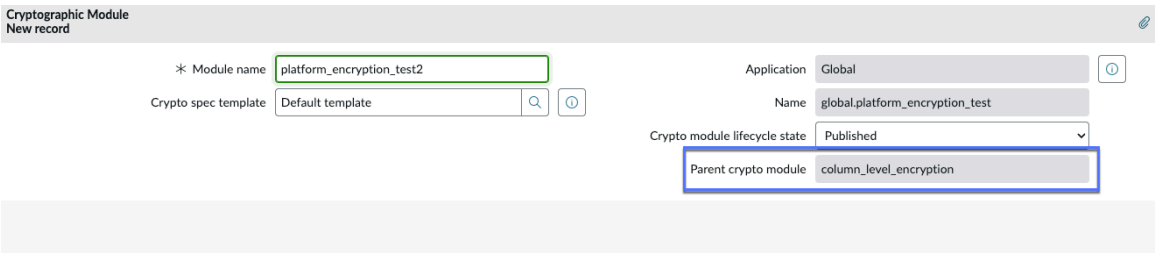

2. On the form, fill in the fields.

Cryptographic Module form

Field	Description
Module name	Alphanumeric string to be referenced when running scripts.
Crypto spec template	Default template used to create the cryptographic module that contains mappings of many crypto purposes to crypto specifications and recommended algorithms.
Application	The selected application scope.
Name	Encryption module name is prepended with application scope name to avoid conflict with other scoped applications on module creation. For example, if you created a module with the name `my_crypto_module` in the global application scope, the name is saved as `global.my_crypto_module`.
Crypto module lifecycle state	The term lifecycle refers to the creation, use, and deactivation of a cryptographic module. Set to **Draft** initially during configuration. When using the module, set to **Published**. The Default template is automatically set to **Published**.

Field	Description
Parent crypto module	The parent is populated automatically as **column_level_encryption**.

3. Click **Submit**.

   After submitting successfully, your cryptographic module is listed in the Cryptographic Modules table.

   Warning:

   For legacy encryption support users:

   If you're using the non-enterprise version of Field Encryption, you're limited to five fields. If you've exceeded this limit, you receive the following warning:

   This insertion exceeds the number of published fields limit for Field Encryption #entitled with the Subscription Product. The Enterprise subscription for Field Encryption is required for additional fields. Please reach out to your Account team.

   A default cryptographic specification is created with the crypto purpose set to Symmetric Data Encryption/Decryption and the algorithm as AES 256 CBC. Select the algorithm for updates.

4. To open the configuration options, click the newly created cryptographic module.

   Note: A maximum of five Field Encryption fields are allowed before upgrading to Field Encryption Enterprise. An error message displays and you are prevented from adding additional cryptographic modules.

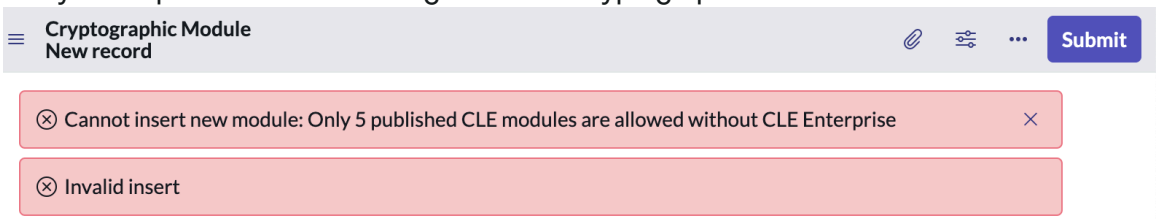

What to do next

Create a cryptographic specification for Field Encryption.

Using multiple encryption modules

Multiple encryption modules enable data to be encrypted with more than one encryption module. If each module has its own access policy based on a role, for example, users with different roles can encrypt data on the same table but used to help prevent them from viewing each other's encrypted data.

There are two ways to encrypt data using multiple encryption modules in the platform: Row Conditions and multiple module.

Row Conditions

Enables admins to define the encryption crypto module to be used. This option provides a deterministic approach and permits multiple Encryption Field Configurations (EFCs) to a single column. You can encrypt different records for a field or attachment by configuring distinct encryption keys to specific fields. This is achieved by assigning multiple Encryption Field Configurations (EFCs). Row Conditions is the preferred

method of encrypting data with multiple modules. See Encrypt data using Row Conditions.

Multiple module

Enables users to the crypto module to use to encrypt data. Since multiple module depends on the user selecting the correct crypto module to encrypt the data, this configuration is nondeterministic and is considered legacy and not preferred, even though the feature is still currently available. See Encrypt data using the Multiple Modules feature.

Encrypt data using Row Conditions

Encrypt fields with multiple Field Encryption modules using Row Conditions to define the data being encrypted and the associated encryption keys. Row Conditions can also be used to define the users that have access using the condition builder.

Before you begin
Role required: sn_kmf.cryptographic_manager or sn_kmf.admin

About this task

Note: The following guidelines apply when using Row Conditions:

- Row Conditions is only supported on encrypted columns and attachments.
- Mass encryption isn't available when using the multiple encryption modules method.
- You can't change a field using multiple encryption modules to use a single encryption module.
- Row Conditions enables data to be encrypted on the main record. Dot-walking isn't supported in the condition builder.
- Row Conditions isn't supported by the following service catalog tables:
  - Options [sc_item_option]
  - Question Answer [question_answer]
  - Multi Row Question Answer [sc_multi_row_question_answer]

The field is encrypted by the encryption module of the first user to enter data. Because the encryption module is set on a per record basis, fields in a list can have different encryption modules. Within a single record, the field can be encrypted by only one module.

Procedure

1. Create multiple cryptographic modules and an access policy for each one.

   Make sure that you grant different roles to the different cryptographic modules through the access policies.

2. Navigate to **System Security** > **Field Encryption** > **Encrypted Field Configurations** > **New**.

   If you need more information on Encrypted Field Configurations, see Set encrypted field configurations.

3. Select the **Table** and the **Column** in the table that you want to encrypt.
4. Select **Encrypt by default** to encrypt saved records that don't meet a defined row condition.

   If you select this option, make sure to select the **Crypto Module** to be used as a default encryption module.

   If this box isn't selected and records that are added don't meet a row condition, they aren't encrypted.

   For example, a Row Condition can be defined to encrypt records to declare that the "Department" field on the record equals "IT." Then, the records in which the "Department" field doesn't equal "IT" wouldn't be encrypted unless **Encrypt by default** is selected. This example utilizes a custom created "Department" Choice field on the sn_customerservice_case table, but other custom or default fields, if they are a supported data type of Row Conditions.

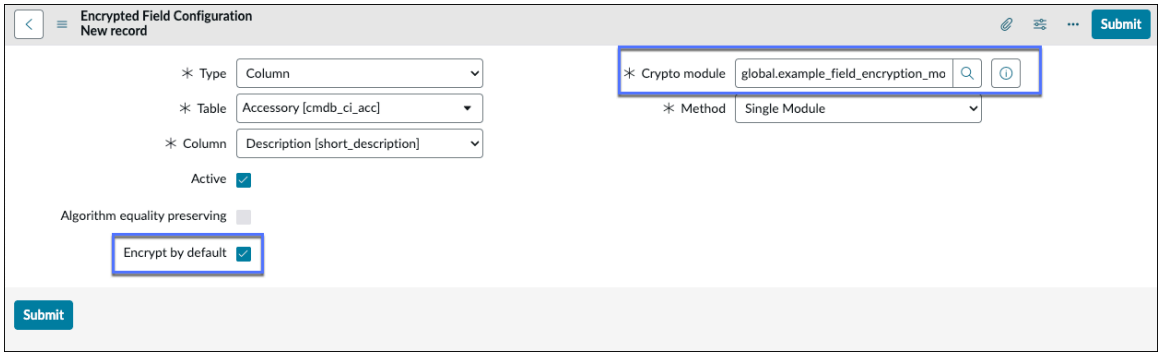

5. Select **Submit**.
6. Navigate to the condition builder **All** > **System Security** > **Field Encryption** > **Encrypted Field Configurations** and select the newly created Encrypted Field Configuration (EFC).
7. Select **New** under the **Encrypted Row Configurations** related list.
8. Define Row Conditions by filling out the required fields.
9. Select **Submit**.

Result

Note: Repeat this process as needed to ensure that the number of Row Conditions meet your required Encryption Crypto Modules (ECMs). When selecting Encrypt by default in the ECM setup, a Row Condition must also be established to define the default ECM selected.

Newly created data for the specified field is encrypted with the key for the relevant module. When a user with the role specified in module A's access policy writes to the specified table, the data is encrypted with module A's key. Only users with the same role can read the data.

Example:

This example explains how to encrypt the Short Description field on the Incident table. This example works similarly if you want to encrypt a different field in a table.
For encrypting the Short Description, you would do the following:

1. Create two Field Encryption modules A and B.
2. For each module, create a Module Access Policy (MAP) and define access as follows:
   a. Module A – to users with an HR role.
   b. Module B – to users with a Sales role.
3. Create an Encrypted Field Configuration (EFC) record.
   a. Navigate to **System Security** > **Field Encryption** > **Encrypted Field Configurations** > **New**.
   b. Select **Incident** in the Table field.
   c. Select **Short Description** in the Column field.
   d. Select the **Encrypt by default** box if you must verify any records that fall outside of the condition builder criteria are still encrypted by the default field encryption module. Not selecting this option would mean that any records that fall outside of the condition builder criteria won't be encrypted.

      Enter the default crypto module in the related **Crypto module** field.

   e. Select **Submit**.
4. Create the Row Conditions.
5. Run the appropriate encryption job:

- Mass Encryption - Run this job when a new Row Condition is created.
- Mass Rekeying - Run this job when an existing Row Condition is modified.

See Schedule mass encryption, decryption, and rekeying jobs.

6. Have a user from Module A and a user from Module B create an incident with a short description. Have both users view the list of incidents.

The short description for the incident created by the user with the HR role is encrypted by the key for module A. The short description for the incident created by the user with the Sales role is encrypted by the key for module B.

All users with the HR and Sales roles have access to incidents. However, only a user with the HR role can decrypt and view the short description for those incidents created by another user with the HR role. Likewise, only users with the Sales role can decrypt and view the short descriptions for those incidents created by the user B, who had the Sales role.

What to do next

Perform one of the following operations:

- Schedule a **Mass Encryption** job to update the encryption of the condition field.
- If modifying an existing Row Condition run a **Mass Rekeying** job to encrypt necessary data with the updated encryption module.

See Schedule mass encryption, decryption, and rekeying jobs

Encrypt data using the Multiple Modules feature

Encrypt data with more than one encryption module permitting the user to determine which keys are used for specific rows within the encrypted data.

Before you begin

Role required: sn_kmf.cryptographic_manager or sn_kmf.admin

About this task

The Multiple Modules option is considered non-deterministic and isn't the preferred method because the user determines which key to use for a given record. The ability to use multiple modules for a column is being replaced by Row Conditions. See Using multiple encryption modules. This non-deterministic implementation is still supported because it was created first and is still in use, but it's preferred to use Row Conditions for any new multiple modules use cases.

Note: Only encryption on columns supports multiple modules. Attachment encryption doesn't. Mass encryption isn't available when using the multiple encryption modules method.

You can't change a field using multiple encryption modules to use a single encryption module.

The field is encrypted by the encryption module of the first user to enter data. Because the encryption module is set on a per record basis, fields in a list can have different encryption modules. Within a single record, the field can be encrypted by only one module.

Procedure

1. Create multiple Field Encryption modules and a Module Access Policy (MAP) for each one.

   Make sure that you grant different roles to the different cryptographic modules through the access policies.

2. Navigate to **System Security** > **Field Encryption** > **Encrypted Field Configurations** > **New**.

   If you need more information on Encrypted Field Configurations, see Set encrypted field configurations.

3. In the **Type** field, you must select **Column**.
   Attachment encryption doesn't support multiple modules.

4. Select **Multiple Modules** in the **Method** field.

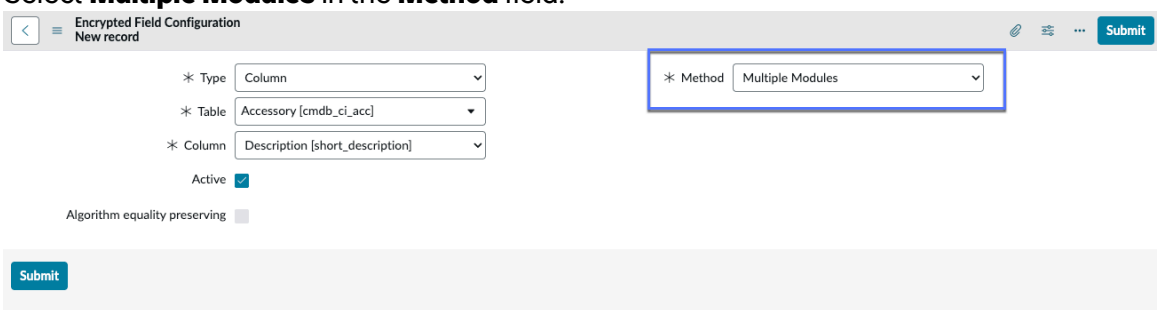

5. Select the **Table** and the **Column** in the table that you want to encrypt.
6. Select **Submit**.

Result

Newly created data for the specified field is encrypted with the key of the relevant module. When a user with the role specified in module A's access policy writes to the specified table, the data is encrypted with module A's key. Only users with the same role can read the data.

Example:

To encrypt the Short Description column on the Incident table. You would do the following:

1. Create two cryptographic modules A and B.
2. For each module, create a module access policy.

   For module A, give users with an HR role access. For module B, give users with a Sales role access.

3. Create an Encrypted Field Configuration record specifying the Short Description column on the Incident table, and make sure that you select **Multiple Modules** in the **Method** field.
4. Have two users, one with the HR role (user A) and one with the Sales role (user B), create an incident with a short description, and then have both users look at the list of incidents.

   The short description for the incident created by the user with the HR role is encrypted by the key for module A. Likewise, the short description for the incident created by the user with the Sales role is encrypted by the key for module B.

   Although all users with the HR and Sales roles have access to incidents, only a user with the HR role can decrypt and view the short description for those incidents created by user A, who had the HR role. Likewise, only users with the Sales role can decrypt and view the short descriptions for those incidents created by the user B, who had the Sales role.

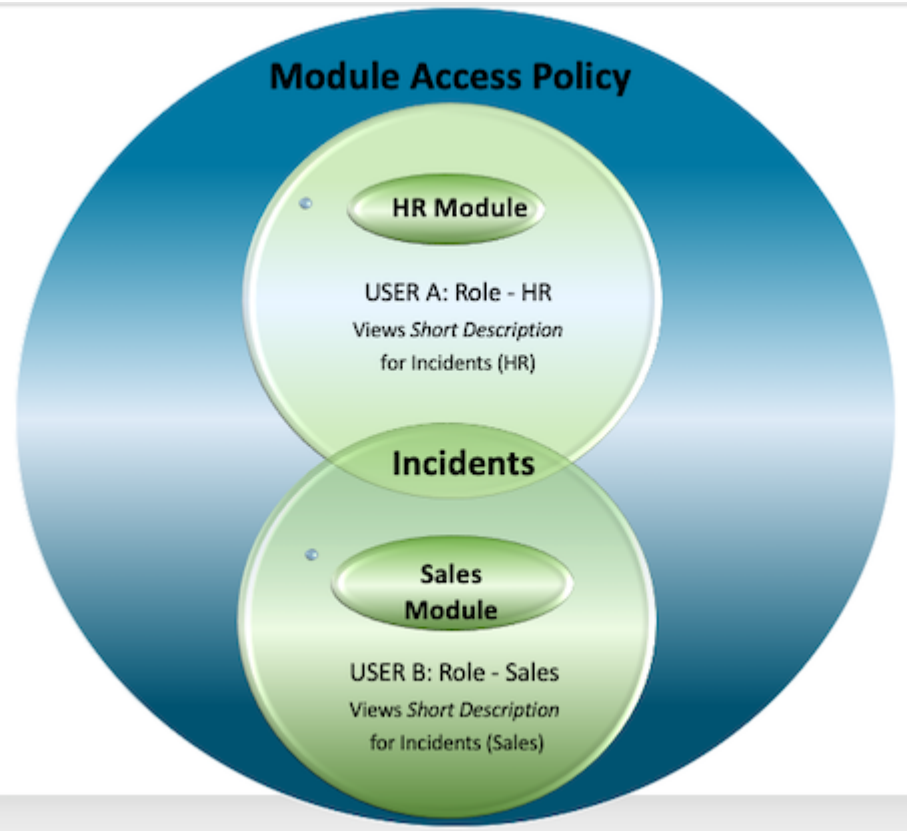

Create a cryptographic specification for Field Encryption

After you create a cryptographic module, access the corresponding cryptographic specification to define the algorithm.

Before you begin

Role required: sn_kmf.cryptographic_manager or sn_kmf_admin and security_admin or admin

About this task

This procedure describes options that are available with Field Encryption with the base system and additional configuration options that become available with Field Encryption Enterprise functionality. Field Encryption Enterprise functionality is available with a paid subscription. Refer to Encryption and Key Management subscription bundle for supported features and options available with each offering. See Activate Field Encryption for more information on obtaining Field Encryption Enterprise.

A cryptographic specification will be created by the system when you create a cryptographic module for Field Encryption Enterprise.

Procedure

1. Navigate to **System Security** > **Field Encryption Modules** > **All**.
2. Select the cryptographic module to open the configuration options.
   Cryptographic module information is displayed at the top of the screen. A Symmetric Data Encryption/Decryption crypto specification is auto-created with an AES 256 CBC algorithm.
3. Select the crypto specification from the table to open the Algorithm Definition.
   For Field Encryption Enterprise see Configure advanced algorithms for Field Encryption Enterprise.
4. Click **Next** to access the Key Lifecycle.

What to do next

Perform one of the following operations:

- Select an entry in the Key Lifecycle table to define key lifecycle behavior. See Configure key lifecycle states for details to complete the lifecycle definition for the key.
- Click **Next** to create a cryptographic key. For details on this process, see Generate a ServiceNow cryptographic key.

Configure advanced algorithms for Field Encryption Enterprise

Create a cryptographic specification to define the algorithm for a cryptographic module. Customize the encryption specifications with advanced options that are available for Field Encryption Enterprise.

Before you begin

Role required: admin

Procedure

1. On the **Crypto Specifications (#)** tab, click **New**.

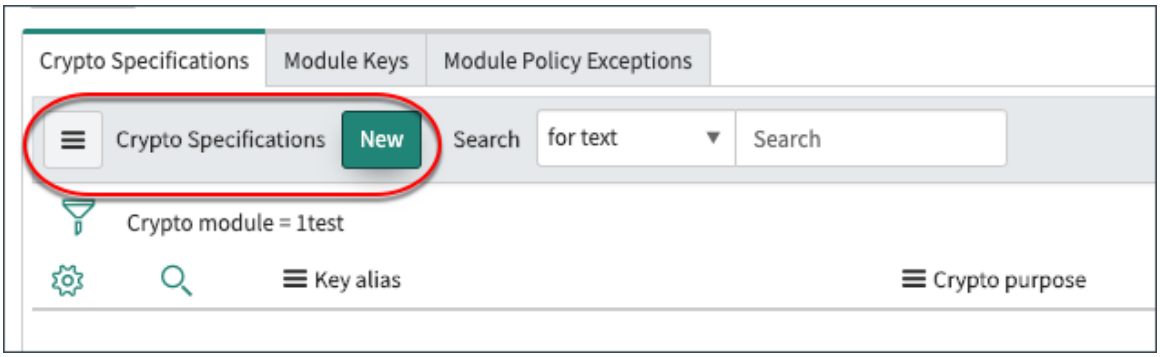

2. On the form, fill in the fields.

Algorithm Definition form

Field	Description
Crypto module	Name of the selected cryptographic module populates.
Crypto purpose	The value is **Symmetric Data Encryption/ Decryption** for Field Encryption Enterprise.
Algorithm	The value is **AES** for Field Encryption Enterprise.
Operation mode	The value is **CBC** for Field Encryption Enterprise.
Size	Possible values are **256** and **128**. **Note:** 256-bit size is most secure for encryption and is used for Symmetric Data Encryption/Decryption for Field Encryption Enterprise.
Equality preserving	Option to enable deterministic encryption. **Note:** Selecting this option means that the encrypted value of a field should be the same when the field value remains the same. Option to enable **Symmetric Data Encryption/Decryption** with **AES** in **Cipher Block Chaining (CBC)** mode.

Field	Description
Integrity	Option to provide Integrity in GCM operation and does not apply for Field Encryption Enterprise functionality.

3. Click **Submit**.

   The following example shows AES CBC-256 encryption. When Field Encryption Enterprise is active and the parent module is column_level_encryption, only Symmetric Data Encryption/Decryption AES CBC-256 applies as the crypto purpose. See Cryptographic specification overview for details.

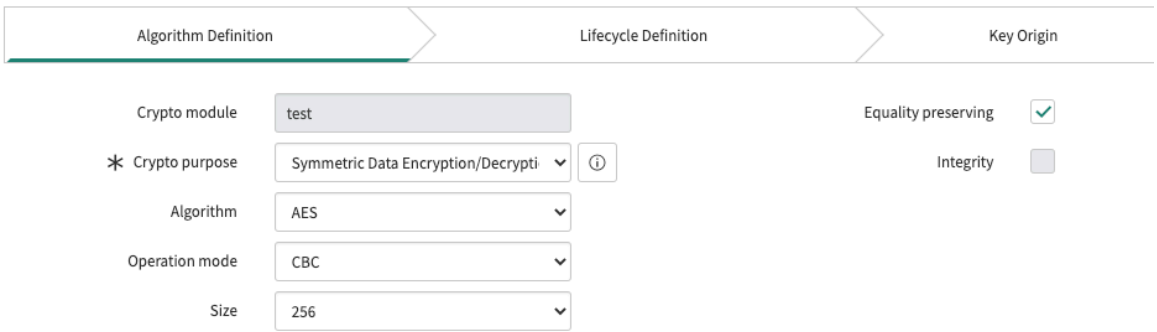

What to do next

Perform one of the following operations:

- Select an entry in the Key Lifecycle table to define key lifecycle behavior. See Configure key lifecycle states for details to complete the lifecycle definition for the key.
- Select **Next** to create a cryptographic key. See one of the following tasks for key generation:
  - Generate a ServiceNow cryptographic key.
  - Configure properties for customer-supplied keys.
  - Import the wrapping / unwrapping key pair.

Using customer supplied keys with Field Encryption Enterprise

You can use your own customer-supplied key instead of using the ServiceNow® system-generated keys.

Important: These topics only apply instances using Field Encryption Enterprise, which is only available with the `com.glide.now.platform.encryption` plugin. See Activate Field Encryption for more information on obtaining this plugin.

With Field Encryption Enterprise you can use your own keys for encryption. Administrators have the choice to use ServiceNow® supplied keys or your own customer-supplied keys (CSK) for encryption on the ServiceNow AI Platform.

Important: To make use of the customer supplied key option, you must have your own cryptographic key.

Once you have your key, you can being using it on your instances by following these steps.

1. **Configure properties for customer-supplied keys**

   There are three system properties which define the size, padding algorithm, and validity period of the wrapping RSA key pair. Review these properties and adjust their values if the defaults to not fit your needs.

2. **Wrap your customer-supplied key**

   Use a cryptographic tool to wrap your key like OpenSSL to wrap the symmetric key to use for encryption with the downloaded public key.

Configure and upload your customer supplied key

Upload your wrapped your customer supplied key and configure cryptographic module to begin using your key for encryption on your instance.

Configure properties for customer-supplied keys

If the Field Encryption Enterprise plugin is enabled, you can use system properties to define key padding, ephemeral key pair size, and a key validity period of your customer-supplied keys.

Field Encryption Enterprise with Key Management lets you manage the full key lifecycle of your data encryption keys. Optionally, you can securely exchange data encryption keys generated within your environment.

Platform Encryption with Key Management lets you manage the full key life cycle of your data encryption keys. Optionally, you can securely exchange data encryption keys generated within your environment.

System properties for defining key-pair attributes

When you provide your own key, you must wrap it with the RSA public key. Three properties define the size, padding algorithm, and validity period of the wrapping RSA key pair:

- *glide.kmf.ephemeral_key.key_padding* controls the key padding scheme for the ephemeral key. The default scheme is OAEP SHA256, but SHA1 is also supported.
- *glide.kmf.ephemeral_key.key_size* controls the key size of the ephemeral key pair. The default is 4096 bits, but 2048 bits are also supported.
- *glide.kmf.ephemeral_key.key_validity_period* defines the period for which the ephemeral key pair is valid. The default value is two hours.

After the data encryption key is imported to the instance, a secure wrapping key protects new module keys on the instance. The wrapping key is an instance key encryption key (IKEK) generated by a hardware security module (HSM) on SafeNet KeySecure. See Instance level keys in the Key Management Framework for details in key types.

Continue to Wrap your customer-supplied key.

Wrap your customer-supplied key

Wrap the symmetric key to use for encryption with the downloaded public key.

Before you begin

Note: This procedure describes options that are available with KMF base system and options to be used with Field Encryption Enterprise functionality. Field Encryption Enterprise functionality is available only when the *com.glide.now.platform.encryption* plugin is active. See Activate Field Encryption for more information on obtaining Field Encryption Enterprise.

Some of the steps in this document require the use of a cryptographic tool installed on your local device. The examples in this task use the OpenSSL tool. For more information on this tool

see https://www.openssl.org. If you are using other cryptographic tools, such as LibreSSL or GnuTLS, refer to the documentation for those products for similar steps.

- Modify optional properties that control the size, padding algorithm, and validity period of the key. See Configure properties for customer-supplied keys.
- You must have your symmetric key (.BIN) for encryption.

> **Important:** Your key must be in binary format. If another format is used, a `Token failed validation. Please reattach the unmodified token.` error message displays.

- You must have a cryptographic tool to wrap your key. This example uses OpenSSL 1.1.

Role required: sn_kmf.cryptographic_manager or sn_kmf.admin

Procedure

1. Navigate to **All** > **Key Management Framework** > **Cryptographic Modules** > **All**.
2. Select the cryptographic module that you created for the customer supplied key from the Crypto Specifications related list.
3. You will be directed to the **Key Creation** step.
4. If you have not previously downloaded the wrapping key, click the link to download the *token_publickey<id>.zip* file and save it to the same location as your key.

   > **Note:** Do not rename the downloaded *token_publickey<id>* file.

5. Unzip the file to your local network.
   The zip file contains two files, an import token and a public key .PEM certificate. Wrap your symmetric key with the public key to encrypt it.
6. Copy the name of the *token_publickey* file to your clipboard.
7. From a command line, use the copied *token_publickey* file name to open the folder of the unzipped files as a placeholder for the wrapped key.
8. Edit this script by replacing the examples with the names of your crypto files.

```
"downloads user.name$ cd token_publickey_<token>
openssl pkeyutl -encrypt -pubin -inkey publickey_<keyname>.PEM
-in <keyname.bin>
-out wrapped_key_material -pkeyopt rsa_padding_mode:oaep
 -pkeyopt rsa_oaep_md:sha<128 or 256> "
```

Review the key wrapping commands in the following table for more information.

Key wrapping commands

Directions	Command	Example
Open the file directory where you downloaded the wrapping token.	`cd`	`cd token_publickey123456789`

Directions	Command	Example
Paste the name of the `publickey.PEM` certificate.	`openssl pkeyutl -encrypt -pubin -inkey`	`publickey_586798643ffff.PEM`
Paste the name of your key here.	`-in`	`mykey.bin`
Enter the <-out> command and specify if the key is 128 bit or 256 bit.	`-out wrapped_key_material -pkeyopt rsa_padding_mode:oaep -pkeyopt rsa_oaep_md:sha256`	N/A

9. Run the command.
   A system message displays `token_publickey_<keynumber>`. The key will be generated and a *`wrapped_key_material`* file added to the directory.
10. Upload the wrapped key.

What to do next
Return to Configure and upload your customer supplied key to upload your wrapped key.

Configure and upload your customer supplied key

You can use your own customer-supplied key instead of using the ServiceNow® system-generated keys.

Before you begin
Roles required: security_admin, sn_kmf.cryptographic_manager

If you're NOT supplying your own keys, you don't need to perform this procedure. To create a cryptographic module with ServiceNow® keys, go to Create a cryptographic module or Create cryptographic module for Field Encryption.

Note: This procedure only applies to Field Encryption Enterprise functionality. See Activate Field Encryption for more information.

Important: You can't revoke a customer supplied key.

Procedure

1. Navigate to **All** > **System Security** > **Field Encryption Settings** and verify that **Customer Supplied Keys** is selected.

Key source selection

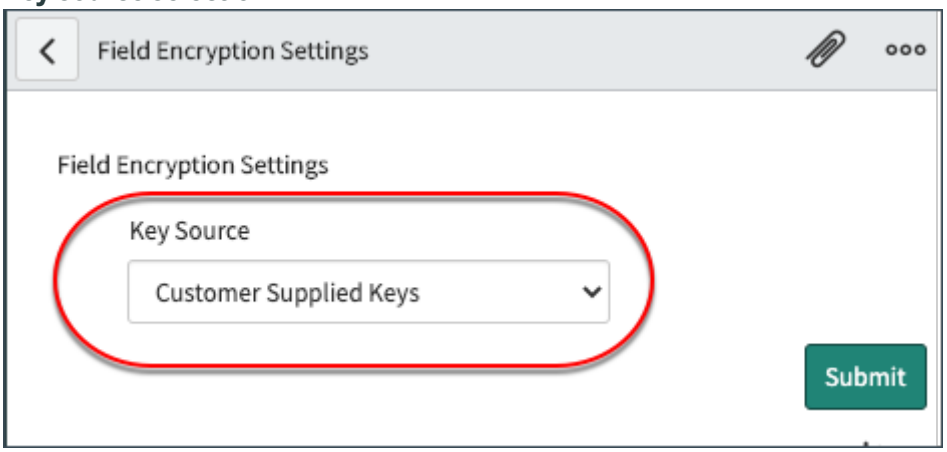

2. Select **Submit**.
3. Return to **System Security** > **Field Encryption Modules** > > **Create New**.

Create new cryptographic module

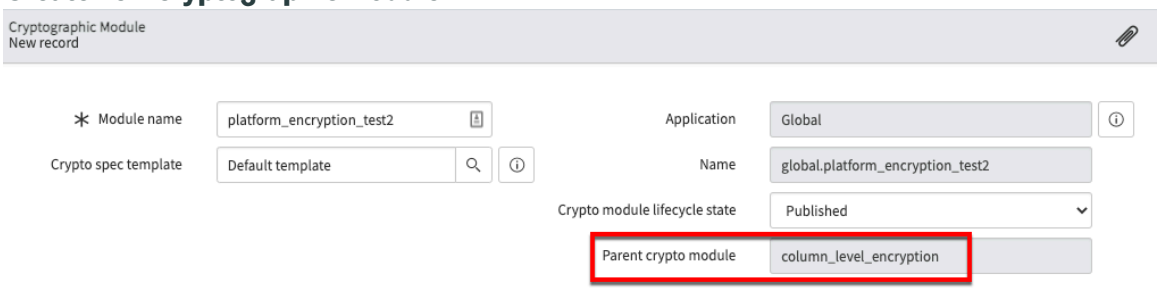

4. Complete the Cryptographic Module form as follows:

Cryptographic Module fields

Field	Description
Module Name	Enter a name for the module.
Crypto spec template	The default cryptographic template is selected.
Name	Auto-populates based on the module name and prepends the name with the scope to ensure which application is being applied. In this case, the global scope is applied.
Crypto module lifecycle state	Select **Published** to activate the crypto module.

Field	Description
Parent crypto module	The parent module **column_level_encryption** is selected automatically when using customer-supplied keys and encryption modules.

5. Select **Submit**.
6. Select the newly created cryptographic module from the table.
   In the **Crypto Specifications** related list, select the auto-generated key alias with the AES 256 CFB algorithm.

   The system populates the Crypto purpose and the Algorithm for Field Encryption automatically and jumps to the **Key Origin** stage.
7. Notice that **Upload customer supplied key** is the **Origin** and the **Key alias** is already populated.

Key origin

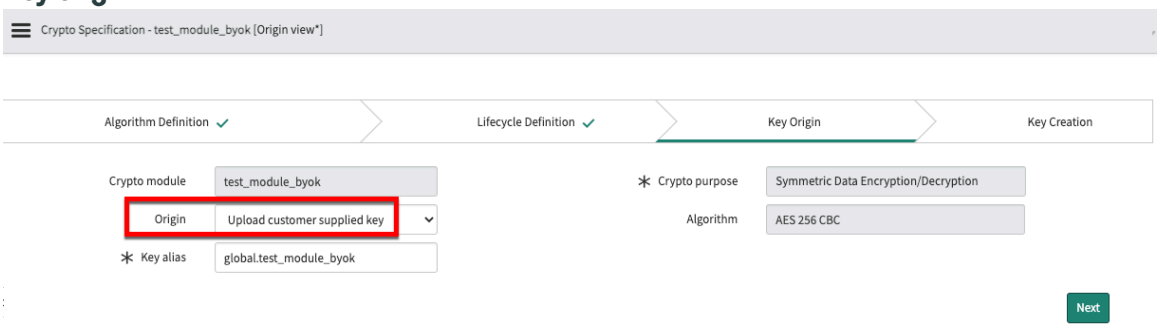

8. Select **Next** to move to the **Key Creation** stage.

   There are two links:

   - **Download wrapping key** downloads the key in a zip file containing an import token and a public key certificate, `.PEM` file. Use the import token to verify successful key wrapping according to security specification for the instance. Use the public key certificate `.PEM` file to wrap your customer supplied key securely before uploading it along with the token.
   - **Upload customer supplied key** opens the file browser to select the token and the encrypted key that you wrapped.

Key creation upload links

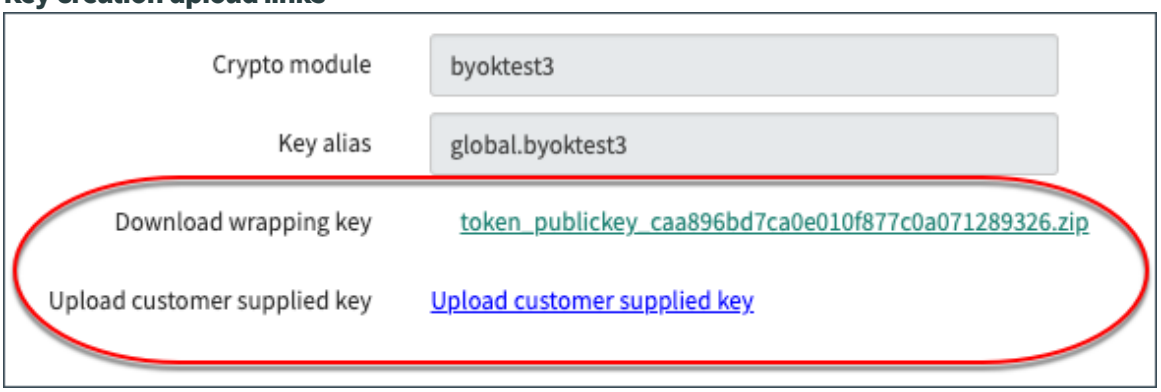

9. Select **Download wrapping key** to save the token.
   Save the token to the same destination location as the key is saved on your system. Don't rename the downloaded token.

10. Run the BYOK command on a terminal to wrap the key.
    For more information, refer to Wrap your customer-supplied key.
11. Select **Upload customer supplied key**.
12. Select **Browse** to select the two files, the wrapped key and the token file.
    The Attachments window displays the two files.

Wrapped key attachments upload

Select a file to remove and reupload, if necessary.

13. Select **OK**.
    You're returned to the Cryptographic Module screen. A confirmation message displays for a successful upload of the customer key. The key is also listed in the Module Keys related list.

Confirmation of key upload

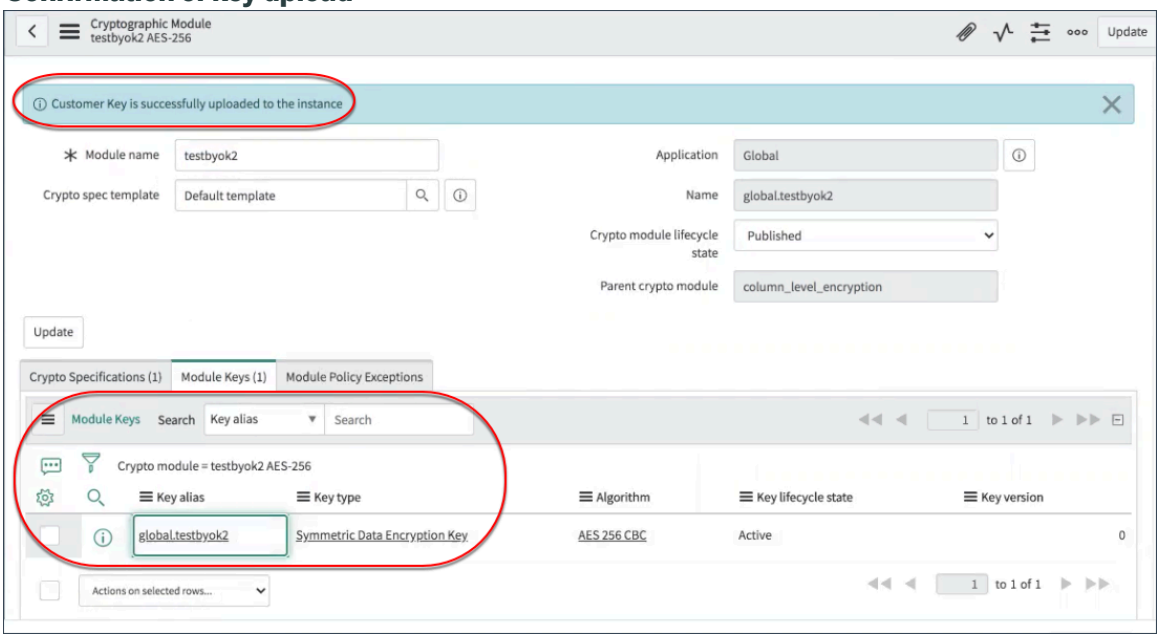

What to do next

Now that you have finished configuring your cryptographic module with your customer-supplied key, move on to Create a module access policy

Encrypting fields and attachments

Once cryptographic modules are created, a security admin can define the encrypted fields configuration (EFC) and opt to encrypt a field or attachment on a table.

How to encrypt fields

Note: Encrypted fields aren't audited by design. This behavior isn't configurable.

1. Specify the key source: ServiceNow generated keys or your customer-supplied keys (bring your own key) in **System Security** > **Field Encryption Settings**.
2. After specifying the key source, create a cryptographic module or use an existing cryptographic module. Start with Create a cryptographic module for instructions.

   Note: If you use customer-supplied keys, follow the directions in Create cryptographic module for Field Encryption and Configure properties for customer-supplied keys.

3. Create an encrypted field configuration to define where the encryption is applied. Here, you specify the target table and choose whether to encrypt a column or attachments within the table. See Set encrypted field configurations to get started.

Note: See Field Encryption Enterprise examples that illustrates how to encrypt fields and attachments using customer-supplied keys.

Set encrypted field configurations

Configure which table columns or attachments that the system encrypts using a preconfigured cryptographic module.

Before you begin

Role required: sn_kmf.cryptographic_manager and security_admin or elevate role to security admin.

About this task

Make sure you are in the correct application scope so you can see the tables in that scope.

Only users with access to the cryptographic module used in this configuration can read the data in the encrypted table column or access the attachment.

- If a user has write access but not read access, the field displays in edit mode and the data entered displays as asterisks.
- If a user has read access but not write access, the field displays the decrypted data in read-only mode.
- If a user has all access, both read/write functionality is available on the encrypted field.

See Create a cryptographic module or Create cryptographic module for Field Encryption to begin.

Important:

After encrypting a column, any new data inserted into the column is encrypted automatically. However, data that existed in the column before the encryption was active is not automatically encrypted.

In order to encrypt data that existed before the column was encrypted, you must run a separate mass encryption job. Learn more about mass encryption in Run mass encryption or decryption.

Procedure

1. Navigate to **All** > **System Security** > **Field Encryption** > **Encrypted Field Configurations** > **New**.
2. Select **New**.
3. Complete the form.

Field	Description
Type	**Column** to encrypt a table column or **Attachment** to encrypt all of a table's attachments. Types of data encrypted are: ○ String text (Full UTF-8) ○ Attachments ○ Date, Date/Time: **Note:** You can create encrypted field configurations to encrypt existing Date and Date/Time fields. You can add a new encryption configuration to a parent table only. You can't add a new encryption configuration to a child table. ○ URL ○ HTML ○ Journal ○ Translated
Table	Table whose fields or attachments are to be encrypted.
Column	Column (field) to be encrypted if you selected **column** as the type.
Active	Select to mark the configuration active. Deselect if the configuration isn't yet in use.
Crypto module	The cryptographic module that the encrypted field configuration applies to.
Method	Select **Single Module** to set the field configuration across one module. Select **Multiple Modules** for role-based access that spans across more than one cryptographic module. **Single Module** Use this option to encrypt all attachments using a single module. Your users need access to this module, otherwise they aren't able to upload attachments. **Multiple Modules** Use this option to allow users to choose a module when uploading attachments. Users with access to at least one module can select a module to use for encryption. Users with no module access can upload unencrypted attachments.

Field	Description
Algorithm Encrypted Preserving [read-only]	Indicates if the crypto module that you selected is already configured to support non-deterministic encryption. This means that if the same data is encrypted more than once, the encryption is different each time.

4. Select **Submit**.

Script access for cryptographic modules

Scripts can be run to access a cryptographic module policy for a cryptographic purpose.

For Key Management Framework, policies can be based scripts. When an access policy is triggered for script access, the backend script can execute the module policy actions from the script.

Cryptographic modules can support one or more encryption purposes, such as Asymmetric Data Decryption and Symmetric Data Decryption. Each cryptographic purpose requires the generation of an encryption key and defined cryptographic purpose.

Consider the following when executing an encryption script request:

- The referenced cryptographic purpose must be defined in the cryptographic module.
- An active generated key must exist for the cryptographic module.
- The Module Access Policy type must be set to **script**.

Check script version

When creating a module access policy that is set to the script type, there is an option available to validate the integrity of the script version being accessed. Only the assigned version of the script is allowed access to the encryption modules. When the **Check script version** check box is selected in the module access policy, anytime the script is run, the system performs a version comparison. If the script has been changed, the user is notified.

Check script version check box

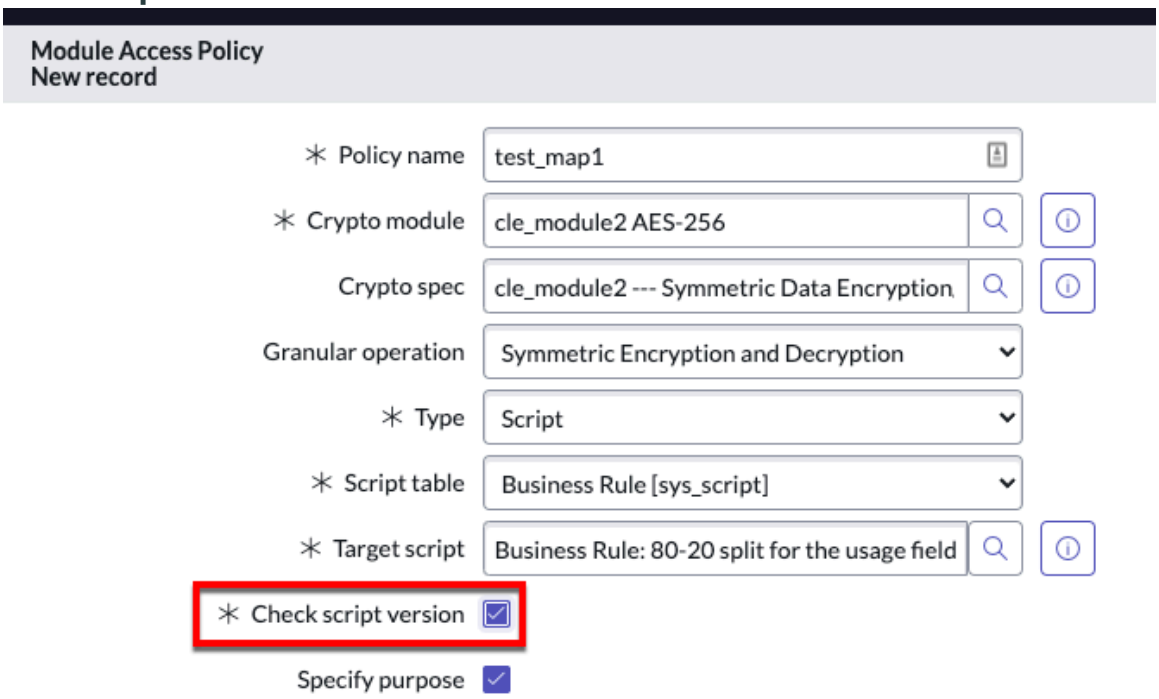

Configure script access to encrypted data

Execute a script to run the cryptographic module policy for a cryptographic purpose. Specific read (decrypt/unwrap) or write (encrypt, wrap) access can be defined based on the module access policy operation granularity.

Before you begin

Role required: sn_kmf.cryptographic_manager

About this task

Examples of uses are for Business Rules and Script Includes. This procedure uses a script for Business Rules.

Procedure

1. Create a cryptographic module with the symmetric data encryption/decryption algorithm. Refer to Create a cryptographic module for details. Specific access to the data or attachment is controlled with a module access policy with the following characteristics:
   - Symmetric encryption: The script is able to encrypt data but unable to decrypt the data.
   - Symmetric decryption: The script is able to decrypt uploaded encrypted data or attachment but unable to encrypt data or attachments.
   - Symmetric encryption and decryption: The script is able to both encrypt and decrypt data or attachments.
2. Navigate to **System Definition** > **Business Rules**.
3. Click **New**.

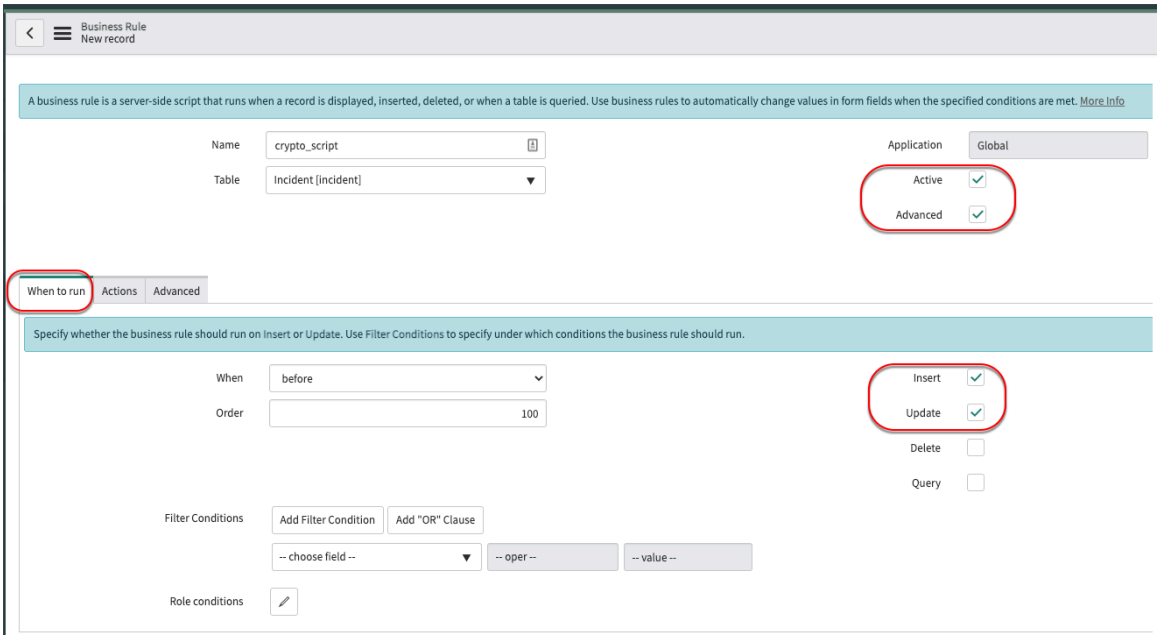

4. Complete the form on the **When to run** tab and enter the script on the **Advanced** tab:

Business Rule fields

Field	Description
Name	Enter a name for the business rule.
Table	Select **Incident [incident]**from the drop-down list.
Application	**Global** is selected by default.
Active	Mark the rule as **Active**.
Advanced	Select the check box to display advanced options.
When to run tab	On the **When to run** tab, enable **Insert** and **Update** fields.
Advanced tab	On the **Advanced** tab, paste the following script text at line 3: `var gc = global.GlideCryptoModule.getModule('global.acme_mod');` `var value = 'test';` `var encrypted = gc.encryptData(value);` `gs.info('value: ' + value);` `gs.info('Encrypted: ' + encrypted);` `var decrypted = gc.decryptData(encrypted);` `gs.info('Decrypted: ' + decrypted);` `gs.info(decrypted == value);` **Note:** Refer to the "Business Rules Advanced Tab" image for details.

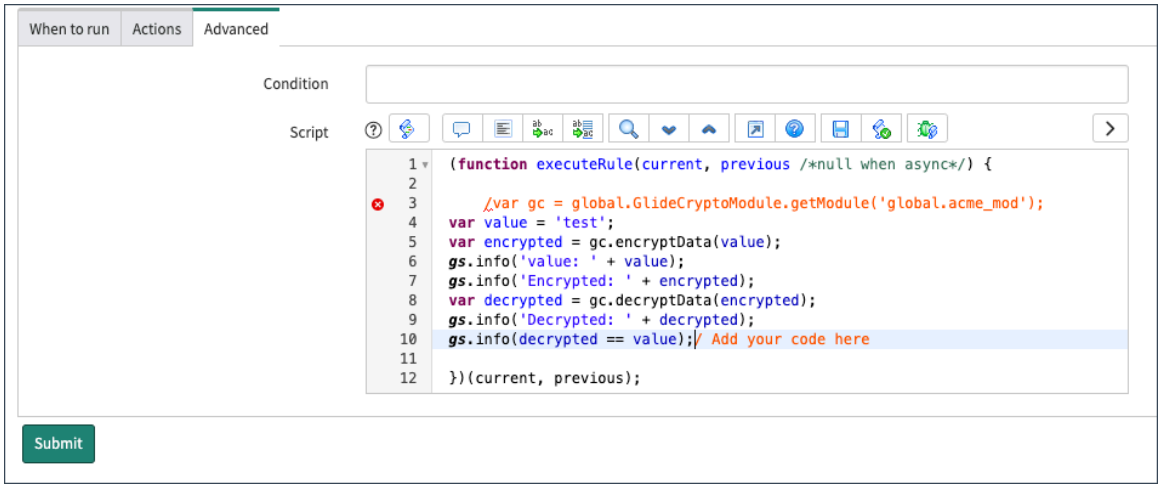

5. Click **Submit**.
6. Navigate to **Key Management** > **Module Access Policies** > **All**.

> **Note:** For additional information, refer to Create a module access policy.

7. Click **New**.
8. Complete the form.

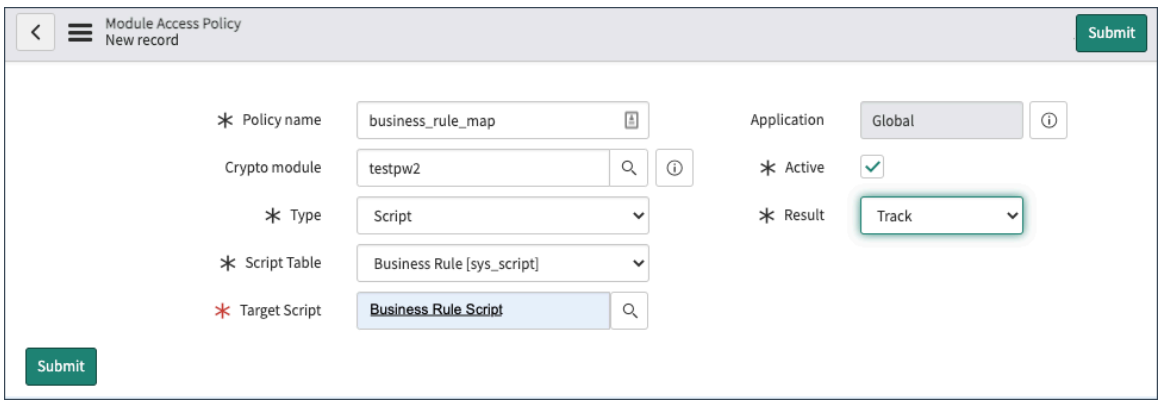

Module Access Policies fields

Field	Description
Policy name	Enter a name for the policy.
Crypto module	Click the search icon to select a module with the symmetric data encryption/decryption algorithm.
Type	Select **Script** to control access by script.
Script Table	Select a value from the script table drop-down list. For this example, select **Business Rule [sys_script]**.
Target Script	Select the script document for the policy. Select the **Table name** and then the related document for the policy. For this example,

Field	Description
	select the Business Rule that you created in previous steps.
Active	Select to activate the policy.
Result	To give the script access to the module, select **Track** in the **Result** field.

9. Click **Submit**.
   The Module Access Policy for the script is now available in the system.

View declined cryptographic module usage requests

View cryptographic modules that rejected encryption requests made by scripts because of unsupported encryption mechanisms.

Before you begin
Role required: sn_kmf.cryptographic_manager

About this task
Cryptographic modules can support one or more encryption purposes, such as Asymmetric Data Decryption and Symmetric Data Decryption. Encrypted data can only be accessed based on the module access policy. If a script tries to use a cryptographic module for a purpose not defined in the module, the script cannot access to the encrypted data.

In the following example, a cryptographic purpose was assigned to a cryptographic module, but a key was never generated for it.

Procedure

Navigate to **All** > **Key Management** > **Module Key Policies** > **Module Key Rejections**.

A list of cryptographic modules that rejected requests displays along with the encryption key used in the corresponding script.

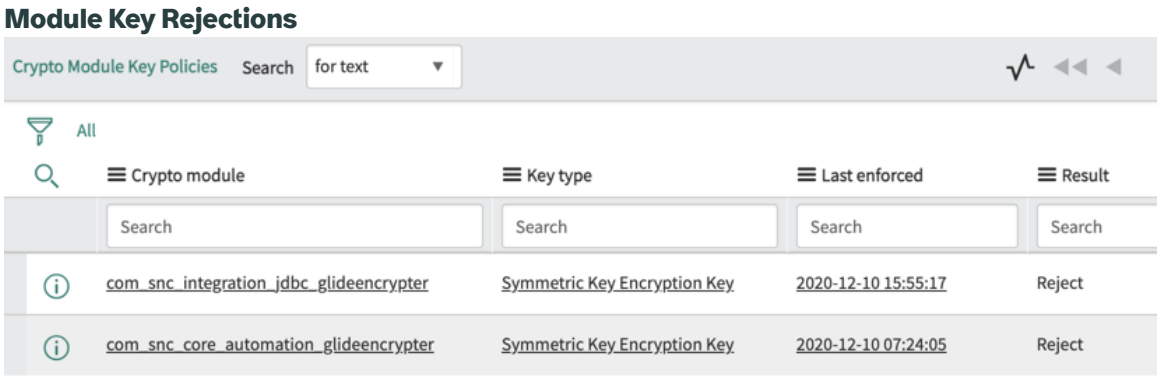

Note: If a different script attempts to use the same cryptographic module using the same key type, the value for **Last enforced** updates. Another row does not generate.

In this example, at 2020-02-10_15:55:17, the first module rejected a request because module1's key is compromised. At 2020-02-10_07:24:05, the second module rejected a request because the second module's key is suspended.

To grant scripts permission to use the encryption module the next time they run, create a module access policy for script encryption. For more information, refer to Configure script access to encrypted data.

Schedule mass encryption, decryption, and rekeying jobs

Schedule encryption, decryption, and rekeying jobs to run at a time that is best for your instance.

Before you begin

Encryption, decryption, and rekeying jobs can be time and resource intensive, so consider scheduling at non-peak hours. Also ensure that the user scheduling the job has the appropriate access for each job.

Role required: sn_kmf.cryptographic_manager

About this task

Mass encryption and decryption is also available from the Encrypted Field Configurations form. See Run mass encryption or decryption for instructions.

Procedure

1. Navigate to **All** > **System Security** > **Security Jobs**.
2. Click **New**.
3. Complete the scheduling form.

Field	Description
Name	Name of the encryption, decryption, or re-keying job.
Type	Job type: ○ **Key Migration Context to Module:** Mass migration of Encryption Context keys to Encryption Modules, including creation of Module Access Policies records for access controls on the Encryption Modules ○ **Data Migration Context to Module**: Migrates data encrypted by Encryption Contexts to Encryption Modules ○ **Mass Decryption Attachment**:Decrypts all encrypted attachments in records for a single table you define in the **Table** field. ○ **Mass Encryption Attachment**: Encrypts all attachments in records for a single table you define in the **Table** field. ○ **Mass Encryption**: Encrypts any pre-existing value in the defined column/field used in the Field Encryption Configuration ○ **Mass Decryption Module**: Decrypts any pre-existing value in the defined column/field used in the Field Encryption Configuration with Single Module. ○ **Mass Decryption Multi Module**: Decrypts any pre-existing value in the defined column/field used in the Field Encryption Configuration with Multiple Module. ○ **Mass Rekeying**: Re-encrypts any pre-existing value in the defined column/field used in the Field Encryption Configuration using the current active key for the module. ○ **Migrate Attachment Context to Module**: Encrypts any pre-existing attachment on the table defined in the Field Encryption Configuration. Any attachment previously encrypted with a context is re-encrypted with the module.

Field	Description
State	The initial job state is New. After the job has been executed as scheduled, the state will update accordingly.
Time window start	Start time for the job in 24-hour format.
Time window end	End time for the job in 24-hour format.
Table	Table to be encrypted or decrypted.
Field	Field to be encrypted or decrypted.
Summary	Job status information when the job is running, has completed, or has errors.

Note: Because of system overhead, you should schedule mass encryption, decryption, and rekeying jobs to run at non-peak hours. The ServiceNow AI Platform runs the job between the **Time window start** and **Time window end**. If the job is not complete in one processing window, it continues during the next specified processing window until all processing is complete.

4. Click **Submit**.
5. After you schedule a job, you can do the following.
   - Click **Cancel Job** to cancel a running job.
   - Click **Start** to start a job immediately.
   - Click **Update** to save any changes you make to the job schedule.
   - Click **Delete** to delete the scheduled job.

Run mass encryption or decryption

You can run mass encryption on encryption configurations, as well as a mass decryption to decrypt previously encrypted values.

Before you begin

Role required: security_admin

About this task

You can also create scheduled jobs for mass encryption and decryption. See Schedule mass encryption, decryption, and rekeying jobs for instructions.

Mass encryption and decryption are available only when an encrypted field configuration uses the single cryptographic module. Mass decryption is available for both the single and multiple encryption method.

Note: You should run mass encryption and decryption only during non-peak hours because the operations are resource and time intensive.

Procedure

1. Navigate to **All** > **System Security** > **Field Encryption** > **Encrypted Field Configurations**.
2. Open the encrypted field configuration for the field you would like to mass encrypt or decrypt.
3. Under Related Links, select an available option.

- **Schedule Mass Decryption job**
- **Schedule Mass Encryption job**

4. Confirm your selection in the dialog.

Result
If running a mass encryption, all values are encrypted with the encryption module defined in the encrypted field configuration record. If running a mass decryption, only fields encrypted with an encryption module you have access to are decrypted.

Upload attachments for encryption

Protect sensitive files by encrypting record attachments using Field Encryption and Row Conditions.

Before you begin
Role required: Any role that aligns with the module access policy (MAP) created by the admin.

Procedure

1. Navigate to a record.
2. Select the **Manage Attachments** paperclip icon.
3. In the **Attachments** window, select the module to encrypt with from the **Encrypt with Module** drop-down options.

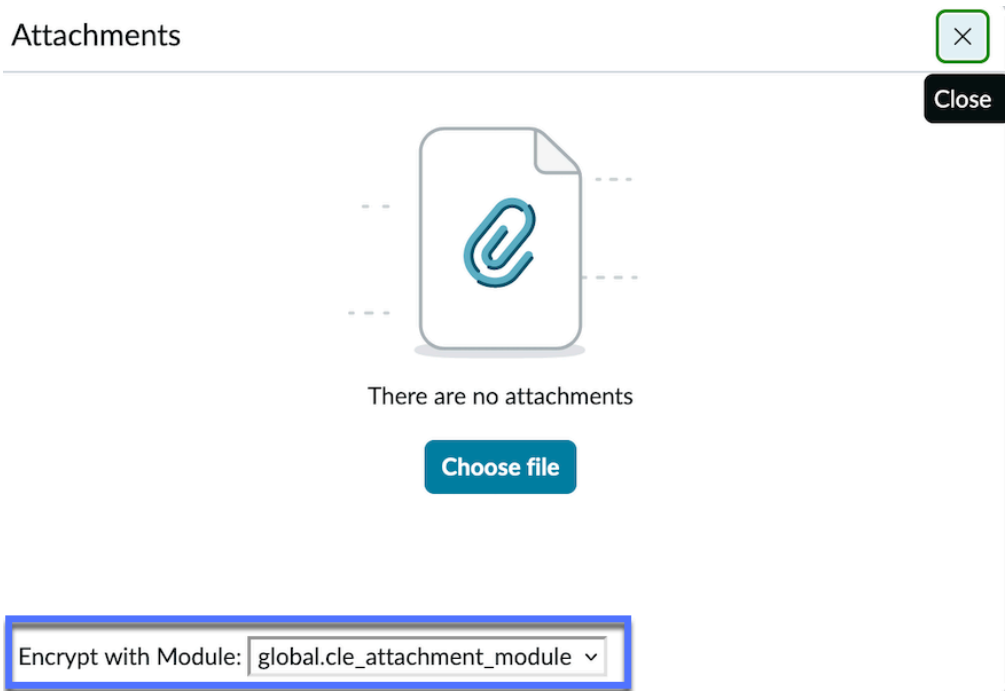

Note: If Row Conditions have been added to the Encryption Field Configuration, this option isn't displayed. The attachment would then be automatically encrypted with the module defined in the Row Condition. Similarly, this drop-down option only displays when using the Manage Attachments icon and not if the attachment is added by dragging and dropping.

4. Select **Choose file** to locate the attachment and select **Open** to attach it.

Result
Attached files display at the top of the form. Encrypted attachments are denoted by a lock icon. Only users with the encryption module to view encrypted files will see them listed.

Field Encryption Enterprise examples

These examples walk you through the encryption of fields and attachments using customer-supplied keys.

Field Encryption Enterprise walkthrough

This walkthrough shows you how to encrypt a field in your instance using Field Encryption Enterprise with the Key Management Framework (KMF). It also shows you how to use your own key.

Before you begin

Note: This procedure only applies to Field Encryption Enterprise functionality. See Activate Field Encryption for more information on obtaining Field Encryption Enterprise.

Role required: admin or security_admin

Note: security_admin is a privileged role, for details on using privileged roles, see Elevate to a privileged role

About this task
This walkthrough starts with an instance where you have already created and uploaded your personal cryptographic key. You could use the ServiceNow key, but this example uses a customer-supplied key.

After the key has been stored in a cryptographic module, you can start configuring fields in your instance, such as salary or social security numbers that have limited access from certain users. In the Encrypted Field Configuration, specify which authorized personnel can access sensitive data.

This task demonstrates two scenarios. One example encrypts the **Short Description** field in an Incident for users who are not authorized to view the sensitive data.

Attachments can also be encrypted and only visible to users who are granted access, or is visible to all users that are not restricted from viewing the data. See Attachment encryption walkthrough to encrypt an attachment.

Procedure
1. Make sure that Field Encryption Enterprise is enabled.
2. Create a cryptographic module for column_level_encryption.
   See Create cryptographic module for Field Encryption Create a cryptographic module for more information.
3. Navigate to **System Security** > **Encrypted Field Configurations**.
4. Click **New**.
5. On the form, fill in the fields.

Encrypted field configuration form

Field	Description
Type	**Column** is required to use your personal key.

Field	Description
Table	Table that stores the sensitive information. For this example, select **Incident [incident]**.
Column	Column, or specific information, that represents the sensitive date to be encrypted. For this example, select **short_description**.
Active	Option to mark **Active** to use the field configuration.
Algorithm Equality Preserving	The option is automatically selected.
Crypto module	Module that you created to use with the personal key.
Method	The **Single Module** option is used to apply the policies for one module. **Multiple Modules** is used to apply the policies across multiple modules.

Encryption field configuration example

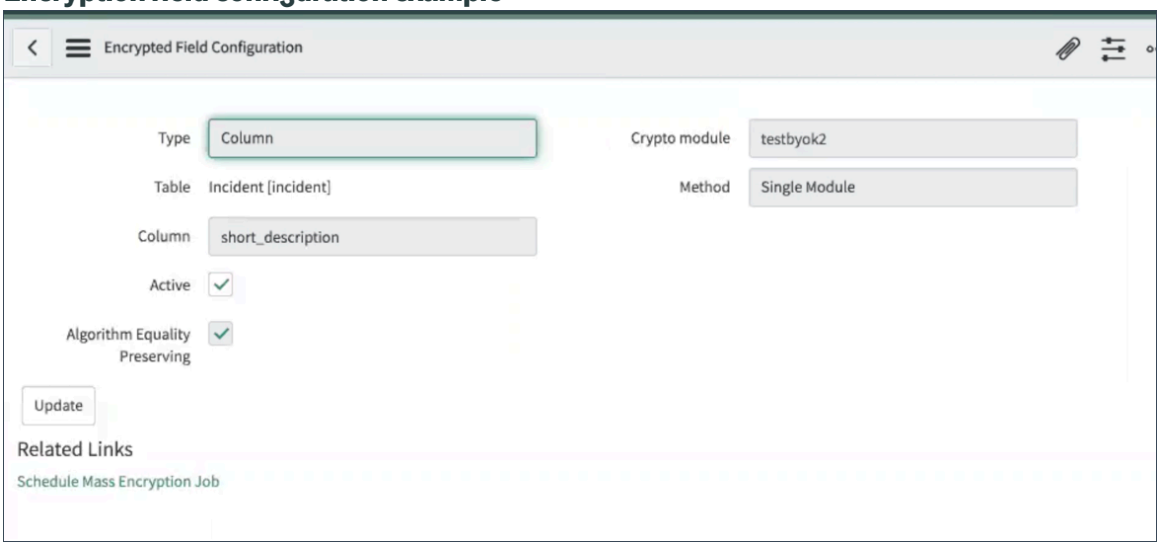

6. Click **Submit**.

   Establish a Module Access Policy to assign access to the cryptographic module. See Create a module access policy for additional information.

7. Navigate to **Key Management** > **Module Access Policies** > > **Create New** > **.**
8. On the form, fill in the fields.

Module access policy form

Field	Description
Policy name	Name for the policy, such, as `short description`.
Crypto module	Crypto module that you created to encrypt your key.
Type	Type of access designation for the crypto policy. Use **Role** to grant access to the encrypted field to only those users that have the assigned role.
Target Role	The role that has access to the encrypted field. For this example, select **Admin**.

Field	Description
Active	Option to activate the Module Access Policy.
Result	The **Track**option enables the access to the field for the selected role. (To restrict access to that field for the selected role, select **Reject** or**Strict Reject**.)

Module access policy example

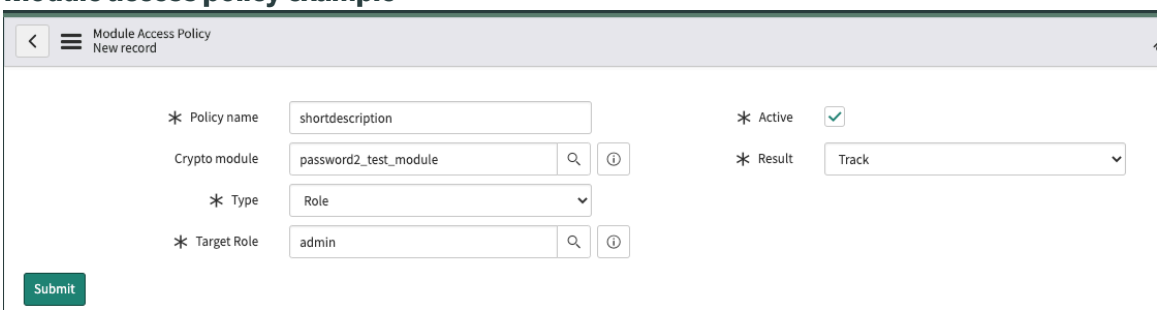

9. Click **Submit**.
10. As a user with the sn_kmf.admin role, navigate to **Incident** > **New**.

Example of encrypted field visible

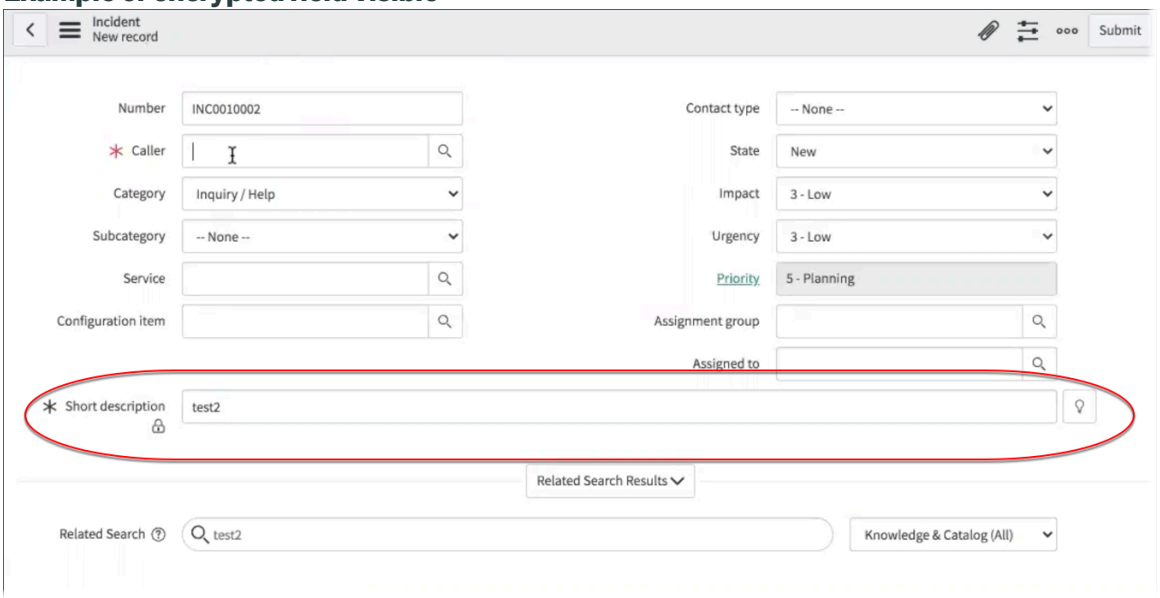

You can now view the Short description field based on the module access policy configuration.

Note: The sn_kmf.admin role was granted user access to the encrypted field, Short description, by setting the module access policy to **Track**. Notice the lock icon ( ) under the field name indicating that the field is an encrypted field.

You can now access the **Incidents** module as an end user to test the encrypted field configuration.

11. Log in as a user to be restricted from viewing the encrypted data in the configured field.

Encrypted field level data

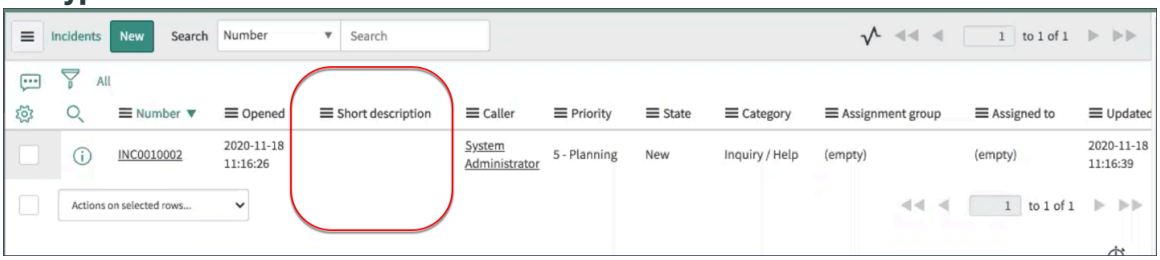

When you access the incident number, the data in the Short description will not be visible.

Result

You have successfully used your symmetric key to control access to a specific field using Field Encryption Enterprise.

Attachment encryption walkthrough

This walkthrough shows you how to encrypt an attachment in your instance using Field Encryption Enterprise with the Key Management Framework (KMF). It also shows you how to use your own key.

Before you begin

Note: This procedure only applies to Field Encryption Enterprise functionality. See Activate Field Encryption for more information on obtaining Field Encryption Enterprise.

Role required: kmf cryptographic manager

About this task

This walkthrough starts with an instance where you have already created and uploaded your customer-supplied cryptographic key. You could use the key, but this example uses a customer-supplied key.

Upload confidential attachments in your instance and limit access from certain users. Use Encrypted Field Configuration to specify which authorized personnel can access sensitive data.

We show you how to encrypt attachments to only be visible to users who are granted access, or be visible to all users that are not restricted from viewing the data. In this example, we restrict a certain role from being able to access an attachment in the **Incidents** module.

Note: Although you can use multiple modules with Field Encryption Enterprise, attachment encryption must use single modules.

Procedure

1. Make sure that Field Encryption Enterprise is enabled.
2. Create a cryptographic module.
   See Create cryptographic module for Field Encryption for more information.
3. Navigate to **System Security** > **Encrypted Field Configurations**.
4. Click **New**.
5. Complete the form:

Encrypted field configuration fields

Field	Description
Type	Select **Attachment** to use your personal key for encrypting an attachment from the selected **Table** For this example, select **Incident**.
Table	Select the table to access the sensitive information. For this example, select **Incident [incident]**.
Active	Mark **Active** to be able to use the field configuration.
Algorithm Equality Preserving	When selecting Field Encryption Enterprise, this field is visible based on the table selected.
Crypto module	Select the module that you created to use with the personal key.
Method	The **Single Module** option is used to apply the policies for one module. **Multiple Modules** is used to apply the policies across multiple modules.

Encrypted Field Configuration table

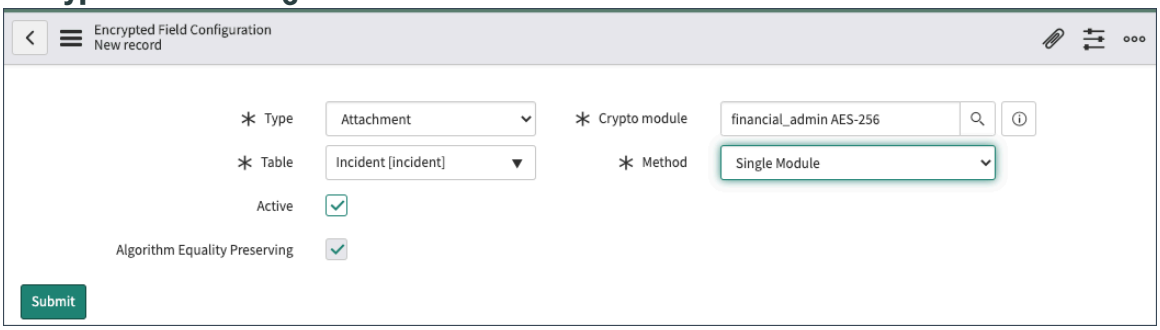

6. Click **Submit**.

   Establish a Module Access Policy to assign access to the cryptographic module. Refer to Create a module access policy for additional information.

7. Navigate to **Key Management** > **Module Access Policies** > **All**.
8. Click **New**.
9. Complete the form:

Module access policy fields

Field	Description
Policy name	Enter a name for the policy, such, as "Attachment policy."
Crypto module	Select the crypto module that you created to encrypt your key.
Type	Select **Role** to restrict access to the encrypted field from users with the assigned role.
Target Role	Select the role that will not have access to the encrypted field. For this example, select **itil**.
Active	Select this check box to be able to use the Module Access Policy.

Field	Description
Result	Select **Strict Reject** to control the access to the attachment from the selected role. (To grant access for the selected role, select **Track**.)

Module Access Policy form

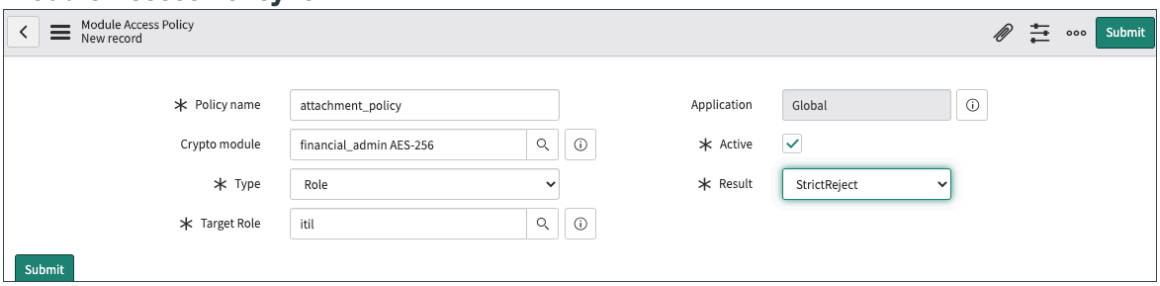

10. Click **Submit**.
11. As admin or as a person that created the incident, navigate to **Incidents** and add an attachment to **Activities** on the **Notes** related list.

Attachment available per role

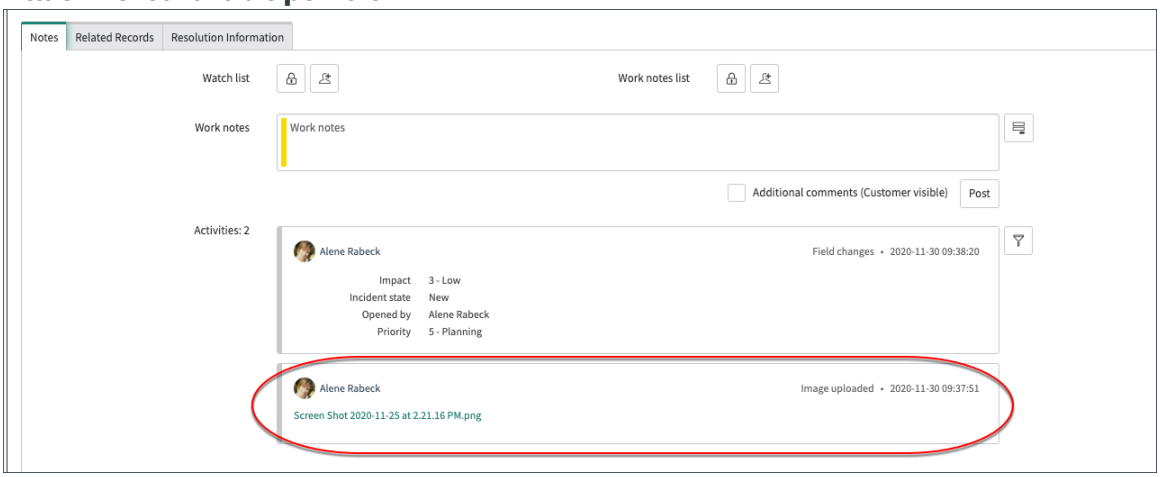

12. Log in as a user that restricted from accessing the encrypted attachment.
13. Open the incident and scroll to the **Activities:** section.
    The link to open the attachment is not accessible for users with the restricted role.
14. You have now successfully used your customer-supplied key to control access to a specific attachment using Field Encryption Enterprise.

Field Encryption Enterprise

Field Encryption Enterprise uses the Key Management Framework (KMF) to enable you to customize and manage how fields and attachments are encrypted and decrypted on your instance. A subscription is required to use Field Encryption Enterprise.

Important: This topic covers the enterprise version of Field Encryption. For information on the standard version of Field Encryption, or to learn the differences between the two versions, see Exploring Field Encryption.

Field Encryption Enterprise is premised with Field Encryption and uses the Key Management Framework and its full support of key management functions. Field Encryption Enterprise provides key-protection and key life-cycle management for application-level field encryption. All

keys are protected with a key-wrapping hierarchy ultimately rooted in FIPS (Federal Information Processing Standards) 140-2-L3 Hardware Security Modules (HSM).

Field Encryption Enterprise gives you the ability to manage how supported fields are encrypted and decrypted in accordance with NIST 800-57 practices. It also uses the most updated version of field-level encryption, including integration for proper key protection and management.

Specifically, Field Encryption Enterprise uses the KMF encryption modules, granting you more control of server-side encryption. KMF verifies proper data encryption key protection using key hierarchy and envelope encryption. Your instance encrypts data through cryptographic modules that you configure. You can create an access policy for each module then configure cryptographic specifications and access policies and control key life-cycle management control.

Field Encryption Enterprise supports module access policies based on:

- Scope
- Role
- Script
- Resource Exchange
- System User

See Create a module access policy for additional information.

Note: For details on the supported features of Field Encryption and how to upgrade and subscribe to the Field Encryption Enterprise entitlement refer to Encryption and Key Management subscription bundle.

Encryption terms

Term	Description
Key management 	Support for key management Fundamental to Field Encryption Enterprise is the Key Management Framework (KMF). Gain the following capabilities: • Key life-cycle management. • Key rotation. See Rotate keys for details. • Key protection and key generation with FIPS 140-2-L3 Hardware Security Modules (HSMs). • Segregation of roles and duties. • The secure transfer of data encryption keys between instances, such as production and non-production instances. • Customer Supplied Keys (CSK) with key-wrapping. • Non-deterministic encryption. • Mass encryption/decryption. • Auditing of key access/use.

Term	Description
	See Key Management Framework Reference for details.
Customer-supplied key 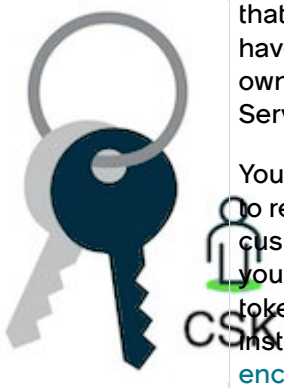	Support for customer supplied keys One of the biggest benefits of Field Encryption Enterprise is that you can use your own keys for encryption. Administrators have the choice to use ServiceNow supplied keys or your own customer-supplied keys (CSK) for encryption on the ServiceNow AI Platform®. You can also manage the key life cycle and decide when to revoke, rotate, and inactivate the keys. After you enable customer-supplied keys and create a cryptographic module, you download a token and public ephemeral key. You use the token and public key to wrap your key and then upload to the instance. To use customer-supplied keys, see Configure field encryption settings to select key type and Using customer supplied keys with Field Encryption Enterprise.
Field Encryption 	Support for both field encryption and attachment encryption Both field encryption and attachment encryption use cryptographic modules and access policies through Encrypted Field Configurations. The Encrypted Field Configuration form is used to choose an encryption type of column or attachment encryption. See Set encrypted field configurations for more information and supported field types.
Non-deterministic encryption 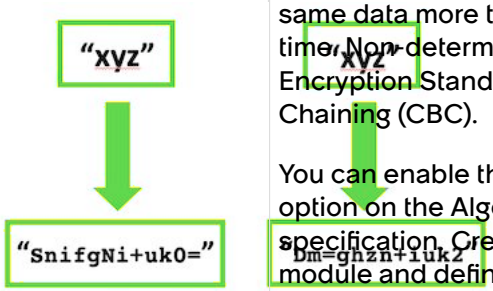	Support for non-deterministic encryption Field Encryption Enterprise supports non-deterministic encryption for enhanced security. If the system encrypts the same data more than once, the ciphertexts are different each time. Non-deterministic encryption is available with Advanced Encryption Standard (AES) encryption with Cipher Block Chaining (CBC). You can enable this feature through the Equality Preserving option on the Algorithm Definition stage of the cryptographic specification. Create a cryptographic specification for a crypto module and define an algorithm for encryption and generate the key. See Create a cryptographic module to define the mechanisms used for cryptographic operations and for more information on enabling non-deterministic encryption.

Term	Description
Resource Exchange 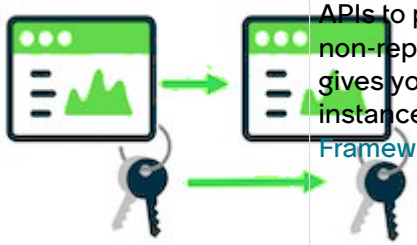	Resource Exchange Field Encryption Enterprise keys instance to instance in a secure manner using the KMF cryptographic APIs to provide confidentiality, integrity, authentication, and non-repudiation. Resource Exchange is a KMF feature that gives you the capability to exchange resources between instances in a secure manner. See Key Management Framework Resource Exchange for details.

Note: If you choose not to activate Field Encryption Enterprise, you can still use Field Encryption. See Exploring Field Encryption for information.

Field Encryption Enterprise supports on-premise customers. It doesn't support Domain Separation.

Support for additional encrypted fields

The standard version of Field Encryption is limited to five encrypted columns. Field Encryption Enterprise supports an unlimited number of encrypted columns.

Supported field information

The following field types can be encrypted:

- Attachments
- Date
- Date/Time
- Email
- HTML
- Journal
- Journal Input
- Journal List
- Phone
- String text
- Translated Field
- Translated HTML
- Translated Text
- URL

Attachment Encryption

Attachment encryption by default

Customers using Field Encryption have attachments encrypted by default in tables that have an active Encrypted Field Configuration (EFC) type of *Attachment*.

This default encryption defined by the EFC configuration means that it's not necessary for admins to manually declare that an attachment should be encrypted on upload for these tables.

Administrators can disallow users from attaching unencrypted files

For details, see Prevent users from attaching unencrypted files.

Opt out of default encryption

If you don't want attachments encrypted by default based on EFC configuration, you can opt out of this option by contacting ServiceNow support.

To opt out of this feature, create a support case with ServiceNow support, and include this statement in a comment on the case record:

```
"I [customer name], understand that I am asking ServiceNow to turn off a recommended security best practice for attachments, and that [customer company] assumes any additional risk related to their configuration and use of unencrypted attachments in the ServiceNow application."
```

API support

Field Encryption Enterprise enables the following APIs.

Note: The API behavior described in the following table represents the default configuration for the latest base system package. If you are working with older package versions, you may experience different functionality.

Field Encryption APIs

API	Description	Parameters	Return type
`changeEncryptionContext()`	Updates an active Encryption Context (EC) used to encrypt an attachment. When CLE is enabled with the CLE Starter plugin using KMF Crypto Module (CM), the API locates the CM for the EC and uses it to encrypt the attachment. **Note:** This API is only available in the Global scope.	• sourceTable – Name of table that has the attachment. • sourceID – Table record system id. • attachmentID – The sys_attachment record system id. • newEncryptionContextID - System ID of the new context.	Boolean
`changeCryptoModule()`	Updates an active encryption module used to encrypt an attachment.	• sourceTable – Name of table that has the attachment. • sourceID – Table record system id.	Boolean

Field Encryption APIs (continued)

API	Description	Parameters	Return type
	Note: This API is only available in the Global scope.	• attachmentID – The sys_attachment record system id. • newCryptoModuleId - System ID of the new encryption module to encrypt the attachment.	
`disableEncryption()`	Disable active encryption on an attachment.	• sourceTable – Name of table that contains the attachment. • sourceID – Table record system id. • attachmentID – The sys_attachment record system id.	Boolean
`getDisplayValue()`	Returns the cleartext display value of an encrypted field.		String
`getValue()`	Returns the cleartext value of an ecrypted field when glide_encryption.set_value_support_cle.disabled is false (requires Module Access Policy (MAP)). Returns the encrypted value of an encrypted field when glide_encryption.set_value_support_cle.disabled is true.		String
`setDisplayValue()`	Inserts encrypted data into an encrypted field for display purposes.	• name – Field name. • value – Field value.	Boolean
`setValue()`	Inserts encrypted data into an encrypted field, controlled by a system property. Encrypts data when glide_encryption.set_value_support_cle.disabled is false (requires	• name – Field name. • value – Field value.	Boolean

Field Encryption APIs (continued)

API	Description	Parameters	Return type
	MAP); writes unencrypted data when set to true (no MAP required), when glide_encryption.set_value_support_cle.disabled is true.		

The following script illustrates API changes when the Incident short description is encrypted:

```
var gr = new GlideRecord('incident'); //creates a new incident
gr.setValue('short_description','test123'); //sets the value
 to test123
var sys_ID = gr.insert(); //inserts the record in the Incident
 table.
gs.info(gr.getValue('short_description')); //displays the
 unencrypted value
```

When the Field Encryption plugin is installed, `glide_encryption.set_value_support_cle.disabled` is set to false by default.

Column Level Encryption

Column Level Encryption permits and denies access to encrypted data based on user role. Column Level Encryption includes basic key management using encryption modules.

Important:

Starting with the Zurich release, Column Level Encryption (CLE) and Column Level Encryption Enterprise (CLEE) are being prepared for future deprecation. They will be hidden and no longer activated on new instances but will continue to be supported. Field Encryption and Field Encryption Enterprise provide the latest experience for this functionality.

For details, see the Deprecation Process [KB0867184] article in the Now Support knowledge base.

With Column Level Encryption, you can encrypt specific fields within your tables, as opposed to encrypting the entire table or database. Use this method to help ensure that your sensitive data remains protected without the need to encrypt and entire table. The ability to encrypt only the portions of your tables that require it helps to reduce the time spent encrypting and decrypting data.

Column Level Encryption grants access to encrypted data based on a user's role. Because of this approach, users must be associated with a role to view data encrypted by Column Level Encryption. Users can be associated with a role directly, or they can be assigned to a group that is associated with a role. This role-based approach simplifies the process of making sure that your data is visible only to users who need it.

Role-based encryption example

In this example, you can see four users attempting to access data stored in two fields on a form. These fields are encrypted by an encryption context, which is only accessible to users who are associated with a specific role (Role 1).

- User 1 is a member of Role 1, which provides access to encryption module 1. User 1 can see the contents of Field A and Field B.
- User 2 and User 3 are members of Group 1. Group 1 is a member of Role 1, which enables everyone in Group 1 access to encryption module 1 and enables User 2 and User 3 to see the contents of Field A and Field B.
- User 4 isn't a member of any group or role and has no access to encryption module 1. User 4 does note access to Field A or Field B. User 4 also doesn't see these fields on a form. In a list view, these fields are visible, but the values are be empty.

Get started

Explore	Configure	Use
		
Learn the benefits of the Standard and Enterprise editions of Field Encryption.	Learn how to activate and configure Column Level Encryption Enterprise, and manage migration from Encryption Support	Use Field Encryption to manage access to encrypted data on your instances

Troubleshoot and get help

- ServiceNow community developer forum
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Column Level Encryption

Learn more about Column Level Encryption.

Column Level Encryption overview

Column Level Encryption is a base system feature that permits encryption of data stored within an instance using AES128, or AES256.

Column Level Encryption enables you to encrypt selected database fields and stored files through the use of encryption contexts. In these contexts you define what is encrypted, choose which algorithm to use, and supply the encryption key, which is stored within your instance.

After the context is created, you can associate it to a user role. Users assigned to this role, either directly of through a group, are able to access the encrypted data.

Because Column Level Encryption bases access to data on role assignment, it's important to be familiar with administering roles on your instance. For more information, see Managing roles.

Field Encryption benefits

Benefit	Feature	Required Roles
Configure access to your encrypted data based on assigned user roles.	Role-based access to encrypted data	security admin
Protect your data using the Advanced Encryption Standard (AES). You can choose to use either the AES-128 or AES-256 encryption algorithms.	AES Encryption	security admin
Create up to 5 modules and module access policies (MAP)s using the standard version of Column Level Encryption. MAPs expand on role-based access to allow considerations for: • System users • Scripts • KMF Resource Exchange Column Level Encryption Enterprise supports additional MAPs.	Support for up to 5 modules and module access policies (MAP)s	security admin
Encrypt common field types using the standard version of Column Level Encryption. Column Level Encryption Enterprise supports additional field types.	Encryption for String text, Date and Date/Time fields, attachments, and URLs	security admin
Choose between standard and equality preserving encryption. When enabled, equality preserving encryption ensures that the encrypted value of a field is the same when the field value remains the same. This type of encryption enables equality comparisons and group by operations on a field. **Note:** Non-deterministic encryption isn't supported.	Equality preserving encryption support	security admin
Use `getDisplayValue()` and `setDisplayValue()` APIs to return cleartext values and insert encrypted data for encrypted fields.	`getDisplayValue()` and `setDisplayValue()` APIs	security admin, developer

Column Level Encryption Enterprise benefits

Column Level Encryption Enterprise builds on the existing Column Level Encryption framework and provides these additional features after you purchase a subscription.

Benefit	Feature	Required Roles
Encrypt additional field types.	Support for additional field types: • HTML • Journal • Translated	security admin
Column Level Encryption Enterprise supports more than 5 modules and module access policies to provide more options for access to secured data.	Support for additional modules and MAPs	security admin
Keys from a key vault can be rotated on an automated schedule you configure. Using automatic key rotation can improve security while reducing administrative overhead.	Configurable automatic key rotation	security admin
Manage the full life cycle of your data encryption keys. Optionally, you can securely exchange data encryption keys generated within your environment.	Customer supplied keys	security admin
Ephemeral keys are cryptographic keys that are generated for each execution of a cryptographic process. These keys more secure because they're generated for use in a single session.	Ephemeral cryptographic keys	security admin
Updated `setDisplayValue()` and `setDisplayValue()` APIs can insert encrypted data for encrypted fields.	Updated `getDisplayValue()` and `setDisplayValue()` APIs	security admin, developer

Column Level Encryption Guided Tour

The tour gives a brief overview of the Column Level Encryption setup needed to encrypt table fields or attachments. Steps for the creation of Field Encryption Modules, Module Access Policies, and Encrypted Field Configurations are also covered. The tour includes links to detailed documentation and the ServiceNow University Column Level Encryption Overview course.

Before you begin

Role required: sn_kmf.crypto_manager or security_admin

Note: The guided tour is not yet available when Next Experience is enabled.

Procedure

1. Navigate to the **$pa_dashboards_overview.do** UI page on your instance.
   For example: `http://yourcompany.servicenow.com/nav_to.do?uri=%2F$pa_dashboards_overview.do` Replace `yourcompany` with the name of your instance.
2. Open the **Show Help** icon (in the upper left corner.
3. Select **Take a Tour** at the bottom of the sidebar.

4. Select **Next** to progress through the tour.

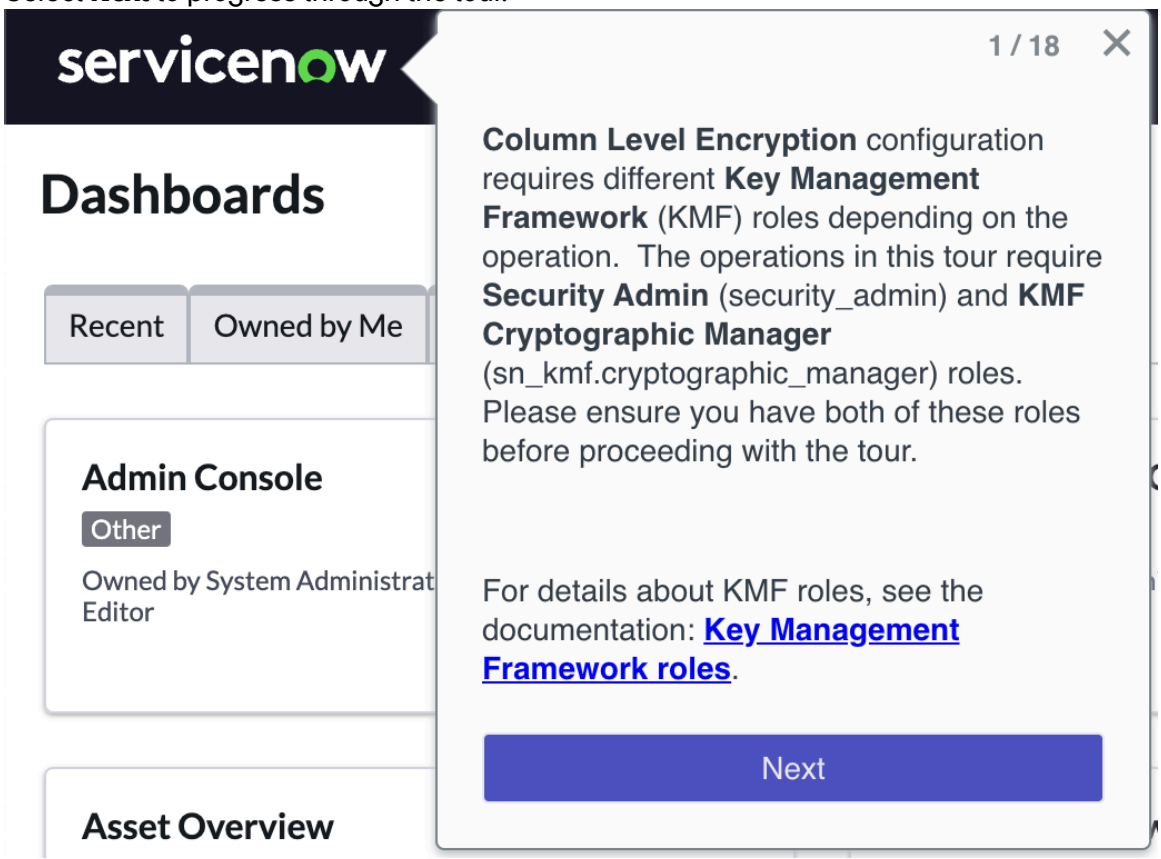

Configuring Column Level Encryption

Learn how to activate and configure Column Level Encryption and manage migration from Encryption Support.

Activate Column Level Encryption Enterprise

Learn how to active either Column Level Encryption or Column Level Encryption Enterprise.

Roles Required for Configuring Field Encryption

Learn about the roles required to configure Column Level Encryption.

Migration from encryption support

Use Scheduled jobs to migrate your keys and encrypted data from legacy Encryption Support to Column Level Encryption Enterprise. See details for this process at Migrating to Field Encryption

Change attachment encryption settings

Improve security by preventing users from attaching unencrypted files. For details, see Prevent users from attaching unencrypted files.

Activate Column Level Encryption Enterprise

With subscription to Column Level Encryption Enterprise, an admin can activate the com.glide.now.platform.encryption plugin.

Before you begin

Role required: admin

To purchase a subscription, contact your ServiceNow account manager. The account manager can arrange to have the plugin activated on your organization's production and non-production instances, generally within a few days.

About this task

Activating the Column Level Encryption Enterprise plugin (`com.glide.now.platform.encryption`) makes the following changes to your instance:

- The Encryption Support (com.glide.encryption) plugin is also activated.

  Note: The Key Management Framework (com.glide.kmf.global) plugin is already active by default.

- The `glide_encryption.set_value_support_cle.disabled` property is set to **false**, which turns on the SetValue feature. The SetValue support enables both `setDisplayValue()` and `setValue()` APIs to support encrypted data. It also enables both `getDisplayValue()` and `getValue()` to return clear text values.
- Two scheduled jobs are enabled:

  - **autoKeyMigration**: Migrates encryption context keys to Key Management Framework (KMF) cryptographic module keys.
  - **autoDataMigration**: Migrates data that you already encrypted to use the KMF cryptographic module key.

  Administrators can modify when these scheduled jobs run, and can pause or restart them at any time.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the plugin, *`Platform Encryption - com.glide.now.platform.encryption`* using the filter criteria and search bar. You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install**, and then in the Activate Plugin dialog box, select **Activate**.

   Note: When domain separation and delegated admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

Migrating to Column Level Encryption Enterprise

Scheduled jobs migrate your keys and encrypted data from Encryption Support to Field Encryption Enterprise.

You can review the scheduled jobs by navigating to **System Security** > **High Security Settings** > **Security Jobs**:

- **autoKeyMigration**: Migrates encryption context keys to Key Management Framework (KMF) cryptographic module keys.
- **autoDataMigration**: Migrates data that you already encrypted to use the KMF cryptographic module key.

You can modify when these scheduled jobs run, and can pause or restart them at any time.

Verify that the encrypted field configurations are using your newly migrated module keys by navigating to **System Security** > **Field Encryption** > **Encrypted Field Configurations**. Look for the following items:

- The **Method** field is **Single Module**.
- The **Crypto module** field is populated with the name of the cryptographic module that the system automatically creates. You can review that module and the module access policy, both of which are active and published.

Column Level Encryption Enterprise and system clones

If Column Level Encryption Enterprise is installed on your instance, a new field encryption module encryption key is automatically generated on the target clone instance as a part of clone process. These keys are generated for all modules to which the user has access, and that does not have a key already.

Because of this, field encryption modules on the target clone instance may have two module encryption keys present:

- An active module encryption key. This is the new key generated after clone, as long as the module is accessible to the user and has no prior keys.
- A deactivated encryption module key (from the automated key exchange transfer

  The active module encryption key is used to encrypt inserted data as needed on the target clone instance. The deactivated module is used to decrypt existing data that was cloned over as part of the system clone.

  To use a single key to decrypt and encrypt all data, you can run a module rekeying job. For more information about module rekeying jobs, see Schedule mass encryption, decryption, and rekeying jobs.

Column Level Encryption migration status page

Use the migration status page to track the migration of encryption contexts to encryption modules.

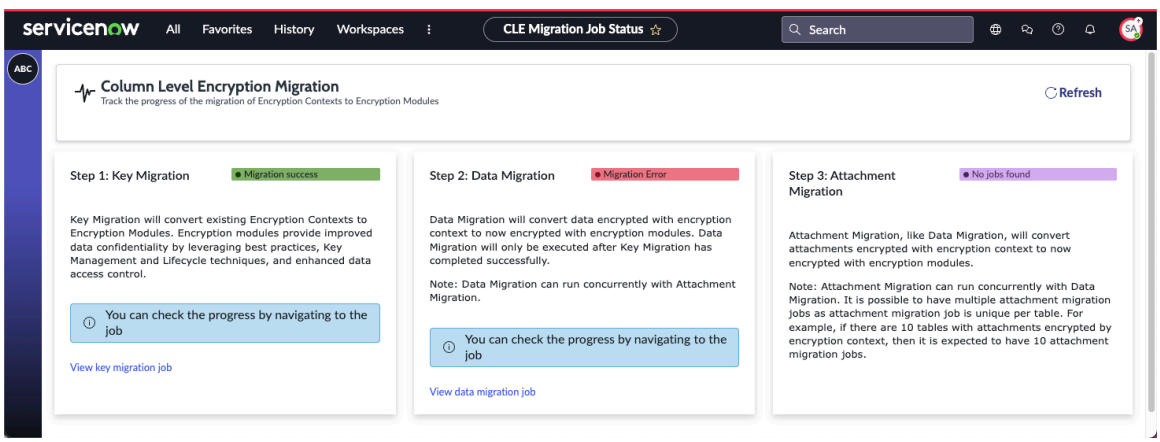

The Column Level Encryption Migration page displays the status of the steps involved in migrating encryption contexts to encryption modules. Each of the three sections displays the status of a specific step in the process.

Page section cards

The page contains three cards representing the steps in the migration progress. These cards display:

- (1) The status of the current step. This status will display whether the step has completed successfully, or if there are no jobs to process.
- (2) A description of the listed step.
- (3) A link to the relevant encryption job [sys_mass_encryption_job] record.

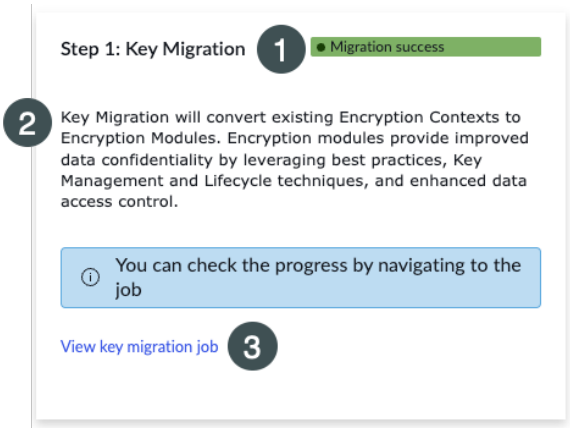

Prevent users from attaching unencrypted files

Modify the com.glide.encryption.enable_attachment_key_ui property to prevent your users with access to an encryption module key from attaching unencrypted attachments.

Before you begin

Role required: security_admin

You must elevate to the security_admin role performing these steps. For instructions, see Elevate to a privileged role

By default, users who have access to an encryption module key are able to upload unencrypted attachments. Use the *com.glide.encryption.enable_attachment_key_ui* system property to change this behavior.

When attaching, your users see a UI picker on records that have a multi-module encrypted field configuration. When this property is set to false, users no longer see an option not to encrypt an attachment.

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. In the system properties list, find and open the system property.
3. Set the **value** of the property to `false`.

Using Column Level Encryption

Use Column Level Encryption to manage access to encrypted data on your instances.

Use the related links to find information on common Field Encryption tasks.

Related topics

Create cryptographic module for Column Level Encryption

Create a cryptographic specification for Column Level Encryption

Configure advanced algorithms for Column Level Encryption Enterprise

Configure properties for customer-supplied keys

Encrypting fields and attachments

Column Level Encryption Enterprise

Create cryptographic module for Column Level Encryption

Create a Column Level Encryption cryptographic module to define the mechanisms used for cryptographic operations.

Before you begin

Role required: sn_kmf.cryptographic_manager or sn_kmf_admin, security_admin, admin

About this task

This procedure describes options that are available with Column Level Encryption with the base system and additional configuration options that become available with Column Level Encryption Enterprise functionality. Column Level Encryption Enterprise is available with a paid subscription. Refer to Encryption and Key Management subscription bundle for supported features and options available with each offering. See Activate Column Level Encryption Enterprise for more information on obtaining Column Level Encryption Enterprise.

Procedure

1. Navigate to **All** > **System Security** > **Field Encryption Modules** > **New**.

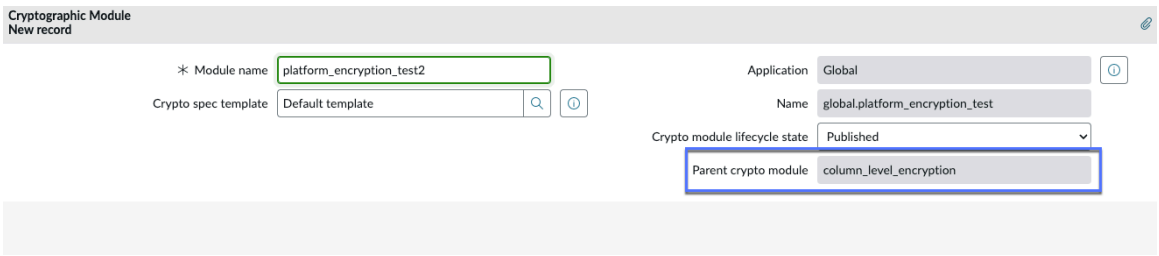

2. On the form, fill in the fields.

Cryptographic Module form

Field	Description
Module name	Alphanumeric string to be referenced when running scripts.
Crypto spec template	Default template used to create the cryptographic module that contains mappings of many crypto purposes to crypto specifications and recommended algorithms.
Application	The selected application scope.
Name	Encryption module name is prepended with application scope name to avoid conflict with other scoped applications on module creation. For example, if you created a module with the name `my_crypto_module` in the global application scope, the name is saved as `global.my_crypto_module`.
Crypto module lifecycle state	The term lifecycle refers to the creation, use, and deactivation of a cryptographic module. Set to **Draft** initially during configuration. When using the module, set to **Published**. The Default template is automatically set to **Published**.

Field	Description
Parent crypto module	The parent is populated automatically as **column_level_encryption**.

3. Click **Submit**.

   After submitting successfully, your cryptographic module is listed in the Cryptographic Modules table.

   Warning:

   For legacy encryption support users:

   If you're using the non-enterprise version of Column Level Encryption, you're limited to five fields. If you've exceeded this limit, you receive the following warning:

   This insertion exceeds the number of published fields limit for Column Level Encryption #entitled with the Subscription Product. The Enterprise subscription for Column Level Encryption is required for additional ields. Please reach out to your Account team.

   A default cryptographic specification is created with the crypto purpose set to Symmetric Data Encryption/Decryption and the algorithm as AES 256 CBC. Select the algorithm for updates.

4. To open the configuration options, click the newly created cryptographic module.

   Note: A maximum of five Column Level Encryption fields are allowed before upgrading to Column Level Encryption Enterprise. An error message displays and you are prevented from adding additional cryptographic modules.

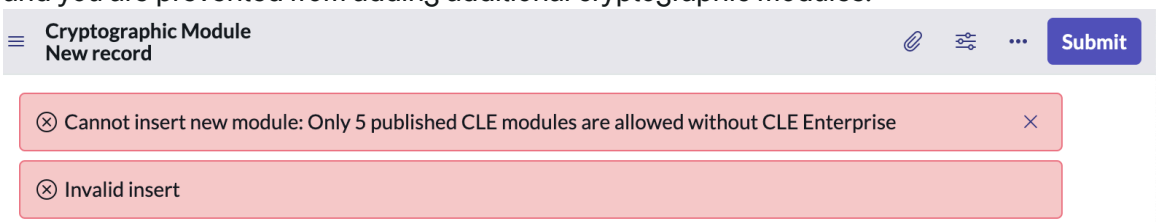

What to do next

Create a cryptographic specification for Column Level Encryption.

Encrypt data using the Multiple Modules feature

Encrypt data with more than one encryption module permitting the user to determine which keys are used for specific rows within the encrypted data.

Before you begin

Role required: sn_kmf.cryptographic_manager or sn_kmf.admin

About this task

The Multiple Modules option is considered non-deterministic and isn't the preferred method because the user determines which key to use for a given record. The ability to use multiple modules for a column is being replaced by Row Conditions. See Using multiple encryption modules. This non-deterministic implementation is still supported because it was created first and is still in use, but it's preferred to use Row Conditions for any new multiple modules use cases.

Note: Only encryption on columns supports multiple modules. Attachment encryption doesn't. Mass encryption isn't available when using the multiple encryption modules method.

You can't change a field using multiple encryption modules to use a single encryption module.

The field is encrypted by the encryption module of the first user to enter data. Because the encryption module is set on a per record basis, fields in a list can have different encryption modules. Within a single record, the field can be encrypted by only one module.

Procedure

1. Create multiple Field Encryption modules and a Module Access Policy (MAP) for each one.

   Make sure that you grant different roles to the different cryptographic modules through the access policies.

2. Navigate to **System Security** > **Field Encryption** > **Encrypted Field Configurations** > **New**.

   If you need more information on Encrypted Field Configurations, see Set encrypted field configurations.

3. In the **Type** field, you must select **Column**.
   Attachment encryption doesn't support multiple modules.
4. Select **Multiple Modules** in the **Method** field.

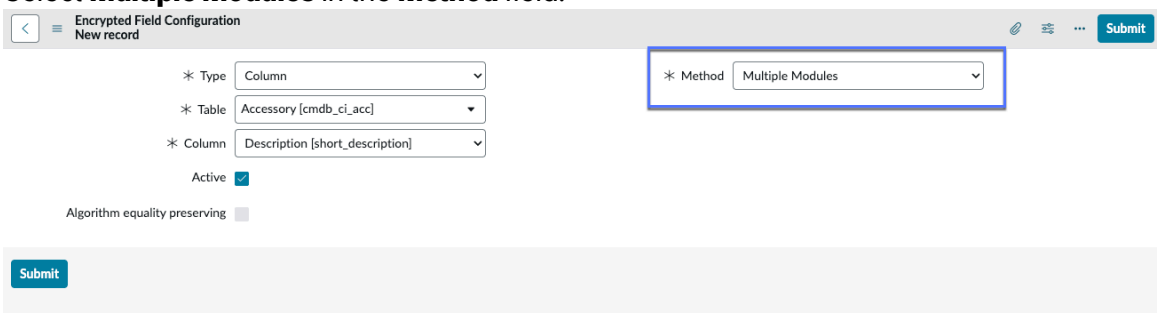

5. Select the **Table** and the **Column** in the table that you want to encrypt.
6. Select **Submit**.

Result

Newly created data for the specified field is encrypted with the key of the relevant module. When a user with the role specified in module A's access policy writes to the specified table, the data is encrypted with module A's key. Only users with the same role can read the data.

Example:

To encrypt the Short Description column on the Incident table. You would do the following:

1. Create two cryptographic modules A and B.
2. For each module, create a module access policy.

   For module A, give users with an HR role access. For module B, give users with a Sales role access.

3. Create an Encrypted Field Configuration record specifying the Short Description column on the Incident table, and make sure that you select **Multiple Modules** in the **Method** field.
4. Have two users, one with the HR role (user A) and one with the Sales role (user B), create an incident with a short description, and then have both users look at the list of incidents.

The short description for the incident created by the user with the HR role is encrypted by the key for module A. Likewise, the short description for the incident created by the user with the Sales role is encrypted by the key for module B.

Although all users with the HR and Sales roles have access to incidents, only a user with the HR role can decrypt and view the short description for those incidents created by user A, who had the HR role. Likewise, only users with the Sales role can decrypt and view the short descriptions for those incidents created by the user B, who had the Sales role.

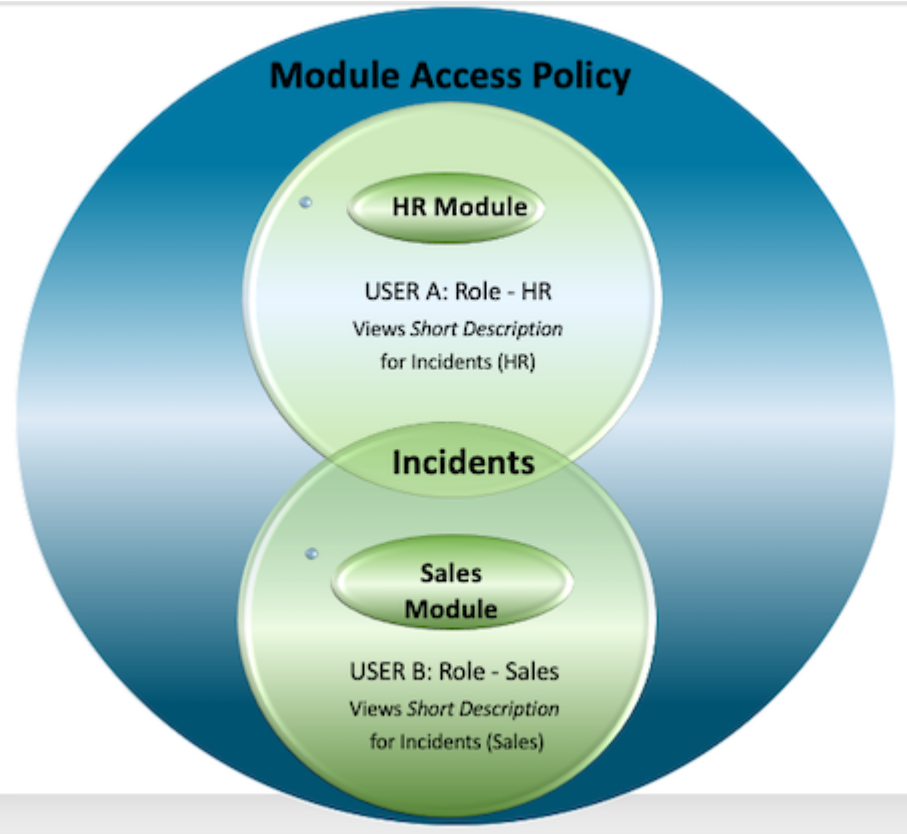

Create a cryptographic specification for Column Level Encryption

After you create a cryptographic module, access the corresponding cryptographic specification to define the algorithm.

Before you begin

Role required: sn_kmf.cryptographic_manager or sn_kmf_admin and security_admin or admin

About this task

This procedure describes options that are available with Column Level Encryption with the base system and additional configuration options that become available with Column Level Encryption Enterprise functionality. Column Level Encryption Enterprise functionality is available with a paid subscription. Refer to Encryption and Key Management subscription bundle for supported features and options available with each offering. See Activate Column Level Encryption Enterprise for more information on obtaining Column Level Encryption Enterprise.

A cryptographic specification will be created by the system when you create a cryptographic module for Column Level Encryption Enterprise.

Procedure

1. Navigate to **System Security** > **Field Encryption Modules** > **All**.
2. Select the cryptographic module to open the configuration options.
   Cryptographic module information is displayed at the top of the screen. A Symmetric Data Encryption/Decryption crypto specification is auto-created with an AES 256 CBC algorithm.
3. Select the crypto specification from the table to open the Algorithm Definition.
   For Column Level Encryption Enterprise see Configure advanced algorithms for Column Level Encryption Enterprise.
4. Click **Next** to access the Key Lifecycle.

What to do next

Perform one of the following operations:

- Select an entry in the Key Lifecycle table to define key lifecycle behavior. See Configure key lifecycle states for details to complete the lifecycle definition for the key.
- Click **Next** to create a cryptographic key. For details on this process, see Generate a ServiceNow cryptographic key.

Configure advanced algorithms for Column Level Encryption Enterprise

Create a cryptographic specification to define the algorithm for a cryptographic module. Customize the encryption specifications with advanced options that are available for Column Level Encryption Enterprise.

Before you begin

Role required: admin

Procedure

1. On the **Crypto Specifications (#)** tab, click **New**.

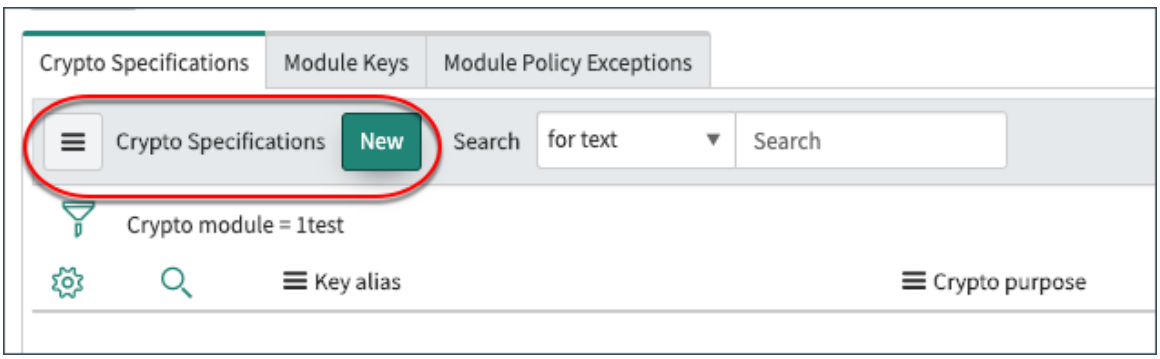

2. On the form, fill in the fields.

Algorithm Definition form

Field	Description
Crypto module	Name of the selected cryptographic module populates.

Field	Description
Crypto purpose	The value is **Symmetric Data Encryption/ Decryption** for Column Level Encryption Enterprise.
Algorithm	The value is **AES** for Column Level Encryption Enterprise.
Operation mode	The value is **CBC** for Column Level Encryption Enterprise.
Size	Possible values are **256** and **128**. **Note:** 256-bit size is most secure for encryption and is used for Symmetric Data Encryption/Decryption for Column Level Encryption Enterprise.
Equality preserving	Option to enable deterministic encryption. **Note:** Selecting this option means that the encrypted value of a field should be the same when the field value remains the same. Option to enable **Symmetric Data Encryption/Decryption** with **AES** in **Cipher Block Chaining (CBC)** mode.
Integrity	Option to provide Integrity in GCM operation and does not apply for Column Level Encryption Enterprise functionality.

3. Click **Submit**.

The following example shows AES CBC-256 encryption. When Column Level Encryption Enterprise is active and the parent module is column_level_encryption, only Symmetric Data

Encryption/Decryption AES CBC-256 applies as the crypto purpose. See Cryptographic specification overview for details.

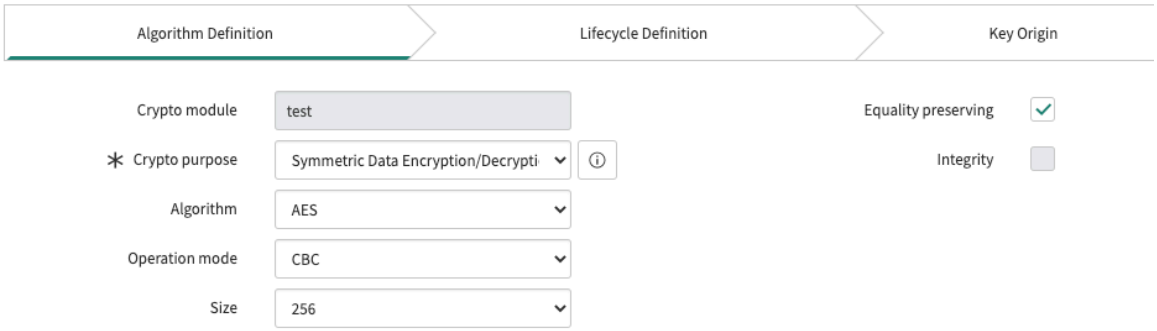

What to do next

Perform one of the following operations:

- Select an entry in the Key Lifecycle table to define key lifecycle behavior. See Configure key lifecycle states for details to complete the lifecycle definition for the key.
- Select **Next** to create a cryptographic key. See one of the following tasks for key generation:
  - Generate a ServiceNow cryptographic key.
  - Configure properties for customer-supplied keys.
  - Import the wrapping / unwrapping key pair.

Using customer supplied keys with Column Level Encryption Enterprise

You can use your own customer-supplied key instead of using the ServiceNow® system-generated keys.

Important: These topics only apply instances using Column Level Encryption Enterprise, which is only available with the `com.glide.now.platform.encryption` plugin. See Activate Column Level Encryption Enterprise for more information on obtaining this plugin.

With Column Level Encryption Enterprise you can use your own keys for encryption. Administrators have the choice to use ServiceNow® supplied keys or your own customer-supplied keys (CSK) for encryption on the ServiceNow AI Platform.

Important: To make use of the customer supplied key option, you must have your own cryptographic key.

Once you have your key, you can being using it on your instances by following these steps.

1. **Configure properties for customer-supplied keys**

   There are three system properties which define the size, padding algorithm, and validity period of the wrapping RSA key pair. Review these properties and adjust their values if the defaults to not fit your needs.

2. **Wrap your customer-supplied key**

   Use a cryptographic tool to wrap your key like OpenSSL to wrap the symmetric key to use for encryption with the downloaded public key.

Configure and upload your customer supplied key

Upload your wrapped your customer supplied key and configure cryptographic module to begin using your key for encryption on your instance.

Configure properties for customer-supplied keys

If the Field Encryption Enterprise plugin is enabled, you can use system properties to define key padding, ephemeral key pair size, and a key validity period of your customer-supplied keys.

Field Encryption Enterprise with Key Management lets you manage the full key lifecycle of your data encryption keys. Optionally, you can securely exchange data encryption keys generated within your environment.

Platform Encryption with Key Management lets you manage the full key life cycle of your data encryption keys. Optionally, you can securely exchange data encryption keys generated within your environment.

System properties for defining key-pair attributes

When you provide your own key, you must wrap it with the RSA public key. Three properties define the size, padding algorithm, and validity period of the wrapping RSA key pair:

- *glide.kmf.ephemeral_key.key_padding* controls the key padding scheme for the ephemeral key. The default scheme is OAEP SHA256, but SHA1 is also supported.
- *glide.kmf.ephemeral_key.key_size* controls the key size of the ephemeral key pair. The default is 4096 bits, but 2048 bits are also supported.
- *glide.kmf.ephemeral_key.key_validity_period* defines the period for which the ephemeral key pair is valid. The default value is two hours.

After the data encryption key is imported to the instance, a secure wrapping key protects new module keys on the instance. The wrapping key is an instance key encryption key (IKEK) generated by a hardware security module (HSM) on SafeNet KeySecure. See Instance level keys in the Key Management Framework for details in key types.

Continue to Wrap your customer-supplied key.

Wrap your customer-supplied key

Wrap the symmetric key to use for encryption with the downloaded public key.

Before you begin

Note: This procedure describes options that are available with KMF base system and options to be used with Field Encryption Enterprise functionality. Field Encryption Enterprise functionality is available only when the *com.glide.now.platform.encryption* plugin is active. See Activate Field Encryption for more information on obtaining Field Encryption Enterprise.

Some of the steps in this document require the use of a cryptographic tool installed on your local device. The examples in this task use the OpenSSL tool. For more information on this tool see https://www.openssl.org. If you are using other cryptographic tools, such as LibreSSL or GnuTLS, refer to the documentation for those products for similar steps.

- Modify optional properties that control the size, padding algorithm, and validity period of the key. See Configure properties for customer-supplied keys.
- You must have your symmetric key (.BIN) for encryption.

Important: Your key must be in binary format. If another format is used, a `Token failed validation. Please reattach the unmodified token.` error message displays.

- You must have a cryptographic tool to wrap your key. This example uses OpenSSL 1.1.

Role required: sn_kmf.cryptographic_manager or sn_kmf.admin

Procedure

1. Navigate to **All** > **Key Management Framework** > **Cryptographic Modules** > **All**.
2. Select the cryptographic module that you created for the customer supplied key from the Crypto Specifications related list.
3. You will be directed to the **Key Creation** step.
4. If you have not previously downloaded the wrapping key, click the link to download the *token_publickey<id>.zip* file and save it to the same location as your key.

   Note: Do not rename the downloaded *token_publickey<id>* file.

5. Unzip the file to your local network.
   The zip file contains two files, an import token and a public key `.PEM` certificate. Wrap your symmetric key with the public key to encrypt it.
6. Copy the name of the *token_publickey* file to your clipboard.
7. From a command line, use the copied *token_publickey* file name to open the folder of the unzipped files as a placeholder for the wrapped key.
8. Edit this script by replacing the examples with the names of your crypto files.

```
"downloads user.name$ cd token_publickey_<token>
openssl pkeyutl -encrypt -pubin -inkey publickey_<keyname>.PEM
-in <keyname.bin>
-out wrapped_key_material -pkeyopt rsa_padding_mode:oaep
 -pkeyopt rsa_oaep_md:sha<128 or 256> "
```

Review the key wrapping commands in the following table for more information.

Key wrapping commands

Directions	Command	Example
Open the file directory where you downloaded the wrapping token.	`cd`	`cd token_publickey123456789`
Paste the name of the `publickey.PEM` certificate.	`openssl pkeyutl -encrypt -pubin -inkey`	`publickey_586798643ffff.PEM`
Paste the name of your key here.	`-in`	`mykey.bin`

Directions	Command	Example
Enter the <-out> command and specify if the key is 128 bit or 256 bit.	`-out wrapped_key_material -pkeyopt rsa_padding_mode:oaep -pkeyopt rsa_oaep_md:sha256`	N/A

9. Run the command.
   A system message displays `token_publickey_<keynumber>`. The key will be generated and a *`wrapped_key_material`* file added to the directory.
10. Upload the wrapped key.

What to do next
Return to Configure and upload your customer supplied key to upload your wrapped key.

Configure and upload your customer supplied key

You can use your own customer-supplied key instead of using the ServiceNow® system-generated keys.

Before you begin
Roles required: security_admin, sn_kmf.cryptographic_manager

If you're NOT supplying your own keys, you don't need to perform this procedure. To create a cryptographic module with ServiceNow® keys, go to Create a cryptographic module or Create cryptographic module for Field Encryption.

Note: This procedure only applies to Field Encryption Enterprise functionality. See Activate Field Encryption for more information.

Important: You can't revoke a customer supplied key.

Procedure
1. Navigate to **All** > **System Security** > **Field Encryption Settings** and verify that **Customer Supplied Keys** is selected.

Key source selection

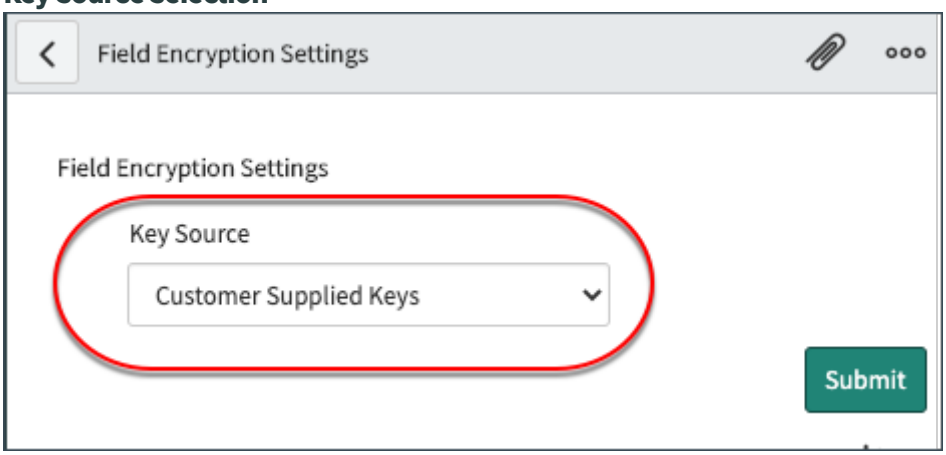

2. Select **Submit**.

3. Return to **System Security** > **Field Encryption Modules** > > **Create New**.

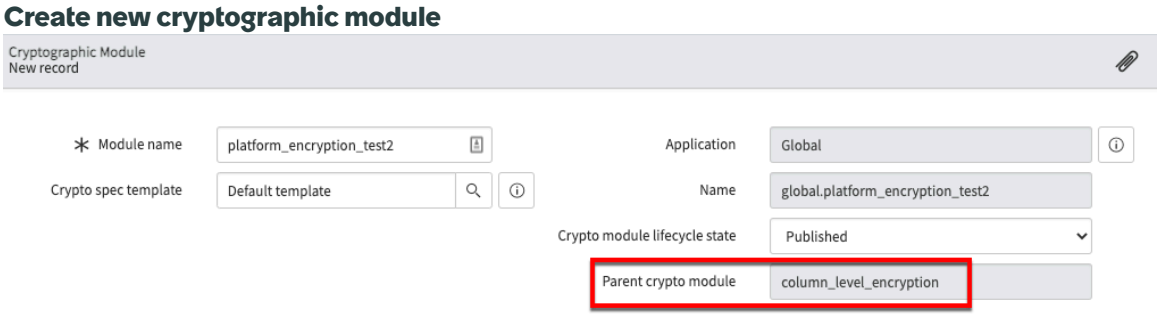

4. Complete the Cryptographic Module form as follows:

Cryptographic Module fields

Field	Description
Module Name	Enter a name for the module.
Crypto spec template	The default cryptographic template is selected.
Name	Auto-populates based on the module name and prepends the name with the scope to ensure which application is being applied. In this case, the global scope is applied.
Crypto module lifecycle state	Select **Published** to activate the crypto module.
Parent crypto module	The parent module **column_level_encryption** is selected automatically when using customer-supplied keys and encryption modules.

5. Select **Submit**.
6. Select the newly created cryptographic module from the table.
   In the **Crypto Specifications** related list, select the auto-generated key alias with the AES 256 CFB algorithm.

   The system populates the Crypto purpose and the Algorithm for Field Encryption automatically and jumps to the **Key Origin** stage.
7. Notice that **Upload customer supplied key** is the **Origin** and the **Key alias** is already populated.

Key origin

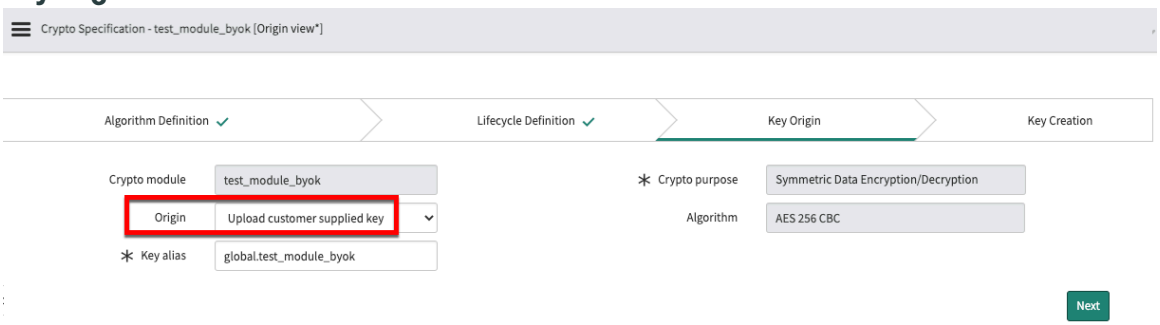

8. Select **Next** to move to the **Key Creation** stage.

   There are two links:

   - **Download wrapping key** downloads the key in a zip file containing an import token and a public key certificate, `.PEM` file. Use the import token to verify successful key wrapping according to security specification for the instance. Use the public key certificate `.PEM` file to wrap your customer supplied key securely before uploading it along with the token.
   - **Upload customer supplied key** opens the file browser to select the token and the encrypted key that you wrapped.

Key creation upload links

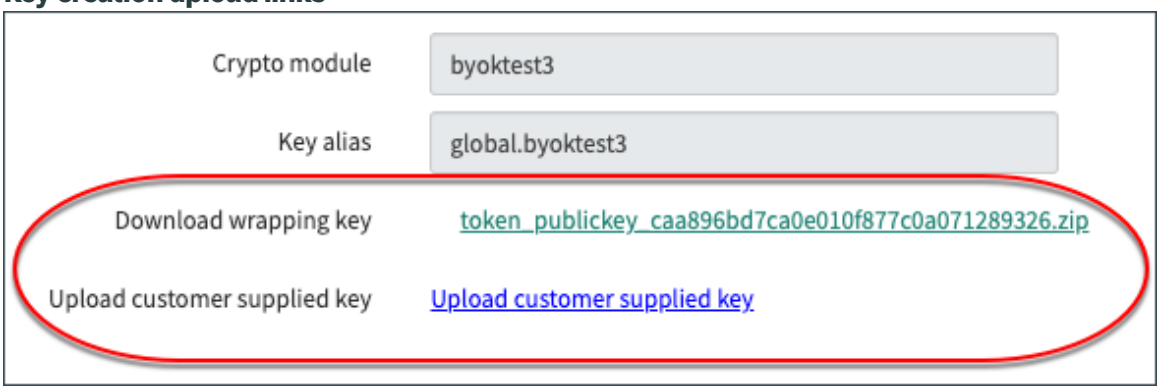

9. Select **Download wrapping key** to save the token.
   Save the token to the same destination location as the key is saved on your system. Don't rename the downloaded token.
10. Run the BYOK command on a terminal to wrap the key.
    For more information, refer to Wrap your customer-supplied key.
11. Select **Upload customer supplied key**.
12. Select **Browse** to select the two files, the wrapped key and the token file.
    The Attachments window displays the two files.

Wrapped key attachments upload

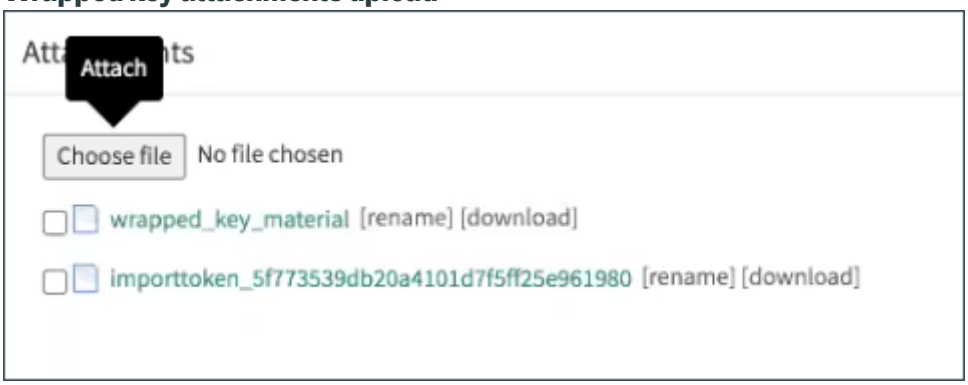

Select a file to remove and reupload, if necessary.

13. Select **OK**.
    You're returned to the Cryptographic Module screen. A confirmation message displays for a successful upload of the customer key. The key is also listed in the Module Keys related list.

Confirmation of key upload

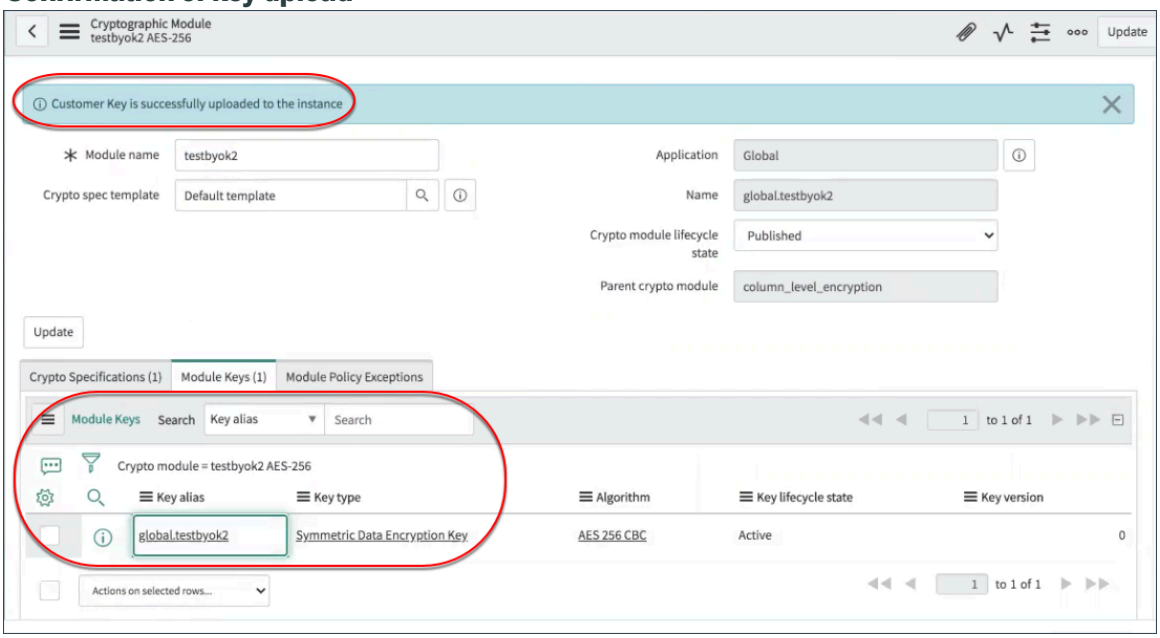

What to do next

Now that you have finished configuring your cryptographic module with your customer-supplied key, move on to Create a module access policy

Encrypting fields and attachments

Once cryptographic modules are created, a security admin can define the encrypted fields configuration (EFC) and opt to encrypt a field or attachment on a table.

How to encrypt fields

Note: Encrypted fields aren't audited by design. This behavior isn't configurable.

1. Specify the key source: ServiceNow generated keys or your customer-supplied keys (bring your own key) in **System Security** > **Field Encryption Settings**.
2. After specifying the key source, create a cryptographic module or use an existing cryptographic module. Start with Create a cryptographic module for instructions.

   Note: If you use customer-supplied keys, follow the directions in Create cryptographic module for Field Encryption and Configure properties for customer-supplied keys.

3. Create an encrypted field configuration to define where the encryption is applied. Here, you specify the target table and choose whether to encrypt a column or attachments within the table. See Set encrypted field configurations to get started.

Note: See Field Encryption Enterprise examples that illustrates how to encrypt fields and attachments using customer-supplied keys.

Set encrypted field configurations

Configure which table columns or attachments that the system encrypts using a preconfigured cryptographic module.

Before you begin

Role required: sn_kmf.cryptographic_manager and security_admin or elevate role to security admin.

About this task

Make sure you are in the correct application scope so you can see the tables in that scope.

Only users with access to the cryptographic module used in this configuration can read the data in the encrypted table column or access the attachment.

- If a user has write access but not read access, the field displays in edit mode and the data entered displays as asterisks.
- If a user has read access but not write access, the field displays the decrypted data in read-only mode.
- If a user has all access, both read/write functionality is available on the encrypted field.

See Create a cryptographic module or Create cryptographic module for Field Encryption to begin.

Important:

After encrypting a column, any new data inserted into the column is encrypted automatically. However, data that existed in the column before the encryption was active is not automatically encrypted.

In order to encrypt data that existed before the column was encrypted, you must run a separate mass encryption job. Learn more about mass encryption in Run mass encryption or decryption.

Procedure

1. Navigate to **All** > **System Security** > **Field Encryption** > **Encrypted Field Configurations** > **New**.
2. Select **New**.
3. Complete the form.

Field	Description
Type	**Column** to encrypt a table column or **Attachment** to encrypt all of a table's attachments. Types of data encrypted are: ○ String text (Full UTF-8) ○ Attachments ○ Date, Date/Time: **Note:** You can create encrypted field configurations to encrypt existing Date and Date/Time fields. You can add a new encryption configuration to a parent table only. You can't add a new encryption configuration to a child table. ○ URL ○ HTML ○ Journal ○ Translated
Table	Table whose fields or attachments are to be encrypted.
Column	Column (field) to be encrypted if you selected **column** as the type.
Active	Select to mark the configuration active. Deselect if the configuration isn't yet in use.
Crypto module	The cryptographic module that the encrypted field configuration applies to.
Method	Select **Single Module** to set the field configuration across one module. Select **Multiple Modules** for role-based access that spans across more than one cryptographic module. **Single Module** Use this option to encrypt all attachments using a single module. Your users need access to this module, otherwise they aren't able to upload attachments. **Multiple Modules** Use this option to allow users to choose a module when uploading attachments. Users with access to at least one module can select a module to use for encryption. Users with no module access can upload unencrypted attachments.
Algorithm Encrypted Preserving [read-only]	Indicates if the crypto module that you selected is already configured to support non-deterministic encryption. This means that if the same data is encrypted more than once, the encryption is different each time.

4. Select **Submit**.

Script access for cryptographic modules

Scripts can be run to access a cryptographic module policy for a cryptographic purpose.

For Key Management Framework, policies can be based scripts. When an access policy is triggered for script access, the backend script can execute the module policy actions from the script.

Cryptographic modules can support one or more encryption purposes, such as Asymmetric Data Decryption and Symmetric Data Decryption. Each cryptographic purpose requires the generation of an encryption key and defined cryptographic purpose.

Consider the following when executing an encryption script request:

- The referenced cryptographic purpose must be defined in the cryptographic module.
- An active generated key must exist for the cryptographic module.
- The Module Access Policy type must be set to **script**.

Check script version

When creating a module access policy that is set to the script type, there is an option available to validate the integrity of the script version being accessed. Only the assigned version of the script is allowed access to the encryption modules. When the **Check script version** check box is selected in the module access policy, anytime the script is run, the system performs a version comparison. If the script has been changed, the user is notified.

Check script version check box

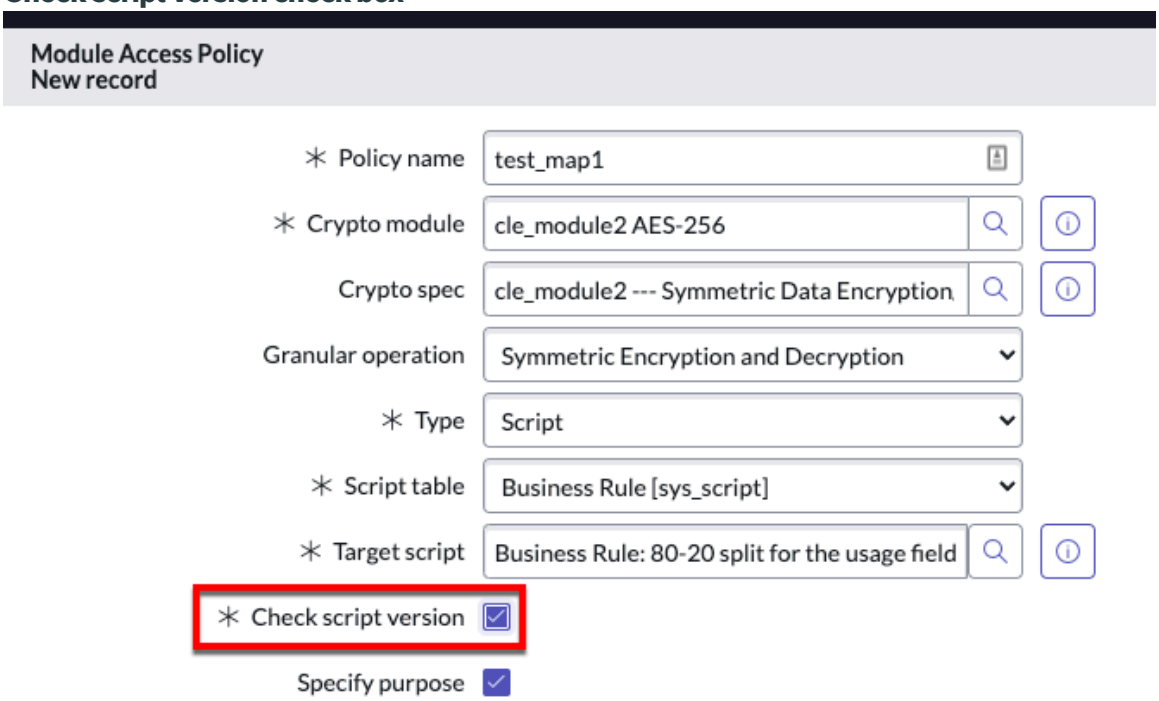

Configure script access to encrypted data

Execute a script to run the cryptographic module policy for a cryptographic purpose. Specific read (decrypt/unwrap) or write (encrypt, wrap) access can be defined based on the module access policy operation granularity.

Before you begin

Role required: sn_kmf.cryptographic_manager

About this task

Examples of uses are for Business Rules and Script Includes. This procedure uses a script for Business Rules.

Procedure

1. Create a cryptographic module with the symmetric data encryption/decryption algorithm. Refer to Create a cryptographic module for details. Specific access to the data or attachment is controlled with a module access policy with the following characteristics:
   - Symmetric encryption: The script is able to encrypt data but unable to decrypt the data.
   - Symmetric decryption: The script is able to decrypt uploaded encrypted data or attachment but unable to encrypt data or attachments.
   - Symmetric encryption and decryption: The script is able to both encrypt and decrypt data or attachments.
2. Navigate to **System Definition** > **Business Rules**.
3. Click **New**.

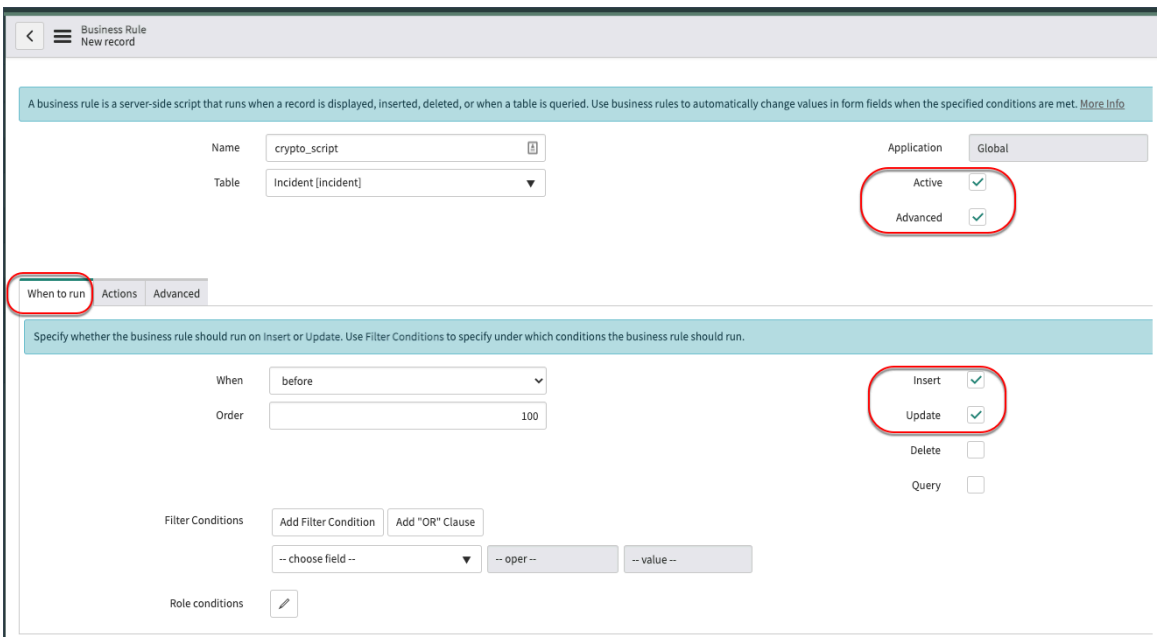

4. Complete the form on the **When to run** tab and enter the script on the **Advanced** tab:

Business Rule fields

Field	Description
Name	Enter a name for the business rule.
Table	Select **Incident [incident]**from the drop-down list.
Application	**Global** is selected by default.
Active	Mark the rule as **Active**.
Advanced	Select the check box to display advanced options.
When to run tab	On the **When to run** tab, enable **Insert** and **Update** fields.
Advanced tab	On the **Advanced** tab, paste the following script text at line 3:

Field	Description
	`var gc = global.GlideCryptoModule.getModule('global.acme_mod');` `var value = 'test';` `var encrypted = gc.encryptData(value);` `gs.info('value: ' + value);` `gs.info('Encrypted: ' + encrypted);` `var decrypted = gc.decryptData(encrypted);` `gs.info('Decrypted: ' + decrypted);` `gs.info(decrypted == value);` **Note:** Refer to the "Business Rules Advanced Tab" image for details.

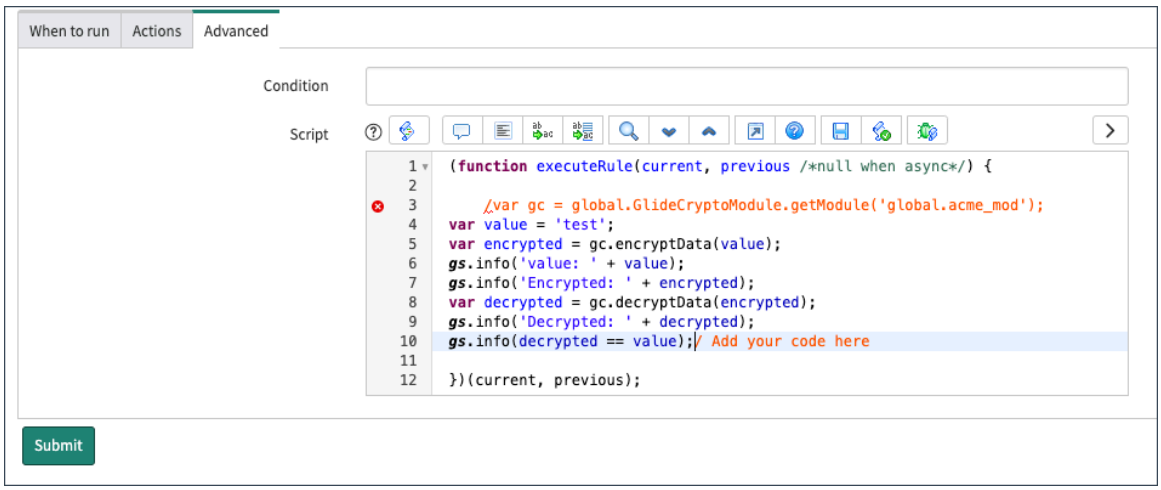

5. Click **Submit**.
6. Navigate to **Key Management** > **Module Access Policies** > **All**.

   Note: For additional information, refer to Create a module access policy.

7. Click **New**.
8. Complete the form.

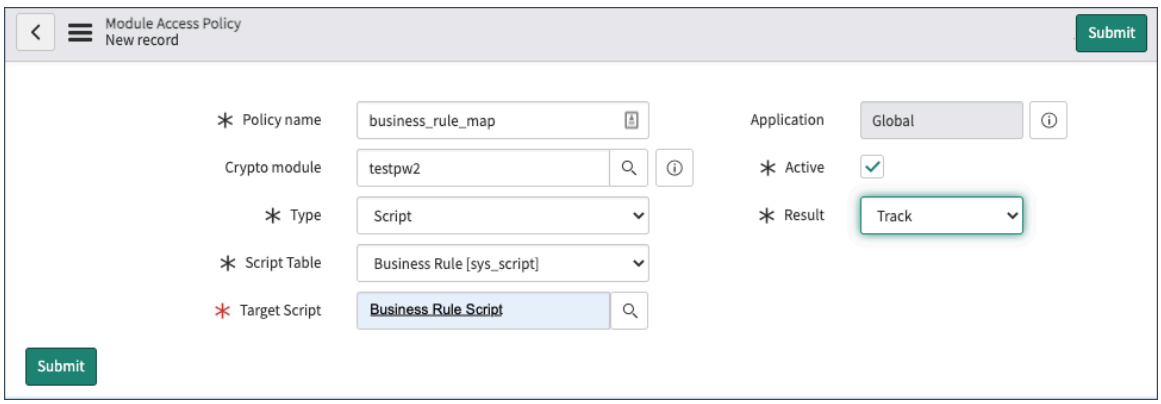

Module Access Policies fields

Field	Description
Policy name	Enter a name for the policy.
Crypto module	Click the search icon to select a module with the symmetric data encryption/decryption algorithm.
Type	Select **Script** to control access by script.
Script Table	Select a value from the script table drop-down list. For this example, select **Business Rule [sys_script]**.
Target Script	Select the script document for the policy. Select the **Table name** and then the related document for the policy. For this example, select the Business Rule that you created in previous steps.
Active	Select to activate the policy.
Result	To give the script access to the module, select **Track** in the **Result** field.

9. Click **Submit**.
   The Module Access Policy for the script is now available in the system.

View declined cryptographic module usage requests

View cryptographic modules that rejected encryption requests made by scripts because of unsupported encryption mechanisms.

Before you begin
Role required: sn_kmf.cryptographic_manager

About this task
Cryptographic modules can support one or more encryption purposes, such as Asymmetric Data Decryption and Symmetric Data Decryption. Encrypted data can only be accessed based on the module access policy. If a script tries to use a cryptographic module for a purpose not defined in the module, the script cannot access to the encrypted data.

In the following example, a cryptographic purpose was assigned to a cryptographic module, but a key was never generated for it.

Procedure

Navigate to **All** > **Key Management** > **Module Key Policies** > **Module Key Rejections**.

A list of cryptographic modules that rejected requests displays along with the encryption key used in the corresponding script.

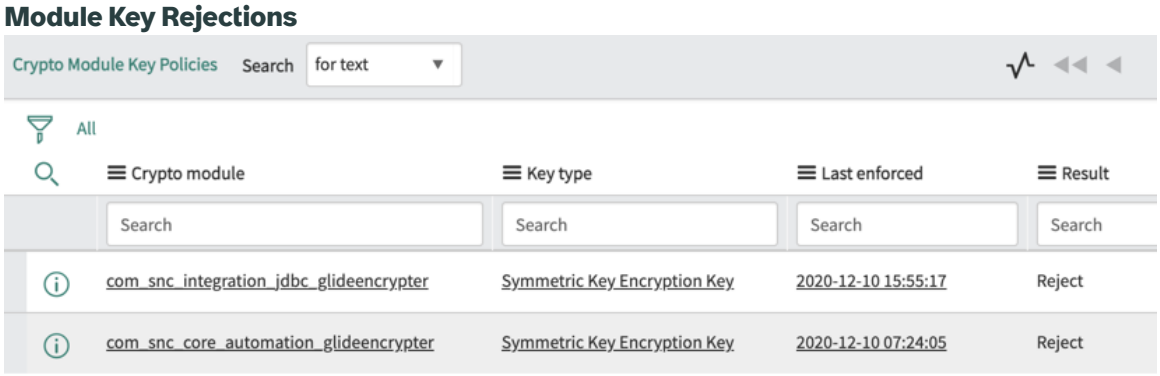

Note: If a different script attempts to use the same cryptographic module using the same key type, the value for **Last enforced** updates. Another row does not generate.

In this example, at 2020-02-10_15:55:17, the first module rejected a request because module1's key is compromised. At 2020-02-10_07:24:05, the second module rejected a request because the second module's key is suspended.

To grant scripts permission to use the encryption module the next time they run, create a module access policy for script encryption. For more information, refer to Configure script access to encrypted data.

Schedule mass encryption, decryption, and rekeying jobs

Schedule encryption, decryption, and rekeying jobs to run at a time that is best for your instance.

Before you begin

Encryption, decryption, and rekeying jobs can be time and resource intensive, so consider scheduling at non-peak hours. Also ensure that the user scheduling the job has the appropriate access for each job.

Role required: sn_kmf.cryptographic_manager

About this task

Mass encryption and decryption is also available from the Encrypted Field Configurations form. See Run mass encryption or decryption for instructions.

Procedure

1. Navigate to **All** > **System Security** > **Security Jobs**.
2. Click **New**.
3. Complete the scheduling form.

Field	Description
Name	Name of the encryption, decryption, or re-keying job.
Type	Job type: ○ **Key Migration Context to Module:** Mass migration of Encryption Context keys to Encryption Modules, including creation of Module Access Policies records for access controls on the Encryption Modules ○ **Data Migration Context to Module**: Migrates data encrypted by Encryption Contexts to Encryption Modules

Field	Description
	○ **Mass Decryption Attachment**:Decrypts all encrypted attachments in records for a single table you define in the **Table** field. ○ **Mass Encryption Attachment**: Encrypts all attachments in records for a single table you define in the **Table** field. ○ **Mass Encryption**: Encrypts any pre-existing value in the defined column/field used in the Field Encryption Configuration ○ **Mass Decryption Module**: Decrypts any pre-existing value in the defined column/field used in the Field Encryption Configuration with Single Module. ○ **Mass Decryption Multi Module**: Decrypts any pre-existing value in the defined column/field used in the Field Encryption Configuration with Multiple Module. ○ **Mass Rekeying**: Re-encrypts any pre-existing value in the defined column/field used in the Field Encryption Configuration using the current active key for the module. ○ **Migrate Attachment Context to Module**: Encrypts any pre-existing attachment on the table defined in the Field Encryption Configuration. Any attachment previously encrypted with a context is re-encrypted with the module.
State	The initial job state is New. After the job has been executed as scheduled, the state will update accordingly.
Time window start	Start time for the job in 24-hour format.
Time window end	End time for the job in 24-hour format.
Table	Table to be encrypted or decrypted.
Field	Field to be encrypted or decrypted.
Summary	Job status information when the job is running, has completed, or has errors.

Note: Because of system overhead, you should schedule mass encryption, decryption, and rekeying jobs to run at non-peak hours. The ServiceNow AI Platform runs the job between the **Time window start** and **Time window end**. If the job is not complete in one processing window, it continues during the next specified processing window until all processing is complete.

4. Click **Submit**.
5. After you schedule a job, you can do the following.
   - Click **Cancel Job** to cancel a running job.
   - Click **Start** to start a job immediately.
   - Click **Update** to save any changes you make to the job schedule.
   - Click **Delete** to delete the scheduled job.

Run mass encryption or decryption

You can run mass encryption on encryption configurations, as well as a mass decryption to decrypt previously encrypted values.

Before you begin
Role required: security_admin

About this task
You can also create scheduled jobs for mass encryption and decryption. See Schedule mass encryption, decryption, and rekeying jobs for instructions.

Mass encryption and decryption are available only when an encrypted field configuration uses the single cryptographic module. Mass decryption is available for both the single and multiple encryption method.

Note: You should run mass encryption and decryption only during non-peak hours because the operations are resource and time intensive.

Procedure
1. Navigate to **All** > **System Security** > **Field Encryption** > **Encrypted Field Configurations**.
2. Open the encrypted field configuration for the field you would like to mass encrypt or decrypt.
3. Under Related Links, select an available option.
   - **Schedule Mass Decryption job**
   - **Schedule Mass Encryption job**
4. Confirm your selection in the dialog.

Result
If running a mass encryption, all values are encrypted with the encryption module defined in the encrypted field configuration record. If running a mass decryption, only fields encrypted with an encryption module you have access to are decrypted.

Column Level Encryption Enterprise

These examples walk you through the encryption of fields and attachments using customer-supplied keys.

Column Level Encryption Enterprise walkthrough

This walkthrough shows you how to encrypt a field in your instance using Column Level Encryption Enterprise with the Key Management Framework (KMF). It also shows you how to use your own key.

Before you begin

Note: This procedure only applies to Column Level Encryption Enterprise functionality. See Activate Column Level Encryption Enterprise for more information on obtaining Column Level Encryption Enterprise.

Role required: admin or security_admin

Note: security_admin is a privileged role, for details on using privileged roles, see Elevate to a privileged role

About this task
This walkthrough starts with an instance where you have already created and uploaded your personal cryptographic key. You could use the ServiceNow key, but this example uses a customer-supplied key.

After the key has been stored in a cryptographic module, you can start configuring fields in your instance, such as salary or social security numbers that have limited access from certain users.

In the Encrypted Field Configuration, specify which authorized personnel can access sensitive data.

This task demonstrates two scenarios. One example encrypts the **Short Description** field in an Incident for users who are not authorized to view the sensitive data.

Attachments can also be encrypted and only visible to users who are granted access, or is visible to all users that are not restricted from viewing the data. See Attachment encryption walkthrough to encrypt an attachment.

Procedure

1. Make sure that Field Encryption Enterprise is enabled.
2. Create a cryptographic module for column_level_encryption.
   See Create cryptographic module for Column Level Encryption Create a cryptographic module for more information.
3. Navigate to **System Security** > **Encrypted Field Configurations**.
4. Click **New**.
5. On the form, fill in the fields.

Encrypted field configuration form

Field	Description
Type	**Column** is required to use your personal key.
Table	Table that stores the sensitive information. For this example, select **Incident [incident]**.
Column	Column, or specific information, that represents the sensitive date to be encrypted. For this example, select **short_description**.
Active	Option to mark **Active** to use the field configuration.
Algorithm Equality Preserving	The option is automatically selected.
Crypto module	Module that you created to use with the personal key.
Method	The **Single Module** option is used to apply the policies for one module. **Multiple Modules** is used to apply the policies across multiple modules.

Encryption field configuration example

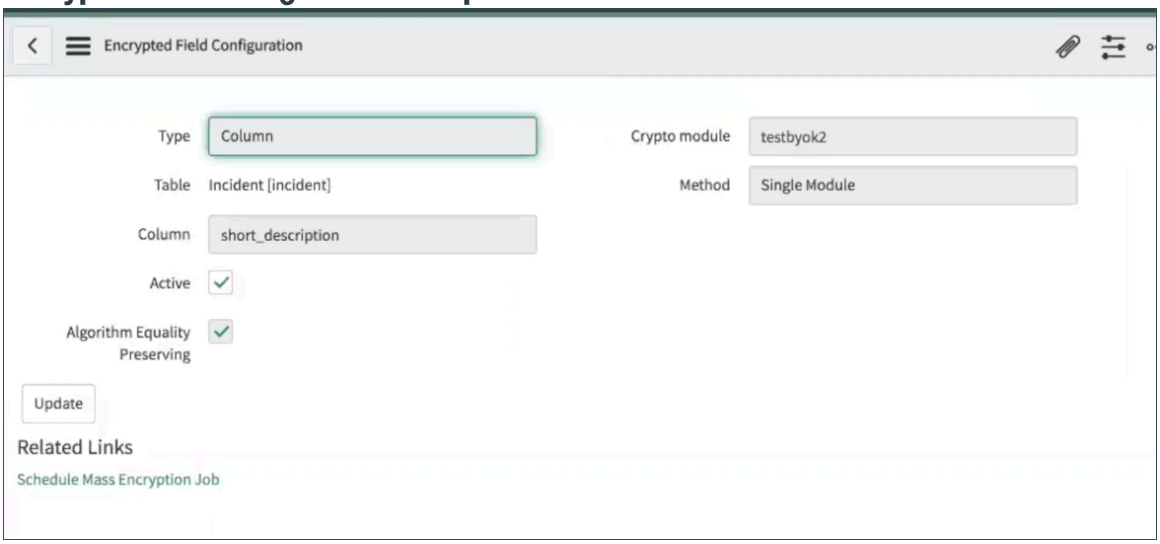

6. Click **Submit**.

   Establish a Module Access Policy to assign access to the cryptographic module. See Create a module access policy for additional information.

7. Navigate to **Key Management** > **Module Access Policies** > > **Create New** > **.**
8. On the form, fill in the fields.

Module access policy form

Field	Description
Policy name	Name for the policy, such, as `short description`.
Crypto module	Crypto module that you created to encrypt your key.
Type	Type of access designation for the crypto policy. Use **Role** to grant access to the encrypted field to only those users that have the assigned role.
Target Role	The role that has access to the encrypted field. For this example, select **Admin**.
Active	Option to activate the Module Access Policy.
Result	The **Track**option enables the access to the field for the selected role. (To restrict access to that field for the selected role, select **Reject** or**Strict Reject**.)

Module access policy example

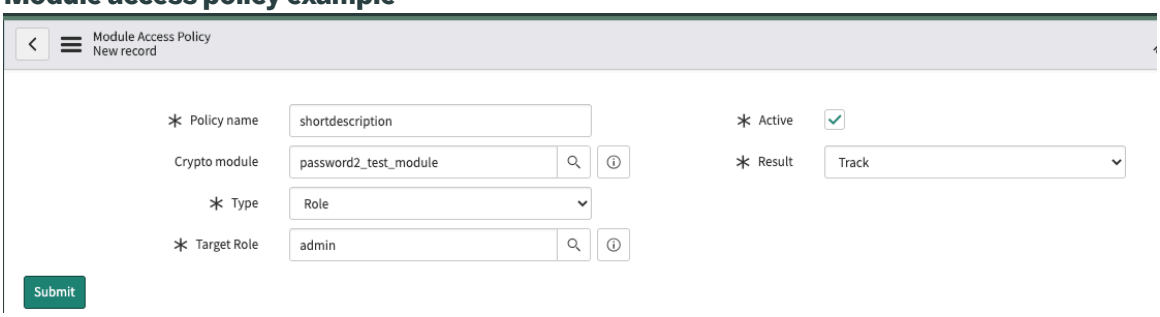

9. Click **Submit**.
10. As a user with the sn_kmf.admin role, navigate to **Incident** > **New**.

Example of encrypted field visible

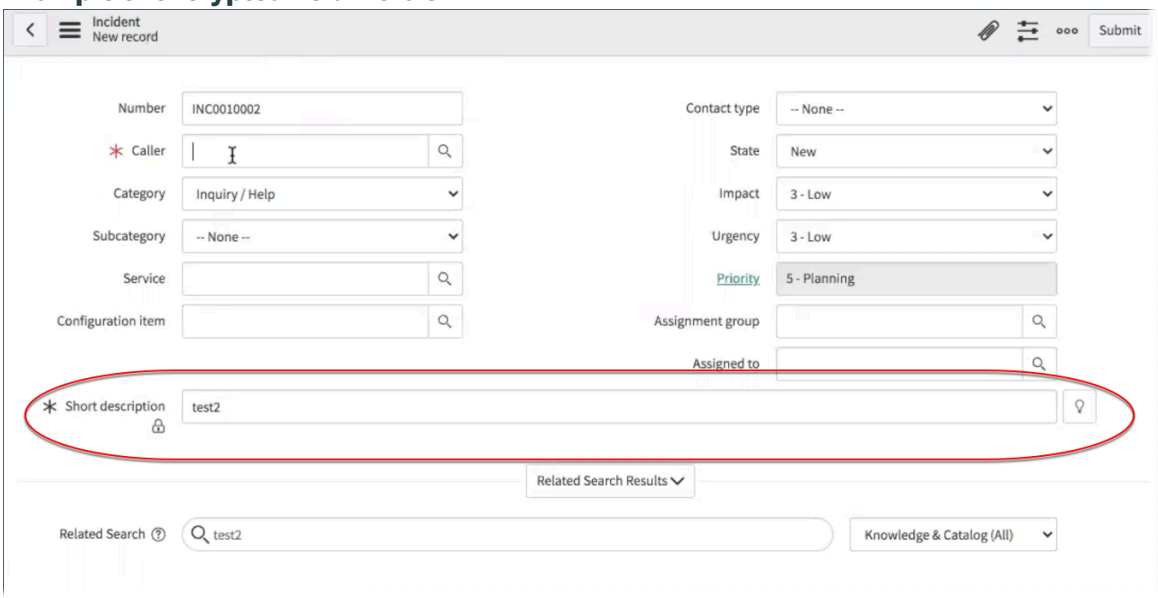

You can now view the Short description field based on the module access policy configuration.

Note: The sn_kmf.admin role was granted user access to the encrypted field, Short description, by setting the module access policy to **Track**. Notice the lock icon ( ) under the field name indicating that the field is an encrypted field.

You can now access the **Incidents** module as an end user to test the encrypted field configuration.

11. Log in as a user to be restricted from viewing the encrypted data in the configured field.

Encrypted field level data

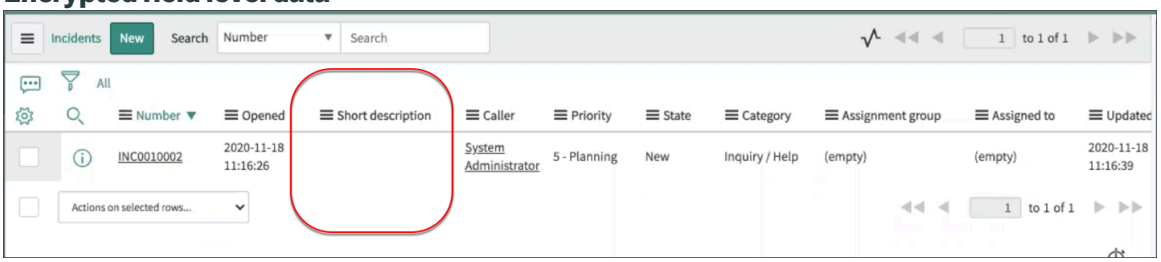

When you access the incident number, the data in the Short description will not be visible.

Result

You have successfully used your symmetric key to control access to a specific field using Column Level Encryption Enterprise.

Attachment encryption walkthrough

This walkthrough shows you how to encrypt an attachment in your instance using Field Encryption Enterprise with the Key Management Framework (KMF). It also shows you how to use your own key.

Before you begin

Note: This procedure only applies to Field Encryption Enterprise functionality. See Activate Field Encryption for more information on obtaining Field Encryption Enterprise.

Role required: kmf cryptographic manager

About this task

This walkthrough starts with an instance where you have already created and uploaded your customer-supplied cryptographic key. You could use the key, but this example uses a customer-supplied key.

Upload confidential attachments in your instance and limit access from certain users. Use Encrypted Field Configuration to specify which authorized personnel can access sensitive data.

We show you how to encrypt attachments to only be visible to users who are granted access, or be visible to all users that are not restricted from viewing the data. In this example, we restrict a certain role from being able to access an attachment in the **Incidents** module.

Note: Although you can use multiple modules with Field Encryption Enterprise, attachment encryption must use single modules.

Procedure

1. Make sure that Field Encryption Enterprise is enabled.
2. Create a cryptographic module.
   See Create cryptographic module for Field Encryption for more information.
3. Navigate to **System Security** > **Encrypted Field Configurations**.
4. Click **New**.
5. Complete the form:

Encrypted field configuration fields

Field	Description
Type	Select **Attachment** to use your personal key for encrypting an attachment from the selected **Table** For this example, select **Incident**.
Table	Select the table to access the sensitive information. For this example, select **Incident [incident]**.
Active	Mark **Active** to be able to use the field configuration.
Algorithm Equality Preserving	When selecting Field Encryption Enterprise, this field is visible based on the table selected.
Crypto module	Select the module that you created to use with the personal key.
Method	The **Single Module** option is used to apply the policies for one module. **Multiple Modules** is used to apply the policies across multiple modules.

Encrypted Field Configuration table

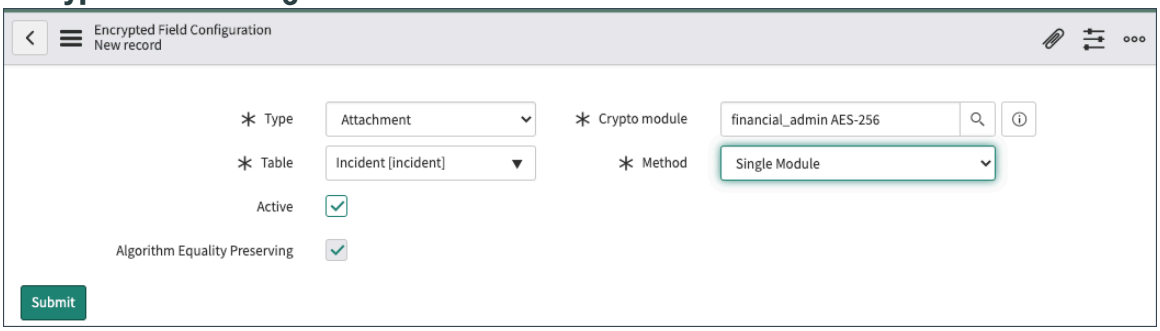

6. Click **Submit**.

   Establish a Module Access Policy to assign access to the cryptographic module. Refer to Create a module access policy for additional information.

7. Navigate to **Key Management** > **Module Access Policies** > **All**.
8. Click **New**.
9. Complete the form:

Module access policy fields

Field	Description
Policy name	Enter a name for the policy, such, as "Attachment policy."
Crypto module	Select the crypto module that you created to encrypt your key.
Type	Select **Role** to restrict access to the encrypted field from users with the assigned role.
Target Role	Select the role that will not have access to the encrypted field. For this example, select **itil**.
Active	Select this check box to be able to use the Module Access Policy.
Result	Select **Strict Reject** to control the access to the attachment from the selected role. (To grant access for the selected role, select **Track**.)

Module Access Policy form

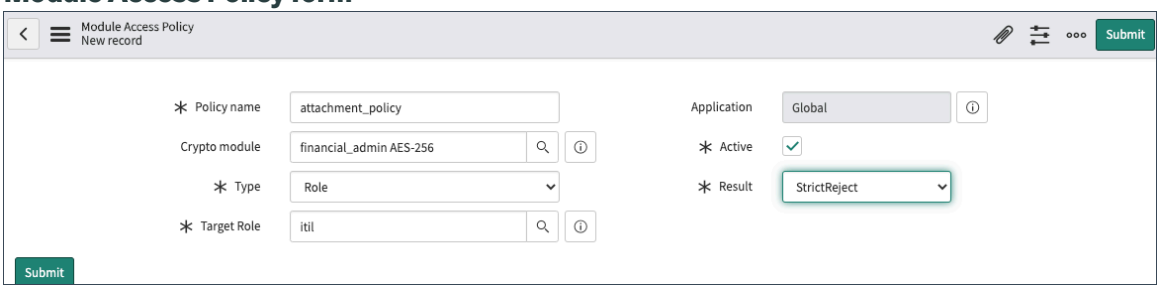

10. Click **Submit**.
11. As admin or as a person that created the incident, navigate to **Incidents** and add an attachment to **Activities** on the **Notes** related list.

Attachment available per role

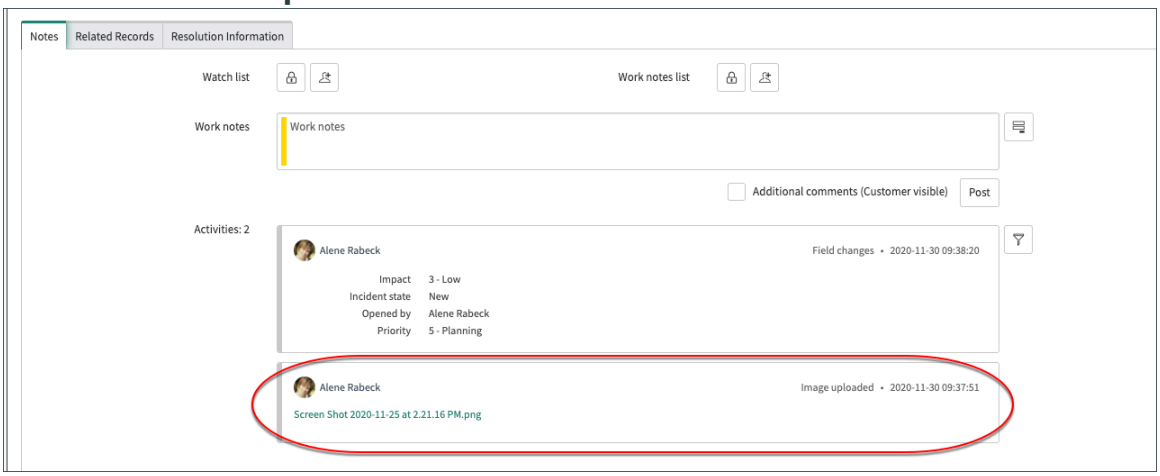

12. Log in as a user that restricted from accessing the encrypted attachment.
13. Open the incident and scroll to the **Activities:** section.
    The link to open the attachment is not accessible for users with the restricted role.
14. You have now successfully used your customer-supplied key to control access to a specific attachment using Field Encryption Enterprise.

Column Level Encryption Enterprise

Column Level Encryption Enterprise uses the Key Management Framework (KMF) to enable you to customize and manage how fields and attachments are encrypted and decrypted on your instance. A subscription is required to use Column Level Encryption Enterprise.

Important:

Starting with the Zurich release, Column Level Encryption (CLE) and Column Level Encryption Enterprise (CLEE) are being prepared for future deprecation. They will be hidden and no longer activated on new instances but will continue to be supported. Field Encryption and Field Encryption Enterprise provide the latest experience for this functionality.

For details, see the Deprecation Process [KB0867184] article in the Now Support knowledge base.

Column Level Encryption Enterprise is premised with Column Level Encryption and uses the Key Management Framework and its full support of key management functions. Column Level Encryption Enterprise provides key-protection and key life-cycle management for application-level field encryption. All keys are protected with a key-wrapping hierarchy ultimately rooted in FIPS (Federal Information Processing Standards) 140-2-L3 Hardware Security Modules (HSM).

Important: This topic covers the enterprise version of Column Level Encryption. For information on the standard version of Column Level Encryption, or to learn the differences between the two versions, see Exploring Column Level Encryption.

Column Level Encryption Enterprise gives you the ability to manage how supported fields are encrypted and decrypted in accordance with NIST 800-57 practices. It also uses the most updated version of field-level encryption, including integration for proper key protection and management.

Specifically, Column Level Encryption Enterprise uses the KMF encryption modules, granting you more control of server-side encryption. KMF verifies proper data encryption key protection using key hierarchy and envelope encryption. Your instance encrypts data through cryptographic modules that you configure. You can create an access policy for each module then configure cryptographic specifications and access policies and control key life-cycle management control.

Column Level Encryption Enterprise supports module access policies based on:

- Scope
- Role
- Script
- Resource Exchange
- System User

See Create a module access policy for additional information.

Note: For details on the supported features of Column Level Encryption and how to upgrade and subscribe to the Column Level Encryption Enterprise entitlement refer to Encryption and Key Management subscription bundle.

Encryption terms

Term	Description
Key management 	Support for key management Fundamental to Column Level Encryption Enterprise is the Key Management Framework (KMF). Gain the following capabilities: • Key life-cycle management. • Key rotation. See Rotate keys for details. • Key protection and key generation with FIPS 140-2-L3 Hardware Security Modules (HSMs). • Segregation of roles and duties. • The secure transfer of data encryption keys between instances, such as production and non-production instances. • Customer Supplied Keys (CSK) with key-wrapping. • Non-deterministic encryption. • Mass encryption/decryption. • Auditing of key access/use. See Key Management Framework Reference for details.

Term	Description
Customer-supplied key 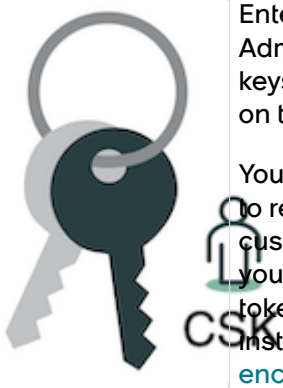	Support for customer supplied keys One of the biggest benefits of Column Level Encryption Enterprise is that you can use your own keys for encryption. Administrators have the choice to use ServiceNow supplied keys or your own customer-supplied keys (CSK) for encryption on the ServiceNow AI Platform®. You can also manage the key life cycle and decide when to revoke, rotate, and inactivate the keys. After you enable customer-supplied keys and create a cryptographic module, you download a token and public ephemeral key. You use the token and public key to wrap your key and then upload to the instance. To use customer-supplied keys, see Configure field encryption settings to select key type and Using customer supplied keys with Column Level Encryption Enterprise.
Column Level Encryption 	Support for both field encryption and attachment encryption Both field encryption and attachment encryption use cryptographic modules and access policies through Encrypted Field Configurations. The Encrypted Field Configuration form is used to choose an encryption type of column or attachment encryption. See Set encrypted field configurations for more information and supported field types.
Non-deterministic encryption 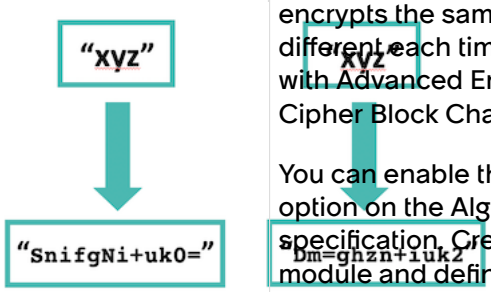	Support for non-deterministic encryption Column Level Encryption Enterprise supports non-deterministic encryption for enhanced security. If the system encrypts the same data more than once, the ciphertexts are different each time. Non-deterministic encryption is available with Advanced Encryption Standard (AES) encryption with Cipher Block Chaining (CBC). You can enable this feature through the Equality Preserving option on the Algorithm Definition stage of the cryptographic specification. Create a cryptographic specification for a crypto module and define an algorithm for encryption and generate the key. See Create a cryptographic module to define the mechanisms used for cryptographic operations and for more information on enabling non-deterministic encryption.

Term	Description
Resource Exchange 	Resource Exchange Column Level Encryption Enterprise keys instance to instance in a secure manner using the KMF cryptographic APIs to provide confidentiality, integrity, authentication, and non-repudiation. Resource Exchange is a KMF feature that gives you the capability to exchange resources between instances in a secure manner. See Key Management Framework Resource Exchange for details.

Note: If you choose not to activate Column Level Encryption Enterprise, you can still use Column Level Encryption. See Exploring Column Level Encryption for information.

Column Level Encryption Enterprise supports on-premise customers. It doesn't support Domain Separation.

Support for additional encrypted fields

The standard version of Column Level Encryption is limited to five encrypted columns. Column Level Encryption Enterprise supports an unlimited number of encrypted columns.

Supported field information

The following field types can be encrypted:

- Attachments
- Date
- Date/Time
- Email
- HTML
- Journal
- Journal Input
- Journal List
- Phone
- String text
- Translated Field
- Translated HTML
- Translated Text
- URL

Attachment Encryption

Attachment encryption by default

Customers using Column Level Encryption have attachments encrypted by default in tables that have an active Encrypted Field Configuration (EFC) type of *Attachment*.

This default encryption defined by the EFC configuration means that admins don't need to manually declare that an attachment should be encrypted on upload for these tables.

Administrators can disallow users from attaching unencrypted files

For details, see Prevent users from attaching unencrypted files.

Opt out of default encryption

If you don't want attachments encrypted by default based on EFC configuration, you can opt out of this option by contacting ServiceNow support.

To opt out of this feature, create a support case with ServiceNow support, and include this statement in a comment on the case record:

```
"I [customer name], understand that I am asking ServiceNow to turn off a recommended security best practice for attachments, and that [customer company] assumes any additional risk related to their configuration and use of unencrypted attachments in the ServiceNow application."
```

API support

Column Level Encryption Enterprise enables the following APIs.

Note: The API behavior described in the following table represents the default configuration for the latest base system package. If you are working with older package versions, you may experience different functionality.

Field Encryption APIs

API	Description	Parameters	Return type
`changeEncryptionContext()`	Updates an active Encryption Context (EC) used to encrypt an attachment. When CLE is enabled with the CLE Starter plugin using KMF Crypto Module (CM), the API locates the CM for the EC and uses it to encrypt the attachment. **Note:** This API is only available in the Global scope.	• sourceTable – Name of table that has the attachment. • sourceID – Table record system id. • attachmentID – The sys_attachment record system id. • newEncryptionContextID - System ID of the new context.	Boolean
`changeCryptoModule()`	Updates an active encryption module used to encrypt an attachment.	• sourceTable – Name of table that has the attachment. • sourceID – Table record system id.	Boolean

Field Encryption APIs (continued)

API	Description	Parameters	Return type
	Note: This API is only available in the Global scope.	• attachmentID – The sys_attachment record system id. • newCryptoModuleId - System ID of the new encryption module to encrypt the attachment.	
`disableEncryption()`	Disable active encryption on an attachment.	• sourceTable – Name of table that contains the attachment. • sourceID – Table record system id. • attachmentID – The sys_attachment record system id.	Boolean
`getDisplayValue()`	Returns the cleartext display value of an encrypted field.		String
`getValue()`	Returns the cleartext value of an ecrypted field when glide_encryption.set_value_support_cle.disabled is false (requires Module Access Policy (MAP)). Returns the encrypted value of an encrypted field when glide_encryption.set_value_support_cle.disabled is true.		String
`setDisplayValue()`	Inserts encrypted data into an encrypted field for display purposes.	• name – Field name. • value – Field value.	Boolean
`setValue()`	Inserts encrypted data into an encrypted field, controlled by a system property. Encrypts data when glide_encryption.set_value_support_cle.disabled is false (requires	• name – Field name. • value – Field value.	Boolean

Field Encryption APIs (continued)

API	Description	Parameters	Return type
	MAP); writes unencrypted data when set to true (no MAP required), when glide_encryption.set_value_support_cle.disabled is true.		

The following script illustrates API changes when the Incident short description is encrypted:

```
var gr = new GlideRecord('incident'); //creates a new incident
gr.setValue('short_description','test123'); //sets the value
 to test123
var sys_ID = gr.insert(); //inserts the record in the Incident
 table.
gs.info(gr.getValue('short_description')); //displays the
 unencrypted value
```

When the Column Level Encryption plugin is installed, `glide_encryption.set_value_support_cle.disabled` is set to false by default.

Cloud Encryption with Key Management

ServiceNow® Cloud Encryption offers encrypted storage for the database using block encryption, along with enhanced key management. Cloud Encryption is available with the ServiceNow® Platform Encryption subscription bundle.

Cloud Encryption offers:

- Segregation of duties.
- Rotation of ServiceNow Managed keys.
- Customer-Managed keys option.

> **Note:** With customer-managed keys, ServiceNow holds the encryption key on its infrastructure, but you perform key operations on it. Managing your key means you can bring your own key material (BYOK), rotate ServiceNow-managed or customer-managed keys, and withdraw your key. Keys aren't hosted on your own infrastructure. See Key management operations for details.

The following diagram shows how Cloud Encryption works.

The Cloud Encryption Key Management module consists of the following submodules:

- Key management operations:
  - Access the list of keys.
  - Perform key rotation operations.
  - Withdraw customer-managed key.
- Key management transactions:

  Reference all transactions that have occurred for the keys that have been used. Bring your own encryption key (BYOK) for use with Cloud Encryption.

  Use your own customer-managed key for encryption.

In certain circumstances, you may opt for a key withdrawal request when using a customer-managed key. To do so, you must license the Cloud Encryption Withdraw and Resupply optional add-on SKU and then request the key withdrawal functionality be activated by a Customer Service and Support team member.

The Quorum Control Policy Settings option becomes available when the withdrawal feature is activated, otherwise the module isn't visible on the menu. This feature can be activated only when using customer-managed keys. This policy enables settings to be configured regarding quorum when the withdrawal feature is activated. For more details on this feature, see Quorum Control Policy.

Cloud Encryption supports production and non-production instances for MariaDB and RaptorDB databases. Cloud Encryption is supported in the ServiceNow Commercial Cloud, Government Customer Cloud (GCC) pod 101, and ServiceNow Protected Platform – Australia (SPP-AU).

Licensing and enabling Cloud Encryption

For information about licensing Cloud Encryption, see Encryption and Key Management subscription bundle.

For licensed customers with new instances, the new instance provisioning will include Cloud Encryption.

For licensed customers with existing instances, to request an instance be moved to Cloud Encryption, follow the instructions in KB1117369. You must have the customer admin or partner admin role to request the Service Catalog item to Enable Cloud Encryption on your instance. Enabling this feature requires a one-hour maintenance window.

Cloud Encryption UI

When Cloud Encryption is enabled, the Cloud Encryption user interface (UI) is visible to the security_admin user when this user has the sn_kmf.admin role.

To access the Cloud Encryption UI by searching for **Cloud Encryption Key Management** in the navigation bar. Navigate to the **Key Management Operations** section to see information about encryption keys, such as details of the active key, and whether Cloud Encryption is enabled for the instance.

Key management operations

The Key Management Operations submodule provides access to view and manage all encryption keys used with ServiceNow Cloud Encryption.

Key life cycle states

There's only one active key in the system at any given time. When selecting a key, you access the activity for the selected key, such as which keys were rotated or withdrawn and the corresponding timestamp.

The key life-cycle state updates according to the key management operation performed.

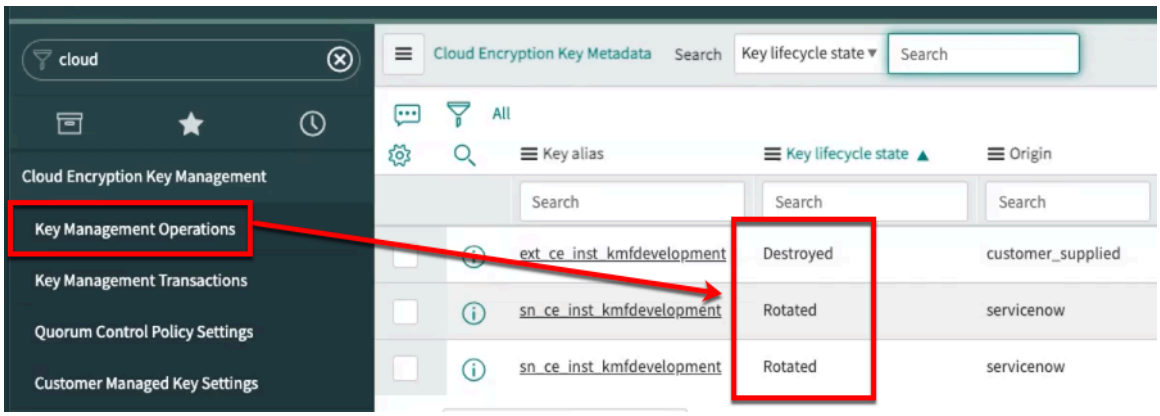

See Rotate a ServiceNow managed key or Rotate a customer managed key for details.

Note: The key rotation process may take up to 20 minutes to complete.

Rotate a ServiceNow managed key

Rotate the active Cloud Encryption ServiceNow managed key.

Before you begin
Roles required: sn_kmf.admin or sn_kmf.cryptographic_manager

Important: If you're using customer-managed keys, see Rotate a customer managed key.

Procedure

1. Navigate to **All** > **Cloud Encryption Key Management** > **Key Management Operations**.
   The Cloud Encryption Key Metadata list loads. All keys used in your Instance are listed. Only one key can be active at a given time.

   If accessing the Cloud Encryption module for the first time, key entries will be available after an initial key rotation is performed. A ServiceNow managed key is the default in the system.
2. Select the active key from the table.
   The Key Definition table displays with general information about your ServiceNow generated key.
3. Select the **Rotate Key** button.
   A notification displays with the option to continue the key rotation or to cancel the operation.
4. Select **OK** to rotate the key.
   A confirmation message displays at the top of the Key Definition page.
5. Return to the Key Management Operations screen to refresh the Cloud Encryption Key Metadata table.
   Entries are listed for the current active key and the key that is being generated to rotate in place of the current active key. See Key Management Framework key life-cycle states for the different available states.

   The active key is listed with a key version of `0` and the generated key has a version of `1`.
6. Open the entry for the original key to view the Key Management Transactions.
   For more information, see Key management transactions for more information.
   The previously active key, version `0`, updated to the key life cycle state of `Rotated` and the new key, version `1` is `Active`.

Prepare your customer managed key

Follow these steps to prepare your customer managed key for upload to your instance.

Before you begin
Roles required: sn_kmf.admin or sn_kmf.cryptographic_manager

About this task
For a customer managed key, you may use any cryptographic library or HSM to generate your key. This key must be an AES 256-bit key and be wrapped by a Cloud Encryption wrapping certificate with an RSAES_OAEP_SHA_256 encryption schema.

Note:

If you choose to use the OpenSSL cryptographic tool to generate your key, the OpenSSL version must be version 1.1.1x or later.

If you're creating and wrapping your customer-managed key using Windows, you must generate the wrapped key via Bash shell support applications such as Git Bash.

Procedure

1. Generate a random value to use as your AES-256 bit symmetric key using OpenSSL.

   For example, using openSSL you can generate this key with the `openssl rand 32` command.

   For compatibility, your symmetric key must have the following attributes:

Attribute	Value
Key type	Advanced Encryption Standard (AES) algorithm-based symmetric key.
Key size	256 bits (32 bytes)
Key wrapping requirement	○ RSA encryption algorithm ○ Optimal Asymmetric Encryption Padding (OAEP) ○ SHA-256 hash function (RSAES_OAEP_SHA_256) ○ Encoded using Base64 algorithm

2. Save the key to a file.
   For example, the openSSL command `openssl rand 32 > plaintext_key.bin` generates a 32-byte key and saves it to a file named `plaintext_key.bin`.

   Important: Save this file securely for future reference. This key is wrapped with the public key for upload.

3. Extract the public key from the downloaded wrapping certificate file from your instance:

```
openssl x509 -pubkey -noout -in wrapping_cert.pem >
 public_key.pem
```

   Note: Refer to for information to download the wrapping certificate.

4. Wrap the generated key with the public key downloaded with the wrapping certificate using the RSAES_OAEP_SHA_256 algorithm:

```
cat plaintext_key.bin | openssl pkeyutl -encrypt -inkey
 public_key.pem -pubin -pkeyopt rsa_padding_mode:oaep -pkeyopt
 rsa_oaep_md:sha256 | openssl base64 -A -out wrapped_key.txt
```

   A file specified on this command contains a wrapped customer-managed key that can be provided to SN for the CMK process.

Switch between ServiceNow and customer-managed keys

Switch between customer-managed key or a ServiceNow managed keys for use in ServiceNow Cloud Encryption.

By default, your instance is configured to use ServiceNow managed keys, and ServiceNow encryption key generation is active. However, administrators can choose to use customer-managed keys. They may also choose to return to ServiceNow managed keys.

Rotate a customer managed key

Rotate your customer managed key to your instance after you've wrapped your customer managed key for Cloud Encryption.

Before you begin

Roles required: sn_kmf.admin or sn_kmf.cryptographic_manager

Procedure

1. Navigate to **All** > **Cloud Encryption Key Management** > **Key Management Operations**.
   The **Cloud Encryption Key Metadata** list loads. All keys that have been used in your instance are listed.
2. In the **Cloud Encryption Key Metadata** list, open the record for your active key.
   If you have multiple keys, select the key that has a **Key Lifecycle state** of `Active`. There's only one active key on your instance.
3. In the key definition record, select the **Rotate Key** button.
4. In the **Upload Customer Managed key** window perform the listed steps.

   a. Select **Download Wrapping Certificate**.
   The *`public_certificate....zip`* downloads to your local machine and is used to wrap your customer managed key.

   > ⚠ **Warning:** Avoid potential certificate-related issues by downloading the wrapping certificate each time you rotate to or switch to a customer managed key.

   b. Select **Browse** to upload your customer manged key, then locate and select your wrapped encryption key.
   To choose a different file, select the file and select **Remove**.

   c. Close the Attachments window.

   d. Select **OK** to upload your key.
   If the key is in the proper format, a confirmation message appears, otherwise an error message displays. The key file is attached to the Key Definition record.

   In the Key Management Transactions table, the certificate download and key upload steps are listed. See Key management transactions for details on the request steps.

5. Navigate to **All** > **Cloud Encryption Key Management** > **Key Management Operations** to see the list of keys.
   In the list of keys, you can see a new record for your customer managed key. This new key has a **Origin** value of **customer_supplied** and is in the **Active** state. Your previous key has be in the **Rotated** state.

Switch to a customer managed key

Use your customer managed key for ServiceNow Cloud Encryption.

Before you begin

Role required: sn_kmf.admin or sn_kmf.cryptographic_manager

To switch to a customer managed key, you must have a wrapped customer managed key ready to upload as part of these steps. For details on preparing this key for upload, see Prepare your

customer managed key. After uploading your key, this process will initiate a key rotation to your new key.

Procedure

1. Navigate to **All** > **Cloud Encryption Key Management** > **Key Management Operations**.
2. In the **Cloud Encryption Key Metadata** list, open the record for your active key.
   If you have multiple keys, select the key that has a **Key Lifecycle state** of `Active`. There's only one active key on your instance.
3. In the **Related Links** section of the form, select the **Switch To Customer Managed Key** link.
4. In the **Switch to Customer Managed Key** dialog box, select the **Upload Managed Key** button.
5. In the **Upload Customer Managed Key** dialog box, perform the listed steps.

   a. Select **Download Wrapping Certificate**.

   > ⚠ **Warning:** Avoid potential certificate-related issues by downloading the wrapping certificate each time you rotate to or switch to a customer managed key.

   b. Select **Browse**, and follow the prompts to select and upload your key from your device.

   c. Select **Switch to Customer Managed Key**.

   A request is generated by the instance to switch to your customer managed key. In the current form, you can see that the **Key lifecycle state** of the originally active key has changed to **Rotated**.
6. Navigate to **All** > **Cloud Encryption Key Management** > **Key Management Operations** to see the list of keys.
   In the list of keys, you can see a new record for your customer managed key. This new key has a **Origin** value of **customer_supplied** and is in the **Active** state.

Result

Your instance now uses your customer managed key for ServiceNow Cloud Encryption.

> ℹ **Important:** Ensure that a copy of your encryption key is always available in a secure location for key management operations. Without this key, your instance may be rendered inaccessible.

Switch to a ServiceNow managed key

Switch from a customer managed key back to a managed key for ServiceNow Cloud Encryption.

Before you begin

Role required: sn_kmf.admin or sn_kmf.cryptographic_manager

Procedure

1. Navigate to **All** > **Cloud Encryption Key Management** > **Key Management Operations**.
2. In the **Cloud Encryption Key Metadata** list, open the record for your active key.
   If you have multiple keys, select the key that has a **Key Lifecycle state** of `Active`. There's only one active key on your instance.
3. In the **Related Links** section of the form, select the **Switch To ServiceNow Managed Key** link.
4. In the **Switch to ServiceNow Managed Key** dialog box, select the **Switch to ServiceNow Managed Key** button.

A request is generated by the instance to switch to a ServiceNow managed key. In the current form, you can see that the **Key lifecycle state** of the originally active key has changed to **Rotated**.

5. Navigate to **All** > **Cloud Encryption Key Management** > **Key Management Operations** to see the list of keys.
   In the list of keys, you can see a new record for a ServiceNow managed key is. This new key has a **Origin** value of **ServiceNow** and is in the **Active** state.

Schedule key rotation

Set a schedule for automatic rotation of your ServiceNow managed keys. This process automatically retires an encryption key and replaces the old key with a newly generated cryptographic key. If you're using a customer managed key, this schedule can provide a reminder to rotate your custom keys manually.

Before you begin

Role required: sn_kmf.admin

Procedure

1. Navigate to **All** > **Cloud Encryption Key Management** > **Scheduled Key Rotation Settings**.
2. Select the **Enable scheduled key rotation** check box.
3. Fill in the remaining fields based on your business needs.

Scheduled key rotation settings

Field	Description
Number of months between key rotations (Maximum of 60 months)	Number of months between key rotations. This value is 12 by default and can have a maximum of 60 months.
Day of the week to perform key rotation	Day of the week the key rotation is performed.
Time of day to perform key rotation	Time of day the key rotation is performed.
Date and time of next key rotation	Date and time of the next scheduled key rotation. This value isn't editable directly, and is automatically calculated based on your choices.
Number of days before key rotation to send reminder (Maximum of 15 days)	Number of days before the date of your key rotation that your instance sends notifications.
Email notifications are sent to the following list of your approved security administrators	List of users who receive notifications for key rotation. The **System Administrator** is on this list by default.

4. Select **Submit**

After selecting submit, you can see a notification at the top of the form. The notifications confirm your key rotation and notification schedule.

Warning:

Each scheduled key rotation has a unique signature, which ensures the integrity of a job and detect any unauthorized modification. The signature for a scheduled job is unique on each instance. Cloning a scheduled key rotation job from a source instance A to a target instance B, the scheduled job on the instance B will fail the signature validation. If this occurs, you can re-create the signature by de-selecting and then reselecting the **Enable scheduled key rotation** check box. For more details on this issue, see KB1247113.

Withdraw a customer managed key

After the customer managed key withdrawal functionality is activated, a withdrawal operation becomes available in the Key Management Operations page. Withdraw key and quorum approval operations can also be managed.

Before you begin

Roles required: sn_kmf.admin or sn_kmf.cryptographic_manager

This section applies only if you've licensed Cloud Encryption Withdraw and Resupply, an optional add-on to Cloud Encryption.

Procedure

1. Navigate to **All** > **Cloud Encryption Key Management** > > **Key Management Operations**.
2. Select the active customer managed key from the table.
   The Key Definition table displays with general information about your customer key. A withdraw key function is now available.
3. Select the **Withdraw Key** to trigger the withdrawal process.

Warning:

A Withdraw Key warning message displays. The withdrawal of the key triggers a shutdown of your instance until a restore operation is performed with the withdrawn key.

Danger: You can only perform a restore operation with the same key that was withdrawn. If you want to rotate to another key, you must do so after restoring the key that was withdrawn.

If the withdrawn customer managed key isn't restored within the time frame for which ServiceNow retains backups (see the Backup and Restoration SOP [Standard Operating Procedure] for details), your instance database backups will no longer be accessible. Backup data lost in this way isn't recoverable.

4. Select **OK** to withdraw the key.
   Select **Cancel** if there's any doubt about the key withdrawal function.
   You're returned to the Key Definition screen and a confirmation message displays.
5. Refresh the Key Definition page to view the pending withdrawal request.

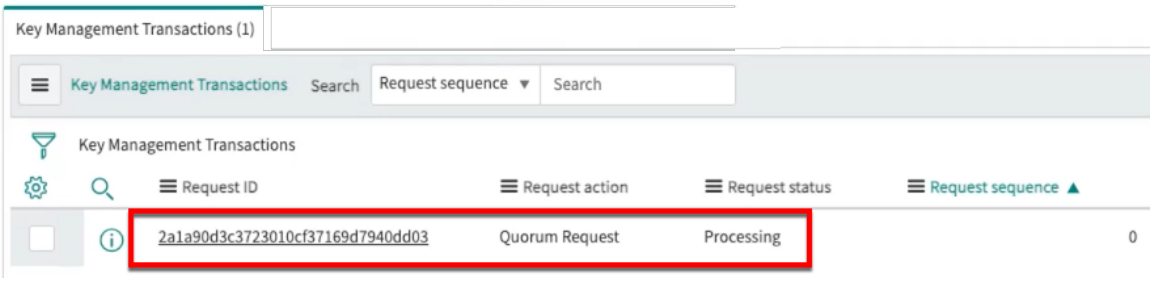

If the Quorum Control Policy has been activated, the approval workflow must be completed successfully to complete the key withdrawal. See Manage Quorum Control for details.

Resupply a customer managed key

After a key withdrawal operation has completed, your customer managed key must be resupplied into your instance.

Before you begin

Role required: sn_kmf.admin or sn_kmf.cryptographic_manager

Note: This section applies only if you've licensed Cloud Encryption Withdraw and Resupply.

Procedure

1. Go to Now Support and navigate to **Service Catalog** > **Catalog** > **Instance Management** > **Instance Restore - Resupply Managed Key**.
2. Click **Request**.
3. In the **Instance Restore - Resupply Managed Key** window, select your instance in the **Select Instance** drop-down.
4. Download the wrapping certificate by clicking on the **wrapping certificate** text.

   Warning: You must download a new wrapping certificate each time you rotate or upload a customer managed key.

5. Prepare your key for upload.
   For details on this process, see Prepare your customer managed key.
6. In the **Step 4** section, click **Browse and upload** to upload your wrapped key from your local device.
   After your key is uploaded, you can see the key under **Uploaded below file successfully**. If you need to re-upload the key, click **Remove file** and upload your key again as described in previous steps.
7. Click **Rotate key** to complete the re-supply.

Quorum Control Policy

The Quorum Control Policy specifies the minimum number of approvals required among the total number of selected approvers to reach quorum for customer managed key withdrawal.

Warning: You must sign a legal addendum to activate the key withdrawal functionality. The Quorum Control Policy Settings option becomes available when the withdrawal feature is enabled, otherwise the module isn't visible on the app menu. After key withdrawal, your instance is no longer available until the encryption key is resupplied and is active again.

In certain circumstances, you may want to create a key withdrawal. You must first request the key withdrawal functionality from Customer Service and Support.

The Quorum Control Policy specifies the minimum number of approvals required among the total number of selected approvers to reach quorum for customer managed key withdrawal. For instance, there are a total of five approvers, but only four approvals are required to reach quorum. When four approvals are obtained, the withdrawal request is processed and the key withdrawn.

Whenever a withdrawal operation is performed in a group enabled for quorum approval, a workflow for quorum approval is triggered. An email notification is sent to all users who can grant approval. Tasks are also generated in the approvers' accounts, which they see on the dashboard when they log in to their ServiceNow account.

The users can grant approvals from the Instance, the email, or the Key Management Operations page. The key withdrawal operation is blocked until the quorum is met.

Once the minimum number of approvers is reached, quorum is reached and the key withdrawal will trigger. The withdrawal is performed and is logged, including the names of users who approved the request.

See Configure Quorum Control Policy Settings for setup details.

Configure Quorum Control Policy Settings

Follow these steps to configure Quorum Control Policy Settings.

Before you begin

Roles required: sn_kmf.admin

About this task

> ⚠ **Warning:** You must sign a legal addendum to activate the key withdrawal functionality. The Quorum Control Policy Settings option becomes available when the withdrawal feature is enabled, otherwise the module is not visible on the app menu. After key withdrawal, your instance is no longer available until the encryption key is active again.

Procedure

1. Request the key withdrawal functionality from Customer Service and Support.
2. Navigate to **Cloud Encryption Key Management** > **Quorum Control Policy Settings**.
3. Select the **Quorum control enabled** check box.

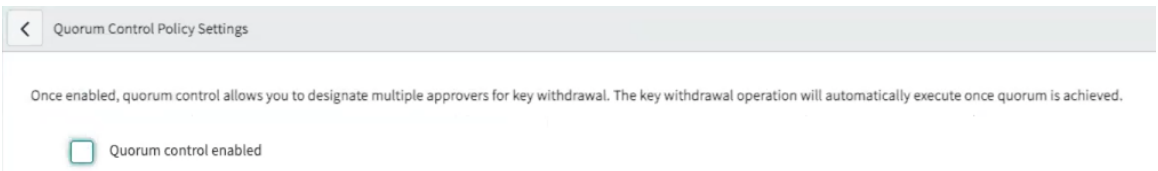

Additional fields appear that are required to configure quorum

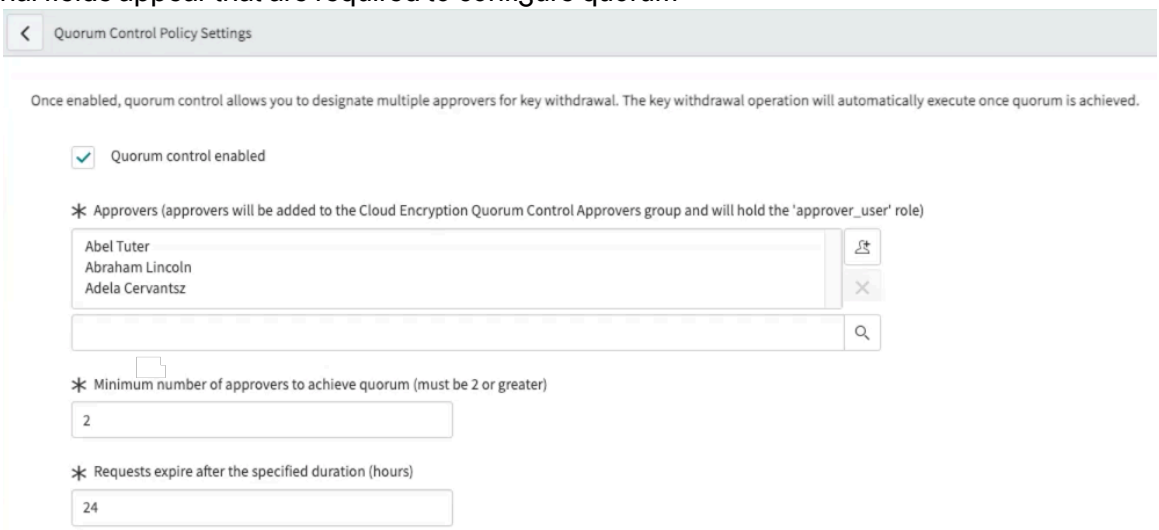

control.

4. Fill in the fields to complete the form.

Field	Description
Approvers	Designate the members of the quorum from the list of users. Select the lock icon to open the user directory. There is no limit to the number of approvers that can be select ed.
Minimum number of approvers to achieve quorum	Designate the minimum number of approvers required to achieve quorum. For example, if there are nine approvers selected, a mini mum of five may be configured for quorum. When five approvals are received in the sys tem, quorum is reached and the withdraw op eration starts. **Note:** The minimum number of re quired approvers is two.
Requests expire after the specified duration (hours)	Set a numeric value in hours that is the maxi mum time allotment for the minimum number of approvals to be obtained. After the time frame expires, the quorum requests also ex pire. A new quorum request is required to con tinue with a withdrawal request.

5. Click **Submit**.
   A confirmation message is displayed.

What to do next

The withdrawal actions are available in Key management operations.

Manage Quorum Control

After a withdrawal operation workflow is triggered, quorum actions can be managed from the Key Management Operations page. The key withdrawal operation is blocked until the quorum is met.

Before you begin

Role required: sn_kmf.admin or sn_kmf.cryptographic_manager

When the quorum has been either approved or rejected, the requestor of the key withdrawal will receive an email notating if quorum was achieved or denied.

Procedure

1. Perform the steps to withdraw a customer managed key found in Key management operations.
2. View the *Quorum Control Requests* and *Quorum Control Approvers* tabs that activated.

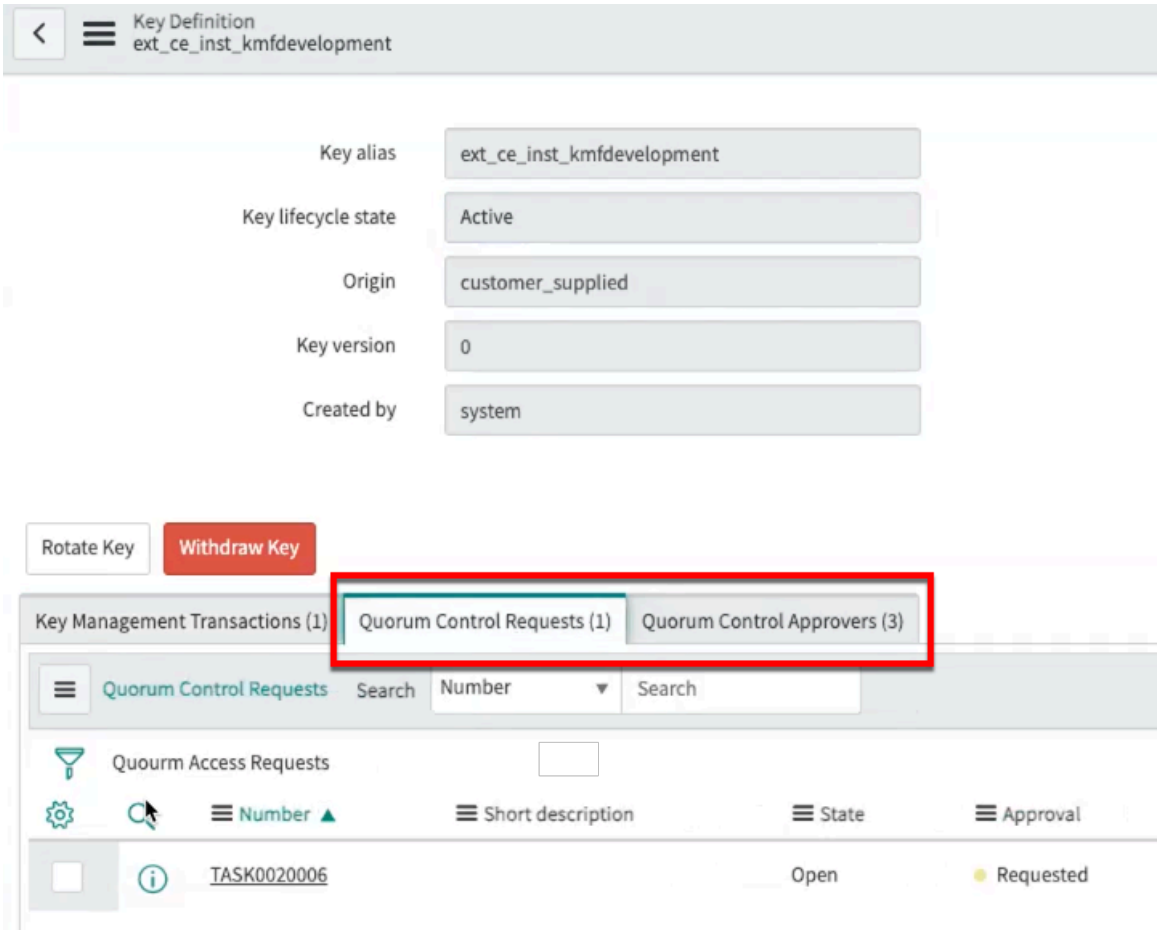

3. Open the **Quorum Control Requests** tab to view the actual request that is created.
   - State:
     - Open: The key withdrawal action is pending the quorum being met.
     - Closed Complete: The quorum has been met and can be no further action on this particular quorum request.
   - Approval:
     - Requested: Approval emails have been sent and the workflow has been triggered to reach quorum.
     - Approved: The key will be withdrawn and the instance will be shut down.
     - Denied: The quorum request is canceled and no further action is taken with this request. A new withdrawal request will be required to withdraw the key.

4. Open the Quorum Control Approvers tab to view the list of approvers and the state of the approval request.

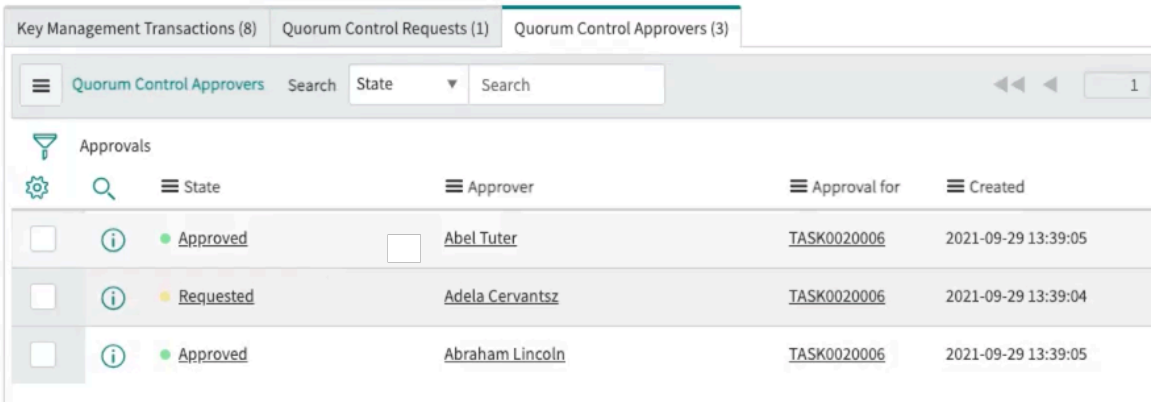

   State:

   - Requested: The approver has not yet taken action on the approval request.
   - Approved: The request has been approved either from the email or the approvals page.

5. Select the **Key Management Transactions** tab to view the progress of the request step for the key withdrawal.
   - Step 0 - Quorum Request: The actual quorum request. The quorum request must be completed in order to trigger the key withdrawal steps.
   - Step 1 - Key Withdrawal: The key withdrawal step. This is composed of steps two through seven.
   - Step 2 - Request_preparation: Creates a request to trigger and the wrapping and rotation.
   - Step 3 - request_integrity_check: Validates that the request is legitimate and secure.
   - Step 4 - request_validation: Validates that there is a request in progress, only one rotate request can process at a time.
   - Step 5 - hsm_key_delete: Makes the call to KeySecure to delete the active key.
   - Step 6 - key_metadata_withdraw: Converts the active key metadata lifecycle state to "destroyed."
   - Step 7 - post_withdraw: Makes a call to shut down the instance.

Approve or deny a quorum control request

Approve or deny a quorum control request from Key Management Transactions.

About this task

When a quorum request has been created, the minimum number of approvals is required by the members. After a withdrawal operation workflow is triggered, quorum actions can be managed using several methods. The users can grant approvals from the Key Management Operations page, My Approvals in the Instance, or directly from the request email. The key withdrawal operation is blocked until the quorum is met.

This procedure describes how to approve or deny a quorum request from the Key Management Operations page.

Before you begin

Roles required: sn_kmf.admin or sn_kmf.cryptographic_manager

Procedure

1. Navigate to **All** > **Cloud Encryption Key Management** > **Key Management Transactions** > **Quorum Control Approvers**.
2. Select your user name from the table.
3. Approve or deny the request.

Approve or deny a quorum request

When a quorum request has been created, the minimum number of approvals is required by the members. After a withdrawal operation workflow is triggered, quorum actions can be managed using several methods. The users can grant approvals from the Key Management Operations page, **My Approvals** in the Instance, or directly from the request email. The key withdrawal operation is blocked until the quorum is met.

Before you begin

Roles required: sn_kmf.admin or sn_kmf.cryptographic_manager

Key management transactions

The Key Management Transactions submodule displays all transactions that have occurred for the keys in your ServiceNow instance.

- A key transaction is defined by the following:
  - composed of several request steps.
  - A single *Request ID* is shared across all request steps.
  - The initial step, request sequence 0, of a transaction provides the current state of the overall transaction.

    As seen in the image below, the initial step 0 has an overall *Request Status* of Completed.
- The following can be identified for the transaction by the individual request step:
  - The order of each step in a transaction can be identified by the sequence number for the step.
  - The status of each transaction is visible through the status of the request step.
  - If any steps beyond the initial step fail, the overall transaction has a status of Failed. If all steps are completed, the transaction status is also completed.

The following screen is a sample of the type of information that displays with a ServiceNow key rotation.

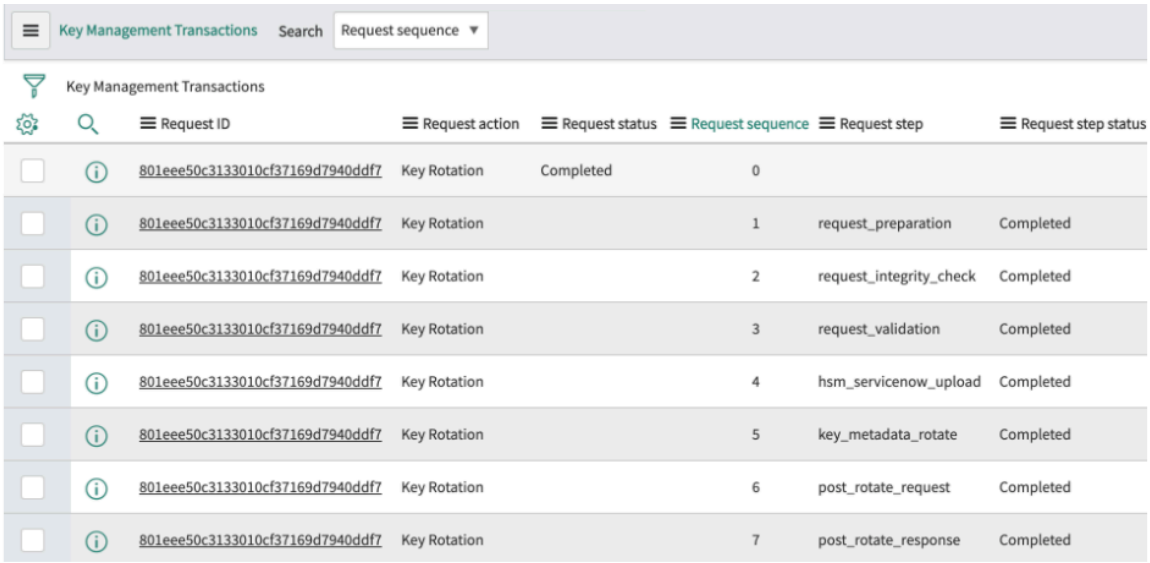

The following table displays the field information available on the Key Managements Transactions page.

Key Management Transactions

Field	Description
Request ID	Unique system-generated Id for the action being performed One request ID is shared across all request steps.
Request action	Displays the action for the key operation being performed.
Request status	• Processing: A request has been entered but hasn't yet been completed. • Completed: The request has been completed successfully. • Failed: An issue occurred and the process hasn't been completed. **Note:** Contact Customer Service and Support and provide the request number where the failure occurred.
Key alias	Alphanumeric entry.
Key life-cycle state	See Key Management Framework key life-cycle states for definitions.
Origin	• ServiceNow key • Customer-managed key
Key version	When a key rotates, the version number increments.
Request sequence	Displays the order in which a request is being processed in the system.
Request step	Displays whether a step is being processed in the system during key rotation. The quantity and content of the steps vary based on the type of key operation performed.

Key Management Transactions (continued)

Field	Description
	1. request_preparation: Creates a request to trigger and the wrapping and rotation. **2.** request_integrity_check: Validates that the request is legitimate and secure. **3.** request_validation: Validates that there's a request in progress, only one rotate request can be processed at a time. **4.** attachment_process: Extracts the wrapped key material from the attachment. (Additional step when rotating a Customer Managed key.) **5.** hsm_<key type>_upload: Uploads the wrapped key material to the HSM, KeySecure. **6.** key_metadata_rotate: Generates the new key metadata. **7.** post_rotate_request: Sends a request to perform the key rotation. **8.** post_rotate_response: Response to perform the key rotation based on the request from the customer instance. **Note:** Provide the request step to Customer Service and Support to analyze the status progression in case a request step doesn't complete.
Request step status	• Completed: Rotation is successful. • Failed: Rotation isn't successful. **Note:** Provide the request step to Customer Service and Support to analyze the status progression in case a request step doesn't complete.

Cloud Encryption logging

Learn about logging options for Cloud Encryption.

Cloud Encryption logging tables

Use these tables to find logging information related to Cloud Encryption transactions on your instance.

Table	Description
Cloud Encryption Metadata [dare_key_metadata]	Cloud Encryption Metadata captures key life-cycle management metadata. On this table you can find key life-cycle, state, and version information. This table is updated after each key operation.
Key Management Transactions [dare_key_request]	Key Management Transactions captures key management transaction information. On this table you can find logging for each step of a transaction. The table records any error information for a transaction in the **error message** field.

Table	Description
Sys Audits[sys_audit]	The Sys Audits table captures inserts and updates to all audited records on your instance. On this table you can find changes to records on your instance, when the changes were made, and which user account initiated the change.

Monitor key rotation operations

Use the Cloud Encryption Key Metadata [dare_key_metadata] table to find information on the life-cycle of your key. In this table you can find information like the origin, activation date, state, and version of your keys.

Use the Key Management Transactions [dare_key_request] table to monitor transactions of key operations. In this table you can find all requests relating to your keys, including the state, status, and which step in the process the request is in. Completed requests are retained on this table with the **Completed** status.

This example shows a key rotation operation. During this operation, the old key life- cycle state updates from active to rotated, and the version state updates from active to retired.

Key definition for a rotated key

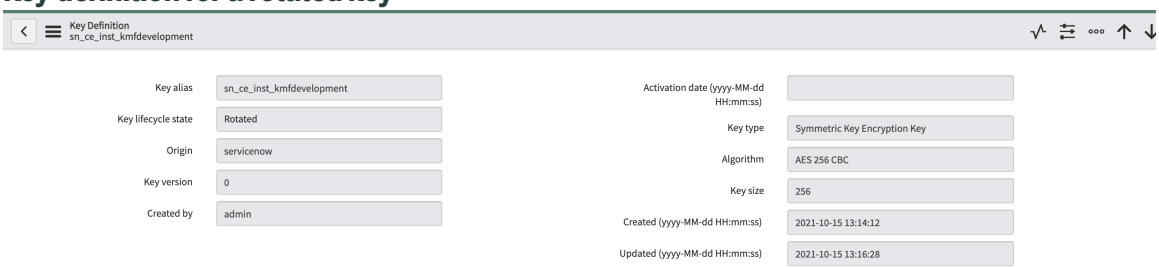

Looking at the Sys Audits[sys_audit] table, admins can see changes made to records on the Cloud Encryption Key Metadata [dare_key_metadata] table. Admins can see which records were updated and when. The log entries also record the field that was changed, and the old and new values.

Audit logs for a withdrawn key

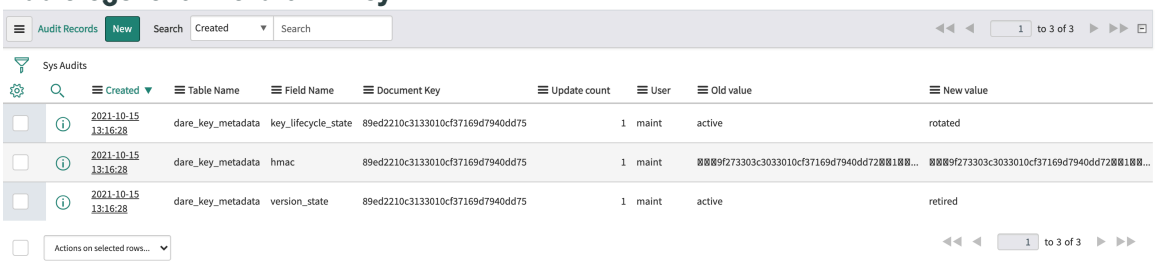

Admins can view the records on the Cloud Encryption Key Metadata [dare_key_metadata] table. In the audit records below, the request status was changed from processing to completed.

Audit logs for a withdrawn key

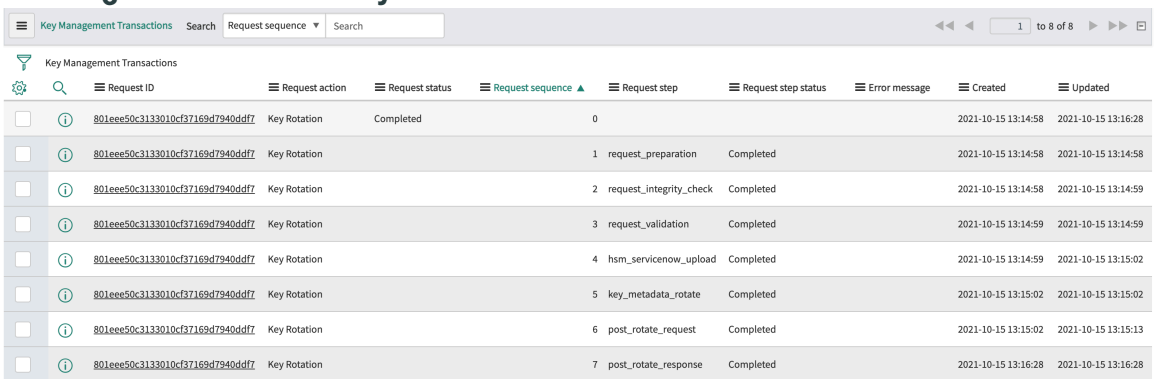

Key Management Transactions | Search | Request sequence | Search | 1 to 8 of 8

Key Management Transactions

Request ID	Request action	Request status	Request sequence ▲	Request step	Request step status	Error message	Created	Updated
801eee50c3133010cf37169d7940ddf7	Key Rotation	Completed	0				2021-10-15 13:14:58	2021-10-15 13:16:28
801eee50c3133010cf37169d7940ddf7	Key Rotation		1	request_preparation	Completed		2021-10-15 13:14:58	2021-10-15 13:14:58
801eee50c3133010cf37169d7940ddf7	Key Rotation		2	request_integrity_check	Completed		2021-10-15 13:14:58	2021-10-15 13:14:59
801eee50c3133010cf37169d7940ddf7	Key Rotation		3	request_validation	Completed		2021-10-15 13:14:59	2021-10-15 13:14:59
801eee50c3133010cf37169d7940ddf7	Key Rotation		4	hsm_servicenow_upload	Completed		2021-10-15 13:14:59	2021-10-15 13:15:02
801eee50c3133010cf37169d7940ddf7	Key Rotation		5	key_metadata_rotate	Completed		2021-10-15 13:15:02	2021-10-15 13:15:02
801eee50c3133010cf37169d7940ddf7	Key Rotation		6	post_rotate_request	Completed		2021-10-15 13:15:02	2021-10-15 13:15:13
801eee50c3133010cf37169d7940ddf7	Key Rotation		7	post_rotate_response	Completed		2021-10-15 13:16:28	2021-10-15 13:16:28

Logging for key withdrawal operations

Logging information on key withdrawal is stored in the Audits [sys_audit] table. This logging information contains information on who initiated the key withdrawal and when the withdrawal took place.

This example shows a key withdrawal operation. During this operation, the key lifecycle state updates from generated, to active, to destroyed. The key version updates from unknown, to active, to retired.

Key definition for a withdrawn key

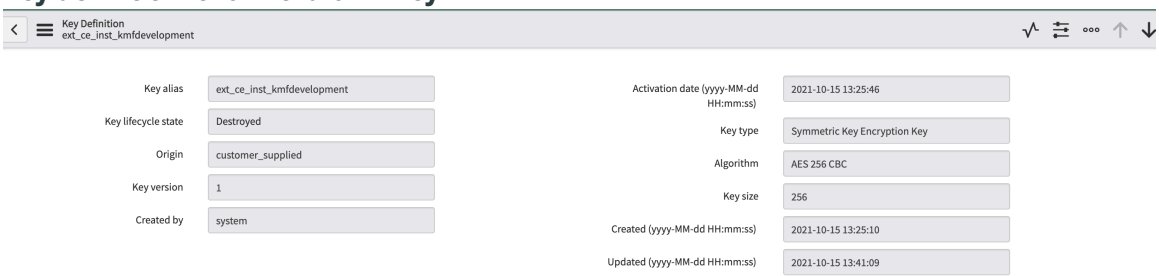

Key Definition
ext_ce_inst_kmfdevelopment

Key alias	ext_ce_inst_kmfdevelopment	Activation date (yyyy-MM-dd HH:mm:ss)	2021-10-15 13:25:46
Key lifecycle state	Destroyed	Key type	Symmetric Key Encryption Key
Origin	customer_supplied	Algorithm	AES 256 CBC
Key version	1	Key size	256
Created by	system	Created (yyyy-MM-dd HH:mm:ss)	2021-10-15 13:25:10
		Updated (yyyy-MM-dd HH:mm:ss)	2021-10-15 13:41:09

Looking at the Sys Audits[sys_audit] table, admins can the Cloud Encryption Key Metadata [dare_key_metadata] table changes.

Audit logs for a withdrawn key

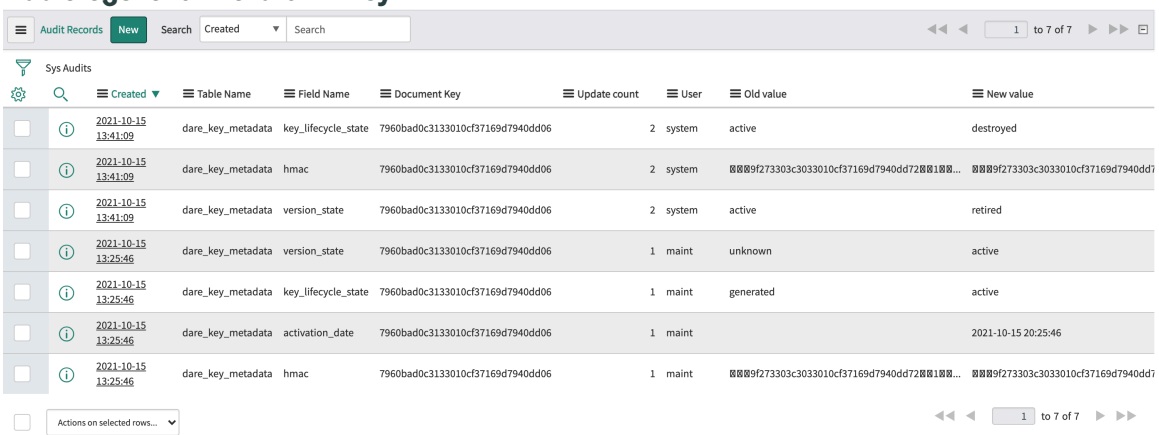

Audit Records | New | Search | Created | Search | 1 to 7 of 7

Sys Audits

Created ▼	Table Name	Field Name	Document Key	Update count	User	Old value	New value
2021-10-15 13:41:09	dare_key_metadata	key_lifecycle_state	7960bad0c3133010cf37169d7940dd06	2	system	active	destroyed
2021-10-15 13:41:09	dare_key_metadata	hmac	7960bad0c3133010cf37169d7940dd06	2	system	▯▯▯9f273303c3033010cf37169d7940dd72▯▯1▯▯...	▯▯▯9f273303c3033010cf37169d7940dd7
2021-10-15 13:41:09	dare_key_metadata	version_state	7960bad0c3133010cf37169d7940dd06	2	system	active	retired
2021-10-15 13:25:46	dare_key_metadata	version_state	7960bad0c3133010cf37169d7940dd06	1	maint	unknown	active
2021-10-15 13:25:46	dare_key_metadata	key_lifecycle_state	7960bad0c3133010cf37169d7940dd06	1	maint	generated	active
2021-10-15 13:25:46	dare_key_metadata	activation_date	7960bad0c3133010cf37169d7940dd06	1	maint		2021-10-15 20:25:46
2021-10-15 13:25:46	dare_key_metadata	hmac	7960bad0c3133010cf37169d7940dd06	1	maint	▯▯▯9f273303c3033010cf37169d7940dd72▯▯1▯▯...	▯▯▯9f273303c3033010cf37169d7940dd7

Actions on selected rows... | 1 to 7 of 7

Tamper Detection

Use tamper detection to improve security by detecting unauthorized changes to your quorum control settings.

Tamper detection process

When enabled, tamper detection validates your quorum control settings by checking for any unauthorized modifications (tampering). Tamper detection uses hash-based message authentication code (HMAC).

1. When a setting is changed or created, your instance creates an HMAC. The HMAC is based on the value of the setting (dare_property) record.
2. Whenever your instance uses these settings, tamper detection validates it using the HMAC.
3. If the setting validates successfully, it can be used by the platform, otherwise it cannot.

Tamper detection runs daily on your instance

Tamper detection checks your settings for tampering using a daily scheduled job, and reports validation failures in your node and security logs. Tamper detection sends a notification to Security and KMF admins for validation failures.

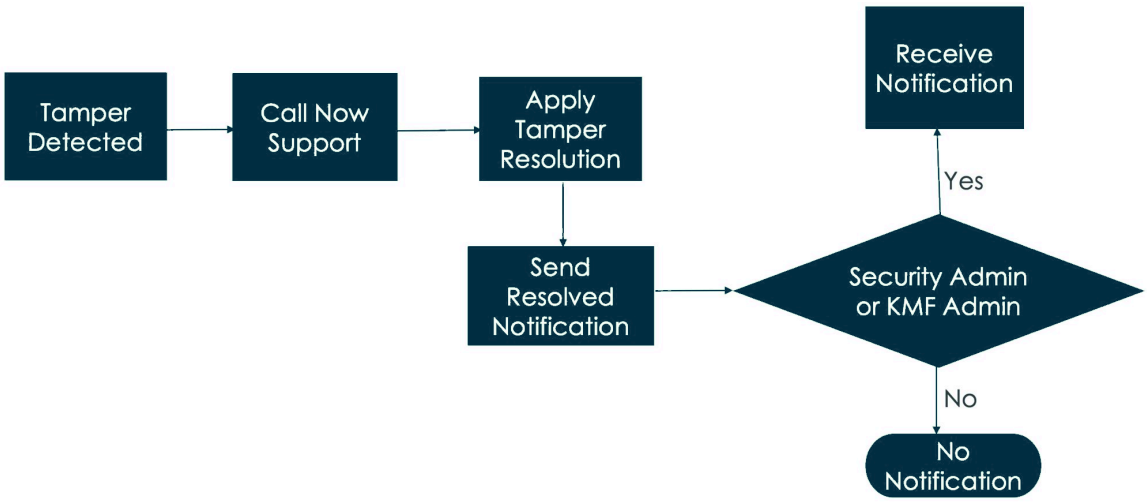

Tamper detection runs before executing a key withdrawal

Tamper detection also validates your properties when you request a key withdrawal. If your settings do not pass validation, the key withdrawal does not execute. In this case, you must resolve any validation issues before key withdrawal can compete.

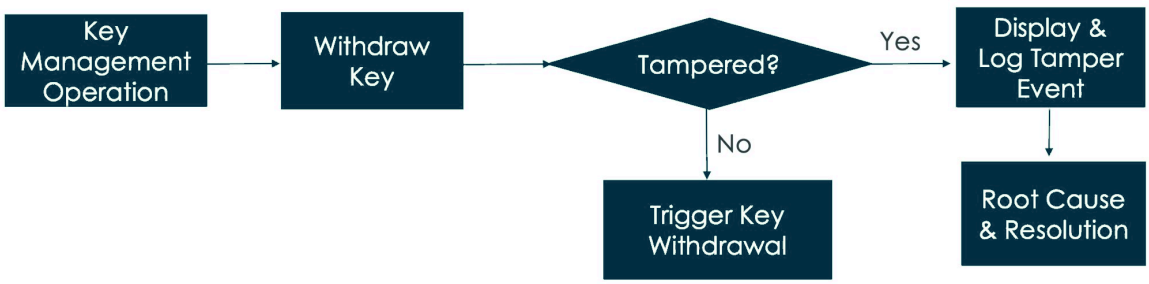

Identifying tampering

Tamper detection updates your logs when validation fails.

If tamper detection fails to validate any of your quorum control settings, these failures appear in your node and security logs. The log entry includes the sys_id of the settings (dare_property) record that failed validation.

```
2022-06-28 13:45:46 (582) Default-thread-5
 B6FAC1F6C3D01110CF37169D7940DD6E txid=231c4d72c310
```

```
SEVERE HMAC_VALIDATION_FAILED:The dare_property record
with sys_id: 776e3200c3210110900b169d7940dd76 failed
HMAC validation

2022-06-28 13:47:35 (264) Default-thread-8
B6FAC1F6C3D01110CF37169D7940DD6E txid=8e8cc972c310
SEVERE HMAC_VALIDATION_FAILED:The dare_property record
with sys_id: 758b3200c3210110900b169d7940dd76 failed
HMAC validation
```

Logging displays information similar to these examples when validation fails. Successful validations don't appear in the logs.

Tamper detection displays a warning message on your quorum control settings page

If a quorum control setting has failed validation, you can see a warning when you view the Quorum Control Policy settings page on your instance. The warning includes the sys_id of the settings (dare_property) record that failed validation.

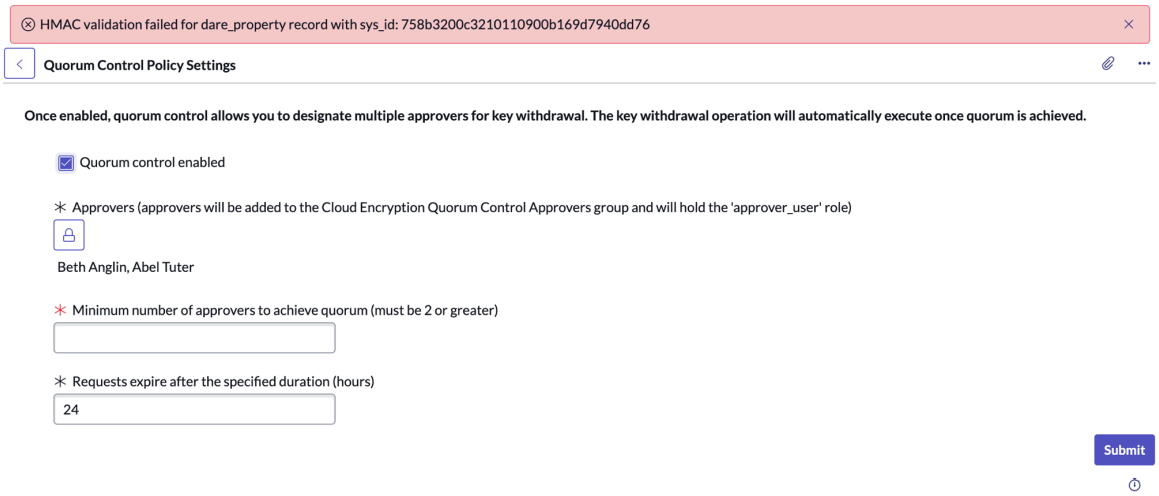

Tamper detection sends notifications to users with the *Security Admin* and *KMF Admin* roles

If tamper detection fails to validate any of your quorum control settings, your security admins and KMF admins receive a notification similar to this example.

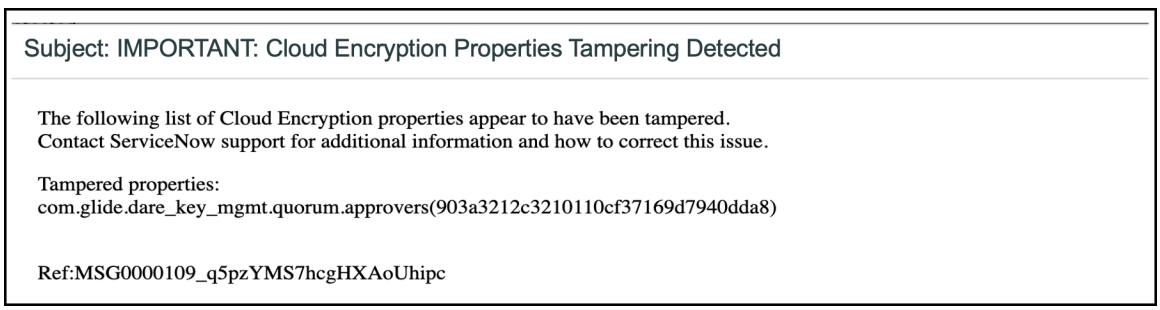

Resolving tampering issues with ServiceNow support

Important: Tamper detection validation failures can only be resolved with assistance of ServiceNow support.

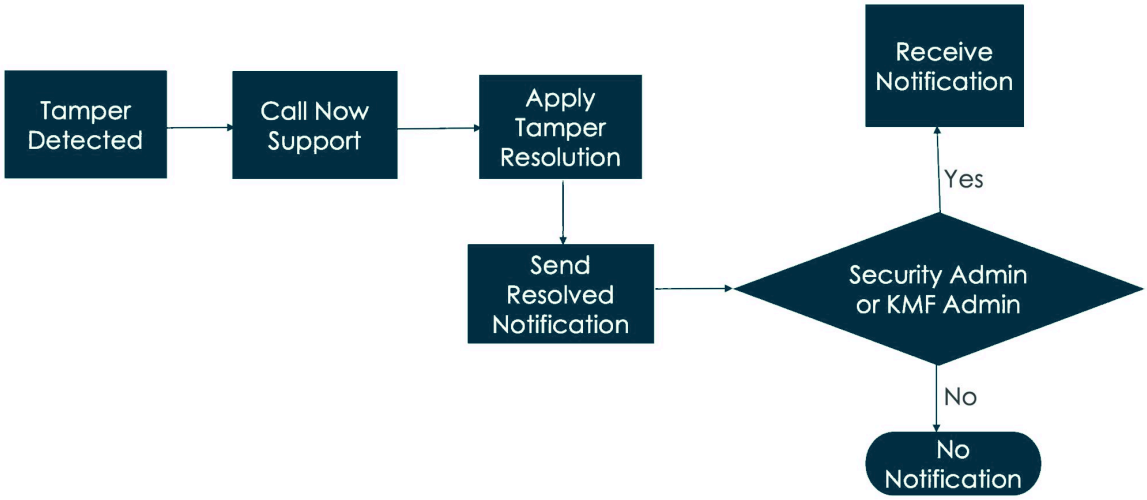

If tamper detection fails to validate any of your quorum control settings, contact ServiceNow support for assistance in resolving the issue. After a support agent has resolved the validation failure, security and KMF admins receive a notification indicating that the issue has been resolved.

Subject: IMPORTANT: Cloud Encryption Properties Reset

At your request, ServiceNow support has reset the following list of Cloud Encryption properties to their default values.
If you did not request this action, contact ServiceNow support immediately.

Reset properties:
com.glide.dare_key_mgmt.quorum.approvers(903a3212c3210110cf37169d7940dda8)

Ref:MSG0000116_0HAz8U1r59OmppqV8ysg

Full disk encryption

Full disk encryption (FDE) applies encryption to the entire storage system within the database server only, because this is the only customer data-storing component. FDE protects only against physical loss or theft of storage devices. When encrypted disk servers are powered on and providing data, the encryption provides no additional protection.

Full disk encryption

Full disk encryption may be relevant to heavily regulated organizations, but can add significant cost to a customer's ServiceNow deployment. Measures in place by ServiceNow to mitigate loss or theft of storage devices may also be a factor in its selection.

From the ServiceNow AI Platform perspective, all data flows are decrypted.

Commercial environments use full disk encryption (FDE) with FIPS 140 validated hardware or storage devices that are in the process of validation along with a ServiceNow dedicated hardware option at extra cost. FDE applies to the hardware itself, and therefore provides encryption at rest for all data stored in every instance assigned to you.

For further details on selecting FDE and dedicated hardware options, contact your ServiceNow representative.

Edge Encryption

ServiceNow® Edge Encryption ™ encrypts sensitive data on your company premises before sending it over the Internet to your ServiceNow instance (encrypted in flight), where it remains encrypted at rest.

Important:

Starting with the Zurich release, Edge Encryption and Edge Encryption Core are being prepared for future deprecation. They will be hidden and no longer activated on new instances but will continue to be supported. Field Encryption and Field Encryption Enterprise provide the latest experience for this functionality.

For details, see the Deprecation Process [KB0867184] article in the Now Support knowledge base.

Explore

Learn the features and business value of Edge Encryption.

Plan

Understand how to plan for Edge Encryption.

Install

Install Edge Encryption.

Upgrading Edge Encryption

Upgrade

Know about how to upgrade Edge Encryption.

Configure

Understand how to configure Edge Encryption.

Exploring Edge Encryption

Edge Encryption is a network encryption system that resides on your network and that encrypts and decrypts sensitive data as it travels between your data center and the ServiceNow cloud.

What is Edge Encryption

Also referred to as 'client-side' encryption, Edge requires all bi-directional user traffic to pass through proxies that are maintained on your infrastructure. You have the full control over your key management because the keys are stored within your proxy on your infrastructure. The ServiceNow AI Platform can't decrypt your ciphertext to access your keys.

The Edge Encryption feature is an additional cost option that provides you with the ability to control the end-to-end encryption of your data and key management. Edge Encryption uses a proxy application, provided by ServiceNow and installed by you within your own network. This proxy application tokenizes specified data patterns or encrypts string fields, Date fields, Date/Time fields, and attachment data before it's sent from your environment to your instance. The proxy application also decrypts the same data, again only within your own network, using keys stored only within your own network.

The relevant encryption keys and configuration exist only on the Edge proxy within your network and aren't visible to ServiceNow. The data is encrypted from the moment that it leaves your environment and is only decrypted upon retrieval. At no point is the data accessible in plaintext by ServiceNow systems or personnel.

Who uses Edge Encryption

Only a user logged in to the instance through a proxy server on your network can view encrypted data in clear text. Likewise, only a security_admin user logged in to an instance through a proxy server in your network can configure and administer Edge Encryption.

Because the proxy server resides in your network, you own and manage the encryption keys, and they're never sent to the instance. As a result, ServiceNow never shows sensitive data in clear text.

In addition to the Edge proxy configuration and management of rules, you're responsible for the usual requirements of operating a server within your environment (including hosting, routing, backup, DNS configuration, and so on) to enable and support your Edge proxies.

Encryption and tokenization

Edge Encryption supports both encryption (through encryption configurations) and tokenization (through encryption patterns) as a means of protecting your sensitive information.

Encryption configurations

You can encrypt individual fields using encryption configurations. Edge Encryption supports AES 128-bit and AES 256-bit encryption keys. Edge Encryption supports standard, equality-preserving, and order-preserving encryption types.

In addition to attachments, you can encrypt the following field types:

- Date
- Email

- Date/Time
- IP Address
- Journal
- Journal Input
- Multi-line text
- Single-line text
- String
- URL

Note:

If a Journal field marked for encryption is added to the activity stream, all user input to the field is encrypted in the activity stream.

Multi-byte characters within supported field types can be encrypted.

You can also encrypt the following service catalog variable types:

- String Types
  - Single-line text
  - Multi-line text
  - Wide single-line text
- Date
- Date/Time
- URL
- Email
- HTML
- IP Address

Encryption patterns

You can use encryption patterns to tokenize strings that match regular patterns such as social security and credit card numbers. While encryption configurations should be the primary method of encryption, use encryption patterns as a supplement to secure sensitive information found outside of encrypted fields.

Note: The Edge Encryption proxy server requires a MySQL database in your network only if using order preserving encryption or encryption patterns. Clear text values are stored in the proxy database in your network. For this reason, it is critical that you secure and regularly back up your proxy database. For recommendations, see Edge Encryption components.

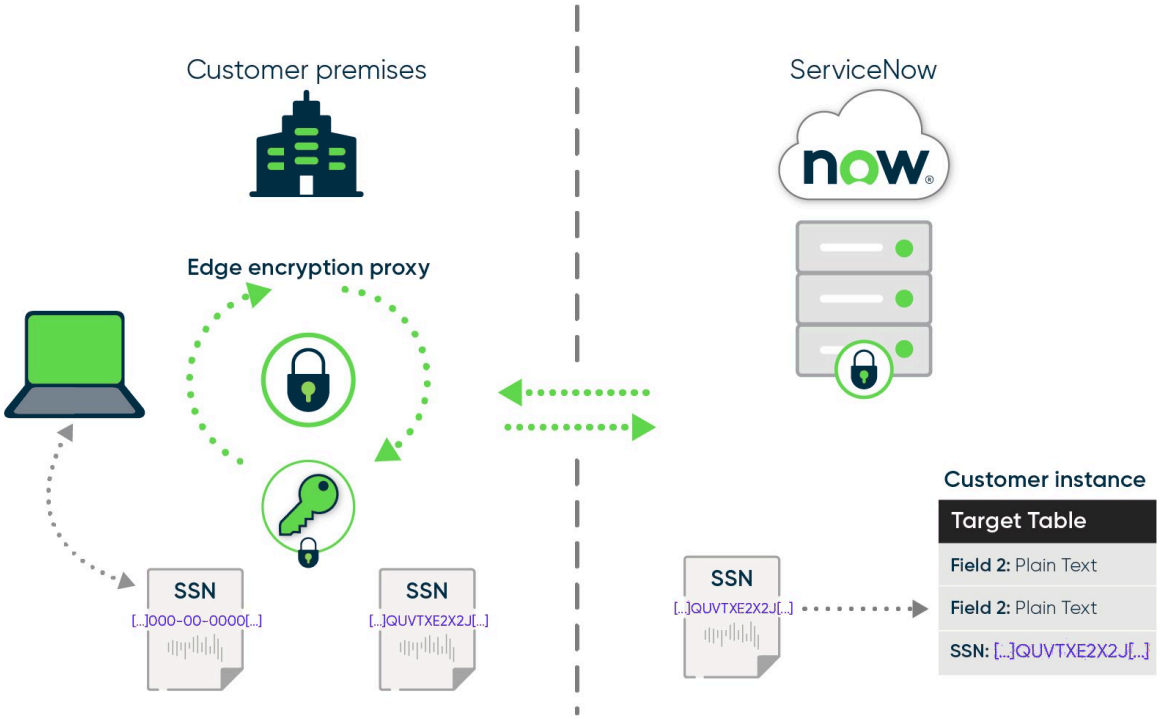

Edge Encryption on the ServiceNow AI Platform

Edge Encryption acts as a gateway between your browser and your ServiceNow instance. Traffic from your browser passes through the gateway on its way to the ServiceNow instance. The gateway, in turn, is configured to encrypt outbound data that is marked for encryption. Inbound traffic is decrypted through the gateway, and the end user sees clear text in the browser. The advantage of this implementation from a security control perspective is that the encryption and key management are handled externally from ServiceNow.

Pros and cons

As with Field Encryption Enterprise and Field Encryption, Edge Encryption imposes some functional limitations within an instance as a result of the additional security. The local Edge proxy does, however, also provide some additional functionality relating to sorting when compared to column-level encryption.

Pros:

- Edge Encryption provides absolute control of who sees your information and prevents data breaches.
- Information remains on your proxy server and never leaves your network unencrypted.
- Information is encrypted in transit, before it even reaches the ServiceNow instance.
- You hold and manage all your own encryption keys. No one else, not even ServiceNow personnel, can access your keys.
- You can choose the strength of the encryption algorithm: AES-128 or AES-256.
- Edge Encryption includes the ability to encrypt String text, Date and Date/Time fields, attachments, URLs, and journals.
- Edge Encryption provides Standard, Equality Preserving, and Order Preserving encryption of data at rest within the database and instance.

- Encryption rules enable you to write custom scripts that tell the proxy server specifically what to encrypt and where to put that encrypted information in the instance. These scripts are useful when the data structure doesn't exactly match the ServiceNow instance.
- Encryption patterns enable you to tokenize information such as passwords.

Cons:

- Edge Encryption requires an extra network hop through the Edge proxies cluster, and extra processing, which can add delay to traffic. The added processing delay of the Edge Encryption application is negligible compared to the network hop.
- Maintaining your own encryption keys can be complex and time-consuming.
- A maximum of two encryption keys can be active at any given time. You cannot assign different keys to specific subsets of columns, data categories, user roles, or access scopes. However, one key can encrypt certain columns while a different version of the same key can be used to decrypt others.
- Edge Encryption has the side effect that the server or platform can't decrypt the data to perform any manipulation of the decrypted data. As a consequence, functionality and data processing on the ServiceNow AI Platform may be restricted when encrypting columns with Edge Encryption.

What to know before you begin

Because encryption and tokenization change the nature of your data, Edge Encryption can affect other instance processes. Before using Edge Encryption, carefully consider the impact on your instance.

Because the proxy server is installed and maintained in your network, Edge Encryption requires network administration and management. Review the network requirements to ensure a smooth implementation.

Review the following topics to understand the impact of Edge Encryption on your instance:

- Planning for Edge Encryption
- Edge Encryption system requirements
- Sizing your Edge Encryption environment
- Calculate the order-preserving and tokenization database size
- Edge Encryption limitations
- Key management for Edge Encryption

Edge Encryption components

Edge Encryption is composed of the Edge Encryption proxy server that runs on a server in your network, and the Edge Encryption plugin that must be installed on your ServiceNow instance. If using order-preserving encryption types or encryption patterns, a proxy database must also be installed in your network.

Proxy application

When going through the Edge Encryption proxy server, the Edge Encryption plugin enables you to specify which fields, patterns, and attachments should be encrypted. You can also manage encryption rules to encrypt specific requests and can schedule mass encryption jobs.

Proxy server

The Edge Encryption proxy server uses encryption rules to identify in an HTTP request what, if anything, must be encrypted and encrypts it before forwarding the request to the instance. For decryption, the Edge Encryption proxy server looks at the HTTP responses for any encrypted data and decrypts it before sending the response back to the client. For this decryption to happen, all HTTP requests and responses must go through the Edge Encryption proxy server. These HTTP requests include any requests originating from a browser, as well as any SOAP or REST requests.

Proxy database

If using order preserving encryption or encryption patterns, your proxy servers rely on a MySQL database located in your network. All proxy servers in your network must use the same database.

The proxy database contains these tables.

Proxy database tables

Name	Description
db_id	Unique database ID
edge_token_map	Encryption pattern data
token_map	Order-preserving encryption data

Backing up your proxy database

Because encryption patterns rely on tokenization, clear text values are stored in your proxy database. If the database is lost, clear text values can't be restored. It's critical that you maintain regular backups. To avoid data loss, back up your proxy database according to ServiceNow recommendations.

- Back up your database every 24 hours.
- Retain MySQL database binary log files for at least two days. After a backup has been restored, use the binary log to regenerate any data lost since the most recent backup.

Edge Encryption clients

Edge Encryption uses three clients to inform the instance that the proxy is running, to synchronize requests between the proxy and the instance, and to forward all end user requests to the instance after any potential encryption.

Client	Description
heartbeat/keepalive client	In charge of sending a request to the ServiceNow instance every 5 seconds to let the instance know that this proxy is up and running. The requests drive the last_response_on field on the Edge proxy table, and as a consequence drive the state of the proxy. If your system has issues sending the requests, or if the request or request processing is delayed, the instance may mark the proxy as unresponsive, even if the other

Client	Description
	clients (including the one for user traffic) are up and running. This client also controls the online status of the proxy on the instance. The *edgeencryption.proxy.keepalive.interval* property controls the polling rate for this client. The default is 5 (seconds).
polling/sync client	In charge of various requests the proxy sends to the instance to synchronize on the Edge Encryption configuration (for example, which table, column, or attachment to encrypt, keys, jobs, rules, and tokenization patterns). The *edgeencryption.config.poll.interval* property controls the polling rate for this client. ⚠ **Warning:** Do not change this setting. Changing the default setting of the proxy poll interval may result in synchronization delays when updating Edge Encryption settings on the instance.
default/user traffic client	For everything else, this client handles all end user requests and forwards them to the ServiceNow instance after any potential encryption. This client also handles responses from the instance, forwarding them to the end user after any potential decryption.

Key management for Edge Encryption

You are responsible for providing and managing the encryption keys used by Edge Encryption.

This topic refers to keys for the Edge Encryption product. If you are looking for information on the Key Management Framework, which can be used with Field Encryption, see Key Management Framework.

When obtaining and creating encryption keys to support the encryption types used by Edge Encryption, consider the following:

- Whether to use AES 128-bit or AES 256-bit. You must define a default AES 128-bit encryption key, even if it is not used.
- Whether to use file system, Java KeyStore, or Enterprise Key Management (EKM).
- When to rotate encryption keys.
- When and if to use a mass encryption job to re-encrypt data using the new key.

Before removing a key from the proxy configuration files and the keystore, it is critical that you decrypt all data on the instance that uses the key. You can do this by adding a new encryption key and scheduling a mass key rotation job.

Keystores

Edge Encryption supports the following types of key storage.

File store

Keys are stored in a file in a file system that is accessible by the Edge Encryption proxy. Encryption keys stored in a file are not encrypted, so it is your responsibility to protect these files.

Java KeyStore

Keys are stored in Java's JCEKS KeyStore. A Java KeyStore is protected by a password, so it is more secure than storing keys in a file in the file store. A single Java KeyStore can store multiple keys, and the keys are identified by a key alias, making it easier to manage multiple keys.

Enterprise Key Management (EKM)

Keys are stored and retrieved with the SafeNet KeySecure or Unbound Technology key management systems.

The Edge Encryption proxy ships with the Java JCEKS KeyStore file named `keystore.jceks` in the `keystore` directory. This keystore file contains the ServiceNow public key used to validate encryption rules signed by ServiceNow.

Note: If using a keystore other than the base system Java JCEKS KeyStore, you must import the ServiceNow public key into your keystore. The public key alias is *`servicenow`*.

In addition to the encryption keys, the Java JCEKS KeyStore is used to store the RSA key pair for digitally signing the encryption configuration and encryption rules that are stored in the instance, and the digital certificate that the Edge Encryption proxy uses to establish a secure connection with the browsers and any other clients.

SafeNet key versioning for Edge Encryption

Use SafeNet key versioning to simplify changing keys. Instead of creating an alias for every new key, SafeNet key versioning keeps the same alias and increments the version.

You must set up key versioning in SafeNet before you can configure SafeNet key versioning on the Edge proxy server.

Note: Edge proxies installed before the London release support SafeNet keys, but do not support SafeNet key versioning. If you mistakenly use a versioned key on a Kingston or earlier proxy, when you upgrade to a London or later release, the London or later proxy detects this problem, and to prevent potential data loss the proxy does not start.

You must first schedule a mass key rotation job or a single key rotation job to replace the old SafeNet versioned key with a non-versioned key, and then create a new SafeNet versioned key, if needed. This new versioned key is safe to use with the London or later proxy, and you can restart the proxy.

Encryption key configuration

If using SafeNet versioned keys, the Change Default Keys section of the Encryption Key Configuration form includes new fields for the **Key version** of the default 128-bit and 256-bit keys. **Key version** fields are grayed out and cannot be edited.

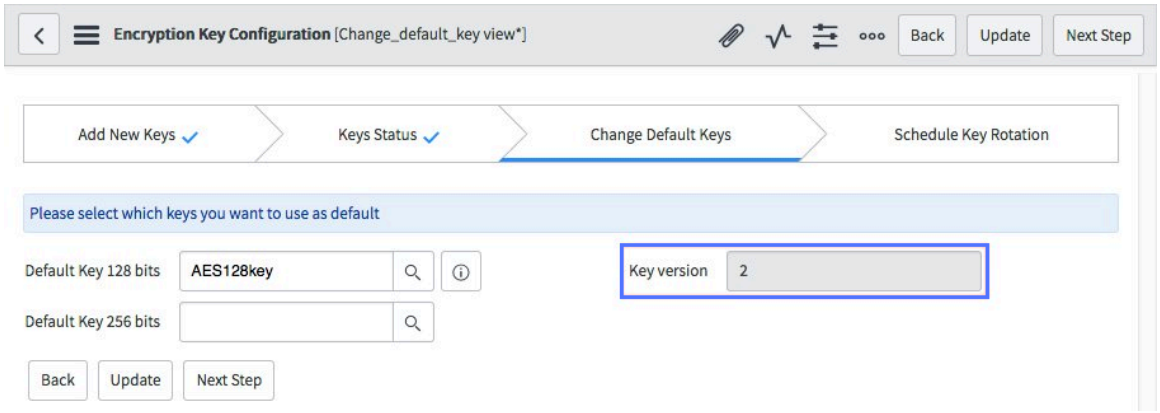

For procedures, see Configure encryption keys on the instance.

Versioned keys

If using SafeNet versioned keys, when you navigate to **Edge Encryption Configuration** > **Encryption Key Configuration** > **All Keys**, versioned keys include the **Key version**.

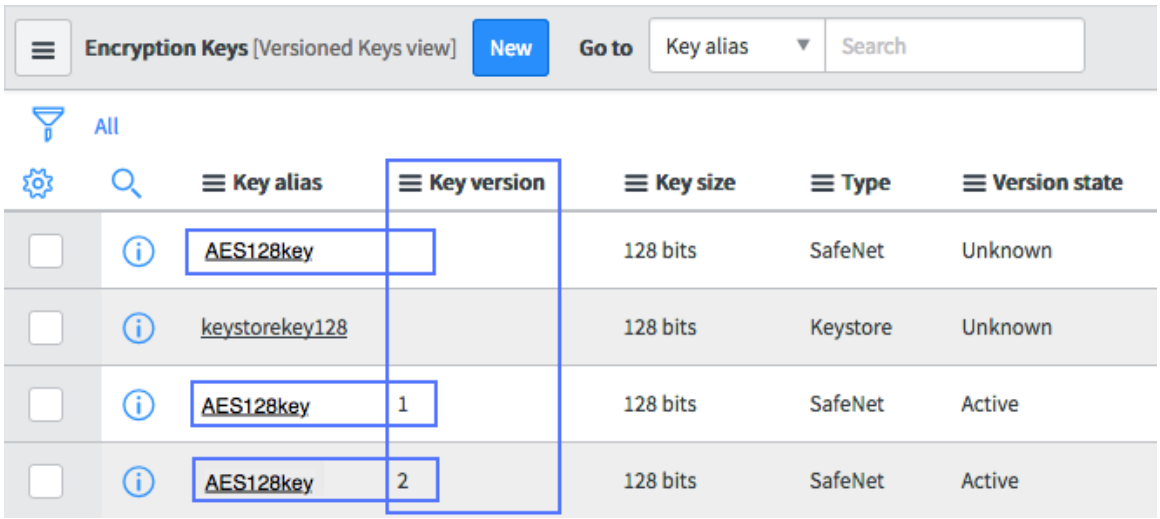

A version number does not appear for the initial entries you make in the Change Default Keys section of the Encryption Key Configuration form. When the proxy server requests a key from SafeNet, the system adds a new line for the alias and adds the **Key version**.

In the above example, **AES128key** is listed three times:

- The first listing, with no **Key version** indicated, is the initial entry.
- The second listing, with **1** in the **Key version** column, is the first version of the key returned from SafeNet.
- The third listing, with **2** in the **Key version** column, is the second version of the key returned from SafeNet.
- As other versions of the key are returned from SafeNet, new lines are added to record the **Key version** now in use.

Encryption configurations and patterns

With Edge Encryption, you can encrypt fields and tokenize strings.

Encryption configurations

You can encrypt individual fields using encryption configurations. Edge Encryption supports AES 128-bit encryption keys. If the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy files are installed, Edge Encryption supports AES 256-bit encryption keys for each encryption type. Edge Encryption supports the following types of encryption configurations.

Standard encryption

The encrypted value of a field is different each time the field is encrypted, even when the field value remains the same. Standard encryption is the most robust form of encryption. Fields using standard encryption cannot be sorted, grouped by, or filtered on.

Equality-preserving encryption

The encrypted value of a field is the same when the field value remains the same. Supports equality comparisons and group by operations on a field.

Note: When equality-preserving encryption is selected for a field that already contains data, performing a group by action on the field may not group fields with the same value if one is encrypted and the other is not.

Order-preserving encryption

Uses tokens and encryption to secure data in your proxy database. Supports equality comparisons, group by operations, and the ability to sort data. The order preserving encryption type is only supported if there is a MySQL database configured for the Edge Encryption proxy server.

Note: When using order-preserving encryption and the proxy database is down, updates can be made to fields using order-preserving encryption. However, the sort order will not be correct when trying to sort data based on those fields. Groups also will not work as expected. When the proxy database is again operational, schedule an order token repair job to repair missing tokens.

Encryption types

The following encryption types are listed in decreasing security quality.

Encryption type	Description
Standard AES 256	Fields cannot be filtered, sorted, or compared.
Standard AES 128	Fields cannot be filtered, sorted, or compared.
Equality preserving AES 256	Fields can be filtered using equality comparisons.
Equality preserving AES 128	Fields can be filtered using equality comparisons.
Order preserving AES 256	Fields can be sorted and equality comparison filtering can be used. Requires the use of a MySQL database in your network.

Encryption types

The following encryption types are listed in decreasing security quality.

(continued)

Encryption type	Description
Order preserving AES 128	Fields can be sorted and equality comparison filtering can be used. Requires the use of a MySQL database in your network.

Encryption Patterns

You can secure sensitive data found in strings using encryption patterns. Once an encryption pattern is stored and activated, the Edge Encryption proxy server identifies strings that match the pattern in requests. Once located, the clear text string is stored in the proxy database and replaced on the instance with a token. Use encryption patterns to tokenize strings that match regular patterns such as social security and credit card numbers. While we recommend that encryption configurations be the primary method of encryption, use encryption patterns as a supplement to locate and secure sensitive information found outside of encrypted fields.

Note: The Edge Encryption proxy server requires a MySQL database in your network only if using order preserving encryption or encryption patterns. Clear text values are stored in the proxy database in your network. For this reason, it is critical that you secure and regularly back up your proxy database. For recommendations, see Edge Encryption components.

Related topics

Encrypt fields using encryption configurations

Tokenize strings using encryption patterns

Installed with Edge Encryption

Edge Encryption installs tables to store encryption-related data, system properties to configure default behavior, and the edge_encryption role to administer Edge Encryption.

Tables installed with Edge Encryption

Edge Encryption adds the following tables.

Edge Encryption Configuration [sys_encryption_configuration]

Contains encrypted fields and tables for which attachments are encrypted.

Edge Encryption Rule [sys_encryption_rule]

Contains a record for each rule. A rule has a name, the condition when it is used, a script, and an order field.

Edge Encryption Invalid Insert Log [sys_edge_encryption_invalid_insert_log]

Contains log messages created for attempts to save unencrypted data to an encrypted field.

Edge Encryption Proxy [sys_encryption_proxy]

Contains information about the encryption proxy application.

Edge Proxy Encryption Type [sys_proxy_encryption_type]

Used for enabling and disabling encryption types on the encryption form.

Encryption Job Execution [sys_encryption_job_execution]

Supports mass encryption jobs.

Encryption Job Execution Chunk [sys_encryption_job_execution_chunk]

Supports mass encryption jobs.

Scheduled Encryption Job [sysauto_encryption_job]

Lists scheduled jobs for encryption, decryption, key rotation, order token repair, and database recovery.

Encryption Key Configuration [sys_encryption_key_configuration]

Lists default encryption keys.

Encryption Key [sys_encryption_key]

Lists available keys and key attributes.

Proxy Encryption Key [sys_encryption_proxy_key]

Lists proxy encryption keys.

Properties installed with Edge Encryption

Edge Encryption adds the following properties.

Note: To open the System Properties [sys_properties] table, enter `sys_properties.list` in the navigation filter.

glide.edge.pattern.disallowed.chars

A list of characters that are not allowed in patterns.

- **Type**: a string of a comma-separated list of values
- **Location**: System Properties [sys_properties] table

glide.edge.pattern.min.size

The minimum pattern size allowed. Allowing smaller patterns means finding more matches, which increases overhead.

- **Type**: number
- **Default value**: 5
- **Location**: System Properties [sys_properties] table

sn_edge_encryption.logging.destination

Where messages are logged.

- **Type**: string
- **Default value**: file
- **Location**: System Properties [sys_properties] table

sn_edge_encryption.logging.verbosity

The logging level to use.

- **Type**: string
- **Default value**: info
- **Location**: System Properties [sys_properties] table

sn_edge_encryption.encryption.proxy.buildtag

The proxy version registered with your instance.

- **Type**: string
- **Location**: System Properties [sys_properties] table

sn_edge_encryption.cleartext.allowed

When true, allows clear text to be saved in an encrypted field. This happens when a user is accessing the instance without going through the Edge Encryption proxy. When false, the system prevents clear text from being saved in an encrypted field.

- **Type**: Boolean
- **Default value**: false
- **Location**: System Properties [sys_properties] table

Planning for Edge Encryption

Successful implementation of Edge Encryption requires planning and preparation.

Answer the following questions in the planning stage.

- Which fields are to be encrypted?
- Which encryption types are to be used?
- How many Edge Encryption proxies are needed? See Sizing your Edge Encryption environment for recommendations and considerations.
- If an order preserving encryption type or encryption patterns are to be used, where is the MySQL database located?
- Which key management system is to be used?

System administrators, network administrators, and security team members have different tasks to fulfill for implementing Edge Encryption.

- System administrators need the security-admin role. The system administrator must :
  - Download the Edge Encryption proxy application.
  - Set up an Edge Encryption user account for the proxies to use to connect to the instance. The user must be assigned the edge_encryption role.
  - Configure encryption keys, and set the default keys.
  - Configure Edge Encryption on the instance.
  - Schedule encryption jobs.
  - Monitor Edge Encryption.
  - Create and edit encryption rules.
- Your network administrator must:
  - Install the Edge Encryption proxy application.
  - Know the network addresses for the proxy servers and the proxy database used for order-preserving encryption and encryption patterns.
  - Install the proxy database to be used for order-preserving encryption and encryption patterns.
  - Start and stop the proxy applications.
  - Perform encryption key management.
  - Determine how to map users to encryption proxy applications. This can be done with DNS settings or routing rules, and is specific to each network.

  - Manage multiple proxy servers.
  - Configure load balancer pools and settings.
- Your security administrator must determine the encryption types to be assigned to each field.

Edge Encryption system requirements

You can run the Edge Encryption proxy application on servers or virtual machines that run on Microsoft Windows or Linux operating systems. For optimum performance, ensure that your configuration meets these requirements.

Java requirements

The host machine installing or running the Edge Encryption proxy server must maintain a supported version of Java. Current supported versions are Java 11.0.6 or later in the 11.x version series

Note: Java 8 is no longer be supported as of the Utah release. Upgrade your environment with the Edge Encryption proxy to Java 11 before you attempt to install the Utah version of the Edge Encryption proxy.

Note: Java does not automatically allow unlimited strength keys. You must specifically enable the use of AES 256-bit encryption.

Support for OpenJDK

The ServiceNow AI Platform supports OpenJDK version 11.

Proxy server minimum configuration

A proxy server requires this minimum configuration:

- 4 GB of RAM per proxy server (6 GB is recommended for most deployments).

  Note: The proxy server host requires at least 1 GB of RAM more than the proxy server. The proxy server host needs the extra 1 GB for operating system services. For example, if you configure a proxy server to use 4 GB of RAM, you must install at least 5 GB of RAM on the proxy server host.

  Because the proxy server requires at least 4 GB of memory, 32-bit JREs and 32-bit operating systems are no longer supported starting with the London release.

- 3 or more GHz CPU (4-core CPU preferred for optimum performance).
- Multiple proxy servers behind a load balancer. The number of proxy servers you need depends on the number of application nodes, the number of simultaneous users, and the number of servers needed for failover. See Sizing your Edge Encryption environment for more information.
- Ability to run concurrently with other services, depending on the server utilization and resource availability.

Proxy server supported systems

The following systems are supported:

Supported System	Description
Windows Server 2012, 2012-R2, 2016, and 2019 editions	• Virtual machines or physical hardware • 64-bit systems
Linux	• Virtual machines or physical hardware • 64-bit systems On 64-bit Linux systems, you must install the 32-bit GNU C library (glibc). The installation command for CentOS is `yum install glibc.i686`.

Proxy server version requirements

Keep your Edge Encryption proxy version in sync with your ServiceNow instance version (same major release, for example Tokyo). To eliminate downtime during the upgrade process, the Edge Encryption proxy is backwards compatible. However it is important to upgrade as soon as possible to avoid ensure users can access new features and important bug fixes.

Proxy server connection requirements

The proxy server that runs the Edge Encryption application must be able to communicate with machines in your network. Make sure that the proxy server has these network privileges:

Network Privilege	Description
Firewall access	Configure any firewalls between the proxy server and the client devices to allow a connection. If your network uses a DeMilitarized Zone (DMZ) to add an extra layer of security to your Local Area Network (LAN), and if your network security protocols limit port access from within the network to the DMZ, you might have to deploy a proxy server to a machine within the DMZ.
Network access	Configure each client to enable the proxy server to connect with it. If network security prevents you from configuring new machines that can connect to the clients, install the proxy server on an existing machine with connection privileges.
Instance access	Ensure that the proxy server has network access to the instance. Make sure that you configure the proxy server network to allow traffic over TCP port 443.
Network account	Install the proxy server with either a local or domain administrator.

Order-preserving and tokenization database system requirements

Order-preserving encryption and encryption patterns require that you configure an Oracle MySQL database for the Edge Encryption proxy server. Order-preserving encryption allows any comparison operation to be directly applied on encrypted data, without first decrypting the data. Encryption patterns let you replace string patterns with tokens (called tokenization) before they are sent to and stored in the database. Because of the size of the MySQL database, use a dedicated proxy server to run the order-preserving and tokenization database.

The minimum database system requirements include:

MySQL Database	Requirement
Version	MySQL database versions 5.7 and 8.0 **Note:** MySQL versions 5.5 and 5.6 are no longer tested and have reached the end of support.
OS	64-bit systems
CPU	2 or more GHz CPU (4-core CPU preferred for optimum performance)
RAM	16 GB
Disk	Storage Area Network (SAN) or local storage (RAID 10 recommended)
Size	Determined by the number of potential records multiplied by the record size. See Calculate the order-preserving and tokenization database size.
Configuration	High availability cluster. If you are unsure of how to configure your MySQL server, contact MySQL for configuration information.

Sizing your Edge Encryption environment

Choosing the number of proxy servers for your environment is an important task. Consider the number of users, redundancy needs, and acceptable latency.

Redundancy

Maintain redundant proxy servers in case of hardware failure. Proxy servers should be located behind a load balancer to provide a functional path for all users if a proxy server is unreachable. At a minimum, ensure that two proxy servers are always available.

Size

Size refers to the number of proxy servers required to avoid additional latency that the encryption of data produces. Depending on use, you may want to reduce the amount of latency by adding additional proxy servers. For example, if regular mass encryptions are run, add additional proxy servers to handle the load, or run the mass encryptions when the user load is light. In addition, the hardware that the proxy server runs on influences performance and latency. Proxy servers running on hardware with faster CPUs, more CPUs, and more RAM have higher throughput than slower, limited systems.

The following guidelines assume that your proxy server is running on at least the minimum hardware requirements. To determine the number of proxy servers:

- Consider setting up one proxy server for every two application nodes on the instance.
- For redundancy, set up a minimum of two proxy servers behind a load balancer.
- Add an extra proxy server for every 500 simultaneous users.
- Depending on the desired redundancy, add additional proxy servers for failover.

For example, for an instance with 2,000 users, you should have at least five proxy servers behind a load balancer. This calculation includes one proxy server for every 500 users, with an extra proxy server for failover. Determine ahead of time when you will approach a threshold of 500 users and place another proxy server in the load balancer pool.

Load balancers

To balance requests and improve server response time, distribute proxy servers in a load balancer pool. Configure load balancers to use the "least connections" method. This method connects requests to the proxy server with the fewest active connections, preventing the overloading of a single proxy.

CPU utilization

Because data encryption and tokenization are CPU intensive operations, CPU spikes while encrypting data are normal and expected. When CPU utilization is over 80% for several minutes at a time, it likely means that the proxy server has too much work to do. When this happens, latency increases for the period that the CPU utilization is high. If latency persists, adding another proxy server may help decrease the latency.

Memory

The proxy server must have a minimum of 4 GB of RAM available (6 GB recommended). Set the proxy server initial and upper bound memory limits to the recommended settings.

Calculate the order-preserving and tokenization database size

If using order-preserving encryption or encryption patterns, determine the size of your MySQL database by multiplying the number of potential records by record size.

Before you begin

Role required: admin

About this task

Use a dedicated machine to run the order-preserving and tokenization database. Do not run the database on the same hardware as the proxy server.

Procedure

1. Determine the potential number of records that could include fields encrypted with order-preserving encryption.
   a. Multiply the number of encryption configurations using order-preserving encryption by the number of records each configuration is applied to.
   b. To allow for growth, multiply the result by three.
2. Multiply the result of step 1 by 1,536.

1,536 is the average size of a record in bytes.

3. If using encryption patterns, perform steps 1–2 for tokenized records and add the result to the total.

Result
The calculated value is the recommended size in bytes for your order-preserving and tokenization database.

Edge Encryption limitations

Edge Encryption impacts system functions. Carefully evaluate the impact of encrypting a field.

Field type restrictions

You can encrypt only the following field types:

- Date
- Email
- Date/Time
- IP Address
- Journal
- Journal Input
- Multi-line text
- Single-line text
- String
- URL

You can't encrypt the following field types:

- Choice fields
- HTML
- Virtual fields
- Fields in system tables, except for certain fields in sys_user
- System fields in tables
- Number fields or fields associated with an auto-numbering scheme
- Any other field type not listed above

Additional restrictions:

- When a Journal field is encrypted, the **Post** button is inactive, even if there are multiple Journal fields and only one of those fields is encrypted.
- Encrypted fields aren’t available in **Go to** and header filter boxes.
- When encrypting fields used as an index, you can use only order-preserving and equality-preserving encryption types. Indexed fields can’t be encrypted using the standard encryption type.

For more information, see Field types.

Filtering and searching restrictions

Standard encryption

When you select a String, Date, Date/Time, or URL field with a standard encrypted field configuration as the left operand in a filter, no filtering options are available.

Equality-preserving encryption

When you select a String, Date, Date/Time, or URL field with an equality-preserving encrypted field configuration as the left operand in a filter, the following operators are available:

- **is**
- **is not**
- **is empty**
- **is not empty**

Order-preserving encryption

When you select a String field with an order-preserving encrypted field configuration as the left operand in a filter, the following operators are available, in addition to **is**, **is not**, **is empty**, and **is not empty**:

- **greater than**
- **less than**

When you select a Date or Date/Time field with an order-preserving encrypted field configuration as the left operand in a filter, the following operators are available, in addition to **is**, **is not**, **is empty**, and **is not empty**:

- **after**
- **before**
- **after or on**
- **before or on**

Date and Date/Time pickers

For Date fields, use the date picker to specify the date:

For Date/Time fields, use the date and time picker to specify the date and time:

List condition filters

The **Show Matching** and **Filter Out** options are supported in lists. Only exact matches are returned or filtered out.

Note: Adding encrypted fields in condition filters is supported in scripts such as UI policies and business rules.

Configuration restrictions

Restrictions and behavior of encryption configurations:

- After you add a field to the Edge Encryption Configuration table, you can't delete the configuration record. If you no longer want a field to be encrypted, deactivate the record in the Edge Encryption Configuration table and schedule an encryption job to decrypt the data.
- If a field in a parent table is marked to be encrypted, the field is also encrypted in all inherited tables. For example, if the **Short description** field in the Task table is encrypted, then the contents of the **Short description** field in the Incident table are encrypted.
- If a field inherited from a parent table is marked to be encrypted, the field in the parent table can't be encrypted. For example, if the **Short description** in the Incident table is marked to be encrypted, then the **Short description** in the Task table can't be encrypted. In this example, you can encrypt the **Short description** in the Problem table.
- When a field with an encryption configuration defined is exported to any format, the output includes encrypted values even when exported through the proxy server.
- You can't import data to a field with an encryption configuration defined.
- You can't encrypt inherited Date and Date/Time fields. Date or Date/Time fields inherited from a parent table aren't listed on the **Column** field drop-down list, and you can't create Date or Date/Time encryption configurations for those fields.
- You can encrypt a String or URL field only from a parent table or a child table, but not both.

Instance restrictions

Impact of using Edge Encryption on the instance:

- Back-end logic can't process encrypted data. When the instance contains encrypted data, any business rule, back-end script, or back-end feature that relies on evaluating the data in the encrypted field doesn't run correctly.

  Note: Data encrypted with equality-preserving or order-preserving encryption still passes equivalence checks when compared against an identical encrypted value.

- Since email processing goes from the mail systems straight to the instance and can't pass through the Edge proxy, data sent in or out via email can't be encrypted or decrypted by the Edge proxy.

  - Data and attachments in inbound emails aren't encrypted.
  - Data and attachments in outbound emails remain encrypted and can't be decrypted.
- Scripts run on the server can't change encrypted data.
- Global search isn't supported. Because global search attempts to search both encrypted and clear text data, the results may not be as expected.
- Encrypted data can't be copied and pasted into a record where the field isn't encrypted.
- Depending on the type of encryption selected, the user interface functionality for the encrypted fields is reduced. For example, being able to compare, group by, sort, and search may be impacted. Generally, the stronger the encryption selected, the more functionality is reduced.
- Except for Java KeyStore, SafeNet, and Unbound Technology, no third-party software or hardware encryption key management is supported.
- Although multiple proxy servers connected to a single instance are supported, encryption proxy cluster management and monitoring aren't available. Each proxy must be managed separately.
- System configurations such as workload and the number of encrypted fields can impact the performance of encrypted fields.
- The Edge Encryption proxy server can only connect to a single instance.
- If your instance uses an Oracle database and the String field you're marking to be encrypted is greater than 2925 characters, that field can't be sorted even when order preserving encryption is selected.
- If your instance uses an Oracle database, Unicode AL32UTF8 is the only supported character set.
- Encrypted data can't be used in reports.
- Edge Encryption can't be used with Data Archiving.
- Edge Encryption proxies cannot encrypt requests that use the batch REST request API. If you are using Edge Encryption proxies, disable REST batching by setting the `glide.uxf.disable_rest_batching` system property to true.

Installing Edge Encryption

You can install an Edge Encryption proxy manually or using the Edge Encryption interactive installer.

Java requirements

The host machine installing or running the Edge Encryption proxy server must maintain a supported version of Java. Current supported versions are Java 17.0.3 or later in the 17.x version series.

> **Note:** Java 11 is no longer be supported as of the Yokohama release. Upgrade your environment with the Edge Encryption proxy to Java 17 before you attempt to install Yokohama or later versions of the Edge Encryption proxy.

Installing the proxy server

Installing Edge Encryption includes these steps.

- Install the Edge Encryption proxy application on a server in your network using the interactive installer or the manual installer.
- Generate the RSA key pair for digitally signing encryption configurations and encryption rules.
- Install the Java Cryptography Extension (JCE), if you plan to use AES 256 encryption.
- If you are using a secure SSL connection, obtain a server certificate and import it to the Java KeyStore.
- Set up your keystore and encryption key.
- If order preserving encryption types or encryption patterns are to be used, set up a MySQL database on a machine in your network.
- Set the desired properties. Properties are located in the `edgeencryption.properties` configuration file.
- Specify that a proxy server is a trusted source so that Edge Encryption can process requests coming from that proxy server.

Accessing the proxy server

Once installation is complete, point each user's browser to an Edge Encryption proxy using the URL format: `<host>:<port>`. Values are determined by the host and port properties in the `edgeencryption.properties` file.

As an example with the following values:

Property	Example value
edgeencryption.proxy.host	hostname.mycompany.com
edgeencryption.proxy.http.port	8081

A client will access the proxy server using the following address: `http://hostname.mycompany.com:8081/`.

Note: DNS settings and routing rules may be used. Host and port values are determined by your network administrator.

Request Edge Encryption

The Edge Encryption plugin (com.glide.edgeencryption) is available as a separate subscription.

Before you begin

To purchase a subscription, contact your ServiceNow account manager. The account manager can arrange to have the plugin activated on your organization's production and subproduction instances, generally within a few days.

If you don't have an account manager, decide to delay activation after purchase, or want to evaluate the product on a subproduction instance without charge, follow these steps.

Role required: none

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Select **Request plugin** to open the **Activate Plugin** form on Now Support.
3. On the **Activate Plugin** form, provide the following information.

Activate Plugin form

Field	Description
What is your target instance	Select the instance that you want to activate the plugin on.
Which plugin would you like to activate	Select the name of the plugin to activate. **Note:** If the system doesn't list the plugin you want or if you're activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	Select the date and time to activate the plugin.

Example

For example, see the following form to activate the Event Management plugin on an instance named SNC Instance.

4. Select **Submit**.
   After the maintenance window, the system installs the plugin on your instance. To confirm the installation, go to the Installed tab in the Application Manager.

Set up an Edge Encryption user account

The Edge Encryption proxies connect to the instance as a user to obtain and update encryption configuration information. Create a user account for this purpose and give the edge_encryption role to the user.

Before you begin

The Edge Encryption plugin must be installed before you can assign the role.

Role required: admin

Procedure

1. On your ServiceNow instance, create a user account to be used by the Edge Encryption proxy applications.
2. Assign the edge_encryption role to the user.

Download the Edge Encryption proxy server

Download the Edge Encryption proxy server application from your instance, and then copy the file to each computer that is to run the Edge Encryption proxy server.

Before you begin

Before starting this procedure, the Edge Encryption plugin must be installed and activated on your instance.

Note: The Edge Encryption proxy is officially supported only on Oracle JRE.

Role required: security_admin

About this task

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Installation & Downloads** > **Downloads**.
2. To use the interactive installer, click **Download Interactive Installer**.

   If manually installing the proxy server, select the OS version for your proxy server.

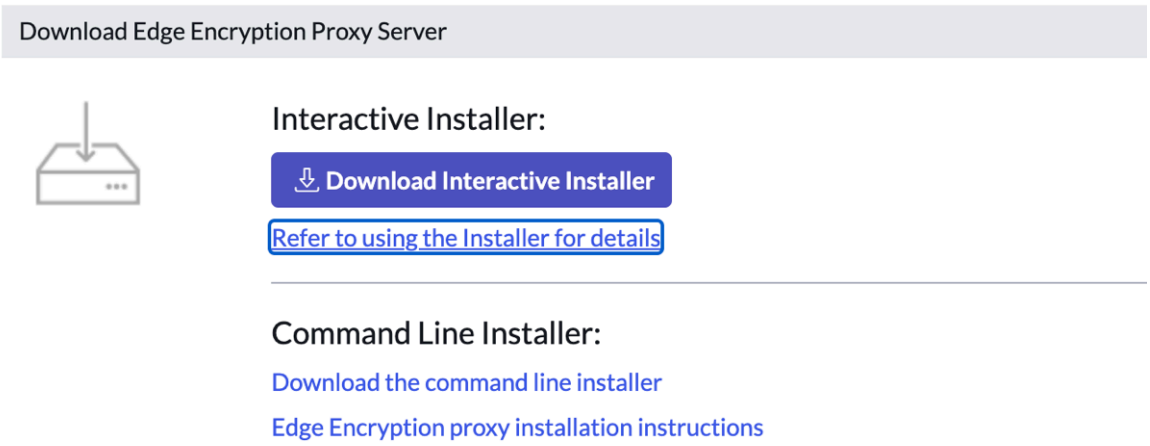

Note: Because the proxy server requires at least 4 GB of memory to run, 32-bit JREs and 32-bit operating systems are no longer supported starting with the Washington DC release.

3. Copy the installer to each computer that is to run the Edge Encryption proxy server.

Note: If you are manually installing the Edge Encryption proxy server, copy the ZIP file to each computer that is to run the Edge Encryption proxy server.

What to do next

After downloading the Edge Encryption installer, Install the Edge Encryption proxy server using the interactive installer. If installing manually, Install the Edge Encryption proxy server using the command line installer.

Install the Edge Encryption proxy server using the interactive installer

Install the Edge Encryption proxy server on a Windows or Linux computer using the interactive installer.

Before you begin

Note: SafeNet KeySecure keystore files are not supported by the Edge Encryption installer. To use a SafeNet KeySecure keystore, Install the Edge Encryption proxy server using the command line installer.

The Edge Encryption plugin must be installed and activated on your instance before you start this procedure. Ensure that Java version 11.0.6 or later is installed on the machine running the Edge Encryption installer.

Role required:

- security_admin on your ServiceNow instance
- local or domain administrator on a Windows host
- service user with full file system access on a Linux host

About this task

After installing a new proxy server, you can run the installer again to perform tests to detect issues with an installation or modify current settings. Your options include:

- **Install New**: Install a new proxy server.
- **Verify Installation**: Perform tests to detect and fix issues in a previous installation.
- **Reinstall Existing**: Perform tests to detect and fix issues in a previous installation and view or modify existing settings.

Note: If installing the proxy server on a Linux machine on a privileged port (port 80 or 443), you must run the installer as a root user with full file system access. To restrict file system access after the proxy server is installed, you can use the SetUID feature in the proxy installer. To enable this feature, start the installer as root or sudo root. When prompted by the installer, provide the username and usergroup of an unprivileged user. The proxy server will install with file system privileges of the given user. You can skip this step to continue the default installation with root privileges.

Procedure

Use the installer to install multiple proxies for your instance on multiple machines, ensuring that the following criteria apply:

- All proxies must have the same encryption keys and the same RSA key pair used to digitally sign encryption configurations and rules.
- The encryption key must be the default key configured on the instance.
- When a proxy database is set up as part of the installation, all proxies must use the same proxy database.

  You may need a proxy database for equality-preserving encryption, order-preserving encryption, or tokenization. If you do not use any of these features, you do not need a proxy database.

What to do next

To use NVDA, an Assistive Technology screen reader designed to read accessibility-enabled Java applications built for keyboard users, see Configure a Windows 64-bit host to use 32-bit NVDA with Java applications.

After installing the Edge Encryption proxy server, Set the proxy server initial memory limit and upper bound memory limit.

Install the Edge Encryption proxy server (interactive installer)

Install the Edge Encryption proxy on a Windows or Linux computer.

Before you begin

Role required: admin

Procedure

1. Download the Edge Encryption proxy server installer.
2. Open the Edge Encryption proxy installer.

   Note: If installing on a Windows machine, you must run the installer as Administrator.

   a. To run the installer as Administrator on a Windows machine, right-click the Command Prompt and select **Run as administrator**.

   b. From the command line, navigate to the directory that contains the downloaded .jar file.

   c. Run the following command: `java -jar <file name>.jar`.

ServiceNow Edge Encryption proxy installer

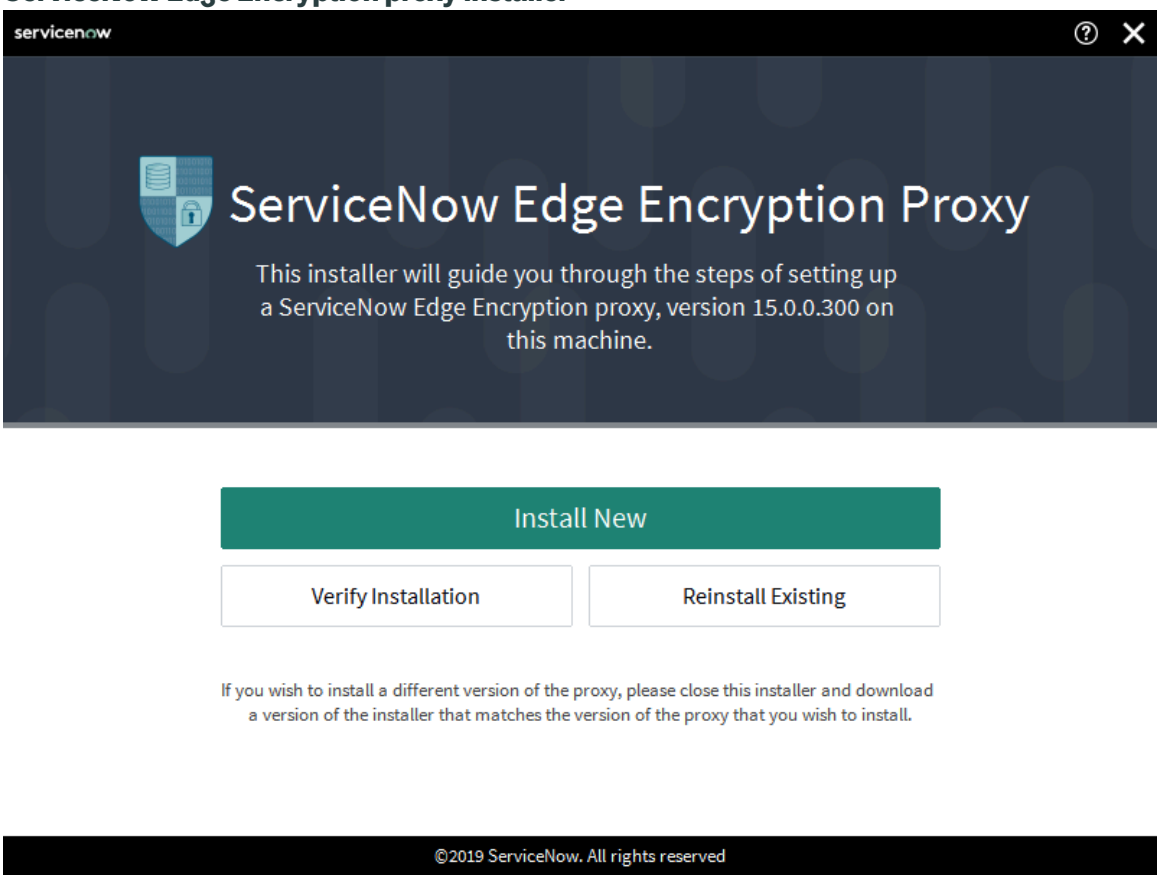

3. To install a new proxy server, select **Install New**.
   If a proxy is already installed, you can run the installer to:
   - **Verify Installation**: Perform tests to detect and fix issues in a previous installation.
   - **Reinstall Existing**: Perform tests to detect and fix issues in a previous installation and view or modify existing settings.
4. Configure the **Installation Location** and **Target ServiceNow Instance**.

   a. Click **Browse** to select an installation location or manually enter an installation path.

   b. Enter the URL of the target ServiceNow instance.
   Include the protocol and port number.

Example
`https://example.servicenow.com:443`

c. Enter the user name and password for a user with the edge_encryption role on the target ServiceNow instance.

5. Click **Next**.
6. Configure the **Connection Settings** and **Proxy Settings**.

Setting	Description
Proxy Host	Fully qualified domain name of the machine on which you are installing the proxy server. **Note:** Click **Detect FQDN** to look up the machine's fully qualified domain name and auto-populate the **Proxy Host** field. Along with the port, this property defines the URL used by the client to access the proxy server.
HTTP Port	Port on the proxy for HTTP communication.
HTTPS Port	Port on the proxy for HTTPS communication.
Proxy Name	Name of the proxy and the service. The proxy name must be unique.
Proxy Poll Interval	Poll interval in seconds. With the default setting, it takes 5 seconds for the proxy to learn of encryption configuration changes. Larger values cause the instance to take longer to detect proxies that have come online. **Note:** Changing the default setting of the Proxy Poll Interval can result in detection delays when a proxy comes online.
Proxy Keep-Alive Ping Interval	Time in seconds between pings issued by the proxy to the instance. Pings are issued periodically to verify connectivity between the proxy and the instance. The default value is 10. The minimum value is 5.

7. Click **Install**.
The Edge Encryption proxy server installs. The installation may take a few minutes.

Configure CyberArk properties protection

Optionally, configure CyberArk properties protection to securely store Edge Encryption passwords in a centralized and secure digital vault.

Before you begin
Role required: admin

About this task
You must purchase and configure CyberArk AIM (Application Identity Management) before you can configure CyberArk connection parameters and protected credentials for a proxy server. As part of the installation of the AIM client, the `JavaPasswordSDK.jar` file is installed in the AIM client installation directory. The CyberArk vault is installed on an independent hardened server, and the AIM clients allow secure access to that server.

Note: You must install the CyberArk AIM client on every host computer where an Edge proxy is installed.

In the Edge installer, you must specify the location of the `JavaPasswordSDK.jar` file to set up the CyberArk connection to the Edge proxy. You must also enter other values you defined during the AIM client installation.

Setting up CyberArk password storage is optional. If you do not want to set up CyberArk password storage, click **Skip** through the CyberArk screens.

Procedure
1. On the CyberArk Connection page of the Edge Encryption installer, enter the CyberArk connection parameters.

CyberArk connection parameters

Setting	Description
Path to PasswordSDK.jar	The path to the `JavaPasswordSDK.jar` file installed on the host Windows machine during CyberArk configuration.
App ID	The **App ID** entered during CyberArk configuration.
Safe Name	The **Safe Name** entered during CyberArk configuration.

2. Click **Next**.
3. On the CyberArk Protected Credentials page of the installer, enter the credentials to be protected by CyberArk.
   - To use a single credential name for all protected passwords, select the **Apply one Credential Name to all Credentials** check box, enter the credential name, and click **Apply**.
   - Enter the credential name for one or more of the following fields. Credential names are the usernames entered for the SSH keys during CyberArk configuration.

CyberArk protected credentials

Setting	Description
Edge Encryption User	The CyberArk credential name for an Edge Encryption user.
Signature Key Keystore	The CyberArk credential name for the signature key keystore.

Setting	Description
HTTPS Cert Keystore	The CyberArk credential name for the HTTPS certification keystore.
Encryption Key Keystore	The CyberArk credential name for the encryption key keystore.
Database	The CyberArk credential name for the database keystore.
SafeNet HTTPS Cert Keystore	The CyberArk credential name for the SafeNet HTTPS certification keystore.
SafeNet Server	The CyberArk credential name for the SafeNet server.
Forward Proxy	The CyberArk credential name for the forward proxy.

4. Click **Next**.

Configure the signature key

Configure the signature key after installing the proxy server through the Edge Encryption proxy installer.

Before you begin
Role required: admin

About this task
The signature key signs changes to configurations and properties made by the proxy server. The signature key must be an asymmetric RSA key pair in a JCEKS KeyStore.

Note: If installing multiple proxies, each proxy must use the same signature key.

Procedure
1. On the Signature Key page of the Edge Encryption installer, select the keystore on the host machine to store the signature key.
   - **Create New Java KeyStore**: Enter the directory location, name, and password for the new keystore.
   - **Use Existing Keystore**: Enter the keystore file location and password.
2. Click **Next**.
3. Select or create a signature key.
   - **New Key**: Create a signature key for this proxy.
   - **Use Existing Key**: Use an RSA key-pair from the selected keystore.
   - **Import Existing Key**: Import an RSA key-pair from a different keystore. Browse to the keystore file, enter the password for the keystore, and select the key alias. Provide a new alias for the key.
4. Click **Next**.

Configure the HTTPS certificate

To enable clients to connect to the Edge Encryption proxy server using a secure SSL connection, import the HTTPS certificate to the proxy server.

Before you begin
Role required: admin

About this task
The Edge Encryption proxy provides the HTTPS certificate to clients trying to connect.

Procedure
1. On the HTTPS Certificate page of the Edge Encryption installer, select the keystore to store the certificate.
   - **Create New Java KeyStore**: Enter the directory location, name, and password for the new keystore.
   - **Use Existing Keystore**: Enter the keystore file location and password.
2. Click **Next**.
3. Select or import a certificate.
   The key alias is the given alias for the certificate.
   - **Use Existing Certificate**: Use an existing certificate in the selected keystore.
   - **Import from File or KeyStore**: Import a certificate from a different keystore or a .cer file. Browse to the keystore or .cer file, enter the password, and select the alias. You must provide a new alias for the certificate.
4. Click **Next**.

Configure the AES 128-bit encryption key

After you configure the HTTPS certificate through the Edge Encryption proxy installer, configure the AES 128-bit encryption key to encrypt your data.

Before you begin
Role required: admin

About this task
The encryption key is either a plain text file inside the `/keys` directory or a secret key inside a keystore. If you use a keystore for your AES 128-bit and AES 256-bit encryption keys, they must both use the same keystore.

If you are updating an SSL certificate on an Edge proxy server, see Update SSL certificate.

Procedure
1. Select the encryption key location.

Option	Description
File Store	Use a file to store a single encryption key. You can use an existing file in the `/keys` directory, or you can generate a new file. To generate a new file, enter an alias and click **Generate**. A file containing an encryption key is created. **Note:** This choice designates both the storage location and the encryption key. If you select **File Store**, click **Next** and go to step 5.

Option	Description
Create New Java KeyStore	Create a keystore to store the encryption key.
Java KeyStore File	Store the encryption key in an existing Java KeyStore file.

2. Click **Next**.
3. Select or create the encryption key.

Option	Description
New Key	Create an encryption key and alias. **Note:** You must use lowercase letters and numbers for the alias name (key name, key alias), per Java KeyStore re quirements. To find out more about the keytool utility, see the Java SE Documen
Use Existing Key	Use an existing encryption key in the selected keystore.
Import Existing Key	Import an encryption key from a different key store.

4. Click **Next**.
5. Configure the key on the instance according to the requirements defined in your installer. To configure the key on the instance, navigate to the instance and define a default key. See Configure encryption keys on the instance. Ensure that the key alias, size, and type match the requirements defined in the installer.
6. Once the key is configured on the instance, return to the installer and click **Next**.

Configure the AES 256-bit encryption key

After you configure the AES 128-bit key through the Edge proxy installer, you can optionally configure an AES 256-bit encryption key to encrypt your data.

Before you begin
Role required: admin

About this task
The encryption key is either a plain text file inside the `/keys` directory or a secret key inside a keystore. If you use a keystore for your AES 128-bit and AES 256-bit encryption keys, both keys must use the same keystore. If you don't want to configure an AES 256-bit encryption key, select **Skip** to continue installing the proxy server.

If you're updating an SSL certificate on an Edge proxy server, see Update SSL certificate.

Procedure

1. Select the encryption key location.

Option	Description
File Store	Use a file to store a single encryption key. You can use an existing file in the /keys directory, or you can generate a new file. To generate a new file, enter an alias and select **Generate**. A file containing an encryption key is created. **Note:** This choice designates both the storage location and the encryption key. If you select **File Store**, select **Next** and go to step 5.
Create New Java KeyStore	Create a keystore to store the encryption key.
Java KeyStore File	Store the encryption key in an existing Java KeyStore file.

2. Select **Next**.
3. Select or create the encryption key.

Option	Description
New Key	Create an encryption key and alias. **Note:** You must use lowercase letters and numbers for the alias name (key name, key alias), according to Java Key Store requirements. To find out more about the keytool utility, see the Java SE
Use Existing Key	Use an existing encryption key in the selected keystore.
Import Existing Key	Import an encryption key from a different key store.

4. Select **Next**.
5. **Optional:** If you want to use AES 256-bit encryption, see Configure the AES 256-bit encryption key.
6. To use AES 256-bit encryption, you must also configure the AES 256-bit default encryption key on the instance.

   Do this by navigating to the instance and defining a default key. See Configure encryption keys on the instance. Ensure that the key alias, size, and type match the requirements defined in the installer.

7. After the key is configured on the instance, return to the installer and select **Next**.

Update SSL certificate

When updating an SSL certificate on an Edge proxy server, you must delete the old one.

Before you begin
Role required: admin

About this task
When updating the SSL certificate on the Edge proxy server, you must also delete the old certificate. If you don't, the old certificate (in the form of an alias in the KeyStore file) continues to be used even though the Edge proxy server is configured to use the new certificate.

Procedure

1. On the Edge proxy server, list the entries in the Java KeyStore:

```
keytool -list -keystore keystore.jceks -storetype jceks
 -storepass MY_SUPER_PASSWORD
```

2. Remove the old SSL certificate:

```
keytool -delete -alias MY_OLD_ALIAS -keystore keystore.jceks
 -storetype jceks -storepass MY_SUPER_PASSWORD
```

3. Add the new SSL certificate into the Java KeyStore.

Configure the Edge Encryption proxy database

If using order-preserving encryption types or encryption patterns, you can optionally configure the Edge Encryption proxy database properties.

Before you begin
Role required: admin

About this task
To use order-preserving encryption types or encryption patterns, a MySQL database running in your network is mandatory. This task connects the proxy to the database, but it does not install or configure the database.

Note: If using multiple proxy servers, all proxy servers must use the same proxy database. The values entered in the installer must be the same for all proxy servers.

Procedure

1. Confirm or change the database URL, which is the location of the proxy database.
2. In the **Name** field, enter the name of the proxy database.
   The default value is *edgeencryption*.
3. Enter the username and password for accessing the proxy database.
4. Click **Next**.

Launch the Edge Encryption proxy server

After an Edge Encryption proxy is installed and configured, you can start the proxy from the installer.

Before you begin
Role required: admin

Procedure

1. After configuring keys on the instance and configuring the proxy database, return to the Edge Encryption proxy installer and click **Launch**.
2. If an issue is detected, or to check the status of your proxy server, you can click **Check Status** to verify that the proxy is running.
   A message displays the proxy

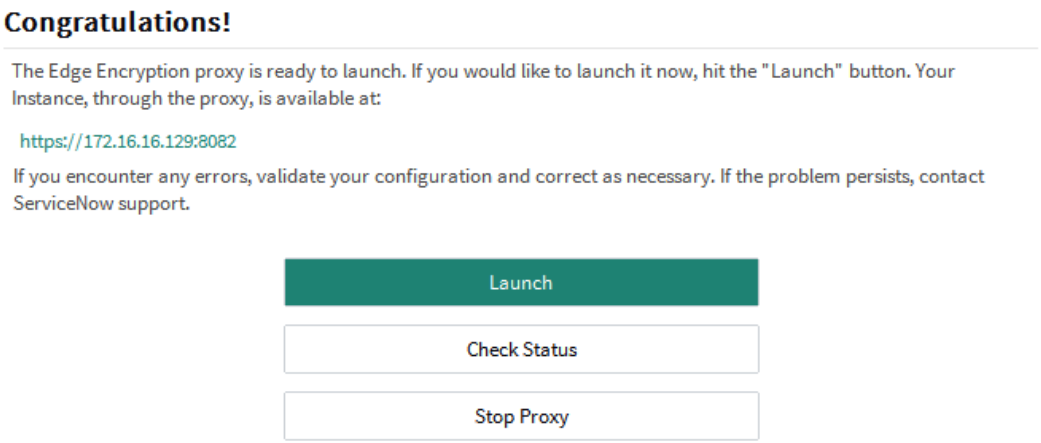

   status.

What to do next

After successfully installing the Edge Encryption proxy server, Set the proxy server initial memory limit and upper bound memory limit.

Verify and troubleshoot the Edge Encryption proxy server installation

After your Edge Encryption proxy is installed, you can verify the installation to locate problems or start and stop the proxy.

Before you begin

Role required: admin

Procedure

1. Open the Edge Encryption proxy installer.
2. Select **Verify Installation**.
3. Click **Proxy Directory** and select the proxy directory.
4. Click **Run Tests**.

Test results

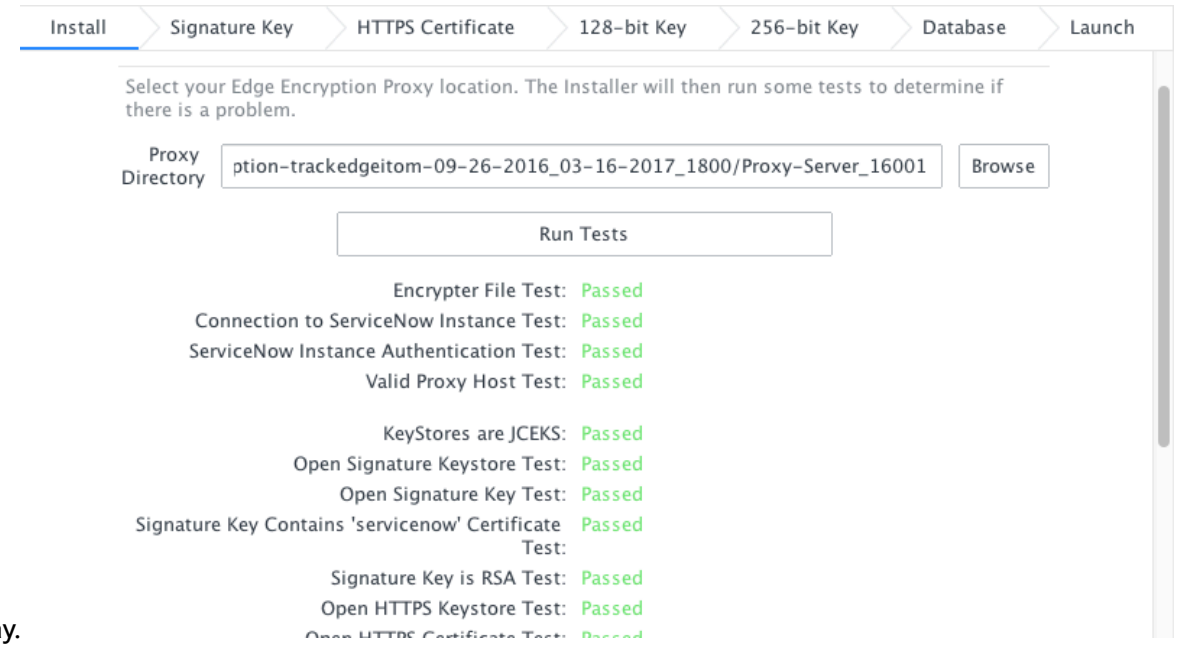

display.

5. Select **Next**.
   If an issue is encountered, you can move through the installer to correct the configuration. If no issues are encountered, the installer jumps to the **Launch** page. You can check the proxy status, stop the proxy, or start the proxy from the **Launch** page.

Install the Edge Encryption proxy server using the command line installer

Manually install multiple Edge Encryption proxy servers in your network.

Before you begin

Roles required: security_admin on your ServiceNow instance and local administrator on the host machine.

If order preserving encryption types or encryption patterns are to be used, set up a MySQL database on a machine in your network if not already present.

Note: If using Unbound Technology encryption keys with Edge Encryption, install the proxy server using the command line installer on the Unbound client machine. The Edge Encryption proxy server must run on the same machine as the Unbound technology client.

About this task

First, set up a single Edge Encryption proxy server. After your first proxy server is successfully running, add additional proxy servers for one instance to ensure an optimal environment. See Sizing your Edge Encryption environment to determine the number of additional proxy servers needed.

Install the Edge Encryption proxy server (command line installer)

Install an Edge Encryption proxy on a 64-bit Windows or Linux computer.

Before you begin

Role required: admin

Java version 11.0.6 or later in the 11.x version series is required to run the installer.

About this task

Install the Edge Encryption proxy server on a machine in your network using the appropriate command for your target machine. If installing the Edge Encryption proxy server on a Windows machine, you must additionally install the proxy server as a Windows service.

When you upgrade the Edge Encryption proxy server, the system backs up the old proxy in the `backup.dist-upgrade-<timestamp>` directory under the current installation directory. The backup directory is generated during the upgrade process and stores the old proxy information.

When you run an upgrade via the command line, a dist-upgrade.log may be generated in the directory where the command runs. The dist-upgrade.log contains logs for the upgrade process.

In case of a failed upgrade, the system creates a `failed-backup.dist-upgrade-<timestamp>` directory. In addition, `logs/wrapper.log` in the original proxy directory may also contain failure information.

Procedure

1. Create the installation directory.
2. Download the Edge Encryption proxy archive file to the installation directory.
3. Open the terminal and change to the installation directory.

   Note: If installing on a Windows machine, you must start the Windows Command Prompt with administrator privileges.

4. Run this command for the target machine and change the variables according to your configuration: `java -jar edgeencryption-<version>-all.jar -m install -n <proxy_name> --instancehost <host> -p <port> --protocol https -s <install_path>`

Option	Variable	Description
none	version	Version number of the Edge Encryption proxy being used to perform the current operation.
-m	mode	Runtime mode. Options are "install" for a new Edge proxy server or “dist-upgrade” to upgrade an existing Edge proxy server.
-n	proxy_name	Name of the installed Edge Encryption proxy server. Use a unique proxy_name to be able to identify specific proxy instances.
--instancehost	host	The host name of your ServiceNow instance (for example, `mycompany.service-now.com`).
-p	port	Port your ServiceNow instance listens on. Typically

Option	Variable	Description
		secure HTTPS connections listen on port **443** and HTTP connections listen on port **80**.
--protocol	protocol	Protocol the installed Edge proxy uses when connecting to the backend ServiceNow host. This is typically **HTTPS** (preferred for secure TLS connections) or **HTTP** (connections without TLS) depending on which protocol the host instance supports.
-s	install_path	Path to the directory or folder where the new Edge proxy is installed (the destination directory). If the directory does not already exist, this command will create it. If it does exist it must not contain an existing installation. If this option is skipped, the default folder name is derived from the proxy_name and port (for example, `EdgeProxy_443`), in the current directory.

Note: Do not copy and paste commands from the browser. Occasionally, copy/paste operations cause unexpected characters to be pasted to the target machine and results in the command being executed incorrectly. It is best to type out the command by hand using documentation as a reference.

To see the help screen, execute this command with the `–help` option: `java -jar edgeencryption-<version>-all.jar --help`

5. If installing on a Windows machine, install the Edge Encryption proxy as a Windows service.

   a. **Optional:** Change the name of the service by opening the `conf/wrapper.conf` file on the new proxy and setting the properties in the following table.

Property	Description
wrapper.ntservice.name	Unique name of the Edge Encryption proxy service.
wrapper.ntservice.displayname	Edge Encryption proxy service display name.
wrapper.ntservice.description (Optional)	Proxy server description.

If this step is not performed, the Edge Encryption proxy service is installed under the name **Edge Encryption**.

b. Save and close the file.

c. Open the Windows Command Prompt and `cd` to `ServerName_port/bin`.

d. Execute `edgeencryption.bat install`.

Result

The `ProxyName_port` directory is created in the current directory. The `edgeencryption.properties` file is updated with the host, port, and protocol values from the command line.

Create and configure the RSA key pair for the digital signature

Create an RSA key pair that the proxy server can use to create the digital signature for signing changes to the encryption properties and configuration.

Before you begin

Role required: admin

To generate and validate the digital signature, an RSA key pair must be generated and stored in the JCEKS Java KeyStore and each proxy must be configured to use this key pair. Generate an encryption key pair using the keytool command.

To use the keytool utility with a proxy installed on SElinux (CentOS), you must enable loading of shared libraries from the proxy java-installation directory. To do this, run the following command as root.

```
chcon -R -t texrel_shlib_t proxy_install_dir/java/jre /lib
```

You must use the Java 1.8 version of the keytool utility. A copy of the utility can be found in `<proxy install dir>java/jre/bin/keytool`.

Procedure

1. Change to the KeyStore directory in the proxy download directory.
2. Change the default password.

   The default password is *changeme*.

   ```
   keytool -keystore keystore.jceks -storetype jceks -storepasswd
    -new <newpassword>
   ```

3. Create an encryption key pair.

   Note: Do not enter a password for the key when the keytool utility prompts for one.

   Enter this command on a single line.

   ```
   keytool -genkeypair -alias <key alias> -keyalg rsa -keystore
    keystore.jceks
   -storetype jceks -storepass <keystore password> -keysize 2048
   ```

4. Update the encryption proxy property file (`edgeencryption.properties`).

a. Change to the `<installation directory>/conf/` directory.

b. Open the `edgeencryption.properties` file.

c. Enter the properties for the digital signature.

These properties must be the same for all proxies.

5. Save and close the `edgeencryption.properties` file.

Import and configure the certificate for secure SSL connection

To use a secure SSL connection, import a server certificate and add it to the Java KeyStore.

Before you begin

Role required: admin

You must obtain the server certificate and matching private key before adding it to the Java KeyStore.

Procedure

1. Generate a Certificate Signing Request (CSR) using the `openssl` command.

```
openssl req -newkey rsa:2048 -keyout PRIVATEKEY.key -out
 MYCSR.csr
```

2. Send your CSR (MYCSR.csr in the example above) to your certificate authority to have it signed.
3. Create a P12 keystore for import using the `openssl` command.

```
openssl pkcs12 -export -in MYSIGNEDCERT.pem -inkey
 PRIVATEKEY.key -name shared > MY_SERVER.p12
```

4. Store your certificate and private key into a jceks file.

```
keytool -importkeystore -destkeystore keystore.jceks
 -deststoretype jceks -srckeystore MY_SERVER.p12 -srcstoretype
 pkcs12 -alias MYALIAS
```

The alias, shown in the example as `MYALIAS` can be any value. You will use this alias in the `edgeencryption.proxy.https.cert.alias` property in the `edgeencryption.properties` file located in the `<installation directory>/conf/` folder.

5. Stop and restart the edge proxy.

Note: During a restart, the proxy server is offline for a short time. The amount of time is determined by your environment and how long it takes to stop and restart the proxy service.

Set up a keystore and encryption keys

Set up the keystore and encryption keys used by the Edge Encryption proxy server.

Before you begin

Role required: security_admin

Procedure

1. Carefully determine the appropriate type of keystore to use based on your organization's needs.

Supported keystore	Description
File store	Keys are stored in a file in a file system accessed by the Edge Encryption proxy server. Because encryption keys stored in a file are not encrypted, it is your responsibility to protect these files.
Java KeyStore	A Java KeyStore: ○ Stores keys in a Java JCEKS KeyStore. ○ Is password protected and more secure than storing keys in a file in the file system. ○ Can store multiple keys. A key alias represents each key, making it easier to manage multiple keys. The Edge Encryption proxy ships with the Java JCEKS KeyStore file named `keystore.jceks` in the `keystore` directory. This keystore file contains the ServiceNow public key used to validate encryption rules signed by ServiceNow.
Enterprise Key Management (EKM)	**SafeNet KeySecure** Keys are stored and retrieved with SafeNet KeySecure key management. You must secure a license with Gemalto, download the libraries, and install the SafeNet KeySecure keystore on a host machine in your network before configuring the keystore on the Edge Encryption proxy server. **Unbound Technology** The base64-encoded wrapped encryption key is stored as text file on the Edge Encryption proxy server. The Unbound Technology implementation (previously Dyadic Security) maintains control of the wrapping key. The Edge Encryption proxy server must run on the same machine as the Unbound technology client.

Note: If using a keystore other than the base system Java JCEKS KeyStore, you must import the ServiceNow public key into your keystore. The public key alias is *`servicenow`*.

2. Set up the keystore and encryption keys in your local network.

Set up a Java KeyStore keystore

You can use a Java KeyStore keystore to store encryption keys.

Before you begin

Role required: admin

You must use the Java 1.8 version of the keytool utility. A copy of the utility can be found in `<proxy install dir>/java/jre/bin/keytool`.

About this task

The Edge Encryption proxy ships with the Java JCEKS KeyStore file named `keystore.jceks` in the `keystore` directory. This keystore file contains the ServiceNow public key used to validate encryption rules signed by ServiceNow.

Procedure

1. Set up the keystore properties.
   a. Change to the `<installation directory>/conf/` directory.
   b. Open the `edgeencryption.properties` file.
   c. Enter the properties for the Java KeyStore.
2. Save and close the `edgeencryption.properties` file.

What to do next

After setting up the Java KeyStore, Create encryption keys using the Java KeyStore keytool.

Create encryption keys using the Java KeyStore keytool

You can use the keytool shipped with the encryption proxy distribution to create AES 128-bit and AES 256-bit encryption keys.

Before you begin

Role required: admin

You must use the Java 1.8 version of the keytool utility. A copy of the utility can be found in `<proxy install dir>/java/jre/bin/keytool`.

To find out more about the keytool utility, see the Java SE Documentation.

About this task

Note: The Java KeyStore requires that the alias name (key name, key alias) use lowercase letters and numbers.

Procedure

1. Change to the keystore directory, `<installation directory>/keystore/`.
2. To create the encryption key, run one of the following commands.

Note: If you choose to run these commands from a directory other than the keystore directory, that is you skipped the previous step, you must change the *-keystore* option to include the path from your current directory to the keystore directory. For example, if you were in the `<installation directory>\bin` directory, the option would be `-keystore ../keystore/keystore.jceks`.

Option	Description
AES 128	`keytool -genseckey -alias 128bitkey -keyalg aes -keysize 128 -keystore keystore.jceks -storetype jceks`
AES 256	`keytool -genseckey -alias 256bitkey -keyalg aes -keysize`

Option	Description
	`256 -keystore keystore.jceks -storetype jceks`

You add the alias on the instance when you assign default keys.

Note: The key password must be the same as the keystore password.

Set up a SafeNet KeySecure keystore

If you are using a SafeNet keystore, copy a set of libraries into the proxy distribution directory.

Before you begin

Role required: admin

You must install and set up the SafeNet keystore before performing this step. Secure a license with Thales in order to download the libraries.

Note: For IngrianNAE version 8.12, you must also download the `commons-collections.jar` file

About this task

Note: On Linux, file paths use a forward slash (/).

Procedure

1. Change to the `<installation directory>/conf/` directory, and open the `edgeencryption.properties` file.
2. Enter the properties for the SafeNet keystore.

Note: You may configure SafeNet keystore using with username/password authentication or client certificate authentication, but not a combination of both.

Example

An example for a SafeNet keystore using username and password authentication.

```
edgeencryption.nae.retries = 3
edgeencryption.nae.enabled = true
edgeencryption.nae.server = url
edgeencryption.nae.port = 9000
edgeencryption.nae.protocol = ssl
edgeencryption.nae.keystore.path = keystore/safenet_truststore
edgeencryption.nae.keystore.password = password
edgeencryption.nae.user = safenet_user
edgeencryption.nae.password = safenet_password
```

Example

An example for a SafeNet keystore using client certificate authentication. This authentication method eliminates the need to store the SafeNet server username and password in the properties file.

```
edgeencryption.nae.retries = 3
edgeencryption.nae.enabled = true
edgeencryption.nae.server = url
edgeencryption.nae.port = 9000
```

```
edgeencryption.nae.protocol = ssl
edgeencryption.nae.keystore.path = keystore/safenet_clientcert
edgeencryption.nae.keystore.password = password
edgeencryption.nae.client.certificate = cert_name
```

3. Add or create a key in the SafeNet keystore.
   You add the key name (alias) on the instance when you assign default keys.
4. Save and close the `edgeencryption.properties` file.

Upgrade from Kingston or lower to London or higher

If you use a SafeNet NAE server for key storage with Edge, before upgrading the proxy from Kingston or lower to London or higher, you must copy Gemalto SafeNet client ProtectApp JAR files and add new properties.

Before you begin
Role required: admin

About this task
Note: On Linux, file paths use a forward slash (/).

Procedure
1. Copy the following files from `<installation directory>/lib` to the `<installation directory>/nae directory`:
   - `commons-collections<version>.jar`
   - `ingrianlog4j-api-<version>.jar`
   - `ingrianlog4j-core-<version>.jar`
   - `ingrianNAE-<version>.jar`
2. On the current version (not upgraded) of the proxy, update the `<installation directory>/conf/edgeencryption.properties` file by adding the following two properties:
   - *`edgeencryption.ekm.provider.classname`* = `com.snc.edgeencryption.encryption.CloudEdgeNaeKeyProvider`
   - *`edgeencryption.thirdparty.vendor.library.path`* = `<directory path to the directory where you copied the jar files in step 1>`

     Note: *`edgeencryption.thirdparty.vendor.library.path`* for Java 11.
3. Save the changes.
4. Proceed with the upgrade to London or higher.

Set up Unbound Technology keys

Use Unbound Technology (previously Dyadic Security) keys with Edge Encryption by storing the base64-encoded wrapped encryption key as text file on the Edge Encryption proxy server and providing the wrapping key alias. The Unbound Technology implementation maintains control of the wrapping key.

Before you begin
Role required: security_admin

In your Unbound Technology implementation, identify both the wrapping key and the wrapped key. Use the RSA/ECB/OAEPWITHSHA-256ANDMGF1PADDING algorithm for wrapping and padding. Export the wrapped key in base64-encoded text format. Save the file using the key alias as the name with no file extension.

Note: If using Unbound Technology encryption keys with Edge Encryption, install the proxy server using the command line installer on the Unbound client machine. The Edge Encryption proxy server must run on the same machine as the Unbound technology client.

Procedure

1. Add the wrapped encryption key in base64-encoded text format to the `<proxy-installation-directory>/keys` directory.
   The name of the file must be the key alias with no file extension.
2. Update the `edgeencryption.properties` file.

   a. Change to the `<proxy-installation-directory>/conf` directory.

   b. Open the `edgeencryption.properties` file.

   c. Enter the File store properties and set the value of *`edgeencryption.keyfile.directory`* to `keys`.

   This property directs the proxy server to look for the encryption key in the `<Java-home-directory>/keys` directory.

   For more information on Edge Encryption properties, see Edge Encryption proxy server properties.

   d. Uncomment the properties for the Dyadic provider configuration and set the value of *`edgeencryption.ekm.provider.rsa.wrapping.key.alias`* to the wrapping key alias in your Unbound implementation.

   e. Save and close the file.

What to do next

Add the encryption key alias to the instance. The encryption key alias is the file name of the wrapped encryption key added to the `<proxy-installation-directory>/keys` directory. For example, if the file in the directory is named `myunboundkey`, add this name to the **Key alias** field. See Configure encryption keys on the instance.

Create an encryption key stored in a file

You can use a simple text file as a keystore. Each file holds a single encryption key.

Before you begin

Role required: admin

About this task

This step creates both the key storage and the encryption key.

Note: The name of the key file must match the key alias specified in the encryption keys table in the instance. See Configure encryption keys on the instance.

Procedure

1. Create a file in the `/keys` folder of the proxy server installation directory.
2. Add the encryption key to the file.

Option	Description
AES 128	Place the encryption key, exactly 16 bytes, in to the file.
AES 256	Place the encryption key, exactly 32 bytes, in to the file.

3. Update the `edgeencryption.properties` file.
   a. Change to the `<installation directory>/conf/` directory.
   b. Open the `edgeencryption.properties` file.
   c. Enter the properties for the file store.
   d. Save and close the file.

Configure encryption keys on the instance

Edge Encryption provides the tools to manage encryption keys without taking the proxy offline.

Before you begin

Role required: security_admin

Before setting up new encryption keys on the instance:

1. Create the encryption key.
2. Make the new key available to all encryption proxies. Either copy the file or Java KeyStore file to each proxy, or ensure that each proxy has access to the Java KeyStore or Enterprise Key Management (EKM) device.

About this task

Key aliases must be unique. Each key alias must have the same key size and type on each proxy, or the key cannot be assigned as the default.

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Encryption Key Configuration** > **Set Up Keys**.
2. On the Add New Keys section of the form, complete the following steps to add a new key.

> **Important:** If using SafeNet versioned keys, an additional column appears for the **Key version**. The **Key version** cannot be edited. Click the **Retrieve latest key versions** link in the Related Links to retrieve the latest version of each key from the Edge proxy.

Rows in the list with an **X** in the left column can be deleted. Keys that have been used as the default or that have a **Status** of **Available** cannot be deleted.

   a. Double-click in the row that says **Insert a new row.**.

   b. In the edit box, enter a name for the key, then click the check mark.

Key aliases are lowercase letters and numbers. Capital letters are changed to lowercase letters when you click **Update**. Key aliases must be unique.

> **Note:** If using Unbound technology keys, add the encryption key alias. The encryption key alias is the file name of the wrapped encryption key added to the `<proxy-installation-directory>/keys` directory. For example, if the file in the directory is named `myunboundkey`, add this name to the **Key alias** field.

c. In the same row, double-click in the **Key size** column.

d. In the select box, select a key size, either **128 bits** or **256 bits**, then click the check mark.

e. In the same row, double-click in the **Type** column.

f. In the select box, select a key type, either **File**, **Keystore**, **SafeNet**, or **Unbound**, then click the check mark.

g. When you are done adding keys, click **Next Step**.
You must specify an alias, key size, and key type for each key before moving on.

3. On the Keys Status section of the form, check the **State** of the key and ensure that it is **Available**.
4. When the key is **Available**, click **Next Step**.
This might take a few minutes.

> **Note:** If using SafeNet versioned keys, an additional column appears for the **Key version**. The **Key version** cannot be edited.

The instance tracks the status of every encryption key available to any proxy. When a key is available on all proxies, its state becomes **Available**. If the state does not change after a few minutes, check to ensure that the key is available on all proxies. If the state remains **Unavailable**, one or more of the proxies does not have the key.

Encryption key states

Status	Description
Available	All online proxies have the key.
Unavailable	This is a new key and the proxies have not yet loaded the key, or at least one proxy failed to load the key.

5. On the Change Default Keys section of the form, do one of the following:
   - Type in the key alias.
   - Click the magnifying glass icon and select an alias.

Note: If using SafeNet versioned keys, an additional field appears for the **Key version**. The **Key version** is grayed out and cannot be edited. Choose only the most recent key version. If you choose an earlier version, the following message appears when you click **Update** or **Next Step**.

```
One of the default keys chosen is not the latest version
 available for the key. Please use the latest version.
```

If the default keys are not the latest versions of the SafeNet keys, an **Update default keys to latest version** link appears in the Related Links. Click the link to update the default keys to use the latest version.

6. On the Schedule Key Rotation section of the form, schedule a mass key rotation job or single key rotation job to encrypt existing data using the new encryption key.

   If you do not run a mass key rotation job or single key rotation job, existing data remains encrypted with the old key until the data is accessed again.

Configure additional properties in the Edge Encryption properties file

After installing the Edge Encryption proxy server in your network and setting up your keystore and keys, configure the additional Edge Encryption properties.

Before you begin
Role required: admin

Procedure

1. Open the `<installation directory>/conf/edgeencryption.properties` file and configure the following Edge Encryption proxy server properties:
   - Target (instance) properties
   - User account properties
   - Proxy properties
   - If using order preserving encryption types or encryption patterns, configure the Proxy database properties
   - Clear text and static IV properties
2. Save and close the file.

Configure a web proxy

If your network uses a web proxy, you can set up the Edge Encryption proxy to use the web proxy.

Before you begin
Role required: admin

About this task
If your network does not use a web proxy, leave the web proxy properties in the configuration file commented out.

The Edge Encryption proxy server supports HTTP connection to and basic authentication with the web proxy.

Procedure

1. Change to the `<installation directory>/conf/` directory.
2. Open the `edgeencryption.properties` file.
3. Configure the web proxy properties.
4. Save and close the `edgeencryption.properties` file.
5. If the web proxy is using a customer-specific server certificate, add this certificate to the JVM used by the Edge Encryption proxy server to establish trust between the web proxy and the Edge Encryption proxy server.
   a. Use the `cd` command to navigate to `<Java home directory>/jre/lib/security/cacerts`
   b. Execute the command: `keytool -keystore cacerts -importcert -alias <chooseAlias> -file <certificateFile>`

Set the proxy server initial memory limit and upper bound memory limit

Set the initial memory limit and upper bound memory limit to specify how much memory the proxy server can consume. Set these limits to avoid performance issues in your Edge Encryption implementation.

Before you begin

Role required: admin

About this task

As a guideline, set both the initial memory limit and the upper bound memory limit to the same value. On any machine, allocate 2 GB of the physical memory to the operating system (OS). Then allocate the rest of the physical memory to the heap using the initial memory limit and upper bound memory limit properties. For example, on a machine with 8 GB of memory, allocate 2 GB to the OS, and allocate the remaining 6 GB (6144 m) to the initial and upper bound memory.

Important: If your Edge Encryption proxy server is running, you must stop and restart the proxy server after updating these properties.

Procedure

1. In your proxy server directory, open `<install dir>/conf/wrapper.conf`.
2. To set the initial memory limit, add the following line at the end of the file:

```
wrapper.java.additional.<number>=-Xms<min_memory_in_MB>m
```

   Set `<number>` to the next available `<number>` in the sequence of `wrapper.java.additional.<number>` properties defined in the `wrapper.conf` file.

   Example

   For example, you have the following list of `wrapper.java.additional.<number>` properties:

```
wrapper.java.additional.1=
wrapper.java.additional.2=
```

   The maximum `<number>` in the above list is **2**. When you add the `wrapper.java.additional.<number>=-Xms<min_memory_in_MB>m` line, set `<number>` to **3**, the next available number.

   Important: Do not leave gaps in the numbering sequence.

Set `<min_memory_in_MB>` to the number of megabytes of memory remaining after allocating 2 GB of memory to the OS.

3. Set the upper bound memory limit.
   Because an upper bound memory limit is not set in the base system, the proxy server can use all available memory. If other services are running on the server, you may want to set the upper bound memory limit.

   Add the following line at the end of the file:

   `wrapper.java.additional.<number>=-Xmx<max_memory_in_MB>m`

   Set `<number>` to the next available `<number>` in the sequence of `wrapper.java.additional.<number>` properties defined in the `wrapper.conf` file.

   Example
   For example, you have the following list of `wrapper.java.additional.<number>` properties:

   ```
   wrapper.java.additional.1=
   wrapper.java.additional.2=
   ```

   The maximum `<number>` in the above list is **2**. When you add the `wrapper.java.additional.<number>=-Xmx<max_memory_in_MB>m` line, set `<number>` to **3**, the next available number.

   Note: Do not leave gaps in the numbering sequence.

   Set `<max_memory_in_MB>` to the number of megabytes of memory remaining after allocating 2 GB of memory to the OS.

4. Save and close the file.

Example: Example: Setting proxy server initial and upper bound memory limits

```
wrapper.java.additional.1 = -Djava.io.tmpdir=../tmp
wrapper.java.additional.2 = -Dcloudedge.home.dist=..
# must ensure UTF8 encoding when running on Windows
wrapper.java.additional.3 = -Dfile.encoding=UTF8
# additional properties for heap settings
wrapper.java.additional.4 = -Xms6144m
wrapper.java.additional.5 = -Xmx6144m
```

What to do next

Start the Edge Encryption proxy.

Start the Edge Encryption proxy

After an Edge Encryption proxy is installed and configured, you can start the proxy from the command line.

Before you begin

Role required: admin

Before starting the encryption proxy, verify the following:

- The Edge Encryption plugin is activated on the instance.
- The `edgeencryption.properties` file on this machine has been configured.
- If using an order preserving encryption type or encryption patterns, the proxy database is running.

Note: The first time you set up the `edgeencryption.properties` file or change properties, you may not want to set the password encryption property. After you have verified that everything is working, you can set the password encryption property, shut down the proxy, and then restart the proxy.

Procedure

1. Run the proxy server.

Option	Description
On a Linux machine	**a.** cd to `ServerName_port` **b.** Execute `./startup.sh`
On a Windows machine	Perform the following steps from the com mand line as admin: **a.** cd to `ServerName_port/bin` **b.** Execute `edgeencryption.bat start`

2. Check the log on the proxy server to verify that the proxy is running.

Obfuscate passwords in the properties file

Obfuscate passwords in the `edgeencryption.properties` file to be able to share the properties file without revealing clear text passwords.

Before you begin

Role required: admin

Make sure that the Edge Encryption proxy server is set up and successfully running before you set this property. Before setting this property, Stop the Edge Encryption proxy.

About this task

Setting this property may make it difficult to debug connection and access issues during initial startup. Only set this property in production environments after the proxy has been set up and tested successfully.

Procedure

1. Change to the `<installation directory>/conf/` directory.
2. In the `conf` directory, create a text file containing a complex string or phrase that can be used as a passphrase which the proxy uses to obfuscate the passwords in the `edgeencryption.properties` file.
   This passphrase should be a random and complex phrase not related to the passwords themselves.
3. Open the `edgeencryption.properties` file.

4. Set the password encryption property.
5. Save and close the `edgeencryption.properties` file.

What to do next
After setting this property, you can Start the Edge Encryption proxy.

Manually add an additional proxy

After the first Edge Encryption proxy is properly configured and tested, you can set up additional proxies on a Linux or Windows machine. Installing multiple proxies on the same machine is not recommended.

Before you begin
Role required: admin

About this task
Add additional proxy servers on additional machines to ensure an optimal environment. See Sizing your Edge Encryption environment to determine the number of additional proxies needed.

Note: Make sure that all proxies have the same encryption keys and the same RSA key pair used to digitally sign encryption configuration and encryption rules. If a proxy database was set up as part of the installation, all proxies must use the same proxy database.

Procedure
1. Install the proxy using the appropriate command.
   For more information, see Install the Edge Encryption proxy server (interactive installer).
2. Copy all the encryption keys and the `edgeencryption.properties` file from the first proxy to the new proxy.
   Encryption keys may be located in the proxy keystore, in the `/keys` directory, or in a SafeNet KeySecure keystore.
3. Open the `edgeencryption.properties` file on the new proxy.
4. Change the following properties:

Property	Description
edgeencryption.proxy.name	Unique name of the proxy server
edgeencryption.proxy.host	The server name, IP address, or fully-qualified domain name of the computer running the proxy.
edgeencryption.proxy.http.port	Port on the proxy for HTTP communication. Must be unique across all processes on the machine.
edgeencryption.proxy.https.port	Port on the proxy for HTTPS communication. Must be unique across processes on the machine.

5. If installing the proxy server on a Windows machine, you must change the name of the service. by opening the `conf/wrapper.conf` file on the new proxy and adding the properties listed in the following table.

Note: You must perform this step before launching the proxy server.

Property	Description
wrapper.ntservice.name	Unique name of the Edge Encryption proxy service.
wrapper.ntservice.displayname	Edge Encryption proxy service display name.
wrapper.ntservice.description (Optional)	Proxy server description.

6. Save and close the file.
7. Launch the proxy using the appropriate command.
   For more information, see Start the Edge Encryption proxy.

Authenticate an Edge Encryption proxy server

Specify that a proxy server is a trusted source so that Edge Encryption can process requests coming from that proxy server.

Before you begin

If a proxy server is not authenticated, the console log includes the following message:

```
WARN This Edge Encryption proxy has not yet been authenticated
 by the instance.
Please navigate to the matching Proxy record on your ServiceNow
 instance and authenticate it.
```

If you attempt to access the proxy, you receive the following message: `This site can't be reached`.

To maintain the proxy in an operational state during the upgrade process, authentication is not required until after the proxy update is successful.

Role required: admin

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Proxies**.
2. Select the proxy and click **Authenticate**.

Result

The proxy moves from **Unauthenticated** to **Pending** to **Authenticated**. The status changes from **Unauthenticated** to **Pending** when you start the authentication. When authentication is complete, the status changes from **Pending** to **Authenticated**, and you can access the proxy and Edge Encryption can accept requests from the proxy.

Note: If you stop and restart the proxy, the proxy remains **Authenticated** and restarts successfully.

Stop the Edge Encryption proxy

You can stop an Edge Encryption proxy from the command line.

Before you begin

Role required: admin

Procedure

1. Stop the proxy server.

Option	Description
On a Linux machine	Execute `./shutdown.sh`
On a Windows machine	Execute `edgeencryption.bat stop` To remove the Windows service, execute `edgeencryption.bat remove`

2. Check the log on the proxy server to verify that the proxy has stopped.

Uninstall the Edge Encryption proxy on Linux

You can uninstall the Edge Encryption proxy. If you are upgrading the proxy, it is not necessary to shut down and uninstall the current version.

Before you begin

Role required: admin

You must have access to the computer running the Edge Encryption proxy.

About this task

Before shutting down the Edge Encryption proxy, ensure that no users are connected to the instance using the proxy.

The encryption proxy running on Linux operates as a single process. You can end this process to accommodate such tasks as redeploying the encryption proxy to another host machine, updating the proxy version, updating the Java version, or changing the unique name of the encryption proxy when deploying the encryption proxy on multiple proxy servers.

Procedure

1. You may want to save the `edgeencryption.properties` file before deleting the distribution directory.
2. Execute the `shutdown.sh` shell script.
3. Check the log on the proxy server to verify that the proxy server is shut down.
4. Delete the files in the distribution folder.

Uninstall the Edge Encryption proxy on Windows

You can uninstall the Edge Encryption proxy. If you are upgrading the proxy, it is not necessary to shut down and uninstall the current version.

Before you begin

Role required: admin

You must have access to the computer running the Edge Encryption proxy.

Before shutting down the Edge Encryption proxy, ensure that no users are connected to the instance using the proxy.

Procedure

1. You may want to save the `edgeencryption.properties` file before deleting the distribution directory.
2. Execute `edgeencryption.bat stop`
3. Execute `edgeencryption.bat remove`
4. Check the log on the proxy server to verify that the proxy server is shutdown.
5. Delete the files in the distribution folder.

Set up multiple provider SSO with Edge Encryption

Set up multiple provider SSO to enable logging in through the Edge Encryption proxy server URL or the instance URL. If you are implementing multiple provider single sign-on (SSO) with Edge Encryption enabled, some users might need to log in to your instance through the Edge Encryption proxy server, while other users might not.

Before you begin

- Enable the Integration - Multiple Provider Single Sign-On Installer plugin (com.snc.integration.sso.multi.installer).
- Enable the Edge Encryption plugin (com.glide.edgeencryption) and ensure that one or more proxy servers are set up in your network.
- Determine the URL for the Edge Encryption proxy server that users will log in through using multiple provider SSO. To determine the URL of an Edge Encryption proxy server, see Installing Edge Encryption.

Role required: admin

About this task

The user logging in will need to use the appropriate URL to log in, either using the Edge Proxy or not using the Edge Proxy.

- If routing all users through the Edge Encryption proxy server, set up your identify provider record and define the proxy server URL in the **ServiceNow Homepage**, **Entity ID / Issuer**, and **Audience URI** fields.
- To route some users through the proxy server and some users to the instance, create two identify provider records. Both records use the same value in the **Identity Provider URL** field. However, one of the records routes through the proxy server, while the other routes to the instance.
  - Login via instance name: `https://<instance name>.service-now.com/login_with_sso.do?glide_sso_id=<sys_id` of IdP record for non-Edge Proxy
  - Log in via Edge Proxy: `https://<edge hostname>:<port>/login_with_sso.do?glide_sso_id=<sys_id` of the IdP record for the Edge Proxy

Procedure

1. Enable the duplication of identity provider URLs in identity provider records.
   A unique constraint prevents duplication of the identity provider URL in two different identity provider records. You can enable duplication of the identity provider URL in multiple IdP records by setting a field to false.
   a. Navigate to **System Definition** > **Dictionary**.
   b. Open the definition record for the **idp** field of in the Identity Providers table [saml2_update1_properties].

c. Configure the form to add the **Unique** field.

d. Ensure that the value of the **Unique** field is set to **false**.

2. Navigate to **Multi-Provider SSO** > **Identity Providers**.

3. Create two identity provider records for the same identity provider: one using the instance URL and one using the Edge Encryption proxy server URL.
To create an identity provider record, see Create an external identity provider.

a. For the Edge Encryption proxy server URL, complete the form using these values.

Field	Value
Identity Provider URL	Imported from IdP metadata.
ServiceNow Homepage	The URL for your proxy server homepage. For example: `https://<proxy hostname>:<port>/navpage.do`
Entity ID / Issuer	`https://<proxy hostname>:<port>`
Audience URI	`https://<proxy hostname>:<port>`

b. Click **Submit**.

c. For the instance URL, complete the form using these values.

Field	Value
Identity Provider URL	Imported from IdP metadata.
ServiceNow Homepage	`https://<instance>.service-now.com/navpage.do`
Entity ID / Issuer	`https://<instance>.service-now.com/navpage.do`
Audience URI	`https://<instance>.service-now.com/navpage.do`

d. Click **Submit**.

4. **Optional:** If using more than one identity provider, modify the MultiSSO installation exit.

a. Navigate to **System Definition** > **Installation Exits**.
The system displays the current list of installation exits.

b. Open the **MultiSSO** installation exit.

c. Locate the following statement in the **Script** field.

```
var samlResponseTxt = request.getParameter("SAMLResponse");
if (!GlideSession.get().isLoggedIn() &&
 GlideStringUtil.notNil(samlResponseTxt)) {
    var idpRecord = this.getIdPRecord(request);
    if (idpRecord) {
```

```
        SSO_Helper.debug("IdP found based on SAML response: "
 + idpRecord.getUniqueValue());
        return new SSO_Helper(idpRecord.getUniqueValue(),
 false, null, true);
    }
}
```

d. Replace the statement with the following code.

```
var samlResponseTxt = request.getParameter("SAMLResponse");
if (!GlideSession.get().isLoggedIn() &&
 GlideStringUtil.notNil(samlResponseTxt)) {
   /* // You have two profiles that use the same IdP entity
 id it cannot use
   // the IdP issuer / entity id from the response otherwise
 it may result in the
   // wrong IdP profile. IdP initiated login will not work
   var idpRecord = this.getIdPRecord(request);
   if (idpRecord) {
      SSO_Helper.debug("IdP found based on SAML response: " +
 idpRecord.getUniqueValue());
      return new SSO_Helper(idpRecord.getUniqueValue(),
 false, null, true);
      }*/
   return new SSO_Helper(null, true);
   }
```

Note: IdP initiated login does not work in this configuration.

e. Click **Update**.

5. **Optional:** If using more than one company, configure users for multi-provider SSO and update sys_id of the identity provider record depending on the user.

   (Optional) For more information, see Configure users for Multi-Provider SSO.

   - To configure a user to log in through the Edge Encryption proxy server, use the sys_id of the identity provider record that uses the Edge Encryption proxy server URL.
   - To configure a user to log in to the instance, use the sys_id of the identity provider record that uses the instance URL.

   Login URLs

URL	Login destination
`https://<proxy hostname>:<port>/login_with_sso.do?glide_sso_id=<sys_id of the IdP record for the proxy server URL>`	Logs in through the proxy server.
`https://<instance name>.service-now.com/`	Logs in through the instance.

URL	Login destination
`login_with_sso.do?glide_sso_id=<sys_id of IdP record for the instance URL>`	

Edge Encryption proxy server properties

The `edgeencryption.properties` configuration file located in the `<installation directory>/conf/` folder contains properties used to configure your environment.

You must restart the proxy server after making changes to any proxy server properties.

Clear text and static IV properties

edgeencryption.customer.assigned.known.cleartext	Clear text to let the instance verify that all proxies are using the same keys. At startup, the proxy encrypts the clear text and sends the encrypted text to the instance. The instance does not know the clear text, nor are keys sent to the instance. This property must be the same for all proxies.
edgeencryption.encrypter.static.iv	Static IV (initialization vector) used in equality-preserving and order-preserving encryption. This property must be the same for all proxies and must be exactly 16 bytes (16 ASCII characters).

Digital signature properties

edgeencryption.proxy.signature.keystore.path	Path and Java KeyStore file name.
edgeencryption.proxy.signature.keystore.password	Password. The default password is <changeme>. Change the password after installing the Java KeyStore.
edgeencryption.proxy.signature.keystore.keyalias	The key alias given as the -alias argument when the RSA key pair is generated.

File store property

edgeencryption.keyfile.directory	The directory specifies where key files are stored. If using the Java KeyStore or a SafeNet KeySecure keystore, leave this property commented out. Example: `edgeencryption.keyfile.directory=keys` If using Unbound Technology keys, uncomment this property and set the value to the keys directory.

General configuration properties

edgeencryption.config.poll.interval	Poll interval in seconds. The default setting means that it takes 5 seconds for the proxy to learn of encryption configuration changes. Larger values cause the instance to take longer to detect an offline proxy. ⚠ **Warning:** Do not change this property. Changing the default setting of the Proxy Poll Interval can result in detection delays when a proxy comes online.
edgeencryption.rules.dir	Folder where the encryption rules are stored on the proxy.
edgeencryption.encryption.order_preserving.caching.enable	Setting determines whether caching is used to support order-preserving encryption types.
edgeencryption.encryption.order_preserving.cache.size	Maximum cache size, in bytes.
edgeencryption.jobs.concurrency	Maximum number of mass encryption jobs that can run concurrently on this proxy.
edgeencryption.jobs.requests_per_second	Number of http job requests per second that can be sent to the instance by this proxy.
edgeencryption.attachments.request.timeout.seconds	Attachment upload request timeout in seconds.
edgeencryption.request.buffer.size	Size of an encryption request. If an encryption request is larger than this size, the excess is saved to disk. ⚠ **Warning:** Do not change this property.
edgeencryption.httpclient.request.buffer.size	Size of the client request. If the client request is larger than this size, the excess is saved to disk. ⚠ **Warning:** Do not change this property.
edgeencryption.httpclient.header.size	Size of the request/response header. • Minimum value: 8K • Maximum value: 32K ⚠ **Warning:** Do not change this property.
edgeencryption.proxy.idle.timeout	Time in seconds after which a transaction times out. Default value: 300 (seconds)
edgeencryption.proxy.keepalive.interval	Time in seconds between pings issued by the proxy to the instance. Pings are issued

	periodically to verify connectivity between the proxy and the instance. • Default value: 10 (seconds) • Minimum value: 5 (seconds)
edgeencryption.register.retry.count	Maximum number of times the proxy pings the instance to try to register. Default value: 0 (no limit)
edgeencryption.tokenization.exclusion.list	Encryption patterns cannot tokenize strings found in these fields.

Java KeyStore properties

edgeencryption.keystore.path	Path to the Java KeyStore. If using a file store or a SafeNet KeySecure keystore, leave this property commented out. Example: `edgeencryption.keystore.path = keystore/keystore.jceks`
edgeencryption.keystore.password	Password the proxy uses to connect to the Java KeyStore. If using a file store or a SafeNet KeySecure keystore, leave this property commented out.

Logging properties

Logging properties are found in the lo4gj2.properties file found in the `<installation directory>/conf/` directory. These properties are only changed for troubleshooting or when directed by ServiceNow support. For details see How to increase debug logging for the Edge Encryption proxy.

NAE device keystore properties

edgeencryption.nae.retries	Number of retries to make.
edgeencryption.nae.enabled	Setting indicates whether an NAE device is available.
edgeencryption.nae.server	Name of the NAE server.
edgeencryption.nae.port	Port used by the NAE server.
edgeencryption.nae.protocol	Protocol used by the NAE server.
edgeencryption.nae.keystore.path	Path to the keystore on the NAE server.
edgeencryption.nae.keystore.password	NAE keystore password.
edgeencryption.nae.username	User name to use to authenticate with the NAE device.

edgeencryption.nae.password	Password to use to authenticate with the NAE device.
edgeencryption.nae.client.certificate	Certificate located in the keystore on the NAE server. Set this property to authenticate using a certificate instead of a username and password.

Password property

edgeencryption.encrypter.properties.password	Name of the file in the `conf` folder that contains a string used within a secure process to obfuscate passwords in the `edgeencryption.properties` file. **Note:** Name of the file in the `conf` folder that contains a string used within a secure process to obfuscate passwords in the `edgeencryption.properties` file.

Proxy properties

edgeencryption.proxy.host	Server name, IP address, or fully qualified domain name of the computer running the proxy. Along with the port, this property defines the URL used by the client to access the proxy server.
edgeencryption.proxy.name	Proxy name. Must be unique for each proxy.
edgeencryption.proxy.http.port	Port on the proxy for HTTP communication.
edgeencryption.proxy.https.port	Port on the proxy for HTTPS communication.

Proxy configuration locked property

edgeencryption.proxy.locked	When true, the proxy does not accept encryption configuration changes or encryption rule changes from the instance. Set this property on the production instance after all encryption configurations and rules are final.

Proxy database properties

edgeencryption.db.url	Proxy database location. Must be the same for all encryption proxies connecting to the same instance.

edgeencryption.db.user	User name for accessing the proxy database. Must be the same for all encryption proxies connecting to the same instance.
edgeencryption.db.password	Password to access the proxy database. Must be the same for all encryption proxies connecting to the same instance.
edgeencryption.db.name	Proxy database name. Must be the same for all encryption proxies connecting to the same instance. The default for this property is `edgeencryption`.
edgeencryption.db.bootstrap.file	Bootstrap file for the proxy database. The file is relative to the `sql/` directory. Must be the same for all encryption proxies connecting to the same instance. ⚠ **Warning:** Under normal circumstances, do not change this parameter.

Proxy server performance properties

Proxy server performance properties are not present in the configuration file by default. To change the default values, you must add the properties and restart the proxy server. For more information, see Edge Encryption diagnostics and performance.

edgeencryption.stat.collection.enabled	Enables the collection of statistics used by the Edge Encryption proxy server performance dashboard. Default value: true Add this property and set the value to false to disable the collection of statistics used by the Edge Encryption proxy server performance dashboard.
edgeencryption.stat.collection.interval	Interval length in seconds during which the Edge Encryption proxy server collects statistics. The value cannot be less than 30 seconds. Default value: 30 (seconds)

SSL certificate properties

Restart your proxy if you change the value of any SSL certificate property. The system uses the HTTPS keypair on startup to establish the proxy server connection and determine how the proxy answers client requests.

edgeencryption.proxy.https.cert.alias	Alias of the certificate provided by the proxy server to connecting clients.

edgeencryption.proxy.https.keystore.path	Path to the keystore that contains the HTTPS certificate.
edgeencryption.proxy.https.keystore.password	Password for the keystore that contains the HTTPS certificate.

Target (instance) properties

edgeencryption.target.host	Host name for the instance. Must be the same for all encryption proxies connecting to the same instance. This property is set when the proxy is installed. For example, `instancename.servicenow.com`
edgeencryption.target.port	Instance port. Must be the same for all encryption proxies connecting to the same instance. This property is set when the proxy is installed.
edgeencryption.target.protocol	Instance protocol. Must be the same for all encryption proxies connecting to the same instance. This property is set when the proxy is installed. Options include: • http • https

Unbound Technology provider properties

edgeencryption.ekm.provider.classname	Internal class name for the implementation. **Warning:** Do not change this property.
edgeencryption.thirdparty.vendor.library.path	Path to the Unbound API JAR file on the Unbound client machine.
edgeencryption.ekm.provider.rsa.wrapping.key.alias	Wrapping key alias in the Unbound Technology implementation. Must be the same for all proxies.

User account properties

edgeencryption.target.username	User name that the proxy uses to log in to the instance. The user must have the edge_encryption role. See Set up an Edge Encryption user account.
edgeencryption.target.password	Password that the proxy uses to log in to the instance.

Web proxy properties

edgeencryption.webproxy.host	Web proxy name or IP address.
edgeencryption.webproxy.port	Port on the web proxy.
edgeencryption.webproxy.user	User name used to connect to the web proxy. If your web proxy does not use authentication, leave this property commented out.
edgeencryption.webproxy.password	Password to use to connect to the web proxy. If your web proxy does not use authentication, leave this property commented out.

Deprecated proxy encryption properties

edgeencryption.encrypter.default.key128

Specifies the name of the current AES 128 key. An AES 128 key must be available even if it is not used. Must be the same for all proxies.

Perform maintenance of these keys on the instance.

edgeencryption.encrypter.default.key256

Specifies the name of the current AES 256 key. Must be the same for all proxies.

Perform maintenance of these keys on the instance.

edgeencryption.encrypter.key

Specifies the key name for each key and is used to specify the default keys. This is the key alias integrated with the metadata that is included with each encrypted item and, therefore, is stored on the instance. The key name must use lowercase letters.

edgeencryption.encrypter.type

Specifies the type of encryption keystore system.

edgeencryption.encrypter.file

Specifies the path and file name of the text file associated with the key.

edgeencryption.encrypter.password

Specifies the password for accessing the keystore.

CyberArk integration with the Edge proxy server

Use CyberArk to store passwords in a centralized and secure digital vault to secure passwords that were previously stored in clear text and secured by file access, or that were previously encrypted via a second file.

CyberArk AIM (Application Identity Management) prevents unauthorized access by eliminating hard-coded and visible passwords. AIM stores passwords in a digital vault on an independent hardened server, where the passwords are represented as digital credentials. The AIM clients (the Edge proxy servers) use CyberArk digital credentials and access the independent server to retrieve the secured passwords. No passwords are stored on the Edge proxy servers or in the instance.

CyberArk digital vault credentials

You must purchase and configure CyberArk before you can set up CyberArk integration with the Edge proxy server.

To add a credential to CyberArk (which is read by the Edge proxy), set the **Platform Name** of the credential to **Unix via SSH** and make sure you either create a **Custom** credential **Name** or write down the **Auto-generated** credential **Name**. When you configure the Edge proxy to use this credential, the proxy server matches this credential **Name** to the setting in the proxy.

Each credential entry holds a **Password** that is being secured, as well as a credential **Name** used by an application to access that password.

Note: CyberArk credentials are not encryption keys.

Adding CyberArk during an Edge proxy installation

The proxy installer includes a new configuration page for a CyberArk integration. This page is optional if you do not want to include CyberArk when installing your proxy with the proxy installer. You can also manually set up and configure CyberArk integration in the configuration file.

The proxy installer also includes a new page for CyberArk protected credentials. This page allows configurations of different properties using a single credential name or multiple credential names. This page is optional if you do not want to include CyberArk when installing your proxy with the proxy installer.

CyberArk password protection

Any password field in the Edge proxy installer that has a CyberArk credential configured in the CyberArk vault and specified on the CyberArk Protected Credentials page of the installer is grayed out and contains the message `Protected by CyberArk`.

Using a load balancer with the Edge proxy server

You can use a load balancer to balance the load across the proxy servers in your Edge Encryption proxy setup. If the load balancer and proxy servers are using different ports, specify the host name and HTTPS port of the load balancer to enable users to view responses on their browser.

Important: All production environments should include at least two Edge Encryption proxy servers for redundancy.

Edge request processing without a load balancer

If you are not using a load balancer, a request is processed as described below.

1. The user issues a request from a browser.
2. The browser sends the request to the Edge proxy server.
3. The proxy server sends the request to the ServiceNow instance.
4. The ServiceNow instance returns the response to the proxy server.
5. The proxy server adds its own port number in the response header before returning the response to the user's browser.

The request is completed successfully because the user can view the response from the proxy server at the port number specified in the response header.

Edge request processing with a load balancer

However, if you are using a load balancer, the user's browser communicates directly with the load balancer, not with the proxy server. A request is processed as described below.

Note: The following example uses 1025 as the proxy server port number.

1. The user issues a request from a browser.
2. The browser sends the request to a load balancer Virtual IP (VIP), also known as a Virtual Server.
3. The VIP is configured to point to the proxy server (for example, 10 . 2 . 200 . 148 : 1025), so the load balancer forwards the request to the proxy server.
4. The proxy server sends the request to the ServiceNow instance.
5. The ServiceNow instance returns the response to the proxy server.
6. The proxy server rewrites the location header in the response with values configured in the properties for `risk - servicenow . dev . echonet : 1025`.
   - **Host**: `edgencryption . proxy . host`
   - **HTTP port**: `edgeencryption . proxy . http . port`
   - **HTTPS port**: `edgeencryption . proxy . https . port`
7. The proxy server forwards the response to the load balancer with the location header pointing to the proxy server port.

The outcome depends on whether the load balancer and proxy servers are using the same port.

- If the load balancer and proxy servers are using the same port, the request succeeds because the user receives the response from the same port identified in the response header.
- If the load balancer and proxy servers are using different ports, the request fails because the user's browser communicates only with the load balancer, but the response is on the proxy server.

Solution

You could resolve the issue by simply using the load balancer and all Edge proxy servers on the same port, but this is not an ideal solution. A better solution is to enable the system to know which port the load balancer uses.

The following properties enable the Edge proxy server to reroute response messages to the load balancer if the proxy server and load balancer are using different ports.

- *`edgeencryption.proxy.rewrite.location.host`* specifies the host name used to access ServiceNow through the load balancer.
- *`edgeencryption.proxy.rewrite.location.https.port`* specifies the HTTPS port used to access ServiceNow through the load balancer.

Configure the load balancer

If the load balancer and proxy servers are using different ports, specify the host name and HTTPS port of the load balancer to enable users to view responses on their browser.

Before you begin
Roles required:

- local or domain administrator on a Windows host
- service user with full file system access on a Linux host

Procedure

1. Login to the proxy server host as admin, domain admin, or a service user.
2. Navigate to the installation directory for the Edge proxy and select `conf/edgeencryption.properties`.
3. Set the following properties:

Property	Description
edgeencryption.proxy.rewrite.location.host	If your Edge configuration includes a load balancer to balance the load among proxy servers, rewrites responses to the load balancer so requests can be completed. ○ If there is a load balancer in the proxy setup, specify the host name used to access ServiceNow through the load balancer. ○ Optional: If there is no load balancer in the proxy setup, you can set this value to the host name used by the proxy server.
edgeencryption.proxy.rewrite.location.https.port	If your Edge configuration includes a load balancer to balance the load among proxy servers, specifies the HTTPS port used to access ServiceNow through the load balancer. ○ If there is a load balancer in the configuration, specify the HTTPS port used to access ServiceNow through the load balancer. ○ Optional: If there is no load balancer in the configuration, you can set this value to the HTTPS port used by the proxy server.

4. Save the file.

Result

Requests can be completed because users can now view responses on their browser.

Upgrading Edge Encryption

Both instance upgrades and proxy server upgrades require special consideration in an Edge Encryption environment.

Instance upgrades

Instance upgrades in an Edge Encryption environment require caution to ensure that Edge controls work properly after the instance upgrade.

During an instance upgrade, you should not add, edit, or delete the following:

- Edge Encryption configurations
- Edge Encryption rules
- Edge Encryption tokenization patterns
- Edge Encryption scheduled jobs
- Edge Encryption key configurations
- Edge Encryption scheduled upgrades
- Edge Encryption denylist IP configurations

Any scheduled job running during the instance upgrade will not complete. To complete the interrupted job, rerun the job once the instance is upgraded. When you reschedule the job, the processing that occurred before the instance upgrade is not lost, and the job continues to process only the data that has not yet been processed.

Proxy server upgrades

Schedule a proxy upgrade to enable the instance to upgrade the Edge Encryption proxy server, or manually upgrade the proxy server at any time.

Warning: For an upgrade on Windows, you may encounter file lock issues and the upgrade may fail. For the upgrade to succeed, you should not have any files open under the installation directory. Also, there should be no existing shell in the installation directory. In particular, if you start the proxy from the command line (via `bin\edgeencryption.bat install/start`) while in the installation directory, you should close that shell or move it out of the installation directory afterward. No files under the installation directory should be opened by an editor or by any other application.

Third party libraries

Third party libraries, such as Gemalto, are lost during instance and proxy server upgrades if they are kept in the same directory. You can perform the following to prevent the loss of third party libraries during upgrades:

1. Manually add the following property to edgeencryption.properties:

   ```
   edgeencryption.ekm.provider.classname =
   com.snc.edgeencryption.encryption.CloudEdgeNaeKeyProvider
   ```

2. Add the *edgeencryption.thirdparty.vendor.library.path* vendor library location property and set it to `/path/to/jars`.

   For example:

   ```
   edgeencryption.thirdparty.vendor.library.path = /app/servicenow/
   libs
   ```

3. Copy the SafeNet JARs into that path.

After you install the third party libraries outside of the Edge Encryption installation, they are no longer lost during upgrades.

Scheduled upgrades

Important: During ServiceNow Instance upgrades, also upgrade your proxy server version to align with your Instance version and reduce the chance of compatibility issues.

Schedule an upgrade to allow the instance to upgrade the proxy server at the scheduled time. This functionality is available by default after upgrading. A scheduled upgrade includes these events:

1. The proxy server checks with the instance to see if there is a new version available for upgrade. New versions generally become available when the instance is upgraded.
2. The administrator receives a notification upon logging in when a new version of the proxy server is available.
3. The administrator can Schedule an Edge Encryption proxy server upgrade for each proxy server.

   Note: Only users with the security_admin role can create an upgrade schedule through the proxy server.

4. Once the upgrade is scheduled, the proxy server automatically upgrades at the scheduled time. During the upgrade, the proxy server is offline for only a short time.

   Note: Because the proxy server restarts during the upgrade, it is offline for a short time. The amount of time is determined by your environment and how long it takes to stop and restart the proxy service.

5. During the scheduled upgrade, a new proxy directory is created and your configuration files are copied to the new directory. New properties are written to your existing properties file. The following files or directories in your old proxy directory are copied to the new proxy directory.
   - `/conf` directory
   - `/keys` directory
   - `/keystore` directory
   - `java/jre/lib/security/cacerts` file

   As a result, your keys, keystores, settings, and certificates are preserved.

   Note: Only the above files are copied to the new proxy directory. Any other customized files in the proxy server directory are not preserved during a scheduled upgrade. The upgrade log file can be found in the original proxy directory in the following folder: `<original-proxy-directory>/tmp/upgrade-wrapper/bin`.

Prerequisites for scheduled upgrades

Before scheduling an upgrade for an Edge Encryption proxy, ensure the following:

1. The *JAVA_HOME* environment variable points to a java installation on the machine that is outside the Edge Encryption proxy's directory structure.
2. The *JAVA_HOME* environment variable points to a java installation that is at version 1.8_u144 or higher.
3. The *-Djava.io.tmpdir* parameter in the wrapper.conf file of the Edge Encryption proxy points to a directory that is OUTSIDE the Edge Encryption proxy's directory structure, and the proxy has read/write/execute permissions on the directory. Optionally, you could comment out the parameter entirely so that Java uses its default tmp location.

Manual upgrades

Instead of creating an upgrade schedule, you can manually upgrade each proxy server through the command line. See Manually upgrade an Edge Encryption proxy server running on Linux or Manually upgrade an Edge Encryption proxy server running on Windows.

Proxy build status

You can easily identify whether a proxy server is out of date by navigating to **Edge Encryption Configuration** > **Proxies** > **All**. The status of your proxy build is indicated in the **Proxy build** column by the following colors:

Green

Your proxy server is up-to-date.

Yellow

Your proxy server is out-of-date and an upgrade is needed.

Orange

Upgrade failed. Your proxy server reverts to the old version to ensure that there is no downtime.

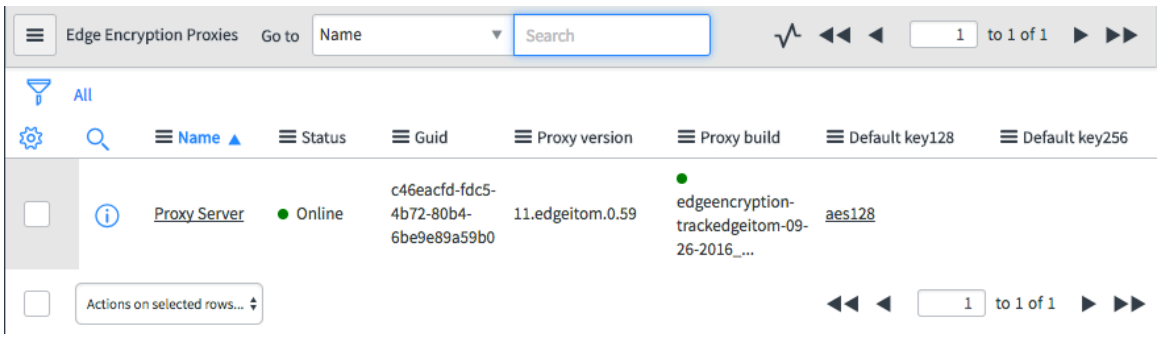

Troubleshoot a failed scheduled proxy upgrade

When a scheduled proxy upgrade fails, the proxy server reverts to the version you are upgrading from. All original data, keys, and configuration files are preserved. This process may take several minutes. Contact Customer Service and Support to ensure a successful upgrade.

To determine the reason for the failure, you can check the **Failure Reason** in the upgrade schedule. In addition, the installation directory for the failed upgrade is maintained so that log files are available for troubleshooting.

Note: Before deleting any extra proxy directories, always confirm which directory is current by reviewing the log files. If the log files have recent activity, the proxy might be connected to your instance.

If a scheduled proxy upgrade fails repeatedly, you can manually upgrade your proxy server. See Manually upgrade an Edge Encryption proxy server running on Linux and Manually upgrade an Edge Encryption proxy server running on Windows.

Java minimum requirements

The host machine installing or running the Edge Encryption proxy server must maintain a supported version of Java. Current supported versions are Java 17.0.3 or later in the 17.x version series.

Note: Java 11 is no longer be supported as of the Yokohama release. Upgrade your environment with the Edge Encryption proxy to Java 17 before you attempt to install Yokohama or later versions of the Edge Encryption proxy.

If using AES 256-bit encryption with Java 8 update 141 (8u141) or lower, you must install the Java Cryptography Extension (JCE) jurisdiction policy files by copying them into the system Java home directory of each Edge Encryption proxy server host. Add these files to the <Java -

home-directory>/jre/lib/security folder before performing a scheduled or manual upgrade. To install the AES 256-bit encryption policy files, see Configure the AES 256-bit encryption key.

Mixed proxy-version environments

Although an environment running old versions of the proxy server with up-to-date versions of the proxy server is not recommended, it is supported if all proxy servers are within the same version family as your instance. For example, if you have an instance on the Zurich release, your environment supports proxy servers from any Zurich patch or hot fix. However, the following limitations apply.

- If one proxy server supports functionality that another proxy does not support, you will see inconsistent behavior, depending on which proxy server is used.
- If a proxy server is out-of-date, it may not include recent security enhancements.

If a proxy server from a previous release is registered with a newer release of the instance, you will receive regular notifications that the proxy server is out-of-date. To ensure an optimal and secure environment, ServiceNow recommends always upgrading your proxy server to the most recent version of the software supported by your instance.

Schedule an Edge Encryption proxy server upgrade

Create an upgrade schedule to enable the instance to upgrade an out-of-date proxy server.

Before you begin

To schedule an upgrade, you must be logged in to your instance through the proxy server.
If using AES 256-bit encryption with Java 8 update 141 (8u141) or lower, you must install the Java Cryptography Extension (JCE) jurisdiction policy files by copying them into the system Java home directory of each Edge Encryption proxy server host. Add these files to the <Java-home-directory>/jre/lib/security folder before performing a scheduled or manual upgrade. To install the AES 256-bit encryption policy files, see Configure the AES 256-bit encryption key.

Role required: security_admin

About this task

Once the upgrade is scheduled, the proxy server automatically upgrades at the scheduled time. During the upgrade, the proxy server is offline for only a short time.

Note: Because the proxy server restarts during the upgrade, it is offline for a short time. The amount of time is determined by your environment and how long it takes to stop and restart the proxy service.

During the scheduled upgrade, a new proxy directory is created and your configuration files are copied to the new directory. New properties are written to your existing properties file. The following files or directories in your old proxy directory are copied to the new proxy directory.

- /conf directory
- /keys directory
- /keystore directory
- java/jre/lib/security/cacerts file

As a result, your keys, keystores, settings, and certificates are preserved.

Note: Only the above files are copied to the new proxy directory. Any other customized files in the proxy server directory will not be preserved during a scheduled upgrade. The upgrade log file can be found in the original proxy directory in the following folder: `<original-proxy-directory>/tmp/upgrade-wrapper/bin`.

If multiple proxy servers are out-of-date, you must schedule an upgrade for each proxy server individually.

Note: Avoid hosting multiple proxy servers on the same machine. However, if your environment includes this configuration, do not schedule upgrades to multiple proxies on the same machine at the same time.

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Proxies** > **Upgrade Schedules**.
2. Click **New**.
3. Complete the form.

Edge Encryption Proxy Upgrade Schedule form

Field	Description
Proxy server	Proxy server being upgraded.
Target version	Version to which you are upgrading your proxy server. This value is read-only and set to the most up-to-date proxy version available for your instance.
Scheduled Start Time	Date and time on which to start the upgrade.
Active	Whether the scheduled upgrade is active. If this field is not selected, the upgrade will not perform on the scheduled date and time.
Status	The status of the upgrade. This value is read-only. Possible statuses include: ○ Pending ○ Running ○ Complete ○ Failed

4. Click **Submit**.

What to do next

The typical time for an upgrade is less that 15 minutes. After an upgrade is executed, you can review the upgrade details to learn more about it. If your upgrade failed, review the **Failure Reason** to determine next steps.

Upgrade details

Field	Description
From Version	The version that the server was upgraded from.

Upgrade details (continued)

Field	Description
To Version	The version that the server was upgraded to.
Actual Start Time	Time that the upgrade began.
End Time	Time that the upgrade ended.
Failure Reason	Reason that the upgrade failed.

Manually upgrade an Edge Encryption proxy server running on Linux

Update a proxy running on Linux.

Before you begin

If using AES 256-bit encryption with Java 8 update 141 (8u141) or lower, you must install the Java Cryptography Extension (JCE) jurisdiction policy files by copying them into the system Java home directory of each Edge Encryption proxy server host. Add these files to the `<Java-home-directory>/jre/lib/security` folder before performing a scheduled or manual upgrade. To install the AES 256-bit encryption policy files, see Configure the AES 256-bit encryption key.

Role required: security_admin or local administrator on the host machine

Procedure

1. Download the Edge Encryption proxy-update archive file to the installation directory.
   a. Navigate to **Edge Encryption Configuration** > **Installation & Downloads** > **Downloads**
   b. Select the **Download the command line installer** link.
   c. When the download begins, select the your installation directory as the download location.
2. Change to the installation directory.
3. Run the following command:

```
java -jar edgeencryption-dist-<version>-linux-x86-64.jar -m
dist-upgrade -c <proxy directory>
```

Option	Description
proxy directory	The directory in the installation directory where the proxy was initially installed. This directory is created by the install.

If you want to see the help screen, execute this command without arguments: `java -jar edgeencryption-dist-<version>-linux-x86-64.jar`

A new proxy directory is created with a current timestamp. A backup of the old proxy directory is maintained as backup.dist-upgrade_timestamp in the new proxy installation directory. The old proxy shuts down and the new proxy starts up. Any open connections/transactions on the old proxy server are terminated.

4. Check the proxy log in the new directory and the instance to verify that the new proxy is running.

Manually upgrade an Edge Encryption proxy server running on Windows

Update a proxy running on Windows.

Before you begin

If using AES 256-bit encryption with Java 8 update 141 (8u141) or lower, you must install the Java Cryptography Extension (JCE) jurisdiction policy files by copying them into the system Java home directory of each Edge Encryption proxy server host. Add these files to the `<Java-home-directory>/jre/lib/security` folder before performing a scheduled or manual upgrade. To install the AES 256-bit encryption policy files, see Configure the AES 256-bit encryption key.

Role required: security_admin or local administrator on the host machine

Procedure

1. Download the Edge Encryption proxy-update archive file to the installation directory.
   a. Navigate to **Edge Encryption Configuration** > **Installation & Downloads** > **Downloads**
   b. Select the **Download the command line installer** link.
   c. When the download begins, select the your installation directory as the download location.
2. Start the Windows cmd terminal program with administrator privileges.
3. Change to the installation directory.
4. Run the following command:

```
java -jar edgeencryption-dist-<version>-all.jar -m dist-upgrade
-c <proxy directory>
```

Option	Description
proxy directory	The directory in the installation directory where the proxy was initially installed. This directory is created by the install.

If you want to see the help screen, execute this command without arguments: `java -jar edgeencryption-dist-<version>-all.jar`

A new proxy directory is created with a current timestamp. A backup of the old proxy directory is maintained as backup.dist-upgrade_timestamp in the new proxy installation directory. The old proxy shuts down and the new proxy starts up. Any open connections/transactions on the old proxy server are terminated.

5. Check the proxy log in the new directory and the instance to verify that the proxy has been updated and is running.

Roll back an Edge Encryption proxy server upgrade

If a proxy upgrade is unsuccessful, you can go back to the earlier version.

Before you begin

Role required: admin

About this task

If an upgrade fails when using the scheduled upgrade feature in the Zurich release, the proxy server will automatically roll back to the old version. The old proxy server is stored unmodified in a backup directory.

If you would like to roll back a manual upgrade, you can follow these steps.

Procedure

1. Shut down the proxy.
2. Delete the new proxy directory.
3. Rename the backup directory to the proxy name.
   The backup directory is in the proxy installation directory with the name `<proxy name>_backup`
4. Start the proxy.
5. Check the proxy log and the instance to verify that the proxy is online.

Configuring Edge Encryption

After the Edge Encryption proxy server is installed and running, manage Edge Encryption through the proxy server.

You must complete all the steps in Installing Edge Encryption before creating encryption configurations and encryption patterns on the instance.

Note: To access the Edge Encryption configuration, you must log in through the proxy server and elevate to the security_admin role.

Rotate encryption keys

Perform encryption key rotation from the instance. Add a new key, change the default key assignment, and then schedule a mass key rotation or a single key rotation.

Before setting an encryption key as the default key, make the key available to each proxy. This ensures that the proxies have the key to encrypt data when the key is assigned as the default key. All proxies must have access to a key before that key can be assigned as the default key.

Warning: Before deleting a key from the proxy, set up and run a mass key rotation job to ensure that no data on the instance uses the key. If any information is still encrypted with that key, you cannot decrypt the information after you delete the key.

Edge filtering and sorting behavior

Whenever you change default keys, be sure to perform a key rotation (either mass or single key rotation). Otherwise, you may receive unexpected results when sorting and filtering records. For example, consider the following scenario:

1. You create encrypted records using one encryption key.
2. You create a new key and set it as default.
3. You create a new set of encrypted records using the new encryption key.

If you filter by any encrypted field when connected through the Edge proxy, all records may not be filtered out correctly, or records may appear unexpectedly. The filter works only for records encrypted using the current default key. The records encrypted using the previous default key still appear in the list view.

If you sort by any encrypted field when connected through the Edge proxy, you see two groups of records with the same human readable text in the encrypted field.

Schedule a single key rotation job

Schedule a job to find data encrypted using a specified key alias and then re-encrypt the data with the current default encryption key. The data is decrypted before it is re-encrypted with the default key.

Before you begin

Role required: security_admin
Before scheduling this job, update the default key in **Edge Encryption Configuration** > **Encryption Key Configuration** > **Set Default Keys**.

Procedure

1. Navigate to **Edge Encryption Configuration** > **Maintenance** > **Schedule Single Key Rotation**.
2. Fill in the fields on the form as appropriate.

Field	Value
Name	Enter a descriptive name.
Job Type	Select **Single Key Rotation**.
Key	Enter the key to be retired. Verify that this key is no longer the default key in **Edge Encryption Configuration** > **Encryption Key Configuration** > **Set Default Keys**.
Estimate record count	Total estimated number of records to process. Not available when running a single key rotation.
Process Historical Records	Select to process historical records in the Audit table if the field is audited. When encrypting historical records for a field in the Audit table, both new values and old values are encrypted. This field is read only and active. To learn more about audited fields, see Auditing.
Estimate Maximum Audit Record Count	Estimated maximum number of audited records to process. Not available when running a single key rotation.
Active	Clear this check box if you want to deactivate this job.
Run	Select the period between job executions.
Starting	Enter the date and time to run the job for the first time.

3. Click the menu icon in the form header and select **Save**.
   Estimate Record Count is not supported when processing audited fields.

Schedule a mass key rotation job

Schedule a job to find data encrypted with any previous key, and then re-encrypt the data with the current default encryption keys. The data is decrypted before it is re-encrypted with the current default key.

Before you begin

Role required: security_admin

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Maintenance** > **Schedule Mass Key Rotation**.
2. Fill in the fields on the form as appropriate.

Field	Value
Name	Enter a descriptive name.
Job Type	Select **Mass Key Rotation**.
Estimate record count	Total estimated number of records to process. Not available when running a mass key rotation.
Process Historical Records	Select to process historical records in the Audit table if the field is audited. When encrypting historical records for a field in the Audit table, both new values and old values are encrypted. This field is read only and active. To learn more about audited fields, see Auditing.
Estimate Maximum Audit Record Count	Estimated maximum number of audited records to process. Not available when running a mass key rotation.
Active	Clear this check box to deactivate this job.
Run	Select the period between job executions.
Starting	Enter the date and time to run the job for the first time.

3. Click the menu icon in the form header and select **Save**.
   Estimate Record Count is not supported when processing audited fields.

Schedule an attachment key rotation job

Schedule a job to find attachments encrypted using a specified key alias, and then re-encrypt the attachments with the current default encryption key. The attachment is decrypted before it is re-encrypted with the default key.

Before you begin

Role required: security_admin

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Maintenance** > **Schedule Attachment Key Rotation**.
2. Fill in the fields on the form as appropriate.

Field	Value
Name	Enter a descriptive name.
Job Type	Select **Attachment Key Rotation**.
Active	Clear the check mark if you want to deactivate this job.

Field	Value
Table	Select a table.
Run	Select the period between job executions.
Starting	Enter the date and time to run the job for the first time.

3. Click the menu icon in the form header and select **Save**.
4. To see an estimated count of records to be updated, click **Estimated Record Count**.
5. To run the job immediately, click **Execute Now**.

Encrypt fields using encryption configurations

Encrypt fields by creating encryption configurations.

To configure Edge Encryption, you must be connected to the instance through the proxy. Test all changes on a non-production instance before applying them to the production instance.

Define encryption keys

After setting up one or more proxies and configuring a default encryption key, the instance verifies that the keys are available to all proxies. You can't make an encryption key the default key unless all proxies have the key. After a default key is defined, you can create encryption configurations.

Assign fields and attachments to be encrypted

Assigning fields and attachments to be encrypted means assigning an encryption type to the field or attachment. Before marking a field as encrypted, evaluate these issues.

- Determine what system features might be impacted.
- Examine all scripts for use of the field.
- Make any desired adjustments to the field size. After a field has been configured for encryption, the field size can't be changed.

Marking a field to be encrypted expands the field size to store the encrypted data. The process of expanding the field size can take a long time, depending on the number of records in the table.

API support

Field Encryption updates the `setDisplayValue()` and `setValue()` APIs so they can insert encrypted data for encrypted fields. It also enables `getDisplayValue()` and `getValue()` to return cleartext values.

The following script illustrates these API changes when the Incident short description is encrypted:

```
var gr = new GlideRecord('incident'); //creates a new incident
gr.setValue('short_description','test123'); //sets the value to
 test123
var sys_ID = gr.insert(); //inserts the record in the Incident
 table.
gs.info(gr.getValue('short_description')); //displays the
 unencrypted value
```

When using `getValue()` to get encrypted text, your script no longer returns the ciphertext. Your script returns the plaintext, assuming that the user has access to the cryptographic module. `getValue()` returns the ciphertext if the user doesn't have access to the cryptographic module.

Create a field encryption configuration

Select the fields to be encrypted and identify the encryption type.

Before you begin

Role required: security_admin

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Encryption Configurations** > **Create New**.
2. Complete the form.

Field	Description
Table	The table that contains the field to be encrypted.
Type	Whether to encrypt a table column or attachments for the table. Select **Column**.
Column	The field to be encrypted. Appears only when the **Type** is **Column**. Only String, Date, Date/Time, Journal, Journal Input, and URL fields are supported. ○ **String and URL fields**: You can add an encryption configuration to either a parent table or a child table. ○ **Date and Date/Time fields**: You can add an encryption configuration to a parent table only. You cannot add a new encryption configuration to a child table. **Note:** Depending on the number of records affected by the Date and Date/Time fields you are encrypting, it may take up to a few minutes to create the encryption configuration. Make sure that you create the encryption configuration for Date and Date/Time fields when transaction volume on the instance is low.
Encryption type	The encryption type to use.

Note: A specific table and field combination can have only one active configuration at a time.

3. Click **Submit**.

What to do next

After you add the encryption configuration record, you can create an encryption job to encrypt existing data. If you do not run an encryption job, Edge encrypts the existing data the next time the data changes. For details, see Schedule an encryption job.

Create a variable encryption configuration

Select service catalog variables to be encrypted and identify the encryption type.

Before you begin

Role required: security_admin

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Variable Encryption Configuration**.
2. In the **Edge Encryption Variable Configurations** list, click **New**.
3. Complete the form.

Field	Description
Variable	The variable to be encrypted.
Encryption type	The encryption type to use.

4. Click **Submit**.

What to do next

After you add the encryption configuration record, you can create an encryption job to encrypt existing data. If you do not run an encryption job, Edge encrypts the existing data the next time the data changes. For details, see Schedule an encryption job.

Deactivate an encryption configuration

After configuring a field or a table's attachments to be encrypted, you can stop encryption by deactivating the encryption configuration. After deactivating encryption, you can run a Decryption job for fields or an Attachment Decryption job for attachments to remove the encrypted data from the instance.

Before you begin

Role required: security_admin

About this task

⚠ **Warning:** Deactivating an encryption configuration does not delete the encryption record and the encryption type cannot be changed.

Procedure

1. Navigate to **Edge Encryption Configuration** > **Edge Encryption Configurations** > **All**.
   The **Edge Encryption Configurations** list is shown.
2. Click on the encryption configuration to be deactivated.
   The **Edge Encryption Configuration** form is shown.
3. Click on the **Active** box.
   The **Active** box is clear.
4. Click **Update**.
   The **Edge Encryption Configurations** list is shown.

What to do next

You can run a Decryption or Attachment Decryption job to decrypt data on the instance. If you do not run a job, the encrypted data is decrypted the next time it is changed.

Schedule an encryption job

You can schedule a job to find and encrypt any unencrypted data in a specified field, using the default encryption key configured for the field. If you do not create an encryption job after configuring a field for encryption, only new values are encrypted.

Before you begin

Role required: security_admin

Procedure

1. Navigate to **Edge Encryption Configuration** > **Encryption Configurations** > **All** to create a job for a field or **Edge Encryption Configuration** > **Variable Encryption Configuration** to create a job for a variable.
2. Click the field that you want to schedule an encryption job for.
3. Under **Related Links**, click **Schedule Mass Encryption Job**.

   The **Scheduled Encryption Job** form is shown with all fields populated. The bottom of the form shows records for any previous job executions.
4. Fill in the fields on the form, as appropriate.

Field	Value
Name	Enter a descriptive name.
Active	Clear this check box if you want to deactivate this job.
Job Type	Select **Encryption**.
Table	Select a table.
Column	Select a column.
Estimated record count	Total estimated number of records to process. Populates after selecting **Estimate Record Count**.
Process Historical Records	Select to process historical records in the Audit table if the field is audited. When encrypting historical records for a field in the Audit table, both new values and old values are encrypted. To learn more about audited fields, see Auditing.
Estimate Maximum Audit Record Count	Estimated maximum number of audited records to process. Populates after selecting **Estimate Record Count**. This field is only visible when **Process Historical Records** is selected. **Note:** The estimate may be larger than the actual number of records processed.
Run	Select the period between job executions.
Starting	Enter the date and time to run the job for the first time.

5. Click the menu icon in the form header and select **Save**.
6. To see an estimated count of records to be updated, click **Estimate Record Count**.
7. To run the job immediately, click **Execute Now**.

Schedule a decryption job

You can schedule a job to decrypt data in an encrypted field, to store clear data in the instance.

Before you begin

Note: You must mark the encryption record for the field as inactive (clear the **Active** box) in order to run the decryption job.

Role required: security_admin

Procedure

1. Navigate to **Edge Encryption Configuration** > **Encryption Configurations** > **All** to create a job for a field or **Edge Encryption Configuration** > **Variable Encryption Configuration** to create a job for a variable.
2. Click the field that you want to decrypt.
3. Under **Related Links**, click **Schedule Mass Decryption Job**.

   The **Scheduled Encryption Job** form is shown with all fields populated. The bottom of the form shows records for previous job executions.
4. Fill in the fields on the form, as appropriate.

Field	Value
Name	Enter a descriptive name.
Job Type	Select **Decryption**.
Active	Clear this check box if you want to deactivate this job.
Table	Select a table.
Column	Select a column.
Estimated record count	Total estimated number of records to process. Populates after selecting **Estimate Record Count**.
Process Historical Records	Select to process historical records in the Audit table if the field is audited. When encrypting historical records for a field in the Audit table, both new values and old values are encrypted. To learn more about audited fields, see Auditing.
Estimate Maximum Audit Record Count	Estimated maximum number of audited records to process. Populates after selecting **Estimate Record Count**. This field is only visible when **Process Historical Records** is selected. **Note:** The estimate may be larger than the actual number of records processed.
Run	Select the period between job executions.
Starting	Enter the date and time to run the job for the first time.

5. Click the menu icon in the form header and select **Save**.
6. To see an estimated count of records to be updated, click **Estimate Record Count**.
7. To run the job immediately, click **Execute Now**.

Encrypt attachments using standard encryption

You can encrypt attachments for specific tables.

All attachments to a table use the same encryption type. Encrypted attachments are not searched when performing a text search. Only the standard encryption types are allowed for attachments. The order preserving or equality preserving encryption types are not allowed.

For a session bypassing the Edge Encryption proxy:

- On a record with attachment encryption activated:
  - The user can see that there are attachments and the attachment names.
  - The user cannot add new attachments.
- On a record without attachment encryption activated:
  - The user can open and download existing attachments.
  - The user can add new attachments.

For a session using the encryption proxy, the user can open and download existing attachments and add new attachments.

Configure attachment encryption

Select the tables whose attachments are to be encrypted and identify the encryption type.

Before you begin

Role required: security_admin

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Edge Encryption Configurations** > **Create New**.
2. Fill in the fields on the form, as appropriate.

Edge Encryption configuration

Field	Description
Table	Select a table whose attachments are to be encrypted.
Type	Whether to encrypt a table column or attachments for the table. Select **Attachment**.
Column	The table field to be encrypted. This field appears when the **Type** is **Column**, and not when **Type** is **Attachment**.
Encryption type	The encryption type to use. For attachments, only Standard AES128 and Standard AES256 are allowed.

3. Click **Submit**.

What to do next

After the encryption record has been added, you can create an attachment encryption job to encrypt existing attachments. If you do not run an attachment encryption job, the system encrypts new attachments when you attach them.

Note: If you mark the *edge_encryption_clear_attachment_allowed attribute* as **True** in the table's Collection Dictionary entry, unencrypted attachments are added to a table using Edge Encryption to encrypt attachments. If enabling this attribute, you should setup an 'Attachment Encryption Job so that any unencrypted attachments added will be encrypted.

Schedule an attachment encryption job

You can schedule a job to find and encrypt any unencrypted attachments for a specified table, using the default encryption key configured for the table.

Before you begin

Role required: security_admin

Procedure

1. Navigate to **Edge Encryption Configuration** > **Encryption Configurations** > **All**.
2. Click the table you want to schedule an encryption job for.
3. Under **Related Links**, click **Schedule Mass Encryption Job**.

   The **Scheduled Encryption Job** form is shown with all fields populated. The bottom of the form shows records for previous job executions.
4. Fill in the fields on the form, as appropriate.

Field	Value
Name	Enter a descriptive name.
Active	Clear this check box if you want to deactivate this job.
Job Type	Select **Attachment Encryption**.
Table	Select a table.
Run	Select the period between job executions.
Starting	Enter the date and time to run the job for the first time.

5. Click the menu icon in the form header and select **Save**.
6. To see an estimated count of records to be updated, click **Estimate Record Count**.
7. To run the job immediately, click **Execute Now**.

Schedule an attachment decryption job

You can schedule a job to decrypt any encrypted attachments for a specified table, to store clear attachments in the instance.

Before you begin

Note: You must mark the encryption record for the table as inactive (clear the **Active** box) before the decryption job runs, otherwise, nothing happens.

Role required: security_admin

Procedure

1. Navigate to **Edge Encryption Configuration** > **Encryption Configurations** > **All**.
2. Click the table with the attachments that you want to decrypt.
3. Under **Related Links**, click **Schedule Mass Attachment Decryption Job**.

   The **Scheduled Encryption Job** form is shown with all fields populated. The bottom of the form shows records for previous job executions.
4. Fill in the fields on the form, as appropriate.

Field	Value
Name	Enter a descriptive name.
Job Type	Select **Attachment Decryption**.
Active	Clear the check mark if you want to deactivate this job.
Table	Select a table.
Run	Select the period between job executions.
Starting	Enter the date and time to run the job for the first time.

5. Click the menu icon in the form header and select **Save**.
6. To see an estimated count of records to be updated, click **Estimate Record Count**.
7. To run the job immediately, click **Execute Now**.

Change a field or attachment's encryption type

You can change a field or attachment's encryption type by selecting a new encryption type in the existing encryption configuration record. A specific table and field combination can only have one active configuration at a time.

Before you begin

Role required: security_admin

Procedure

1. Navigate to **Edge Encryption Configuration** > **Encryption Configurations** > **All**.
   The **Edge Encryption Configurations** list is shown.
2. Open the record for the encryption configuration to be changed.
3. Click the **Encryption type** dropdown and select a new encryption type.

   Note: For attachments, only Standard AES128 and Standard AES256 are allowed.
4. If needed, run an encryption or attachment encryption job.
   It is not necessary to run an encryption job. If you do not run an encryption job, the field or attachment is encrypted using the new encryption type the next time the field or attachment is changed.

Tokenize strings using encryption patterns

You can replace string patterns with tokens before they are sent to and stored in the instance.

Before you begin

To use encryption patterns, you must install and set up a MySQL database in your network. This is the same database used for order-preserving encryption. To create or edit encryption patterns, you must be connected to the instance through the proxy.

Role required: security_admin

About this task

You can use base system patterns, or create your own patterns. Base system patterns are advanced patterns. Encryption patterns include the following limitations.

- A pattern of all alpha characters is not allowed.
- The minimum pattern size is five characters. You can change this setting using a system property.
- The * and + quantifiers are forbidden in encryption patterns.
- Encryption patterns match complete words, not parts of strings embedded in a larger string. Words are defined by spaces and characters not available for inclusion in a pattern.
- If the same string is sent to the instance multiple times, it is replaced with the same token.
- Text search on exact matches is supported. The query string is exchanged with a token when sent to the instance, the search is performed on tokens, and when the search results are returned to the proxy server, the tokens are replaced with the clear text. Features such as stemming are not supported.

When using patterns, the clear text never leaves your network. When the proxy server matches a pattern in a request going to the instance, the proxy replaces the string with a token of the same size. The token is sent to instance instead of the clear text string. When the response is sent from the instance to the proxy server, the proxy replaces the token with the string. When viewed through the proxy server, the string displays as clear text.

Note: Encrypted fields are not checked for encryption patterns.

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Encryption Patterns** > **Create New**. Alternatively, you can navigate to **Advanced Patterns** to activate or edit a preconfigured pattern.
2. Enter the pattern name.
3. Define the **Edge pattern input type**.

Option	Description
Basic	A series of character types. In the **Basic Pattern Input** tab, click **Add** and select a character type. The **Sample pattern** displays the pattern as characters are added. Click **New Block** to move the next character to the next line. This enables you to group characters in a long pattern. Click **X** to delete the last character in the pattern.

Option	Description
Advanced	A Java RegEx expression. If advanced is se lected, you cannot change the input type back to basic. In the **Sample match** field, enter a sample pattern to test the RegEx expression. In the **Pattern** field, enter a Java RegEx expression. Click **Validate** to verify that the expression matches the sample pattern.

The input type defines how you are going to enter the pattern. It does not impact how the pattern is used.

4. Click **Submit**.

Repair or recover order-preserving encrypted data

If you have the security-admin role, you can schedule jobs performed by the Edge Encryption proxy to repair or recover fields that use order preserving encryption.

Schedule jobs to:

- Repair order tokens.
- Recreate the proxy database.

Running these jobs can be a time-consuming operation that might impact the performance of the Edge Encryption proxy. Schedule these jobs at a time when no users or a minimum set of users are using the system, such as midnight on the weekend.

Schedule an order token repair job

You can schedule a job to find and repair fields where the order token is missing.

Before you begin
Role required: security_admin

About this task
Use these jobs to repair individual fields in a table or to repair all fields using order preserving encryption. Run this job when the proxy database has been offline while the instance has been running, which results in order preserving fields that are missing order tokens.

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Maintenance** > **Schedule Order Token Repair**.
2. Fill in the fields on the form, as appropriate.

Field	Value
Name	Enter a descriptive name.
Job Type	Select **Order Token Repair**.
All fields	Select this check box to repair all tables.
Table	Select a table.

Field	Value
Column	Select a column.
Active	Clear this check box if you want to deactivate this job.
Run	Select the period between job executions.
Starting	Enter the date and time to run the job for the first time.

3. Click the menu icon in the form header and select **Save**.
4. To see an estimated count of records to be updated, click **Estimated Record Count**.

Schedule a proxy-database recovery job

Run this job when the proxy database has lost data. This job finds all records that have been encrypted with a token (order preserving encryption type) and sends them to the proxy so that the proxy database can be rebuilt.

Before you begin

Role required: security_admin

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Maintenance** > **Schedule Database Recovery**.
2. Fill in the fields on the form, as appropriate.

Field	Value
Name	Enter a descriptive name for this job.
Job Type	Select **Database Recovery**.
Active	Clear this check box if you want to deactivate this job.
Run	Select the period between job executions.
Starting	Enter the date and time to run this job for the first time.

3. Click the menu icon in the form header and select **Save**.
4. To see an estimated count of records to be updated, click **Estimate Record Count**.

Configure the IP address deny list

Prevent an IP address in your network from sending requests to your instance

Before you begin

Role required: security_admin

Because the Edge Encryption proxy server resides in your network, it may be subject to vulnerability scans by your network software. To prevent IP scanner or other requests from being forwarded to your ServiceNow instance, you can add IP addresses, IP ranges, or network masks to a deny list. Any connection to the proxy server from a deny listed address is terminated and is not forwarded to your instance.

To place an IP address on a deny list, you must be logged in to your instance through the proxy server.

Important: Ensure that you understand your network topology before adding IP addresses in your network to a deny list. If an IP address is added to the deny list, any user with that IP address will be blocked from accessing the Edge Encryption proxy server.

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Maintenance** > **Denylist IP Addresses**. The Encryption Proxy IP Denylists [edge_encryption_ip_blacklist] list view opens.
2. Click **New**.
3. Complete the form.

Field	Description
Proxy server	The Edge Encryption proxy server that is prevented from forwarding requests from addresses on the deny list.
IP, IP range, or net-mask	Requests from this IP address, range, or network mask are not forwarded to your ServiceNow instance. Example values include: ○ IP address: 10.10.10.5 ○ IP range: 10.10.10.1-15 ○ Network mask: 10.10.10.0/24 **Note:** You may use either IPv4 or IPv6 addresses
Active	Whether the record is active. Only IP addresses from active records are prevented from sending requests to the instance.
Description	Description of the deny list record.

4. Click **Submit**.
5. Repeat these steps for all other proxies for which an IP address should be denied.

Result

The Edge Encryption proxy server terminates any connection from IP addresses, ranges, or network masks on the deny list and cannot forward the request to the instance.

Encrypt data from a record producer

Configure your Edge Encryption proxy server to allow inserts from a record producer by creating encryption rules from the record producer record.

Before you begin

Role required: security_admin

Record producers allow end users to create task-based records, such as incident records, from the Service Catalog and Service Portal. If a record producer attempts to insert data into a field marked for encryption, an invalid insert message displays and the data is not saved to the field.

Encrypting data from a record producer requires an encryption configuration defined for the target field. Check that you have created an encryption configuration for the target field and table before creating an encryption rule from a record producer. See Create a field encryption configuration. To encrypt attachments from a record producer, Configure attachment encryption.

Procedure

1. Log in to your instance through the Edge Encryption proxy server.
2. Navigate to **Service Catalog** > **Catalog Definitions** > **Record Producers**.
3. Create a record producer record or open an existing record producer record.
4. Under **Related Links**, select **Create Edge Encryption Rule**.
   Two inactive encryption rules are automatically created to encrypt data sent from the record producer to the field marked for encryption.

Encryption rule	Description
<RecordProducerName>	Rule created to process POST parameters from the Service Catalog and map variables to fields in the instance.
<RecordProducerName>Json	Rule created to process a JSON payload from the Service Portal and map variables to fields in the instance.

5. Activate the necessary encryption rules created by the record producer.

   a. Navigate to **Edge Encryption Configuration** > **Rules** > **All**.

   b. Depending on where the record producer will be used, open the associated encryption rule created by the record producer and select the **Active** flag.
   If using the record producer in the Service Catalog, activate the <RecordProducerName> encryption rule. If using the record producer in the Service Portal, activate the <RecordProducerName>Json encryption rule.

6. **Optional:** Examine the Encryption rule **Action** field and add any necessary field names or statements.

   (Optional) If a record producer directly maps a variable to a field in a table, the encryption rule automatically maps the variable to the correct field. However, if a variable is indirectly mapped through various scripts on the platform, you may need to update the rules to map each variable to the correct field.

   #### Example

   (Optional) The below encryption rule was created from the Report Outage record producer and processes POST parameters from the Service Catalog to map variables to fields in the instance. Replace `'FILL ME IN'` with the target field.

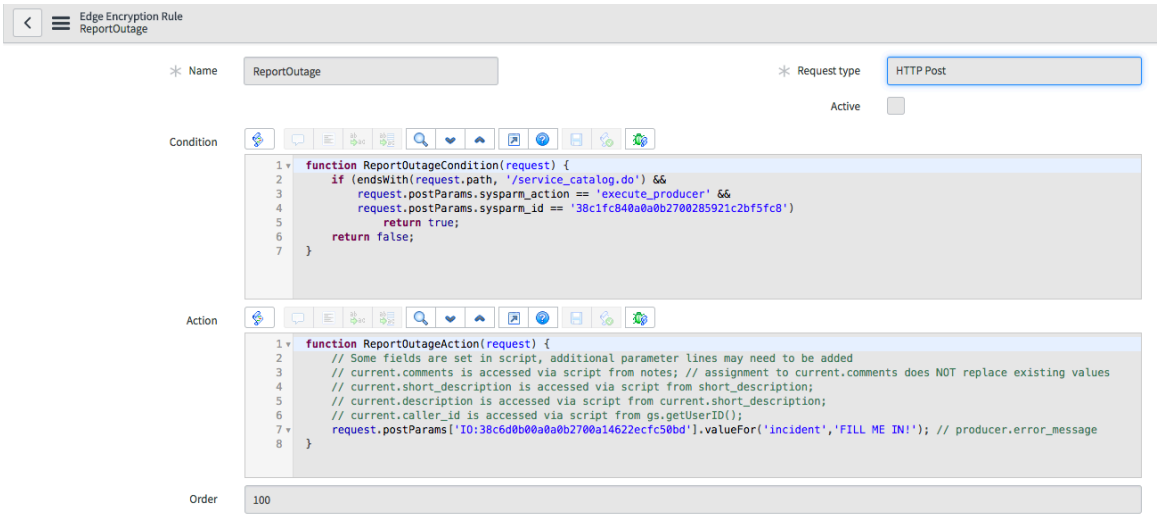

The below encryption rule was created from the Report Outage record producer and processes a JSON payload from the Service Portal to map variables to fields in the instance. Add additional statements to map any scripted variables to the target fields.

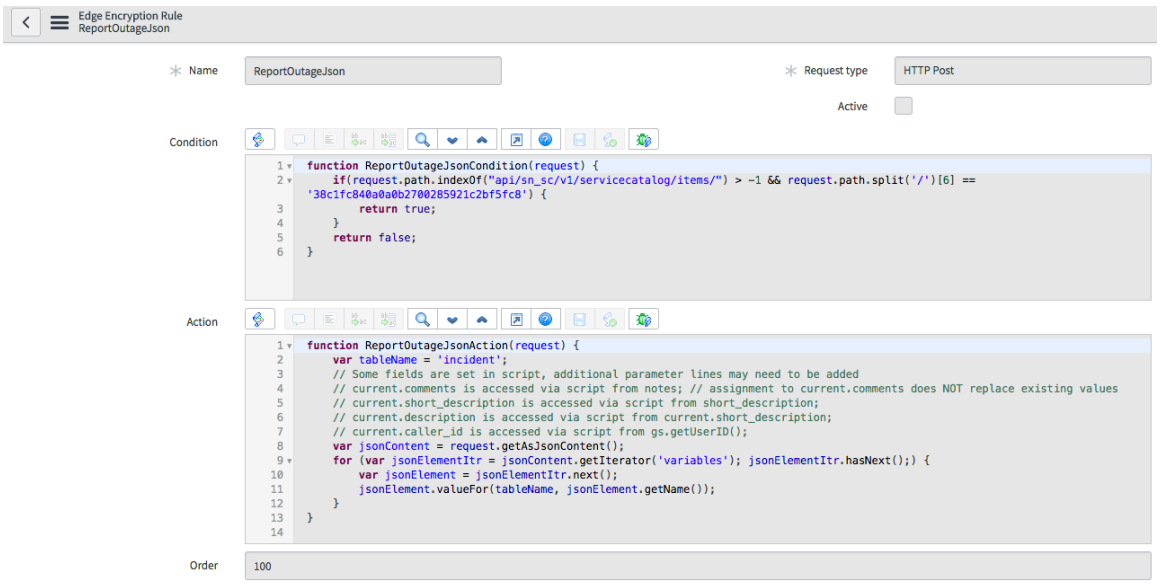

When the payload from the record producer is examined, the error_message element contains the value for the short_description field. By adding the following statement, you can map the scripted variable error_message to the short_description field.

```
if (jsonElement.getName() == 'error_message')
        jsonElement.valueFor(tableName, 'short_description');
```

The value of the **Action** field becomes:

```
function ReportOutageJsonAction(request) {
 var tableName = 'incident';
 // Some fields are set in script, additional parameter lines
 may need to be added
 // current.comments is accessed via script from notes; //
 assignment to current.comments does NOT replace existing
 values
```

```
// current.short_description is accessed via script from
short_description;
// current.description is accessed via script from
current.short_description;
// current.caller_id is accessed via script from
gs.getUserID();
var jsonContent = request.getAsJsonContent();
for (var jsonElementItr =
jsonContent.getIterator('variables');
jsonElementItr.hasNext();) {
 var jsonElement = jsonElementItr.next();
 if (jsonElement.getName() == 'error_message')
  jsonElement.valueFor(tableName, 'short_description');
            } else {
        jsonElement.valueFor(tableName,
jsonElement.getName());
 }
}
```

Result
The two encryption rules enable the record producer to insert values into fields marked for encryption from either the Service Catalog or Service Portal.

Define a custom encryption rule

It may be necessary to identify and encrypt sensitive information in HTTP requests on the way to your instance. You can write encryption rules to identify, interpret, and encrypt data in such requests, mapping fields in the request to table-field names on your instance.

What is an encryption rule

Encryption rules are scripts executed on the Edge Encryption proxy server to map fields in a request to fields in a table on your ServiceNow instance. An encryption rule tells the Edge Encryption proxy server how to encrypt data in custom payloads.

Note: Encryption rules only support ECMAScript 3 and below.

When to use custom rules

A set of encryption rules is included as part of the Edge Encryption plugin. These rules handle many core platform use cases, such as:

- Editing a field from the list edit form
- Updating a record from the record form
- Managing direct web service
- Processing data from the REST Application Program Interface (API)

Applications created using standard forms and lists should work without custom encryption rules.

If you develop scripts that contain data that should be encrypted, create encryption rules to find and map that data to Glide table-field names. For example:

- Scripted processors
- Scripted web services
- Scripted REST APIs, UIs, or Ajax scripts

Format of an encryption rule

Rules include three parts:

- **Condition**: Identifies the type of request.
- **Action**: Maps fields in the request to fields in a table, encrypting values that map to fields with encryption configurations defined.
- **Order**: Priority of the rule. The lowest priority rule with a satisfied condition is the only rule that runs. Like business rules, rules run from lowest to highest.

Except for attachment requests, HTTP requests are evaluated by the Edge Encryption proxy server. The Edge Encryption proxy server evaluates all encryption rule conditions in priority order until either all conditions return false, or one condition returns true. When a condition returns true, the action is executed on the request, and the result is forwarded to the instance. No other conditions are evaluated. As a result, encryption rule conditions should be as specific as possible. A generic rule might evaluate as true for a request meant to be processed by another rule, causing the request to be processed by the wrong action. If a generic condition is unavoidable, the rule should be marked with a high-order value so that more specific rules are evaluated first.

Guidelines for creating encryption rules

Creating efficient, optimized encryption rules can reduce processing time for script validation.

Overall guideline: When rules get long, do your best to minimize the number of blocks and break the rules apart whenever possible. Ideally, custom rules should apply to specific use cases, rather than encompassing several cases, with `if`s or switch statements in the action script.

1. Split rules whenever possible. For example,
   - Create different rules for different tables and ensure that each rule runs only on its respective table.
   - Create different rules for each record producer that you're targeting, or at least for each subset of record producers. Instead of one rule targeting dozens of `sys_ids`, you could create several different rules targeting smaller subsets of record producers, or even create one rule per `sys_id`.

     Note: Creating multiple rules requires more maintenance. The trade-off is that multiple, simpler rules can be validated more efficiently than longer, more complex rules.

2. Minimize the number of blocks. Because the processing engine scans each block while evaluating scripts, a large number of blocks causes delays in validation. For example,
   - Replace all `if` blocks with an array lookup, and replace all blocks in the array lookup with just one `if` block.
   - Combine `if` blocks whenever it's possible to group them.

Encryption rule APIs

Encryption rules are written in JavaScript and utilize Edge Encryption APIs to locate and encrypt sensitive information in the body of a request. The API uses expressions similar to xPath to navigate through both JSON and XML content.

Edge Encryption APIs process the request off the stream as it is being written to the output stream. Stream parsing enables encryption rules to be network performant. However, fetching and parsing content from the body multiple times could lead to unexpected results. To void this potential problem, requests should be processed by the action in a single pass.

When creating encryption rules, you can't use Glide APIs, script includes, business rules, or any global parameters such as `current`. Because the rules are created for HTTP objects, a global `request` object is available.

When creating encryption rules, you can't use APIs from the allow list manager or scoped applications.

Error handling

If an encryption rule condition or action throws an exception, check the proxy log for troubleshooting information.

Inspect the client request

Before creating a custom encryption rule, you must determine the format of the client request entering the Edge Encryption proxy server.

Before you begin
Role required: admin

About this task
Because encryption rules iterate over client requests and determine what, if anything, needs to be encrypted, you must understand the type of request you are creating a rule for. The format of the client request determines the structure of your encryption rule and the APIs available for use in the rule.

Procedure

1. Inspect the client request.

   Depending on the source of the request, the following tools are available to inspect the request and determine the format.

Source of request	Available tools
Client browser	Use the developer console in your browser to inspect the client request. Useful tools include: ○ Firefox Network Monitor ○ Chrome Network Panel

Source of request	Available tools
Third-party/external source	Use an HTTP protocol analyzer to inspect the request. Useful tools include: ○ Wireshark ○ HTTP Scoop Alternatively, you can often use documentation for the external source to determine the format of the request.

2. From the client request, inspect the packet and determine:
   - The client request method
   - The URL path of the request
   - The URL parameters
   - The POST parameters, if any
   - The format of the request body, if included

Example

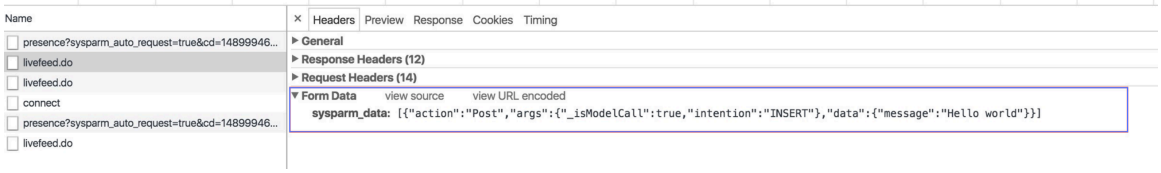

Result
Inspecting the request provides an understanding of the fields you need to filter for and iterate over in your encryption rule. To understand the fields in the `request` object, see request.

Create an encryption rule

Encryption rules are used by the proxy to find content in HTTP requests that should be encrypted.

Before you begin
Role required: security_admin
Before creating an encryption rule, you must Inspect the client request to determine the format.

About this task
To create or edit encryption rules, you must be connected to the instance through an encryption proxy.

Procedure

1. Navigate to **All** > **Edge Encryption Configuration** > **Rules** > **Create New**.
2. In the **Name** box, enter a name.
3. In the **Request Type**, select an HTTP method.
   - **HTTP Post**
   - **HTTP Get**
   - **HTTP Put**

- **HTTP Patch**
- **HTTP Delete**

Note: Pre-Jakarta instances allow only **HTTP Get** and **HTTP Post** methods.

4. In the **Condition** box, enter a JavaScript statement defining when the rule should run.
5. In the **Action** box, enter a JavaScript function to be executed when the condition is true.
6. In the **Order** box, enter the relative priority of the rule.
7. Click **Submit**, or save the form.

Encryption rule conditions

Encryption rule conditions determine if the rule should be executed.

An encryption rule condition must return true if the rule is to handle the HTTP request; otherwise, it must return false.

As you build your condition, keep in mind that only one rule is executed per request. As a result, the condition must be as general or specific as needed to run under the intended circumstances.

Note: Be careful when performing checks on content in the condition. Excessive checks can be expensive for the proxy server and may cause increased latency when handling complex requests.

The condition can use the method type, content type, URL path, or any URL query string parameters to determine if the rule should handle the request. The condition has access to these fields via the request object. Be sure that, prior to creating an encryption rule condition, you have inspected the client request and understand the conditions needed to trigger the rule.

Note: To build efficient rules, consider easy ways to rule out requests that you do not want to be evaluated by a rule. Build your condition to return false for those requests first. This method increases performance and quickly routes the request to the correct rule faster.

Encryption rule objects and APIs are available to encryption rule conditions.

Example using path and postParams

```
/*This condition checks if the request coming in has a path
 ending in
"/sample_processor.do" and if a post parameter exists in that
 request called myPostParam */

function SampleCondition(request) {
   if (endsWith(request.path, "/sample_processor.do") &&
 request.postParams.myPostParam) {
       return true;
   }
   return false;
}
```

Example using urlParams and contentType

```
/* This condition checks if a url parameter exists in the query
 called
myUrlParam and if the content type contains 'xml'
```

```
(if so, you can expect the body to be an XML payload).
Then, it checks if the xml payload contains myXmlTag */

function SampleCondition2(request) {
   if (request.urlParams.myUrlParam &&
 request.contentType.indexOf('xml') > -1 &&
 request.xmlContains('myXmlTag')) {
       return true;
   }
 return false;
}
```

Encryption rule actions

An encryption rule maps fields in a client request to fields in a table on your instance and identifies fields marked for encryption.

An encryption rule action only runs when the encryption rule condition returns true. An encryption rule identifies the data to be encrypted in your request payload. Because the rule iterates over the content in the request object, you must understand the form and structure of your request body and determine what in the request must be encrypted. The data to be encrypted might be located within:

- A POST or URL parameter.
- JSON or XML content within a POST or URL parameter.
- A JSON payload.
- An XML payload.

Before writing an encryption rule action, be sure to:

- Inspect the client request.
- Identify where the sensitive data is located in the request object.
- Determine the field and table name to insert data into, or understand how to dynamically pull this from the request.

Encryption rule objects and APIs are available to encryption rule actions and conditions.

Encryption rule objects and APIs

Use encryption rule APIs to parse and encrypt values in requests moving through the Edge Encryption proxy server to the instance.

The APIs available for your encryption rule depend on the format of the request object. For example, if the contentType parameter of the request object is XML, you can use the XML APIs to parse and encrypt values in the payload. After you determine the type of object in your request, you can build an encryption rule using the available APIs.

Encryption rule APIs are available in both encryption rule condition and action scripts.

request

The `request` object is a global object available in Edge Encryption rule action and condition scripts.

The `request` object is a JavaScript object that represents the client request coming in to the Edge Encryption proxy server. You must build your encryption rule to parse the `request` object,

map *request* object values to fields in a table on the instance, and encrypt any sensitive data in the *request* object.

The *request* object includes the following attributes and data from the client request:

Request object fields

Field	Description
path	The path portion of the URL.
requestMethod	GET, POST, PUT, PATCH, DELETE.
contentType	The Content-Type header field.
urlParams	The parameters in the query string. This can also be evaluated to a String.
postParams	If this is a form post, this contains the post parameters.

request - getAsJsonContent()

Returns the request as an iterable object of type *JsonNode*.

This method is available only in an Edge Encryption rule if the request body is a valid JSON payload. If you are not sure what format the request body includes, check the contentType field on the *request* object.

Once the request is returned as a *JsonNode* object, you can use the JSON APIs to iterate over the object and encrypt fields.

Parameters

Name	Type	Description
None		

Returns

Type	Description
JsonNode	The request as an iterable *JsonNode*.

request - getAsXmlContent()

Returns the request content as an iterable object of type *XMLContent*.

This method is available only in an Edge Encryption rule if the request body is a valid XML payload. If you are not sure what format the request body includes, check the contentType field on the *request* object.

Once the request is returned as an *XMLContent* object, you can use the XML APIs to iterate over the object and encrypt fields.

Parameters

Name	Type	Description
None		

Returns

Type	Description
XMLContent	The request as an iterable object of type *XMLContent*.

request - XMLContains(String path)

Returns true if the given path exists in the XML DOM.

This method is available only if the request body is a valid XML payload. If you are not sure what format the request body includes, check the contentType field on the *request* object.

Parameters

Name	Type	Description
path	String	XPath statement you are searching for.

Returns

Type	Description
Boolean	Whether the given path exists in the XML DOM.

POST and URL parameter APIs

POST and URL parameters can be accessed as properties of the *request* object using `request.postParams` and `request.urlParams`.

Any single parameter can be accessed as a property of the *postParams* and *urlParams* parent objects by calling `request.postParams.myParam`. Any parameter accessed this way is an object of the underlying class *ParameterValue*. Any APIs in this class can be called on any parameter.

After inspecting the client request, it may be necessary to access and encrypt parameter values from the *request* object. Depending on the data in the client request, you can encrypt values and map them to fields on the instance in multiple ways.

Encrypt the value of a known table and field

If you know the name of the instance table and field that will hold the encrypted data, you can explicitly define them in the encryption rule. For example, you may know that the request will be processed on the instance to create an incident and you want to encrypt the **text** parameter in the description field. In this case, you can create the following action.

```
function SampleAction1() {
    request.postParams.text.valueFor('incident', 'description');
}
```

Encrypt the value of a dynamically defined table and field

If, conversely, you do not know the name of the field that the encrypted data will populate, you can dynamically define them using **tableName** and **fieldName**.

The below example processes a generic request that might store data in different task tables (such as incident, problem, and change_request) on the instance.

```
function SampleAction2() {
    var tableName = request.urlParams.table;
    for (var parameter in request.postParams) {
        var currentParam = request.postParams[parameter];
        var fieldName = currentParam.toString();
        if (fieldName == 'text') {
            currentParam.valueFor(tableName, 'description')
        } else {
            currentParam.valueFor(tableName, fieldName);
        }
    }
}
```

This action:

- Gets the destination table from the URL parameters.
- Iterates over the URL parameters.
- Asks the Edge Encryption proxy server to encrypt any URL parameter with a name that matches a field marked for encryption.
- Looks for a specific parameter called *text* and asks the Edge Encryption proxy to encrypt the value based on the encryption configuration for the description field on the incident table.

In this example, the valueFor() method is not actually performing any encryption. Rather, the method asks the Edge Encryption proxy server to check whether the table/field pair in the request object is marked for encryption with an encryption configuration and, if applicable, encrypt it.

Encrypt JSON or XML within a parameter

A POST or URL parameter might include JSON or XML content. In this case, you can process the content within the parameter, iterate over the values, and encrypt required fields. In this example, the **tableName** is still accessed from a POST parameter, but the value of the field is the JSON object **data**.

```
function SampleAction3() {
    var tableName = request.postParams.table;
    var data = request.postParams.data;
    var dataIterator = data.getAsJsonContent().iterator();
    while (dataIterator.hasNext()) {
        var jsonElement = dataIterator.next();
        var fieldName = jsonElement.getName();
        if (fieldName == 'text') {
            jsonElement.valueFor(tableName, 'description');
        } else {
            jsonElement.valueFor(tableName, fieldName);
        }
    }
}
```

An example of an encryption rule action that processes XML within a POST parameter.

```
function SampleAction4() {
    var tableName = request.postParams.table;
    var data = request.postParams.data;
```

```
    var dataIterator =
 data.getAsXmlContent().getIteratorOverAllChildren();
    while (dataIterator.hasNext()) {
        var jsonElement = dataIterator.next();
        var fieldName = jsonElement.getName();
        if (fieldName == 'text') {
            jsonElement.valueFor(tableName, 'description');
        } else {
            jsonElement.valueFor(tableName, fieldName);
        }
    }
}
```

Encrypt a query

You might encounter an encoded query within a parameter in the client request that contains sensitive data. To match a field in a query to an encrypted value in the instance database, you must create an encryption rule that asks the proxy to check whether a field in the query is marked for encryption. The encodedQueryFor() method parses an encoded query on a given table, and checks if any fields in the query have encryption configurations.

In this example, the rule iterates over the parameters looking for the **filter** parameter, which is expected to be a Glide encoded query.

```
function SampleAction5() {
    var tableName = request.urlParams.table;
    for (var parameter in request.postParams) {
        var currentParam = request.postParams[parameter];
        var fieldName = currentParam.toString();
        if (fieldName == 'filter') {
            currentParam.encodedQueryFor(tableName);
        } else {
            currentParam.valueFor(tableName, fieldName);
        }
    }
}
```

For example, if the value of **filter** is: `short_description=My sensitive information^number=INC000056^category=Outage`, the query would become `short_description=<Encrypted(My sensitive information)>^number=INC000056^category=Outage` on the instance.

ParameterValue - toString()

Converts the POST or URL parameter value to a string.

Parameters

Name	Type	Description
None		

Returns

Type	Description
String	The parameter value as a string.

ParameterValue - getAsJsonContent()

Returns the request as an iterable object of type *JsonNode*.

This method is available only in an Edge Encryption rule if the request body is a valid JSON payload. If you are not sure what format the request body includes, check the contentType field on the *request* object.

Once the request is returned as a *JsonNode* object, you can use the JSON APIs to iterate over the object and encrypt fields.

Parameters

Name	Type	Description
None		

Returns

Type	Description
JsonNode	The request as an iterable *JsonNode*.

ParameterValue - getAsXmlContent()

Returns the request content as an iterable object of type *XMLContent*.

This method is available only in an Edge Encryption rule. This method assumes that the request body is a valid XML payload. You can check the contentType to make sure.

Once the request is returned as an *XMLContent* object, you can use the XML APIs to iterate over the object and encrypt fields.

Parameters

Name	Type	Description
None		

Returns

Type	Description
XMLContent	The request as an iterable object of type *XMLContent*.

ParameterValue - encodedQueryFor(String tableName)

Specifies that the value of the element is an encoded query on the specified table.

Calling this function on a parameter tells the proxy that the value of the parameter is an Encoded query strings for the specified table. The proxy parses the encoded query and encrypts the fields in the encoded query that must be encrypted.

Parameters

Name	Type	Description
tableName	String	The table that you expect the query to run on.

Returns

Type	Description
void	

ParameterValue - valueFor(String tableName, String fieldName)

Specifies that the value of the element maps to the specified field in the specified table.

Calling this method on an element value tells the proxy that the value for this element maps to the specified field in the specified table. The proxy then checks if the field must be encrypted.

Parameters

Name	Type	Description
tableName	String	The table name.
fieldName	String	The field name.

Returns

Type	Description
void	

XML APIs

XML APIs can be used after calling `getAsXmlContent()` on either the `request` object or a `ParameterValue` property.

When using XML APIs to write your encryption rule, you can follow a general format:

1. Call `getAsXmlContent()` on the `request` object or `ParameterValue` property. This returns an iterable object of the `XMLContent` underlying class.
2. Call `getIterator()` or `getIterator(String xPath)` on the `XMLContent` object. This returns an `XMLElementIterator` object that can be used to iterate over XML elements.
3. Call the `hasNext()` method on the `XMLElementIterator` object to determine whether another element is available.
4. Call `next()` on the `XMLElementIterator` object to return the next XML element. You cannot call `next()` without first calling `hasNext()`.
5. Call `valueFor(String tableName, String fieldName)` on the XML element. This method tells the proxy that the value for this element maps to the specified field in the specified table. The proxy then checks if the field must be encrypted.

> **Note:** To determine if you want to call `valueFor(String tableName, String fieldName)` on an XML element, you can use the `getName()` method to return the name of the element.

Mapping to a known table-field on the instance

In this example, the XML payload will be processed on the instance to insert records in the incident table. The description field will populate short_description on the incident.

```
<data>
    <record>
        <name>'Test Record 1'</name>
        <description>'Test Record 1 Description'</description>
        <tag>critical</tag>
    </record>
    <record>
        <name>'Test Record 2'</name>
        <description>'Test Record 2 Description'</description>
        <tag>security</tag>
    </record>
</data>
```

The following encryption rule action can apply:

```
function sampleXmlAction1() {
    var xmlContent = request.getAsXmlContent();
    // This loop iterates over all description tags that match
 the given path
    var xmlElementIterator =
 xmlContent.getIterator('data/record/description');
    while (xmlElementIterator.hasNext()) {
        var xmlElement = xmlElementIterator.next();
        xmlElement.valueFor('incident', 'short_decription');
    }
}
```

This action iterates through the **description** tags and asks the proxy server to encrypt the values and insert them into incident.short_description on the instance.

Note: This rule finds all **description** tags within all **record** tags in the XML payload. If there is only one occurrence of a tag to encrypt, the rule still uses the xPath and iterator structure. However, it iterates only once in the loop.

Mapping to an unknown table-field on the instance

In this example, the rule iterates over the **record** tags, but does not know what tags to expect within the **record** tag. The only known is that the tags within the **record** tags match the names of the columns specified in the table URL parameter.

The rule also specifies that, if the table is incident, then the data in the **description** tag should be encrypted and stored in the short_description field on the instance.

```
function sampleXmlAction2() {
    var xmlContent = request.getAsXmlContent();
    var tableName = request.urlParam.table;
    // This first iterator will iterate over all record elements
    var xmlElementIterator =
 xmlContent.getIterator('data/record');
    while (xmlElementIterator.hasNext()) {
        encryptFieldsInRecord(xmlElementIterator.next());
    }
}
function encryptFieldsInRecord(xmlElement) {
    //Then, iterate over all tags representing fields in the
 table
    var fieldIterator = xmlElement.getIteratorOverAllChildren();
```

```
    while (fieldIterator.hasNext()) {
        var field = fieldIterator.next();
        var fieldName = childElement.getName();
        //if table is incident, then description is encrypted
 for the short_description field
        if (tableName == 'incident' && fieldName ==
 'description') {
            field.valueFor(tableName, 'short_description');
        } else {
            //if table is not incident, ask the proxy to check
 if the given field is encrypted for the given table
            field.valueFor(tableName, fieldName);
        }
    }
}
```

In the *encryptFieldsInRecord()* function, the *valueFor()* method is called on a table and a field that are dynamically assigned based on the request. Even though the table and field names can change, the rule asks the proxy to check whether the field in the table must be encrypted based on the encryption configurations defined.

If the field is not configured for encryption, or if the tag does not match a field in the table, the proxy skips that tag. If the tag matches a field marked for encryption, then the Edge Encryption proxy server encrypts the value.

Using an encoded query

In this example, all tags have the **filter** attribute, which indicates whether the tag contains an encoded query.

```
<data>
    <record>
        <name filter="false">'Test Record 1'</name>
        <description filter="false">'Test Record 1
 Description'</description>
        <query filter="true">category=1^name=edge</query>
    </record>
    <record>
        <name filter="false">'Test Record 2'</name>
        <description filter="false">'Test Record 2
 Description'</description>
        <query filter="true">category=2^severity=3</query>
   </record>
</data>
```

The following encryption rule action can apply:

```
function sampleXmlAction3() {
   var xmlContent = request.getAsXmlContent();
   var tableName = request.urlParam.table;
   // This first iterator will iterate over all record elements
   var xmlElementIterator =
 xmlContent.getIterator('data/record');
   while (xmlElementIterator.hasNext()) {
       encryptFieldsInRecord(xmlElementIterator.next());
   }
}
```

```
function encryptFieldsInRecord(xmlElement) {
   //this time we want to iterate over all tags representing
 fields in the table
   var fieldIterator = xmlElement.getIteratorOverAllChildren();
   while (fieldIterator.hasNext()) {
       var field = fieldIterator.next();
       var fieldname = childElement.getName();
       //let's look at the filter attribute, if true, then
 encrypt as encoded query
       if (field.getAttributeValue('filter') == 'true') {
           field.encodedQueryFor(tableName);
       } else {
           //if it is false then check if the field should be
 encrypted
           field.valueFor(tableName, fieldName);
       }
   }
}
```

If the **filter** attribute value is true, the rule asks the proxy server to encrypt the values in the encoded query. If false, the rule asks the proxy to check whether the field should be encrypted.

XMLContent

A global object that provides methods to iterate over the XML content.

You can access an *`XMLContent`* object by calling getAsXmlContent() on a *`request`* object.

You access XML data in a POST or URL parameter by calling `request.postParams.<parameter name>.getAsXmlContent()` or `request.urlParams.<parameter name>.getAsXmlContent()`.

XMLContent - getIterator()

Returns an *`XMLElementIterator`* object for the XML content.

Parameters

Name	Type	Description
None		

Returns

Type	Description
XMLElementIterator	An object that can be used to iterate over elements in the *`XMLContent`* object.

XMLContent - getIterator(String xPath)

Returns an *`XMLElementIterator`* object for the XML content based on the specified parameter.

Parameters

Name	Type	Description
xPath	String	An xPath-like expression that specifies where in the `XMLContent` object to start.

Returns

Type	Description
XMLElementIterator	An object that can be used to iterate over elements in the `XMLContent` object.

XMLElementIterator

Provides methods for iterating over XML elements.

You get an `XMLElementIterator` object by calling the `getIterator()` method of the `XMLContent` class.

XMLElementIterator - hasNext()

Determines if there is another element available.

Parameters

Name	Type	Description
None		

Returns

Type	Description
Boolean	True if another element is available.

XMLElementIterator - next()

Returns the next element in the iterator.

You cannot call `next()` without first calling `hasNext()`.

Parameters

Name	Type	Description
None		

Returns

Type	Description
XMLElement	The next XML element.

XMLElement

Provides methods for iterating through XML elements and mapping values to fields in a table.

You get an *XMLElement* object by calling the *next()* method of an *XMLElementIterator* object.

XMLElement - getIterator(String xPath)

Returns an *XMLElementIterator* object for the XML element based on the specified parameter.

Parameters

Name	Type	Description
xPath	String	An xPath-like expression that specifies where in the *XMLElement* object to start.

Returns

Type	Description
XMLElementIterator	An object that can be used to iterate over elements in the *XMLElement* object.

XMLElement - getIteratorOverAllChildren()

Returns an *XMLElementIterator* object that includes all sub-elements for the XML element based on the specified parameter.

Parameters

Name	Type	Description
None		

Returns

Type	Description
XMLElementIterator	An object that can be used to iterate over elements in the *XMLElement* object.

XMLElement - valueFor(String tableName, String fieldName)

Specifies that the value of the element maps to the specified field in the specified table.

Calling this method on an element value tells the proxy that the value for this element maps to the specified field in the specified table. The proxy then checks if the field must be encrypted. If the table and field names are unknown, you can call the *valueFor()* method on a table and a field that are dynamically assigned based on the request.

Parameters

Name	Type	Description
tableName	String	The table name.
fieldName	String	The field name.

Returns

Type	Description
void	

XMLElement - encodedQueryFor(String tableName)

Specifies that the value of the element is an encoded query for the specified table.

Calling this function on an element tells the proxy that the value of the element is an Encoded query strings for the specified table. The proxy parses the encoded query and encrypts the fields in the encoded query that must be encrypted.

Parameters

Name	Type	Description
tableName	String	The table that you expect the query to run on.

Returns

Type	Description
void	

XMLElement - getName()

Returns the element name.

Parameters

Name	Type	Description
None		

Returns

Type	Description
String	The element name.

XMLElement - getAttributeValue(String attribute)

Returns the value of the specified attribute.

Parameters

Name	Type	Description
attribute	String	Attribute name.

Returns

Type	Description
String	The attribute value.

JSON APIs

JSON APIs can be used after calling `getAsJsonContent()` on either the `request` object or a `ParameterValue` property.

When using JSON APIs to write your encryption rule, you can follow a general format:

1. Call `getAsJsonContent()` on the `request` object. This returns an iterable object of the `JsonNode` underlying class.
2. Call `iterator()` or `getIterator(String xPath)` on the `JsonNode` object. This returns a `JsonNodeIterator` object that can be used to iterate over nodes in the JSON object.
3. Call the `hasNext()` method on the `JsonNodeIterator` object to determine whether another element is available.
4. Call `next()` on the `JsonNodeIterator` object to return the next JSON element. You cannot call `next()` without first calling `hasNext()`.
5. Call `valueFor(String tableName, String fieldName)` on the JSON element. This method tells the proxy that the value for this element maps to the specified field in the specified table. The proxy then checks whether the field must be encrypted.

> **Note:** To determine if you want to call `valueFor(String tableName, String fieldName)` on a JSON element, you can use the `getName()` method to return the name of the element.

Mapping to a known table-field on the instance

In this example, the JSON payload is processed on the instance to insert records in the incident table. The description field populates short_description on the incident.

```
{
    data: {
        records: [
            {
                "name": "Test Record 1",
                "description": "Test Record 1 Description",
                "tag": "security"
            },
            {
                "name": "Test Record 1",
                "description": "Test Record 1 Description",
                "tag": "security"
            }],
        "query":
 "assigned_to=3D4860165813e63a00d00abd322244b092^category=vulner
ability"
    },
    "source": "10.11.13.14"
}
```

The following rule can apply:

```
function sampleJsonAction1() {
    var jsonContent = request.getAsJsonContent();
    // This loop iterates over all description elements in the
 records array
```

```
    var jsonNodeIterator =
 jsonContent.getIterator(’/data/records/description’);
    while (jsonNodeIterator.hasNext()) {
        var jsonNode = jsonNodeIterator.next();
        jsonNode.valueFor('incident', 'short_decription');
    }
}
```

This action iterates through the **description** nodes and asks the proxy server to encrypt the values and insert them into incident.short_description on the instance.

Note: This rule finds all **description** nodes within the JSON payload. If there is only one occurrence of a node to encrypt, the rule still uses the xPath and iterator structure. However, it iterates only once in the loop.

Mapping to an unknown table-field on the instance

In this example, the rule iterates over **records**, but is not sure what nodes to expect. The only known is that for each object within **records**, the nodes match the names of the columns specified in the table URL parameter.

The rule also specifies that, if the table is incident, then the data in the **description** node should be encrypted and stored in the short_description field on the instance.

```
function sampleJsonAction2() {
    var jsonContent = request.getAsJsonContent();
    var tableName = request.urlParam.table;
    // This first iterator will iterate over all record elements
    var jsonNodeIterator =
 jsonContent.getIterator('data/records');
    while (jsonNodeIterator.hasNext()) {
            encryptFieldsInRecord(jsonNodeIterator.next());
    }
}
function encryptFieldsInRecord(jsonNode) {
    //this time we want to iterate over all nodes
    var fieldIterator = jsonNode.iterator();
    while (fieldIterator.hasNext()) {
            var field = fieldIterator.next();
            var fieldname = childElement.getName();
            if (fieldName == 'description') {
                field.valueFor(tableName, 'short_description');
            } else {
                field.valueFor(tableName, fieldName);
            }
    }
}
```

In the *encryptFieldsInRecord()* function, the *valueFor()* method is called on a table and a field that are dynamically assigned based on the request. Even though the table and field names can change, the rule asks the proxy to check whether the field in the table must be encrypted based on the encryption configurations defined.

If the field is not configured for encryption, or if the node name does not match a field in the table, the proxy skips that node. If the node name matches a field marked for encryption, then the proxy encrypts the value.

Using an encoded query

```
function sampleJsonAction3() {
    var jsonContent = request.getAsJsonContent();
    var tableName = request.urlParam.table;
    // This first iterator will iterate over all record elements
    var jsonNodeIterator = jsonContent.getIterator('data');
    while (jsonNodeIterator.hasNext()) {
        var jsonNode = jsonNodeIterator.next();
        if (jsonNode.getName() == 'records')
            encryptRecors(jsonNodeIterator.next());
        else if (jsonNode.getName() == 'query')
            jsonNode.encodedQueryFor(tableName);
    }
}
function encryptRecords(jsonNode) {
    //we iterate over all fields in the node
    var recordIterator = jsonNode.iterator();
    while (recordIterator.hasNext()) {
        encryptFieldsInRecord(recordIterator.next());
    }
}
function encryptFieldsInRecord(jsonNode) {
    //this time we want to iterate over all nodes
    var fieldIterator = jsonNode.iterator();
    while (fieldIterator.hasNext()) {
            var field = fieldIterator.next();
            var fieldname = childElement.getName();
            field.valueFor(tableName, fieldName);
    }
}
```

In this example, the rule iterates over **data**. As it finds **records**, it performs the same logic as in the second example, iterating over fields in each node. When it finds the **query** node, it calls *encodedQueryFor()* to encrypt values that should be encrypted in the query.

JsonNode

A global object that provides methods to iterate over the JSON content.

You can access a *JsonNode* object by calling getAsJsonContent() on a *request* object.

You access JSON content from a POST or URL parameter by calling `request.postParms.<parameter name>.getAsJsonContent()` or `request.urlParms.<parameter name>.getAsJsonContent()`.

JsonNode - getIterator(String xPath)

Returns a *JsonNodeIterator* object for the JSON content.

This method can only be used on the root node, but can be used to traverse deep into the JSON object. Subsequent traversals must use the *iterator()* method.

Parameters

Name	Type	Description
xPath	String	An xPath expression.

Returns

Type	Description
JsonNodeIterator	An object that can iterate over nodes in the JSON object.

JsonNode - iterator()

Returns a `JsonNodeIterator` object that iterates over all child nodes of the current node.

Parameters

Name	Type	Description
None		

Returns

Type	Description
JsonNodeIterator	An object that can iterate over nodes in the JSON object.

JsonNode - getAsString()

Returns the current node value as a string.

Parameters

Name	Type	Description
None		

Returns

Type	Description
String	The current node value.

JsonNode - getAsString(String propertyName)

Returns the string value of the specified property.

Parameters

Name	Type	Description
propertyName	String	Name of the property.

Returns

Type	Description
String	The property value.

JsonNode - getName()

Returns the name of the current JSON node.

Parameters

Name	Type	Description
None		

Returns

Type	Description
String	Name of the current JSON node.

JsonNode - valueFor(String tableName, String fieldName)

Specifies that the JSON property maps to the specified field in the specified table.

Calling this method on a JSON property tells the proxy that the value for this property maps to the specified field in the specified table. The proxy then decides if the field must be encrypted. If the table and field names are unknown, you can call the `valueFor()` method on a table and a field that are dynamically assigned based on the request.

Parameters

Name	Type	Description
tableName	String	The table name.
fieldName	String	The field name.

Returns

Type	Description
void	

JsonNode - encodedQueryFor(String tableName)

Specifies that the value of the JSON property is an encoded query for the specified table.

Calling this function on a JSON node tells the proxy that the value is an Encoded query strings for the specified table. The proxy parses the encoded query and encrypts the values for fields in the encoded query that must be encrypted.

Parameters

Name	Type	Description
tableName	String	The table that you expect the query to run on.

Returns

Type	Description
void	

JsonNodeIterator

You get a `JsonNodeIterator` object by calling the `getIterator()` or `iterator()` methods of the `JsonNode` class.

JsonNodeIterator - hasNext()

Determines if there is another property available.

Parameters

Name	Type	Description
None		

Returns

Type	Description
Boolean	True if another property is available.

JsonNodeIterator - next()

Returns the next property in the iterator.

You cannot call `next()` without first calling `hasNext()`.

Parameters

Name	Type	Description
None		

Returns

Type	Description
JsonNode	The next `JsonNode`.

print(String message)

Prints a message to the wrapper log file: `<proxy server directory>/logs/wrapper_<date>.log`.

This method is available only in an Edge Encryption rule action script.

Parameters

Name	Type	Description
message	String	The message to be written to the wrapper log file.

Returns

Type	Description
void	

Prohibited keywords

The Edge Encryption proxy validates encryption rule scripts before saving the rule. Many JavaScript keywords aren't allowed in encryption rule scripts.

Prohibited keywords

Keyword
__DIR__
__FILE__
__LINE__
__parent__
__proto__
Error
eval
getClass
getPrototypeOf
Java
javax
javafx
JavaImporter
load
loadWithNewGlobal
new
Packages
Object
prototype
RegExp
setPrototypeOf
this
throw

Edge Encryption dictionary attributes

Add dictionary attributes to tables and fields to control how they work with Edge Encryption.

To set a dictionary attribute to true, you must enter `attribute=true` in the **Attributes** field. To add a dictionary attribute to a record, see Altering tables and fields using dictionary attributes.

Edge Encryption Excluded [edge_encryption_excluded]

Determines whether the field is excluded from encryption.

When set to `true`, the field or table can't be encrypted. When set to `false`, the field can be encrypted.

- Value: true/false
- Target element: field or table
- Default value: false

Edge Encryption Enabled [edge_encryption_enabled]

Determines whether the field is eligible for encryption through an encryption configuration.

When set to `true`, the field is eligible for encryption. When set to `false`, the field isn't eligible for encryption. Because this attribute is used by the system and can't be modified, it isn't displayed to the user.

> **Note:** This attribute doesn't indicate that a field is encrypted, nor does it trigger any encryption logic on the field. Rather, the attribute determines the possibility of the field being encrypted by a user.

- Value: true/false
- Target element: field
- Default value: true for String fields

Edge Encryption Clear Text Allowed [edge_encryption_clear_text_allowed]

Determines whether server-side scripts may append non-encrypted data to an encrypted string within the field for user actions performed through the proxy server, or any server-side automated scripts, such as scheduled jobs.

When set to `true`, appending data is allowed. When set to `false`, appending data isn't allowed.

- Value: true/false
- Target element: field
- Default value: false

Domain separation and Edge Encryption

Domain separation is supported in limited circumstances with Edge Encryption. Edge Encryption provides the ability to encrypt data from within the customer's environment through the use of specific configurations, rules, and keys defined on the Edge Encryption proxy. The Edge Encryption proxy is not domain aware and cannot support domain-specific settings. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: No support

- The domain field may exist on data tables but there is no business logic to manage the data.
- This level is not considered domain-separated.

For more information on support levels, see Application support for domain separation.

How domain separation works in Edge Encryption

Edge Encryption can be used where domain-specific keys, configurations, and rules are not required.

Related topics

Domain separation for service providers

Data integration with Edge Encryption

To integrate third-party data with an instance using Edge Encryption, you must route the data through the Edge Encryption proxy server using supported integrations. Supported integrations use base system encryption rules that map data in each payload to fields in a table.

Upload data to fields marked for encryption

Edge Encryption does not support importing data from or exporting data to Excel, CSV, XML, or other file types to or from fields with encryption configurations defined.

ODBC driver

Encrypt requests and query data through the Edge Encryption proxy server using the ODBC driver.

Learn more: Edge Encryption ODBC driver integration

MID Server

You can configure the MID Server to route data through an Edge Encryption proxy server. However, some restrictions apply.

Learn more: Edge Encryption MID Server integration

REST/SOAP web services

Use REST/SOAP web services to update or retrieve record data through the Edge Encryption proxy server.

Learn more: Web services

JSONv2 web service

Use JSONv2 web service APIs to update or retrieve record data through the Edge Encryption proxy server. Base system encryption rules support data retrieval and data modification APIs.

- To insert a single record using the data modification API, use the `insert()` or `insertMutiple()` methods.
- To insert multiple records using the data modification API, use the `insertMultiple()` method.

Learn more: JSONv2 Web Service

To encrypt data from custom third-party integrations not listed above, create custom encryption rules. See Define a custom encryption rule.

Upload attachments to records marked for encryption

Attachments can be uploaded to tables with attachment encryption configured using REST and SOAP web services.

Edge Encryption ODBC driver integration

Configure your ODBC driver to query data encrypted by Edge Encryption. The Edge Encryption proxy server encrypts ODBC driver requests to the ServiceNow instance when Edge Encryption is integrated with the ODBC driver.

Encrypted responses from the instance are decrypted through the Edge Encryption proxy server before passing to the ODBC driver in your network.

For a successful integration, the ODBC driver must trust the Edge Encryption proxy server certificate. If the Edge Encryption proxy server certificate is signed by a Certificate Authority trusted by the ODBC driver, the Edge Encryption proxy server is automatically trusted. However, if a Certificate Authority trusted by the ODBC driver has not signed the Edge Encryption proxy server certificate, you must import the self-signed certificate to the ODBC truststore.

Related topics

ODBC driver

Import a self-signed certificate to the ODBC truststore

If a Certificate Authority trusted by the ODBC driver has not signed the Edge Encryption proxy server certificate, you must import a self-signed certificate to the ODBC truststore. You can export the certificate from the Edge Encryption proxy server and import it into the ODBC truststore.

Before you begin
Role required: admin

To determine whether a Certificate Authority trusted by the ODBC driver has signed the Edge Encryption proxy server certificate, run the following command in the keystore directory in the proxy home directory to view a list of Certificate Authorities trusted by the ODBC driver:

```
keytool -keystore "<ODBC
 directory>\ip\Java\jre\lib\security\cacerts" -list
```

Note:

In most cases, the client connects to the Edge Encryption proxies through a load balancer, so the certificate would be the certificate configured in the load balancer.

In the case where there is no load balancer between the client and the Edge proxy server, the certificate is the certificate presented by the Edge proxy and is configured in the *edgeencryption.properties* file.

```
edgeencryption.proxy.https.port = 9090
edgeencryption.proxy.https.keystore.path = keystore/keystore
edgeencryption.proxy.https.keystore.password = password
edgeencryption.proxy.https.cert.alias = jetty
```

For details on editing properties, see Configure additional properties in the Edge Encryption properties file

Procedure

1. Change to the keystore directory in the proxy home directory.
2. Check the keystore for the self-signed certificate.

   a. To check the keystore for the certificate, you can run the following command to list all the items in the keystore.

   ```
   keytool -list -keystore keystore.jceks -storetype jceks -v
   ```

   b. Locate the key alias in the list of items.

3. Using the key alias, export the certificate to a .cer file.

   ```
   keytool -export -alias <key alias> -keystore keystore.jceks
    -storetype jceks -rfc -file <file name>.cer
   ```

4. Change to your ODBC truststore directory: `ODBC\ip\Java\jre\lib\security\cacerts`.
5. Import the certificate to your ODBC truststore.

   ```
   keytool -keystore cacerts -importcert -alias $<key alias> -file
    <file name>.cer
   ```

Set the ODBC driver properties

Set the ODBC driver properties to route requests through the Edge Encryption proxy server.

Before you begin

Role required: admin

Procedure

1. In Windows, navigate to **Start** > **Programs** > **ServiceNow ODBC Management Console**.
2. Expand the console tree root to: `ServiceNow ODBC Manager\Manager\<installation location>\Services\ServiceNow_ODBC\Data Source Settings\ServiceNow\IP Parameters`.
3. Double-click the **DataSourceIPProperties** attribute.
4. Change the **Value** to the URL of your Edge Encryption proxy server, such as `https://<IP address>:<port>`

DataSourceIPProperties

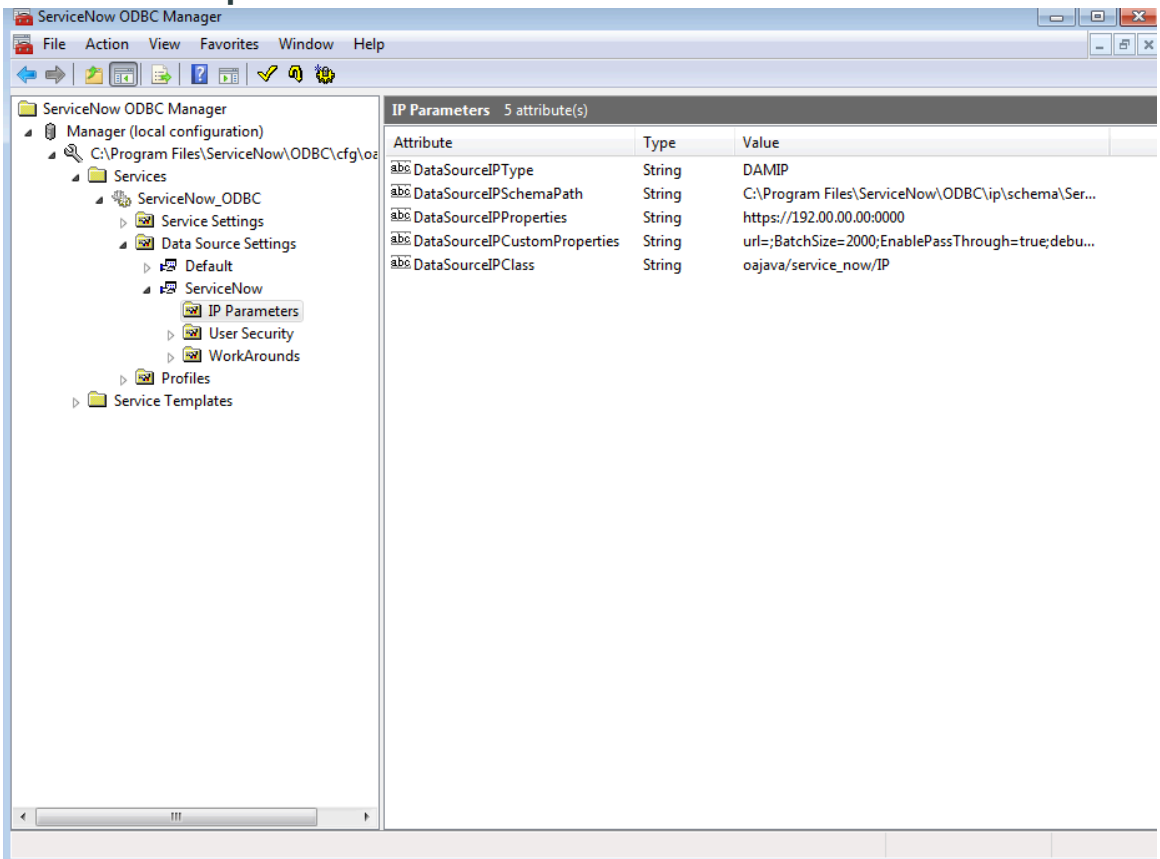

5. Click **OK**.

What to do next

The ODBC driver is now configured to route requests to the instance through the Edge Encryption proxy server.

Edge Encryption MID Server integration

Configure the MID Server to route data through an Edge Encryption proxy server.

When integrated with the MID Server, the Edge Encryption proxy server acts as the MID Server's endpoint. The Edge Encryption proxy server then encrypts and decrypts data passing between the ServiceNow instance and the MID Server.

Limitations when integrating with the MID Server

When MID Server data is configured to pass through the Edge Encryption proxy server, the following limitations apply:

- Encryption of ECC Queue fields is not supported.
- Encrypted data cannot be used with Discovery or Service Mapping.

Point the MID Server to the Edge Encryption proxy server

To pass data from the MID Server through the Edge Encryption proxy server, update the MID Server configuration file to point the MID Server to the Edge Encryption proxy server.

Before you begin

Role required: admin

About this task

When configuring the MID Server to pass through the Edge Encryption proxy server, you cannot use the web proxy properties in the MID Server configuration file to route traffic through the Edge Encryption proxy server to your instance. Instead, you must set the Edge Encryption proxy server as the MID Server's endpoint.

Procedure

1. Navigate to your local MID Server directory and open the `config.xml` file.
2. Find the element `<parameter name="url" value="https://YOUR_INSTANCE.service-now.com" />` and change the value property to the URL of your Edge Encryption proxy server.
   For example, `http://hostname.mycompany.com:8081`
   This step directs the MID Server to pass traffic to the Edge Encryption proxy server instead of the instance. The Edge Encryption proxy server in turn encrypts any necessary fields and passes the payload to the instance.
3. Save and close the file.
4. If running, restart the MID Server.

Edge Encryption diagnostics and performance

Monitor Edge Encryption proxy server performance trends and drill into errors generated by the Edge Encryption proxy server.

Edge proxy performance

View key Edge Encryption proxy server performance trends using the Edge Proxy graph set on the ServiceNow Performance homepage. Monitored trends include:

- Maximum and average response times between the client, proxy server, and instance.
- CPU, disk space, and memory usage of the host machine.
- Maximum and average network latency between the proxy server and the ServiceNow instance.

Note: Edge Encryption proxy servers with duplicate names do not report performance trends.

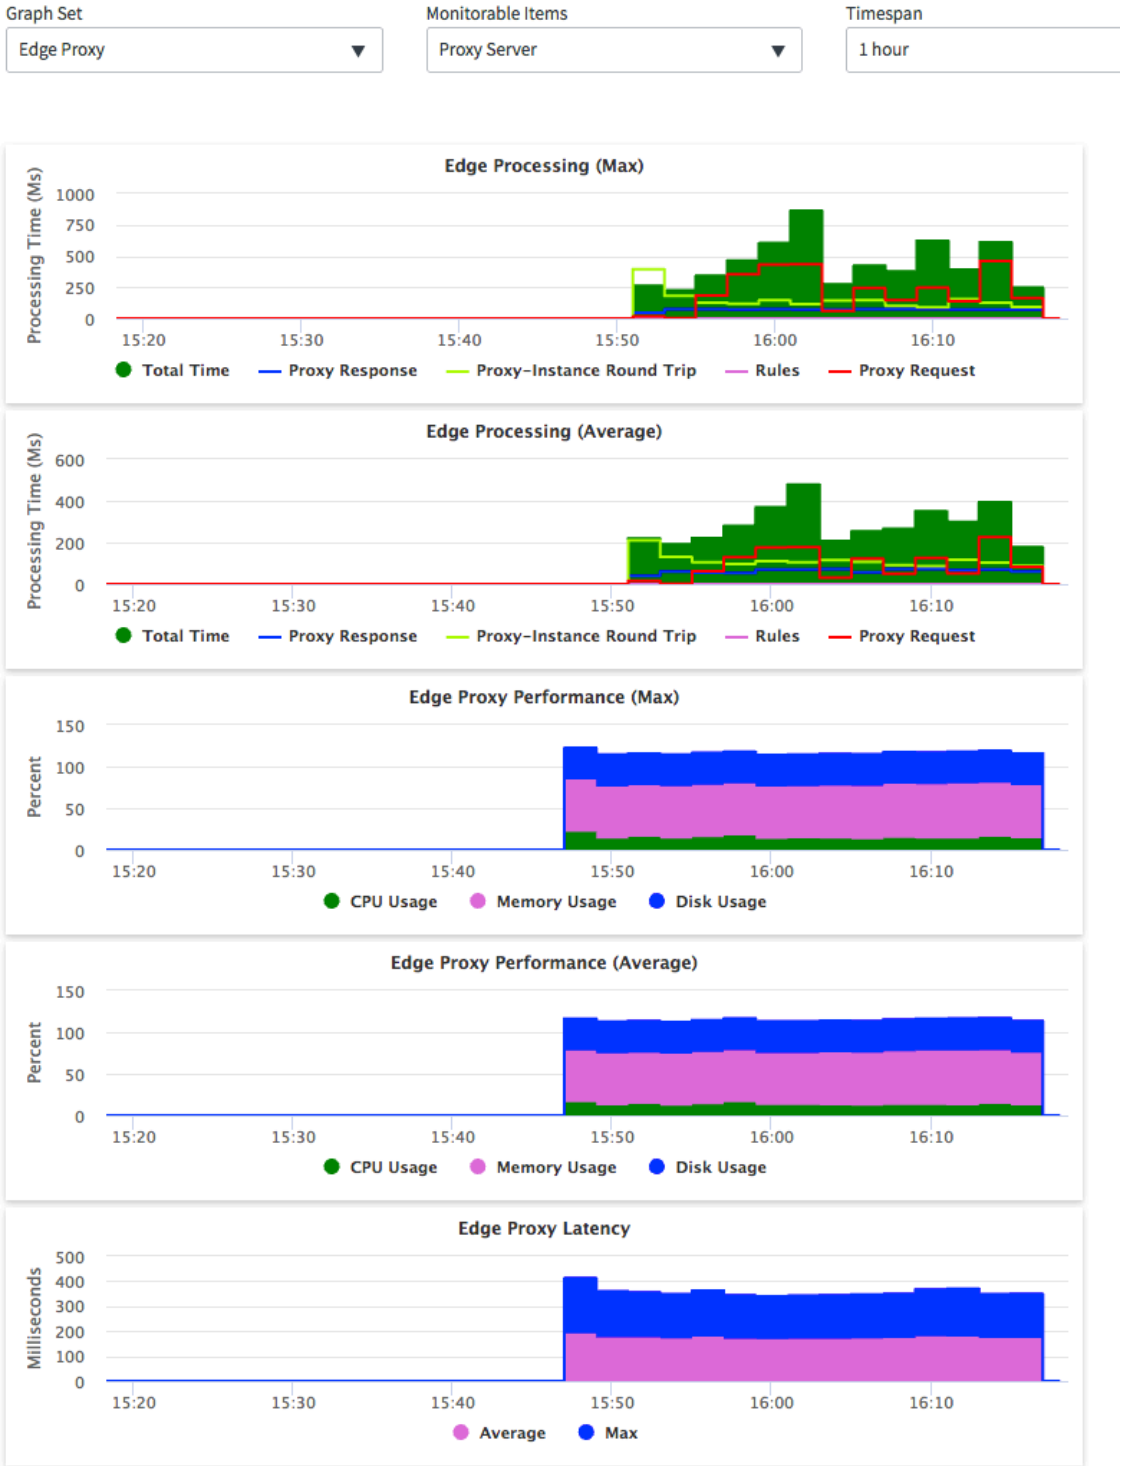

Edge Processing (Max and Average)

Maximum and average time in milliseconds to process a request. These data points are general trends over time.

- **Total Time**: Time for the proxy server to receive a request from a client and send a response. This data point is the sum of the subsequent data points.
- **Proxy Response**: Time for the proxy server to process a response from the instance.

- **Proxy-Instance Round Trip**: Time for the proxy server to send a request to the instance and receive a response. Includes network latency between the proxy server and the instance and time spent by the instance to process the request.
- **Rules**: Time for the proxy server to evaluate a request using defined encryption rules.
- **Proxy Request**: Time for the proxy server to process a client request and forward it to the instance.

Edge Proxy Performance (Max and Average)

Maximum and average percentage of resources used on the host machine.

- CPU Usage
- Memory Usage
- Disk Usage

Edge Proxy Latency

Maximum and average network latency in milliseconds at a given point in time. Latency is determined by round-trip time for a proxy server to send a simple ping to the instance and receive a response.

Proxy Error Reports

Navigate to **Edge Encryption Configuration** > **Diagnostics and Troubleshooting** > **Proxy Error Reports** to view all proxy server errors collected over the past seven days.

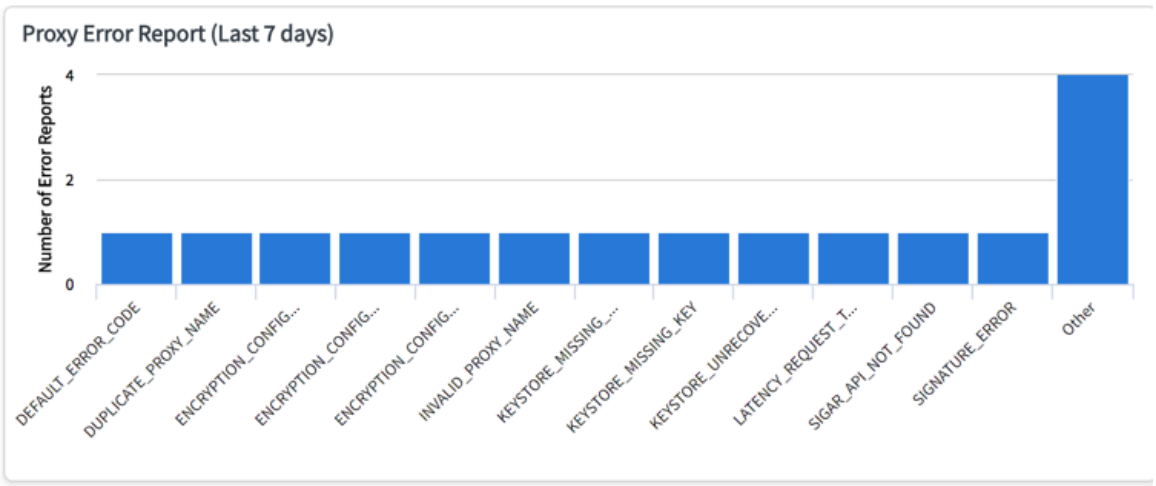

Errors are collected over a one-minute period. Each minute, an error report is generated. The vertical axis displays the number of error reports over the last seven days that include each error. For example, even if the DEFAULT_ERROR_CODE error is thrown multiple times over a one-minute report period, the DEFAULT_ERROR_CODE bar will only reach one on the **Number of Error Reports** axis.

From this view, you can:

- Click each proxy error code bar to see the report on a single error for each proxy server. From this view, you can click the bar again to view the error text in the Edge Encryption Proxy Stat table [edge_encryption_stat]. Follow links in the error text to see more information and possible remediation steps.
- Click **Other** to see page two of the error report.

Note: If you have more than one proxy server with the same name, a single DUPLICATE_PROXY_NAME error appears in the Proxy Error Report. No other errors are reported for proxy servers with duplicate names. If you encounter this error, make sure that all proxy servers have unique names.

Additional monitoring resources

The instance tracks all encryption proxies. Each Edge Encryption proxy server registers when it starts up. The instance is notified when:

- A new Edge Encryption proxy server starts up.
- An Edge Encryption proxy server is intentionally shut down.

If an Edge Encryption proxy server attempts to register with an instance that does not have Edge Encryption installed, the proxy does not start.

All encryption configuration files are audited. Deleted records are audited on all encryption configuration files. Audit records are put in the sys_audit table. To view the history of a specific configuration record, view the record, and click **History** > **List** in the menu. The Mass Encryption Job is not audited.

Use the following additional resources to monitor you proxy servers.

Table	Description
Invalid Insert Attempts [sys_edge_encryption_invalid_insert_log]	List of attempts to save the following data to encrypted fields: • Unencrypted data. • Data that did not come from an Edge Encryption proxy. The instance rejects and then logs any attempts to save this data. If you have the security-admin role, you can view the logs in the Invalid Insert Attempts list.
Job Failures [sys_encryption_job_execution]	A list of jobs that did not execute successfully.
System logs	The instance periodically checks for messages from each registered proxy server. If a proxy server has not sent a message in the required time frame, an error is logged. The log message contains information about the encryption proxy and the last time the proxy pinged the instance. If the instance determines that none of the encryption proxies are online, it logs a message. These messages are added to the system log.

Disable or reduce Edge Proxy statistic collection

Prevent the Edge Encryption proxy server from sending Edge Proxy Graph Set statistics to the ServiceNow Performance homepage, or reduce the frequency of statistic collection.

Before you begin
Role required: admin or security_admin

About this task
By adding properties in the `edgeencryption.properties` configuration file, you can:

- Disable the Edge Proxy graph set.
- Change the interval during which statistics are collected by the Edge Encryption proxy server. By default, statistics are collected every 30 seconds.

Procedure

1. In your proxy server installation directory, open the `edgeencryption.properties` configuration file located in the `<installation directory>/conf/` folder.
2. Add one of the Edge Encryption proxy server properties.
3. Restart the proxy server.

Increase debug logging for the Edge Encryption proxy

Increase the level of logging to interpret the logs and debug issues with the proxy.

There are currently three options for increasing debug logging in the Edge Encryption proxy. Increase the level of logging to debug issues provide technical support with information to look into the issue with the benefit of more verbose log statements.

Depending on the issue being debugged, set up debug logging in one of three ways:

- Debugging issues other than SSL connectivity
- Logging timing metrics for requests through the proxy
- Debugging issues with SSL connectivity between the Edge Encryption Proxy and the instance

For all debug cases, you may view and interpret the logs in your own or open an incident to get an interpretation from ServiceNow technical support providing the description of the issue and how it's reproduced.

Debugging issues with the Edge Encryption application other than SSL connectivity

Use this method to debug issues with the Edge Encryption application, without stopping and restarting the proxy. These steps increase logging level and help troubleshooting the root cause with more verbose log statements.

Before you begin
Role required: admin

Note: Changes made to the `$proxy_installation_location/conf/log4j2.properties` file are taken up by the proxy within about 60 seconds after you make your changes. You don't have to restart the proxies.

Procedure

1. In the `$proxy_installation_location/conf/log4j2.properties` file find the following line.

```
logger.edge.level=info
```

2. Change the above line to the following:

```
logger.edge.level=debug
```

3. Save the change.
   It may take up to 60 seconds for the change to take effect, but this doesn't require a proxy restart.
4. Reproduce your issue.
5. Check for debug log statements related to the application in the `$proxy_installation_location/logs/edgeencryption.log` file.

Result

After making the property change, you can see additional detail in your `$proxy_installation_location/logs/edgeencryption.log` file. When you have finished debugging, revert the change made to the `$proxy_installation_location/conf/log4j2.properties` file.

Logging timing metrics for requests through the proxy

Enable timing metric logging to add a metric statement for each request handled by the Edge Encryption proxy. Each of these timing metric log statements contains useful information about the request, such as processing times and which encryption rule was used.

Before you begin

Role required: admin

Note:

The additional logging settings are added to the `$proxy_installation_location/conf/log4j2.properties` file. Changes made are taken up by the proxy dynamically within about a minute after the changes to the file are made, so you do not have to restart the proxies.

Procedure

1. Modify the `$proxy_installation_location/conf/log4j2.properties` file by adding the following lines at the end of the file

```
appender.timinglog.type=RollingFile
appender.timinglog.name=TimingLog
appender.timinglog.fileName=../logs/edgenetwork.log
appender.timinglog.filePattern=../logs/$
${date:yyyy-MM}/edgenetwork-%d{yyyy-MM-dd-HH}-%i.log.gz
appender.timinglog.layout.type=PatternLayout
appender.timinglog.layout.pattern=%d [%t] %-5p %m%n
appender.timinglog.policies.type=Policies
appender.timinglog.policies.size.type=SizeBasedTriggeringPolicy
appender.timinglog.policies.size.size=500MB
appender.timinglog.strategy.type=DefaultRolloverStrategy
appender.timinglog.strategy.max=4
```

```
logger.timing.name=com.snc.edgeencryption.metrics.EdgeEncryptio
nTimingMetricCache
logger.timing.level=debug
logger.timing.additivity=false
logger.timing.appenderRef.rolling.ref=TimingLog
```

2. Save the file.

Result

After the `log4j.properties` file is saved, the following types of messages appear in the `$proxy_installation_location/logs/edgenetwork.log` log file for network times.

```
2022-07-21 12:56:15,783 [qtp1971991758-7700] DEBUG
 com.snc.edgeencryption.metrics.EdgeEncryptionTimingMetricCache
 -
 request_uri=/api/now/ui/presencesysparm_auto_request=true&cd=16
58433375754 request_method=POST
 client_request_received="2022-07-21 12:56:15,015"
 proxy_request_processing_time=6 all_rules_processing_time=0
 rule_executed="REST JSON" rule_execution_time=1
 proxy_instance_round_trip=14 proxy_response_processing_time=1
 total_time_from_proxy=21 reponse_code=201
 glide_user=SCv3_1:BAz1ZK7ee9XoroG2nvMlixHpgTvsT4fY2bwQvnH2WdU=:
y5HGsTTqo3Pjq6GOxk4LoazCwCiWRJk4/6SpbXuBzqg=:6816f79cc0a8016401c
5a33be04be441 jsessionid_suffix=037A66
```

The values in the log messages are as follows:

```
request_uri: The URI being requested

request_method: The HTTP method being used, for example, GET,
 POST, PUT, PATCH, DELETE

client_request_received: The timestamp noting when the HTTP
 client request arrived at the Edge proxy

proxy_request_processing_time: How long the Edge proxy took to
 process the request in milliseconds

all_rules_processing_time: Total time it took to execute all of
 the Edge Encryption rules for the request in milliseconds

rule_executed: The name of the encryption rule that was executed

rule_execution_time: How long it took to execute listed
 rule_executed in milliseconds

proxy_instance_round_trip: The time from when the Edge proxy
 sent the request to the instance until the instance sent the
 response and was received by the edge proxy in milliseconds

proxy_response_processing_time: How long the Edge proxy took to
 process the response in milliseconds
```

```
total_time_from_proxy: The total time from when the Edge proxy
 received the request from the client and returned the response
 to the client in milliseconds

response_code: HTTP response code

glide_user: The glide_user cookie value

jsessionid_suffix: The JSession cookie suffix associated with
 the request
```

Debug issues with SSL connectivity between the Edge Encryption proxy and the instance

Use this method to debug issues with SSL connectivity between the Edge Encryption proxy and your instance, such as access to the instance fails via the proxy. These steps increase logging and help find the verbose log statements.

Before you begin

Role required: admin

Note: SSL connectivity debugging is only relevant when troubleshooting TLS connectivity type issues. In practice, this is not common and rarely needed.

Procedure

1. Stop the proxy server.
2. Add the following line to the file `$proxy_installation_location/conf/wrapper.conf`, which is a Java startup property:

   ```
   wrapper.java.additional.<next available number in sequence> =
    -Djavax.net.debug=all
   ```

   For example:

   ```
   For example: wrapper.java.additional.4 = -Djavax.net.debug=all
   ```

3. Save the change and restart the proxy server.
4. Reproduce your connectivity issue.

Result

After reproducing the issue debug log statements related to the SSL exchange can be found in the `$proxy_installation_location/logs/wrapper_<current date>.log` file. When you are finished debugging. You can remote the additional logging by removing or commenting out the line created in the previous steps.

Database Encryption

ServiceNow® offers database encryption (DBE) and full-disk encryption methods for customers with statutory obligations for data protection which may require at-rest protection for all data.

Important: Starting with the Washington DC release, Database Encryption is being prepared for future deprecation. Cloud Encryption is the replacement solution for data at rest encryption. For details, see Cloud Encryption with Key Management

Explore	Request	Reference
Learn the key features and business value of Database Encryption.	Learn more about how to request database key rotation.	Database Encryption with Customer-Controlled Switch (DBE-CCS) is an encryption solution that encrypts all data-at-rest when not in use in the database.

Exploring Database Encryption

ServiceNow® offers database encryption (DBE) and full-disk encryption methods for customers with statutory obligations for data protection which may require at-rest protection for all data.

Important: Starting with the Washington DC release, Database Encryption is being prepared for future deprecation. Cloud Encryption is the replacement solution for data at rest encryption. For details, see Cloud Encryption with Key Management

Database Encryption enables all data to be protected with symmetric AES-256 encryption, whether the database is online or offline. From the ServiceNow AI Platform perspective, all data flows in decrypted.

- Database Encryption supports all stored data to be encrypted in real time providing protection for data online and offline with no loss of functionality.
- Full disk encryption protects offline data if there is disk loss or theft.

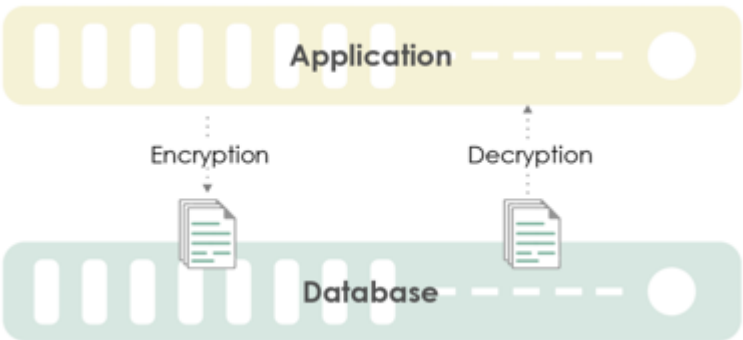

With Database Encryption, all stored data is encrypted and individual records or tables are decrypted in memory while being accessed. New or changed data is encrypted as it is entered into a table and associated activity log files (bin, redo, undo, and error) are also encrypted.

Database Encryption is transparent to users, with no loss of functionality. When using this feature, all instances are encrypted, along with replication traffic and backups. Instance cloning is still available with a minor performance impact for using Database Encryption of up to 5%. Both new and existing instances on supported releases of the ServiceNow AI Platform can take advantage of database encryption.

As illustrated, ServiceNow stores and manages keys using a three-level key hierarchy:

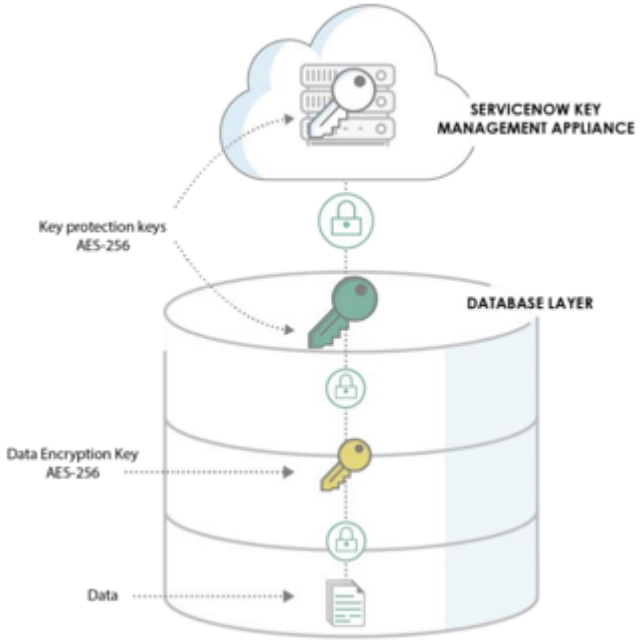

1. A customer specific AES-256 key is created by the database engine and is used to encrypt the data.
2. A second customer specific AES-256 key is created by the database engine and is used to protect the first-level key.
3. A third AES-256 key is created by and stored within FIPS 140 validated key management appliances in the ServiceNow datacenters. This key protects the second-level key and is unique per customer instance.

The ServiceNow AI Platform also supports database encryption with a customer supplied switch, DBE with CCS. This is an encryption solution that encrypts all data-at-rest when not in use in the database. It uses industry standard AES encryption with no impact to functionality. The database encrypts data as it is written to the disk, and decrypts data as it is read from the disk. That means that applications always have the data in an unencrypted state to perform the necessary logic and functions without impact.

Note: Database Encryption is not supported for on-premise instances.

If you are using your own keys for database encryption, see Database Encryption with Customer Controlled Switch.

Requesting database key rotation

Rotate the database key annually or as needed by submitting a request to support.

Before you begin
Role required: admin

Important: Starting with the Washington DC release, Database Encryption is being prepared for future deprecation. Cloud Encryption is the replacement solution for data at rest encryption. For details, see Cloud Encryption with Key Management

About this task

Key rotation occurs at night within 24 hours prior to the expiration date and does not interrupt service for the instance.

Note: Currently, key rotation is only available in ServiceNow Commercial, Government Community Cloud (GCC), France, and Singapore environments.

Procedure

Contact Customer Service and Support to request any of the following key rotation actions:

- Enroll to perform annual key rotation on all designated instances.
- Obtain a historical report of the last three key rotations containing the following:
  - Instance name.
  - Key name and version.
  - Dates and times of rotation.
- Schedule an early key rotation outside of the annual scheduled rotation.

Database Encryption with Customer Controlled Switch

Database Encryption with Customer-Controlled Switch (DBE-CCS) is an encryption solution that encrypts all data-at-rest when not in use in the database.

Important:

Database Encryption with Customer Controlled Switch has begun its End of Life process, and has reached the End Of Sale and End of Renewal milestones as of the Yokohama release. For data-at-rest encryption support, see Cloud Encryption with Key Management.

Overview

Database Encryption with customer controlled switch uses industry standard AES encryption, with no impact to functionality. The database encrypts data as it is written to the disk and decrypted by the database as it is read from the disk. Applications always have the data in an unencrypted state to perform the necessary logic and functions.

DBE-CCS utilizes technology native to the database, often called Tablespace Encryption or Transparent Data Encryption. For more details on the technology, refer to the MariaDB website under "Tablespace Encryption."

DBE-CCS requires you to set up an HTTPS REST service endpoint that periodically provides the secret key to the ServiceNow instance. The CCS endpoint then returns the customer secret key encrypted with the public key of the database instance.

Customer endpoint

Important: Your organization is solely responsible for setting up and maintaining your CCS endpoint. The customer endpoint specification is provided in KB0789788.

A ServiceNow technology partner, Fortanix, is available to implement your customer endpoint for you. Contact the technology partner directly for details of the integration. For details, see Using Fortanix Data Security Manager with ServiceNow.

Multiple ServiceNow version support

Important: Database Encryption is a paid infrastructure offering that is release agnostic. It can be applied to any supported release and to new or existing instances.

Other references

Refer to these references for additional information about DBE with CCS:

Reference	Description
KB0993681	Architecture of Database Encryption Customer Controlled Switch
KB0789788	Implementation guide for DBE with CCS

Note: To access KB articles, you must first authenticate into Now Support.

servicenow

Access Management

Access Management enables you to have access to ServiceNow® instance securely.

Zero Trust Access

Zero Trust Access ensures that all access to applications and data is granted on a least privilege basis, only after the user's identity verification and risk assessment.

Access Analyzer

Access Analyzer is a ServiceNow® Store App which is an access diagnostic tool. It helps to determine who has access to a resource.

Authentication

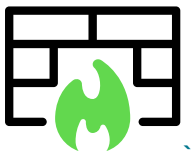

ServiceNow's authentication validates the identity of a user who accesses an instance, and then authorizes the user to features that match the user's role or job function.

ACL

Rules for access control lists (ACLs) restrict access to data by requiring users to pass a set of requirements before they can interact with it.

Data Filtration

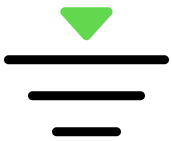

Use data filtration to control access to tables and records based on subject attributes when performing read queries.

Security Roles

Security Roles provide added security, every user must have at least one role so that the instance can distinguish between internal and external users.

Connections and Credentials

ServiceNow Access Control

Domain Separation

Credentials and connection information are required to gain access to a computer or network device for Discovery, Service Mapping, and Cloud Management or to perform work using Orchestration.

The SNC Access Control plugin enables you to control which Customer Service and Support employees can access your instance, and when.

With the ServiceNow AI Platform, service providers (SPs) can provide their customers with faster onboarding, meet compliance, and protect their data using domain separation.

Zero Trust Access (ZTA)

Zero Trust Access (ZTA) is a security model that assumes no user or device is trusted by default.

Explore

Learn the features and business value of Zero Trust Access.

Activate

Understand how to Activate Zero Trust Access.

Configure

Configure Zero Trust Access.

Reference - Properties

Know about the properties in Zero Trust Access.

Explore Zero Trust Access

Zero Trust Access (ZTA) is a security model that assumes that no user or device is trusted by default.

ZTA ensures that all access to applications and data is granted on a least privilege basis, only after the user's identity verification and risk assessment.

Zero Trust - Policy Based Session Access

ServiceNow Zero Trust - Policy Based Session Access (Session Access) enables organizations to dynamically reduce user privilege in a web session based on a variety of factors, including IP address, location, authentication method, user’s role, group, user having MFA and attributes shared by the Identity Provider (IDP). This can help protect organizations from unauthorized

access and data breaches, even when high-privileged users access applications from untrusted devices or locations.

It enables the security admins to reduce or limit user access in a session based on IP address, location, Identity Provider attributes, and user attributes using adaptive authentication policies.

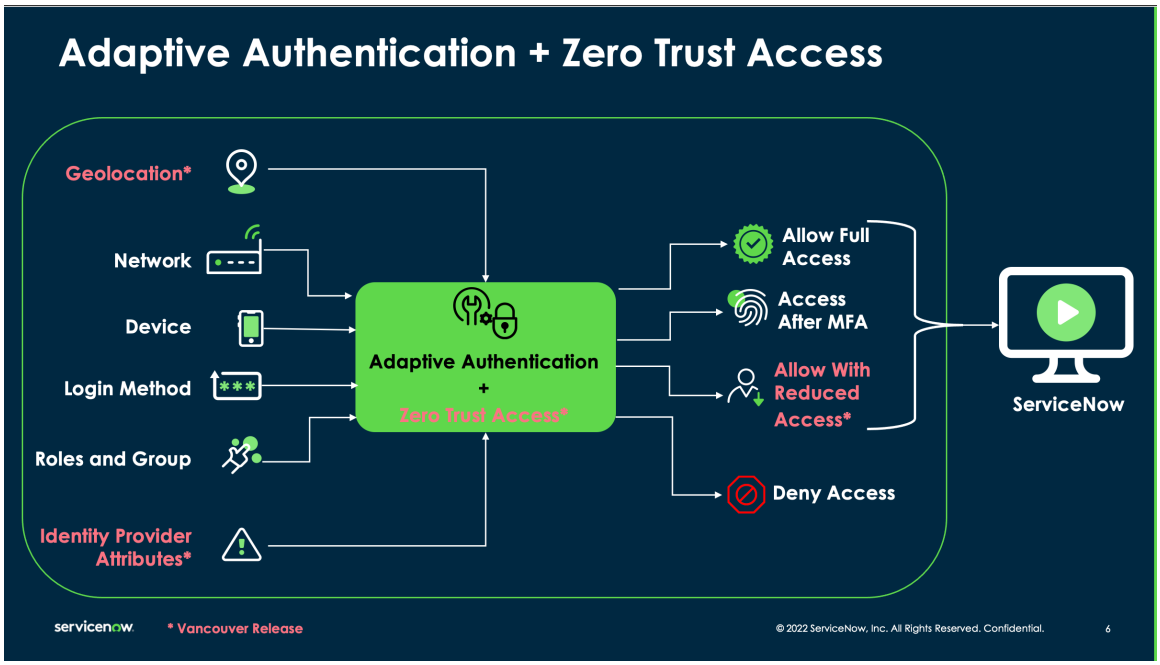

Note:

- Session Access configurations can only be performed with `security_admin` role. You must elevate your role to `security_admin`.
- Session Access doesn't support integrations.
- Session Access has no impact if the reduced or limited role isn't assigned to a user. In this case, there are no changes to the logged in session. The user continues to access the instance with the assigned privileges.
- Session Access has no impact while the user is already logged in to the instance and simultaneously the admin configures the policy. The user has to log out from the session for the policy to be effective.
- Session Access has no impact when the user is in a trusted network and later switches to a VPN (change in location or network) within a session.
- Session Access is enforced at the time of login. Any change in risk parameters during the session won't result in reduced access. For example, a user switching from the corporate network to an untrusted network after establishing the session won't result in reduced access unless the user logs out and logs in again.
- Session Access (Zero trust access - ZTA) feature, roles like `snc_internal` and `snc_external` cannot be removed.
- Session Access (Zero trust access - ZTA) feature does not remove a role from the `sys_user_has_role` or the user group membership table. Based on the ZTA policy, it establishes the user session with reduced or limited roles.
- The scripts running in the system context will not honor the ZTA session roles.

Use case

Following are some of the use cases of Zero Trust Access:

- Reduce privileges based on the risk associated with the session. For example, a fulfiller role user logging from outside the trusted network can be configured to have only the requester role for the session.
- Reduce access based on IDP response for a user session, if the user is using an untrusted device. For more information, see Configure Identity Provider attribute for Session Access.

This role relegation ensures that the user doesn't have any other existing privileges in a session. When the user is logging in from a trusted network, all the existing privileges are assigned for a session.

Multiple IP conditions and multiple role or group assignments can be defined as part of the policy.

Zero Trust Access - Mobile

You can use the Zero Trust Access - Session Access policy within the Adaptive Authentication policy to reduce the roles or privileges of the particular session in mobile.

Zero Trust Access - Session Access mobile can be enabled by enabling the **glide.authenticate.session_access.mobile.enabled** from the system properties table.

To use Zero Trust Access - Session Access mobile with the IDP attributes you can configure the **glide.authenticate.session_access.mobile.refresh_token_interval** field. This enables the administrators to effectively control the session access based on refresh token.

For more information, see Configure Zero Trust Access for mobile.

Activate Zero Trust Access

Activate the **Zero Trust - Policy Based Session Access** `com.snc.zero_trust_session_access` plugin to enable security admins to reduce or limit user access in a session based on IP address, location, Identity Provider attributes, and user attributes using adaptive authentication policies.

Before you begin

Role required: admin

Plugin type: Paid and requires license.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the **Zero Trust - Policy Based Session Access** (`com.snc.zero_trust_session_access`) plugin using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Configure Session Access role

Configure Session Access to reduce user access in a session based on IP, location, Identity Provider attributes, and user attributes using adaptive authentication policies.

Before you begin

Role required: security_admin

Note:

- Session Access configurations can only be performed with `security_admin` role. You must elevate your role to `security_admin`.
- Session Access doesn't support integrations.
- Session Access has no impact if the reduced or limited role isn't assigned to a user. In this case, there are no changes to the logged in session. User will still continue to access the instance with their assigned privileges.
- Session Access has no impact while the user is already logged in to the instance and simultaneously the admin configures the policy. The user has to log out from the session for the policy to be effective.
- Session Access is enforced at the time of login. Any change in risk parameters during the session won't result in reduced access. For example, a user switching from the corporate network to an untrusted network after establishing the session, won't result in reduced access unless the user logs out and logs in again.
- Session Access (Zero trust access - ZTA) feature, roles like `snc_internal` and `snc_external` cannot be removed.
- Session Access (Zero trust access - ZTA) feature does not remove a role from the `sys_user_has_role` or the user group membership table. Based on the ZTA policy, it establishes the user session with reduced or limited roles.
- The scripts running in the system context will not honor the ZTA session roles.

Procedure

1. Navigate to **All** > **Zero Trust Access** > **Session Access Role Configurations**.
2. To create a Session Access role configuration, select **New**.
3. On the form, fill the fields:

Session Access Role Configuration

Field	Description
Name	Name of the configuration
Description	Short description of the configuration.

Field	Description
Policy	Choose the adaptive authentication policy. Use the look-up icon to view the list of policy.
Action	**Remove Roles** or **Limit to Roles**. ○ **Remove Roles**: When the configured user logs in, the list of roles provided in the Role or Group List are removed for the session. ○ **Limit To Roles**: When the configured user logs in, only the selected roles are provided to the user and all other roles are removed for the session.
Role List	Choose roles from the Role List.
Group List	Choose the roles from the Group List that you want to remove or limit to the user.

4. Select, **Submit**.

   The login for users based on the configured countries is as follows:

   - In **Remove Roles**, the users from the configured countries with the selected roles no longer have those roles for the session.
   - In **Limit To Roles**, the users from the configured countries with the selected roles only have those roles for the session.

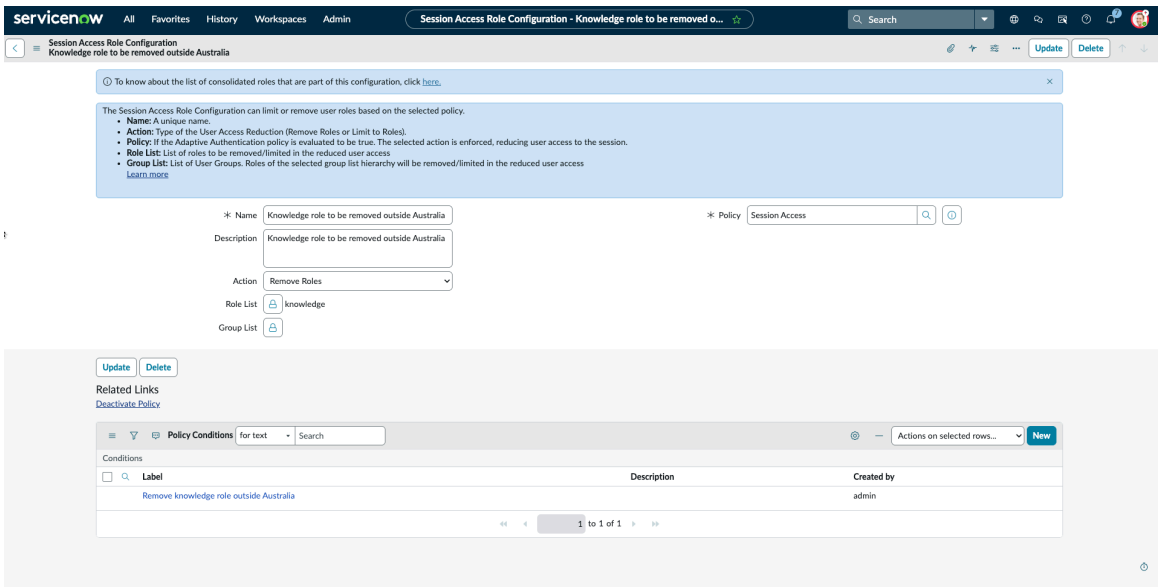

   To know more about how to remove or limit roles for a session explained with a sample use-case, see Tutorial: Use Zero Trust Access.

Zero Trust Access system properties

Use system properties to enable and customize Zero Trust Access to meet your security requirements.

Properties

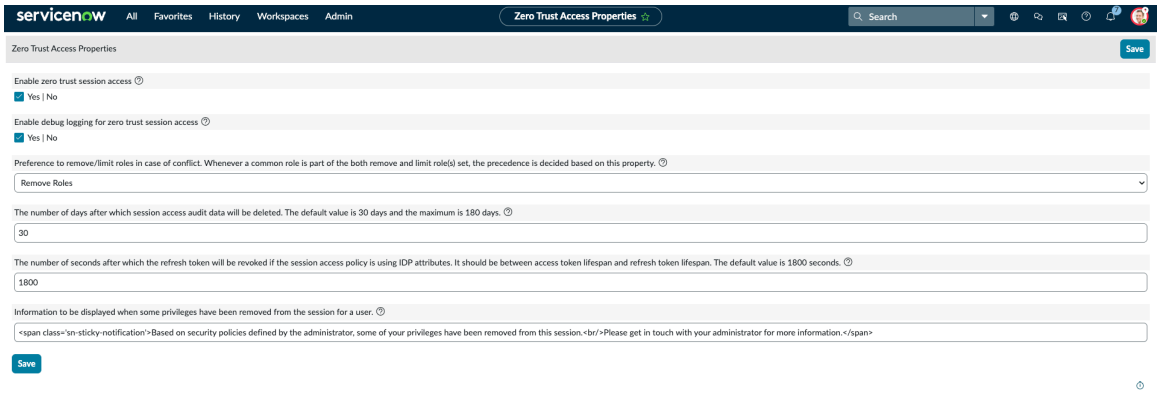

Zero Trust Access system properties

Property	Description
Enable Zero Trust Session Access	Option that enables administrators to use the Zero Trust Session Access feature. By default the value is `false`.
Enable debug logging for Zero Trust Session Access	Option to enable debug logging for Zero Trust Session Access.
Preference to remove/limit roles in case of conflict. Whenever a common role is part of both remove and limit role(s) set, the precedence is decided based on this property.	**Remove Roles** or **Limit Roles**
The number of days after which session access audit data will be deleted. The default value is 30 days and the maximum is 180 days.	By default, it's 30 days.
The number of seconds after which the refresh token will be revoked if the session access policy is using IDP attributes. It should be between access token lifespan and refresh token lifespan. The default value is 1800 seconds.	By default, it's 1800 seconds.
Information to be displayed when some privileges have been removed from the session for a user.	Description that you want to display to your users regarding limiting or removal of privileges. Sample Description: `Based on security policies defined by the administrator, some of your roles have been removed from this session. Please get in touch with your administrator for more information.`

Session Access Audits

The Session Access Audits displays the Session Access logs and information related to a user's session.

Audits

The Session Access audit displays audit information as follows:

Note: Use the `glide.authenticate.session_access.log_audit_event` property to populate the audit information.

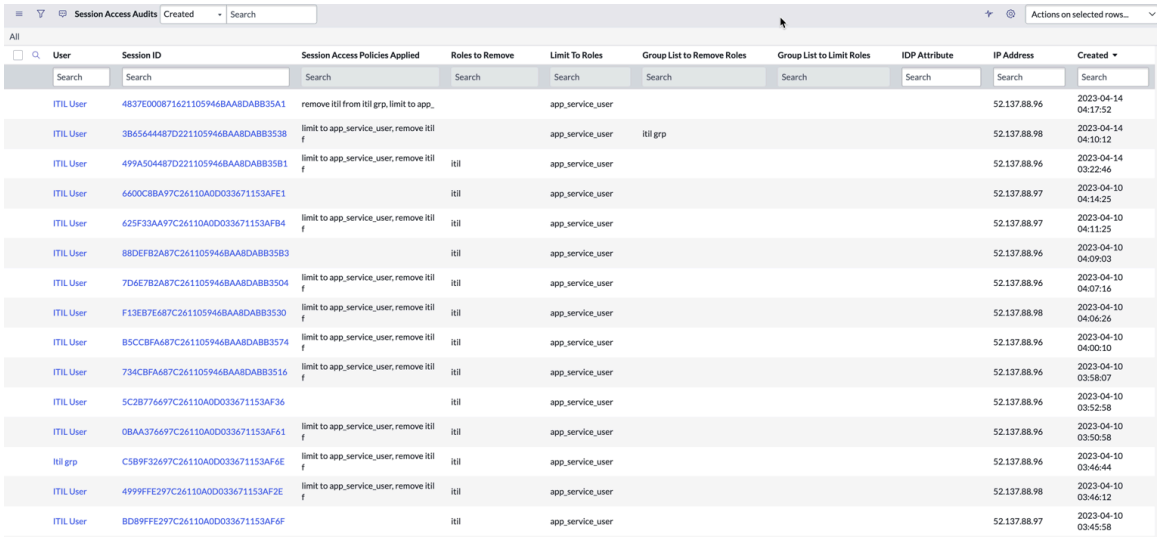

Session Access Audits

Field	Description
User	Details of the user.
Session ID	Details about the session displayed as a unique ID.
Mobile Client	The mobile client details. **ServiceNow Agent** (Now Agent) and **ServiceNow Request** (Now Mobile).
Session Access Policies Applied	The Session Access policy that is applied.
Roles to Remove	Roles that were removed from the user while logging.
Limit To Roles	Roles that were limited to the user while logging.
Group List to Remove Roles	Information about the group that has removed roles for the user.
Group List to Limit Roles	Information about the group that has limited roles for the user.
Mobile Client	Details of the user logged-in through mobile with reduced or removed access.
IDP Attribute	The IDP that was used for that session.

Session Access Audits (continued)

Field	Description
IP Address	Details of the IP address used by the user to log in.
Created	Details of the created user record's date and time.

Tutorial: Use Zero Trust Access

Procedure to use Zero Trust Access feature with an end-to-end use case.

Before you begin

Role required: security_admin

Enable the **Enable Session Access property**.

Note:

- Session Access configurations can only be performed with `security_admin` role. You must elevate your role to `security_admin`.
- Session Access doesn't support integrations.
- Session Access has no impact if the reduced or limited role isn't assigned to a user. In this case, there are no changes to the logged in session. User will still continue to access the instance with their assigned privileges.
- Session Access has no impact while the user is already logged in to the instance and simultaneously the admin configures the policy. The user has to log out from the session for the policy to be effective.
- Session Access is enforced at the time of login. Any change in risk parameters during the session won't result in reduced access. For example, a user switching from the corporate network to an untrusted network after establishing the session, won't result in reduced access unless the user logs out and logs in again.
- Session Access (Zero trust access - ZTA) feature, roles like `snc_internal` and `snc_external` cannot be removed.
- Session Access (Zero trust access - ZTA) feature does not remove a role from the `sys_user_has_role` or the user group membership table. Based on the ZTA policy, it establishes the user session with reduced or limited roles.
- The scripts running in the system context will not honor the ZTA session roles.

Session Access is a feature that enables the administrators to dynamically reduce or restrict a set of roles to the user, when the user is trying to log in to the instance from different environments such as log in from the untrusted network, log in from a different device, and so on.

Session Access can be controlled by the created policy and selected action when performing the configuration. Some of the scenarios are as follows:

- If the Policy is true, and the roles action is set to **Remove Roles**, then the selected roles and its associated child roles are removed for the user when trying to log in to the instance.
- If the Policy is true, and the roles action is set to **Limit To Roles**, then only the selected roles and its associated child roles are assigned to the user when trying to log in to the instance.

The following procedure explains an end-to-end configuration of session access configuration based on which the role is limited to the user who is logging in to the instance. Similarly you can also remove roles by selecting the **Remove Roles** option during the configuration.

Procedure

1. Navigate to **All** > **Session Access** > **Session Access Role Configurations**.
2. On the Session Access Role Configurations page, select **New**.
3. To limit any role for the user, on the form, fill the fields:
   - Name
   - Description
   - Policy
   - Action
   - Role List
   - Group List

   a. Choose **Limit To Roles** to limit roles for the user.
   For example, **itil**.

   b. Choose **knowledge** role from the Role List.

   c. Choose the **Policy**.

   You can create the Session Access policy using an authentication policies and filter criteria (Role, Group, IP, Location) with policy inputs and conditions.

   Use the policy in the Session Access configuration. For example, you want to limit the role (knowledge) to the user logging in outside the Location (Australia).

   d. Choose Action as **Limit To Roles**.
   If the Policy is true, then only the selected roles and their associated child roles are available for the user when trying to log in to the instance.

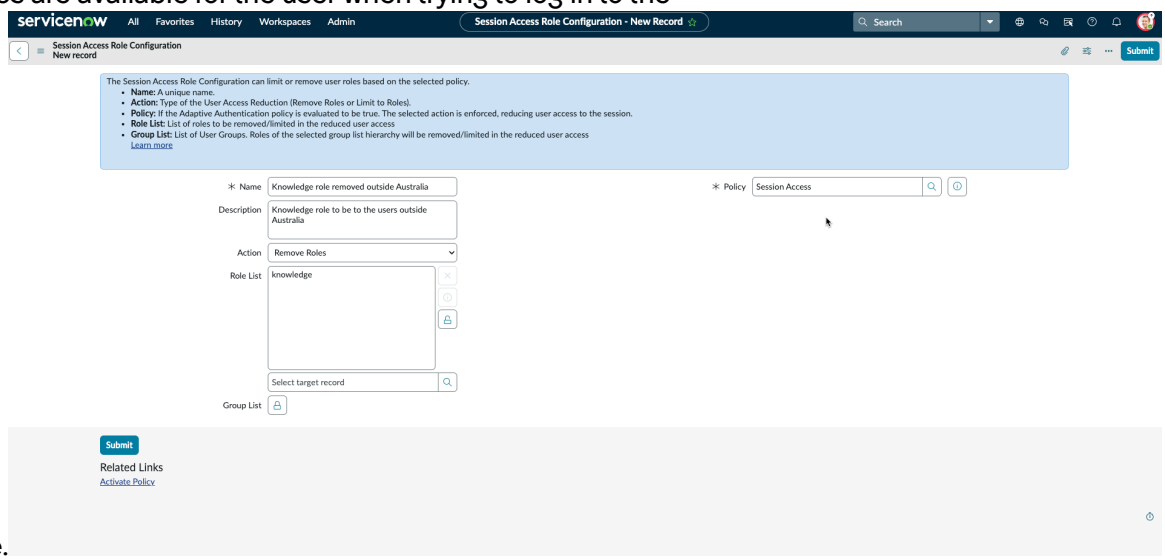

e. Select **Submit**.

Similarly, you can choose the group from the Group List to restrict or remove roles for the users within the group.

When the user logs to the instance outside Australia, only the **Knowledge** role and its associated child roles are assigned for the logged session and other roles to the user are restricted.

After logging in, the user is displayed with the following error message on the platform in their profile section:

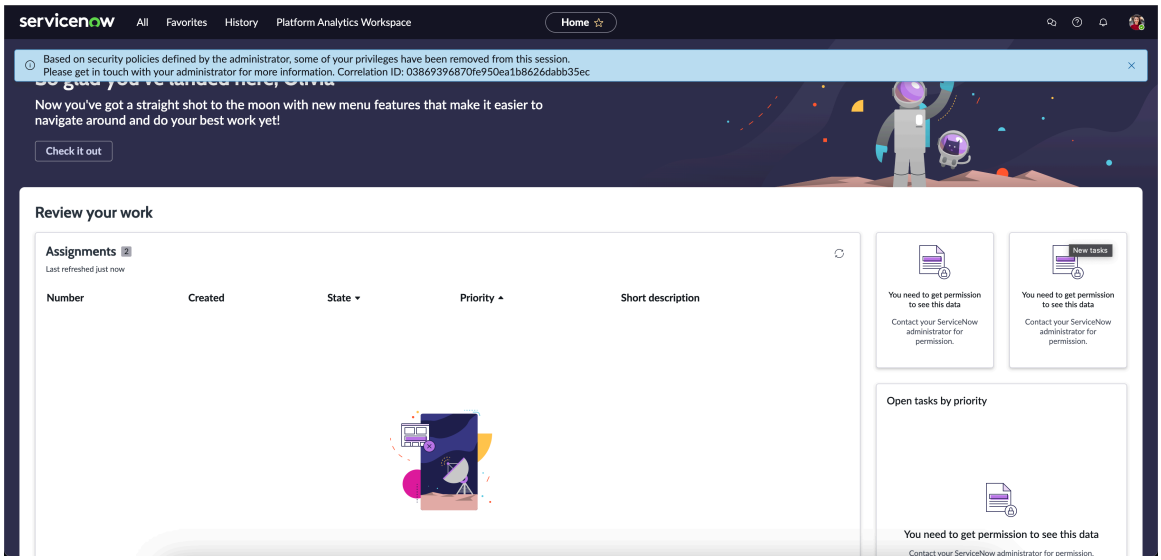

The user can contact the administrators and provide the Correlation ID for investigation.

Note: The correlation ID is the sys_id of the corresponding audit record in the session access audit table.

Configure Identity Provider attribute for Session Access

Use Identity Provider (IDP) attribute created from the Security Assertion Markup Language (SAML) response and OpenID Connect (OIDC) for removing or restricting user session access to the instance.

Before you begin

Role required: security_admin

Enable the **Enable Session Access property**.

Note: To use the Session Access role configuration, you must elevate your role to `security_admin`.

Session Access can be controlled by the created policy and selected action when performing the configuration. Some of the scenarios are as follows:

- If the Policy is true, and the roles action is set to **Remove Roles** along with the IDP attribute input and condition, then the selected roles and its associated child roles are removed for the user when trying to log in to the instance.
- If the Policy is true, and the roles action is set to **Limit To Roles** along with the IDP attribute input and condition, then only the selected roles and its associated child roles are assigned to the user when trying to log in to the instance.

The following procedure shows steps to configure the IDP attribute from the SAML response a policy input to control session access.

Procedure

1. Navigate to **All** > **Session Access** > **Session Access Role Configurations**.
2. On the Session Access Role Configurations page, select **New**.
3. For removing any role for the user, on the form, fill the fields:
   - Name
   - Description
   - Policy
   - Action
   - Role List
   - Group List

   a. Choose **Remove Roles** to remove roles for the user.
   For example: **itil**.

   b. Choose **itil** role from the Role List.

   c. Choose the **Policy**.
   To know more on how to create policy using different filter criteria using Adaptive Authentication policy creation, see Filter criteria.

   d. Choose Action as Remove Roles.
   If the Policy is true, and the roles action is set to **Remove Roles**, then the selected roles is removed for the user when trying to log in to the instance.

   e. In the Policy Input, create the Policy Input and Policy Condition.
   For example:
   - **Policy Input**: Risk score attribute from Okta (IDP).
   - **Policy Conditions**: Risk score from 60 through 80 as a condition.

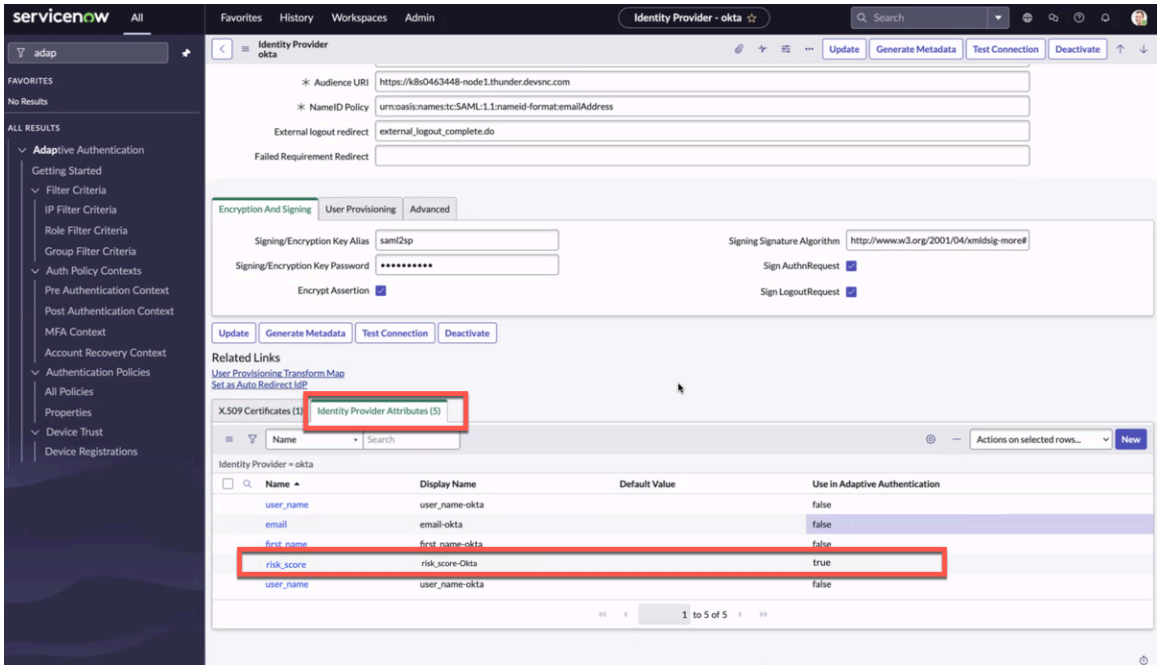

Based on this configuration, when the risk score attribute value from the Okta (IDP) beyond 80, the user isn't authenticated with the roles (**employee**) and its child roles that were removed to the instance, the user is only authenticated with other roles assigned. If the risk score is from 60 through 80, the user is authenticated to the instance with all the roles.

For more information on how to create Post Authentication Context's Policy with Policy inputs and Condition, see Post-authentication context.

Note: If the **Enable Session Access** property is inactive, then the configuration for Session Access doesn't restrict or remove the roles of the user.

f. Select **Submit**.

Zero Trust Access for Mobile

Zero Trust Access (ZTA) is a security model that assumes that no user or device is trusted by default.

You can use the Zero Trust Access - Session Access policy within the Adaptive Authentication policy to reduce the roles or privileges of the particular session in mobile for users.

To enable Zero Trust Access on mobile, you must perform the following tasks:

- Session Access configurations can only be performed with `security_admin` role. You must elevate your role to `security_admin`.
- Activate the **Zero Trust - Policy Based Session Access** `com.snc.zero_trust_session_access` policy.

- Enable the **glide.authenticate.session_access.mobile.enabled** from the system properties

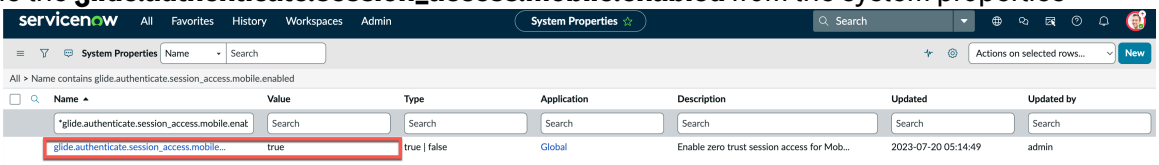

table.

- Configure the **glide.authenticate.session_access.mobile.refresh_token_interval** field to control session access on mobile based on the refresh

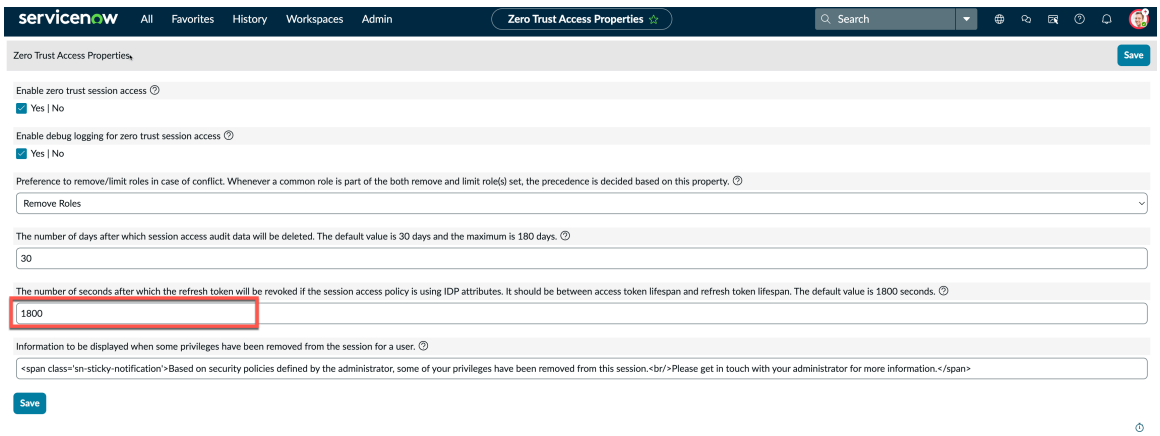

token.

> **Note:** You must configure the refresh token seconds when using an IDP for Mobile App logins. By default, users are logged out from the mobile apps after 1800 seconds (30 minutes).

- Set Enable Zero Trust Access to true under **Application Registries** for the mobile client application (OAuth client). In this case, **ServiceNow Agent** (Now Agent) and **ServiceNow Request** (Now

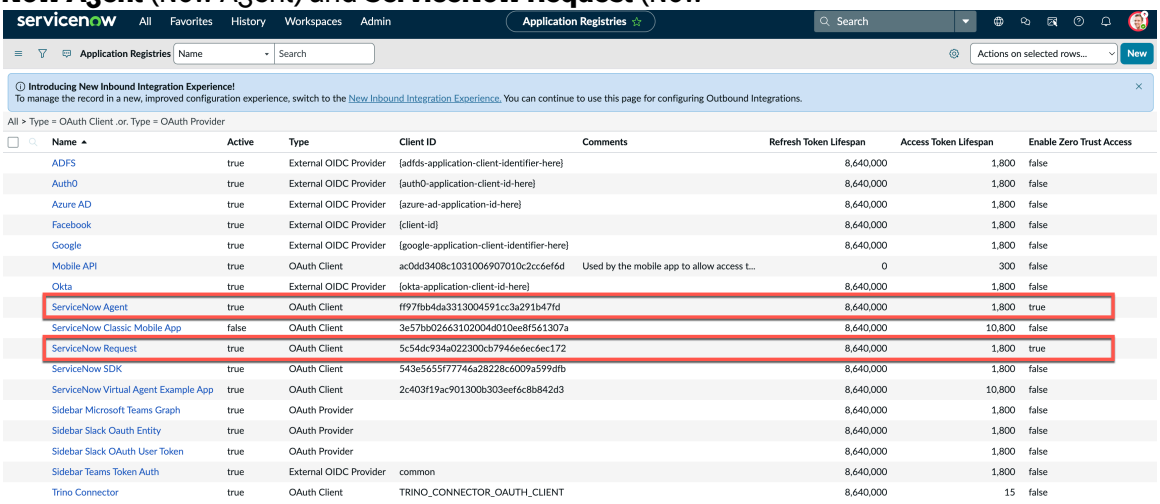

Mobile).

- Configure Session Access role to either reduce or remove roles for the users logging based on the policy inputs and conditions. To learn more about the configuration, see Configure Session Access role.

The configuration evaluates the login to reduce or remove the roles of the users who access your ServiceNow® instance based on the policy filters and conditions. For more information, see Configure Zero Trust Access for mobile.

Continuous Authentication (CA)

ServiceNow's continuous authentication enables you to reverify and authenticate a user if they access resources that are protected by you.

Explore  Learn the features and business value of Continuous Authentication.	Activate  Understand how to Activate Continuous Authentication.
Configure  Configure Continuous Authentication.	Reference - Properties  Know about the properties in Continuous Authentication.

Explore Continuous Authentication

ServiceNow's continuous authentication (CA) enables you to reverify and authenticate a user if they access resources that are protected by you.

ServiceNow's continuous authentication is a security mechanism designed to verify a user's identity not just at the initial login, but throughout the user's entire session. CA is built on ServiceNow's zero trust access security architecture that aims to enhance security by ensuring that the user remains who they claim to be, even after the initial authentication process.

CA works on the following zero trust access principles:

- **Verify explicitly**: No implicit trust for any user, device, or system within a network, regardless of location. Every user and device must be explicitly authenticated and authorized, regardless of location or past access.
- **Use least privilege access**: Grant only the minimum access or permissions needed to perform specific tasks and limit potential damage from compromised accounts or systems."
- **Assume breach**: Instead of relying only on prevention, assume breach and focus on proactive detection, containment, and response.

CA provides the ability to enforce step-up authentication or re-authentication based on the data users are accessing and activities they are performing. It be opted by administrators for creating security policies at a table or data class level.

You can enforce step-up authentication (MFA) or re-authentication (SSO - SAML or OIDC) within a logged-in session whenever there is an attempt by the user to access Personally Identifiable Information (PII) and sensitive data.

Note: You must install the **Zero Trust - Continuous Authentication** (`com.snc.zero_trust_continuous_authentication`) for opting CA which requires a license.

Benefits

Following are the some of the benefits of using CA:

- **Enhanced Security**: By continuously verifying the user's identity, the system can detect and respond to potential security threats more quickly.
- **Reduced Risk of Account Takeover**: Even if an attacker gains access to a user's session, continuous authentication can help prevent them from accessing confidential data.

Use cases

Following are some of the use cases for using CA:

- Enforce re-authentication before allowing access to sensitive data using different policies.
- Enforce periodic re-authentication or step-up authentication using different policies:
  - Use re-authentication that can include IdP's MFA, IdP's SSO.
  - Use step up authentication with ServiceNow's MFA.

Roles in CA

CA has the following roles:

- CA Admin (ca_admin): Ability to create, edit, and view CA policies. Configure CA properties and view dashboards (Metrics) of CA.
- Policy Admin (ca_policy_admin): Ability to create, edit, and view CA policies..
- Auditor (ca_auditor): Ability to view dashboards (Metrics) of CA. And policies, and logs of CA.

To configure CA you must elevate your role to **ca_admin** and perform the policy configurations.

Note: All these 3 roles are elevated roles.

Modules in CA

Following are the different modules within CA:

- Policies: View the different continuous authentication policies that are created.
- Metrics: View the different metrics for continuous authentication for KPI purposes and understand the usage of CA within your organization.
- System Properties: Use system properties to enable and customize continuous authentication (CA) to meet your zero trust access security requirements.

Related topics

Policies

Metrics

System properties

Pre-work for Continuous Authentication

Activate Continuous Authentication

Configure Continuous Authentication

High Assurance session with Continuous Authentication

Policies

View the different continuous authentication policies that are created.

The CA policies form includes the details of the CA policies that are created for the table or data class. You can manage policy records from the custom policies.

To access the policies page, navigate to **All** > **Continuous Authentication**, select **Policies** tab. Following are the policy details that are displayed on the page:

Policies Page

Field	Description
Policy Name	Policy name
Description	Policy description
Active	Policy status
Resource Type	Selected resource type used for the policy
Classification	Data classification selected for the policy
Tables	Tables selected for the policy
Created	Policy creation details

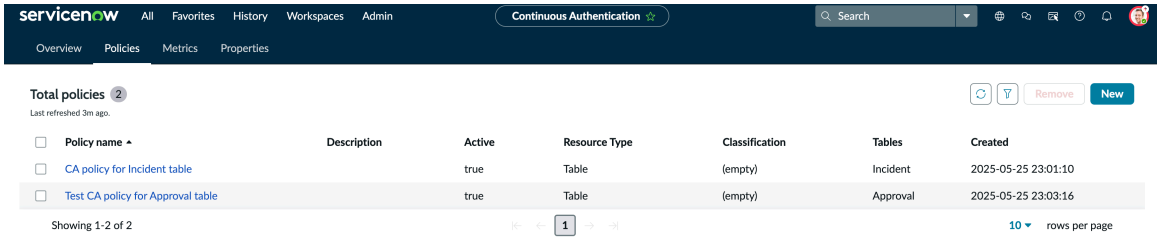

Related topics

- Metrics
- System properties
- Explore Continuous Authentication

Metrics

View the different metrics for continuous authentication.

To access the metrics page, navigate to **All** > **Continuous Authentication**, select **Metrics** tab.

Following are the different KPI details that you can view for continuous authentication:

Metrics

Metrics	Description
Total policies	KPI for number of continuous authentication policy that are created for the users.
Total protected classifications	KPI for number of classification that are protected due to the CA policy creation.
Failed step-up or re-authentication	KPI for the number of failed step-up (MFA) or re-authentication (SSO).
No of times policies invoked	KPI for the total number of times CA policies invoked.

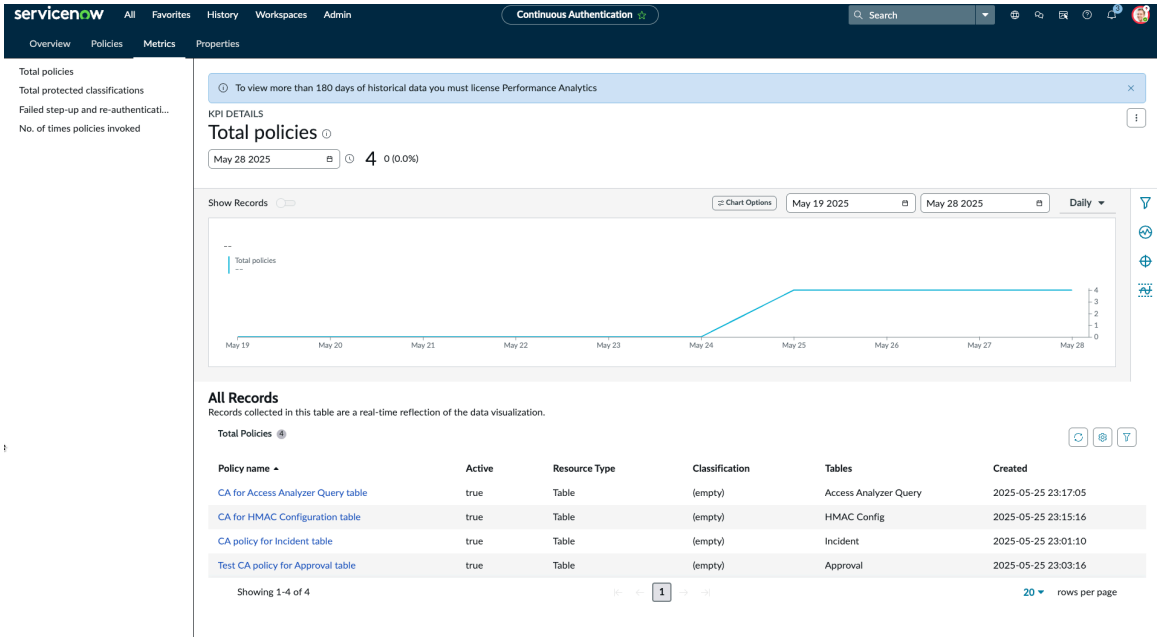

Note: The KPI details are displayed for 180 days. To view more than 180 days of historical data you must license **Performance Analytics**.

Related topics

- Policies
- Metrics
- System properties
- Explore Continuous Authentication

System properties

Use system properties to enable and customize continuous authentication (CA) to meet your zero trust access security requirements.

Properties

To access the properties page, navigate to **All** > **Continuous Authentication**, select **Properties** tab.

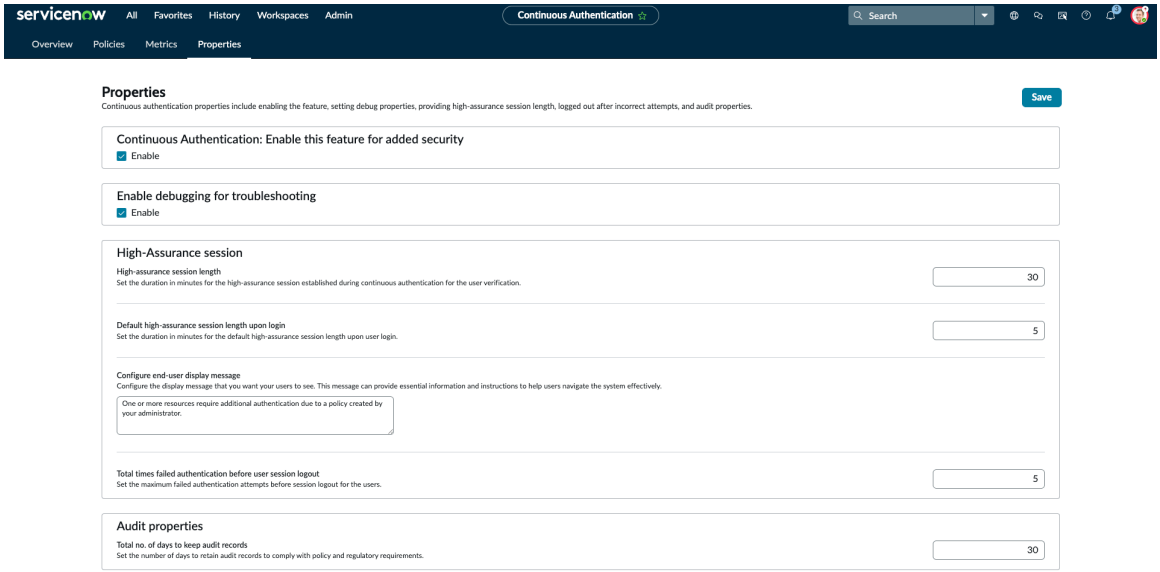

Following are the different system properties for CA:

Continuous Authentication system properties

Property	Description
General Properties	
Continuous Authentication (*glide.zta.continuous_authentication.feature.enabled*)	Enable to use Continuous Authentication feature
Enable Debugging (*glide.zta.continuous_authentication.debug.enabled*)	Enable to view the debugging information related to continuous authentication.
High Assurance	
High Assurance session length (*glide.zta.high_assurance.session.timeout*)	Specify the high assurance session length, after which the end-users should re-authenticate. Default: 10 mins. **Note:** The value must be between 1 and 480.
Default high-assurance session length upon login (*glide.zta.default.high_assurance.session.timeout*)	Specify the duration in minutes for the default high-assurance session length upon user login. Default value: 5 mins. **Note:** The property is only applicable for local login.
Configure end-user display message (*glide.zta.high_assurance.session.message*)	Specify the message that is displayed to the end-user for re-authentication. Default message: `One or more resources require additional authentication due to a policy created by your administrator.`

Continuous Authentication system properties (continued)

Property	Description
Total times failed authentication before user account lock-out (*glide.zta.high_assurance.session.max.login.failed attempts*)	Set the maximum failed authentication attempts before the users are logged out. **Note:** The value must be between 3 and 10.
Audit properties	
Total no of days to keep audit records (*glide.zta.continuous_authentication.audit.records.lifespan*)	Specify the no of days you want to save the audit records for CA. **Note:** The value must be between 1 and 180.
Total no. of days after which policies will be deleted after deactivated (*glide.zta.continuous_authentication.policy.lifespan*)	Specify the no of days after which the CA policies are deleted.

Important:

- By default, high-assurance sessions are not required for mobile app sessions, even when a continuous authentication policy is active on source. To change this behavior and block access from mobile app sessions, update the **glide.zta.high_assurance.mobile.session.allowed** property value to `false`.
- The **sys_properties**, **sys_continuous_auth_policy**, **sys_user** tables are excluded for CA and cannot be added to the CA policy configuration.

Related topics

Explore Continuous Authentication

Policies

Metrics

Pre-work for Continuous Authentication

Pre-work for Continuous Authentication

Ensure to perform the following pre-work before using Continuous Authentication (CA).

CA is built on ServiceNow's zero trust access security architecture that aims to enhance security by ensuring that the user remains who they claim to be, even after the initial authentication process.

Note: You must install the **Zero Trust - Continuous Authentication** (`com.snc.zero_trust_continuous_authentication`) for opting CA which requires a license.

CA configuration can be performed based on the following:

- LDAP or Local login (Username and Password)
- SSO login (SAML or OIDC).

CA for Local login

When the users are performing local login and to verify the Identity of the users, MFA is used as an authentication mechanism. If the users haven't setup MFA, it is compulsory to complete the setup.

Note:

- From Yokohama, MFA is enforced to users for every login to ServiceNow performing local login.
- Make sure the MFA properties are Active and configured based on your requirement. To know more about MFA properties, see Multi-factor Authentication system properties.

To know more, see High Assurance session for non-SSO login.

CA for SSO login

When the users are performing SSO based login (SAML or OIDC). To verify the Identity of the users, the same SSO used during initial login is shown as the re verification with re-authentication or IdP's MFA.

Note:

- You must install the **Zero Trust - Continuous Authentication** (`com.snc.zero_trust_continuous_authentication`) for opting CA which requires a license.
- Activate the Integration - Multiple Provider Single Sign-On Installer (*`com.snc.integration.sso.multi.installer`*) plugin.

You must configure the IDP for CA as follows:

- Enable the check box that is required to be set on a given multi SSO record to validate that set ready for using CA.

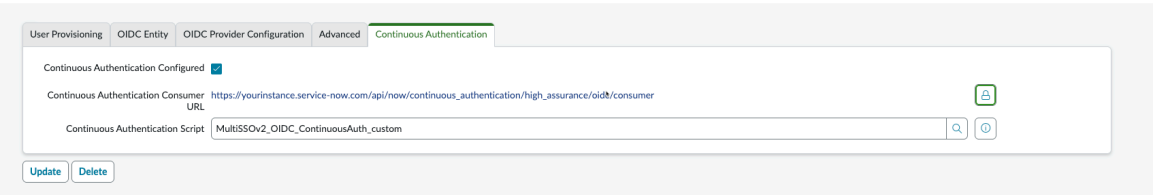

- For OIDC, CA relies on redirecting back to the `api/now/continuous_authentication/high_assurance/oidc/consumer` endpoint, which must be configured on the IDP. Both re-authentication and IdP's MFA options are available.
- For SAML, the SSO records use the default re-athentication script for all Identity Providers (IdP) to support re-authentication.

Note:

- To configure step up for **Okta** you can use the **ContinuousAuth_Okta_StepUp_Script** in the IdP record. To know more, see this documentation.
- To configure step up for **Entra ID** or **Azure** you can use the **ContinuousAuth_Azure_StepUp_Script** and add the required claim. To know more, see this documentation.

To know more, see High Assurance for SSO login.

Related topics

- Explore Continuous Authentication
- Activate Continuous Authentication
- Configure Continuous Authentication

Activate Continuous Authentication

For activating the Continuous Authentication feature on your instance, install the **Zero Trust - Continuous Authentication** (`com.snc.zero_trust_continuous_authentication`) plugin.

Before you begin

Role required: admin

This plugin enables security administrators to define security policies that require step-up authentication (MFA) or re-authentication (SSO) within a logged-in session, based on the data the user is accessing and the activities they are performing.

You must install the **Zero Trust - Continuous Authentication** (`com.snc.zero_trust_continuous_authentication`) for opting CA which requires a license.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the **Zero Trust - Continuous Authentication** (`com.snc.zero_trust_continuous_authentication`) plugin using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Related topics

- Pre-work for Continuous Authentication
- Configure Continuous Authentication
- High Assurance session with Continuous Authentication

Configure Continuous Authentication

Configure continuous authentication (CA) policies to re-authenticate the users if there's an attempt to access resources that are protected by you.

Before you begin

- Role required: Admin (ca_admin)

  Note: You must elevate your role to **ca_admin**.

- You must install the **Zero Trust - Continuous Authentication** (`com.snc.zero_trust_continuous_authentication`) for opting CA which requires a license.
- Enable the Continuous Authentication (*`glide.zta.continuous_authentication.enabled`*) system property. For more information, see System properties.
- Activate the Integration - Multiple Provider Single Sign-On Installer (*`com.snc.integration.sso.multi.installer`*) plugin.
- Understand the pre-work that is required before configuring CA for the instance. For more information, see Pre-work for Continuous Authentication.
- CA policies can be configured for Data Class or Table.

Procedure

1. Navigate to **All** > **Continuous Authentication**.
2. Select **Policies** tab.
3. Select **New**.
4. On the form, fill the fields:

Continuous Authentication

Field	Description
Policy Name	Name of the policy
Description	Generic description to the policy
Select the resources	Options: ○ **Data Class**. You can create data class and use it for CA policy configuration. **Note:** To know more about how to create data class, see Data Classification. ○ **Table**

Field	Description
	Note: ○ Table selected with metadata displays an error. ○ You need to check if you actually want to restrict access to the metadata table, since it can impact configuration access to your users. ○ The **sys_properties**, **sys_continuous_auth_policy**, **sys_user** tables are excluded for CA and cannot be added to the CA policy configuration.

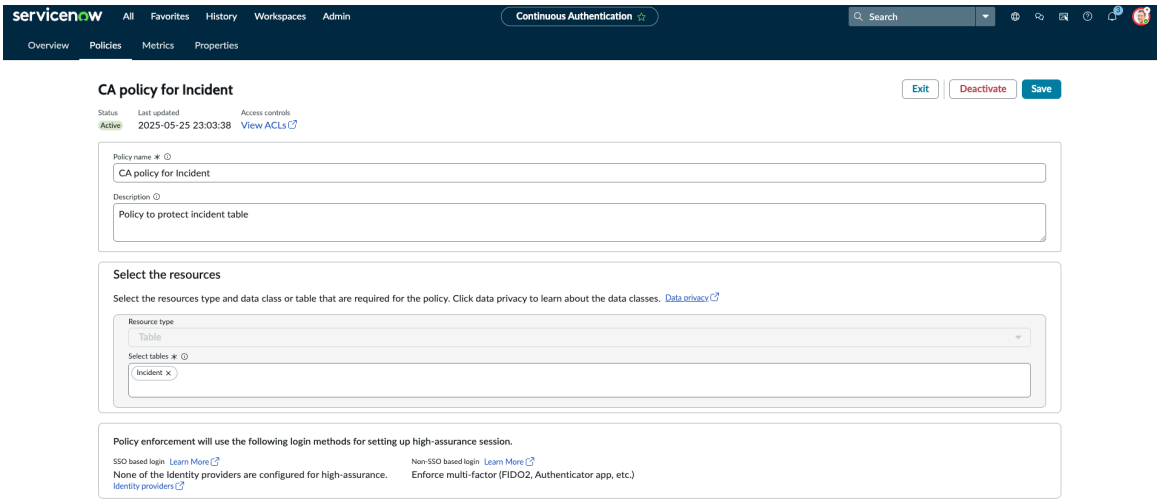

Note: You can use either of the login methods for the CA policy:

- **SSO based login**: Specify the fields in the **Continuous Authentication** tab within the Identity Provider record and the set the Identity Provider record as **Active**.

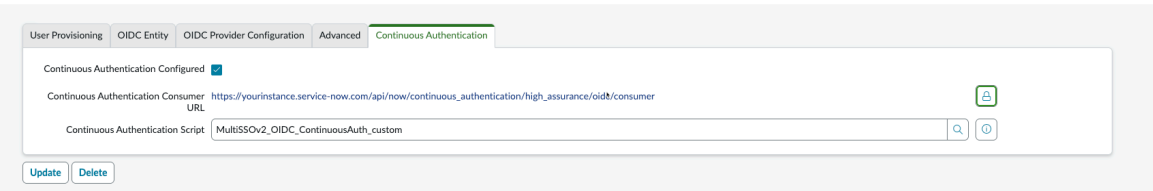

  To know more about Identity Providers configuration, see OIDC and SAML.

- **Non-SSO based login**: By default, if there are no Identity Provider with Continuous Authentication configuration, Multi-factor Authentication (MFA) is used as a login method. Make sure the MFA properties are Active and configured based on your requirement. To know more about MFA properties, see Multi-factor Authentication system properties.

5. Select **Save & Activate**.

Result

Based on the details provided for the configuration, CA policy is created with Access Control List (ACLs) for the selected table or data class. You can view the details of the ACLs that are created by selecting the **View ACLs** on the policy page.

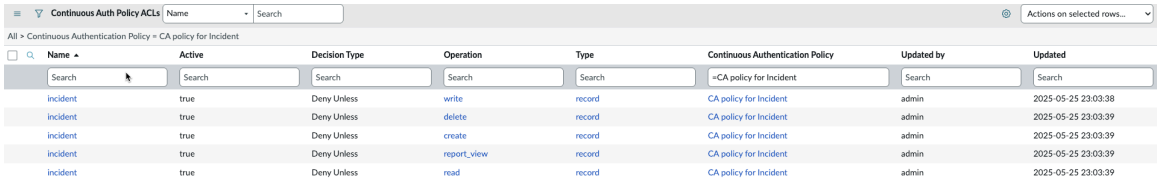

The CA policy created, prompts the user for authentication to access table or data class that you've protected using the policy, based on the following scenarios.

- User who had performed local login to log in to the instance, is displayed with platform MFA for step-up authentication.

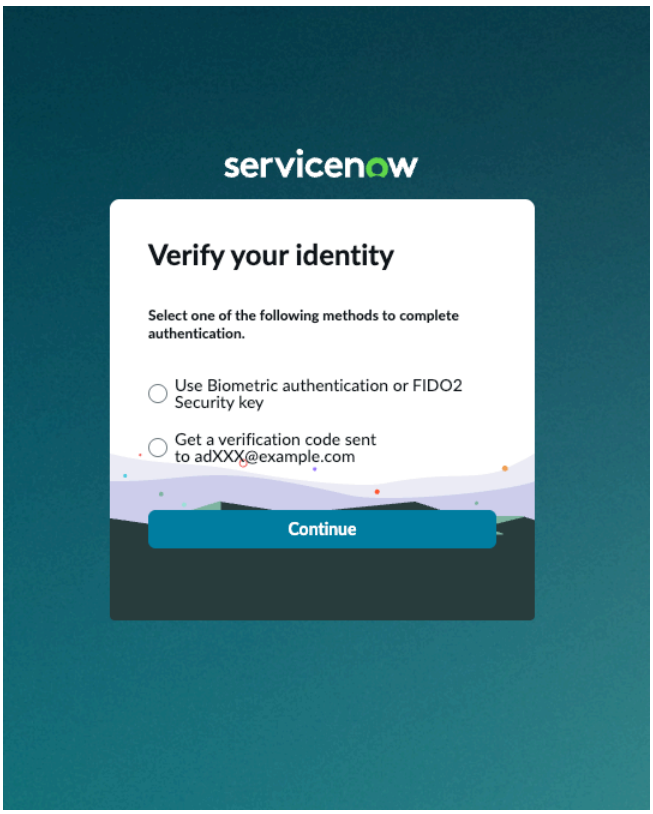

Note: The users recently used MFA factor is displayed for authentication.

- User who had performed SSO login (OIDC or SAML) to log in to the instance is displayed with the SSO for re-authentication.

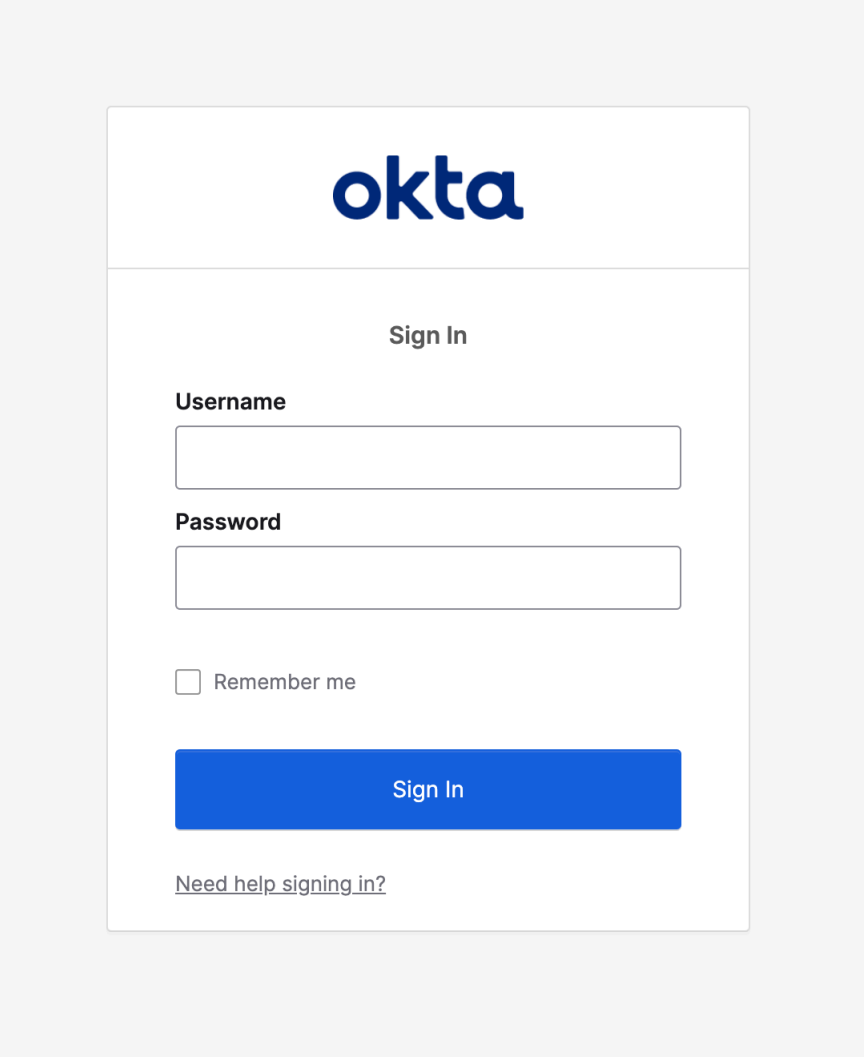

An high assurance session is now established for the user. High assurance session is limited to the High Assurance session length (*glide.zta.high_assurance.session.timeout*) system property. If the high assurance session time exceeds the property length, the user is prompted for re-authentication or step up authentication.

To know more about the end to end configuration of continuous authentication for table or data, see:

- Tutorial: Configure Continuous Authentication for a Table.
- Tutorial: Configure Continuous Authentication for a Data Class.

Related topics

Explore Continuous Authentication

Pre-work for Continuous Authentication

Activate Continuous Authentication

Tutorial: Configure Continuous Authentication for a Table

Procedure that describes end to end configuration of continuous authentication policy for a table and the impacts to the users due to the configuration changes.

Before you begin

- Role required: Admin (ca_admin)

  Note: You must elevate your role to **ca_admin**.

- You must install the **Zero Trust - Continuous Authentication** (`com.snc.zero_trust_continuous_authentication`) for opting CA which requires a license.
- Enable the Continuous Authentication (*`glide.zta.continuous_authentication.enabled`*) system property. For more information, see System properties.
- Activate the Integration - Multiple Provider Single Sign-On Installer (*`com.snc.integration.sso.multi.installer`*) plugin.
- Understand the pre-work that is required before configuring CA for the instance. For more information, see Pre-work for Continuous Authentication.

Procedure

1. Navigate to **All** > **Continuous Authentication**.
2. Select **Policies** tab.
3. Select **New**.
4. On the form, fill the fields:

Continuous Authentication

Field	Description
Policy Name	Name of the policy
Description	Generic description to the policy
Select the resources	Select the **Table**. **Note:** ○ In this example, **Incident** table is selected. You can select as many as tables based on your requirement. ○ Table selected with metadata displays an error. You need to check if you actually want to restrict access to the metadata table, since it can impact configuration access to your users. ○ The **sys_properties**, **sys_continuous_auth_policy**, **sys_user** tables are excluded for CA and cannot be added to the CA policy configuration.

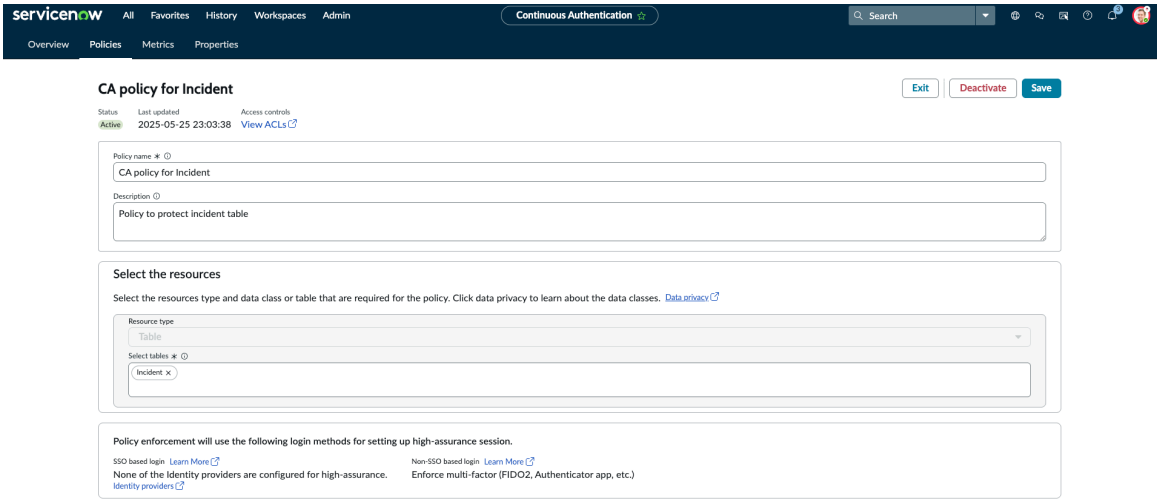

Note: You can use either of the login methods for the CA policy:

- **SSO based login**: Specify the fields in the **Continuous Authentication** tab within the Identity Provider record and the set the Identity Provider record as **Active**.

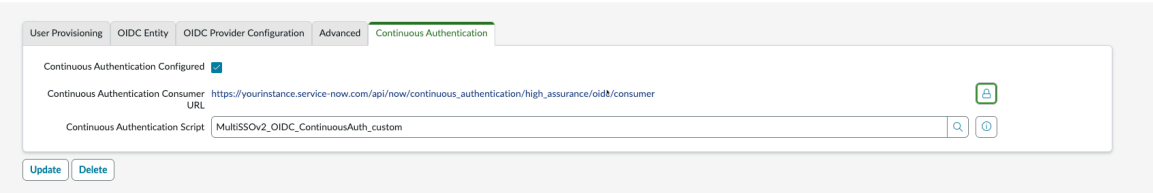

To know more about Identity Providers configuration, see OIDC and SAML.

- **Non-SSO based login**: By default, if there are no Identity Provider with Continuous Authentication configuration, Multi-factor Authentication (MFA) is used as a login method. Make sure the MFA properties are Active and configured based on your requirement. To know more about MFA properties, see Multi-factor Authentication system properties.

5. Select **Save & Activate**.

Result

Based on the details provided for the configuration, CA policy is created with Access Control List (ACLs) for the selected table or data class. You can view the details of the ACLs that are created by selecting the **View ACLs** on the policy page.

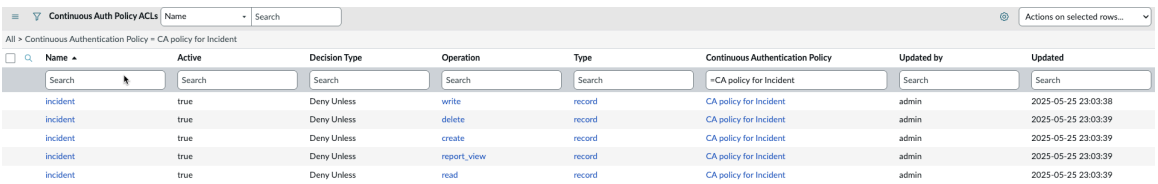

The CA policy created, prompts the user for authentication to access table (in this case **Incident** table) that you've protected using the policy. The users can select **Authenticate** option.

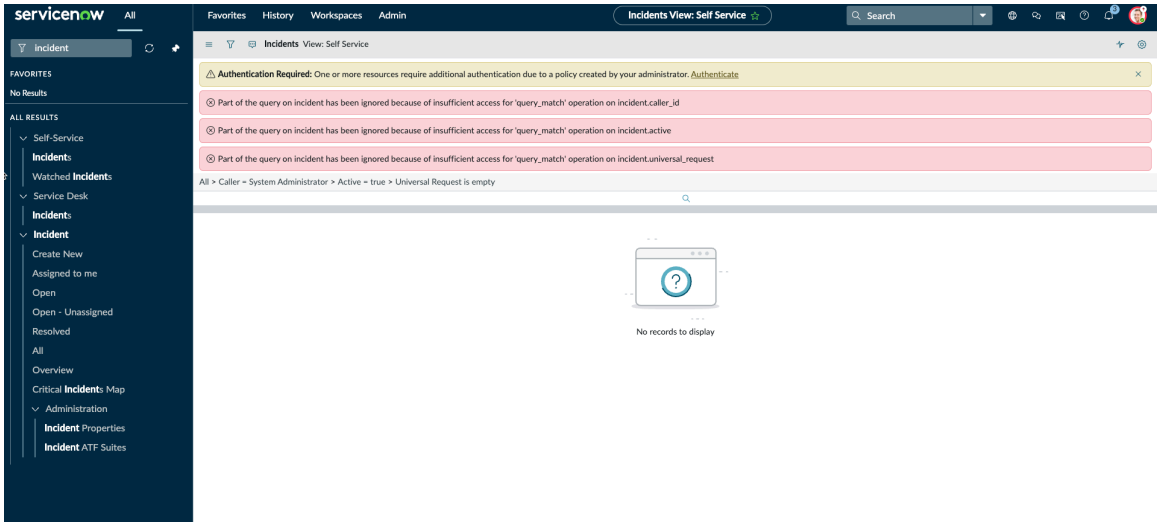

Perform the authentication based on the following:

- User who had performed local login to log in to the instance, is displayed with platform MFA for step-up authentication.

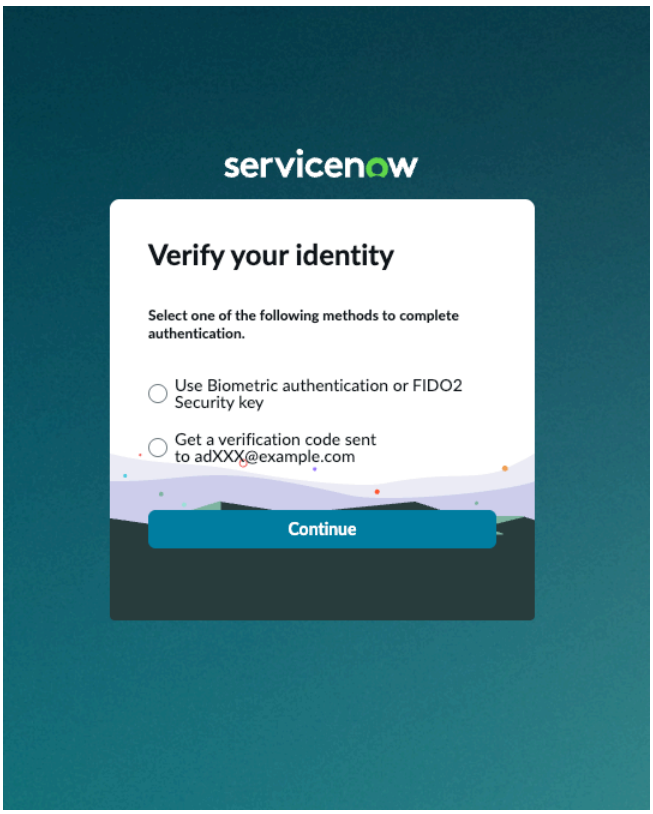

- User who had performed SSO login (OIDC or SAML) to log in to the instance is displayed with the SSO for re-authentication.

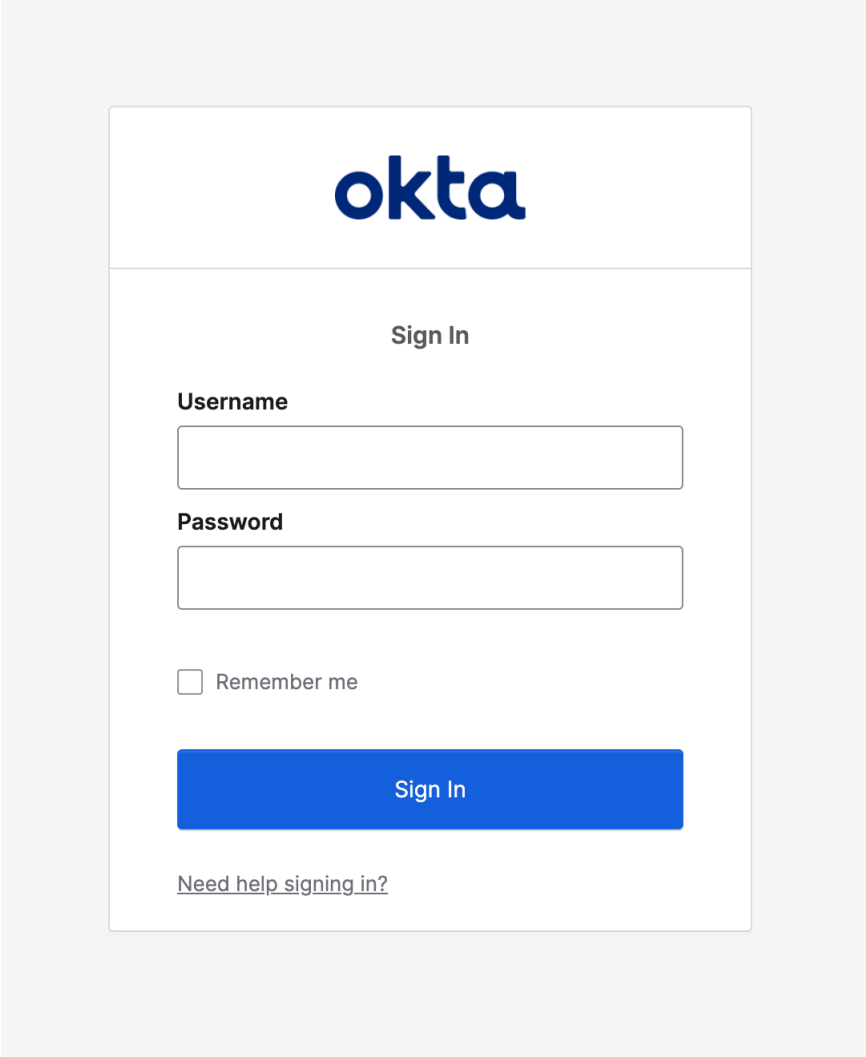

After successful authentication the table is displayed.

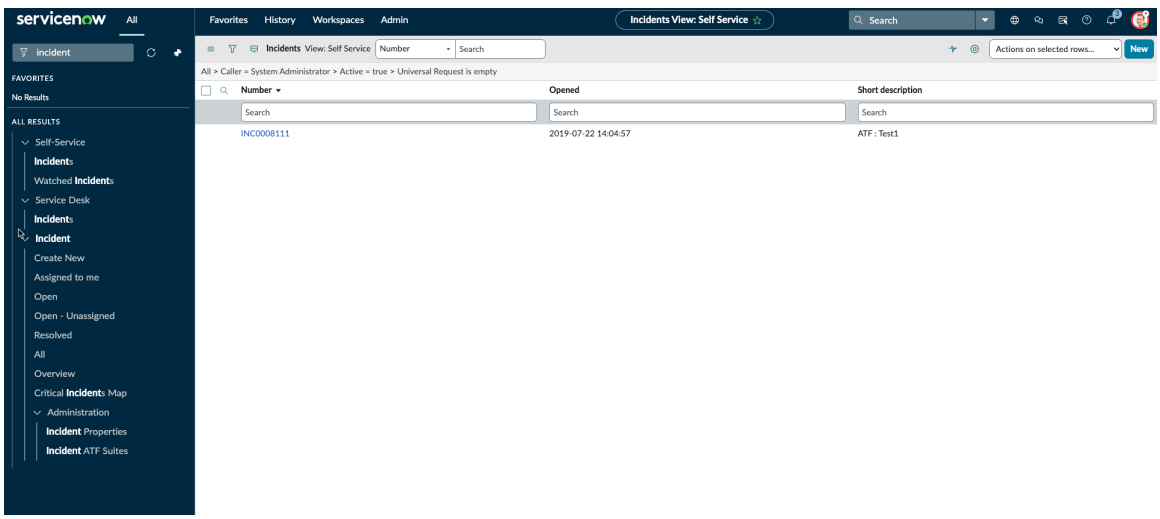

An high assurance session is now established for the user. High assurance session is limited to the High Assurance session length (*glide.zta.high_assurance.session.timeout*)

system property. If the high assurance session time exceeds the property length, the user is prompted for re-authentication or step up authentication.

Related topics

Configure Continuous Authentication

High Assurance session with Continuous Authentication

Explore Continuous Authentication

Tutorial: Configure Continuous Authentication for a Data Class

Procedure that describes end to end configuration of continuous authentication policy for a data class and the impacts to the users due to the configuration changes.

Before you begin

- Role required: Admin (ca_admin)

  Note: You must elevate your role to **ca_admin**.

- You must install the **Zero Trust - Continuous Authentication** (`com.snc.zero_trust_continuous_authentication`) for opting CA which requires a license.
- Enable the Continuous Authentication (*`glide.zta.continuous_authentication.enabled`*) system property. For more information, see System properties.
- Activate the Integration - Multiple Provider Single Sign-On Installer (*`com.snc.integration.sso.multi.installer`*) plugin.
- Understand the pre-work that is required before configuring CA for the instance. For more information, see Pre-work for Continuous Authentication.

Procedure

1. Navigate to **All** > **Continuous Authentication**.
2. Select **Policies** tab.
3. Select **New**.
4. On the form, fill the fields:

Continuous Authentication

Field	Description
Policy Name	Name of the policy
Description	Generic description to the policy
Select the resources	Select the **Data Class**. You can create data class and use it for CA policy configuration. **Note:** To know more about how to create data class, see Data Classification.

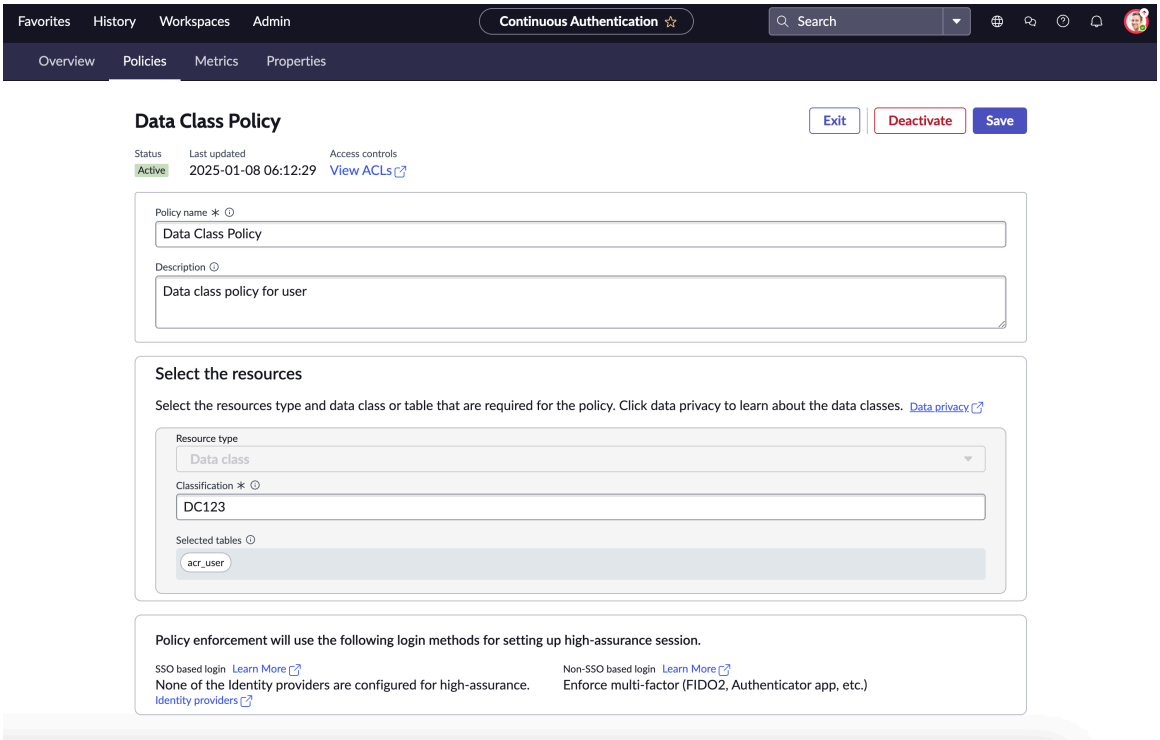

Note: You can use either of the login methods for the CA policy:

- **SSO based login**: Specify the fields in the **Continuous Authentication** tab within the Identity Provider record and the set the Identity Provider record as **Active**.

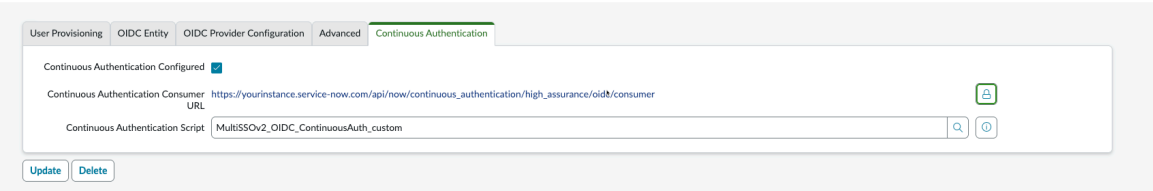

To know more about Identity Providers configuration, see OIDC and SAML.

- **Non-SSO based login**: By default, if there are no Identity Provider with Continuous Authentication configuration, Multi-factor Authentication (MFA) is used as a login method. Make sure the MFA properties are Active and configured based on your requirement. To know more about MFA properties, see Multi-factor Authentication system properties.

5. Select **Save & Activate**.

Result

Based on the details provided for the configuration, CA policy is created with Access Control List (ACLs) for the selected table or data class. You can view the details of the ACLs that are created by selecting the **View ACLs** on the policy page.

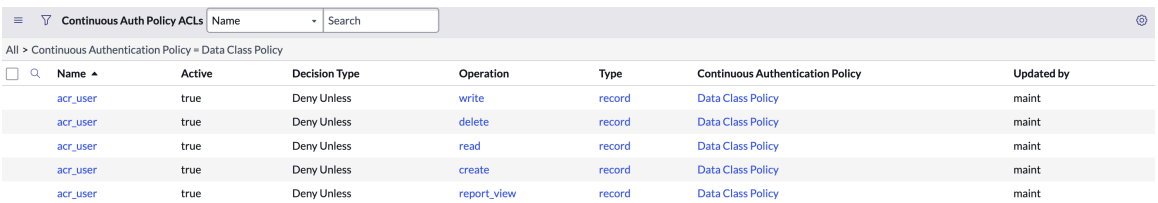

Name	Active	Decision Type	Operation	Type	Continuous Authentication Policy	Updated by
acr_user	true	Deny Unless	write	record	Data Class Policy	maint
acr_user	true	Deny Unless	delete	record	Data Class Policy	maint
acr_user	true	Deny Unless	read	record	Data Class Policy	maint
acr_user	true	Deny Unless	create	record	Data Class Policy	maint
acr_user	true	Deny Unless	report_view	record	Data Class Policy	maint

The CA policy created, prompts the user for authentication to data class (in this case data class set for the table **Account Recovery**) that you've protected using the policy. The users can select **Authenticate** option.

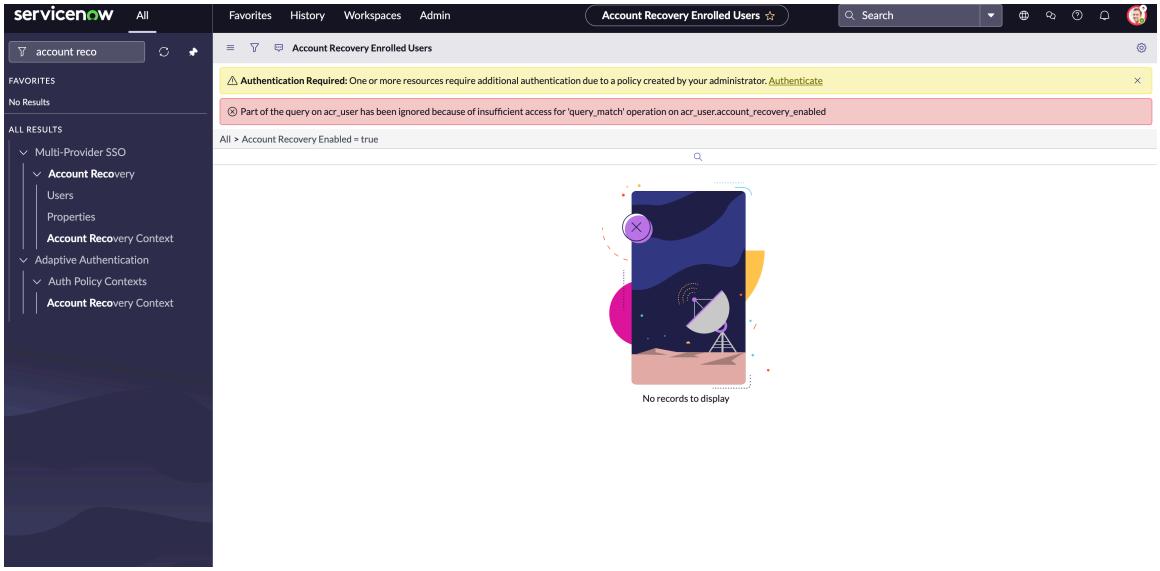

Perform the authentication based on the following:

- User who had performed local login to log in to the instance, is displayed with platform MFA for step-up authentication.

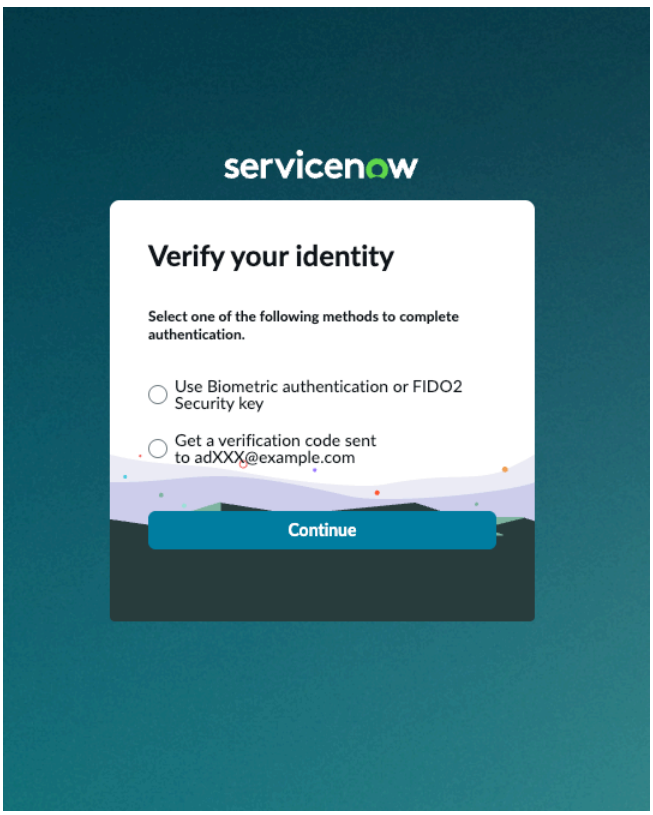

- User who had performed SSO login (OIDC or SAML) to log in to the instance is displayed with the SSO for re-authentication.

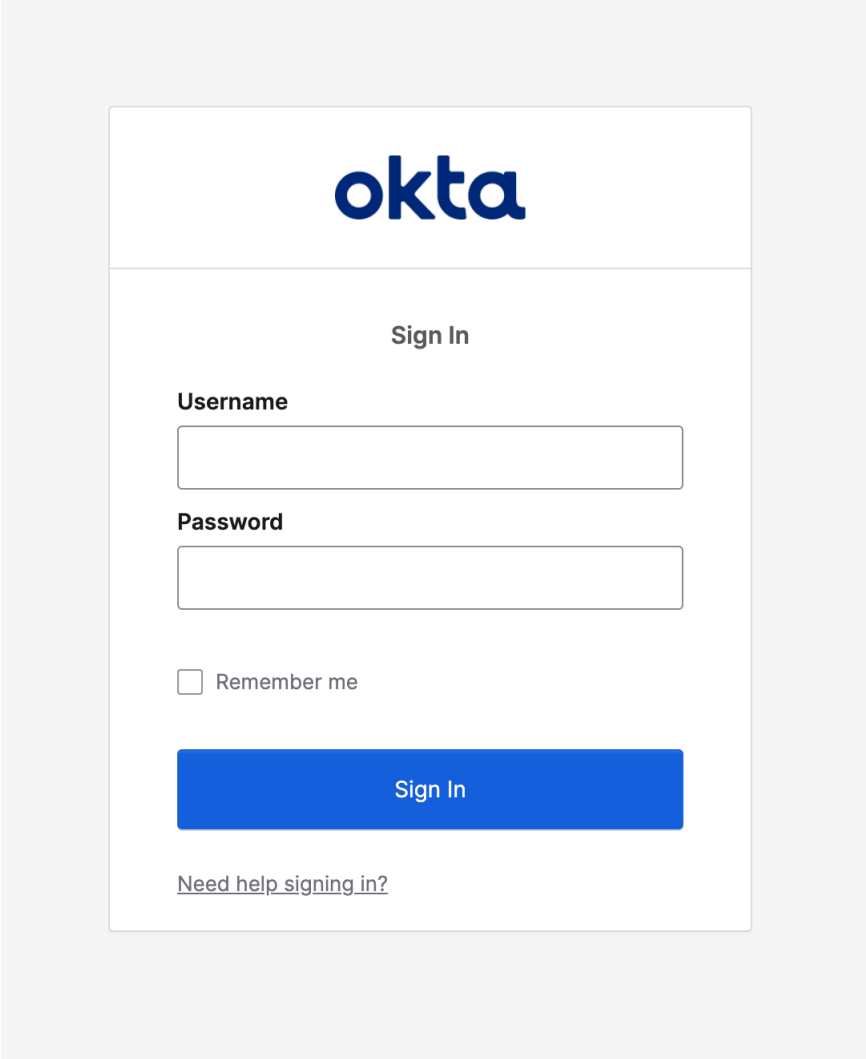

After successful authentication the table with the data class is displayed.

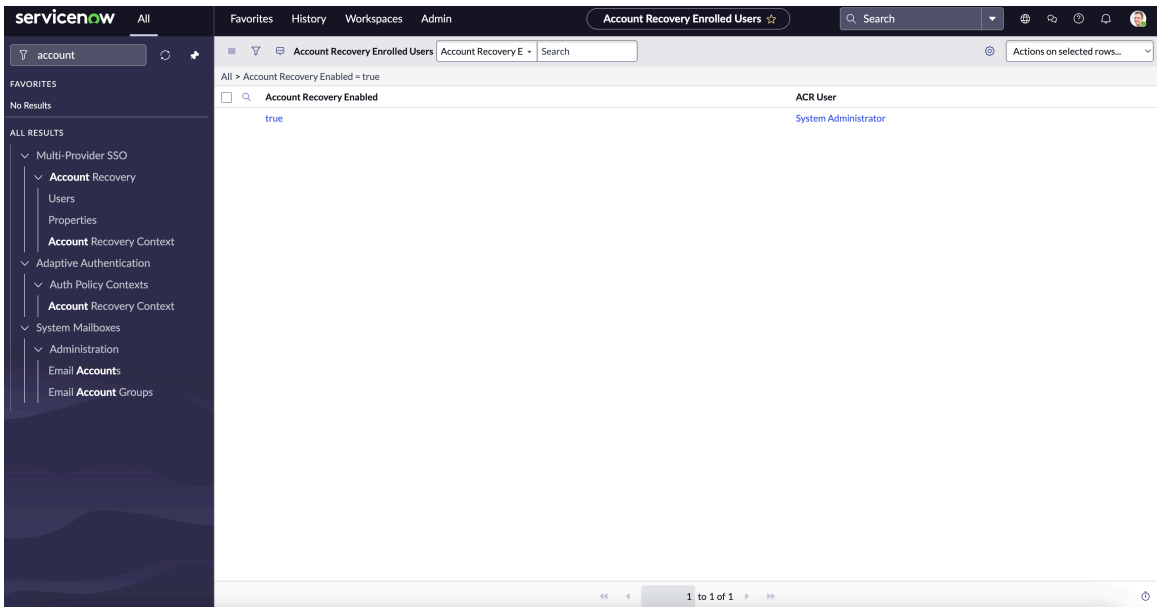

An high assurance session is now established for the user. High assurance session is limited to the High Assurance session length (*glide.zta.high_assurance.session.timeout*) system property. If the high assurance session time exceeds the property length, the user is prompted for re-authentication or step up authentication.

Related topics

Configure Continuous Authentication

High Assurance session with Continuous Authentication

Explore Continuous Authentication

High Assurance session with Continuous Authentication

Establish high assurance session for with ServiceNow's continuous authentication.

A high assurance session is a security measure to establish a secure and trusted connection with the identities (users) who access data and are verified with a high degree of confidence.

ServiceNow's High assurance is achieved through robust authentication methods which enforces re-authentication using methods such as Multi-factor Authentication (MFA) and Single Sign On (SSO) while the users try to access data that are sensitive.

When the user re-authenticates or perform step-up authentication (MFA), there's a high assurance session that is established, which provides the ability for the users to access the data protected by the CA administrator based on the CA policy configuration.

Following are the re-authentication methods used to establish High-assurance based on the type of login:

- High Assurance for SSO login
- High Assurance for non-SSO login

High assurance session created by the user is valid based on the High Assurance session length (*glide.zta.high_assurance.session.timeout*) determined by the CA administrator.

The high assurance session can be customized based on your requirement by setting the **High Assurance** system properties:

High Assurance system properties

Field	Description
High Assurance session length (*glide.zta.high_assurance.session.timeout*)	Specify the high assurance session length, after which the end-users should re-authenticate. Default: 30 minutes. **Note:** The value must be between 1 and 480.
Default high-assurance session length upon login	Specify the duration in minutes for the default high-assurance session length upon user login. Default value: **5** minutes. **Note:** This property is only applicable for non-sso logins.

High Assurance system properties (continued)

Field	Description
Configure end-user display message (*glide.zta.high_assurance.session.message*)	Specify the message that is displayed to the end user for re-authentication. Default message: `One or more resources require additional authentication due to a policy created by your administrator.`
Total times failed authentication before user account lock-out (*glide.zta.high_assurance.session.max.login.failed attempts*)	Set the maximum failed authentication attempts before the users are logged out. **Note:** The value must be between 3 and 10.

High assurance session as a Preemptive measure

Users who work with high privilege data such as financial transactions, government information, PII, can establish high assurance session as a preemptive measure to avoid frequent authentication notification during their logged in session.

High assurance session can be created by the themselves. To create a high assurance session, select **User Profile** > **Profile**. In the **Related Links** section, select **Create High-Assurance Session**. Verify your identity to create a high assurance session.

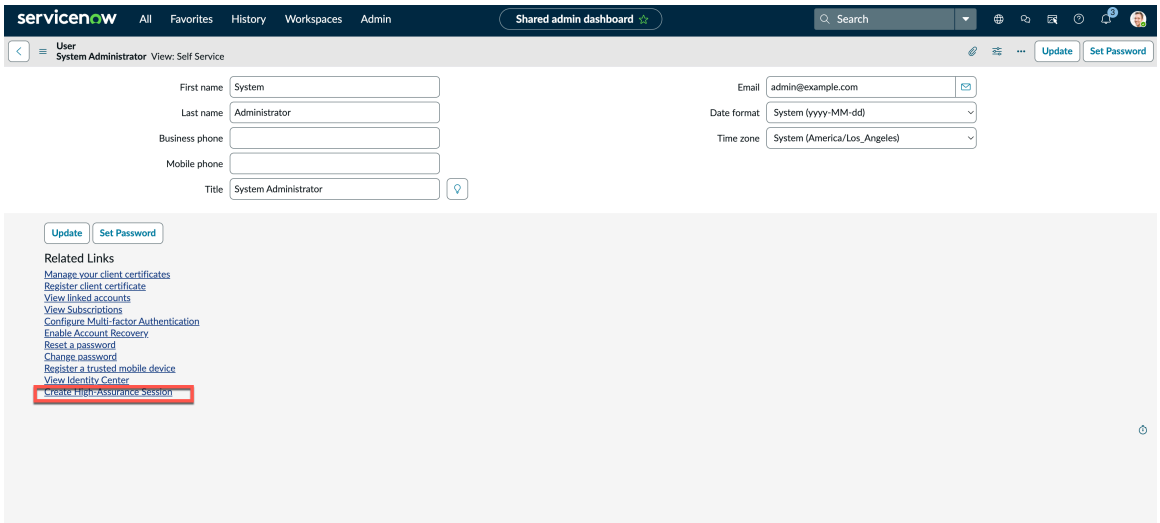

Related topics

- Explore Continuous Authentication
- Pre-work for Continuous Authentication
- Configure Continuous Authentication

High Assurance for SSO login

Establish high assurance session for SSO login using ServiceNow's continuous authentication.

A high assurance session is a session that requires a user to verify their identity and authenticate with a specific identity or Identity Providers for a specific time frame.

ServiceNow's continuous authentication (CA) feature enables you to create policies that creates a high assurance session to the users who access Personally Identifiable Information (PII), sensitive data, or restrict the access to explicit data that you want to protect.

When the user re-authenticates, there's a high assurance session that is established, which provides the ability for the users to access the data protected by the CA administrator based on the CA policy configuration.

You can create CA policies to verify the users identity and authentication the users to access the data that you've protected. You can configure the CA policy and provide the Identity Providers details in the Identity Providers record to establish high assurance session.

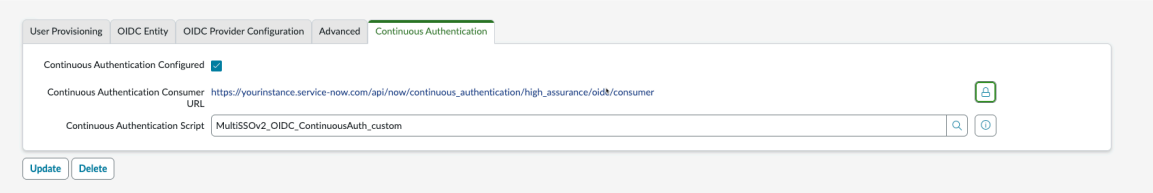

The users who are performing SSO based login (SAML or OIDC) and whenever there is an attempt to access the protected data, re-authentication screen is prompted with the same SSO that was used by the user during the initial login.

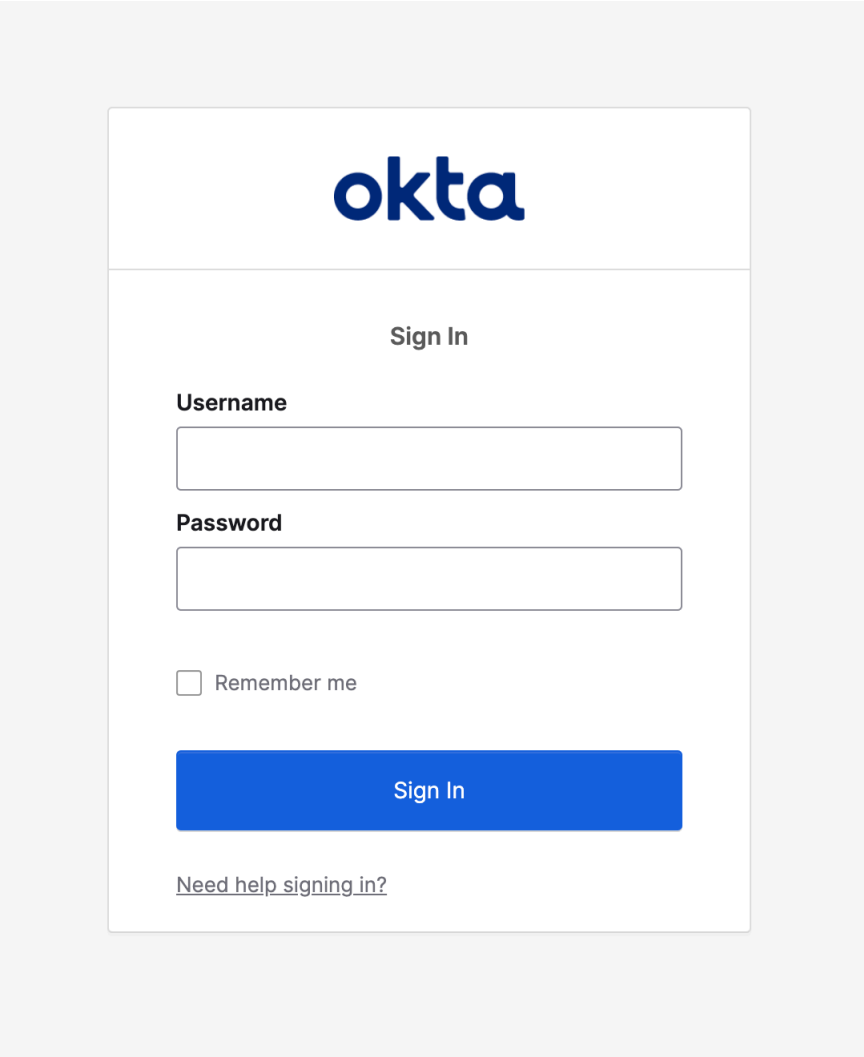

After successful SSO authentication, the protected data is displayed to the users for a certain time frame. You can configure the properties to change the time limit based on your requirement. To know more, see High Assurance session with Continuous Authentication.

Performing re-authentication with SSO login (SAML or OIDC), creates a high assurance session establishing a secure and trusted connection with the identities (users) who are accessing the protected data.

An high assurance session established for the user is limited to the High Assurance session length (*glide.zta.high_assurance.session.timeout*) system property. If the high assurance session time exceeds the property length, the user is prompted for re-authentication or step up authentication.

Related topics

- High Assurance session with Continuous Authentication
- Pre-work for Continuous Authentication
- Configure Continuous Authentication

High Assurance session for non-SSO login

Establish high assurance session for non-SSO logins (local or LDAP) using ServiceNow's continuous authentication.

A high assurance session is a session that requires a user to verify their identity and authenticate with a specific identity or Identity Providers for a specific time frame.

ServiceNow's continuous authentication (CA) feature enables you to create policies that creates a high assurance session to the users who access Personally Identifiable Information (PII), sensitive data, or restrict the access to explicit data that you want to protect.

When the user perform step-up authentication (MFA), there's a high assurance session that is established, which provides the ability for the users to access the data protected by the CA administrator based on the CA policy configuration.

You can create CA policies to verify the users identity and authentication the users to access the data that you've protected. The users who are performing non- SSO based login (local or LDAP) and whenever there is an attempt to access the protected data, step-up authentication (MFA) screen is prompted to the users.

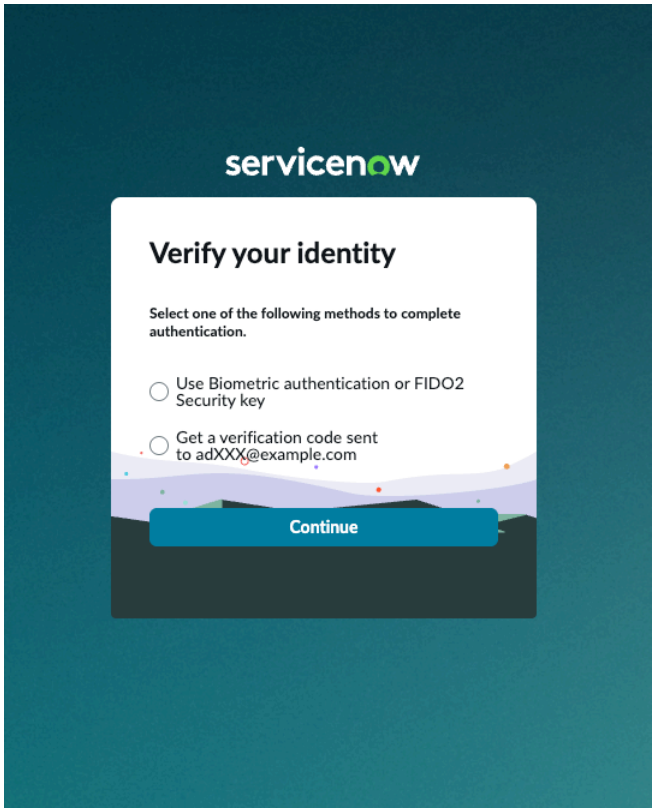

After successful authentication, the protected data is displayed to the users for a certain time frame. You can configure the properties to change the time limit based on your requirement. To know more, see High Assurance session with Continuous Authentication.

Note: If the users haven't setup MFA, it is compulsory to complete the setup.

Performing step up authentication (MFA), creates a high assurance session establishing a secure and trusted connection with the identities (users) who are accessing the protected data.

An high assurance session established for the user is limited to the High Assurance session length (*glide.zta.high_assurance.session.timeout*) system property. If the high assurance session time exceeds the property length, the user is prompted for re-authentication or step up authentication.

Related topics

- High Assurance session with Continuous Authentication
- Pre-work for Continuous Authentication
- Configure Continuous Authentication

Continuous Authentication audit logs

Describes the details about Continuous Authentication (CA) logs. The CA logs displays all the authentication attempts performed by the users.

The CA related log audit information are available in the **continuous_auth_log.LIST** page. You can type **continuous_auth_log.LIST** in the navigation to view the CA logs.

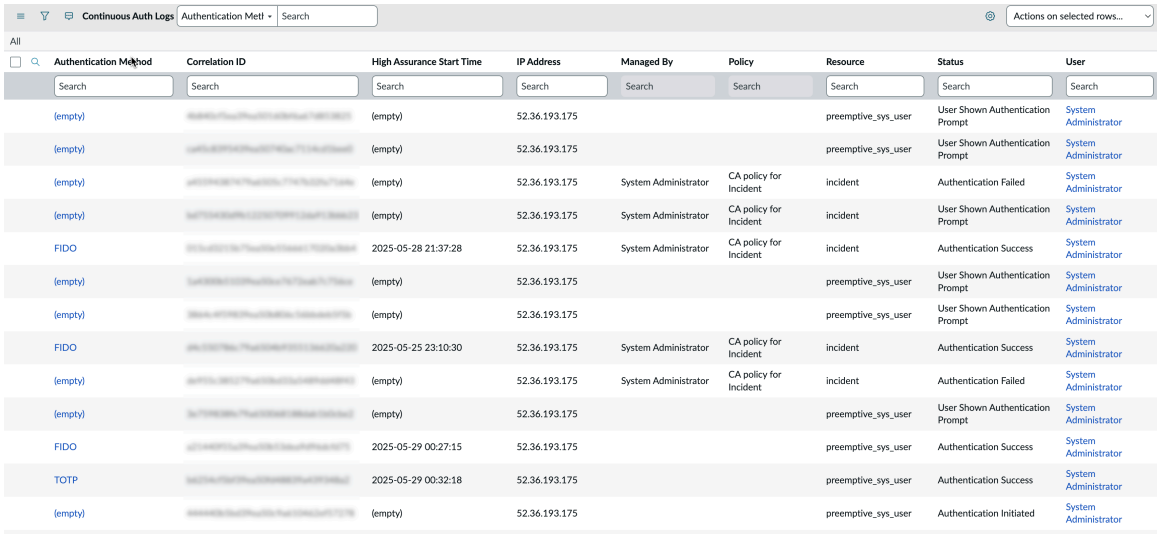

The log page displays information about the authentication method, correlation ID, with other details related to the user's session.

Domain separation for service providers

With the ServiceNow Platform, service providers (SPs) can provide their customers with faster onboarding, meet compliance, and protect their data using domain separation. You can separate client data, processes, and reports into logical groupings called domains. SPs control who sees and accesses what content.

Explore

Learn about Domain Separation.

Configure

Configure Domain Separation.

Reference

Get details about Domain Separation.

Analyze

Learn more about how to analyze Domain Separation

Exploring domain separation

With domain separation you can separate data, processes, and administrative tasks into logically defined domains.

Domain separation is best for those customers who:

- Need to enforce absolute data segregation between business entities (data separation).
- Customize business process definitions and user interfaces for each domain (delegated administration).
- Maintain some global processes and global reporting in a single instance.
- Separate data between service providers, customers, partners, or sub-organizations.
- Have minor or moderate process differences among customers.

Domain separation compared to separate instances

While domain separation provides multi-tenancy support, multi-tenancy is still contained within a single instance. Some global properties, data, and processes are shared across all domains. For example, having the system *Remember me* on the login page of the system is global and cannot be specified per domain.

If you need complete and total separation of all system properties and do not require global reporting or global processes, then separate instances are the best option.

Data separation

Members of a domain see only the data contained within their domain or the child domains that are lower in the domain hierarchy. By default, all users and all records are members of the global domain unless an administrator assigns them to a particular domain. Once you assign a user or a record to a domain, the instance compares the user's domain to the record's domain to determine whether the user can view the record.

ServiceNow applications are defined with the following incremental support levels. These levels are based on the perspective of actual use cases and personas.

Data Separation: Tenants see only data that they have permissions to see. Tenants can be granted access to other tenant data, but cannot query tenant data if they don't have access.

UI Separation: Supports a tenant-specific experience for UI elements such as views, lists, labels, and so on.

Business Logic Separation: You can create tenant-specific system policies such as email notifications, business rules, client scripts, UI policy, and UI actions.

Hierarchical Modeling: Nested-multi-tenancy so parent tenants can access child tenant resources. Business logic for parent tenants runs automatically for child tenants, and can be overridden at any level.

Cross-Tenant Intelligence (Domain Scope): Handles automatically the data, metadata, business logic, and processing context for tenants that have access to additional tenant data.

In general, data defined at a higher level in the domain hierarchy is not visible at lower levels in the hierarchy.

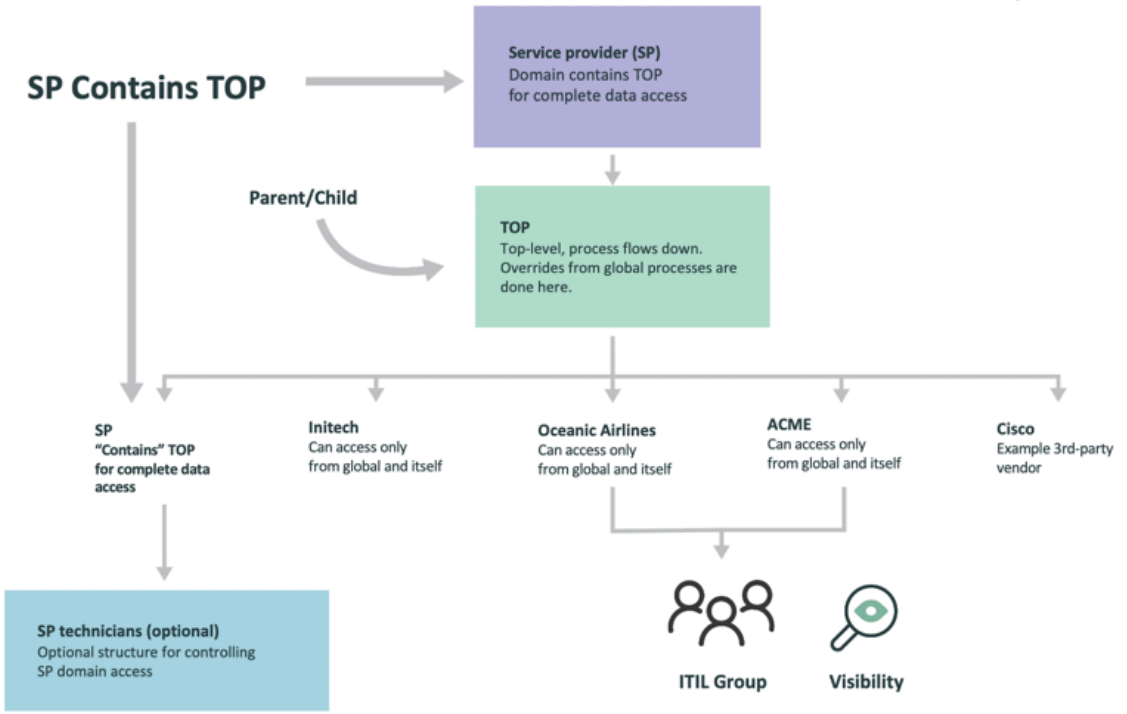

Domain path migration

Domain paths are used for all customers. Domain numbering is not used. Customer Service and Support can assist in the upgrade.

Alternatives to domain separation

Separate instances are a common alternative to domain separation. This provides a great degree of flexibility in meeting the requirements for customers and stakeholders with little to no impact on others.

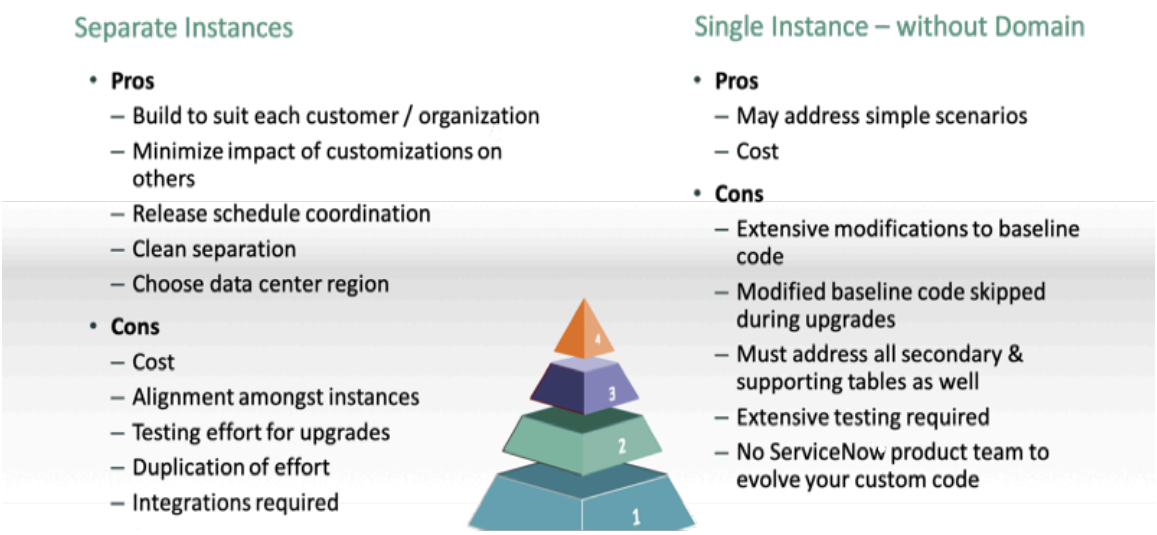

Warning: Before activating domain separation, consult your representative to verify that it is suitable for your environment. Domain separation adds a level of administration overhead. Although it can be disabled, it cannot be removed from an instance.

Related topics

Domain separation recommended practices for service providers

Domain separation plugin

Configuration that can be delegated to internal or external customers

Domain separation is designed to give ServiceNow® service providers (SPs) the ability to configure the services they offer to their customers. It is not designed to enable their customers to administer those services themselves, except in a few areas that this topic details.

overview

It is safe for SP customers, on their own, to manage data contained within their domain that does not affect licensing or other customers. For example, it is safe for a customer to create new reports or manage configuration items, but it's not safe for them to customize fields, choices, business rules, and other processes where they can impact other customers on the same instance.

The ServiceNow base system administrative roles and their access controls on the ServiceNow platform are designed for a single admin team per instance. For example, the domain_admin role is granted to one of the SP's resources to manage all domain setting for the instance, and to create new domains. For any domain-specific admin tasks, the SP should create new "customer admin" roles and access controls as needed to grant specific access to their customers.

The following image is an overview of common admin functions in varying categories of what is safe for a customer to do.

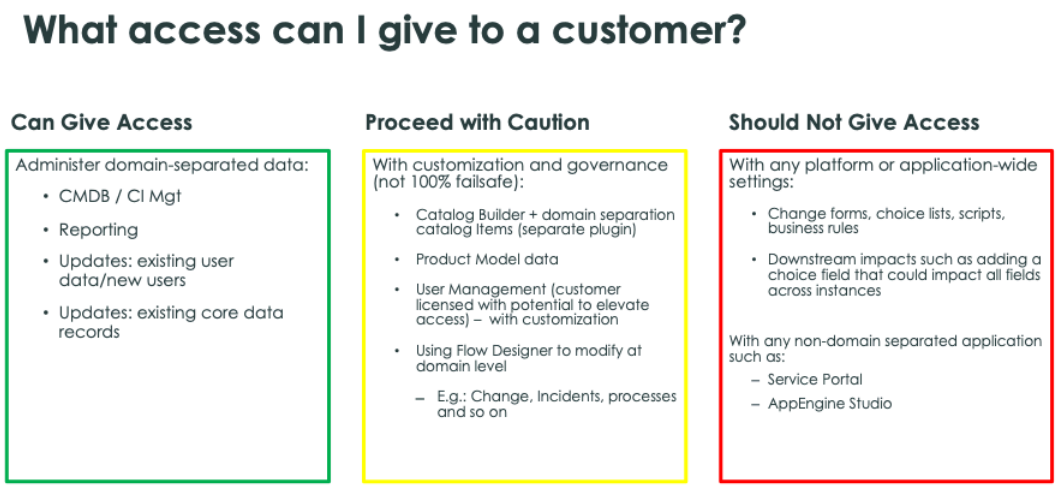

Can give access

Examples:

- CI data management in the CMDB
- Report creation
- Updates to existing user data, or new users without roles
- Updates to existing core data records such as department, group, location,cost center, or new groups without roles, and new departments/ cost centers/ location.

Proceed with caution

Examples:

- Catalog Items: To create customer-specific catalog items that can be updated by the customer, two capabilities can be used together: Domain separation for catalog items (Domain separation and Service Catalog) enables the instance owner to create items in the customer's domain. The instance owner can create a role to allow customers to update safe fields such as price, description, and images. The Catalog Builder (new in the Quebec release), gives the SP admin team the ability to create item templates that are safe to distribute to customers to create new items within their domain from within a prescriptive UI experience.
- User/Group Management: It's safe to create a "customer admin" role that can create and modify user records, but adding and removing roles can affect security and licensing. There is no way in the base system to subdivide roles that are safe for a customer to be able to grant them. The same goes for the creation and modification of groups. While the group itself can be modified, the addition or subtraction of roles should be controlled.
- Flow Designer: ServiceNow Workflow Studio is the building tool used to create process (workflow) for tables. The flow_designer role gives customers script-free access to build flows. They can read and clone every flow in domains above them in the hierarchy. They can create and modify flows in their domain. This cannot happen in a silo, however. Anyone who can affect process must be added to the global admin team for governance so processes do not cancel out each other or cause other conflicts.

Do not give access

Understanding how choice fields work is helpful to understand why only the SP admin team should be managing them.

- Structure of a choice field: Choice field values are stored in the sys_choice table and grouped based on: Table, Domain, and Language.

  For example, the **State** field in a Task is available to every table that extends a Task. That means that each table can have its own values, those values can be multiplied by domain, and the domain values can be multiplied by language.

  All of the values for **State** across all tables, domains and languages are considered the values for the **State** field.

- How changes to choice fields work: When a choice field is updated, a payload is created with all values for that field (Tables, Domains, Languages). When you install this payload on an instance, all existing values for the field are deleted and the new values are loaded. This ensures that there are no duplicate entries or leftover values that are no longer valid.

  For this reason, it's impossible to give a customer in a domain-separated instance the ability to update choice fields directly because that would affect the entire instance. In addition, you can't update choices directly in a production instance because any imported update sets that affect that field would overwrite the existing choices. In some cases, choice fields can drive processes themselves, which would break if a customer were to change those fields.

To learn more, see:

- Exploring user administration
- Create an ACL rule
- Service Provider Learning Path on ServiceNow University
- Domain separation for service providers

- Service Provider Concepts
- Application Support for domain separation
- Domain separation release notes

Domain assignment

By default, domain separation adds a domain field to tables and their extensions.

You can also extend domain separation to any new tables you create by adding a **sys_domain** field to the table's dictionary definition. By default, the system only domain-separates platform and baseline application tables where appropriate.

> ⚠ **Warning:** ServiceNow does not recommend domain-separating platform tables (any table with the sys_ prefix such as the Dictionary Entry [sys_dictionary] and Dictionary Entry Override [sys_dictionary_override] tables) because it can produce unexpected results.

Each record is assigned a single domain. That domain is stored in the **sys_domain** field. Several tables, by default, have the **sys_domain** column and are already domain separated.

The value of the **sys_domain** field contains the domain assigned to the record by any of the following:

- Company to which the user belongs
- Business rule when creating record
- Module used when creating record
- Form template used when creating record
- Domain of the parent record
- Domain assigned to User record
- Domain of the user who creates it

The system prevents the following tables from being domain separated:

- Access Control `[sys_security_acl]`
- Script Include `[sys_script_include]`
- System Property `[sys_properties]`
- Security Exclusion/Inclusion List Entities `[sys_security_restricted_list]`
- Dictionary Entry `[sys_dictionary]`
- Dictionary Entry Override `[sys_dictionary_override]`

Assigning users to companies

Administrators can quickly assign users to a domain by assigning them to a company. After users are assigned to a domain, records automatically inherit the user's domain.

For example, assigning Bow Ruggeri to the ACME company automatically assigns him to the ACME domain. Assigning Don Goodliffe to the Initech company automatically assigns him to the Initech domain. Any records they create are automatically added to the appropriate domain.

Using business rules to assign domains

Administrators can use a business rule to automatically set a domain value when creating a record. The business rule must set a value in the **sys_domain** field. Administrators must

ensure there is a **sys_domain** column available for the record's table. To learn more see Domain separation recommended practices for service providers.

Using modules to assign domains

Administrators can use the **sysparm_domain** URL parameter to automatically assign new records to a particular domain from a module. Administrators must create a module with an **Argument** value of: `sysparm_domain=sys_ID of domain`.

Using form templates to assign domains

Administrators can use a form template to automatically assign new records to a particular domain. Administrators must add the **sys_domain** field to the form and select a domain value. For example, setting the **sys_domain** field to **TOP/ACME domain** automatically assigns all records from this template to the TOP/ACME domain.

Domain inheritance on tables

By default, related records inherit the domain of the parent record. For example:

- A change task record inherits the domain of the parent change request record.
- A problem record inherits the domain of the parent incident record.

Automatic domain assignment based on user domains

If no other domain conditions apply, a record automatically inherits the domain of the user who creates it.

Visibility domains and Contains domains

Visibility domains control what a specific user or group of users can see. "Contains" domains control what an entire domain of users can see.

Visibility domains

The "Visibility domains" element determines whether users from one domain can access records from another domain. Associate this element with User [sys_user] and Group [sys_user_group] records in related lists on those records. Groups grant their members the visibility domains of the group. When a user leaves a group, they lose the group's visibility domains. Granting users a visibility domain grants all the rights to the records in that domain based on ACL (access control list) rules.

A visibility domain:

- Is a user-to-domain relationship and is explicitly granted.
- Is not a child domain.
- Is not controlled by the selection in the domain picker. Users with access to a visibility domain always see data in that domain and its child domains.

Note: Using visibility domains excessively is not recommended. Although visibility is one method to allow users to access records, it's best to use contains domains for more robust control.

Contains domains

Normally parent-child relationships define the domain hierarchy. A contains domain lets you relate domains on an as-needed basis, independent of parent-child relationships. However, contains domains grant visibility only to domain data. Processes remain unaffected by contains relationships.

A contains domain:

- Is a many-to-many, domain-to-domain relationship.
- May have child domains. When a domain is selected, you can see the data from that domain and its children.
- Is controlled by the selection in the domain picker.

Note: When you open the domain record, the scope is set to that record's domain, so you can see only child domains. Choose **Toggle Domain Scope** from the menu to populate the related list.

Contains domain example

When a user's home domain is A, and the A domain contains domains B and C, they all become peer domains. That means the user sees data from domains A, B, and C while in their home domain A. If users change domains with the domain picker to Domain B, they see only data in Domain B. When users interact with a record from Domain B or Domain C directly, they see only data for that domain.

Visibility domain example

Using domain visibility, if Don Goodliffe is in the Database domain, and Bow Ruggeri is in the Network domain, and no incidents are in the global domain, then Don cannot access Bow's incidents because of data separation.

Inheriting visibility domains based on group membership

If you set the domain table to the Group [sys_user_group] table, users can inherit visibility domains based on their group membership.

Related topics

Contains queries and domain access

Domain separation recommended practices for service providers

Domain scope

Domain scope defines what users can and cannot access.

Every user has two domain scopes when establishing a session in a domain-separated instance.

- *Session scope* is set upon session establishment to the domain listed in the user's user record. Users can manually change their session domain scope from the domain picker.
- *Record scope* uses the domain of the record and is active when viewing the form of any record.

By default, the record scope takes precedence over the session scope so that users in higher level domains adhere to each record's data and process constraints. However, these users can choose to expand or collapse the domain scope to show or hide data from other domains. For

example, a user in the Service Provider (SP) domain also has visibility into child domains such as the ACME domain. When looking at an incident record from the ACME domain, the user can choose to expand the domain scope to show values from the SP domain or collapse the domain scope to show only record values that match the record's ACME domain.

Note: Users always have access to data from domains that have been explicitly granted to them by domain visibility.

Users with the domain_expand_scope user role can select the domain scope from the **Toggle Domain Scope** UI action on the form. When record scope is in effect, click the UI action to expand to session scope and display all data available based to the user's domain and child domains. When session scope is in effect, click the UI action to collapse to record scope and display only data that matches the current record's domain.

Note: A record does not display the UI action to toggle the domain scope if the record is in the global domain or if the user's domain matches the record's domain.

Record value selection from other domains

Users who can see multiple domains have the option to select record values from a domain that is different than the record's domain.

For example, service desk agents working for a service provider might want to assign certain incidents to themselves to resolve issues on behalf of their customers. When they do this, the incident **Assigned to** field might contain a user from the SP domain, even though the incident record itself is associated with a child domain such as ACME.

Selecting a record value from another domain does not change the record's domain. The record retains its original domain. When a user views a record with values from multiple domains, the user's domain visibility determines what they see.

Record value selection

When these conditions are met	The user has access to these UI elements
The user has access to the domain of the current record referenced in a field.	The user can: • See reference field display value. For example, sees the user name in the **Assigned to** field. • See the related record from reference icon. For example, sees the user record for the user in the **Assigned to** field. • Select values from any visible domain. For example, can select users from either the SP and ACME domains.
The user does not have access to the domain of the current record referenced in a field.	The user can: • Not see reference field display values. (This is the case if domain separation was activated in Madrid or later releases and the user doesn't have access to the domain of that record.) • Only select values from the record's domain. For example, can only select users from the ACME domain.

Domains and associated companies

With domain separation you can cascade changes you make to a company record to the domain and other records associated to the company.

By default, the system automatically assigns users to the same domain as their company. For example, all users of the ACME company automatically become members of the TOP/ACME domain.

Note: Users with the admin role can change their own user records and therefore change domains. Service Providers may want to either disable delegated administration or set up an approval process to verify that the user needs the admin role.

When you change a company's domain, the instance automatically changes the domain of the following associated records to match the company's new domain.

- Locations
- Departments
- Groups
- Users

Note: The instance does not automatically change the domain of any record where you have selected the **Managed domain** checkbox.

Domain deactivation and associated companies

When you deactivate a domain, the instance also automatically completes the following actions.

- Deactivates all companies in the domain.
- Prevents all users assigned to the inactive company from logging in.

Note: When a user from an inactive company attempts to log in, the user sees an error message similar to `Company inactive - your access to this instance is not authorized`.

For example, if you deactivate the ACME domain from the sample data, the instance also deactivates the ACME company, and the three sample users are locked out.

Concepts for service providers

These concepts work with the existing ServiceNow platform capabilities to help you solve for common use cases.

Global queue v.2

The global queue concept provides a single virtual view of tasks that reside in multiple instances. The concept creates a custom application to provide a fulfiller view of work that resides in multiple instances without having to replicate tasks or data.

Overview

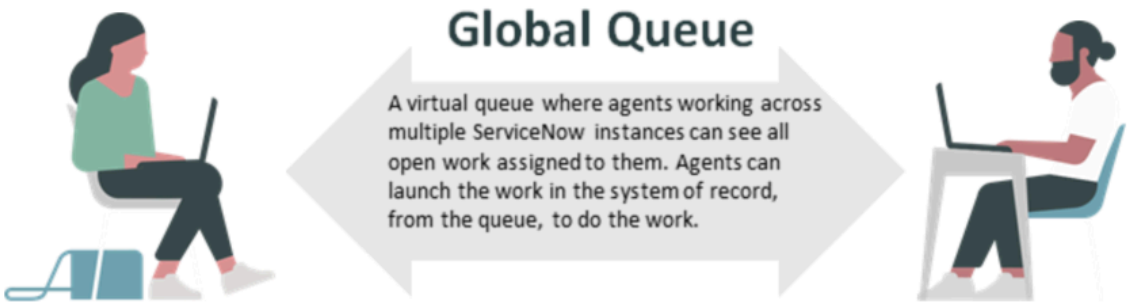

Service Providers with agents working on tasks from multiple systems tend to integrate the data back to a central instance, or a "swivel chair" between instances. While this method might be appropriate in some cases, it can be expensive and time-consuming to build and maintain. This method also opens the provider up to potential auditing and data requirement considerations such as General Data Protection Regulation (GDPR) in all of the instances where the data now lives.

Global queue v.2 is an alternative: With this method, agents can see data assigned to them from a single instance without sovereign data persisting on the instance they are logged into. For example, in cases where clients have data residency requirements, but allow access by agents from other countries, the provider could use a "follow-the-sun" Help Desk using global queue v.2.

Learn more about the Global Queue v.2 Proof of Concept on the ServiceNow Knowledge site.

Note: In the Quebec release forward, the Global Queue Proof of Concept has been upgraded to Global Queue v. 2.

Service provider connector

The service provider connector application is a reference design for creating a ServiceNow Store application for your customers to use to integrate with your systems. Service provider applications help you speed on-boarding and create standardized integrations.

Benefits of service connectors

When service providers (vendors, suppliers, partners) publish connectors, customer on-boarding is faster, which means faster billing. Integrated instances increase productivity, greater visibility in the ServiceNow ecosystem, and partner program benefits. Specific benefits are:

- Eliminates custom integrations, including the cost of services required to deliver and maintain them.
- Services are provider-defined within their ServiceNow instances and remove the need for the complexity and cost of custom integrations.
- Workflows and catalog requests can be synchronized, with the customer's processes and approvals preceding the provider's processes, so customers can follow their own processes.
- Any data created or modified (e.g. CIs) for the customer on the provider's instance can be synchronized back into the customer's instance for their visibility and use in processes.

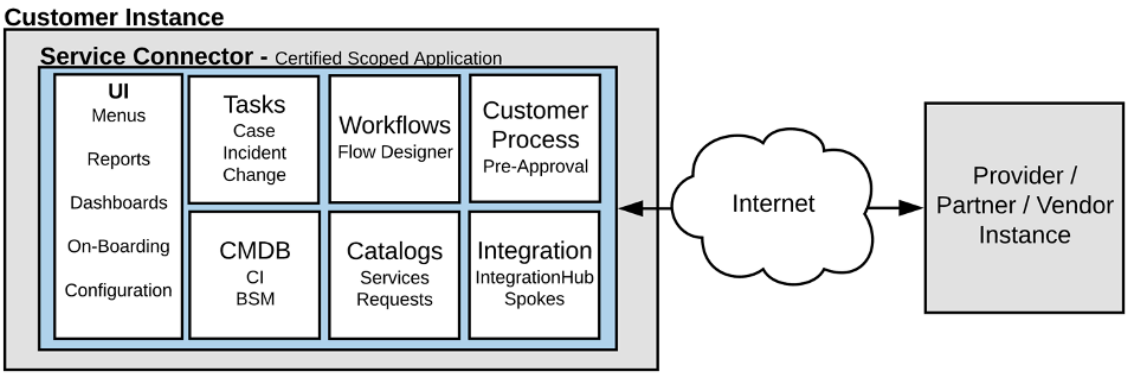

Key features

Feature	Description
User Interface	A service connector should include the following UI components at a minimum: • Reports/Dashboards – Reports and dashboards pre-defined by persona help with visibility. • Menu and modules – The customer should be able to find the application via descriptive Menu **Provider XYZ Services**. • Roles – Every connector requires user and admin roles to ensure access can be controlled by the customer. For example, “x_snc_xyz_user” or “x_snc_xyz_admin”. • On-boarding – Rapid on-boarding is essential to a good customer experience and should be included via playbook, guided setup, or catalog item in such a way that no professional services are required for the customer to go live. (Professional services are available, but on-boarding can be achieved without it). • Configuration – Data points should be built into processes to account for most of the normal process configuration a customer needs to achieve. This eliminates professional services as much as possible. • Documentation – Thorough documentation increases customer satisfaction and ease of use. • Support integration or contact info – The customer needs an easy way to stay in contact with the service provider for any issues, questions, or requests.

Feature	Description
Tasks	Use the application for tasks, such as incidents, case, changes, problems, and so on, should be predefined in the connector. These integrations are built using Workflow Studio and Integration Hub to ensure the highest level of resiliency and performance.
CMDB	Core data synchronize such as the CIs required for proper ITIL processes should be kept in sync between provider and customer instances.
Workflows	All workflows should be designed in Workflow Studio to ensure resiliency and performance.
Catalogs	The provider's catalog that the customer requests from should be included in the application as a record-producing catalog. The request generated by the item in the customer's instance should be eBonded with the provider's instances. The provider's workflows keep the request updated and synchronize back to the customer's instance.
Customer Processes	Any mechanisms for syncing requests with the provider's instance should allow for customer processes to interact with the requests before they are sent to the provider's instance. During provider processing, approvals can be sent to the customer's instance as needed.
Integration	Integrations should be built using Workflow Studio and Integration Hub to ensure the highest level of resiliency and performance.

Possible components you can include in a Service Connector

Component	Description
Instance data replication (IDR)	When replication is the goal: • Can be used for process integration, but can be too rigid depending on the complexity integration logic based on state transition • To learn more, see Instance Data Replication
IntegrationHub	When process integration is the goal: • Easier to inject in the middle of a process, as part of complex or conditional step in the middle of a flow • To learn more, see IntegrationHub

Component	Description
Global Work Queue (Virtual)	When task federation is the goal, and storing data externally is not acceptable: • Used where agents are working across multiple ServiceNow instances and need to see all open work assigned to them • Rows returned should be limited to less than 1000 • To learn more, see Global work queue
Remote Tables	When external data usage without storage is the goal: • Used to connect an instance to third-party sources, or to another instance, to retrieve external data and optionally cache it in the memory. The data is treated as a table in the instance for read-only purposes such as group, sort, aggregate, and filter. • To learn more, see Retrieving external data using remote tables and scripts
Flow Designer	When process design is the goal • Used for automating processes in a single design environment. Process owners can use natural language to automate approvals, tasks, notifications, and record operations without coding. • To learn more, see Flow Designer

Learn more about the Service Connector proof of concept on the ServiceNow Knowledge site.

Installed with domain separation

Several platform components are added or modified with domain separation.

Roles

Role	Description
domain_admin	Can create, edit, and delete domains.

Additions to [sys_domain] fields

The sys_domain field is added to the following tables:

Tables with the sys_domain field

Tables
sys_attachment
sys_user_has_role
sys_group_has_role
sys_email
sys_user_group
core_company
cmn_location
cmn_department
sys_gauge
sys_report
kb_feedback
sysapproval_approver
sys_user_grmember

Field for the Task Table

MSP Extensions add a task_for field to the Task table. This reference field refers to the User table.

Options for the Group Type

MSP Extensions add several new default options to the type field of the Group table. Add to or update these types as needed to support your domains.

Tables
Security
Support
Visibility

Business rules

Name	Table	Description
Domain - Activate/ Deactivate	core_company	Activates the related domain if at least one of its companies is active. Deactivates the related domain if all related companies are inactive.
Domain - Cascade Company	core_company	Keeps a company's domain in sync with its users, groups, departments, and locations.
Domain - Cascade	sys_email	Keeps an email's domain in sync with its attachments.

Name	Table	Description
Domain - Email		
Domain - Cascade Domain - Group	sys_user_group	Keeps a group's domain in sync with its inherited roles (sys_group_has_role records).
Domain - Cascade Domain - Knowledge	kb_knowledge	keeps a knowledge article's domain in sync with its related feedback.
Domain - Cascade Domain - Task	task	Keeps the domain in sync with related tasks for wf_context, wf_executing, wf_history, attachments, emails, task_sla and its workflow, sysapproval_approver and its workflow, and sysapproval_group and its workflow.
Domain - Cascade Domain - User	sys_user	Keeps a user's domain in sync with its group membership (sys_user_grmember) and role (sys_user_has_role) records.
Domain - Cascade Domain - Version	wf_workflow_version	Keeps domains in sync with related workflow versions for wf_activity and wf_transition.
Domain - Deactivate Companies	domain	Deactivates related companies if a domain is deactivated.
Domain - Default - Task	task	Sets the task domain based on the Task for user's domain. If this domain would be global, sets domain to Default instead.
Domain - Default - User	sys_user	Sets a user's domain to Default if the domain otherwise would have been global.
Domain - Disallow Global Domain Record	domain	Prevents creation of a domain with the name global.
Domain - Override Copy	sys_app_application	When an application is overridden for a domain, creates a copy of its modules for the new application.
Domain - Override Copy	sys_data_policy2	When a data policy is overridden for a domain, creates a copy of its data policy rules for the new data policy.
Domain - Override Copy	sys_gauge	When a gauge is overridden for a domain, creates a copy of its gauge counts for the new gauge.

Name	Table	Description
Domain - Override Copy	sys_ui_action	When a UI action is overridden for a domain, creates a copy of its UI action views for the new UI action.
Domain - Override Copy	sys_ui_list_control_embedded	When an embedded list control is overridden for a domain, creates a copy of its client and server scripts for the new embedded list control.
Domain - Override Copy	sys_ui_policy	When a UI policy is overridden for a domain, creates a copy of its UI policy actions for the new UI policy.
Domain - Set Domain - Approvals	sysapproval_approver	Sets the domain based on that of the record being approved.
Domain - Set Domain - Attachment	sys_attachment	Sets the domain based on the parent record's domain.
Domain - Set Domain - CMDB_CI	cmdb_ci	Sets a CI's domain to that of its company.
Domain - Set Domain - Department	cmn_department	Sets a department's domain to that of its company.
Domain - Set Domain - Domain	domain	Sets a domain's domain to itself.
Domain - Set Domain - Email	sys_email	Sets the domain based on the parent record's domain. An email's parent record is the record specified in the instance field.
Domain - Set Domain - Feedback	kb_feedback	Sets a knowledge feedback's domain to that of its knowledge article.
Domain - Set Domain - Group	sys_user_group	Sets a group's domain to that of its company.
Domain - Set Domain - Group Approvals	sysapproval_group	Sets the domain based on that of the record being approved.
Domain - Set Domain - Group Role	sys_group_has_role	Sets a group role's domain to that of its group.
Domain - Set Domain - Location	cmn_location	Sets a location's domain to that of its company.

Name	Table	Description
Domain - Set Domain - Task SLA	task_sla	Sets a task SLA's domain to that of its task.
Domain - Set Domain - User	sys_user	Sets a user's domain to that of its company.
Domain - Set Domain - User Role	sys_user_has_role	Sets a user role's domain to that of its user.
Domain - Set Domain - WF Activity Hist	wf_history	Sets the workflow activity history domain based on the parent workflow context's domain.
Domain - Set Domain - WF Context	wf_context	Sets the workflow context domain based on the referenced record's domain, if it has one.
Domain - Set Domain - WF Exec Activity	wf_executing	Sets the workflow executing activity domain based on the parent workflow context's domain.
Domain - Set task for - Change	change-request	When converting a ticket to a change request, sets the Requested by field to the ticket's Task for value.
Domain - Set task for - Incident	incident	When converting a ticket to an incident, sets the Caller field to the ticket's Task for value.
Domain - Validate Default	domain	Ensures only one domain has the Default check box selected.
Domain - Validate Primary	domain	Ensures only one domain has the Primary check box selected.
Business Rules Installed with Domain Support Plugin		
Change Domain Set	sys_dictionary	Sets the domain set to the current domain.
Domain support properties	sys_properties	Sets the system properties to match the domain query method (domain paths or domain numbering).

Client scripts

Client script	Description
Domain - Set Company	Monitors the incident caller field for changes. If the company and location fields do not already have a value, the script adds this information from the

Client script	Description
and Location (sys_script)	caller record. If the company and location fields already have a value, the script retains the existing values.
Deactivated script	
(BP) Set Location to User	Monitors the incident location field and sets the location field to the caller's location.

Related topics

Domain separation recommended practices for service providers

Application support for domain separation

Many ServiceNow applications support domain separation in the base system but not all. Some supported applications include limitations on the data and administrative settings that can be domain-separated. These definitions delineate the domain separation support levels from the perspective of actual use cases and the people who use them.

Domain separation support levels

ServiceNow applications that support domain separation may support the separation of data and data routing only, have advanced business logic separation, or support tenant (customer) level administration of the application. ServiceNow applications are defined with the following incremental support levels.

No support

- The domain field may exist on data tables, but no logic exists to manage data.
- This level is not considered domain-separated.

Basic

- Business logic: Ensure data goes into the proper domain for the application's service provider (SP) use cases.
- In the application, the user interface, cache keys, reporting, rollups, aggregations, and so on, all use domain at production run time.
- The owner of the instance must be able to set up the application to function across multiple tenants.

Sample use case: When an SP uses chat to respond to a tenant-customer's message, the client must be able to see the SP's response.

Standard

- Includes **Basic** level support.
- Business logic: Processes can be created or modified per customer by the service provider (SP). The use cases reflect proper use of the application by multiple SP customers in a single instance.
- The owner of the instance must be able to configure the minimum viable product (MVP) business logic and data parameters per tenant as expected for the specific application.

Sample use case: An admin must be able to make comments mandatory when a record closes for one tenant but not for another.

Enhanced

- Includes **Basic** and **Standard** levels.
- Data-driven process enables service provider customers to modify business logic that is based on defined use cases. These configurations are UI-based and fail-safe so that configurations by one customer cannot affect another.
- Tenants of the instance must be able to configure minimum viable product (MVP) business logic and data parameters for themselves. This logic and parameters would be expected for the application's normal function.

Sample use case: Tenant-customers of a shared environment must be able to change to the impact, urgency, or priority matrix to set priority within their domain.

Note: Effective Domain (*)

Sometimes, a platform feature or application may effectively support SP use cases even without the domain framework. If so, the use cases must detail its support of domain separation. An asterisk (*) after the support level indicates this kind of configuration.

Supported feature	Basic	Standard	Enhanced
Domain column is present for base system application tables.	●	●	●
Domain-specific configuration is managed by instance owner.	●	●	●
Tenant domains can manage their own application data.		●	●
Application properties are domain aware when needed.		●	●
Business logic and processes can be domain-separated by instance owner.		●	●
Business logic and processes can be administered by the tenant domain.			●

Support levels by application

Product Suite	Application	Support level
App development, vibe coding, and low-code	App Engine Studio	No support
	Automation Center	Basic
	Robotic Process Automation (RPA) Hub	Basic
	ServiceNow Studio	No support
	Workflow Data Fabric Hub / Zero Copy Connectors	No support
	Table Builder	Basic
	App Engine Management Center	No support
	Exploring Decision Tables	Standard

Product Suite	Application	Support level
	Enterprise Resource Planning Integration	No support
	Enterprise Resource Planning Customization Mining	No support
	Next Experience UI Builder	Basic
Customer Service Management	Communities	No support
	Customer Service Management	Basic
	Release Management	Basic*
	Order Management for Customer Service Management	Basic
	Post-Sales Support	Basic
	Domain separation in Workforce Optimization for Customer Service	Basic
	Now Assist for CSM	Basic
DevOps	Domain separation and DevOps Change Velocity	No support
	Domain separation and DevOps Config	
Employee Service Management	HR Service Delivery	Basic*
	Health and Safety	No support
	Contract Management Pro	Basic
	Legal Service Delivery	Basic
	Contract Management Pro	Basic
	Procurement Service Management (PSM)	No support
	Safe Workplace Suite	See application site for individual application support levels
	SharePoint Online Search Connector	Basic
	Universal Request	Basic
	Universal Task	Basic
	Workforce Optimization for HR	Basic
Finance and Supply Chain	Sourcing and Procurement Operations	No support*
	Purchase Order Management	No support*
Operational Sustainability Management (formerly	ESG Management and Reporting	No support*

Product Suite	Application	Support level
Environmental, Social, and Governance)		
Field Service Management	Field Service Management	Basic
Governance, Risk, and Compliance	Business Continuity Management	Basic
	Governance, Risk, and Compliance (GRC)	Basic
	Operational Resilience	Basic
Industry Products		
•Financial Services	Financial Services Card Operations	Basic
	Financial Services Deposit Operations	Basic
	Financial Services Payment Operations	Basic
	Intelligent Servicing for Fraud	Basic
	Property and Casualty Insurance Servicing	Basic
	Life Insurance Servicing	Basic
	Insurance Claims	Basic
	Financial Services Know Your Customer	Basic
	Financial Services Credit Operation	Basic
	Financial Services Document Processor	Basic
• Healthcare and Life Sciences	EMR Help	Basic
	Healthcare and Life Sciences Service Management Core	Basic
	Pre-Visit Management	Basic
	Patient Support Services	Basic
	Vaccine Administration Management	Basic
Manufacturing Commercial Operations	Manufacturing Commercial Operations	Standard
	Retail Core	Basic
	Retail Task Management	Basic
• Manufacturing	Industrial Process Manager	Standard

Product Suite	Application	Support level
	Operational Technology Manager	Standard
	Operational Technology Vulnerability Response	Standard
	Operational Technology Manager	Standard
Telecommunications, Media, and Technology (TMT)	Customer Success Management	Basic
	Customer Service Problem Management	Basic
	Now Assist for Telecommunications, Media and Technology (TMT)	Basic (Inherited from Domain separation in the Now Assist Admin console).
	Proactive Service Experience Workflows	Standard
	Service Bridge	Standard
	Product Support for Technology application	Basic (Inherited from Customer Service Management).
	Telecommunications Network Inventory	Basic
IT Asset Management	Cloud Insights	No support
	Hardware Asset Management	Enhanced
	Software Asset Management	Enhanced
	Enterprise Asset Management	Standard
Strategic Portfolio Management	Agile Development	Basic*
	Alignment Planner Workspace	Basic
	Application Portfolio Management	Basic
	Cost Management	No support
	Demand Management	Basic
	Financial Management	No support
	Investment Funding	Basic
	Project Portfolio Management	Basic*
	Release Management	Basic*
	Scaled Agile Framework (SAFe)	Basic*
	Test Management	Basic*
	Goal Framework	Basic
IT Operations Management	Cloud Provisioning and Governance	Basic

Product Suite	Application	Support level
	Agent Client Collector	Basic
	Discovery	Standard
	Event Management	Basic
	Service Operations Workspace for ITOM	Basic
	Health Log Analytics	Basic
	Metric Intelligence	Basic
	Service Mapping	Basic
	Cloud Migration Assessment	Basic
	Action Library	No support
	Cloud Configuration Governance	No support
	Tag Governance	Basic
	Cloud Insights Billing	No support
	Cloud Provisioning and Governance: Google Cloud	Basic
	Cloud Provisioning and Governance Terraform	Basic
	Cloud Operation Workspace	Basic
	Cloud Discovery	Standard
IT Service Management	Benchmarks	No support
	Change Management	Basic
	Coaching	Basic
	Continual Improvement Management	Basic
	Contract Management	No support
	Domain separation and Digital Product Release	Basic
	Expense Line	No support
	Incident Communications Management	Standard
	Incident Management	Standard
	Facilities Service Management	Standard
	Incident Management	Standard
	On-Call Scheduling	Standard
	Asset Management	Basic
	Problem Management	Standard
	Procurement	Standard*

Product Suite	Application	Support level
	Product Catalog	Standard
	Request Management	Standard
	Service Catalog	Standard
	Service Level Management	Basic
	Service Portfolio Management	Basic*
	Site Reliability Operations	Basic*
	Task outage	Basic
	Domain separation and Vendor Management Workspace	No support
	Walk-up Experience	Basic
Mobile Configuration and Navigation	Mobile	Basic
Now Intelligence	Dashboards	Basic
	Performance Analytics	Enhanced
	Process Optimization	Basic
	Reporting	Basic
	User Experience Analytics	Basic
The ServiceNow AI Platform	Administration	
	Domain separation and Agent Chat	Standard
	ServiceNow AI Platform Capabilities	
	User Interface	
	Advanced Work Assignment	Standard
	AI Search	Searches respect domain restrictions from indexed records
	App Engine Studio	No support
	Application Management	No support
	Assessments	Standard
	Automated Test Framework	Standard*
	ServiceNow Voice	
	Code Signing	Basic support
	Contextual Search	Standard
	Configuration Management (CMDB)	Standard
	Content Management System	No support
	Credentials and Connections	Standard

Product Suite	Application	Support level
	Data Certification	Basic*
	Data Classification	Enhanced
	Data Privacy	No support
	Data Management	Basic*
	Delegated Development	No support
	Dependency Views	Basic
	Document Services	No support
	Dynamic Translation	Basic
	Edge Encryption	Basic support
	External Content Connectors	No support*
	Field Encryption	No support
	Encryption	No support
	Cloud Encryption with Key Management	Basic support
	Field Normalization	No support
	Flow Designer	Standard*
	Guided Setup	No support
	Domain separation and Integration Hub	Standard*
	Integrations with third-party applications and data sources	Basic+Standard
	Knowledge Management	Standard
	Hermes Messaging Service	No support
	Managed Documents	No support
	MetricBase	Basic
	Natural Language Understanding	Basic+Standard
	Notifications	Standard
	ODBC Driver	Basic*
	Orchestration	Standard*
	Password Reset	Standard
	Platform Security	Domain separation landing page
	Data Privacy	No support
	Predictive Intelligence	Standard
	Proactive Triggers	Basic
	Process Automation Designer	Basic
	Remote Tables	No support

Product Suite	Application	Support level
	Schedules	Basic
	Script debugger	Basic
	Search Suggestions	No support
	Service Portal	No support
	Service Graph Connectors	No support
	Domain separation and Sidebar	Standard
	State Flows	No support
	Subscription Management	Basic*
	Survey Management	Basic*
	Task Intelligence	No support
	Domain separation and Time Card	Basic*
	UI Builder	Standard
	Virtual Agent	Basic
	Visual Task Boards	Basic
	Web Services	Standard*
	Workflow	Standard*
	Workspace	Standard
Security Operations	Configuration Compliance	Standard
	Configuration Data Management	Basic
	IBM QRadar Offense Ingestion	Basic
	Microsoft Graph Security API alert ingestion integration	Basic
	Security Incident Response	Standard
	Threat Intelligence	Standard
	Vulnerability Response	Standard
Service Management	Facilities Service Management	Standard
	Planned Maintenance	Standard*
	Proactive Triggers	Basic
	NOW Code Editor	No support
	Workforce Optimization for ITSM	Basic
	Vendor Management Workspace	Basic

Domain separation recommended practices for service providers

You can create, implement, and maintain domain separation for your applications and services.

Domain basics

With domain separation (also known as ServiceNow Multitenant Platform Architecture), you can segregate the application data, user interface, and business logic in a single customer instance that supports hierarchical modeling with cross-tenant intelligence. Business logic describes how the domain separation is configured and what rules are affecting the configuration.

Before you set off on the domain separation journey, here are some good practices to follow. Select topics as you want or follow them in order by clicking the links below the image.

Domain separation explained

With domain separation, you can segregate application data, UI, and business logic, such as rules or workflows, in a single customer instance. Separating these elements into logically defined domains supports specific hierarchies for all customers using your applications.

Domain basics

Domain separation, also known as ServiceNow multitenant platform architecture, adds considerable overhead to the management of an instance. If you use domain separation correctly though, it can improve efficiency, add greater security, and increase the performance of your customers' instances.

You can't separate some global standards and properties, such as system properties and table schema, per tenant.

Before you start separating domains, read the following guidelines.

What you can do with domain separation

- Data separation: Enables tenants of the domain to see only data that they have permissions to see. Tenants can be granted access to other tenant data but can't query tenant data that they don't have access to.
  - When you update data records, they do not generate Update Set records.
  - Users, including the customer accounts that are used for integrations, see only the data in the domains they have permission to access.
  - Customers, agents, and fulfillers see data that pertains to the customers and organizations that they support.
- UI separation: Supports a tenant-specific experience for UI elements such as views, lists, labels, and so on.
  - You can override the browser-based user interface, including application menus, lists, forms, and dashboards. You can also customize them for a specific domain or set of domains while preserving your basic process logic.
  - Service providers can alter the displayed branding and UI elements to meet individual customer needs.
- Business logic separation: Creates tenant-specific system policies such as email notifications, business rules, client scripts, UI policy, and UI actions.
- Hierarchical modeling: Nests your multiple tenants so that parent tenants can access child tenant resources. Business logic for parent tenants runs automatically for child tenants, which you can override at any level.
- Cross-tenant intelligence: Automatically handles data, metadata, business logic, and processing context for tenants with access to additional tenant data.

Domain separation at a glance

The following graphic shows the division of data, process, and UI separation. These concepts are discussed in depth in the Recommended Practices section.

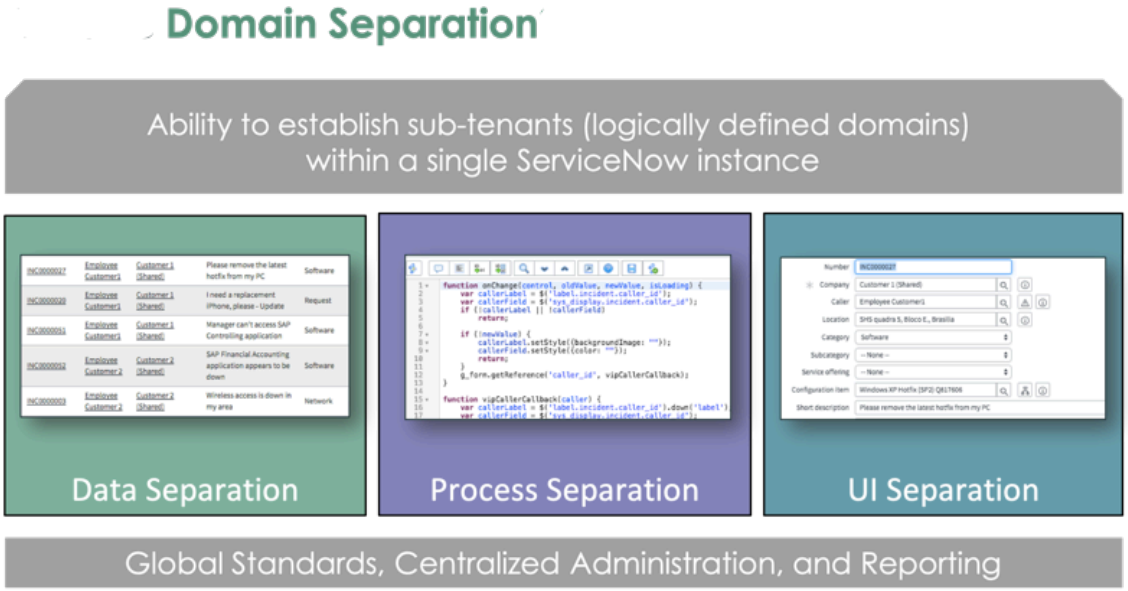

Domain architecture

User records are assigned a domain value that represents the user's home domain. Users have no access to data in parent domains, peer domains, or domains in other branches of the hierarchy.

See Contains queries and domain access for advanced options to grant additional domain visibility.

The following diagram shows how the architecture process flows down to the child

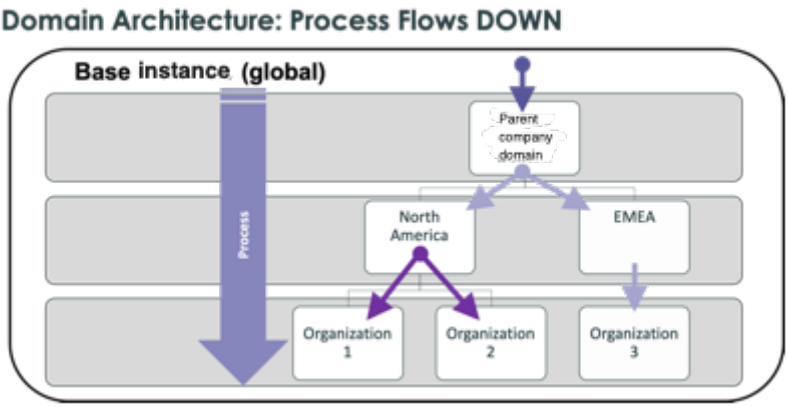

domains.

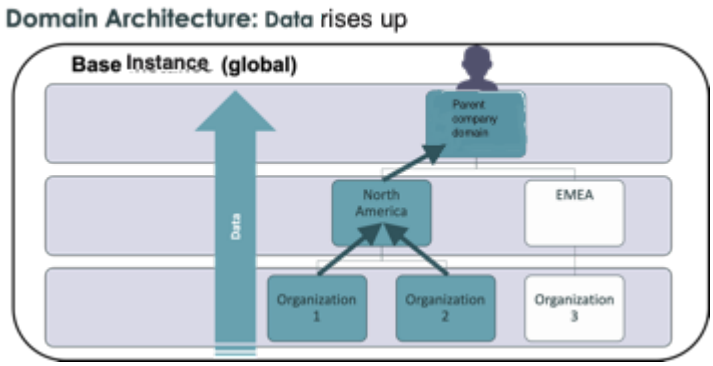

Domain separation value proposition

With domain separation, service providers can have a multitenant instance architecture that delivers offerings efficiently and securely to their clients. Strong universal process standards, data-driven process design, strict governance, and centralized administration help to maximize these benefits.

Benefits of domain separation

The tenants of a domain benefit from a quick return on investment, lower administrative overhead, and leverage of business services provided by instance owners.

Here is an quick view of these benefits.

Instance Owner	Domain Tenants
Service provider employee productivity	Increased security
Only process deltas maintained	Pre-built processes and features
Administration efficiencies	Reduced staff required
Fewer client integrations	Faster onboarding
Upgradability and scalability	Leverage of the latest releases
Data segregation	Services provided by the instance owner
Global reporting	

Definition of domain separation

With domain separation (also known as the ServiceNow® Multitenant Platform Architecture), you can segregate application data, UI, and business logic in a single customer instance that supports hierarchical modeling with cross-tenant (customer) intelligence.

Properties of domain separation

Domain-separated ServiceNow applications are defined with the following properties:

Data separation

Enables customers to see only the data that they have permissions to see. Customers can be granted access to other customer data but cannot query customer data if they don't have access.

UI separation

Supports a customer-specific experience for UI elements such as views, lists, labels, and so on.

Business logic separation

Supports customer-specific system policies such as email notifications, business rules, client scripts, UI policy, and UI actions.

Hierarchical modeling

Supports nested-multitenancy so parent tenants (customers) can access child customer resources. The business logic for parent customers runs automatically for child customers and can be overridden at any level.

Cross-customer intelligence (domain scope)

Handles automatically the data, metadata, business logic, and processing context for tenants that have access to additional tenant data.

Domain separation hierarchies

Create a hierarchy when defining a domain architecture to track your processes and workflows.

Sample domain separation hierarchies

The following diagram is a good starting point for defining domain architecture. It demonstrates the relationship between the TOP and lower domains and how the process, data, and business rules impact parent and child domains.

- In the following example, TOP is a process domain. It should never contain users. Rather, TOP should contain the new processes that instance owners develop and the overrides to these processes from the global domain.
- Only the service provider (SP) has access to the default domain. This domain never contains active users. It contains only the "lost" data that you need to reassign to the correct domain.

  Note: When data is not assigned to a specific domain, it moves to the default domain. It is temporarily "lost" and needs to be assigned to its correct domain.

- Tasks and users without a domain are placed in the default domain automatically when you create or update domains. You can override that action by either clearing the **Default** option on this record or selecting the **Default** option on another domain record. If you have not set a default domain yet, tasks and users with no domain move to the global domain.
  - Don't move data between domains while you are using the instance.
  - If any data ends up in the default domain, that means you have a configuration or procedural problem to address.

You don't see the word "Global" in this diagram because there is no global domain. Remember that "global" is the absence of a domain on a record.

For example, a table that has no domain field means that the table contains all global records. A table that has a domain field means that any record without a domain is a global domain.

The word "global" is in the domain field. It is placed there automatically when the record has no domain.

Global records are available to all users of the instance unless they are restricted by security configurations.

- Use the default domain to make sure that records do not end up in the global domain on tables that should never have global records.
- Instance owners must then triage the records in the default domain and move them to the correct domain.

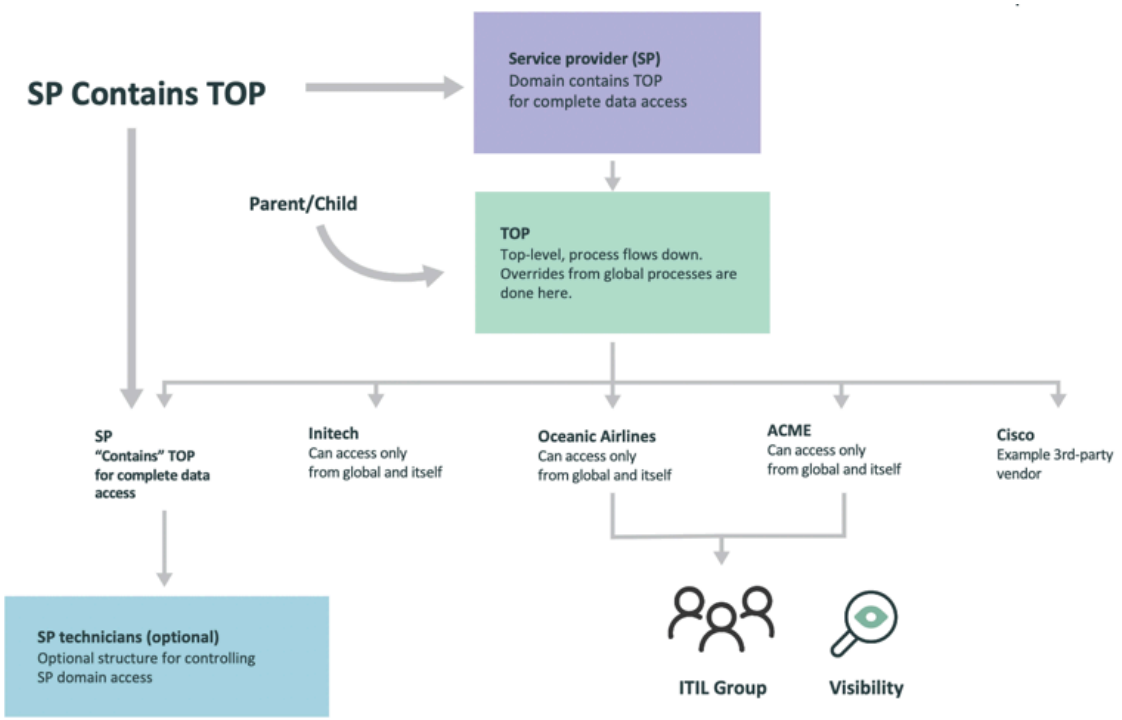

Domain hierarchies

- Parent/child: Process and data affected
  - Design that is based on a process flow.
  - Remember that the parent domains can access all data in the child domains.
- "Contains" domain: Only data is affected. For example, making SP in the diagram contain TOP does not make processes in SP run in the TOP domain and downward.
  - Grants data access rights to individuals in groups that require dedicated access to certain domains.
  - Contains causes or conditions to be added to database queries that can cause performance issues with large domain and data sets.
- Visibility: Hierarchy that is always visible to users once you provide access. Only the data is affected, not the processes.
  - Grants data access of a domain to another domain that did not have that access when the parent-and-child hierarchy was built.
  - Enables users to see all the data in the domains that they have visibility access for, all the time, regardless of the record they are working on.

Note: Use sparingly, as Visibility can allow complete access that you may not intend.

Basic principles of defining a domain hierarchy

Unrestricted and restricted use cases for domain separation.

Many SPs have customers who implicitly state that access to their domains must be tightly regulated, which constrains the use of the “contains” function at the TOP domain. The following diagram explains how to mitigate that regulation by dividing domains into Restricted and Unrestricted domains.

1. Customers exist in a specific "vertical" of the domain separation hierarchy. This means they only consume processes defined in their domain and all parent domains above theirs in the hierarchy. Any processes defined in domains that are not in their linear parent-child hierarchy do not apply.

   Note: Customers or "tenants" are entities that are segregated from each other completely, not like departments or business units that share resources with each other.

2. Super verticals (restricted, manager services, and so on) are allowed as long as the customers only ever belong to one of them.
3. Services, products, or offerings that need to be horizontally available to all customers are not defined within separate domain hierarchies.

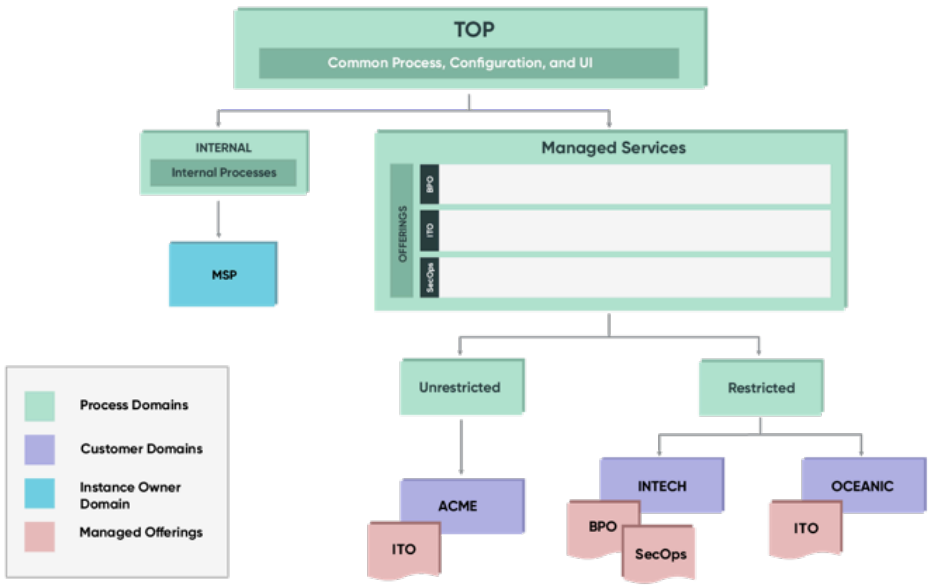

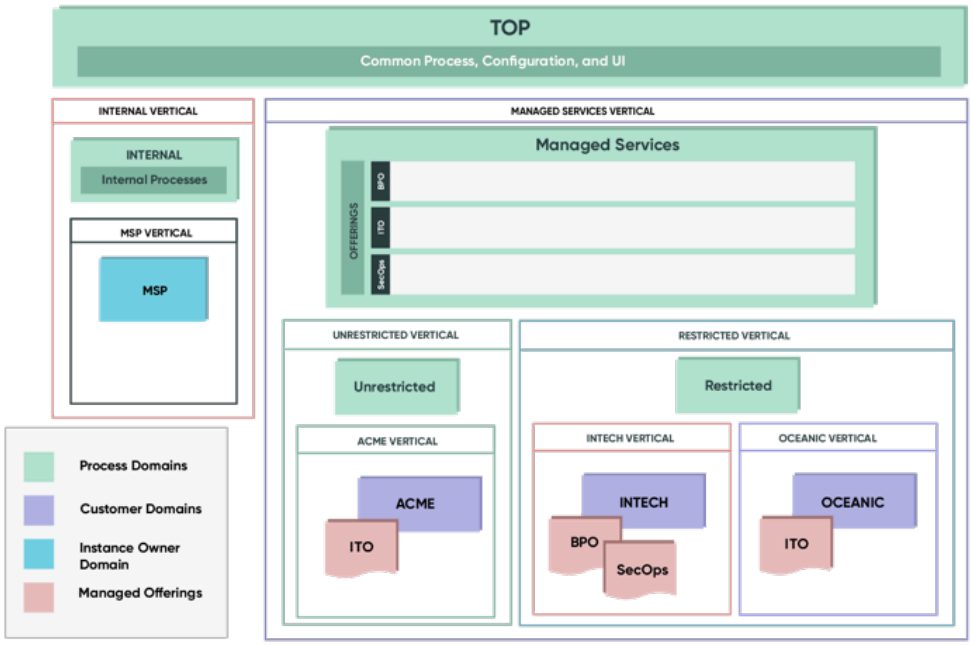

Here are some sample use cases:

- Under TOP, you might create two domains, Unrestricted and Restricted.
  - Place customers and their domains that don't have SP visibility restraints under Unrestricted.
  - Place customers and their domains that do have that requirement under Restricted.
- System admins can then use “contains” and “visibility” functions in an efficient, targeted manner.
  - Apply "contains" to Unrestricted, so a single “contains” can grant visibility to most customers.
  - Apply domain visibility using ”domain visibility groups” to specific domains as needed.

Customer Decision Trees

The following diagram delineates how to choose which hierarchy model is right for you. You can choose separate hierarchies, hybrid, or shared hierarchies depending on which processes and functionality you want in your domain structures.

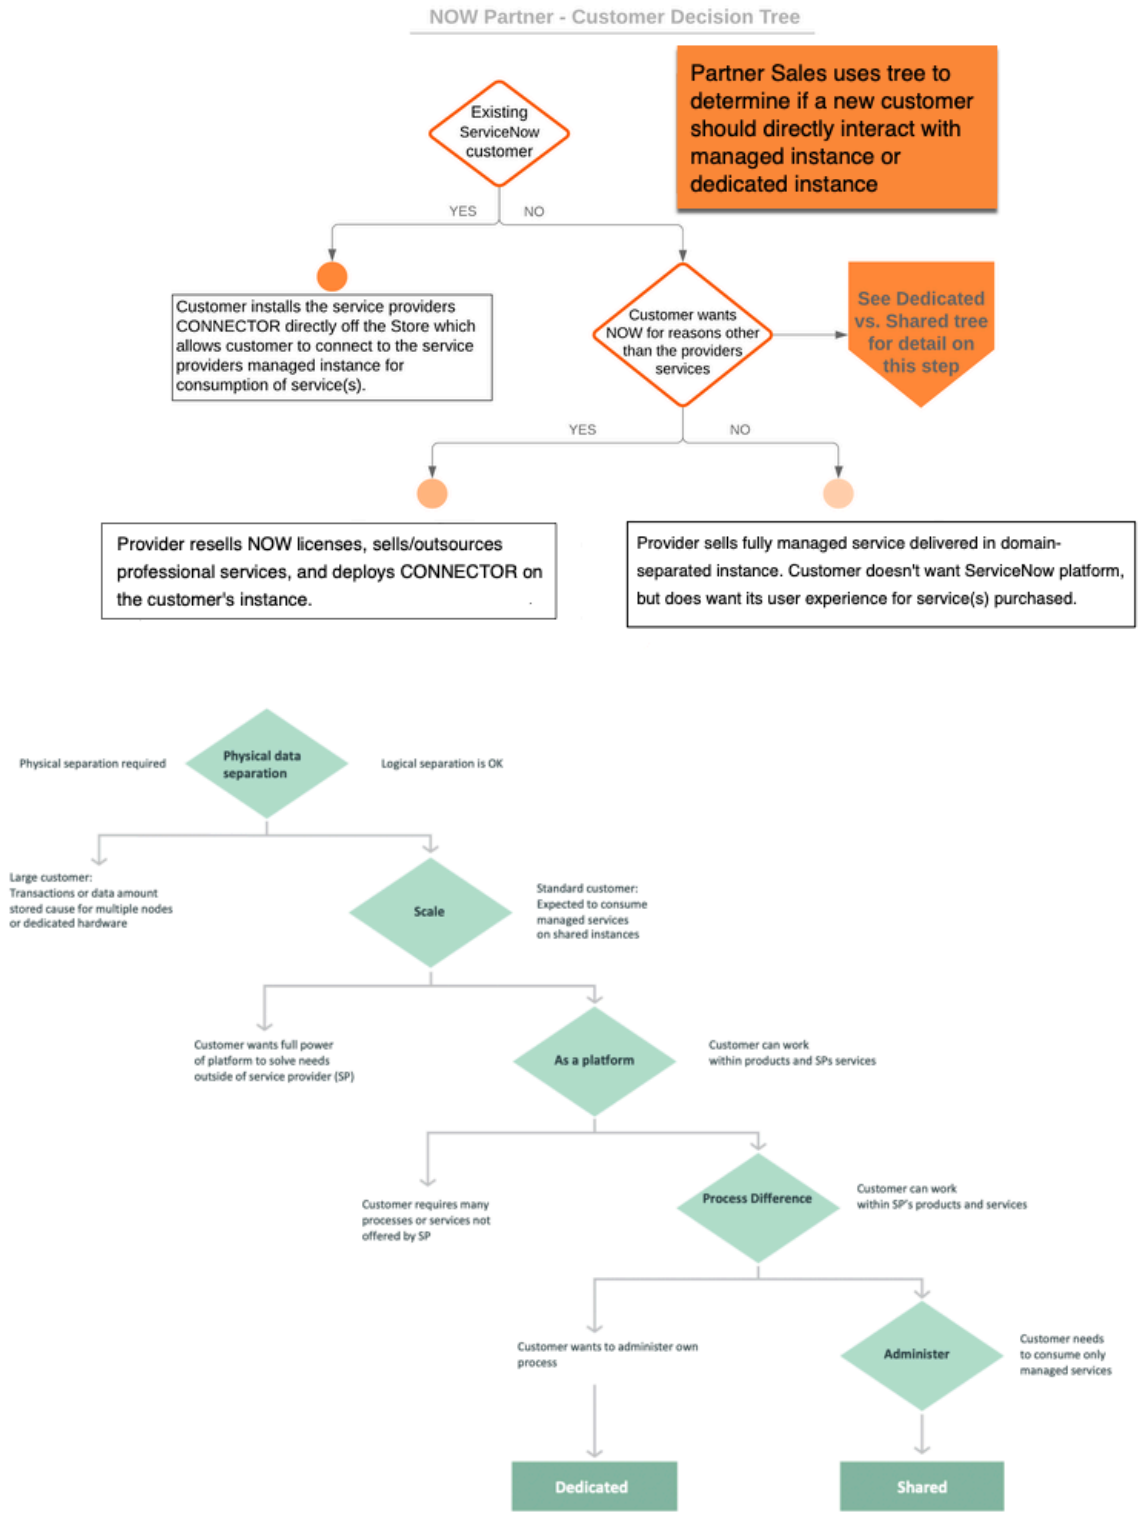

To learn more about hierarchy architecture, see Service provider reference architecture.

Context and domain separation

The context of a user's session determines the processes, data, and user interface (UI) as the user browses through list views, home pages, reports, and knowledge articles. The context is determined by the processes that you create, the business rules that you set, your workflows, and other factors.

User session context

Many factors determine the context of a user session, such as user profiles, groups, company criteria, and so on. In the following diagram, you see that the incidents that a company has created are part of the context.

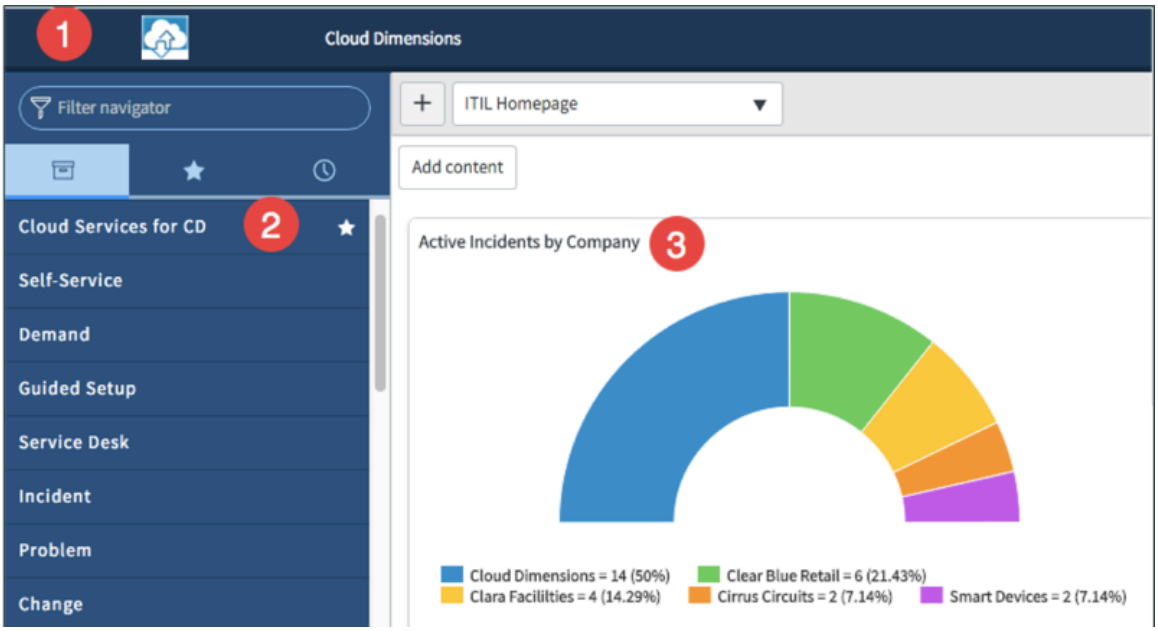

The user in this example has a home domain of Cloud Dimensions.

1. The branding reflects the settings in the Cloud Dimensions domain and company record.
2. The application navigator shows the items that are inherited from higher-level domains as well as the modules that are defined in the Cloud Dimensions domain.
3. The home pages and list data reflect the data that is visible to the user. This data is based on the user's session context. In this case, the user in the Cloud Dimensions domain can see the data in Cloud Dimensions, child domains, and the global domain.

User session context starts in the home domain

In the following diagram, you can see the elements of the context.

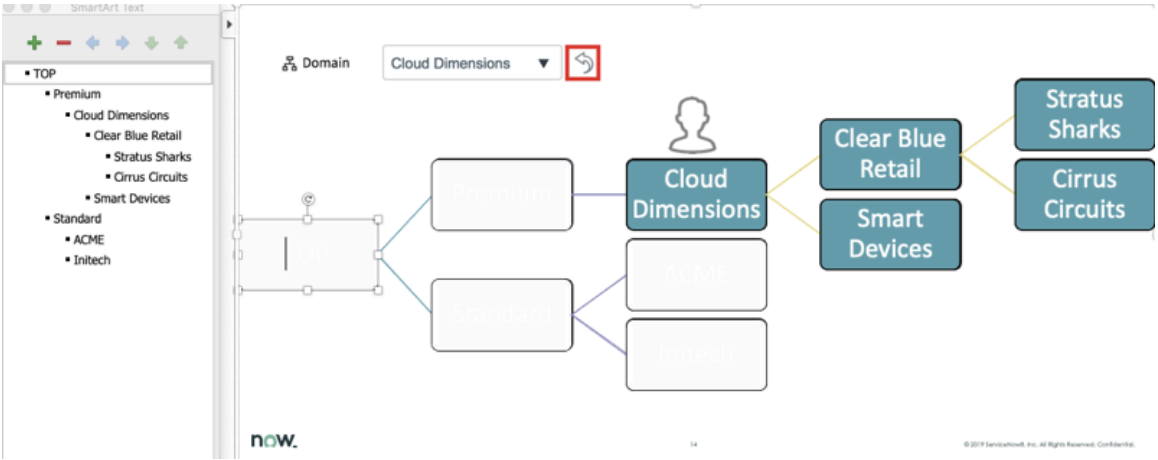

The system administrator sets users' home domains on their user records. Typically, a user's home domain is set to the same domain as their company's domain. When the user logs in, the domain picker sets automatically to the user's home domain. Users can return to their home domain at any time by clicking the arrow icon on the domain picker.

The domain picker's list includes domains within the user's session context. Users may further limit their session context by selecting child domains with the picker.

The context of the user session includes the user's home domain and any child domains. This set of domains in the user's session context is appended automatically to every query that is sent to the database. That way, the results are limited to just the data in these domains and global data. This process is embedded in the compiled code that is not accessible.

Service accounts that are used for integrations also have user session context. There is user context and records context, each with its own data in its own domain. These contexts affect the integrations. Database queries (records) are limited in the same way as interactive users (users), meaning that they work as normal but are limited by whatever constraints the developer has configured.

You can learn about additional ways to add domains to a user's session context in Service provider reference architecture.

Record context

As a user drills into individual records, record context is activated. The record context determines the UI elements and processes to apply to the record.

A record's domain dictates the process, data, and the availability of UI elements within the record.

Note:

- Record context persists even if the user's domain changes.
- Users can view records concurrently in multiple browser tabs, while maintaining their own record context.

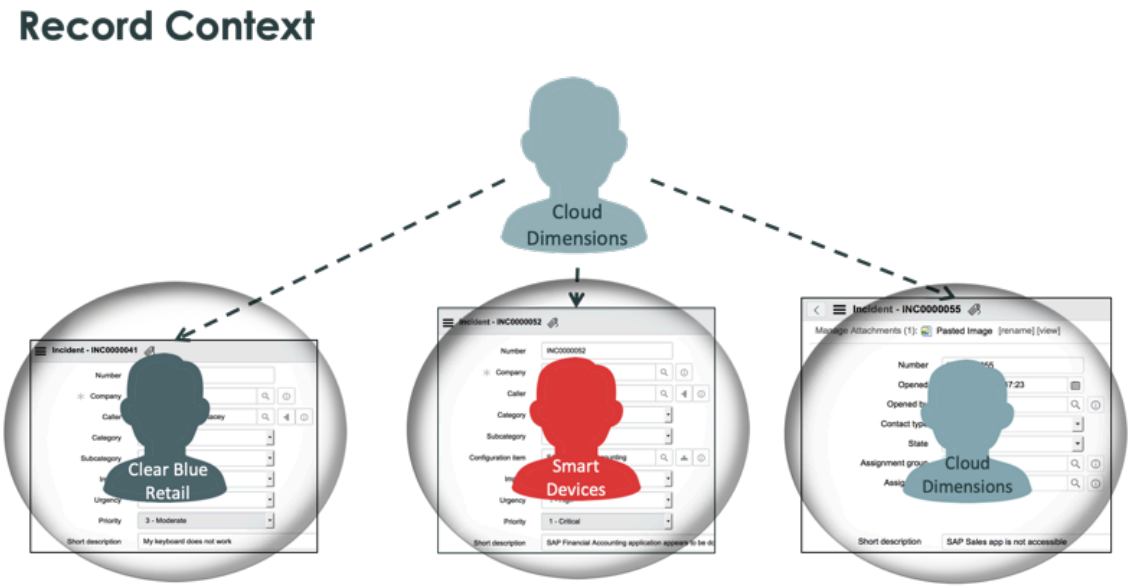

Segregating and securing data with domain separation

You can segregate and secure data on the ServiceNow platform in multiple ways, depending on your customer's needs.

Segregating data in multiple ways

The following diagram shows four ways that you can segregate data. You can use separate instances, domain separation, contextual security and business rules, and the reference architecture itself as ways to segregate data.

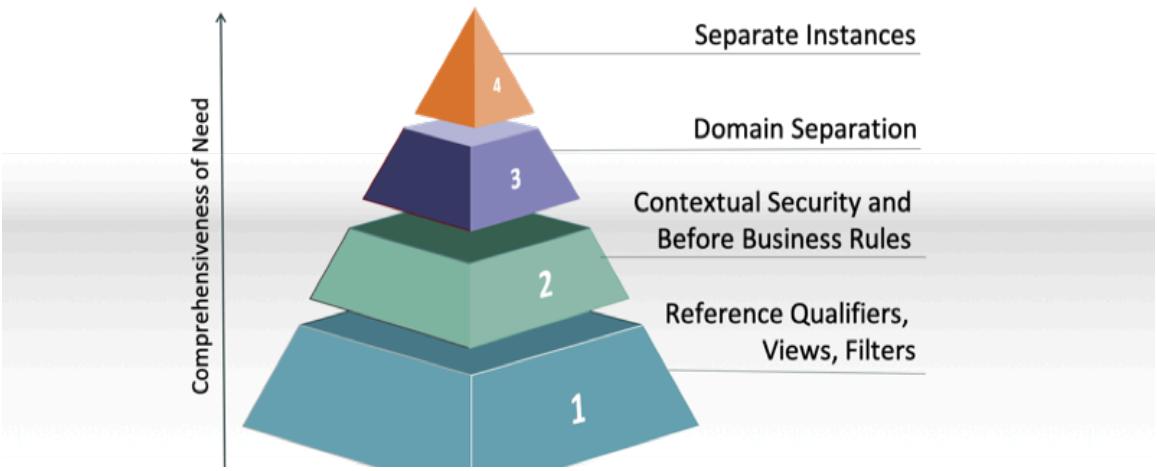

You can segregate data in these four ways:

1. Customizing the reference architecture with qualifiers and filters so that departments and groups within a company can focus on their own work. By segregating the data between these departments or groups, a department or group can't see another department or group's records.
2. Adding contextual security and Before Query business rules as additional layers of security to guard against data breaches. See Context and domain separation and Before Query business rules to learn more about domain separation and business rules.

3. Adding another level of security in a company by using domain separation. The data from every database query is limited to the data that is visible in a domain before contextual security and business rules are executed.
4. Using separate instances to segregate the data at the database and application layer.

Separate instances, domain separation, contextual security and business rules, and the reference architecture are ways to segregate data. These four ways relate to each other as indicated by the Comprehensiveness of need arrow in the diagram. How each layer interacts with the other layers depends on how you set up your domain separation configuration.

Not all organizations require domain separation. You might find other alternatives, such as separate instances or a single instance without a domain. To learn more about these alternatives, see Evaluating the need for domain separation.

Cross tenant intelligence

A multi-tenant architecture is where you have a single instance serving multiple tenants. Data, metadata, business logic, and processing context for tenants is automatically handled with access to additional tenant data.

Single versus multiple tenants

Single-tenant instance

You're a ServiceNow customer, you bought the licenses, and it's up to you to decide what services you want. You can upgrade when you want, see previews of all the great new features, and configure your instance right away. Being a single tenant has these benefits and restrictions:

- You have higher upfront costs and administrative overhead, but you have more freedom to remodel and expand.
- You have higher costs to obtain and maintain the instance and supply administration staff. Although you are free to build out the environment as required, you must comply with the ServiceNow recommended practices and standards.

Multiple-tenant instance

Someone else owns the instance, possibly a service provider with multiple customers. They upgrade when they want and put new services on the instance when they want. If you are a customer of a service provider, you are most likely on their instance because you wanted what they offer. Multiple tenants have these benefits and restrictions:

- A centralized staff administers configurations, integrations, and upgrades.
- The instance owner provides added services.
- Domain tenants have lower upfront costs to use the ServiceNow platform, have lower monthly costs because they are sharing it with many tenants, and do not have to employ a staff to administer the environment.
- Benefits shared from requests or changes that are initiated by other tenants.

Alternatives to domain separation

You can use a separate instance as an alternative to domain separation for your customers. A separate instance allows you the flexibility to meet the requirements for data separation within the groups and departments in an organization with little to no impact on others.

Separate instances

Pros and cons of separate instances

Separate instance	Single instance - without domain
Pros	**Pros**
Built to suit each customer/organization	May address simple scenarios
Minimize impact of customization on others	Cost
Release schedule coordination	**Cons**
Clean separation	Extensive modifications to baseline code
Choose DATACENTER region	Modified base system code skipped during upgrades
Cons	Must address all secondary and supporting tables as well
Cost	Extensive testing required
Alignment of instances	No ServiceNow product team to evolve your custom code
Testing effort for upgrades	
Duplication of effort	
Integrations required	

You can time upgrades and releases separately for each instance. However, if you choose to use separate instances, you need to do a lot of coordination with other people who are administering instances. By configuring an instance with contextual security, form views, reference qualifiers, filters, and robust conditions, you don't have to use domain separation in your company.

With a separate instance, you may address data and process separation but your instance owners must maintain and keep up with the extensive customizations that is required for separate instances.

Related topics

Context and domain separation

Service provider reference architecture

Evaluating the need for domain separation

You may find that domain separation doesn't always work for your customers' organizations. It's best that you base your decision to go with domain separation by looking at your customers' needs.

Evaluating the need for domain separation

Reasons for domain separation

These factors may help you decide to choose domain separation for your customers' organizations:

- Your customers have moderate alignment of processes and general platform requirements.
- Your customers plan to work on tasks as fulfillers rather than as requesters.

- Your customers have a contractual agreement that requires that data records be isolated, but your instance owner has determined that the requirement may be addressed somewhere else in the configuration.
- Your company's instance owners have entire entities that operate as physically separate organizations and do not share data, but full reporting is still required. Separate domains would allow data visibility when configured correctly.

Reasons for no domain separation

These factors can point to reasons why your customers' organizations might not want to set up domain separation:

- Your customers want to administer their environment, have full ownership of it, and set the roadmap for expansion.
- Your customers require that the data and process at the physical or database level be completely isolated.

Note:

Domain-separated instances contain a shared database so this undermines the isolation requirement.

- Departments in your customers' organization want to isolate records. (Access controls may suffice.)
- Your customers all want their own processes, business rules, and workflows.
- The corporate culture is one of non-collaboration between your customers' organizations.
- Your customers interact with the platform as end users only.

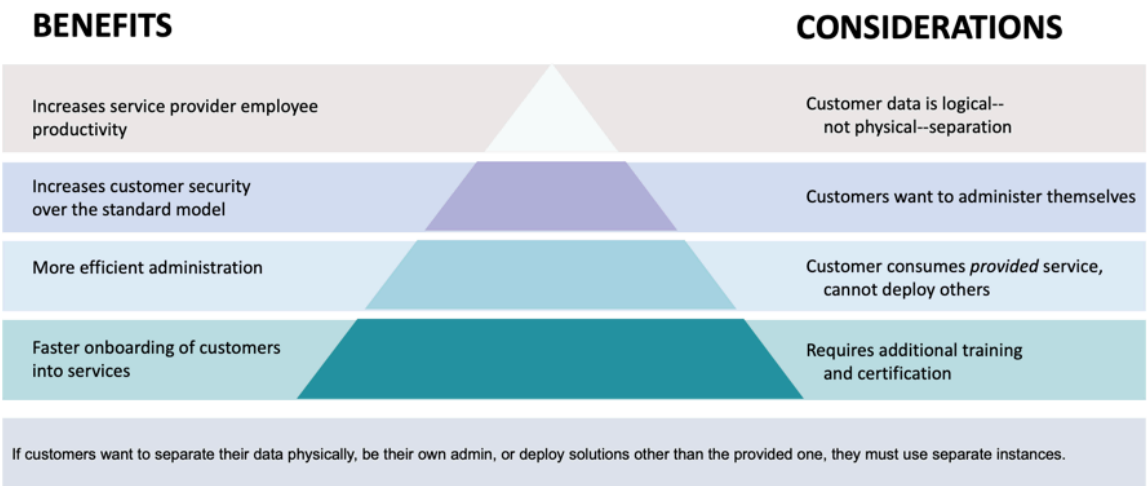

Benefits of domain separation

Domain separation may work better for your customers' organizations than any other method for separating the data between groups and departments.

Domain separation benefits at a glance

You can enable domain separation with a ServiceNow plugin that has functionality built into the core platform. The separate domains configuration is managed by a product manager who is supported by a development team. Enhancements and fixes for domain separation functionality are included with ServiceNow releases and ready to be used by customers. For assistance with

domain separation, your instance owners can use Customer Service and Support resources, such as the Service Portal.

How a database query works with domain separation

Using database queries with domain separation in your customers' applications help them protect their data. These queries then speed up the configuration and build processes.

How domain separation protects data

In the following figure, the Incident table [incident] has a domain field that is inherited from the incident's task. When you see this domain field, you know that the records in the table can have domain assignments.

When users log in, their home domain appears with the set of domains they may access. This is known as the user's session context. For more information about session contexts, see Context and domain separation.

Database query with domain separation

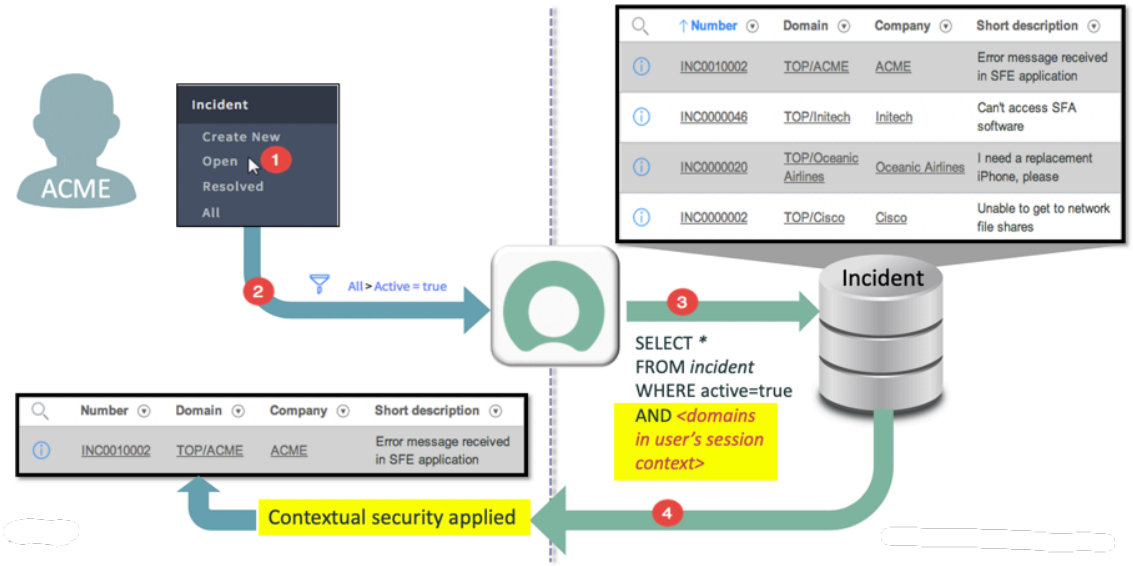

1. In a browser, the user from one of the companies, Acme, selects the Open Incidents module to view all incidents where active=true.
2. The active=true filter is submitted to the application.
3. The application then sends a query to the database by appending a WHERE clause to active=true. The WHERE clause limits the incident records that are returned to those records that are in the user's domain or the domains that the user may access. Only the records in these domains are returned to the application for processing.
4. Contextual security is applied, further limiting the data that is returned to the user. The incident records appear in the Open Incidents list.

Note:

When you apply contextual security, you create limits to the data that are returned to the user. These limits protect other content that you may not want users to see.

To learn more about contextual security, see Context and domain separation.

Note: This processing logic applies for all queries to the database, including those queries that are triggered using integrations.

Domain separation levels of support

Choose from three categories for domain separation of an application for your customers' organizations.

Applications that support domain separation may support the separation of data and data routing only, have advanced business logic separation, or support tenant (customer) level administration of the application. These definitions delineate the support levels from the perspective of actual use cases and the people who implement them.

Incremental ServiceNow support levels

Basic	Standard	Enhanced
• Data is domain-separated • Logic exists to ensure proper data routing, caching, rollups, and aggregations • Global configuration operational for multiple tenants	• Application properties are domain-aware as needed • Business logic can be domain-separated by the instance owner per tenant	• Data-driven process enables failsafe configuration by tenants through the UI to drive business logic

Level	Type	Summary
No support		• The domain field may exist on data tables, but no business logic exists to manage the data. • This level isn't considered domain-separated.
Basic	Customer data management	• Business logic: Ensures that data goes into the proper domain for the application's service provider use cases. • In the application, user interface, cache keys, reporting, rollups, aggregations, and so on, all consider the properties of the domain at run time. • Your instance owners must be able to set up the application to function normally across multiple tenants. Use case: When a service provider uses chat to respond to a customer's message, the client must be able to see the response.

Level	Type	Summary
Standard	Customer process management	• Includes the Basic level • Business logic: Processes can be created or modified per customer by the service provider. The use cases reflect how the application is used by multiple service provider customers in a single instance. • Your instance owners must be able to configure minimum viable product (MVP) business logic and data parameters per customer for the specific application. Use case: Admin must be able to make comments required when a record closes for one customer, but not for another customer.
Enhanced	Customer self-managed configuration	• Includes Basic and Standard levels • Enables service provider customers to modify business logic that is based on defined use cases. These configurations are UI-based and fail-safe so that configurations by one customer can't affect another customer. • The instance customers must be able to configure MVP business logic and data parameters themselves. Use case: Customer of a shared environment must be able to make changes according to impact, urgency, or priority within a domain.
Effective domain*		In some cases, a platform feature or application may support service provider use cases even if the domain framework isn't being used. The use cases must be detailed to support domain separation. An asterisk (*) after the support level indicates this kind of configuration. Use case: Before the New York release, Service Catalog had no domain support but the instance owners could configure separate catalogs and items for each tenant in a domain-separated instance by using user criteria. The result was that each tenant could use Service Catalog at a Standard level.

To view all applications listed by their support level see Application support for domain separation.

Summary

Domain separation is a framework that you must use to make your applications aware of its customers.

Consider the domain framework capabilities, your applications' business use cases, what the personas are, and how they use the application before you can use the framework to make your application supportable.

Service provider reference architecture

Your customers can access service provider (SP) services by using a portal that is designed for them to reach their domain-separated instance.

Basic attributes of service provider reference architecture

- You do not assign fulfillers to a domain. Instead, you share them across domains. This makes it harder to audit how many fulfillers you have per domain.
- You can share and leverage domain administration. This means that there is no overhead and you can optimize licenses.
- The number of users on the instance can change when you get a new customer. A new customer can result in tens or even hundreds of thousands of new users on the system. The number of total users is virtually unlimited in one shared environment.

The portal for SP services is dedicated or shared to the SP shared instance. Service providers use ServiceNow shared instances to manage their service delivery.

Reference hierarchy for domain-separated instances

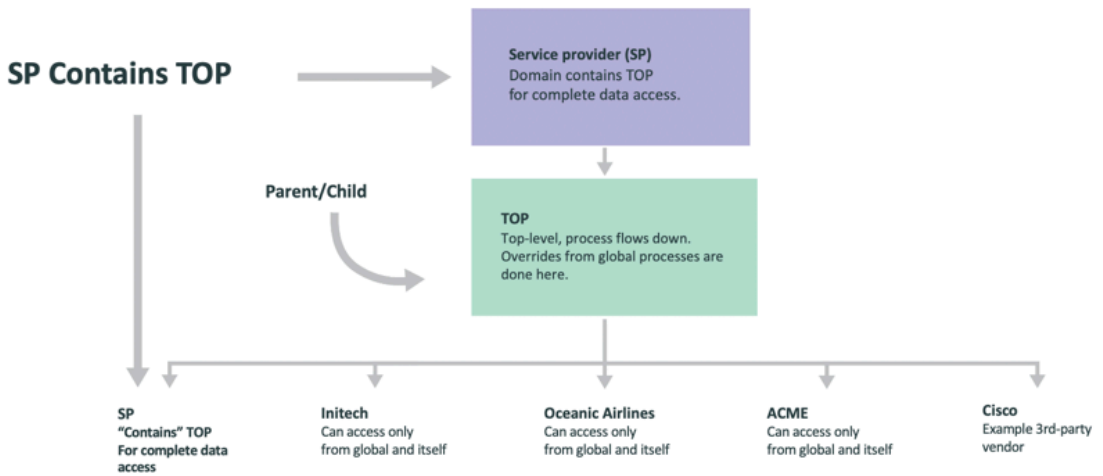

Service provider reference architecture decision trees

You can use decision trees and a comparison chart to determine if a new customer should be added to a shared instance or to their own dedicated instance.

Decision trees

Use these decision trees to help your customers decide whether to use a managed or dedicated instance.

Customer decision tree

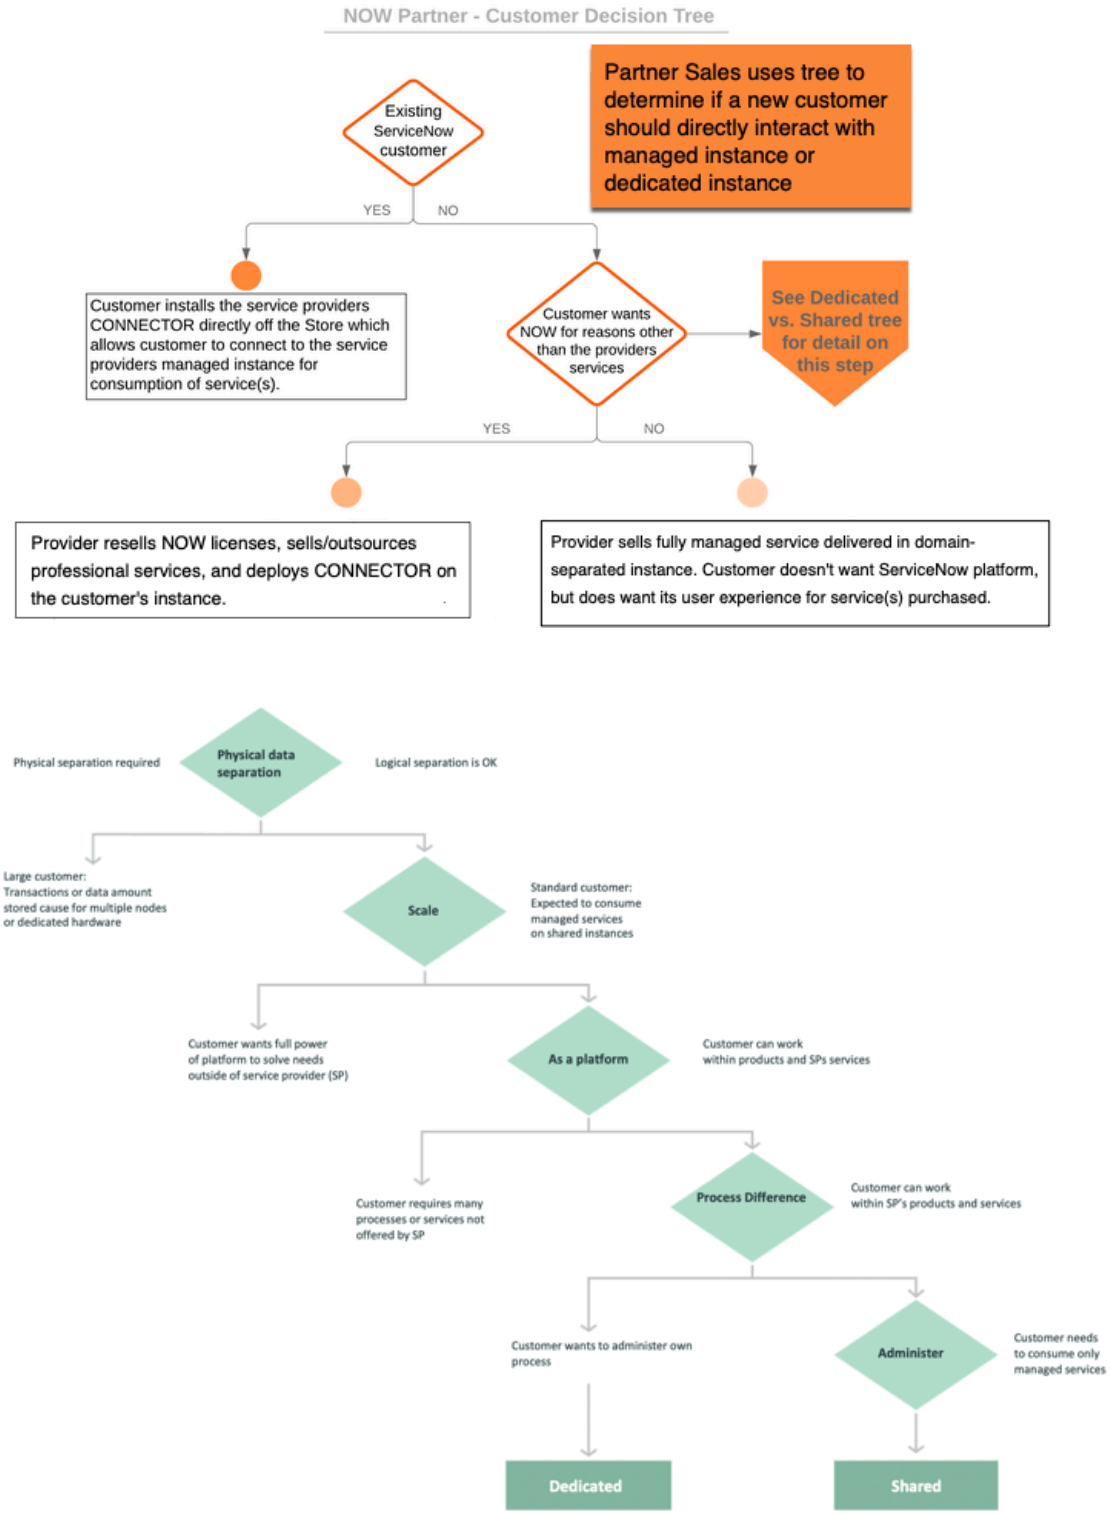

SP reference architecture comparison

Service provider reference architecture for dedicated instances

Service provider (SP) customers can access SP services by using a portal to a dedicated instance. SPs use these dedicated instances to manage their service delivery.

Dedicated instances

Attributes

- Dedicated instances require that you have separate administration and dedicated teams. You need multiple licenses for the administrators and developers who log in to multiple instances.
- Each instance has a finite number of requesters and fulfillers. When you get a new customer, you must procure an instance that is based on the size and scale of your customer's company.

Shared instance

Dedicated vs. shared instance

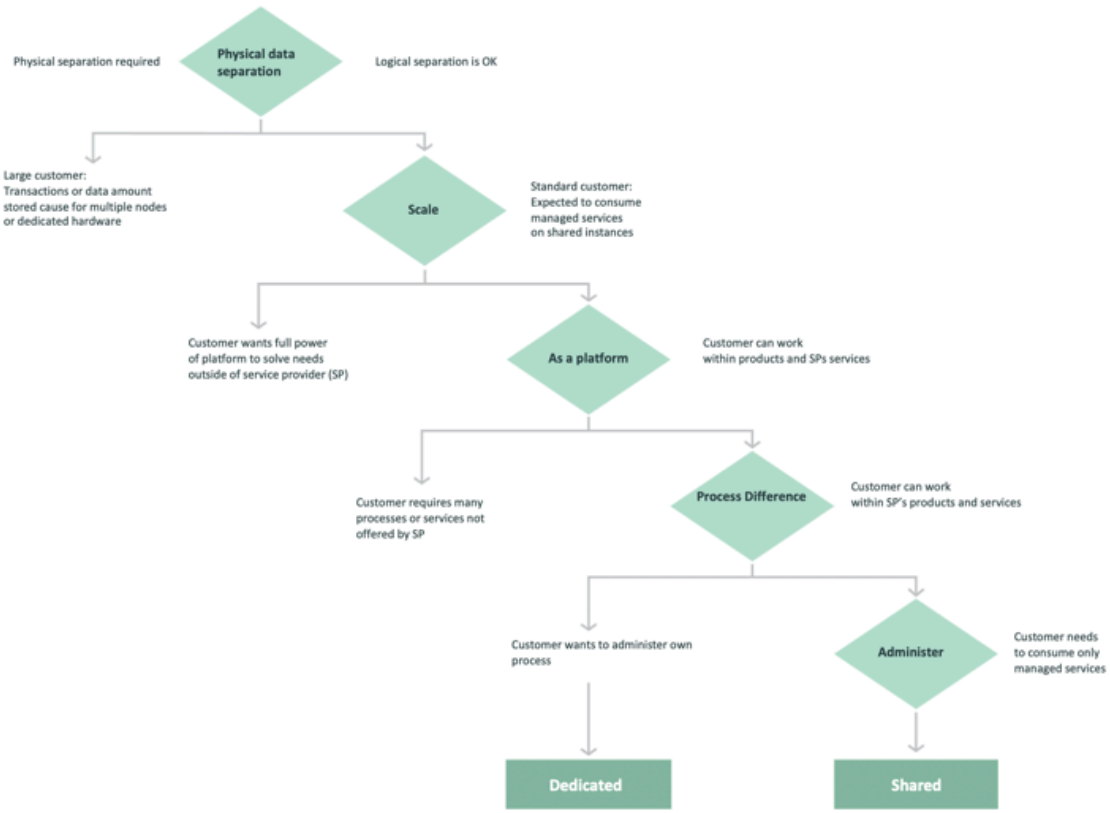

SP reference architecture comparison

Service provider reference architecture for hybrid

Use the hybrid service provider (SP) reference architecture for a customized solution. Your customers require a dedicated instance for a specific service. They can still use the shared SP instance for other services, but it requires integration of each instance.

Hybrid architecture

Your customer may be responsible for delivering this additional service directly. You need to build a hybrid solution of several attributes for your customer to provide.

SP Reference Architecture - Hybrid

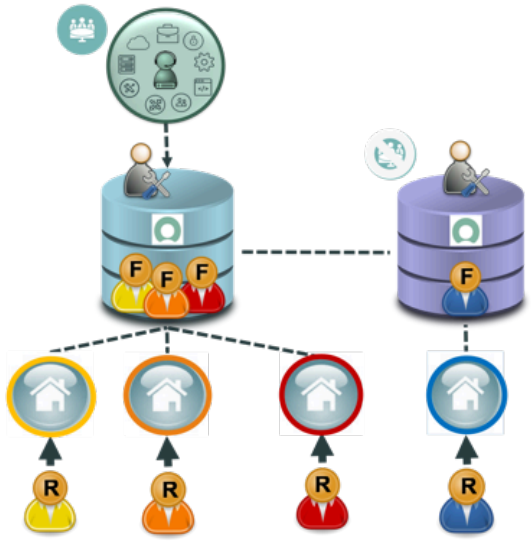

- Some customers require a dedicated instance
  - Service may be for the sole use of that customer
  - Offering Industry or Vertical Solution
  - Government or Regulatory Requirements
  - Application does not support Domain Separation
- Customers can consume some shared services from the SP instance and some via integration with the customers dedicated instance
- SP or customer owns and manages the dedicated instance. The SP could provide operational management of the dedicated instance as a service

Legend:

- Centralized governance at the SP for instances managed by the SP
- Dedicated instance managed independently or governance maintained via a SP "Gold Blueprint"
- Administration (shared or dedicated) can still be owned by the SP
- Fulfillers and requestors are on the shared instance and on the dedicated instance
- Customer requesters on each instance

Attributes

- You can share and leverage the administration of the instance. This means that there is no overhead, and you can optimize licenses.
- If there is a new instance for a decentralized environment, the program team is responsible and funded accordingly as dedicated administrator users for that instance. In a centralized environment where all instances stem from a blueprint, you need duplicate administration licenses.
- You do not assign fulfillers to a domain. Instead, you can share them across domains.
- If a customer is sharing both a shared environment and a dedicated environment, they require a fulfiller in both environments. This means more work for each team because the processes for shared and dedicated instances require different work for each instance.
- The number of users on the instance can change when you get a new customer. A new customer can result in tens or even hundreds of thousands of new users on the system. The number of total users is virtually unlimited in one shared environment.
- Each instance has a finite number of requesters and fulfillers. When you get a new customer, you must procure an instance that is based on the size and scale of your customer's company.

SP reference architecture comparison

Standalone (SA)		Multi-Tenant (MT)		Hybrid (MT+SA)		SIAM (Multi-Vendor)	
Platform as a Service	Enterprise scale	Processes simplified	Cost reduction	Managed Services (MT)	Custom Services (SA)	Service Integration	Multi-supplier governance
Customer wants full power of platform to Solve own needs outside those offered by SP.	Large customer: Transactions or data amount stored cause for multiple nodes or dedicated hardware.	Global governance & admin for business use case & process can be developed & maintained on the instance. Best when used w/ minor-to-moderate process differences among customers or sub-entities.	Lower administration, transaction, & onboarding costs, & ongoing operational costs for resources & licensing. Faster onboarding = more revenue.	Standardized re-usable services. procedures. processes, & resources delivered on a single ServiceNow instance to multiple Customers.	Services for sole use of customer, offerings specific to industry or vertical solutions, government, or regulatory requirements, applications that do not support Domain Separation.	Multiple best in class Service providers for individual services bringing together tooling, reporting, SLAs, strategy, & design for unified customer experience.	Unified governance across all providers to ensure organization requirements are met & that all providers cooperate on behalf of the customer.

Service provider reference architecture for Service Integration Management (SIAM)

The Service Integration Management Service Integration and Management (SIAM) for service provider (SP) architecture integrates services for a unified customer experience.

Attributes of SIAM architecture

- Customers access SP services on a dedicated or shared portal to the SP shared instance.
- SPs use ServiceNow shared instances to manage their service delivery.

SIAM architecture at a glance

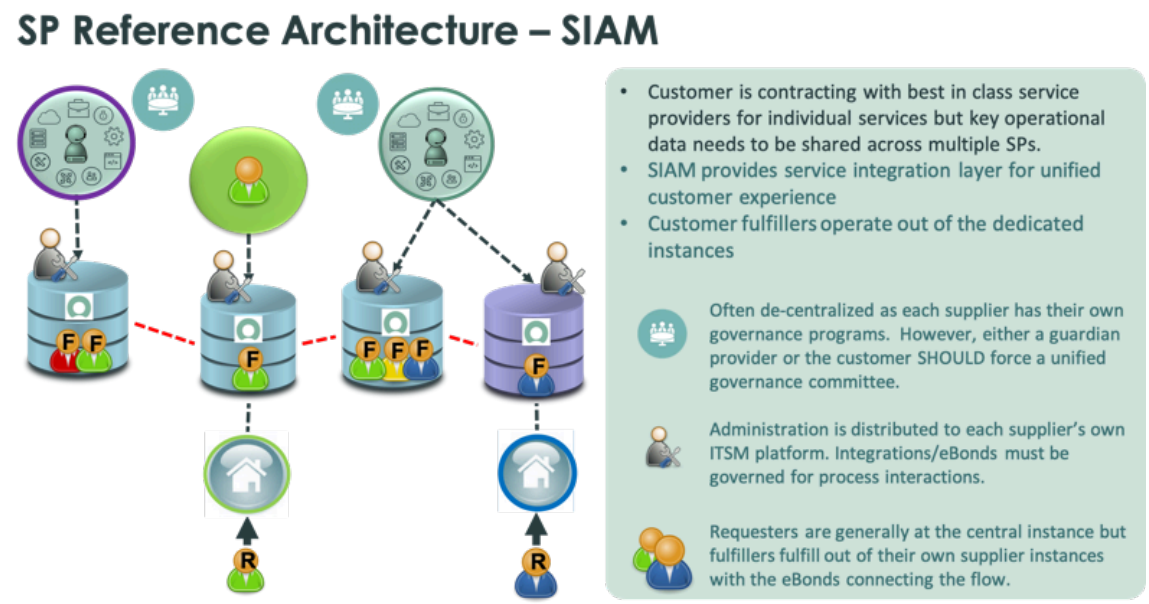

SP reference architecture comparison

Standalone (SA)		Multi-Tenant (MT)		Hybrid (MT+SA)		SIAM (Multi-Vendor)	
Platform as a Service	Enterprise scale	Processes simplified	Cost reduction	Managed Services (MT)	Custom Services (SA)	Service Integration	Multi-supplier governance
Customer wants full power of platform to Solve own needs outside those offered by SP.	Large customer: Transactions or data amount stored cause for multiple nodes or dedicated hardware.	Global governance & admin for business use case & process can be developed & maintained on the instance. Best when used w/ minor-to-moderate process differences among customers or sub-entities.	Lower administration, transaction, & onboarding costs, & ongoing operational costs for resources & licensing. Faster onboarding = more revenue.	Standardized re-usable services. procedures. processes, & resources delivered on a single ServiceNow instance to multiple Customers.	Services for sole use of customer, offerings specific to industry or vertical solutions, government, or regulatory requirements, applications that do not support Domain Separation.	Multiple best in class Service providers for individual services bringing together tooling, reporting, SLAs, strategy, & design for unified customer experience.	Unified governance across all providers to ensure organization requirements are met & that all providers cooperate on behalf of the customer.

Domain separation terms

With a ServiceNow instance, you can improve efficiency, add greater security, and increase performance for your customer organizations. It's helpful to understand some of the most common terms as you create your configurations.

Managed domain

In a managed domain, the **Managed domain** field allows domain administrators to manually select a domain for the user, group, department, location, or CI record, rather than using the domain that is assigned automatically from the company record.

If you want to change those properties, you can override them to further customize the functions of the applications in each of your domains.

Process tables

In process tables, if you see a value in the **Overrides [sys_overrides]** field, a process override record exists. That means that delegated administration, which is how administrators can set domain-specific policies, is in effect. Admins in the global domain can use the **Expand/Collapse Domain Scope** related link to see override records.

Note: Reports are separated into domains and contain an **Overrides** field. To view all reports from the global domain, use the **Expand Domain Scope** related link.

When you view process tables from a domain, you see only the relevant process records for the selected domain. When you view a process table from the global domain, the **Expand Domain Scope** related link is displayed to let you see all process records, including overrides. To view only the relevant process records for global again, use the **Collapse Domain Scope** related link.

The domain scope feature is used only for process tables and causes the visibility of data on the table to shift in the opposite direction. For example, a record in the parent domain can be seen in the child, but a parent cannot see a child record. This allows the process to flow down to child domains.

Types of domains

Different types of domains can help you organize your processes and data and how they function in the application or feature.

Customer Domain

In the customer's domain is the user interface, as well as the process that controls how the data Is used.

The ACME domain in the following image is a customer domain.

Process Domain

You create processes for how the data is used and what it does in the domain. These processes must have these attributes:

- Specific processes and UI settings for a set of domains
- No core data of any kind (such as specific user data).
- The TOP domain in the following image is a process domain.

Data Domain

The data domain holds data that is relevant to multiple customers. That data can be shared without sharing the actual customer domains. Each customer has its own data domain and can access it.

Note:

This kind of domain is not common and can cause performance issues if overused. Consult with an SP architect before use.

Example: The domain may hold tasks that ACME, Cisco, and the SP all need to interact with.

The Default domain in the following image is a data domain.

User Data

User record data never belongs in the global domain or any of the process domains. Users are primarily created in customer domains and can on occasion be created in data domains.

Admin accounts are special as they should not be used as everyday users of the instance and should be in the global domain to facilitate administrative functions.

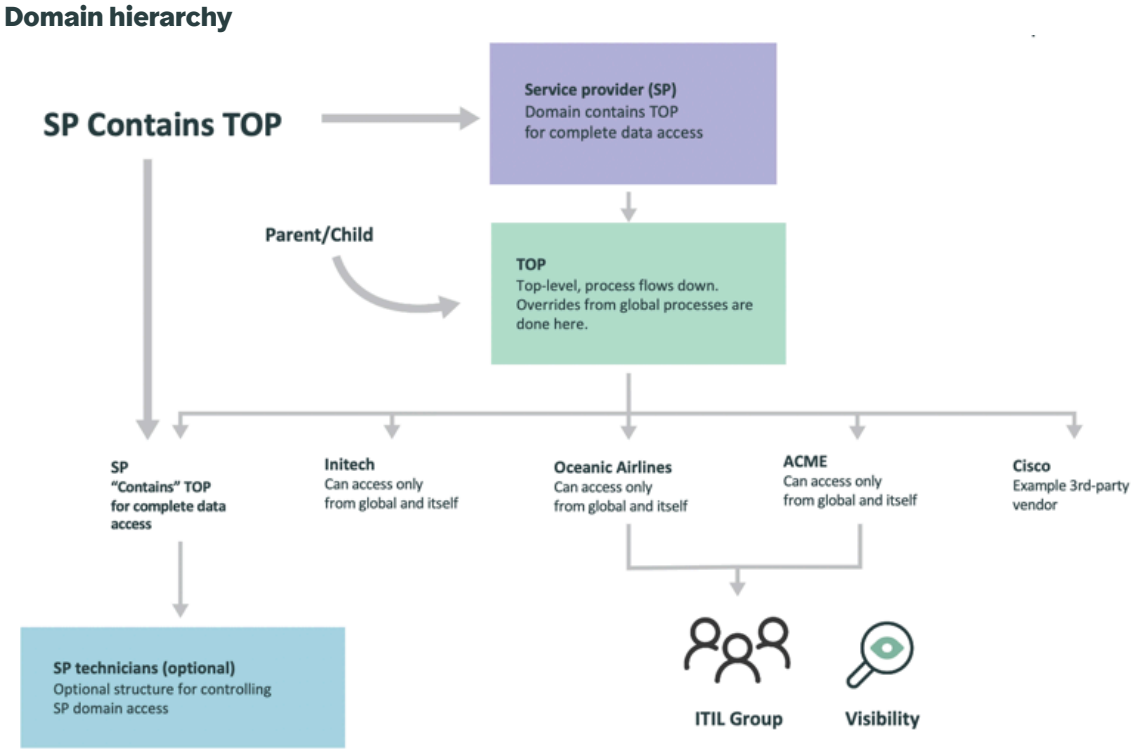

Lists, admin, global process

Lists

From the global domain, if you right-click any choice field's label, select **Configure Choices**, and then add a new choice, the choice pushes automatically to all domain-specific lists for that field. If the new option is marked as **Selected**, it is added as active. If the new option is marked as **Available**, it is added as inactive.

Instance Administration

The instance owner's administrators must handle all normal process creation, modification, and maintenance in a domain-separated instance. Individual domain managers can maintain some parts of data-driven processes. The types of domain managers maintain user administration, support group memberships, and locations, or manage applications that are designed with tenant administration in mind.

Global process/parameters

You can create and maintain the process that affect the global domain as well as set the parameters. These properties are common for all users of a domain-separated instance.

Examples: System properties, dictionary overrides, `sys_documentation` (field labels), the data model (classes, CI types, and so on), tables and fields `[sys_dictionary]` (access can be restricted), indexing (text indexes as well as database), ACLs, installation exits, inbound actions, public pages, and interceptors.

Domain-separate a custom table

You may need to create custom tables in separate domains. This topic covers both the procedure and the concept behind domain-separating a custom table.

1. Create a sys_domain field

Note: If a system table or a table has not been domain-separated by the Domain Separation plugin, it's best not to domain-separate it.

Use these points as a guideline to create a sys_domain field.

- Create a new field as a *domain_id* type.
  - Column Name: sys_domain
  - Other attributes: Defined automatically
- The Sys_domain_path is created automatically.

The column name **sys_domain** is reserved in the ServiceNow AI Platform, which means that the system recognizes it and automatically applies the appropriate field type and attributes for you. This automatic configuration also creates a corresponding **sys_domain_path** field.

- Set the column name to `sys_domain` rather than using the label.
- Domain separation is not appropriate for every table. In general, if a table is part of the base instance and that table does not have a **sys_domain** field, you should leave it that way.

A **sys_domain** field is created automatically when you create a domain_id type field with the name “sys_domain."

2. Add a business rule to set the domain

Without business rules

The domain is set to the current domain of the user who creates the record.

With business rules

The domain is assigned using scripted logic, typically based on the Company field.

In addition to a `sys_domain field`, custom tables need a business rule similar to **Domain - Set Domain – Task**to set the value of the domain field. In addition, you will need **Domain – Default – Task**, which moves records without a domain to the default domain if the first rule fails to assign a domain.

On the task table, review the business rules for Domain. Pay particular attention to the Order field. The priority of execution is given by the Order field from low to high.

The first rule that runs, **Domain – Set Domain – Task**, attempts to set the domain of the record based on the record’s Company’s Domain.

If the first rule fails to find an appropriate domain, the second rule, **Domain – Default – Task**, executes. This rule sets the domain of the record to the default domain.

Finally, if the domain of a task record changes, the **Domain – Cascade Domain – Task** business rule changes the domain on all records related to the task, such as workflows, metrics, SLAs, and attachments.

3. Add a business rule if Step 2 failed

If the initial business rule fails to set a domain and the domain is still empty or global, a second business rule runs. This rule examines the `task_for` field that is based on the caller or `requested_for` field. This rule is checking to see if you can set the domain of the record based on the user’s domain. If not, the business rule sets the domain to the default domain.

Following is a sample script for the business rule:

```
/* essentially
If (task_for is set)
  set the domain to the user's domain
ELSE
  set the domain to the default domain
*/
```

4. Domain – cascade domain – task

Tasks can have many related tables that work together for business objectives. These related records include workflow, SLA, approvals, attachments, and email. If the domain of a task changes, the related records domain must change, too, so they remain visible to users in the new domain.

This Cascade rule is commonly triggered when you clear records out of the default domain.

The related records for a Cascade domain contained in the Script are shown similar to the example:

```
/*
* Keep domains in sync w/related records for:
* workflow context
* workflow history
* approver tables and related workflows
* attachments
* emails
*/
```

Customizing domain properties and themes

You can customize your customers' company properties and themes within the domains that you have configured. Customization makes their instances fit in with their companies' overall look and feel.

Customizing company themes and logos

On the company record, you can customize specific design themes and logos for each company.

Customize the domain

By default, when you have the service provider plugin installed, the standard business rules contained in the core tables cause the record's domain to be set to the domain of the company that is associated with the record. The company can control the tables such as Task, User, Group, Location, Department, and so on.

With all of these tables, except the Task table, you can override the domain that a record is created in. This gives you more customization options.

Managing domain separation for specific uses

You can set up separate domains for email notifications and customize the properties of catalog, tables, users, groups, and views. This enables you to provide more specific behavior in each domain, giving your customers more flexibility.

Emails

You can use separate domains for email notifications and overrides. When you use separate domains for notifications, you can do an override that is based on the domain of just the attached record, not the user's whole domain.

Service Catalog

The Service Catalog is now domain-separated so that your customers can see and access the catalog. Items are processed as OR conditions when multiple items are used. Service providers should administer the categories and Items themselves so they fit their own criteria specifically.

Users and groups

Only use admin accounts in the global domain because admins need access to all domains. Do all your application testing from an actual domain, not in the global domain. Overrides don't process properly in the global domain. Admins should also be given user accounts in production if they are to use the application.

Working with fields

There are several points to consider when you work with fields. Pay close attention to these fields because they can have many variations that affect your configurations.

Lists

There are personal, global, and domain lists, as well as multiple views of each.

Forms

There are global and domain lists as well as multiple views of each.

One database

Any fields that you create exist for all users, in one database. Consider the global impact before you create one.

Note: ACL scripts cannot keep a field from being viewed in a list because they do not run. You can add a READ ACL to hide a field from users if the ACL is only role-based.

Creating tables

When you create a table, you should add a `sys_domain` or an `sys_overrides` field. Any table that contains data that your instance users need to access, needs the `sys_domain` field. Tables that extend or support processes and that need to flow down to children domains also need the `sys_domain` field.

Configuring domain separation with the domain picker

Use the domain picker wisely, and remember the 80/15/5 approach so that you do not customize too much and impact the performance of your instance.

Verify your domain before making changes

The domain picker gathers all the domains into a list for you to choose from.

If your session times out and even if you don't get logged out, your session falls back to the domain on your user record. You also lose any elevated roles at the same time. In this case, your domain picker could still show the last domain that you selected if the top frame of the list hasn't been reloaded. For this reason, you should reload your list completely if you have been away from the instance for any time.

Configuring at the TOP domain or the global domain

Domain separation works best when you are providing services to customers that are mostly standard in their configuration and user and group definition. The more that you customize and create "one-off" solutions, the more you create a margin for error. When you create your processes and business logic, any variations should be in properties that work automatically for each customer. While processes can still be adjusted as needed, use great care when you decide when, and to what degree, to create a unique configuration for a single customer.

You need to use an "80-15-5" approach in configuring your domains to avoid too large a margin for customization, and therefore error.

- Recommended approach for configuration:
  - **80%** or more **Standard**
  - **15%** or more **Parametric**
  - Less than **5% Configuration**
- Determine if a suggested change should be a global or a configurable property.
- Do not overbuild by adding more and more customization that needs to be managed. Instead, do the following:
  - Start with the base system features and verify any gaps before you make any changes.
  - Look for no-code solutions.
  - Use server-side scripts, build modular APIs, and build in domain-separated properties.
  - If you must use client scripting, use only ServiceNow APIs. Limit "synchronous" calls (those that go back and forth from client to server, also called AJAX).
  - Write all scripts logically to keep them simple and effective. Enforce peer reviews of code changes and make sure everyone is following the Domain separation recommended practices for service providers in this section.

Domain separation performance considerations

As you configure domain separation in your application and services, make sure that you consider the number and properties of domains you create. Too many property-heavy domains can impact the performance of your instance.

Limiting property-heavy domains

You can create as many domains as needed, but make sure that you don't create unnecessary domains on the instance. Weighing too many domains on the instance with a lot of properties can impact the performance of your instance.

The number of domains is not what impacts performance but rather what's in them. Too many properties can slow down the domain picker, which then slows down the overall user experience of your customers. If you are loading the domain picker and already have a large number of domains with a lot of properties, the domain picker has to load all the domains before it gives you control in the session. This process could lead to an outage where you can't access anything on the instance until the domain picker finishes. Before you create new domains, navigate to the domain hierarchy under **Domain Admin** > **Domain Map** and make sure that you actually need to create a new domain or whether an existing domain hierarchy can work.

Using Core UI domain picker

The domain reference picker is available in Core UI. With the reference picker, you do not load all the domains at once, but rather the domain is searched as you begin to enter the name of your domain into the domain picker.

Enable the domain reference picker in Core UI by following these steps:

1. In the application navigator, enter `sys_properties.list`.
2. Set the *`glide.ui.domain_reference_picker.enabled`* property to true.
3. Refresh the browser.

Note: Do not upload a large number of domains (over 30) via integration or import sets without doing some testing first or you might bring down your instance.

Setting up domain hierarchies

You can avoid slowdowns and performance impacts in your instance by knowing how domain hierarchies work and by setting them up properly.

Based on the domain hierarchy, users have access to the data in their home domain and any child domains. The process flows down to the child domains and the data rises up.

Make changes to the existing domain hierarchy only when needed. When you update the parent of a domain, the system re-establishes the parent domain with all its child domains that change the domain hierarchy. When the domain hierarchy updates, the system triggers a cascade update on all tables that relate to domains for the records that are created on that domain. As a result, a large number of supporting tables have to be updated too.

For the same reasons, even if you must change the domain hierarchy, never do a mass update. Imagine the number of queries that the system has to run to change the domain hierarchy. Always do an update in small batches. Before you start the next batch of updates, make sure that Domain Work Request (DWR) records are processed. DWRs are reports that display whether there are errors after you've changed the domain hierarchy.

Tracking DWR records

In the syslog_domain table, look for an information entry in the Message column for **DWR execution completed.** to confirm that DWR is completed.

Checking domain logs for errors and warnings

Check the domain logs to find errors or warnings in your domain path processes and hierarchy configurations.

You can find the domain logs in the Domain Log [syslog_domain] table. When the domain hierarchy updates, the system triggers a scheduled job to recalculate the domain paths. The domain logs table captures the results.

Look for any errors and warnings in this table. After reviewing this table, you need to resolve these errors and run the domain path validator again.

In this example of a log, the system has detected ten orphan records in the sys_ui_list table. The errors in these records must be fixed before the domain path can run successfully.

Error	10 records detected in 'sys_ui_list' that are not in any existing domain and Domain Paths. Fix the domain value in these records & run validator again.	com.glide.domain.validator
Error	10 records detected in 'sys_ui_list' that are not in any existing domain and Domain Paths. Fix the domain value in these records & run validator again. Entries causing the error are as follows:sys_id:1d13ae40470002007f47563dbb9a7170, sys_domain:11722b01473231007f47563dbb9a7154 sys_id:22e34b23470002007f47563dbb9a718c, sys_domain:11722b01473231007f47563dbb9a7154 sys_id:4c85fa809f233100fc6cd4b4232e706b, sys_domain:11722b01473231007f47563dbb9a7154 sys_id:94648b23470002007f47563dbb9a711b, sys_domain:11722b01473231007f47563dbb9a7154 sys_id:a0844b23470002007f47563dbb9a71ef, sys_domain:11722b01473231007f47563dbb9a7154 sys_id:a95ff2c09f233100fc6cd4b4232e7096, sys_domain:11722b01473231007f47563dbb9a7154 sys_id:bd25f4719f233100fc6cd4b4232e70da, sys_domain:11722b01473231007f47563dbb9a7154 sys_id:c1263e809f233100fc6cd4b4232e70f6, sys_domain:11722b01473231007f47563dbb9a7154 sys_id:d8a72c45475002007f47563dbb9a71bb, sys_domain:11722b01473231007f47563dbb9a7154 sys_id:e4c3ce39eb10020045e1a5115206fea0, sys_domain:60e014f69f013100fbf01f80a57fcf00	com.glide.domain.validator

To learn more about domain-separation errors, see Troubleshoot domain separation errors.

Importance of the Default domain

Organizing your domains is a crucial part of the domain separation process. If you don't set a default domain, new tasks and user records go to the global domain. Anyone can see the records in the global domain, which means that data can be seen when it is not supposed to.

When you set the default domain, its records are not visible to any user other than an admin.

Note: The defaults access can be changed by granting users visibility to the default domain or the parent domain.

You should always set one default domain for the domain records on your instance. The default domain is where the system automatically assigns task and user records that are not already assigned to a domain.

When you create a default domain from the Domain Administration screen, add the name Default in the **Name** field to differentiate it from other domains. Check the **Default** check box for the record.

Maintain regularly the records that you create in the default domain and move them to the correct domains. If records show up often in the default domain, you may need to investigate why. Ideally, you should make sure that all records are created in their appropriate domains (not global or default domains).

Contains queries and domain access

Use a "contains" query only in special cases, such as when users or groups need to see data from a domain that they don't have access to, but you don't want to move those users to a domain. Creating domain "contains" and user or group access for a domain should be an exception, only when absolutely needed.

"Contains" is a domain-to-domain relationship that is many-to-many, and has no effect on the flow of process. If you create a large number of domain "contains" relationships or provide broad access, you will generate queries with too many OR conditions. OR conditions are slow and impact the performance of your instance. Instead of using too many "contains" relationships, set up your domain hierarchy as follows:

Sample query

```
SELECT ... FROM task task0 ignore index(number) WHERE task0.`sys_class_name`
= 'incident' AND (task0.`sys_domain_path` = '/' OR task0.`sys_domain_path` LI
KE '!!$/!!(/%' OR task0.`sys_domain_path` LIKE '!!$/!!$/!!&/%') ORDER BY task
0.`number` DESC limit 0,20
```

Before you move users to a domain, make sure that they really should have access to that domain. Weigh the benefits and limitations. The query above is for just one contains relationship. If you have a domain that contains another domain, and that domain is the parent of a number

of other domains, you will have many more OR conditions. Be careful when you create a domain map so that you do not impact the performance of your instance.

Domain paths query method

You can create effective queries with domain paths.

Use domain paths instead of domain spooling (`sys_domain`) or domain numbering. Queries that use domain paths are much faster than spooling or numbering.

Domain paths is the default query method for instances that have enabled domain separation.

If you want to verify the query method on your instance, look for the following system properties in the admin dashboard:

- *`If domain path is enabled`*: In the System Properties table, you see `glide.sys.domain.provider=domain_paths` and `glide.sys.domain.paths.installed=true`.
- *`If domain path is not enabled`*: In the System Properties table, you see `glide.sys.domain.provider != domain_paths,glide.sys.domain.paths.installed=false`

Slow queries and SQL debugging

Debugging SQL and slow queries can help you resolve slowness issues in an instance.

When you debug an instance, you can either enable SQL debugging to look for slow queries or you can look for slow queries by checking the Slow Queries [sys_query_pattern] table by navigating to **System Diagnostics > Stats > Slow Queries**. This table stores all the slow queries in the instance.

When you search the table, look for queries that contain `domain_path` to determine if any slow queries are due to the domain path in your instance.

If you do find slow queries, try to analyze why they are slow.

Common reasons for slow queries

- A query has too many OR conditions (for more information, see Contains queries and domain access). In the domain hierarchy, place the user or a domain at a hierarchy level where contains or visibility is not needed.
- The query method is not the domain path query method (for more information, see Domain paths query method): If you are not using the domain path query method, contact Customer Service and Support.
- A query needs a database to be indexed so you can see what is in the database quickly. If you can identify the slow query, run the "explain plan" to see if there are options for indexing available. The "explain plan" is a function of SQL that shows the query and what is going on with it.

Before Query business rules

You can use a Before Query business rule to help support data segregation on an instance. ServiceNow applications that support domain separation may support the separation of data and data routing only, have advanced business logic separation, or support tenant (customer) level administration of the application.

A Before Query business rule is supplementary code that you use to support data segregation within domain-separated environments.

Warning: Do not use the Before Query business rule in place of the Domain Separation plugin. This business rule does not prevent data leakage as securely as the plugin.

Using the Before Query business rule for data segregation

You can use the Before Query business rule with data segregation in these situations:

- When domain separation is not supported by a ServiceNow application and you must grant or restrict table or row access to one or more non-internal customers outside of the service provider organization.

  Note: Before you begin developing, contact ServiceNow Customer Support about the application roadmap for that product; domain support improvements may be planned for upcoming releases.

- When a table is domain-separated but access to its rows must be granted or restricted based on certain conditions that apply only to a set of domains in the system.

  Note: For example, a customer in the X domain has multiple vendors supporting that domain and those vendors are granted access to see only the records that are assigned to them.

Points to consider before creating Before Query business rules

You can script Before Query business rules to prevent parent and child table access based on a combination of user information, group memberships, companies, roles, or record-specific field conditions. Before Query business rules are put into separate domains and created to apply globally, to a specific branch of a domain hierarchy.

- Where you can, create Before Query business rules at the lowest possible part of the domain hierarchy so that the rule runs only for users that it applies to.
- Know that there are scenarios in the system where business rules may not run or where a user-triggered interaction may not trigger a business rule to run. For example, a business rule won't run when you have transform maps with Run business rules turned off, or you have scripts with the workflow disabled.
- Always populate the condition field to specify when the rule runs. For example, you can specify if the business rule applies only to certain vendors in a domain.

  Warning: When designing and coding business rules (especially Query business rules), limit OR clauses and searches in non-indexed fields. Too many OR clauses and searches in non-indexed fields can slow queries or affect how your instance performs.

- Use Before Query business rules only when necessary. Too many Before Query rules can affect how your instance performs.

Before Query business rules run before access control lists (ACLs) and perform better in general. This is true especially when you limit the returned results to those users in service provider (SP) environments who have access to several domains in the system.

Note: Filtering the data is transparent (unlike with ACLs) to users who do not see the message `Data Security restricts...` when interacting with data.

When not to use Before Query business rules and ACLs

Be careful when you use Before Query business rules and ACLs to segregate customer data. By using both business rules and ACLs, you create customizations that you then must maintain. Customizations can potentially cause performance issues. Your development teams should create processes to make sure that they don't break the system.

Domain separation provides both scalability and governance with the current domain path query method (v3), which is a widely supported framework. The ServiceNow Platform and App teams are responsible for maintaining the framework, taking the burden off the customer.

For companies with many customers in many instances, excessive use of Before queries and ACLs may cause the database queries not to perform well.

How domain separation is enabled

You can enable domain separation with a ServiceNow plugin. A product manager, supported by a development team, manages the functionality. Enhancements and fixes for domain separation functionality are included with ServiceNow releases. Instance owners can consult Customer Service and Support resources, such as the Service Portal, at https://support.servicenow.com for assistance with domain separation.

Avoiding domain path in scripts

Domain paths can cause the values of your script to change or even break, so don't use them in scripts.

Your script should not depend on the domain path because if you ever change the domain hierarchy, the domain path recalculates and its value changes. When this happens, your scripts are useless or can throw errors or break. The best strategy is not to write your scripts based on domain paths.

Use the **sys_domain** field in your scripts rather than depending on the domain path. If you ever change the domain hierarchy, the domain path recalculates and its value changes, which can cause your scripts to be useless, throw errors, or break. Look for base system business rules, which use the **sys_domain** field, to get some ideas before creating your own scripts.

The ServiceNow platform does not capture the sys_domain_path values in an update set in order to avoid issues with differences in the domain hierarchy for each instance. Therefore, you should validate the domain hierarchy after you import an update set to ensure that the domain path values for your records are correct.

To learn more about domain path, see Request domain separation and Domain Separation Center.

Domain assignments

How you assign a domain impacts the value of the sys_domain field. The assignments contain designs and business properties that affect how the application functions in each domain.

Value of the sys_domain field

The value of the **sys_domain** field contains the domain that is assigned to the record by any of the following:

- Company to which the user belongs
- Business rule that is used when creating the record
- Module that is used when creating the record

- Form template that is used when creating the record
- Domain of the parent record
- Domain that is assigned to the User record
- Domain of the user who creates it

Make sure that your domain assignment strategies and designs are well documented and tested so that you are creating records as those strategies and designs are inserted into the correct domain. That way you can see that the properties of each domain should you need to duplicate or modify them.

Domain separation and the Customer Service Management (CSM) plugin

For the best outcome, be aware of how the properties in the CSM plugin work. When the plugin is enabled, you can see the status of your records in your domains.

Instance owners should contact Customer Service and Support to enable the *csm_auto_account_domain_generation* property.

Note: This base system property is located in the system properties table and is available after CSM plugins are enabled.

What the property does

Whenever a new account in the Customer Service application is created, a domain is created and placed under the TOP domain. If the parent field on the account form is populated, and a new record is inserted, it creates that account as a subdomain of the parent.

What happens if this property is not true and the domain is enabled

New account records in a domain-separated environment are automatically placed in the default domain.

In the header bar, you can see the status of the records with the plugin enabled.

Domain Separation Help

Additional assistance with Domain Separation

Explore, Learn, Develop

Recommended practices

Tips and tricks for creating and developing your domain structure wisely

Concepts for service providers

Concepts that work with the ServiceNow platform to help you solve for common use cases

Support levels by application

Is your application supported for domain separation? See support levels and use cases.

Classes

- For developers: ServiceNow Developer Site Domain Separation
- For service providers: Domain Separation for Service Providers (ServiceNow University login required)

Safe Workplace suite and domain separation

ServiceNow Safe Workplace applications help you reopen your workplaces and support the health and safety of your employees after emergencies and pandemics such as COVID-19. The suite has many applications to help your organization mobilize, recover, and rebuild.

Setup and administration

- Upgrades
- Request domain separation
- Create a domain
- Configuration that can be delegated to internal or external customers

Troubleshoot

- Search the Known Error Portal for known error articles
- Contact Customer Service and Support
- Ask or answer questions in the community

Domain separation setup and administration

Setting up domain separation involves requesting activation of a plugin, setting options, and assigning users and records to domains.

Do the following to set up domain separation:

1. Request domain separation
2. Create a domain
3. Add a domain field to a table

You can also perform these basic administrative tasks on domains:

- Enable or disable a domain
- View domain relationships
- Expand domain scope
- Create a domain-specific choice list

See Advanced domain separation administration for a list of tasks to perform after you set up domain separation and perform basic administration.

Request domain separation

All domain support features are activated with a plugin called **Domain Support - Domain Extensions Installer**. Administrators can request activation of this plugin.

Before you begin

To purchase a subscription, contact your ServiceNow account manager. The account manager can arrange to have the *com.glide.domain.msp_extensions.installer* plugin activated on your organization's production and sub-production instances, generally within a few days.

If you do not have an account manager, decide to delay activation after purchase, or want to evaluate the product on a sub-production instance without charge, follow these steps.

Role required: admin

About this task

If the Domain Support (com.glide.domain) plugin is already active, content in the Domain Support - Domain Extensions Installer and Domain Support - Domain Extensions plugins will not be installed to avoid potential conflict with an existing implementation.

Domain separation replaces Company Separation. Starting with the Helsinki release, the Company Separation plugin can no longer be activated. However, if company separation is already active when you activate domain separation, both plugins are active at the same time. You can control the company separation activation status with the *glide.db.separation.field* property.

Note: Domain paths are used for all customers on Helsinki and later. Domain numbering is no longer used. Customer Service and Support can assist in the upgrade.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Select **Request plugin** to open the **Activate Plugin** form on Now Support.
3. On the **Activate Plugin** form, provide the following information.

Activate Plugin form

Field	Description
What is your target instance	Select the instance that you want to activate the plugin on.
Which plugin would you like to activate	Select the name of the plugin to activate. **Note:** If the system doesn't list the plugin you want or if you're activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	Select the date and time to activate the plugin.

Example

For example, see the following form to activate the Event Management plugin on an instance named SNC Instance.

4. Select **Submit**.
   After the maintenance window, the system installs the plugin on your instance. To confirm the installation, go to the Installed tab in the Application Manager.

Result
Activating the Domain Extension Installer plugin enables these features:

- Domain separation is based on the Domain `[sys_domain]` table.
- Delegated administration lets each domain have separate policy.
- All records are part of the global domain.
- The current user's domain determines the domain to use when viewing or operating on a record in a different domain.

Related topics

Domain separation plugin

Domain separation plugin

The Domain Support - Domain Extensions Installer plugin activates several domain separation features and properties at once. This plugin is typically referred to as the Domain Separation plugin.

Recommended practice for activating the Domain Separation plugin

As part of domain separation development, administrators must request activation of this plugin. For best results, activate the Domain Separation plugin at the start of the development process, preferably before any other plugin is enabled.

Important: Request to activate the Domain Extensions Installer plugin (com.glide.domain.msp_extensions.installer) before activating Domain Separation (plugin com.snc.pa.domain_support).

If you enable Domain Separation towards the end of ServiceNow implementation or once an instance has gone live, your application is at risk in both performance and process. On established instances, depending on how things were structured in development, the risk to the platform and its usability could be high. To learn more about domain separation process, see Exploring domain separation.

For instance, when the Domain Separation plugin is enabled, the **Domain** (sys_domain) column is added to the task table and every existing record is automatically placed in **global**. To use a script to assign all of the records into the correct domains, an established parent/child hierarchy is required. These types of scripted actions risk data corruption or loss, and possible production down time as large amounts of data are moved. Much of the platform code also is placed into **global**, such as business rules, client scripts, form views, and workflows.

If a customer creates code or modifies ServiceNow code, there is risk to platform performance and usability. Instance owners could severely delay their implementation or experience lengthy down-times with this type of approach.

Features of the Domain Separation plugin

These features are enabled when you activate the plugin:

- Domain separation is based on the Domain `[sys_domain]` table.
- Delegated administration lets each domain have separate policy.
- All records are part of the global domain.
- The current user's domain determines the domain to use when viewing or operating on a record in a different domain.

Related topics

Domain separation recommended practices for service providers

Domain system properties and user preferences

Administrators have access to properties and user preferences that control domain scope.

Properties

New activations of domain separation automatically restrict domain scope to the record's domain for all related data or processes. When the user views a record in a form, the record's related data (such as reference picker and related list data) and applied processes (such as business rules and client scripts) are restricted to the record's domain scope. If there are records in multiple tabs, each tab has its own domain scope based on the record opened within that tab. The following properties restrict domain scope to either the record's domain and the user's current session domain.

Domain system properties

Property	Details
glide.sys.domain.use_record_domain_for_processes	Restricts domain scope to the record's domain for all processes. This property does not apply to business rules. Business rules are always processed from the domain record. • *Type*: true | false • *Default value*: true • *Location*: System Property [sys_properties] table
glide.sys.domain.use_record_domain_for_data	Restricts domain scope to the record's domain for all data. • *Type*: true | false • *Default value*: true in new domain activations from Fuji onwards (upgrades from instances older than Fuji do not have this property in the table) • *Location*: System Property [sys_properties] table

When either the *glide.sys.domain.use_record_domain_for_processes* or the *glide.sys.domain.use_record_domain_for_data* property is set to **true**, the following properties are not used, regardless of their setting:

- *glide.sys.domain.use_record_domain*
- *glide.sys.domain.use_record_domain_for_client_scripts*
- *glide.sys.domain.domain_change_notify*
- *glide.sys.domain.no_change_roles*

For a full list of properties see Available system properties.

Note:

In new activations of domain separation starting with the Jakarta release, the session domain determines the business rules executed on the domain table. In previous versions, business rules executed on the domain table were set based on the newly created domain's hierarchy. This behavior is modified by the *glide.sys.domain.skip_domain_insert_businessrules* property. Setting this property to true significantly improves domain insert performance.

Domain scope properties for business rules executed on the domain table

Property	Details
glide.sys.domain.skip_domain_insert_businessrules	Specifies the domain scope for business rules executed on the domain table. In new activations of domain separation, the property default is true and business rules are determined by the session domain. In existing implementations, the property default is false and the business rules are determined by the newly created domain's hierarchy. • *Type*: true | false • *Default value*: True in new domain activations starting with Jakarta. False in existing implementations.
glide.sys.domain.skip_non_global_businessrule_if_nodomain	Ensures that only bus.rules from global domain are executed when using queryNoDomain() or when table is not domain-separated, so you can skip any other business rules • *Type*: true | false • Setting the property to false restores the old behavior and doesn't align with

Domain scope properties for business rules executed on the domain table (continued)

Property	Details
	ServiceNow® recommended practices. • Recommended: Domain-separate your tables; always try to use the record's domain rather than the session domain.

User preferences

In addition, user administrators can set the following user preference globally or on a per-user basis:

Domain scope user preferences

Preference	Category	Updated By	Details
glide.domain.session_scope	Domain	Admin Only	When true, sets the default scope to the user's session domain rather than the record's domain. When false, the default scope is the record's domain. Users with the domain_expand_scope user role can still change the domain scope as needed. • **Type**: true | false • **Default value**: false
glide.domain.session_scope_notification	Domain	Admin Only	When true, displays a visual cue that record values include an expanded domain scope. When false, the notification is hidden. • **Type**: true | false • **Default value**: true

Related topics

Domain separation application properties

Create a domain

You can create a domain by creating a record in the [domain] table.

Before you begin

Role required: admin

About this task

When creating a new domain, keep the following in mind:

- Only one domain can be the default domain.
- Only one domain can be the primary domain.

Procedure

1. Navigate to **All** > **Domain Admin** > **Domains**.
2. Click **New**.
3. Fill in the necessary fields (see table).
4. Click **Submit**.

Domain form fields

Field	Description
Name	Enter a unique name for the domain.
Type	Select a domain type that describes the domain. By default the domain types are *Vendor*, *Customer*, and *MSP*. You can also add your own choices.
Primary	Select the check box if this domain is to be the top-level domain in the hierarchy. The top-level domain only has child domains and no parent domains.
Default	Select if this domain is to be the default domain for your hierarchy.
Parent	Select the name of the domain higher in the hierarchy that contains this domain. This field must have a value for the domain to appear in the domain map.
Active	Select the check box to make the domain available for use. You must select this option for this domain to appear in the domain map.
Description	Enter a description for the domain.

Each domain record can also have several related records:

- Companies
- Contains Domains
- Contained By

What to do next

To change the domain hierarchy, go to the Contains Domains related list and select the domain records that is the child (contained) domains of the contains relationship.

Make a domain the default

Made a domain the default domain to which the system automatically assigns task and user records that are not already assigned to a domain.

Before you begin

Role required: admin

About this task

Note: If you do not set a default domain, then new tasks and user records are placed in the global domain.

Procedure

1. Navigate to **All** > **Domain Admin** > **Domains**.
2. Open the domain you want to be the default domain, for example, Main.
3. Configure the form layout to add the **Default** field.
4. Select the **Default** check box.
5. Click **Update**.

Manually manage the domain for particular records

By default, the system automatically assigns a domain based on the user's company record. In some cases, however, domain administrators want to manually manage which domain a particular record belongs to.

Before you begin

Role required: admin

About this task

The **Managed domain** field allows domain administrators to manually select a domain for the user, group, department, location, or CI record, rather than using the domain assigned automatically from the company record. The **Managed domain** field is available on these record types.

- User records
- Group records
- Department records
- Location records
- CI records

Procedure

1. Navigate to the record you want to manually manage.
2. Select the **Managed** domain check box.
3. From the **Domain** field, select the domain for the record.
4. Click **Update**.

   Clearing the **Managed domain** check box hides **Domain** field and the record uses the domain value from the record's company.

Related topics

Domain separation recommended practices for service providers

Domain Separated Tables

You can see at a glance which tables are domain-separated in your instance with the Domain Separated Tables feature.

Overview

Use the Domain Separated Tables feature to see which tables are domain-separated. Start typing "domain" in the **All** filter to access **Domain Separated Tables** in the left navigation pane.

You can filter this view to display or remove column names to look for certain properties or attributes that are included in your tables. There are two table types that display on the list:

- Tables that have an explicit *sys_domain* column present, where the **Column** name displays *sys_domain* value in the list.
- Tables that are using the *domain_master* attribute to derive domain separation from a referenced record, where the **Attributes** column includes the *domain_master=ref-field-value* value.

When you use the Show Matching or Filter Out selections on the Column name menu, you can view these two types of tables.

Domain Override Viewer

With the Domain Override Viewer, you can see and manage all your process overrides at once across the entire instance.

Overview

Rather than creating an elaborate search for overrides in your scripts, you can use the Domain Override Viewer to find them quickly. Start typing "domain" in the **All** filter to access **Domain Override Viewer** in the left navigation pane.

The table selector in the Domain Override Viewer displays a list of only the tables that contain record overrides, along with a count of overrides in that table. After you select a table, a list of all parent records with overrides is returned. You can quickly view all **parent records**, the **domain** of the parent record, and a **count** of overrides for each record.

Selecting **View Overrides** loads a new tab where all overrides, including the parent record, display in the standard list view. Select a table from the drop-down list to display all records and overrides. Select **View Overrides** on a record to view all of record specific overrides.

Note: Only tables with overrides are listed.

To learn more, see Create domain-separated property overrides.

Enable or disable a domain

When you activate or deactivate a domain, the activation status cascades to companies within the domain.

Before you begin

Role required: admin

About this task

When you activate a company record, domain separation automatically activates the company's associated domain. For example, if you activate the ACME company, then you also activate the TOP/ACME domain.

Procedure

1. Navigate to the domain record.
2. Clear or select the **Active** check box.
3. Click **Update**.

Warning: Do not delete domains. Deactivate domains that you no longer need instead of deleting them.

Add a domain field to a table

As an administrator, domain-separate a custom table by adding a sys_domain field to it.

Before you begin
Role required: admin

About this task
Note:

Do not add domains to base system tables.

Procedure

1. Navigate to the table's list view.
   For example, type <table name>`.list` in the navigation filter.
2. Right-click the list header and select **Configure** > **List Layout**.
3. In the **Create new field** section, enter `sys_domain` as the **Name** and `Domain ID` as the **Type**.
4. Click **Add**.
5. Click **Save**.

Note:

Any other means of creating a field adds a **u_** prefix to the column name. But with the domain field, the system automatically creates the field without the **u_** prefix. You can use the following functionality as a shortcut: Whenever you create a **sys_domain** field, name it **sys_domain** and leave the field type as is. The system automatically sets the field type to **Domain ID** and the field label to **Domain**, saving you a few clicks.

Adding domains to base system tables requires prohibitively thorough testing, updates and adding new logic. In addition in many cases, the source code is not accessible to the customer.

View domain relationships

The domain map offers domain administrators a read-only representation of the active domains on the instance and how they relate to each other.

Before you begin
Role required: admin

About this task
All domain maps must have one domain set as the primary domain. In addition, each domain in the domain map must meet these criteria:

- The **Parent** field must be filled in (the primary domain is the only exception to this).
- The **Active** check box must be selected.

The domain map does not draw domain relationships for domains that fail to meet the mapping criteria.

Procedure

1. Navigate to **All** > **Domain Admin** > **Domain Map**.
2. Click the plus (**+**) or minus (**-**) icons on the domain headers to show or hide sub domains.

Select a primary domain

The primary domain indicates the top-level domain in the domain map.

Before you begin

Role required: admin

About this task

The primary domain cannot have a parent domain and must have at least one child domain. There can only be one primary domain at a time. If you select another domain as the primary domain, it overrides the previous primary domain.

Procedure

1. Navigate to **All** > **Domain Admin** > **Domains**.
2. Select the domain you want to be the primary domain, for example, TOP.
3. Select the **Primary** check box.
4. Click **Update**.

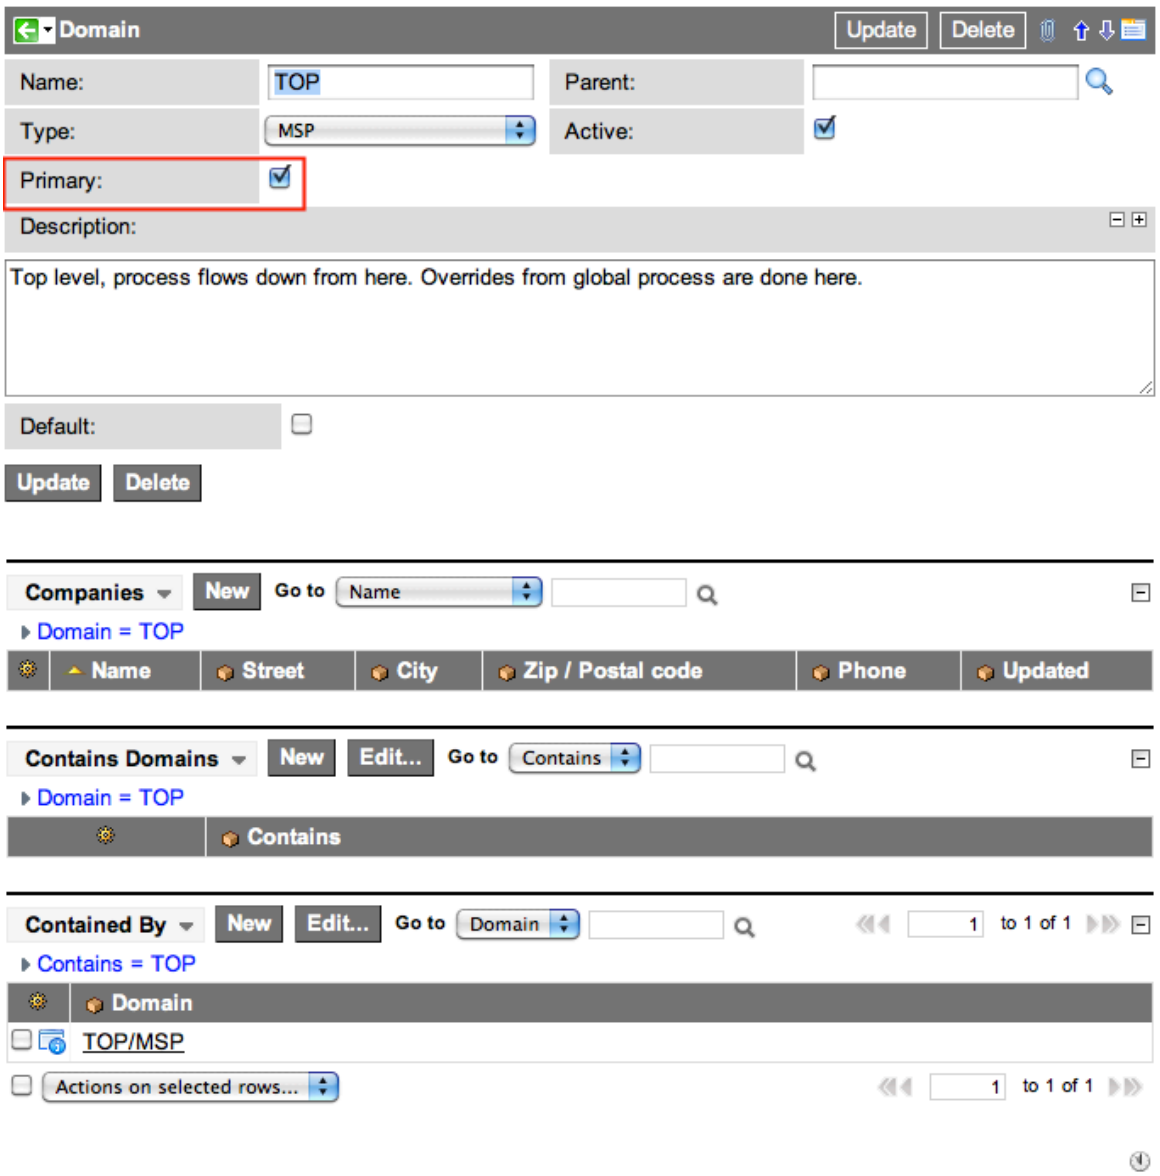

Create Contains relationships between domains

Create a "contains" relationship between domains to change the domain hierarchy.

Before you begin
Role required: admin

About this task
Domains in a contains relationship inherit the settings of the containing domain. The containing domain allows users to see data in the contained domain as well as any of its children. Processes are unaffected by a contains relationship.

Procedure

1. Navigate to the domain table.
2. Select the domain record that is the parent (container) domain of the new contains relationship.
3. Toggle the domain scope to switch between the session scope and record scope, if necessary.
4. From the Contains Domains related list, click **Edit**.
5. Select the domain records that is the child (contained) domains of the contains relationship. Only child domains appear by default when the domain picker is set to Global. Toggle the domain scope to see all domains.
6. Click **Save**, and then click **Update**.

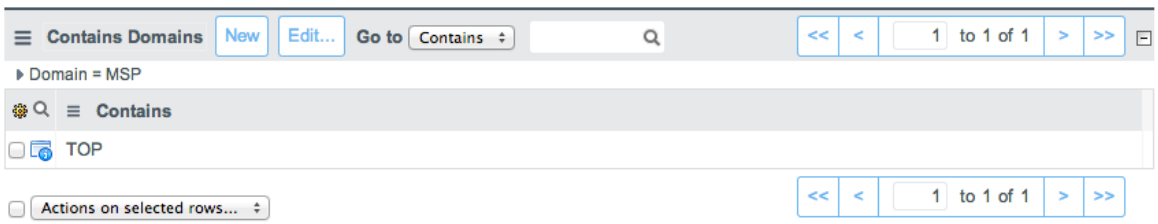

Related topics

Domain separation recommended practices for service providers

Expand domain scope

By default, when a user in the global domain views a table containing a **sys_overrides** column, the user sees records from only the global domain. When an admin in the global domain views a process table, that admin sees only records that are in that process table.

Before you begin
Role required: admin

Procedure

1. Change the *glide.sys.restrict_global_domain_processes* property to **true**.
2. To view records from all domains, click **Expand Domain Scope** under Related Links.
3. To return to viewing records from the global domain only, click **Collapse Domain Scope**.

Add domains to a visibility domains list

Adding a visibility domain allows a user or group to see and potentially edit records from another domain regardless of the user or group's normal domain membership.

Before you begin
Role required: admin

About this task
Assigning visibility domains to all members of a group is preferred over granting them to individual users.

Note: Adding a visibility domain does not change a table or record's access control rule requirements.

Procedure
1. Navigate to the group table.
2. Select the group you want to provide with visibility domains.
3. Add the **Visibility domains** related list to the form.
4. From the **Visibility domains** related list, click **Edit**.
5. Select the domain records you want the group or domain to see.
6. Click **Save**, and then click **Update**.

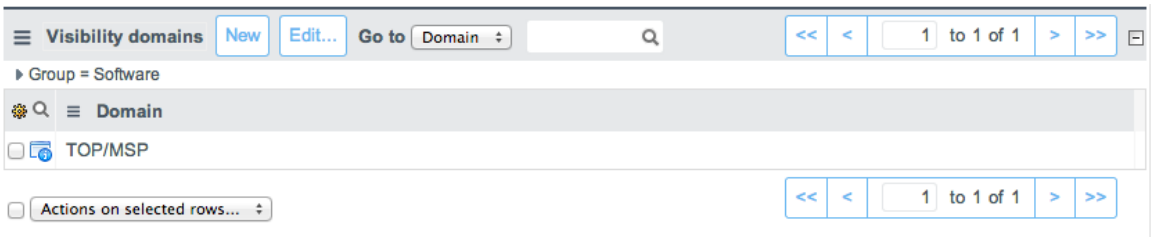

Grant visibility domains to an individual user

While it is possible to add visibility domains for specific users on the User form, it's best to add them only via groups. This controls permissions and access should individuals change departments or leave the company.

Before you begin
Role required: admin

Procedure
1. Navigate to **All** > **User Administration** > **User**.
2. Select the user you want to provide with visibility domains.
3. Add the **Visibility domains** related list to the form.
4. From the Visibility Domains related list, click **Edit**.
5. Select the domains whose records you want the user to see.
6. Click **Save**, and then click **Update**.

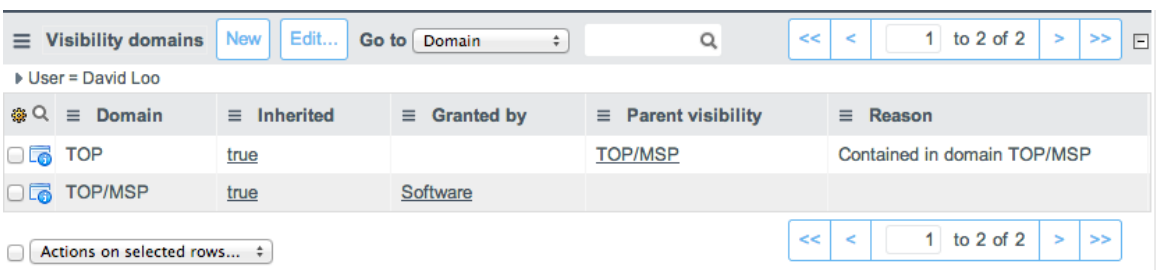

The Visibility domain embedded list contains the following fields.

Field	Description
Domain	Domain that is visible to the group or user.
Inherited	Domain is inherited from domain visibility or a parent domain.
Granted By	Name of the group that granted domain visibility.
Parent visibility	Name of the parent domain and used for grouping records. If the parent record is deleted, then all records with the same parent are deleted as well.

Create a domain-specific choice list

Administrators can configure choice lists to contain entries specific to a particular domain.

Before you begin

Role required: admin

Procedure

1. Select the domain from domain picker where the choice should be added.
2. Right-click the choice field you want to customize and select **Configure Choices**.
3. Update or add choices.
4. Push changes through the normal change process such as update sets.

> **Note:** Administrators should ensure that choices are unique across domains to prevent administrative confusion in the global domain.

If an administrator adds a new choice from the global domain, then users from domains lower in the hierarchy see the new choice at the end of their current choice lists. If the new choice is not active at the global level, then it is available to the domain users via **Configure Choices** but does not show as an active choice.

Advanced domain separation administration

Administrators can view information about domain separation, identify potential issues, and change configuration settings.

You can perform these advanced administrative tasks on domains:

- Use domain selection menus
- View domain relationships

Use domain selection menus

The instance offers domain selection via two menu formats.

- Domain selector: provides a simple drop-down list of available domains.
- Domain reference picker: enables a reference field that offers filtering and an auto-complete, type-ahead entry feature. Use this format for longer lists.

The placement of these pickers and the procedure to show or hide them differ depending on the user interface version.

Enable domain selection menus in Core UI

Displaying the domain picker in Core UI enables the domain selector by default. After enabling the domain selector, you can add a system property to enable the domain reference picker.

Before you begin

Note: The domain separation (plugin com.snc.pa.domain_support) is required to enable the domain reference picker.

Role required: admin

About this task

Procedure

1. Click the gear icon in the header.
2. On the General tab, click the **Show domain picker in header** switch.
   The domain selector appears in Core UI header.
3. **Optional:** Enable the domain reference picker.

   Note: Enabling the domain reference picker removes the global option from the list. To return to your home domain, click the return arrow next to the reference field. Admin users can click the return arrow to return to the global domain.

   a. Enter `sys_properties.list` in the application navigator.
   b. If not already present, add the `glide.ui.domain_reference_picker.enabled` property and set the value to **true**.
   c. Refresh the browser.
   The domain reference picker appears in Core UI header.

Restrict access to the domain picker

Use a system property to restrict access to the domain picker in Core UI and Next Experience.

Before you begin

Role required: admin

About this task

By default, users with the ITIL role, and roles that include the ITIL role (such as the administrator), can access the domain picker in Next Experience. You can grant other roles access by adding them to the property or restrict roles by removing them. It is recommended to restrict the role to admins only.

Admins can grant access to users by creating a system property in the sys_properties table.

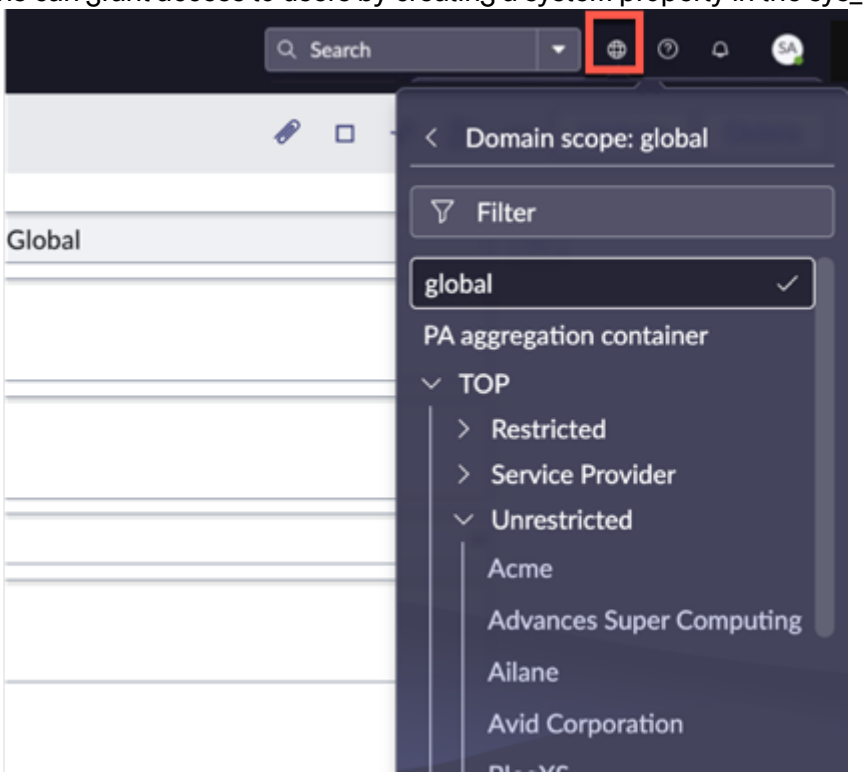

Procedure

1. Open the System Properties [sys_properties] table.
2. Add this property: `glide.ui.polaris.domain_picker.role`
3. Configure the property value as a comma-separated list of roles: **admin, itil**.
   To learn more, see Configure Next Experience pickers

Domain separation application properties

The Domain Separation plugin has two new tables to give service providers more flexibility in customizing their applications that use domain separation. These tables are the System Application Property table [sys_application_property] and the System Application Property Value table [sys_application_property_value].

New tables offer more options

With service provider (SP) applications, certain actions can vary depending on the domain. However, the ServiceNow® base system System Properties [sys_properties] table is not domain-separated, so it doesn't satisfy the requirements for applications that use domain separation.

Each SP customer may want to customize their applications differently. Previously, features that could be customized were defined as only one global value. Application developers need a more flexible table. Now you can modify your application without having to create code every time you want to add or change the functionality.

How overrides work in the new tables

Developers typically use the ServiceNow System Property [sys_properties] table to create various functions in applications. If you wanted to develop an application to behave differently in different domains, you'd have to customize it yourself.

In the Paris release, the new Application Property [sys_application_property] table simplifies that customization. Instead of going directly to the System Property table for a value, the application property table goes to the System Application table first. This new table now stores the logic that you require to configure your application. If it finds a property in the new table, it uses that content. If there is no information in that table, it moves on to the base system properties table.

When you configure support for domain separation, you can add domain logic to this new Application Properties table. This table can contain properties that don't exist in the System Properties table. Or you can add properties to the configuration table that can override any property that you select in the System Properties table.

For example, let's say that you want to configure an application with a First day of the week feature. Sometimes, you might want the first day of the week to be Sunday. In other cases, you might want the first day of the week to be Monday. In the base system table, there might be only one Day 1 option, which is Sunday. With the new table, you can store another property, making Day 1, Sunday, and a child domain, Monday.

This figure shows how the system draws properties from the Application Property table before going to the System Property [sys_properties] table.

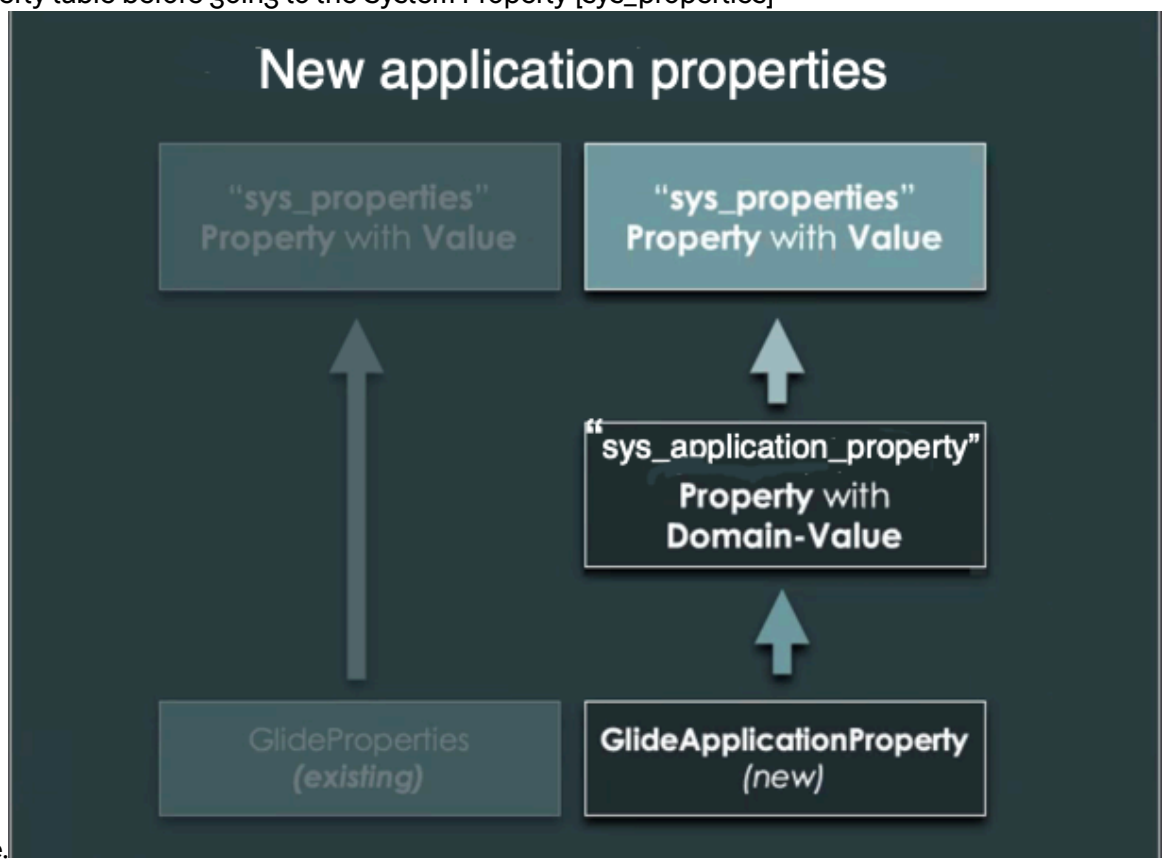

How scoped apps work in the new table

The new Application Properties table is supported from scoped applications. The application property name, similar to the system property name, is unique, which means that it is prefixed with the scope name if it is not global. The scope of an application impacts your configuration. The scope may determine which Day 1 is defined as Sunday and which as Monday. You can use the same property but customize it so that Day 1, Sunday is the parent domain and Day 1, Monday is the child domain. In the new table, there is both a domain column and a scope column, so you can set these properties for each.

You can use the Expand Domain Scope view in the sys_application_property_value table to show all overrides, as shown in the following image.

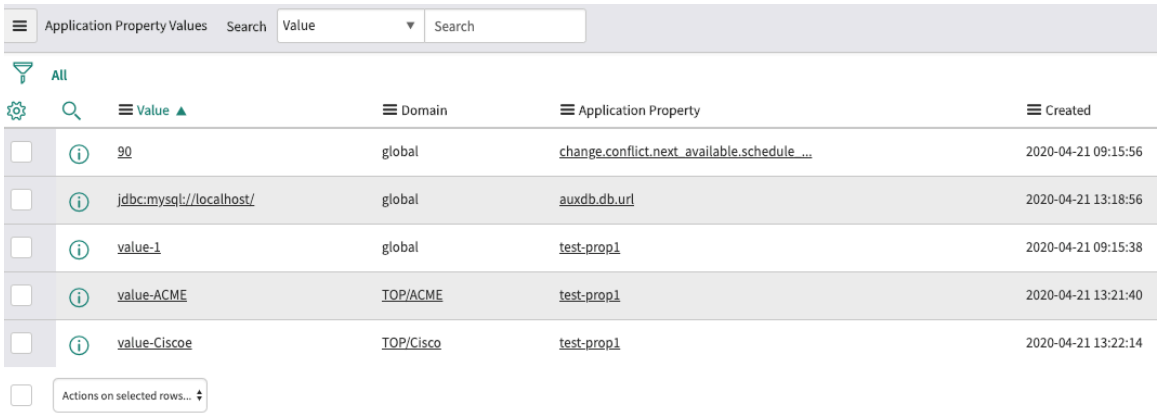

Note:

If those tables are not available, make sure that you have activated the Domain Extension Installer (com.glide.domain.msp_extensions.installer) plugin.

New application property tables

The new System Application Property [sys_application_property] table contains these fields:

- name
- description
- type (choice of string, true|false, integer, time zone, color, and so on)
- default_value
- property (reference to sys_properties)
- usage_notes
- read_roles
- write_roles
- unique key: (name)

The new System Application Property Value [sys_application_property_value] table contains these fields:

- sys_application_property (ref to sys_application_property)
- sys_domain
- sys_overrides
- value
- unique key: (sys_application_property, sys_domain)

New APIs

The new APIs are also supported in scoped apps. Domain-separated application properties have distinct APIs. The GlideApplicationProperty API has two new scriptable methods, available in both global and scoped applications. See GlideApplicationProperty - Scoped, Global to learn more about these new APIs.

Related topics

Domain separation recommended practices for service providers

Domain Migration Tool

Use the Domain Migration Tool to move a customer from a domain-separated environment to their own dedicated instance.

Domain Migration Tool plugin

The Domain Migration Tool plugin (com.glide.domain.migration_tool) simplifies the task of moving a customer from a domain-separated environment to a more flexible dedicated instance. Customers may want to migrate to a separate instance to take greater advantage of the ServiceNow AI Platform capabilities. Although the Domain Separation plugin is installed, the data and process separation properties are turned off.

Note: You must request a cloned instance and request activation of the Domain Migration Tool plugin before you can use it.

The Domain Migration Tool runs only if both data and process separation are enabled in the domain-separated instance:

- The *glide.sys.domain.partitioning* data property must be set to **true**.
- The *glide.sys.domain.delegated_administration* process property must be set to **true**.

What the migration tool does

- Automates much of the migration process, especially cleaning up data.
- Migrates the domain-separated instance to a new dedicated instance.
- Deletes data from the dedicated instance.

  Note: The tool doesn't delete data that is global, in the target domain, or in additional data domains (if specified).

- Collapses process data, or if unable to collapse, deletes process data
- Keeps process records that are visible to the target domain
- Updates the sys_choice, sys_ui_list, and sys_ui_related_list special tables
- Cleans up records added by domain separation plugins:
  - Business rules
  - UI actions
  - Scheduled jobs
  - Installation exits
  - Navigation modules
- Disables domain separation and removes domains from the cloned instance:
  - Sets these properties to **false** in the cloned instance:
    - *glide.sys.domain.partitioning*
    - *glide.sys.domain.delegated_administration*
    - *glide.sys.domain.enabled*
  - Deletes all domains except the target domain and any additional specified data domains.
- Updates the **Status** field in the domain_migration_tool_status table.

Individual status of the tables

Status	Description
Pending	Default status of domain separated tables during migration. The tables are scheduled to be migrated, but have not yet started the migration.
Failure	Table level failure. If the migration process is finished with errors, this status indicates which tables have errors.
Running	The status of the table that is currently migrating. Only one table can have this status and currently migrating.
Successful	The status of the tables that have successfully migrated.
Finished successfully	The migration process is finished without errors.
Finished with errors	The migration process is finished with errors.

- Logs progress and status to syslog_domain

  The source is **MigrationTool** for all log entries associated with migration.

  - Logs each data table and how many data tables remain
  - Logs each process table and which domain's records are currently being inactivated or deleted

What the migration tool doesn't do

- Clone the instance
- Create another domain-separated instance
- Migrate records (data or process) if either data or process separation properties are disabled before you run the tool
- Alter any data on the source instance
- Delete data that is global, in the target domain, or in additional data domains (if specified)

What you must do after running the tool

The Domain Migration Tool automates the removal of data outside the desired domains (the target domain, any additional data domains, and the global domain). You must evaluate all the remaining configurations to ensure that they are appropriate and work for your dedicated instance. For example, if you had a business rule that set the domain field on records, you might want to disable this business rule, since it no longer serves a purpose.

Migrate a domain-separated instance to a dedicated instance

Move a customer from a domain-separated environment into their own dedicated instance environment.

Before you begin
Role required: admin

Procedure

1. Elevate to the security_admin role.
   See Elevate to a privileged role for details.
2. Navigate to **All** > **Domain Admin** > **Domain Migration Tool**.
   Also access with **domain_migration_tool_status.list**.
3. Click **New**.
4. Complete the form.

Field	Description
Target domain	Specify the domain used for both process and data that you want to migrate. Only the target domain is retained, not any of its children, unless specified in the **Additional data domains** field.
Additional data domains	Optionally, specify any additional data domains you want to migrate. If you want to migrate the target domain and all of its children, you must specify all children.

5. Select **Submit**.
6. Open the form you just submitted.
7. Navigate to **All** > **Domain Separation Center** > **Configure Audits** .
   See Domain Separation Center for additional information.
8. Set the **Validate Domain Separated Table Schema** audit to **Active** and assign a schedule.
   This is a precautionary domain audit to pre-check the schema for the health of domain separated tables. This enables you to fix errors before running the migration.
9. Run the **Audit Schedule** that includes the schema.
   See Execute audits immediately for details.
10. Address issues returned from the audit.

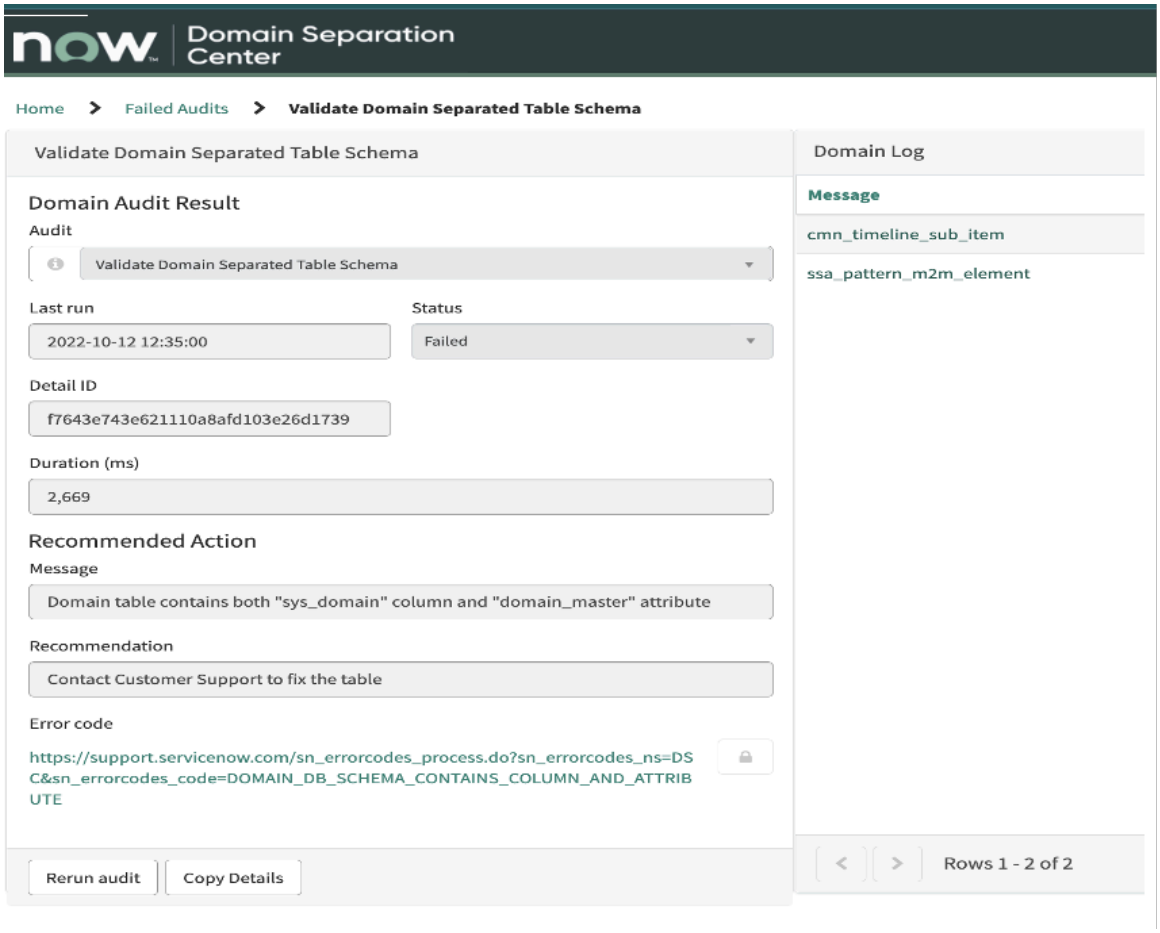

11. Select **Start Migration**.
    - The execution tracker progress bar and the domain migration tool are triggered.

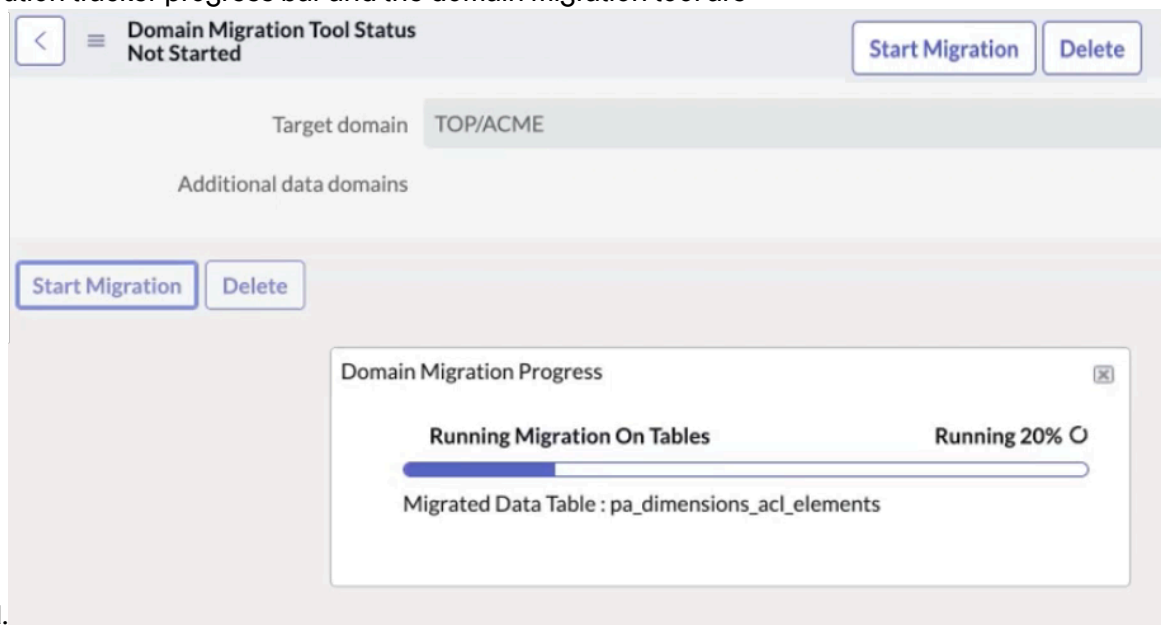

    - The current progressing migrating table is displayed along with the total percentage of successfully migrated tables.
    - The table where all of the domain separated tables are recorded with the migration status, total number of records in each table, and the number of records migrated.
    - The count of tables where migration failed is also recorded.

- The **Status** is updated as the tool performs its functions.
- Domain Migration is still running in the background if you close the progress bar. Access the `sys_execution_tracker` table and look for **Running Migration on Tables** to check the migration process running in the background.

The Domain Migration Tool Status displays the following fields:

Domain Migration Tool Status fields

Field	Description
Status	Displays the migration status. ○ **Migrating data tables...**: In progress status. ○ **Migration successful**: The updated status after a successful migration. ○ **Misconfiguration is observed on one of the domain separated tables**: Indicates a failure of data migration. A schema error has been found, the migration won't start. Run the audit **Validate domain separated table schema**. The audit will fail and show the tables that don't follow the schema standard. ○ **Finished with Errors in Revisit Tables**: The number of tables to address schema migration is listed. Locate the failures in the Revisit tables count and address schema issues. ○ **Domain Migration Failed for Tables...**: Non target records should be deleted manually from failed tables. ○ **Finished Successfully**: When all tables have migrated.
Target domain	The selected target domain for migration.
Additional data domains	Populates if multiple migration domains have been selected.
Revisit Tables Count	This field populates only if a table fails to migrate. If there are no failures, this count is zero. In this case, the tables will be revisited to reattempt the migration. If there are no failures, no need to revisit the tables or reattempt the migration.
Current Progressing Table	Displays the name of the table that is currently being migrated. After successful migration, this field will be empty.

12. Select the **Click here** link on the Domain Migration Tool Status page to open the Execution Tracker.

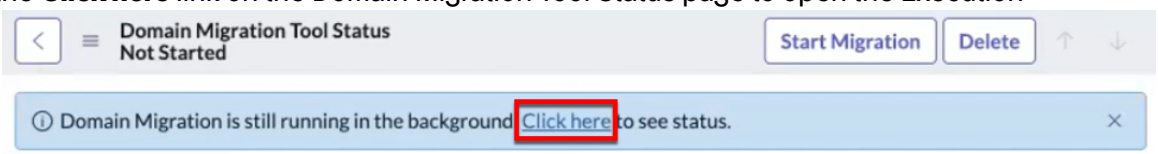

You can also access the `sys_execution_tracker` table and look for **Running Migration on Tables** to check the migration process running in the background.

13. Select **Show status** in the Related Links section to access the progress bar any time during the migration.

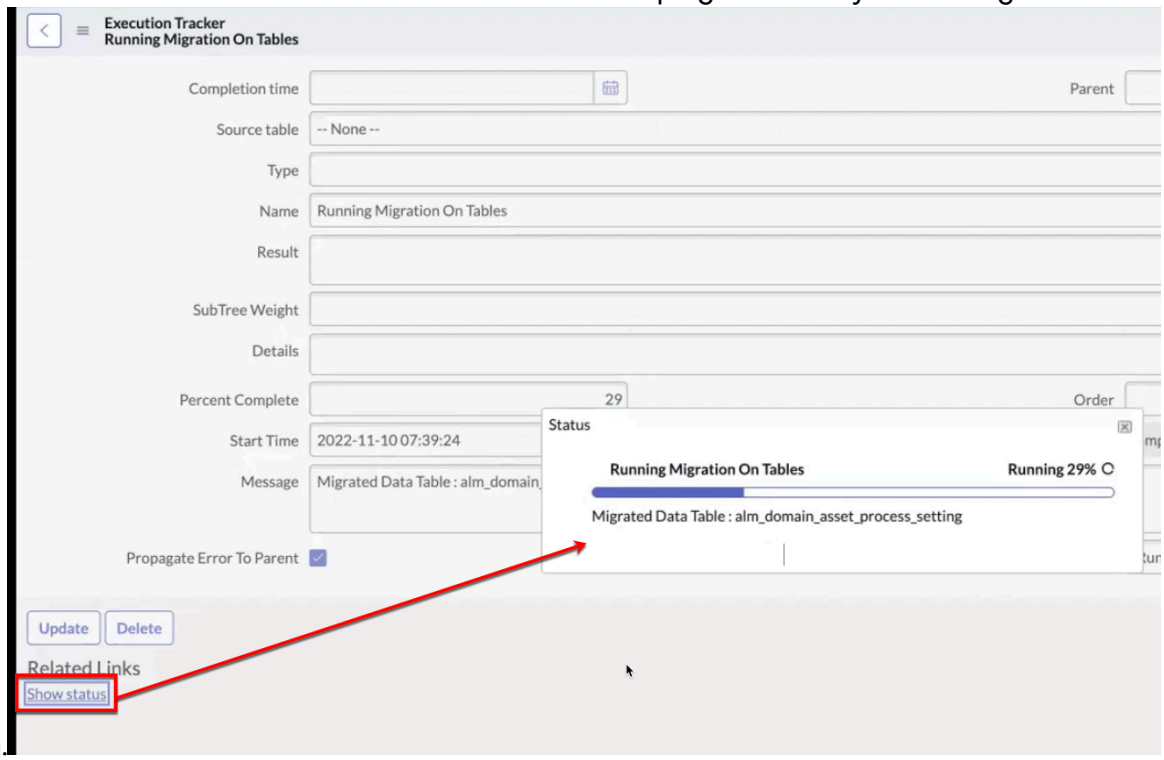

In the case that a table failed the schema check, the Migration Status of Domain Separated Tables overall status will be **Failed**.

There will be **Failed** entries for each corresponding table.

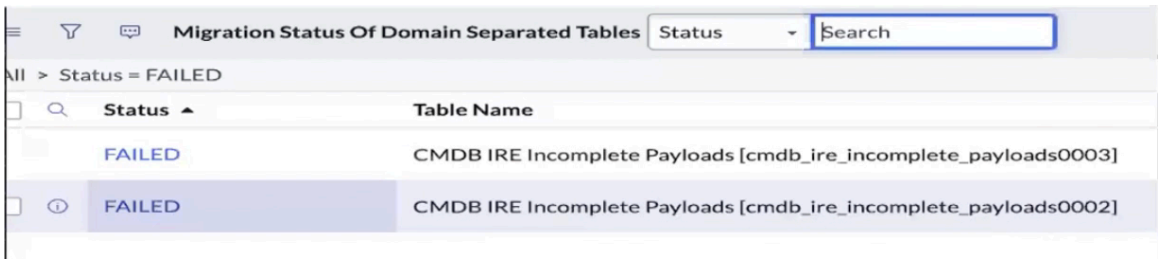

The remainder of the migration continues and a summary of all failed tables and the total number of **Revisit Tables Count** will be populated in the Domain Migration Tool Status.

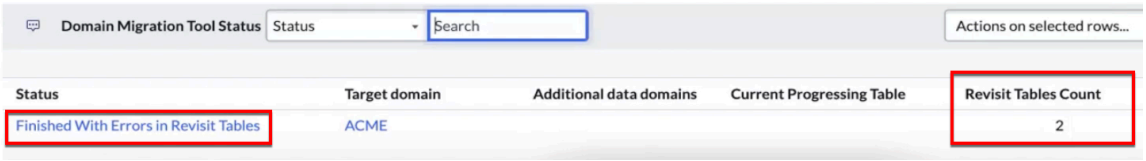

The **Status** is **Finished Successfully** when the migration is complete.

Process administration

Process administration allows administrators to set domain-specific policies.

The policies set lower in the domain hierarchy override policies set higher in the domain hierarchy. While in a domain, administrators can set domain-specific versions of these global policies and settings:

- Client scripts
- System policies
- Application and module names
- Application roles
- Module filters

> ⚠ **Warning:** All users with the admin role have special access to all system features, functions, and data because administrators can override ACL rules and pass all role checks. Grant this privilege carefully.

When users have the **admin** role, then all policies in the instance are available to them regardless of the assigned domain. They can enter a specific domain, and then only policies in that domain or higher are visible and processed during a relevant transaction. When an administrator modifies a policy that is in a higher domain or the global domain, the system automatically creates a new record for that administrator's current domain. It does not modify the original policy, application, or module record. This new record overrides the original.

To make changes to a policy in a lower-level domain, go into that domain and modify the policy. This approach creates the policy record in your domain that overrides the original, higher-level policy record.

Do not change the higher-level policy and then change the **Domain** field on that policy. This approach does not create a policy record in your lower-level domain, nor does it keep the policy record for the higher-level domain.

The **sys_overrides** field indicates that a policy, application, or module at a lower level in the hierarchy overrides a record at a higher level. The system automatically sets this field when an

administrator attempts to modify a policy, application, or module that belongs to another domain higher in the hierarchy.

Rather than changing the higher-level record, the attempted update is changed into an insert, and the **sys_overrides** field is set to indicate the higher-level policy, application, or module that is being overridden. Later when the records for a relevant transaction are loaded, the overriding domain-specific policy, application, or module is used instead of the original.

Domains for process administration

By default, process administration always uses the record's domain to determine what policies to apply.

The record's domain takes precedence over the user's domain. If there are no policies in the record's domain, delegated administration checks for policies in the next highest level of the domain hierarchy. The search for domain policies continues up the domain hierarchy until reaching the global domain. If there are no domain policies lower in the domain hierarchy, processes administration uses the policies for the global domain.

For example, Fred Luddy is a user in the Acme domain who can see records in the child domains of Acme: Atlanta, Acme: San Diego, and Acme: NY child domains. When this user opens a record in the Acme: San Diego domain, process administration first checks for policies in the Acme: San Diego domain. If there are no policies at this level of the domain hierarchy, process administration checks for policies from the Acme domain. If there are no policies in the Acme domain, process administration uses the global domain polices as there are no other domains higher in the domain hierarchy.

Sample process administration with domain specific applications

The following example illustrates process administration with domain-specific applications and modules.

As the administrator of the Oceanic domain, David Loo decides to customize the Configuration application. To start with, David reviews the modules available in the Configuration application module.

Starting view of the Configuration application

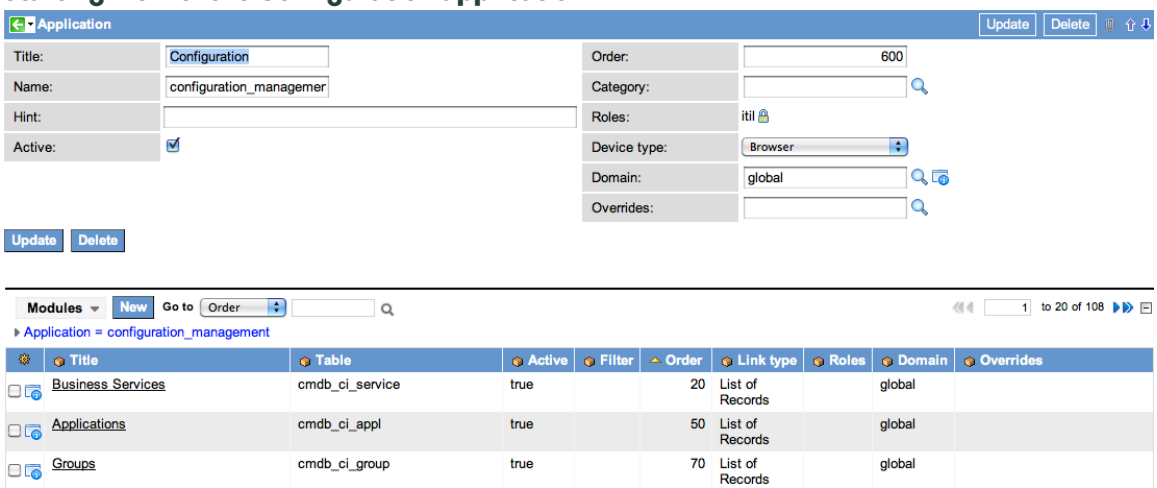

David decides to rename the Configuration application to CMDB and to allow the inventory_admin role to see the application.

Sample domain-specific changes to the Configuration application

Applications | New | Go to Title | 1 to 20 of 46

All > Active = true > Device type != Mobile

Title	Active	Order	Roles	Name	Domain	Overrides
Asset Management	true	900	asset	asset	global	
BSM Map	true		admin itil	bsm_map	global	
Change	true	400	itil	change_management	global	
CMDB	true	600	itil inventory_admin	configuration_management	Database	configuration_management
Content Management	true		content_admin	cms	global	
Contract Management	true	1,000	asset contract_manager	asset_contracts	global	
Domain Admin	true		domain_admin	domain_admin	global	
ECC	true		admin	ecc	global	
Homepage Admin	true		admin	home	global	
Incident	true	200	itil	incident_management	global	
Instance Clone	true		clone_admin	instance_clone	global	
Knowledge Base	true	800	knowledge	km	global	
Metrics	true		itil_admin metric_admin	metrics	global	
MID Server	true		admin	MID	global	
Organization Management	true	875	asset	organization_management	global	
Problem	true	300	itil	problem_management	global	
Reports	true	1,100	itil	reports	global	
SAML 2 Single Sign-on	true		admin	saml_2_single_sign_on	global	
SAML Single Sign-on	true		admin	SAML Single Sign-on	global	

Activate | Deactivate | Actions on selected rows... | 1 to 20 of 46

Next, David decides to change the Incident application by activating the **Open - in "New" State** module and adding a new filter item to show open incidents in the Oceanic category.

Sample domain-specific changes to the Open - "New" State module

Module | = Required field | Update | Delete

Title:	Open - in "New" state	Link type:	List of Records
Table:	Incident [incident]	View name:	
Order:	200	Roles:	
Application:	incident_management		
Hint:			
Active:	☑		
Image:			

Filter:

	Incident state	is	New
and	Active	is	true
and	Category	is	Database

Arguments:

```
incident_state=1^active=true
```

Update | Delete

This creates a new module entry in the application rather than overwriting the existing module in the global domain.

Domain-specific view of the Incident application

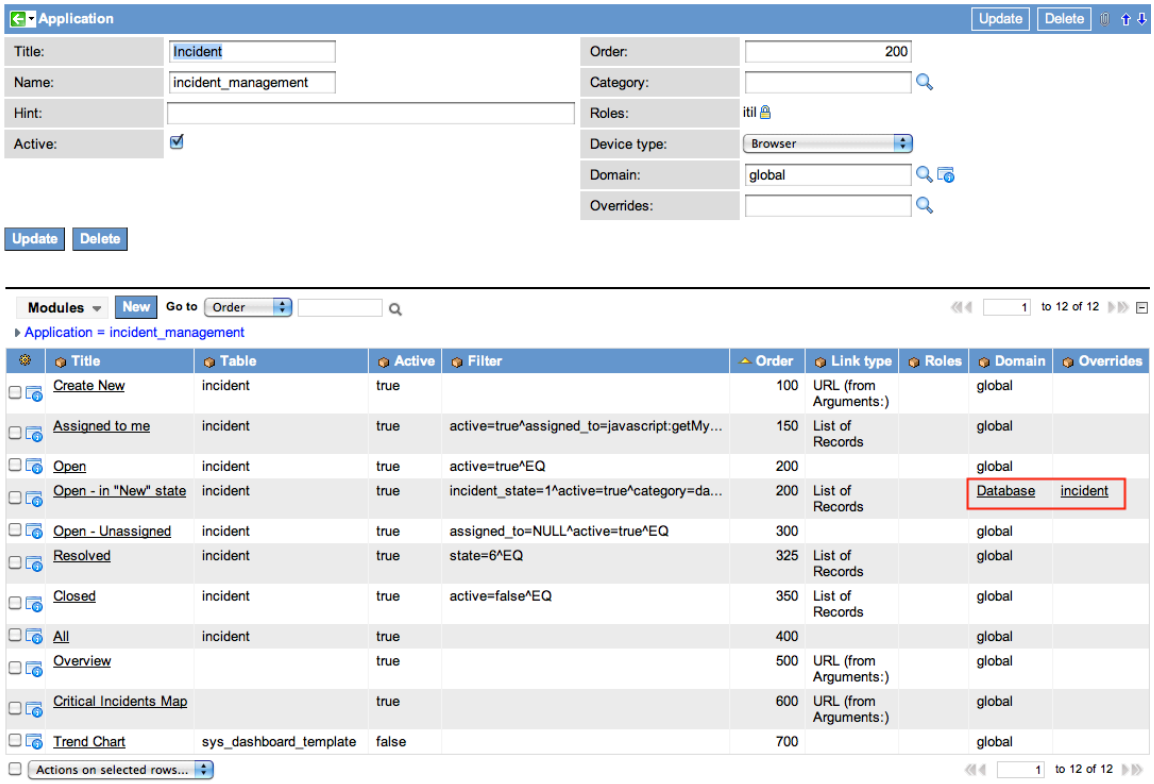

If another administrator from another domain, such as Fred Luddy, logs in and looks at the Configuration application, the settings from the global domain appear.

David Loo's view of applications

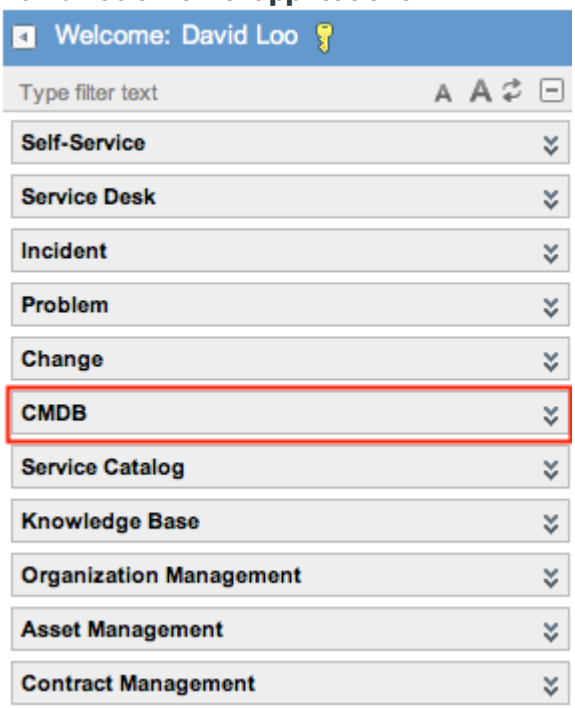

Fred Luddy's view of applications

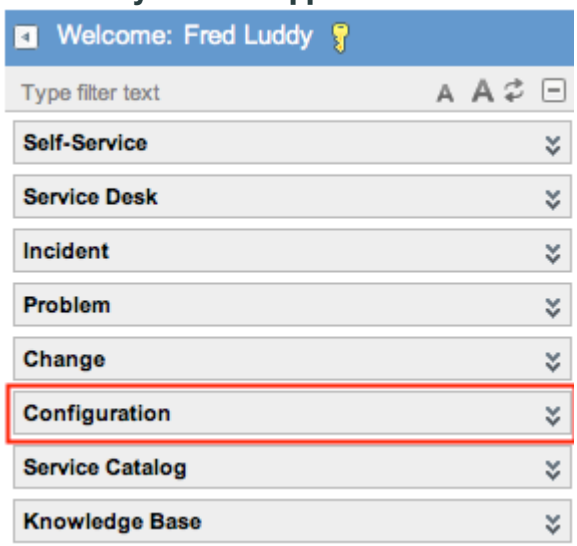

Enable verbose domain logging and debug messages

Domain log and debug messages allow you to troubleshoot domain configuration errors.

Before you begin

Role required: admin
Make sure you are the latest instance for best performance.

Procedure

1. In the Domain Separation Center, navigate to **Domain Admin**.
2. Click **Configure Domain Center**.
3. For **Enable verbose domain logging**, select the **Yes** check box.
4. Click **Update**.

View a real-time domain message

You can view real-time domain messages from the system logs.

Before you begin

Role required: admin

Procedure

1. Enable verbose domain logging: Navigate to **All** > **Domain Admin** > **Domain Separation Center** > **Configure Domain Center** > **Enables detail domain logging** > **Yes** (or set property `glide.sys.domain.verbose` to **True**).
2. Navigate to **System Diagnostics** > **Session Debug** > **Enable All**.
   Because this is a real time review, there is no need to let the debug session run prior to checking the log files.
3. Navigate to the session debug console to view the detailed system domain logs.
4. Search for the text Query against table.

   This query finds log messages in this format:

```
08:36:43.974: [Domain Paths] Query against table
 incident restricted by domain values [Database
 Atlanta[db53580b0a0a0a6501aa37c294a2ba6b],
Database[287ee6fea9fe198100ada7950d0b1b73],
Database San Diego[db53a9290a0a0a650091abebccf833c6], global,
 NY DB[5f74727dc0a8010e01efe33a251993f9]]
```

In this example, the user viewing the Incident table only saw records that matched the Database Atlanta, Database, Database San Diego, global, and NY DB domains.

View a historical domain message

View historical domain messages in the log file to troubleshoot domain separation issues.

Before you begin

Role required: admin

Procedure

1. Enable verbose domain logging: Navigate to **All** > **Domain Admin** > **Domain Separation Center** > **Configure Domain Center** > **Enables detail domain logging** > **Yes** (or set property `glide.sys.domain.verbose` to **True**).
2. Navigate to **System Diagnostics** > **Session Debug** > **Enable All**.
3. Let the debug session run for a time period, such as a day, before checking the log files.
4. Navigate to **System Logs** > **Utilities** > **Node Log File Download**.
5. Open the record for the day you want to view.
   Log files use the naming format `localhost_log.<yyyy-mm-dd>.txt`.
6. Click the **Download** log related link.
7. Open the downloaded log file in a text editor and search for log messages with the following format:

   Query against table incident restricted by domain values `[global, Software[8a4dde73c6112278017a6a4baf547aa7]]`

   In this example, a user only saw records from the Incident table that matched the global and Software domains.

Troubleshoot domain separation errors

If you encounter domain separation issues, review this list of solutions.

Error or symptom	Solution
A domain sys_id points to a non-existent domain	This error occurs when a data record, such as a user or task record, has a `sys_domain` column value whose `sys_id` does not exist in the current domain table. The domain `sys_id` could have been accidentally deleted or it could refer to a previous domain table if you changed the domain table. To fix the error, open a list for the table containing the error, filter on the invalid `sys_domain` value. Then, either manually enter the correct `sys_domain` value or remove it.

Error or symptom	Solution
	Note: You can have invalid domain `sys_ids` in any table that references the domain table. For example, invalid domain IDs can occur in the User Visibility Domain `[sys_user_visibility]`, Group Visibility Domain `[sys_user_group_visibility]`, and Contained Domain `[domain_contains]` tables.
A domain path or domain number sys_id points to the wrong domain	This error occurs when a domain number or domain path query is out of sync with the actual domain name. This error can occur with domain numbers when adding domains requires renumbering or during the conversion from domain numbers to domain paths. To fix the error, check the results in the Domain Separation Center. If the error persists, you can manually edit the value for the `sys_domain_path` or `sys_domain_number` columns to point to the proper domain.
The domain tree structure is corrupt	This error occurs if there is a series of domain contains relationships that create an infinite loop among domains. To fix the error, open a list for the domain table and manually edit the domain contains values to not form a loop.

Post-Production Domain Separation Activation Utility

The post-production Domain Separation activation utility aids in the activation of Domain Separation in a live environment.

Post-Production Domain Separation Activation Utility plugin

The Post-Production Domain Separation Activation Utility plugin (com.glide.domain.activation_utility) simplifies the task of creating a domain-separated environment in a post production live environment. Customers may want to activate Domain Separation on post production environments to take greater advantage of the ServiceNow AI Platform capabilities. The utility provides a step-by-step guided setup for creating domains.

Note: The process of populating domain records in domain separated tables requires downtime or restricted access to the instance.

What the Activation utility does

- Provides a step-by-step guided setup for creating domains.
- Runs a background job to handle the installation of domain separation with both Process and Data Separation properties disabled.
- Keeps process records that are visible to the target(first) domain
- Detects, identifies, and logs errors during the domain creation process
- Presents resolutions for common errors
- Logs all actions performed during the setup process for audit purposes
- Generates a detailed summary, including all actions taken, domains created, and any changes made to the syste

Domain Job Management

Queue multiple Domain Separation updates sequentially into a single background job using the Domain Job Manager.

When a Domain record is modified such as reparenting or a change in hierarchy, by default, a job executes to fix all the records in domain separated tables. You can control when this behavior occurs with the **Domain Job Manager** and monitor ongoing job progress with the **View Job Progress** button.

Use the **Domain Job Manager** to pause domain jobs, queue multiple domain table changes and trigger the queued job when ready with all the changes.

Domain Job Progress

Label	Description
Job Type	The current job type
Status	The status of the job, see **Job state** in the job manager for more information.
Progress	Ongoing job progress

Domain Job Manager

Field	Description
Job Type	Support job types: Hierarchy change
Job State	• Active • Paused **Note:** By default paused jobs will auto resume in 1 hour. Check this option to disable this and manually activate the job later

Delete by domain

Clean up inactive leaf level domains in the domain hierarchy with a controlled and automated tool.

Before you begin

Role required: domain separation admin

Procedure

1. **All** > **Domain Admin** > **Cleanup Queue**
2. Select the **Prepare For Cleanup** button.
3. Select the domains to queue for deletion.
   The list shows all the available inactive leaf level domains that can be selected for deletion.
4. Select **Save**.

Selected domains will now reflect a **Ready for Staging** state in the table. Domains must be staged before they are queued for deletion.

5. Select the **Move to Staging** button.
6. Enter the number of days to retain the staged domain before deletion.

> **Note:** Once the domain moves to the **Staged and Deletion Planned** state, it cannot be removed from the clean up queue.

Before a domain is deleted, domain administrators can access the data footprint of the domain.

7. Select the **Preview Domain Data** button.
   A job will be queued to scan the inactive domains for metadata.
8. After the job has executed, select a domain record to view its metadata.

Domain Separation Center

Audit your domains regularly to reveal problems.

overview

The Domain Separation Center is a dashboard where you can schedule and configure audits for all of the domains stored in your domain_audit_definition table. The dashboard enables you to review audit results and dig deeper into domain errors and warnings. You can find the Domain Separation Center dashboard at `<ServiceNow-instance-name>/domaincenter`.

The Domain Separation Center provides many audits. You cannot create your own. You can, however, configure how often they run. Audits run on all of the domains stored in the domain_audit_definition table.

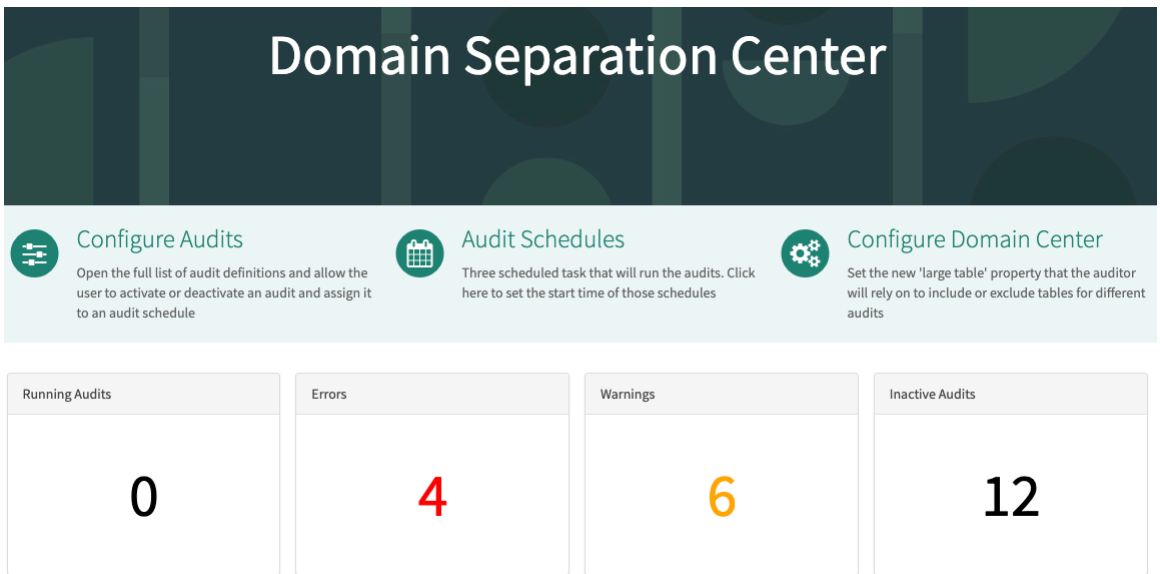

Configure audits

You can configure whether an audit is active and how frequently it runs.

1. In the Domain Separation Center, select **Configure Audits**.

   The **Configure Audits** page displays.

2. Configure each audit you want to be active. Audits are inactive by default.

a. Select an audit.

b. Select the **Active** box to activate the audit.

c. Specify how often the audit runs in the **Frequency** field.

   The Domain Separation Center runs all audits marked with the same frequency at the same time. Do not select **Daily** for audits run on large tables.

d. Repeat these steps for each audit you want to activate.

3. Select **Save**.

Audit schedules

Configure one or more audits to run daily, weekly, or monthly. The scheduler specifies what days and times to run those audits. All audits with the same scheduling frequency run sequentially staring at the time you configure.

1. In the Domain Separation Center, select **Audit Schedules**.

   The **Audit Schedule** page displays.

2. Configure the audit schedules.

   a. Select a schedule name.

   b. Specify the time of day to run the audit.

      Time units are measured with the 24 hour clock, that is, 14 equals 2 PM.

   c. For weekly schedules, select the day of the week to run the audit, where 1 is Sunday.

   d. For monthly schedules, select the day of the month to run the audit.

   e. Repeat the procedure for the other schedules.

3. Select **Save** to save the configuration changes.

4. Select **Execute Now** to run all of the audits scheduled to run at the frequency shown at the top of the right pane, for example, if the pane title is, **Domain Audit Schedule - Daily**, all of the audits that run daily are run.

To see the status of a running audit, in the Domain Separation Center, select the number in the **Running Audits** box.

Execute audits immediately

Audits typically run as scheduled. You can, however run all audits on command.

1. In the Domain Separation Center, select **Audit Schedules**.

   The **Audit Schedule** page displays.

2. Select the name of the audit schedule you want to run, for example, daily, weekly, or monthly.

3. Select **Execute Now**.

To see the status of a running audit, in the Domain Separation Center, select the number in the **Running Audits** box.

Configure the Domain Separation Center

1. In the Domain Separation Center, select **Configure Domain Center**.
2. On the **Configure Domain Center** page, for **Enables detail domain logging**, select **Yes** to store detailed logs that help diagnose domain-related issues. Detailed logging might cause performance issues.

   These logs refer to server-side logs in the syslog_domain table. For information about viewing the logs, see View audits with warnings and failures.
3. In the slushbucket, select and move all of your large tables into the **Selected** column.

   Daily audits should not run on large tables. Grayed-out table names in the **Selected** column are large and cannot be moved to the **Available** column.
4. Select **Update**.

View audits with warnings and failures

On the dashboard, **Errors** and **Warnings** provide detailed information about audits that encountered problems. Logs related to domain separation reside on your server in the syslog_domain table.

1. In the Domain Separation Center, select the number in the **Errors** or **Warnings** box.

   The **Errors** or **Warnings** page displays, respectively.
2. Select one of the audits in the list.

   The page displays detailed information about problems with the audit you selected. The messages in **Audit Result Details** refer to values in the syslog_domain table for table audits that revealed errors or warnings.
3. To see the logs on the server that provide information about warning or error audits:
   a. Copy the **Detail ID** value in the left panel of a warning or error.
   b. In the Filter Navigator of HI, enter *syslog_domain.do*.

      The Domain Log page displays.
   c. In the **Source** search field, enter `=<Detail-ID>` (no space after the equals sign), for example, =f6a00fd29a85b300a9503a81b9169678.

      The Domain Log page displays only the logs related to the audit with the Detail ID you specified. Each row in the table specifies a different record the audit found problems in. Notice that the **Message** field on this page matches the values displayed in the Domain Service Center's **Message** column. The format of the message matches the audit type.
4. Select:
   - **Rerun Audit**—Rerun an audit to see if it still encounters a warning or error.
   - **Deactivate Audit**—Inactivate an audit.
   - **Copy Details**—Copy audit details to the clipboard.

View running and pending audits

Active audits periodically run or are queued to run, as scheduled. You can view their status as they run.

1. Select the number in the **Running Audits** box to view audits that are currently running or pending running.

   The **Running Audits** page displays.

2. Select an audit for more information about it.

View inactive audits

1. Select the number in the **Inactive Audits** box.

   The **Inactive Audits** page displays all of the audits that are currently deactivated.

2. Select one of the audits to display more information about it.
3. To activate an audit, select the **Active** box and use the **Frequency** field to specify how often the audit runs.
4. Select **Update**.

Configure the Domain Separation Center

Specify which tables in domains are large and whether you want detailed logging.

Before you begin

Role required: admin

About this task

You can schedule audits to run on a daily, weekly, or monthly basis. Audits run on large tables can take a significant amount of time. For that reason, avoid running audits daily on large tables. Starting a new audit on one that is taking more than a day to run can have negative consequences.

Detailed logging can provide greater insights into problems found during audits, however might cause performance issues.

Procedure

1. In the Domain Separation Center, select **Configure Domain Center**.
   The Configure Domain Center page appears.
2. Select **Yes** for **Enables detail domain logging**, to store detailed logs that help diagnose domain-related issues.
   These logs refer to server-side logs in the Domain Log [syslog_domain] table. For information about viewing the logs, see the View audits with warnings and failures section.
3. In the list, move all of your large tables into the **Selected** column.
   Daily audits should not run on large tables. Grayed-out table names in the **Selected** column are large and cannot be moved to the **Available** column.
4. Select **Update**.
5. If you have large tables that should never be audited, set the `com.glide.domain.audit.big_tables.additional` system property to a comma-separated list of those table names.

Configure audits

Configure whether an audit is active and how frequently it runs.

Before you begin
Role required: admin

Procedure

1. In the Domain Separation Center, click **Configure Audits**.
   T appears.
2. On the Configure Audits page, configure each audit you want to be active.
   Audits are inactive by default.
   a. Click an audit.
   b. Click the **Active** check box to activate the audit.
   c. Specify how often the audit runs in the **Frequency** field.
      The Domain Separation Center runs all audits marked with the same frequency at the same time. Do not select **Daily** for audits run on large tables.
   d. Repeat steps a-c for each audit you want to activate.
3. Click **Save**.

Schedule audits

Specify the time and day that audits are run.

Before you begin
Role required: admin

About this task
You configured one or more audits to run daily, weekly, or monthly. The scheduler specifies what days and times to run those audits. All audits with the same scheduling frequency run sequentially starting at the time you configure.

Procedure

1. In the Domain Separation Center, click **Audit Schedules**.
   The Audit Schedule page appears.
2. Configure the audit schedules.
   a. Click a schedule name.
   b. Specify the time of day to run the audit.
      Time units use the 24-hour clock (that is, 14 equals 2 PM).
   c. For weekly schedules, select the day of the week to run the audit, where 1 is Sunday.
   d. For monthly schedules, select the day of the month to run the audit.
   e. Repeat the procedure for the other schedules.
3. Click **Save** to save the configuration changes.

4. **Optional:** Click **Execute Now** to run all the audits scheduled to run at the frequency shown at the top of the right pane.
   For example, if the pane title is **Domain Audit Schedule - Daily**, **Execute Now** immediately runs all the audits that are scheduled to run daily.

What to do next
To see the status of a running audit, in the Domain Separation Center, click the number in the **Running Audits** box.

Execute audits immediately

Audits typically run as scheduled. You can, however, run all audits on command.

Before you begin
Role required: admin

About this task
You cannot run individual audits manually. You can, however, run all audits that are configured with the same scheduling frequency. For example, you can run all audits that are configured to run daily.

Procedure

1. In the Domain Separation Center, click **Audit Schedules**.
2. Click the name of the audit schedule you want to run (for example, daily, weekly, or monthly).
3. Click **Execute Now**.

What to do next
To see the status of a running audit, in the Domain Separation Center, click the number in the **Running Audits** box.

View audits with warnings and errors

The Domain Separation Center provides details about audit errors and warnings.

Before you begin
Role required: admin

About this task
On the Domain Separation Center dashboard, **Errors** and **Warnings** provide detailed information about audits that encountered problems. Logs related to domain separation reside on your server in the syslog_domain table. Errors are issues that require immediate attention. Warnings do not lead to failures but present best practices, for example, not making the domain name too long.

Procedure

1. In the Domain Separation Center, click the number in the **Errors** or **Warnings** box.
   The **Errors** or **Warnings** page appears.
2. Click one of the audits in the list.
   The page displays detailed information about problems with the audit you selected. The messages in **Audit Result Details** refer to values in the syslog_domain table on the server for table audits that revealed errors or warnings.
3. To see the logs on the server that provide information about warning or error audits:

a. Copy the **Detail ID** value in the left panel of a warning or error.

b. On the instance running the Domain Separation Center, in the **Filter navigator**, enter *syslog_domain.list*.
The Domain Log page appears.

c. In the search field of the **Source** column, enter `=<Detail-ID>` (no space after the equals sign), for example, `=f6a00fd29a85b300a9503a81b9169678`.
The Domain Log page displays only the logs related to the audit with the Detail ID you specified. Each row in the table specifies a different record the audit found problems in. Notice that the **Message** field on this page matches the values displayed in the **Message** column of the Domain Service Center. The information included in the message matches the audit type.

4. **Optional:** Select one of the following:
   - **Rerun Audit** — Rerun an audit to see if it still encounters a warning or error.
   - **Deactivate Audit** — Inactivate an audit.
   - **Copy Details** — Copy audit details to the clipboard.

View running and pending results

You can view running and pending audits to see their status.

Before you begin
Role required: admin

About this task
Active audits periodically run or are queued to run, as scheduled. You can view their status as they run.

Procedure

1. Click the number in the **Running Audits** box to view audits that are currently running or pending running.
   The Running Audits page appears.
2. Click a running or pending audit to see information about it.

View inactive audits

You can view all inactive audits in one place and optionally activate them.

Before you begin
Role required: admin

Procedure

1. Click the number in the **Inactive Audits** box.
   The Inactive Audits page displays all the audits that are currently deactivated.
2. Click one of the audits to display more information about it.
3. **Optional:** Activate an audit by clicking the **Active** box and specifying how often the audit runs in the **Frequency** field.
4. Click **Update**.

servicenow

Authentication

ServiceNow's authentication validates the identity of a user who accesses an instance, and then authorizes the user to features that match the user's role or job function.

Get started

Authentication factors

Secure your voice agent access by choosing from flexible authentication options such as app-based codes, push prompts, SMS, SoftPIN or knowledge-based checks.

Multi-Provider single sign-on (SSO)

Username and password configured in identity providers, which have a matching user account in the database.

OAuth inbound and outbound

OAuth based authentication validates the identity of the client that attempts to establish a trust on the system by using an authentication protocol.

API access policy

API access policy defines the permissions and access to an API that can be controlled through a policy.

Multi-factor authentication (MFA)

MFA enables you to provide second level of authentication that includes using passcode from an authentication app, hardware key, biometric authenticator, SMS, or Email.

Time limited authentication

Certificate based authentication

Configure link based authentication on the ServiceNow instance. The configured link can be shared with the user through Email or SMS and user can use those links to login to instance.

Unique PEM encoded certificates mapped to users instead of user name and password for certificate based authentication.

LDAP

Integrate with a Lightweight Directory Access Protocol (LDAP) directory to streamline the user login process

Digest token authentication

Username and the secret in the table, perform an hash operation that is user-specific such as SHA1, SHA 256, or MD5. This value has to be appended as part of the URL suffix, which works on the query param.

Self-registration

Use external user self-registration to on-board a large volume of external users to your instance.

You can use several different methods to authenticate users. User credentials are matched to different saved credentials for each method.

Note:

- The Okta SSO plugin is deprecated.
- To learn more about the security properties that affect authorization processing, see Access control in Instance Security Hardening Settings.
- You can use SAML and Digest Authentication through the Multiple Provider SSO application.

Adaptive authentication

Use the Adaptive authentication policy framework to enforce contextual authentication controls to the right users at the right time. Adaptive authentication uses authentication policies to evaluate authentication requests and then either deny or allow access to your instance based on the specified policy conditions.

Use adaptive authentication policies and contexts to restrict the access to your instance for users and APIs based on criteria like IP address, user role, and user group. You can configure the built-in authentication policies according to your security requirements.

For example, an administrator can configure the *`Allow Access Policy`* to allow logins from users only within a trusted range of IP addresses and who are members of a specific role. When

assigned to the *Post-authentication context*, the access policy denies access from untrusted IP addresses.

To set a custom message in the language of your instance you need to add key, value pair in `sys_ui_message.list` and update the `sys_ui_message` record. When you login with an incorrect password, the custom message in the preferred language is displayed.

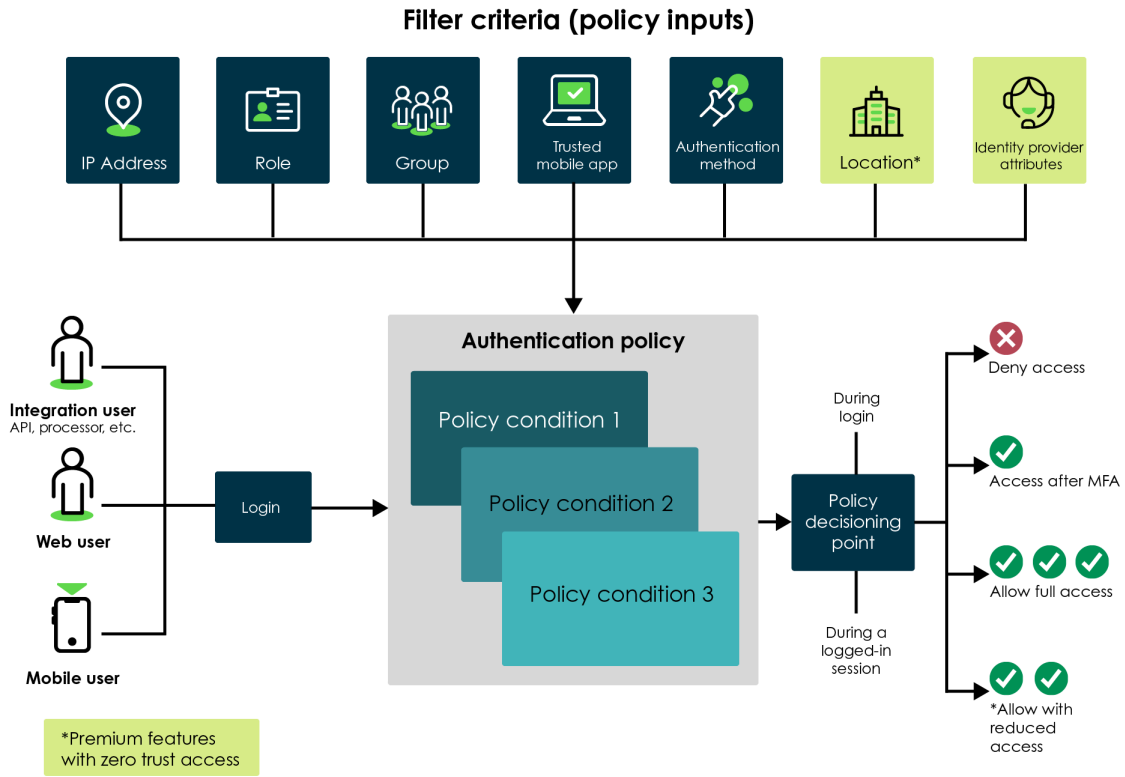

Adaptive authentication components

Authentication policies

Authentication policies evaluate authentication requests based on the specified policy conditions and either allow or deny access depending on the output of policy conditions evaluation. For example, access is allowed only if all the policy conditions specified in **Allow Access Policy** evaluate to true.

Authentication policies use information provided by filter criteria to compare against the policy's conditions to determine whether to grant access to the instance. For example, a filter criteria provides a user's IP address, and a policy condition determines whether this address is within the specific range before granting access. Learn more about authentication policies in Authentication policies.

Authentication policy contexts

Authentication policy contexts define how and when policies are enforced during the login process. The pre-authentication context executes before the user is shown a login screen. The post-authentication context executes after the user enters their credentials. To use a policy, it must be assigned to a policy context. For details on these contexts, see Authentication policy contexts.

Filter Criteria

Filter criteria (also called policy inputs) are used as inputs for policy conditions. Policy conditions use these inputs to verify and meet the requirements of authentication requests. These inputs provide information like user role, IP range, and identity provider. For more detail on filter criteria, see Filter criteria.

Authentication properties

Use authentication properties to control whether adaptive authentication is active on your instance. You can also use properties to enabled debugging, and define the messaging users see when access is blocked. For details on these properties, see Configure adaptive authentication properties.

REST API access policies

You can use the filter criteria of adaptive authentication framework to restrict access to inbound ServiceNow REST APIs. For more information, see REST API access policies.

Domain separation and adaptive authentication

Adaptive authentication is supported on domain separated instances on the authentication policy condition level. Policy conditions affect the domain in the records **Domain** [sys_domain] field. Policy conditions in the global domain affect all domains.

Adaptive Authentication Events

You can use the adaptive authentication events table to know about the events that have occurred specific to the adaptive authentication feature. For more information, see Adaptive Authentication Events.

Activate adaptive authentication

You can activate the Adaptive Authentication plugin (com.snc.adaptive_authentication) for Adaptive Authentication if you have the admin role.

Before you begin

Role required: admin.

About this task

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the **Adaptive Authentication** (com.snc.adaptive_authentication) plugin using the filter criteria and search bar.

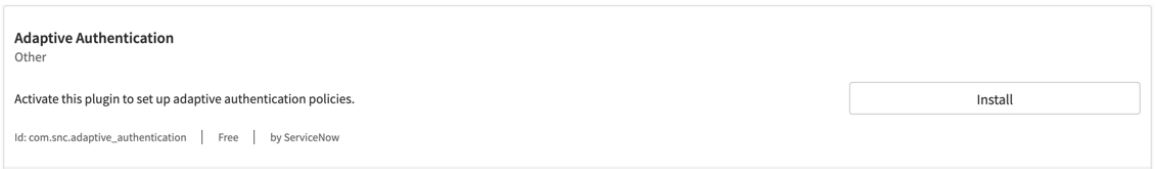

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.

3. Select **Install** to start the installation process.

Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

What to do next

Configure your authentication policies to enforce contextual authentication controls on your instance.

Once your policies are configured. Enable adaptive authentication using the **Enable Authentication Policy** policy. For details on adaptive authentication properties, see Configure adaptive authentication properties.

Filter criteria

Filter criteria (also called policy inputs) are used as inputs for policy conditions to verify and meet the requirements of an authentication request.

Use filter criteria to supply information authentication policies such as a user's IP address, roles, or groups. Add these criteria in the **Policy conditions** section of your policies.

There are seven types of filter criteria used in adaptive authentication. Your authentication policies can use one or more of these criteria to evaluate authentication requests.

Note: Location filter and Identity Provider filter are available with Zero Trust Access feature. For more information, see Zero Trust Access (ZTA).

Filter criteria types

Type	Description
IP filter criteria	Use IP filter criteria to filter users based on the user's IP addresses. Both IPv4 and IPv6 are supported.
Role filter criteria	Use role filter criteria to filter users based on their roles.
Group filter criteria	Use group filter criteria to filter users based on the user group to which the user belongs.
Location filter criteria	Use location filter criteria to filter users based on the user location.
Identity Provider Attribute filter criterias	Use the Identity Provider attributes that are received from SAML response from the IdP as a filter criteria for authentication.

Generic Criteria

In addition to the previously listed types, there are four generic filter criteria. These criteria do not appear in your filter navigator, but you can select them while adding policy inputs to your authentication policies.

Generic filter criteria types

Type	Description
Authentication Scheme	Use to filter based on user's authentication scheme. This criteria is a choice type with two options: • User name and Password, which denotes a local login • SSO, which denotes a Multi-SSO(SAML, OIDC, or Digest) based login. **Note:** This Filter Criteria is available only when the **Integration - Multiple Provider Single Sign-On Installer**[com.snc.integration.sso.multi.installer] plugin is installed.
Identity Provider	Use to filter based on the user's identity provider. Use along with the authentication scheme criteria to have granular control over login process. This criteria is a reference to the **Identity Providers** [sso_properties] table. **Note:** This Filter Criteria is available only when the **Integration - Multiple Provider Single Sign-On Installer**[com.snc.integration.sso.multi.installer] plugin is installed.
Role-based MFA	Use to filter based on the role-based MFA feature. This criteria is a boolean type filter criteria which denotes whether role-based MFA is enabled for the user.
User-based MFA	Use to filter based on the user-based MFA feature. This criteria is a boolean type filter criteria which denotes whether user-based MFA is enabled for the user.
Trusted mobile app	Trusted mobile app filter for enabling instance access from mobile app.

IP Filter

Use IP filter criteria to filter the users based on the user's IP addresses. Both IPv4 and IPv6 are supported.

IP filter criteria allows you to filter users based on the user's IP addresses. You can configure an authentication policy to allow or deny access to a specific address or range of addresses.

Create IP filter criteria

IP filter criteria allows you to filter users based on the user's IP addresses. You can configure an authentication policy to allow or deny access to a specific address or range of addresses.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Filter Criteria** > **IP Filter Criteria**.
2. Click **New**.
3. On the form, fill in these fields.

IP Filter Criteria form

Field	Description
Name	Name to identify the IP network.
Description	Short description of the IP network.
Application	Scope of the application.

Example IP filter criteria record

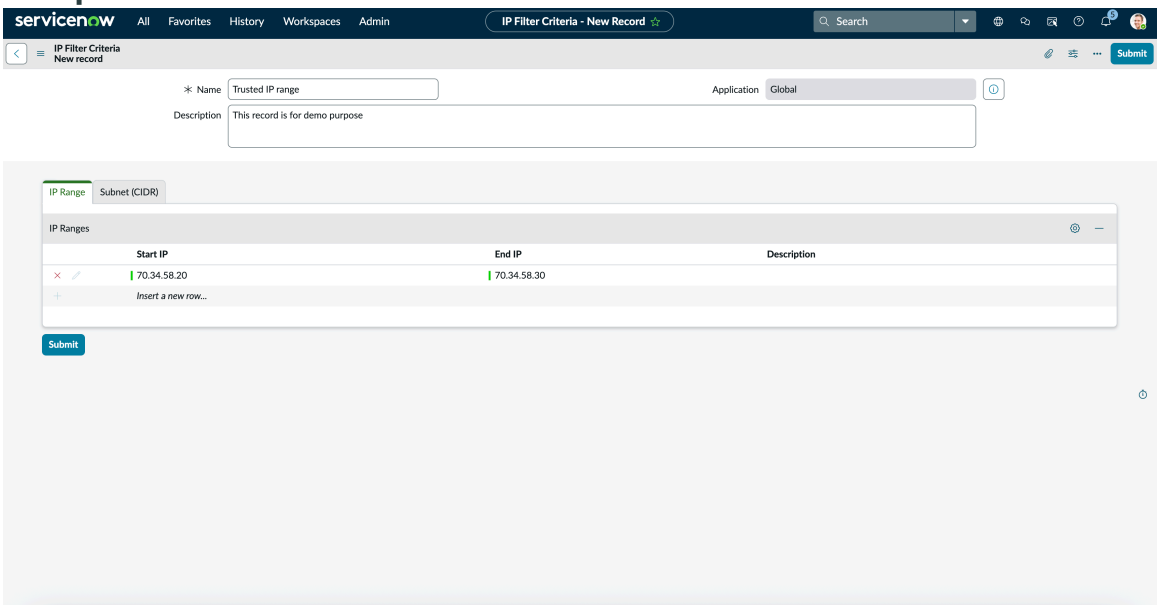

4. Right-click the form header and click **Save**.
5. From the **IP Range** tab, double-click **Insert a new row**.
   You can enter a single IP address or multiple ranges of IP addresses. For example, for a range of IP addresses enter `192 . 0 . 2 . 0` in the **Start IP** column and `192 . 0 . 2 . 255` in the **End IP** column.

   Note: For a single IP addresses, make sure that you enter the same IP address in the **Start IP** address and **End IP** columns. When you enter an IP address in the **Start IP** column, leave the **End IP** column blank and then save the record. The **End IP** column is auto-populated with the same **Start IP** address.

6. From the **Subnet (CIDR)** tab, double-click **Insert a new row**.

Enter the Network IP address and Netmask in the Classless Inter-Domain Routing (CIDR) format. For example, enter 255 . 255 . 255 . 0 as Network IP and 25 as Netmask.

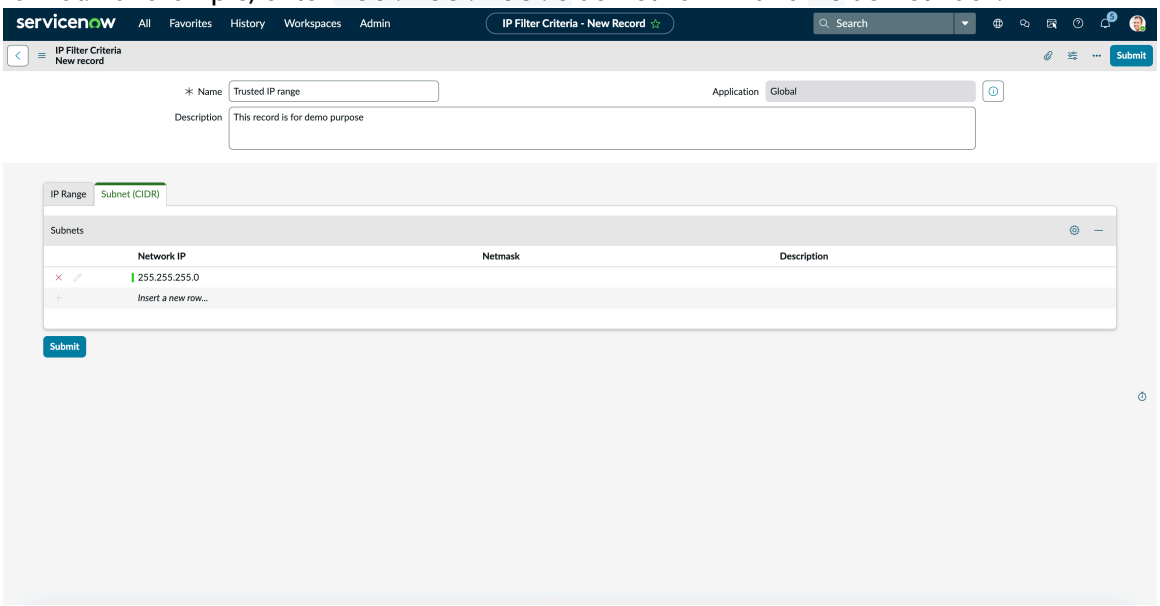

Role Filter

Use role filter criteria to filter users based on their roles.

Role filter criteria allows you to filter users based on the roles. You can configure an authentication policy to allow or deny access to a list of user roles.

Create role filter criteria

Role filter criteria allows you to filter users based on the roles. You can configure an authentication policy to allow or deny access to a list of user roles.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Filter Criteria** > **Role Filter Criteria**.
2. Click **New**.
3. On the form, fill in these fields.

Role Filter Criteria form

Field	Description
Name	Name to identify the role.
Application	Scope of the application.
Description	Short description of the role.

Example role filter criteria record

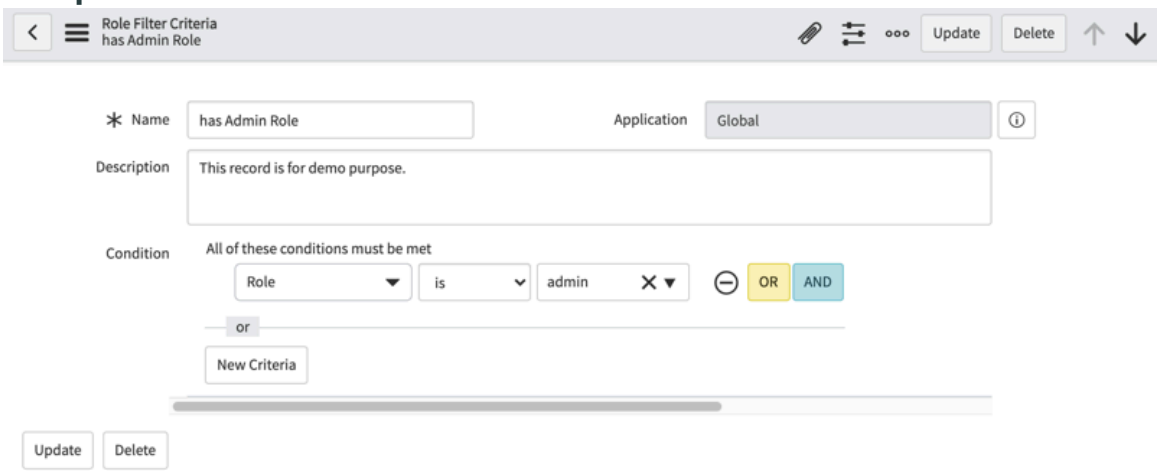

4. From the **Roles for criteria**, double-click **Insert a new row**.
5. Create a condition for a specific role using the Condition Builder.

   For example, you can create a condition that allows only users with admin, itil, or snc_internal roles. For more information about Condition Builder, see Create a condition statement using the condition builder.

   Note:

   - Currently, Dot-walking is not supported in role filter criteria.
   - Following operators are not supported for role filter criteria:
     - Is not
     - Does not contain
     - Is different from
     - Is empty
     - Is same as
     - Is not empty
     - Is empty string

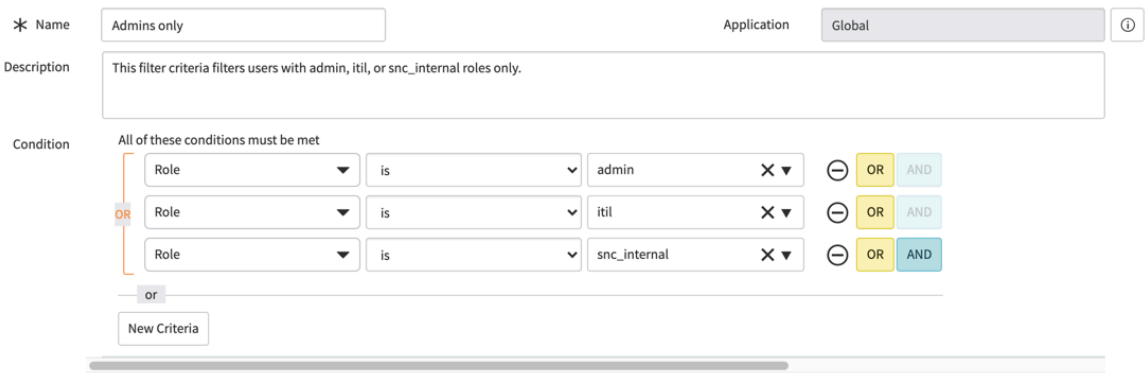

Group Filter

Use group filter criteria to filter users based on the user group to which the user belongs.

Group filter criteria allows or denies user access based on the user group to which the user belongs.

Create group filter criteria

Group filter criteria allows or denies user access based on the user group to which the user belongs.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Filter Criteria** > **Group Filter Criteria**.
2. Click **New**.
3. On the form, fill in these fields.

Group Filter Criteria

Field	Description
Name	Name to identify the group.
Description	Short description of the group.
Application	Scope of the application.

Example group filter criteria record

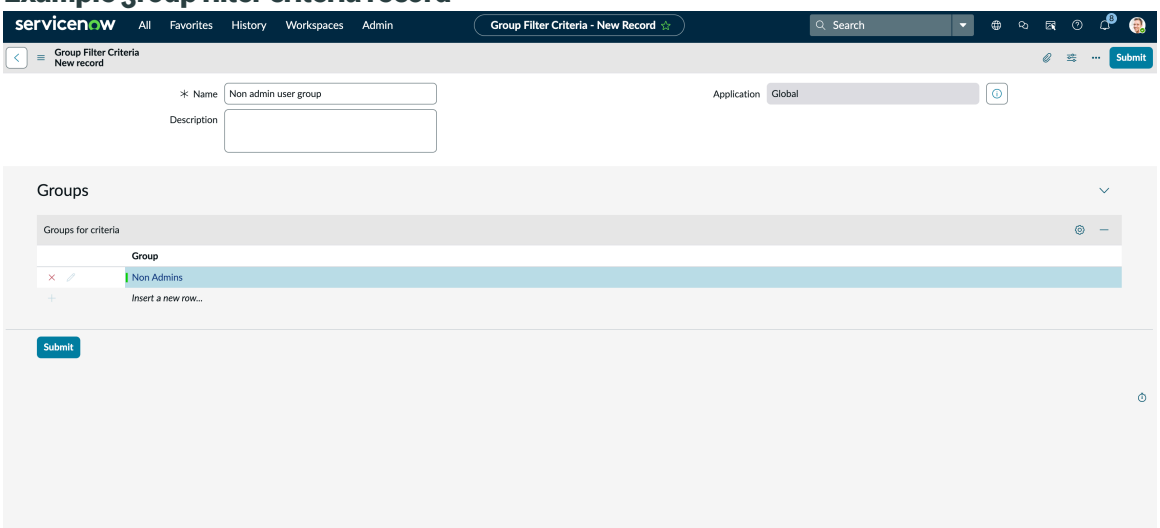

4. From the **Groups for criteria** tab, double-click **Insert a new row**.
5. Click the search icon , look up, and select a user group.
6. Click the save icon.

Location Filter

Location filter criteria can be used as filter input for users based on the user location.

Location filter is a filter criteria that the admins can use while crafting the authentication policies based on the physical location of the device.

Note:

- Location filter is available with Zero Trust Access feature. For more information, see Zero Trust Access (ZTA).
- The instance should be on ADCv2. If the instance isn't on ADCv2, then user location information won't be available.

The location filter criteria enable you to perform the following:

- Act as policy input act for policy conditions to verify and meet the requirements for authentication.
- Provide an ability to create adaptive authentication policies based on the **country**.
- Allow or not allow instance access to a given geography.
- Use the geo location-based pre-authentication or post authentication policy to:
  - Prevent access from sanctioned countries, outside of a tight privacy-controlled region, company discretion.
  - Allow access to areas only within applicable privacy region and company discretion.
  - Configure country-based allow list for authentication.

Use cases

Following are some of the use case for using location filter criteria for Adaptive Authentication:

- Block access to the instance from a country.
- Allow access to the instance only from a particular country.
- Enforce step-up authentication or MFA to log in based on country.
- Reduce or limit roles for the user based on the country.
- Location Filter criteria can be used for MFA, Zero Trust Access (ZTA), Pre-authenication context, and Post authentication context.

Location Identification

The location services for identifying the location of the user are provided by a third-party service - MaxMind.

The location of the user is identified through the VPN, from the x-forwarded-for header. In case if there's no header populated by the service, then only the VPN IP (location) is displayed as the user location.

Note: If there are incorrect locations displayed after the location filter configuration, see the KB article to troubleshoot.

Activate Location Based Access

Activate the **Zero Trust - Location Based Access** (`com.snc.zero_trust_location_access`) to allow admins to configure adaptive authentication policies based on the location of the user.

Before you begin
Role required: admin

- Dependant plugin: Adaptive authentication
- Plugin type: Paid and requires license.
- The instance should be on ADCv2. If the instance is not on ADCv2, then user location information will not be available.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the **Zero Trust - Location Based Access** (`com.snc.zero_trust_location_access`) plugin using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Create location filter criteria

Use location filter criteria to filter input for users authentication based on the user location.

Before you begin

Role required: admin

Plugin required: **Zero Trust - Location Based Access** (`com.snc.zero_trust_location_access`).

Property: Enable the Adaptive authentication property.

Note: Administrators can only create the policy based on location filters if the location is available for the current user session.

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Filter Criteria** > **Location Filter Criteria**.
2. Select **New**.
3. On the form, fill in these fields.

Location Filter Criteria form

Field	Description
Name	Name to identify the criteria.
Description	Short description of the criteria.
Application	Scope of the application.

4. In the **Locations** sections, under the Locations for criteria tab, double-click to **Insert a new row**.
5. Specify the locations.

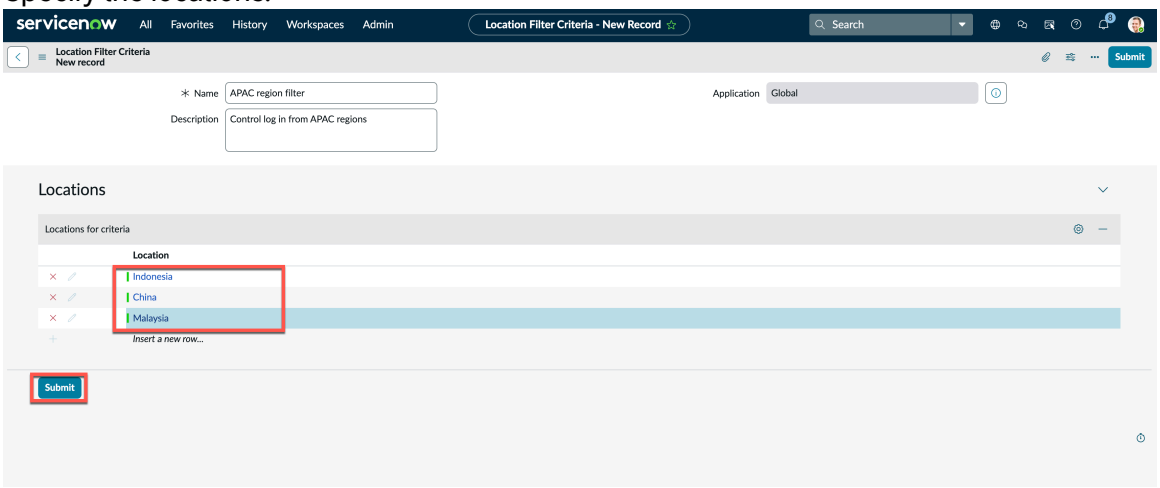

6. Select **Submit**.

Tutorial: Use Location Filter criteria

Describes steps to use location filter criteria in the authentication policy and restrict access to the users based on the location.

Before you begin

Role required: admin

Plugin required: **Zero Trust - Location Based Access** (`com.snc.zero_trust_location_access`).

Property: Enable the Adaptive authentication property.

Note: Administrators can only create the policy based on location filters if the location is available for the current user session.

The following procedure describes how to create and use the location filter criteria in an authentication policy.

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Filter Criteria** > **Location Filter**.
2. Select **New**.
3. On the form, fill in these fields.

Location Filter Criteria form

Field	Description
Name	Name to identify the criteria.
Description	Short description of the criteria.
Application	Scope of the application.

4. In the **Locations** sections, under the Locations for criteria tab, double-click to **Insert a new row**.

5. Specify the locations.

   For example, use some of the APAC regions to control the logins of the users coming from APAC.

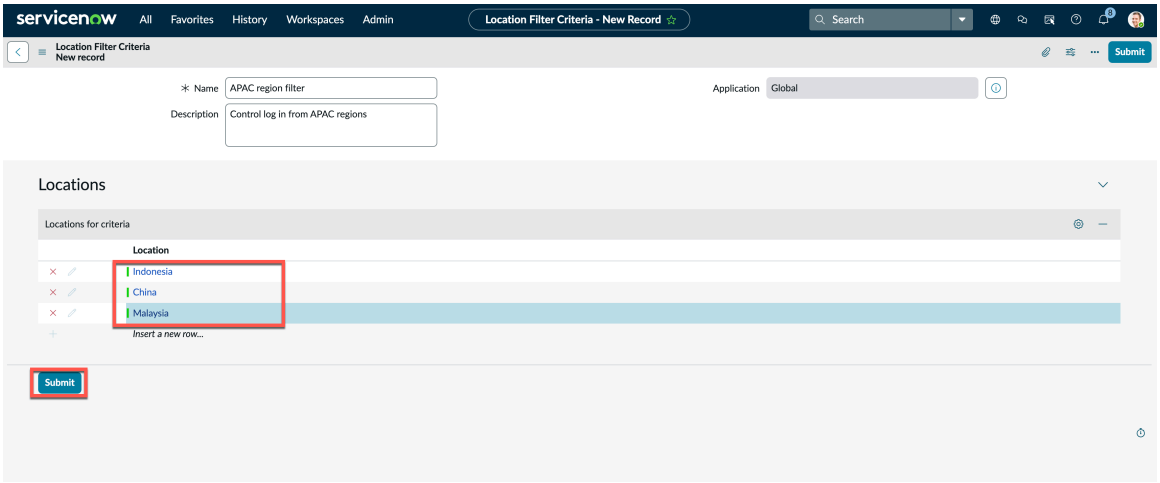

   Based on the criteria set in the example, you can control logins from `Indonesia`, `China`, and `Malaysia` for your instance.

6. Select **Submit**.
7. Use the filter criteria created in any of the authentication context (Pre, Post, MFA) and Session Access.

   To know more about the configuration based on authentication context and session access, see:

   - Location Filter in Pre Authentication Context
   - Location Filter in Post Authentication Context
   - Location Filter in MFA Context
   - Location Filter for Session Access

   You can use the Property ID - Error message to be displayed to the user when login fails due to authentication policy failure (`glide.auth.policy.ui.error.message`) to customize the error message.

Use Location Filter in Pre Authentication Context

Use the location filter criteria created in the Pre Authentication Context.

Before you begin

Role required: admin

Plugin required: **Zero Trust - Location Based Access** (`com.snc.zero_trust_location_access`).

Create a Location Filter with the countries that you want restrict access to the users based on the location. For more information, see Create location filter criteria.

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Auth Policy Context** > **Pre Authentication Context**.

   When a policy is chosen in the pre authentication policy context:

   - Selecting the Deny access policy as the default policy allows the access to all users by default and only denies access when the policy conditions defined in the deny access policy evaluate to true.
   - Selecting the Allow access policy as the default policy denies the access to all users by default and only allows access when the policy conditions defined in the allow access policy evaluates to true.

   The example shows how you can restrict the logins from the specified locations. You can choose the Deny Access and the associated policy (Deny Access) as the Pre Authentication policy and specify the policy inputs and conditions.

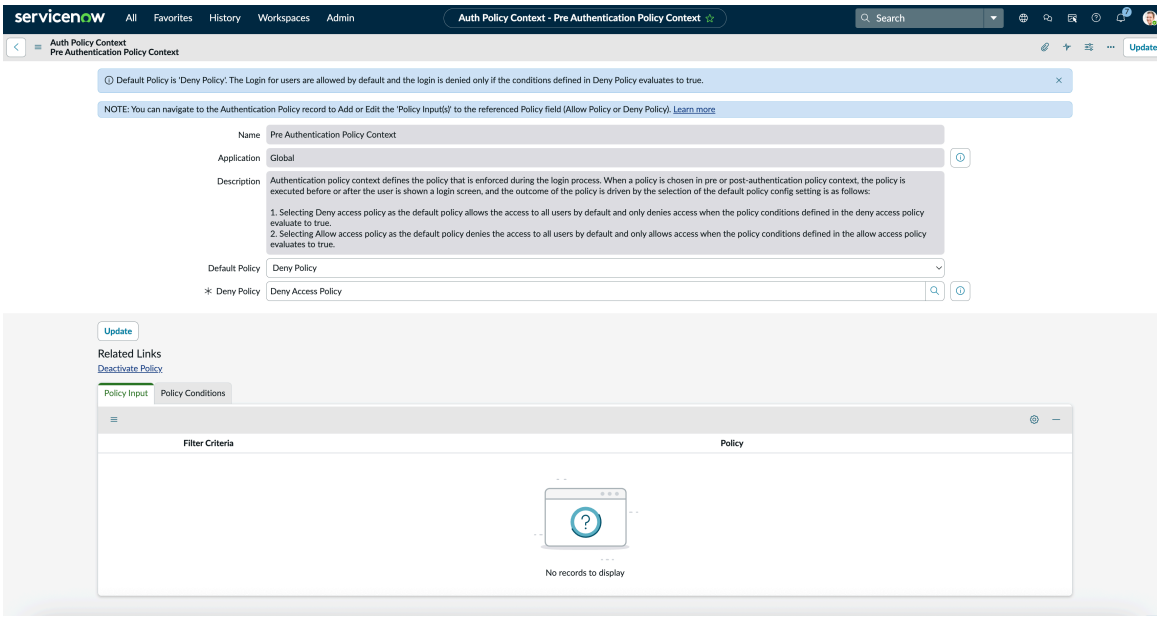

2. Select the information icon and then select **Open Record** to open the **Deny Policy** record.

   Note: The example described in this task is **Deny Policy**. You can also use **Allow Policy** and set the conditions accordingly to control the logins.

3. In the Deny Access Policy, under the Policy Inputs section, select **New**.

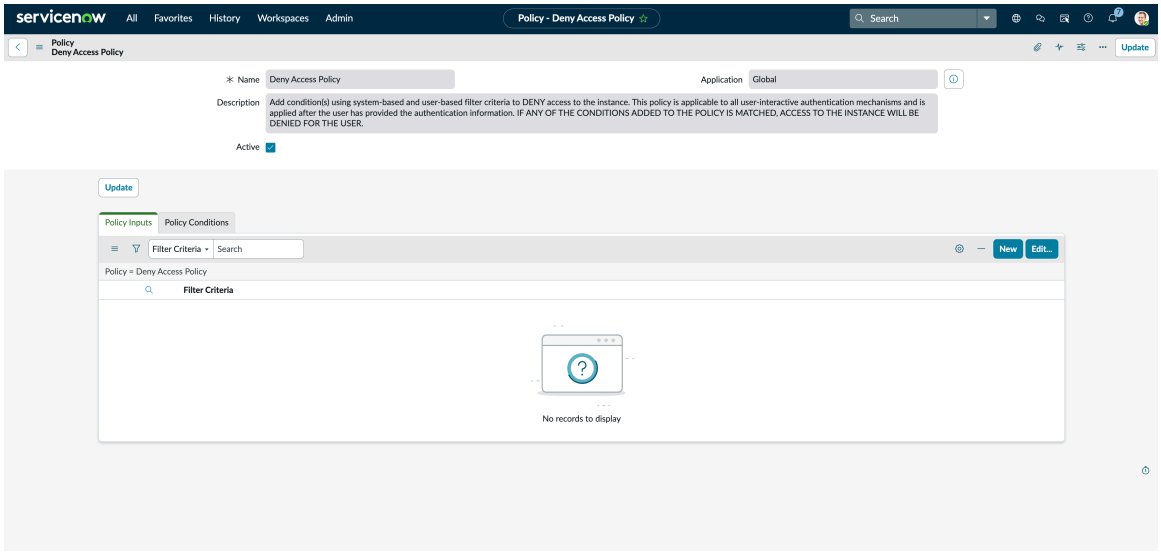

4. Add the Location Filter input and **Save**.

   For example, APAC Region.

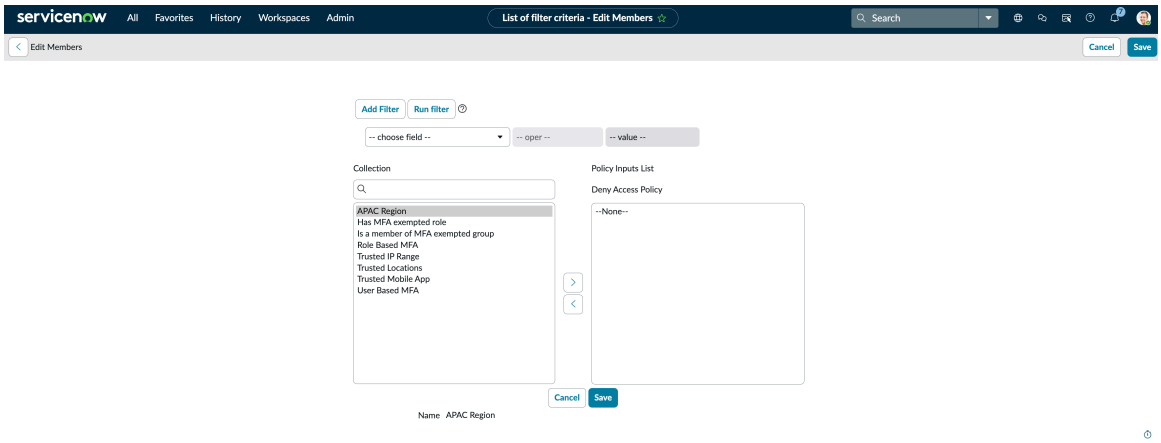

   The filter is added as **Policy Inputs**.

5. Select the Policy Conditions tab and select **New**.
6. In the Conditions page, provide the label, conditions, and set it to true.

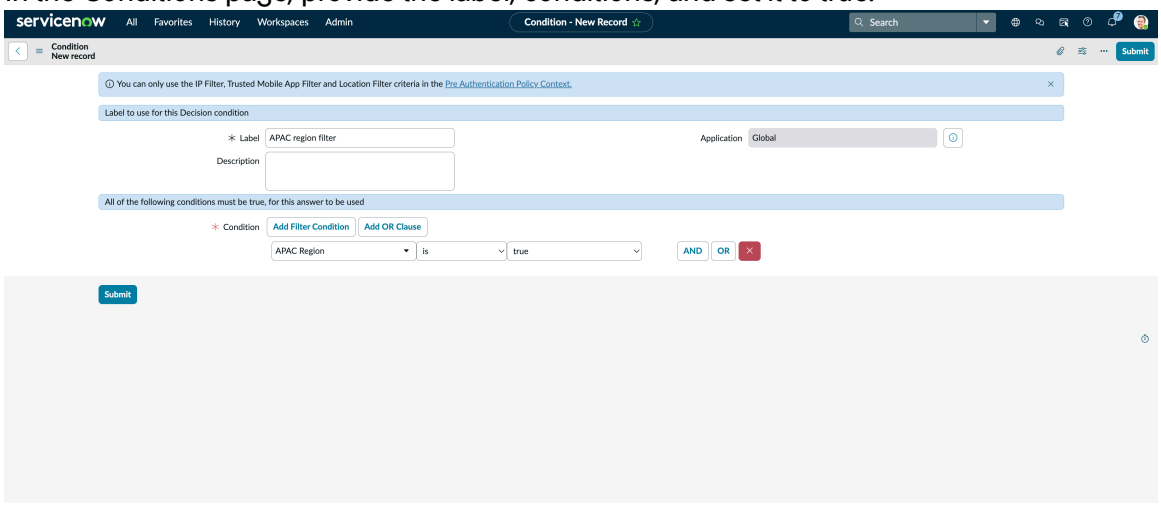

Note:

- In this example, the selection of true as a condition implies that the user logging from the configured regions won't be able to log in to the instance.
- If the condition is set to false, then only users from configured regions are able to log in to the instance and the other users won't be able to log in to the instance.

7. Select **Submit**.
   The users selecting the instance link and logging from the configured countries will be displayed an error message about the access denial (error message configured by their administrators on the policy properties page).

Use Location Filter Post Authentication Context

Use the location filter criteria created in the Post Authentication Context.

Before you begin

Role required: admin

Plugin required: **Zero Trust - Location Based Access** (`com.snc.zero_trust_location_access`).

Create a Location Filter with the countries that you want restrict access to the users based on the location. For more information, see Create location filter criteria.

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Auth Policy Context** > **Post Authentication Context**.

   When a policy is chosen in the post authentication policy context:

   - Selecting the Deny access policy as the default policy allows the access to all users by default and only denies access when the policy conditions defined in the deny access policy evaluate to true.
   - Selecting the Allow access policy as the default policy denies the access to all users by default and only allows access when the policy conditions defined in the allow access policy evaluates to true.

   The example shows on how you can configure a `itil` users with a condition to only access the instance from the specified location (US). And the `itil` users can't log in from other countries.

   After users providing their credentials on the login page. You can choose the Allow Access and the associated policy (Allow Access) as the Post Authentication policy and specify the policy inputs and conditions.

2. Select the information icon and then select **Open Record** to open the **Allow Policy** record.

   Note: The example described in this task is **Allow Policy**. You can also use **Deny Policy** and set the conditions accordingly to control the log in.

3. In the Allow Access Policy, under the Policy Inputs section, select **New**.

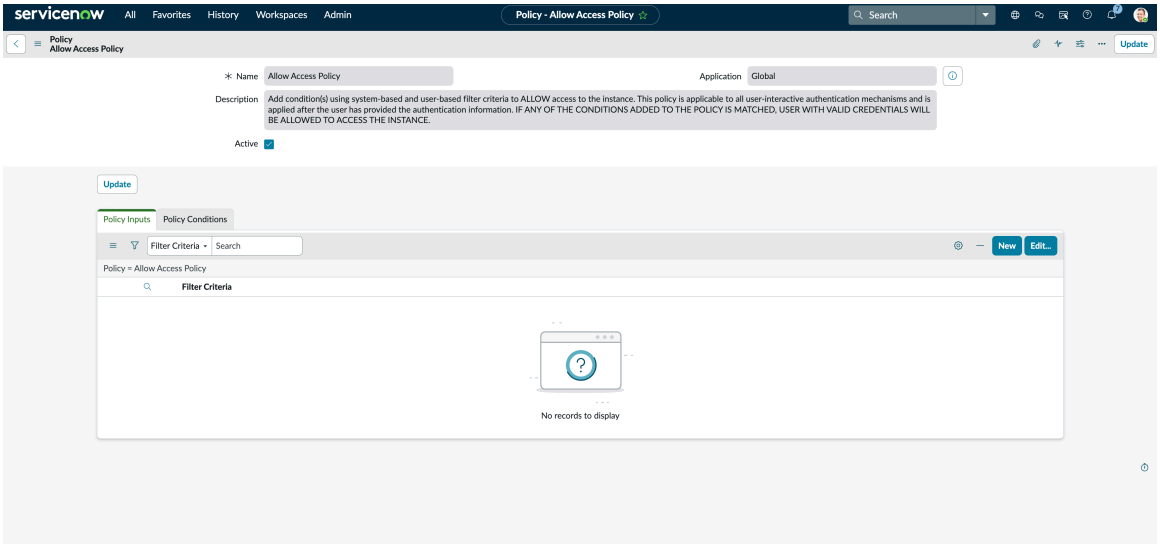

4. Add the Role Filter criteria and Location Filter criteria.

a. Adding Role Filter criteria:

- Create Role Filter Input as

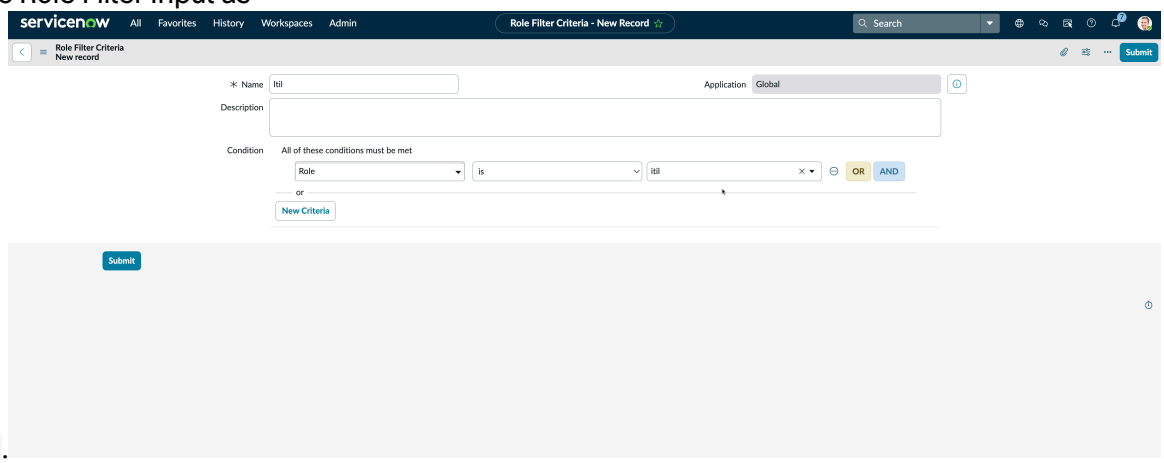

`itil`.

- Create Role Filter Condition and set it to

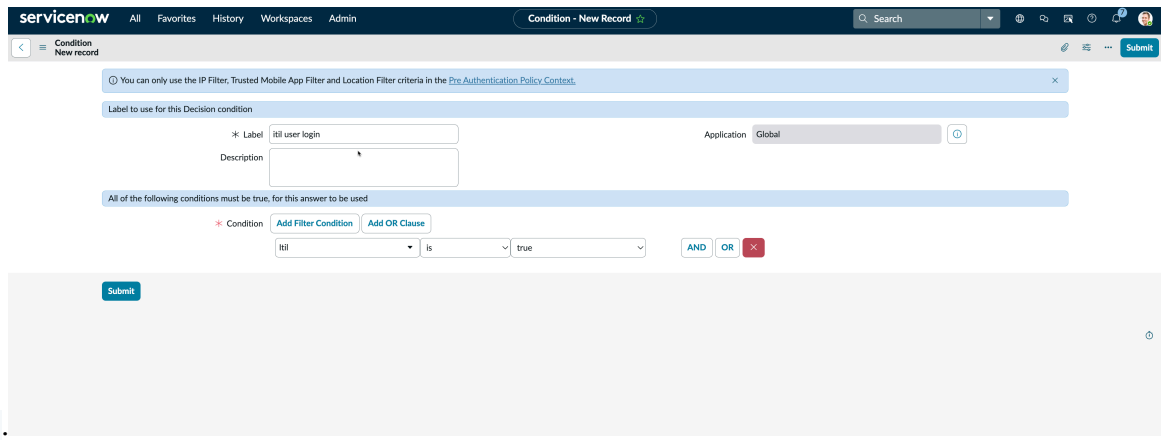

`true`.

For more information on how to create role filter criteria, see Create role filter criteria.

b. Adding Location Filter criteria:

- Create Location Filter Input. Add United States in the Locations.

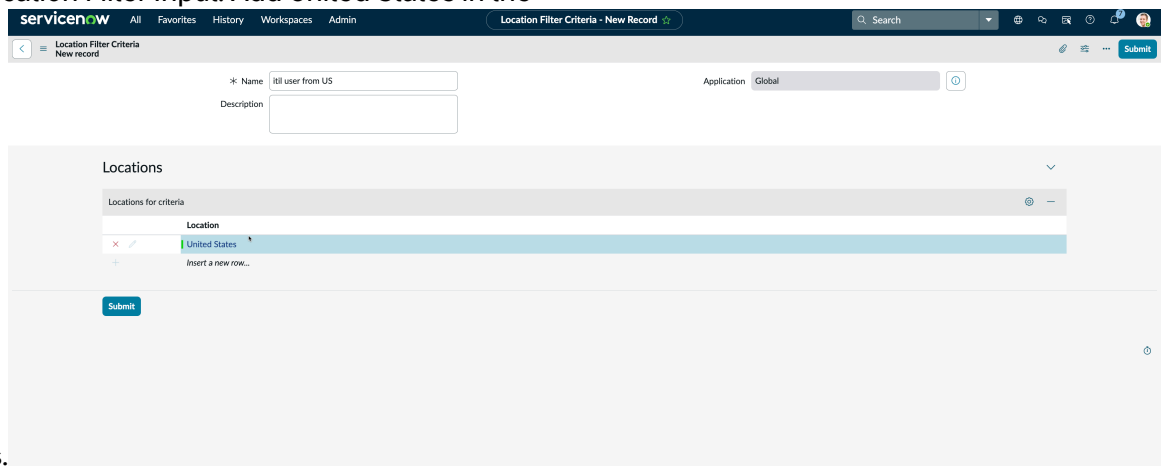

- Create Location Filter Condition and set it to `true`.

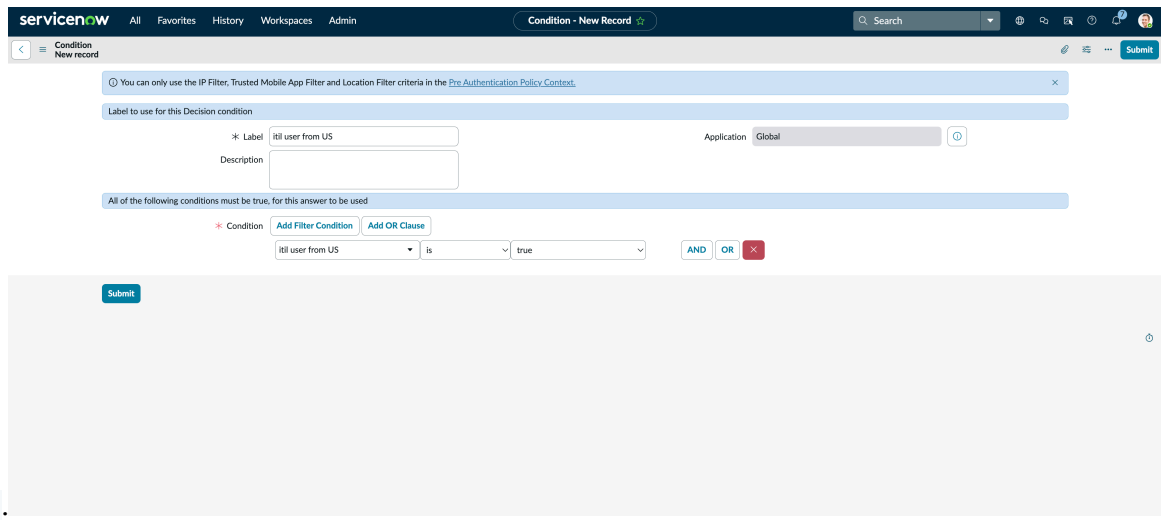

For more information on how to create role filter criteria, see Create location filter criteria.

The Allow Access Policy shows the Policy Inputs and Conditions that are created in the previous steps:

- Policy Inputs: itil, itil user from US
- Policy Conditions: itil user login, itil user from US condition.

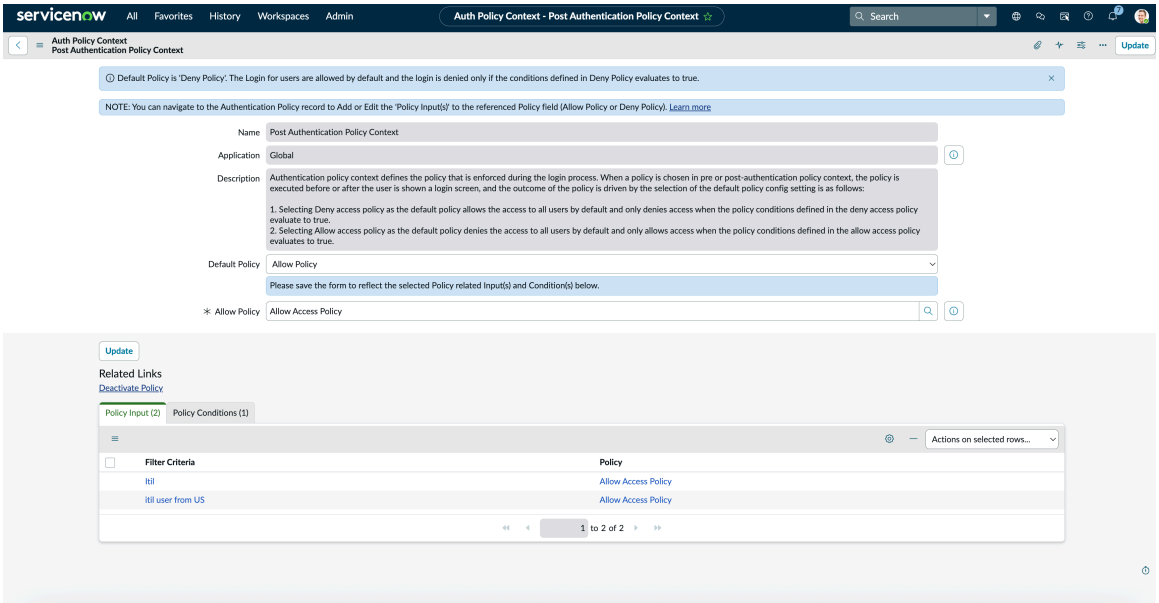

Note:

- In this example, the selection of true as a condition implies that the `itil` users logging from the configured country (US) are able to log in to the instance.
- If the condition is set to false, then the `itil` users from the configured country (US) won't be able to log in to the instance and the other users won't be able to log in to the instance.

5. Select **Submit** or **Update** to update the Post Authentication Context.
   The `itil` users selecting the instance link, specifying their credential and then logging outside the configured country (US) is displayed an error message about the access denial (error message configured by their administrators on the policy properties page).

Use Location Filter in MFA Context

Use the location filter criteria created in MFA Context.

Before you begin

Role required: admin

Plugin required: **Zero Trust - Location Based Access** (`com.snc.zero_trust_location_access`).

The following procedure describes on how to create a Location Filter with the countries that you want to configure MFA as a second factor for authentication to the users based on the location.

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Auth Policy Context** > **MFA Context**.
2. Select the **Step-Up MFA Policy** information icon and then select **Open Record** to open the **MFA Context**.
3. On the Step-Up MFA Policy page, under the Policy Inputs tab, select **New**.
4. Add the Location Filter input and **Submit**.

   For example, you want to display MFA for users logging in to the instance outside Australia.

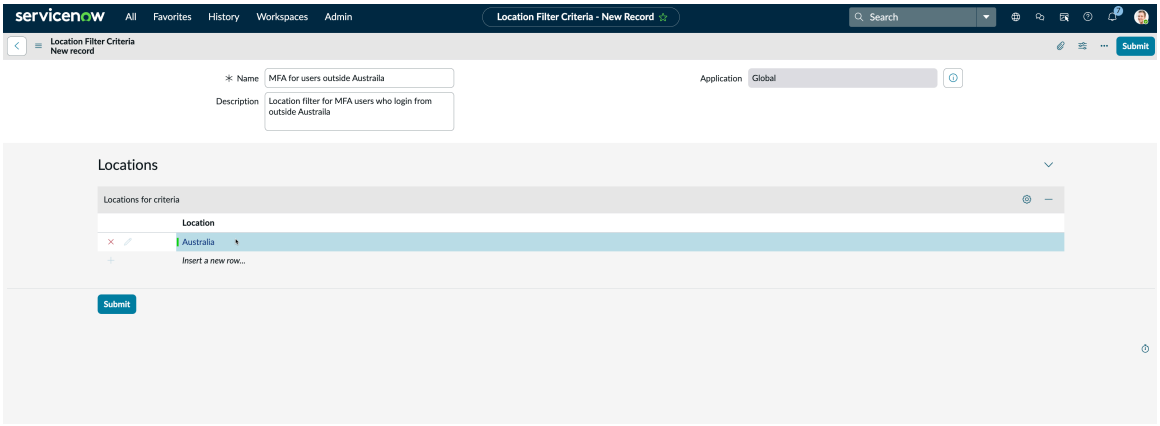

5. On the Step-Up MFA Policy page, select the Policy Conditions tab and select **New**.
6. In the Conditions page, provide the label, conditions, and set it to true.

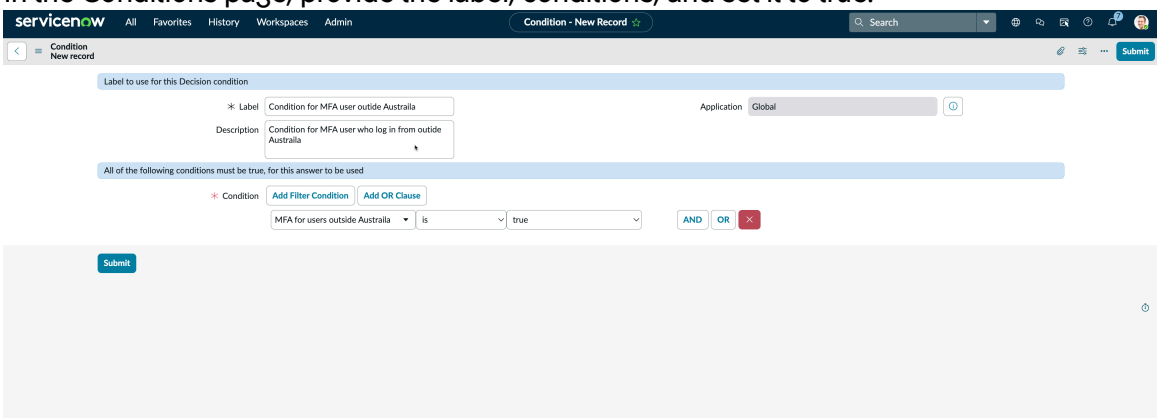

7. Select **Submit**.
8. On the Step-Up MFA Policy page, activate the MFA Policy if it's **Deactivated**.
9. Select **Update** to update the configuration.

   The users outside the configured country (Australia) selecting the instance link, specifying their credential; will be shown with the MFA screen to provide the second factor credentials to log in to the instance.

Use Location Filter for Session Access

Use the location filter criteria created in Session Access to reduce roles based on the location of the user.

Before you begin

Role required: admin

Plugin required: **Zero Trust - Location Based Access** (`com.snc.zero_trust_location_access`).

The following procedure describes on how to create a Location Filter with the countries that you want to remove or limit roles to the users based on the location.

Procedure

1. Navigate to **All** > **Zero Trust Access** > **Session Access Role Configurations**.
2. To create a session access role configuration, select **New**.
3. On the form, fill the fields:

Session Access Role Configuration

Field	Description
Name	Name of the configuration
Description	Short description of the configuration.
Policy	Choose the access policy. Use the look-up icon to view the list of policy. **Note:** You must add the location filter input and conditions by opening the policy record.
Action	**Remove Roles** or **Limit to Roles**. ○ **Remove Roles**: When the configured user logged in, the list of roles provided in the Role or Group List are removed for the logged-in session. ○ **Limit To Roles**: When the configured user logged in only the selected roles is provided to the user and all the other roles are removed for the logged-in session.
Role List	Choose the role from the Role List.
Group List	Choose the role from the Group List.

4. Select, **Submit**.

   The login for users based on the configured countries is as follows:

   - If **Remove Roles**, the users from the configured countries in the location filter will be removed with the roles configured for the session.
   - If **Limit To Roles**, the users from the configured countries in the location filter has only the roles that are configured for the session.

   To know more about how to remove or limit roles for a session, see Tutorial: Use Zero Trust Access.

Identity Provider Attributes Filter

Use the Identity Provider attributes that are received from the Security Assertion Markup Language (SAML) response and OpenID Connect (OIDC) from the Identity Provider (IdP) as a filter criteria for authentication.

To fetch all the attributes from an IdP through the SAML and OIDC response, you should perform a test connection with the IdP. After a successful test connection, the attributes are added in a new tab in the Identity Provider configuration page.

For more information see the following topics:

- Identity Provider attributes for Security Assertion Markup Language
- Identity Provider attributes for OpenID Connect

Identity Provider attributes for Security Assertion Markup Language

Use the Identity Provider attributes that are received from the Security Assertion Markup Language (SAML) response and OpenID Connect (OIDC) from the Identity Provider (IdP) as a filter criteria for authentication.

To fetch all the attributes from an IdP through the SAML response, you should perform a test connection with the IdP. After a successful test connection, the attributes are added in a new tab in the Identity Provider configuration page.

- Identity Provider filter is available with Zero Trust Access feature. For more information, see Zero Trust Access (ZTA).
- IdP attribute filter criteria can be used in Post-authentication context, Zero Trust Access (ZTA) session relegation, and Multi-factor Authentication context.

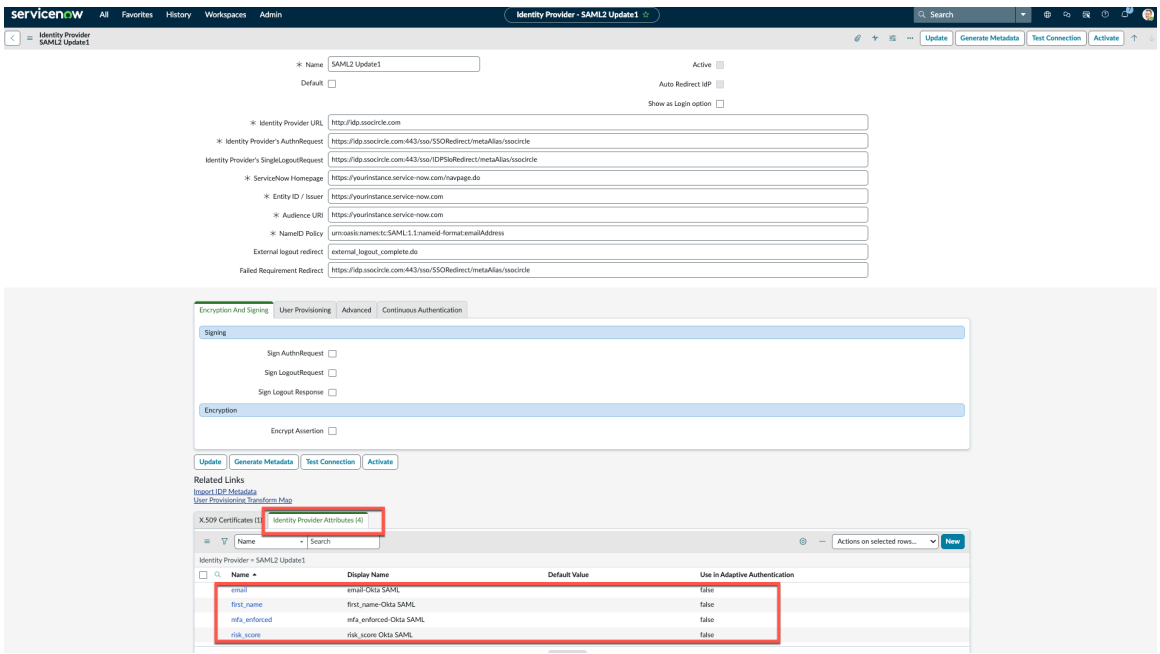

You can also add attributes by selecting **New** from the Identity Provider Attributes section and use those attributes for Adaptive Authentication by setting it to `true`.

The **Identity Provider Attributes** are displayed with the following details:

Location Filter Criteria form

Field	Description
Name	Attribute name that is provided by the Identity Provider.
Display Name	Display Name is the detailed name that is used for the filter criteria. **Note:** You can provide a readable name as a Display Name, in some cases the Display Name provided by the Identity Providers are lengthy and not readable.

Location Filter Criteria form (continued)

Field	Description
Default Value	Default value is used for filter criteria evaluation in case the attribute is missing in the SAML response.
Use in Adaptive Authentication	Option to use the Attribute in the Adaptive Authentication.

Note: Attributes that are populated from Azure IdP have name and display name limited to characters, due to the name length of the attribute.

You can also add new attributes by selecting **New** in the **Identity Providers Attributes** section.

If the Use in Adaptive Authentication is set to true, then the selected attribute is added as filter criteria in the Generic Filter Criteria. For example, **risk_score** set to `true`. The Generic Filter Criteria page has a new filter created.

Use Identity Provider Attribute as Filter Criteria for SAML

Use the Identity Provider (IDP) attribute from the Security Assertion Markup Language (SAML) response as a filter criteria for authentication policy.

Before you begin

Role required: admin

You can create session access policy using policy context (Pre-Authentication, Post Authentication, multi-factor authentication) and filter criteria (Role, Group, IP, Location) with policy inputs and conditions.

The following procedure shows steps to configure the IdP attribute from the SAML response as a policy input to control authentication in the **Post Authentication Context**, **Multi-factor authentication (MFA) Context**, and **Zero Trust - Policy based session access**.

The Okta IDP attributes are as displayed in the following screenshot. You should set the Use in Adaptive Authentication as true to use it in the **Post Authentication Context**, **Multi-factor authentication (MFA) Context**, and **Zero Trust - Policy based session access** policies.

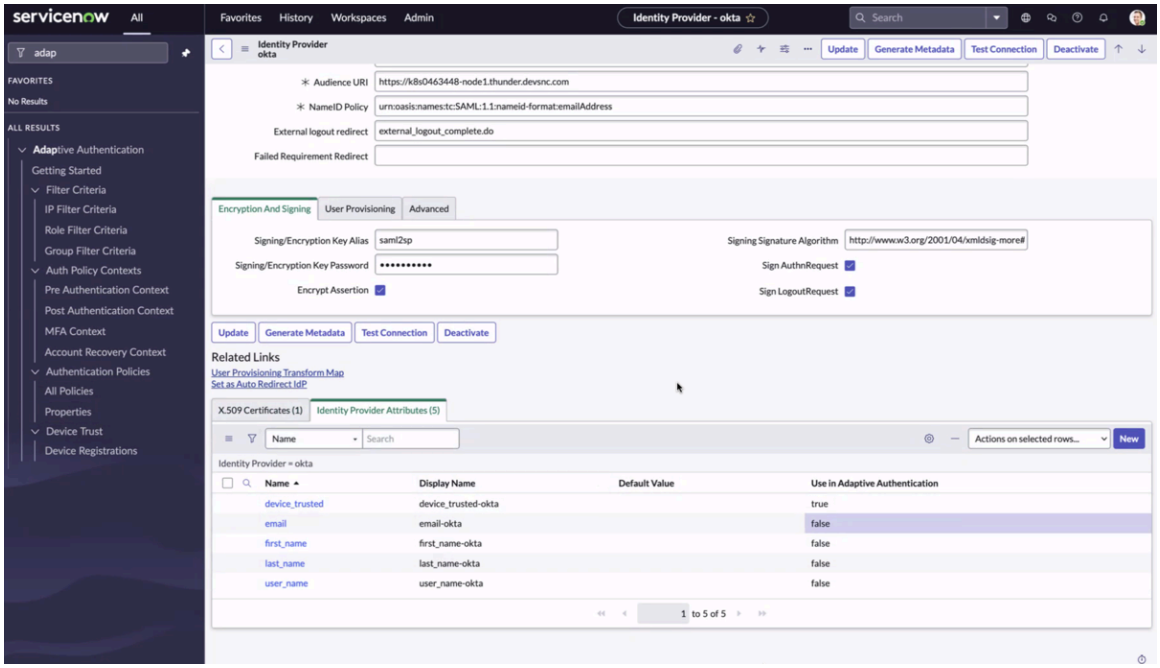

Note: Policies in the post-authorization, MFA, Zero Trust - Policy based session access execute after the users enter the credentials or SSO response.

Procedure

1. Use of IDP attribute in Post Authentication Policy Context.
   Example: Configuring to enable logins from the Okta IDP attributes if the device is trusted.

   a. Navigate to **All** > **Adaptive Authentication** > **Auth Policy Contexts** > **Post Authentication Policy Context.**.

   b. Select **Allow Policy** and open the policy record.

   c. In the Policy Input, create the Policy Input and Policy Condition.

   - **Policy Input**: Add **device_trusted-okta**.

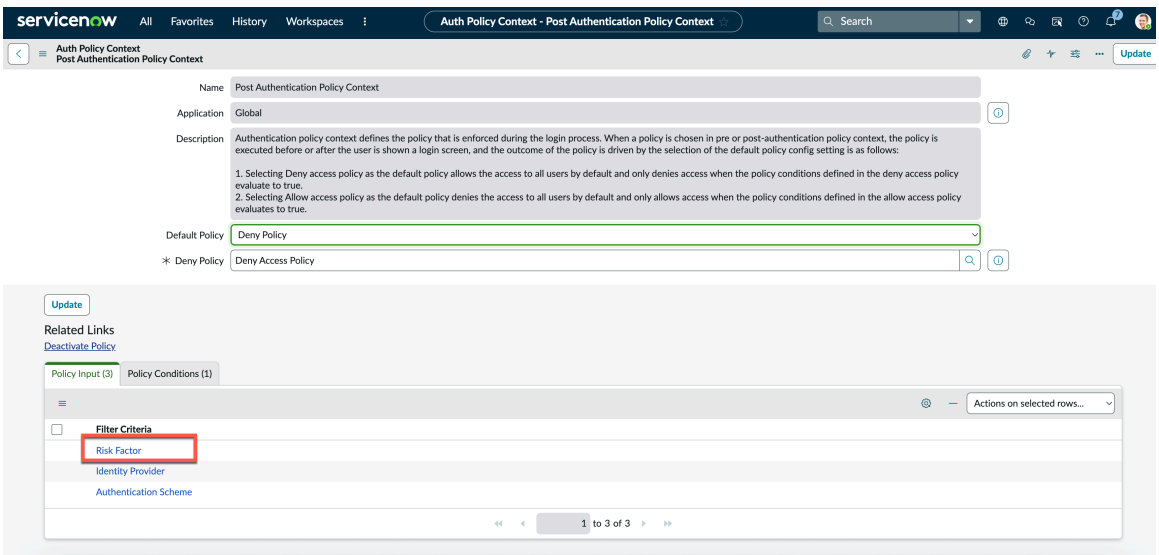

- **Policy Conditions**: **device_trusted-okta** is `trusted` and **Identity Provider** is `okta`.

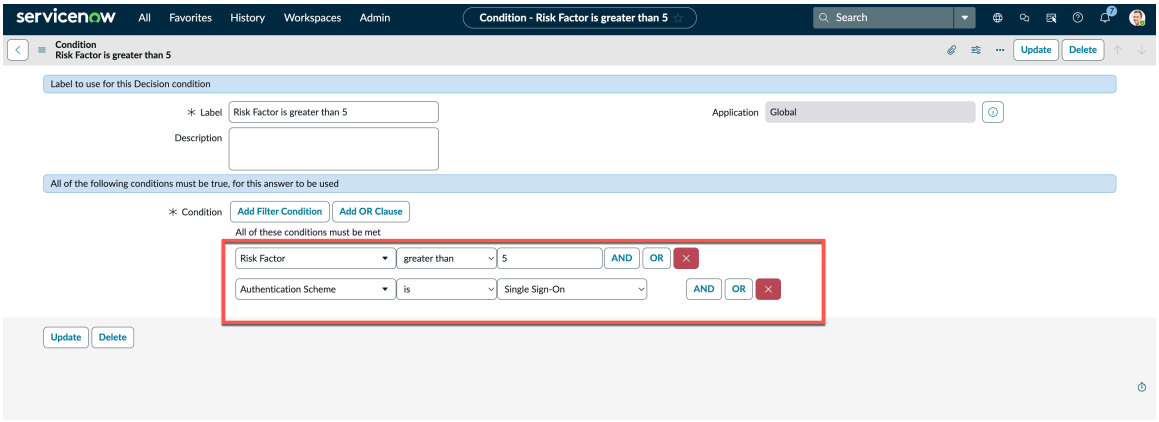

Based on this configuration, when the device is trusted from the Okta (IdP), then the user is authenticated to the instance.

For more information on how to create Post Authentication Context with Policy and Condition, see Post-authentication context.

2. Use of IDP attribute in MFA Policy Context.
   Example: Configuring to enable MFA from the Okta IDP attributes if the device isn't trusted.

   a. Navigate to **All** > **Adaptive Authentication** > **Auth Policy Contexts** > **MFA Authentication Policy Context.**.

   b. In the Policy Input, create the Policy Input and Policy Condition.

   - **Policy Input**: Add **device_trusted-okta**.

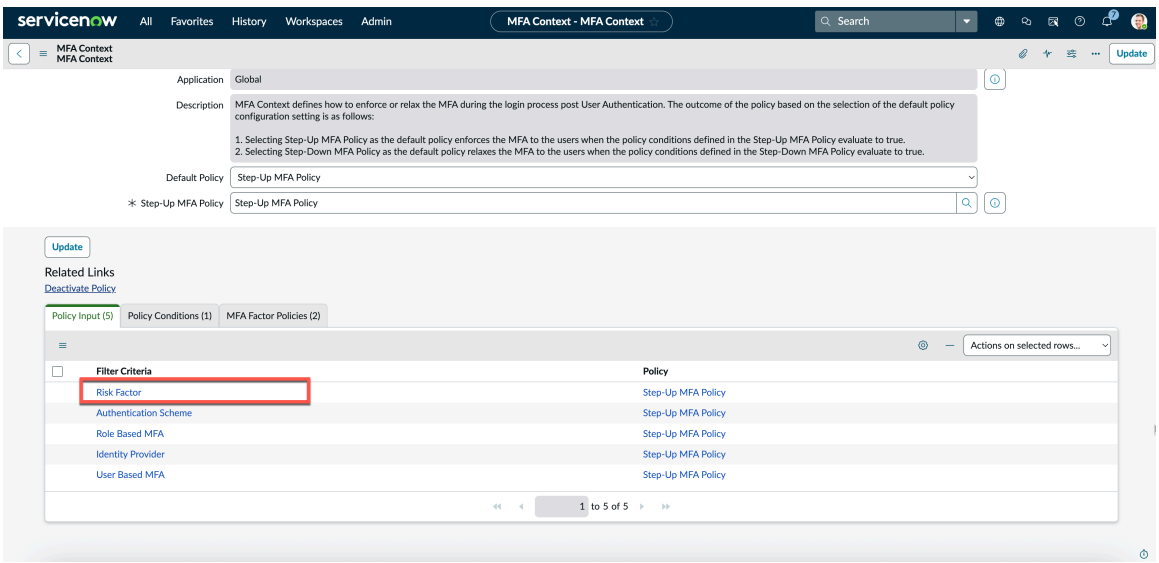

- **Policy Conditions**: **device_trusted-okta** is `not_trusted` and **Identity Provider** is `okta`.

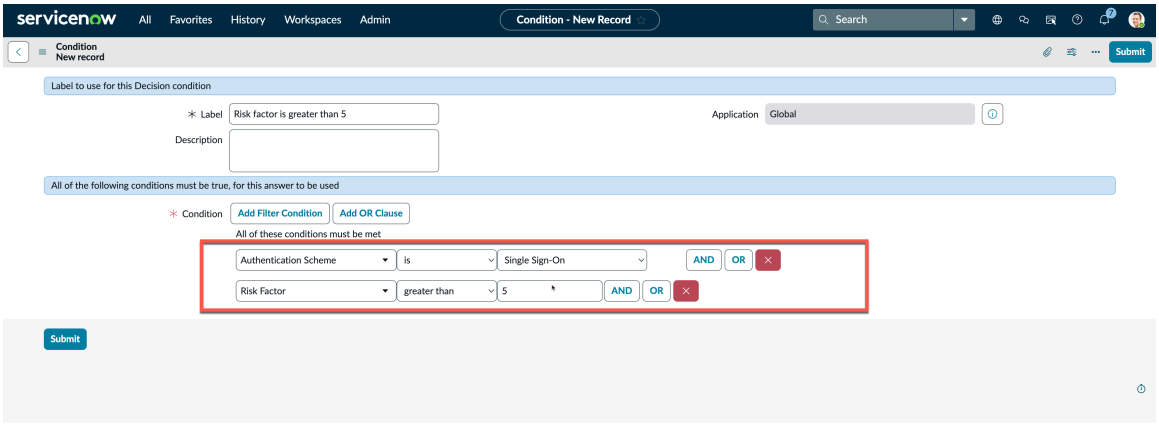

Based on this configuration, when the device is not-trusted from the Okta (IdP), then the user shown a second factor authentication to log in to the instance.

For more information on how to create MFA Context with Policy and Condition, see Multi-factor Authentication context.

3. Use of IDP attribute in Zero Trust - Policy based session access.
   Example: Configuring to reduce the privilege of `Itil` role from Okta IDP attributes if the device isn't trusted.

   a. Navigate to **All** > **Zero Trust Access** > **Session Access Role Configurations**.

   b. Create a Session Access role configuration.

   c. In the Policy Input, create the Policy Input and Policy Condition.

   - **Policy Input**: Add **device_trusted-okta** and **has itil role**.

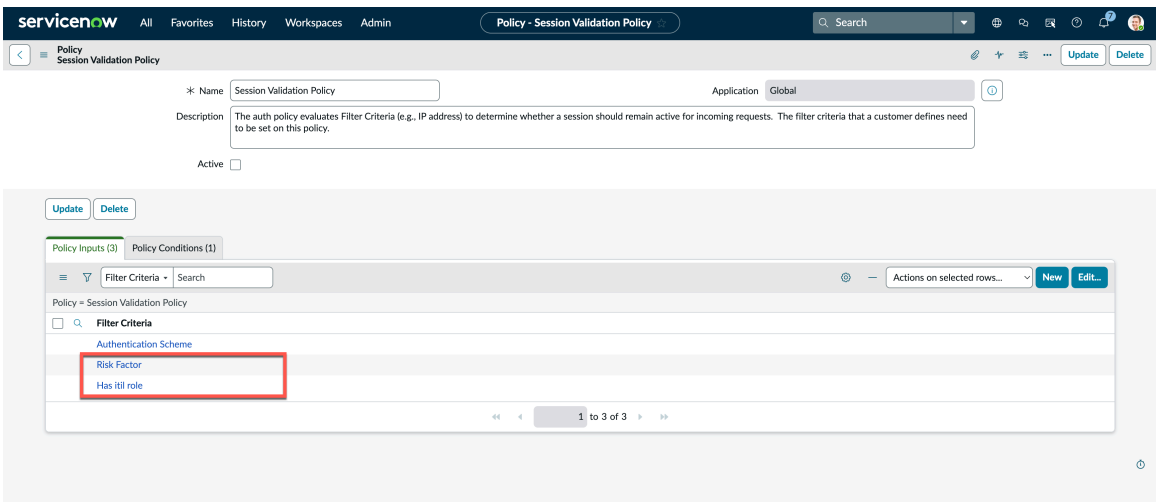

- **Policy Conditions**: **device_trusted-okta** is `not_trusted`, **Identity Provider** is `okta`, and **has itil role** is `true`.

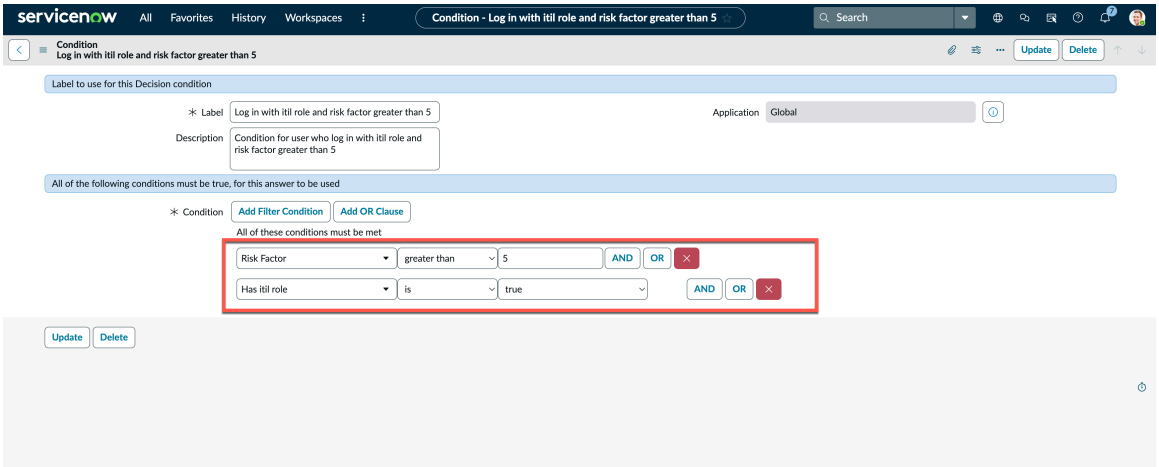

Based on this configuration, when the `itil` user using a device that is not-trusted from the Okta (IdP), then the user's privileges are reduced for the logged in session.

For more information on how to create Zero Trust - Policy based session access with Policy and Condition, see Zero Trust Access (ZTA).

Identity Provider attributes for OpenID Connect

Use the Identity Provider attributes that are received from the OpenID Connect (OIDC) from the Identity Provider (IdP) as a filter criteria for authentication.

You can manually create the IdP attributes based on the claims received as part of the ID token.

Note:

- Identity Provider filter is available with Zero Trust Access feature. For more information, see Zero Trust Access (ZTA).
- IdP attribute filter criteria can be used in Post-authentication context, Zero Trust Access (ZTA) session relegation, and Multi-factor Authentication context.

Start the configuration by adding the IdP attributes by selecting **New** from the Identity Provider Attributes section and use those attributes for Adaptive Authentication by setting it to `true`.

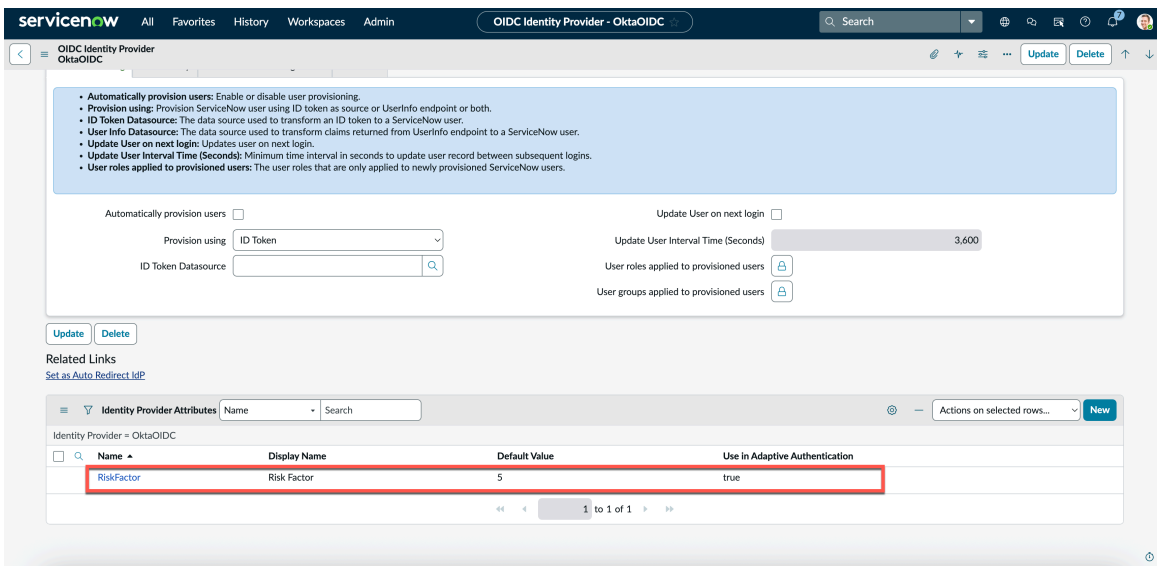

The **RiskFactor** defined in the OIDC configuration in the Identity Provider Attributes is from the ID token claims. This value can be an existing claim or custom claim as configured in the IdP side. Use this claim in various authentication context to customize and control the log in behavior of the user.

The **Identity Provider Attributes** are displayed with the following details:

Location Filter Criteria form

Field	Description
Name	Attribute name that is provided by the Identity Provider.
Display Name	Display Name is the detailed name that is used for the filter criteria. **Note:** You can provide a readable name as a Display Name, in some cases the Display Name provided by the Identity Providers are lengthy and not readable.
Default Value	Default value is used for filter criteria evaluation in case the attribute is missing in the SAML response.
Use in Adaptive Authentication	Option to use the Attribute in the Adaptive Authentication.

Note: Attributes that are populated from Azure IdP have name and display name limited to characters, due to the name length of the attribute.

You can also add new attributes by selecting **New** in the **Identity Providers Attributes** section.

If the Use in Adaptive Authentication is set to true, then the selected attribute is added as filter criteria in the Generic Filter Criteria. For example, **risk_score** set to `true`. The Generic Filter Criteria page has a new filter created.

Use Identity Provider Attribute as Filter Criteria for OIDC

Use the Identity Provider (IDP) attribute from the OpenID Connect (OIDC) response as a filter criteria for authentication policy.

Before you begin

Role required: admin

You can create session access policy using policy context (Pre-Authentication, Post Authentication, multi-factor authentication) and filter criteria (Role, Group, IP, Location) with policy inputs and conditions.

The following procedure shows steps to configure the IdP attribute from the SAML response as a policy input to control authentication in the **Post Authentication Context**, **Multi-factor authentication (MFA) Context**, and **Zero Trust - Policy based session access**.

The Okta IDP attributes are as displayed in the following screen shot. You should set the Use in Adaptive Authentication as true to use it in the **Post Authentication Context**, **Multi-factor authentication (MFA) Context**, and **Zero Trust - Policy based session access** policies.

Note: Policies in the post-authorization, MFA, Zero Trust - Policy based session access execute after the users enter the credentials or SSO response.

Procedure

1. Use of IDP attribute in Post Authentication Policy Context.
   Example: Configuring to enable log in from the Okta IDP attributes if the device is trusted.

   a. Navigate to **All** > **Adaptive Authentication** > **Auth Policy Contexts** > **Post Authentication Policy Context.**.

   b. Select **Allow Policy** and open the policy record.

   c. In the Policy Input, create the Policy Input and Policy Condition.

   - **Policy Input**: Add **Risk Factor**.

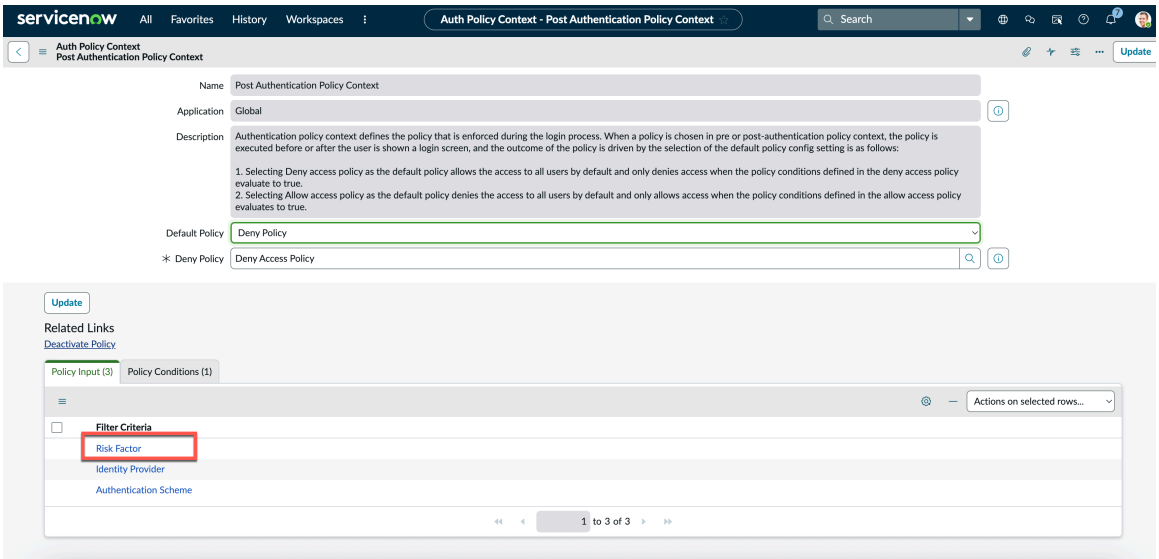

   - **Policy Conditions**: **Risk Factor** greater than `5` and **Authentication Scheme** is `Single Sing-On`.

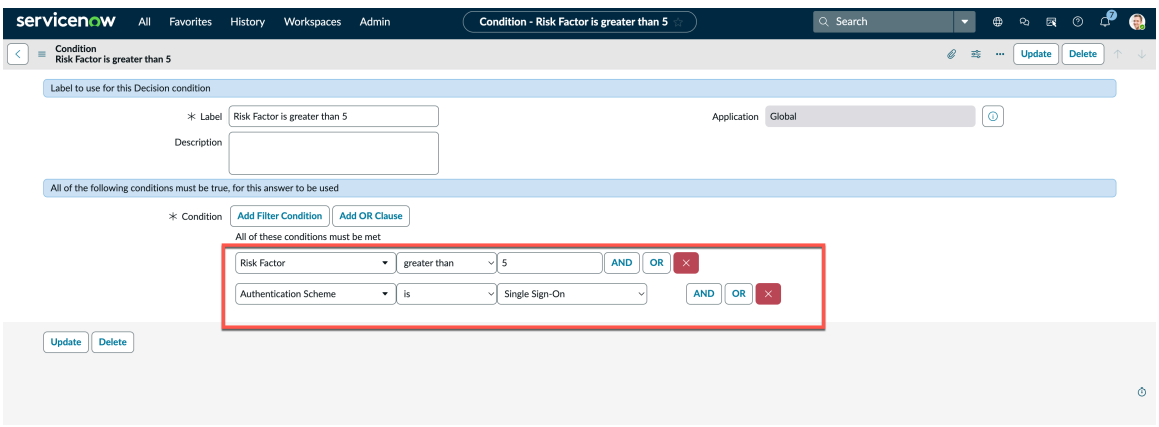

   Based on this configuration, when the device is trusted from the Okta (IdP), then the user is authenticated to the instance.

   For more information on how to create Post Authentication Context with Policy and Condition, see Post-authentication context.

2. Use of IDP attribute in MFA Policy Context.
   Example: Configuring to enable MFA from the Okta IDP attributes if the device isn't trusted.

   a. Navigate to **All** > **Adaptive Authentication** > **Auth Policy Contexts** > **MFA Authentication Policy Context.**.

   b. In the Policy Input, create the Policy Input and Policy Condition.

   - **Policy Input**: Add **Risk Factor**.

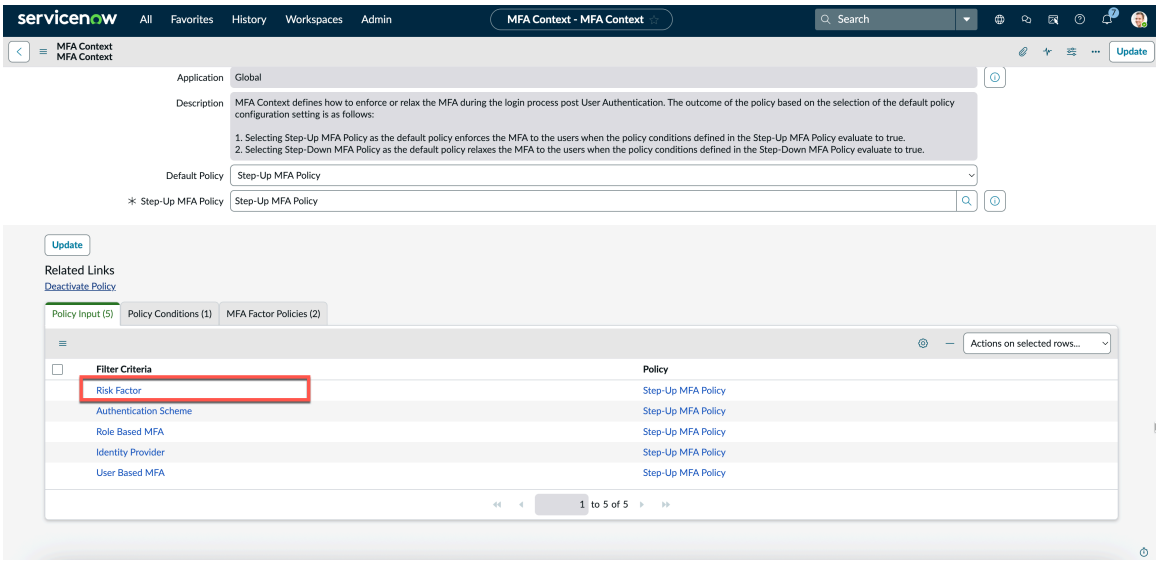

   - **Policy Conditions**: **Risk Factor** greater than `5` and **Authentication Scheme** is `Single Sing-On`.

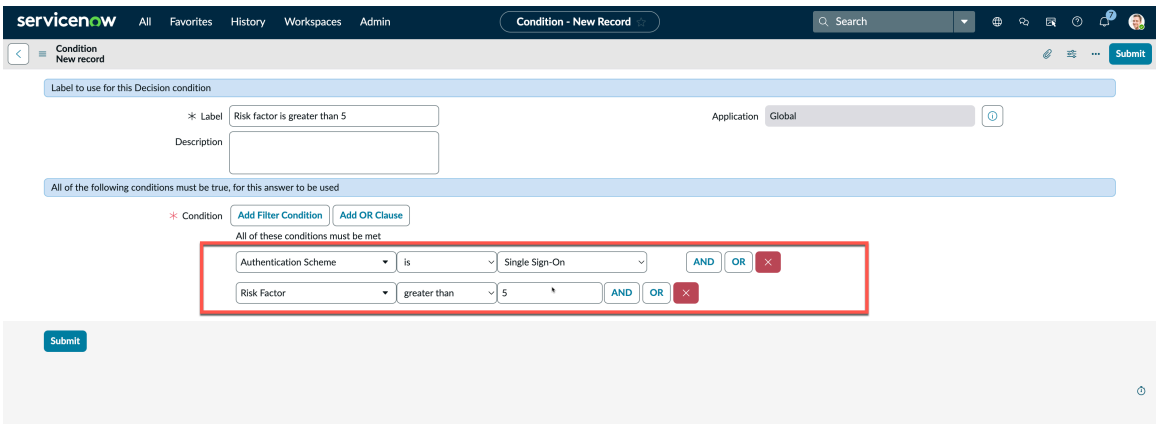

   Based on this configuration, when the device is not-trusted from the Okta (IdP), then the user shown a second factor authentication to log in to the instance.

   For more information on how to create MFA Context with Policy and Condition, see Multi-factor Authentication context.

3. Use of IDP attribute in Zero Trust - Policy based session access.
   Example: Configuring to reduce the privilege of `Itil` role from Okta IDP attributes if the device isn't trusted.

a. Navigate to **All** > **Zero Trust Access** > **Session Access Role Configurations**.

b. Create a Session Access role configuration.

c. In the Policy Input, create the Policy Input and Policy Condition.

- **Policy Input**: Add **Risk Factor** and **Has itil role**.

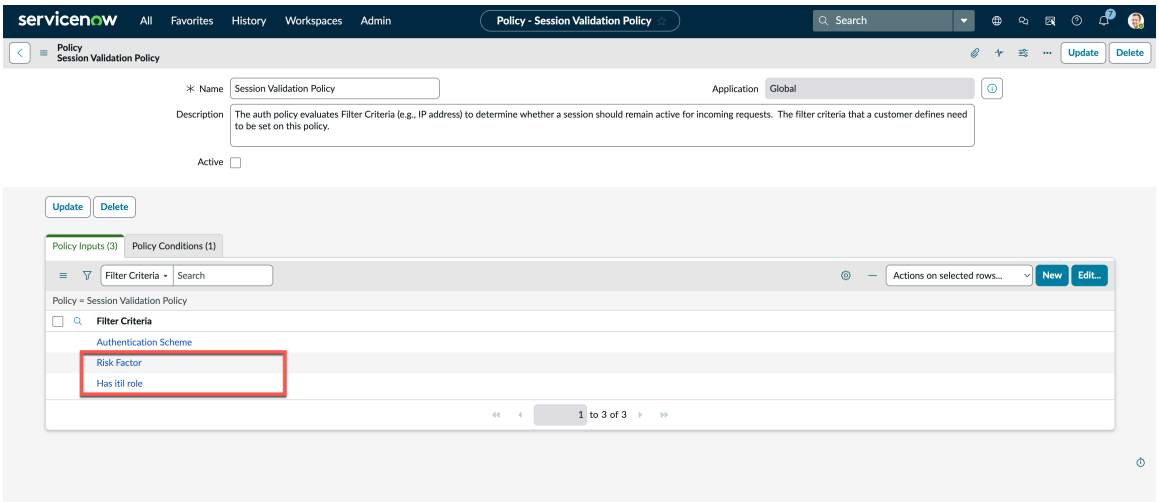

- **Policy Conditions**: **Risk Factor** greater than `5` and **Authentication Scheme** is `Single Sing-On`.

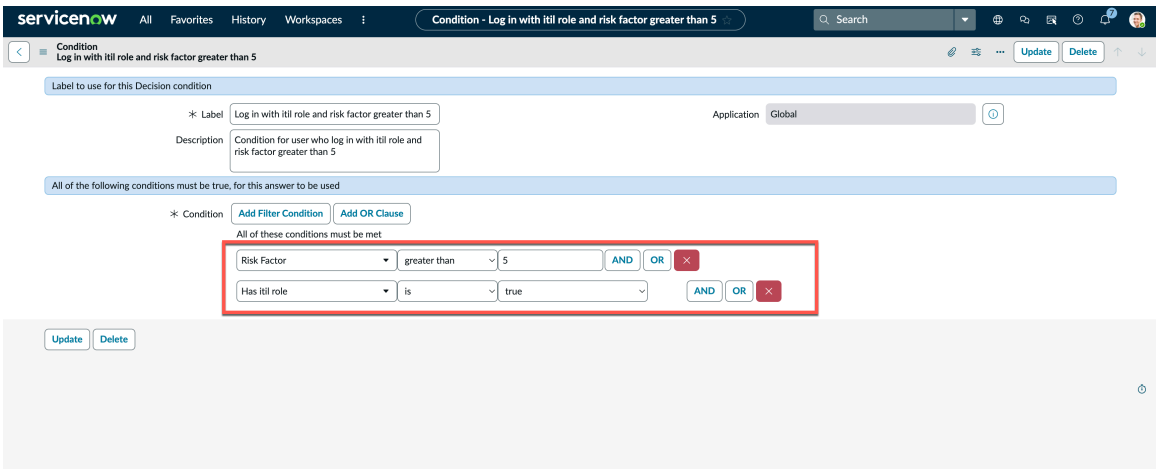

Based on this configuration, when the `itil` user using a device that is not-trusted from the Okta (IdP), then the user's privileges are reduced for the logged in session.

For more information on how to create Zero Trust - Policy based session access with Policy and Condition, see Zero Trust Access (ZTA).

Authentication policy contexts

Use authentication policy contexts to determine how and when your instance enforces authentication policies.

Authentication contexts define how and when a policy is enforced during the login process. Assign a policy to a policy context to define inputs and conditions regarding how your instance handles authentication.

Pre-authentication context

Policies in the pre-authorization context execute when a user first accesses the instance, before the they see a login screen. You can use the pre-authorization context to allow or deny access before your users are prompted for login credentials based on your selected policy. Because these policies evaluate before a user enters any information, those policies cannot take criteria such as a user's roles or groups into account.

For more detail on this context, see Pre authentication context.

Post-authentication context

Policies in the post-authorization context execute after your users enter their credentials or SSO response. Your instance allows or denies access based on your selected policy. Because your users have identified themselves via their login credentials, the policy can use user information to determine whether to grant access.

For more detail on this context, see Post-authentication context.

MFA (Multi-Factor Authentication) context

Policies assigned to the MFA context define whether to enforce MFA during the login process. Whether your instance enforces MFA is determined by the configuration of policies in this context. For more detail on this context, see Multi-factor Authentication context.

Account recovery context

Administrators can configure account recovery (ACR) to perform recovery activities such as addressing SSO misconfiguration or expired certificates. To use account recovery, you must register at least one admin account as an account recovery user. Single sign-on can’t be activated on your instance until there is at least one account configured. For more information about the context that can be set, see Account recovery context.

Session Validation context

The Session Validation context can be used with the Adaptive authentication policy framework. The framework uses authentication policies to evaluate authentication requests (session) and then either deny or allow access based on policy conditions. For more information, see Session validation context.

Default policy

Within the policy context, you can define a default policy in the **Default Policy** field. This default defines how your instance responds to the result of your policy. The available default policy options are determined by which context you are using. Detail on these options can be found in the docs describing these individual contexts.

Pre authentication context

The pre authentication policy context defines how and when a policy is enforced during the login process. The policy used in this context executes before your users see a login screen.

Pre authentication context record

Policies in the pre authentication context execute when a user first accesses the instance, before they see a login screen.

You can use the pre authentication context to allow or deny access before your users are prompted for login credentials based on your selected policy. Because these policies evaluate before a user enters any information, those policies can't take criteria such as a user's roles or groups into account.

Use the fields in the **Pre Authentication policy context** record to define how your instance uses your policy.

Pre Authentication context form

Field	Description
Name	Name of the policy context. This field is static and can't be changed.
Description	Description of the context
Default Policy	Defines the default behavior of this context when evaluating the policy. Select from the following options. **Allow Policy** Denies access to all users by default, and only allows access when the conditions the policy selected in the **Allow Policy** field evaluate to true. **Deny Policy** Allows access to all users by default, and only denies access when the conditions the policy selected in the **Deny Policy** field evaluate to true.
Allow Policy	The policy used for this context uses. This field appears only when the **Default Policy** field is set to **Allow Policy**.
Deny Policy	The policy used for this context uses. This field appears only when the **Default Policy** field is set to **Deny Policy**.

Note:

You can only use the IP Filter, Trusted Mobile App Filter, and Location Filter criteria in the Pre Authentication Policy Context.

Policy inputs and conditions

The **Policy Input** and **Policy Conditions** tabs display the inputs and conditions of the policy selected in the **Allow Policy** or **Deny Policy** field. These tabs serve as a reference, but can't be used to change the policy inputs or conditions. To modify your policy, navigate to the policy using the reference icon ( ) next to the **Allow Policy** or **Deny Policy** field.

This example shows a pre authentication policy context record configured to deny access by default. The context uses a policy called **Deny access policy**. That policy has a set of inputs and conditions that are displayed in the **Policy Input** and **Policy Condition** tabs.

Note:

- Only IP-Based filters, Location based filters, or Trusted Mobile App filter can be used in the pre authentication policy context.
- Whenever there's a pre authentication set with non absolute conditions or filter criteria, you're displayed with an error message stating that the policy or context can't be configured. It's recommended to validate all the inputs for the pre authentication context before executing it to the instance.

  For example: If the administrator is outside the trusted network and configures pre authentication context with IP ranges, if the IP ranges are mismatched with the current session of the admin, the admin is blocked.

Pre authentication policy context form

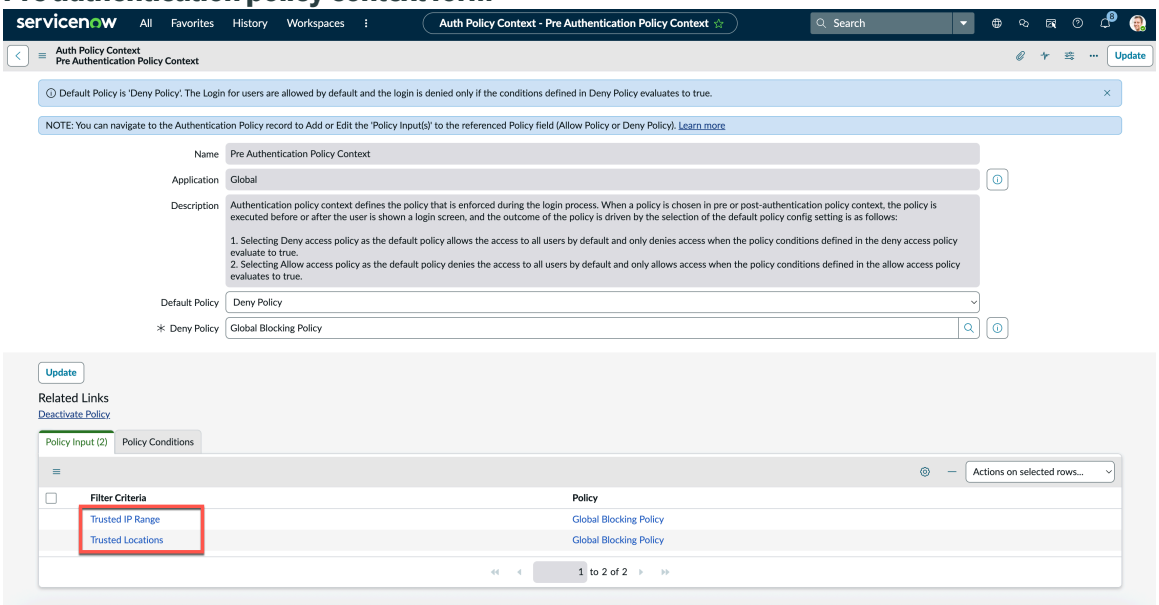

Post-authentication context

The Post Authentication policy context defines how and when a policy is enforced during the login process. The policy used in this context executes after your users see a login screen.

Post-authentication context record

Policies in the post-authorization context execute after your users enter their credentials or SSO response. Your instance allows or denies access based on your selected policy. Because your users have identified themselves via their login credentials, the policy can use user information such as role or group to determine whether to grant access.

Use the fields in the Post-authentication policy context record to define how your instance uses your policy.

Post-authentication context form

Field	Description
Name	Name of the policy context. This field is static and cannot be changed.
Description	Description of the context
Default Policy	Defines the default behavior of this context when evaluating the policy. Select from the following options.

Post-authentication context form (continued)

Field	Description
	Allow Policy Denies access to all users by default, and only allows access when the conditions the policy selected in the **Allow Policy** field evaluate to true. **Deny Policy** Allows access to all users by default, and only denies access when the conditions the policy selected in the **Deny Policy** field evaluate to true.
Allow Policy	The policy used for this context uses. This field appears only when the **Default Policy** field is set to **Allow Policy**.
Deny Policy	The policy used for this context uses. This field appears only when the **Default Policy** field is set to **Deny Policy**.

Policy inputs and conditions

The **Policy Input** and **Policy Conditions** tabs display the inputs and conditions of the policy selected in the **Allow Policy** or **Deny Policy** field. These tabs serve as a reference, but cannot be used to change the policy inputs or conditions. To modify your policy settings, navigate to the policy using the reference icon (ⓘ) next to the **Allow Policy** or **Deny Policy** field.

This example shows a post-authentication policy context record configured to deny access by default. The context uses a policy called **Deny access policy**. That policy has a set of inputs and conditions that are displayed in the **Policy Input** and **Policy Condition** tabs.

Post-authentication policy context form

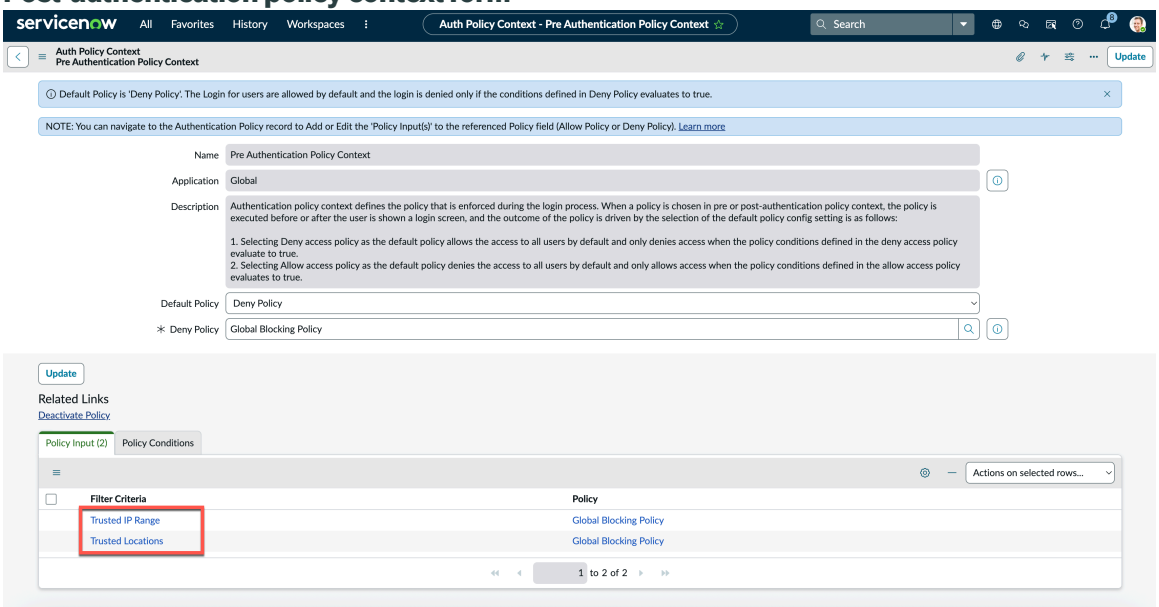

Multi-factor Authentication context

The Multi-factor Authentication (MFA) policy context uses a policy to define how and when MFA is enforced during the login process.

MFA context record

The MFA policy context defines whether your users must provide a second form of authentication when logging in. This context does not deny access to your instance as the post-authentication and pre-authentication policies. The policy you select in this context takes precedence over user or role-based configurations for multi-factor authentication.

To access the MFA context, navigate to **All** > **Multi-factor Authentication** > **MFA Context**.

Use the fields in the Post-authentication policy context record to define how your instance uses your policy.

Note:

- If the default policy is **Step-Up MFA Policy**, users will be shown with Multi-factor Authentication if policy configured in **Step-Up MFA Policy** evaluates to true. Policy takes precedence over the user or role based configuration.
- MFA with SSO login will only be available if `glide.authenticate.mfa.with.multisso.enabled` Property is set to true.
- You can navigate to the Authentication Policy record to Add or Edit the 'Policy Input(s)' to the referenced Policy field (**Step-Up MFA Policy** or **Step-Down MFA Policy**).
- MFA context policy applies only for user log ins. It does not apply for API authentication, basic auth, and OAuth resource owner password credential grant.

MFA context form

Field	Description
Name	Name of the policy context. This field is static and cannot be changed.
Description	Description of the context
Default Policy	Defines the default behavior of this context when evaluating the policy. Select from the following options. **Step-Up MFA Policy** Enforces MFA to users when the policy conditions defined in the **Step-Up MFA Policy** field evaluate to true. **Step-Down MFA Policy** Enforces MFA by default. MFA is not enforced only when the policy conditions defined in the **Step-Down MFA Policy** field evaluate to true.
Step-Up MFA Policy	The policy used for this context uses. This field appears only when the **Default Policy** field is set to **Step-Up MFA Policy**.
Step-Down MFA Policy	The policy used for this context uses. This field appears only when the **Default Policy** field is set to **Step-Down MFA Policy**.

Policy inputs and conditions

The **Policy Input** and **Policy Conditions** tabs display the inputs and conditions of the policy selected in the **Step-Up MFA Policy** or **Step-Down MFA Policy** field. These tabs serve as a reference, but cannot be used to change the policy inputs or conditions. To modify your policy settings, navigate to the policy using the reference icon (ⓘ) next to the **Step-Up MFA Policy** or **Step-Down MFA Policy** field.

Note: Policy conditions can be created from here, but as a good practise it is recommended to add new policy conditions from policy page.

This example shows an MFA context record configured using a step-up MFA policy. This default policy means that MFA is enforced only when the conditions defined in the policy evaluate to true. The context uses a policy called **Step-Up MFA policy**. That policy has a set of inputs and conditions that are displayed in the **Policy Input** and **Policy Condition** tabs.

MFA policy context form

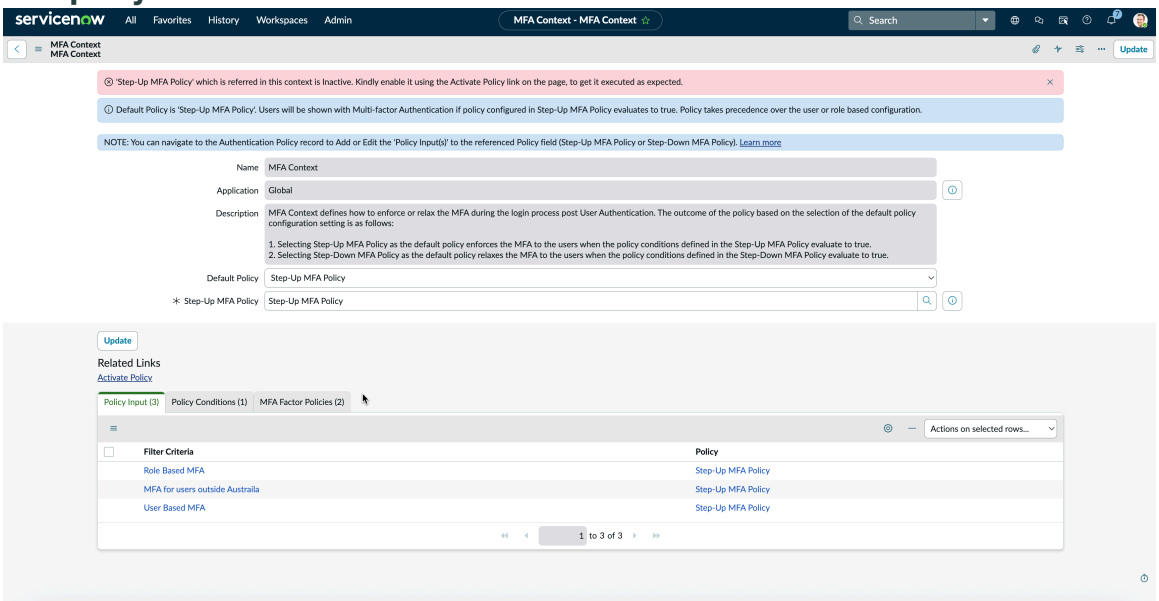

MFA factor policies

MFA factor policies are a critical component of an organization's security posture, enabling you to enforce additional verification steps beyond passwords. These policies define the authentication methods that users must employ to access providing a flexible and customizable approach to authentication. For more information, see Multi-Factor Authentication factor policies.

Account recovery context

The account recovery context uses a policy to define how and when the account recovery can be established.

Administrators can view and modify this context and its associated policy by navigating to **Multi-Provider SSO** > **Account Recovery** > **Account Recovery Context**.

Note: By default the policy is **Allow Policy**. The Login for users are restricted by default and the login is allowed only if the conditions defined in **Allow Policy** evaluates to true.

Use the fields in the account recovery context record to define how your instance uses the policy.

Account recovery context form

Field	Description
Name	Name of the policy context. This field is static and cannot be changed.
Description	Description of the context
Default Policy	Defines the default behavior of this context when evaluating the policy. Select from the following options:

Account recovery context form (continued)

Field	Description
	• Allow Policy • Deny Policy
Allow Policy	This field appears only when the **Default Policy** field is set to **Allow Policy**.
Deny Policy	This field appears only when the **Default Policy** field is set to **Deny Policy**.

Policy inputs and conditions

The **Policy Input** and **Policy Conditions** tabs display the inputs and conditions of the policy selected in the **Allow Policy** or **Deny Policy** field. These tabs serve as a reference, but cannot be used to change the policy inputs or conditions. To modify your policy settings, navigate to the policy using the information icon next to the **Allow Policy** or **Deny Policy** field.

Note: Policy conditions can be created from here, but as a good practice it is recommended to add new policy conditions from policy page.

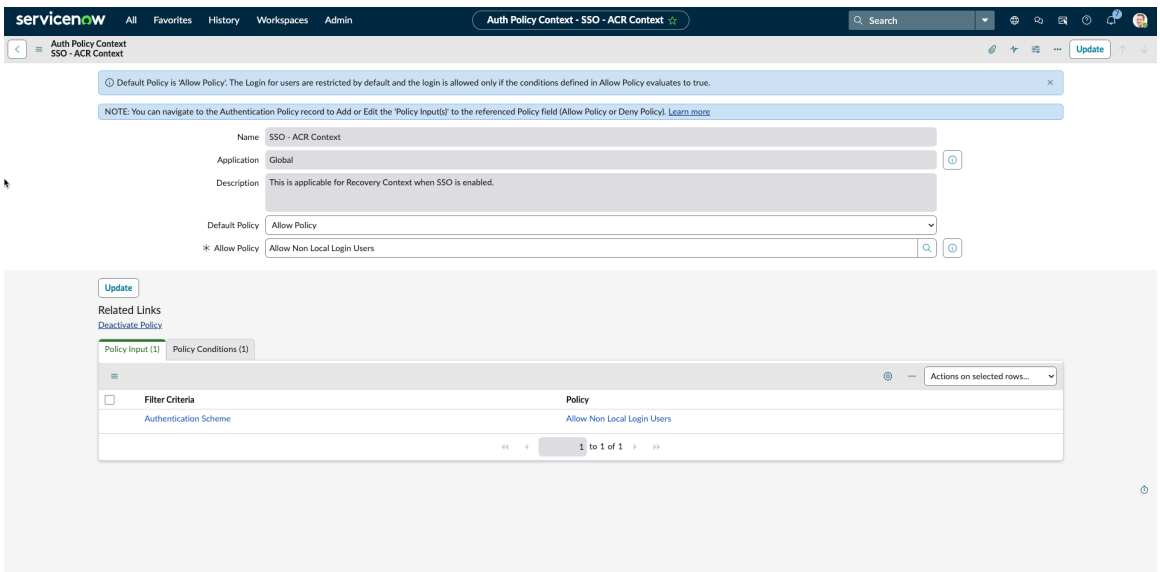

Session validation context

Use the Session Validation Context as an additional layer of protection against session or cookie hijacking.

You can use the Session Validation Context with the Adaptive authentication policy framework. The framework uses authentication policies to evaluate authentication requests and then either denies or allows access based on policy inputs and conditions.

The Session Validation Context policy can be used in conjunction with post auth policy, where an admin can enforce IP restrictions to certain or all users during the logged in session.

The Session Validation Context feature evaluates the IP-addresses based on the conditions you set and allows access to the instance within a session. The Session Validation Context outcome is set based on selecting **Allow Policy** as this policy terminates the user session immediately unless one of the policy conditions defined in the allow access policy evaluates to true.

Note: The Session Validation Context for an authentication policy can only be with **Allow Policy**.

The Session Validation Context works based on the following mechanism:

- Captures the user's IP address on session creation from user request and stores it in the session and database.
- Rejects a request when its IP address differs from that in the session or outside of the customer defined valid IP ranges you defined.

Note: The Session Validation Context is:

- Available only for authenticated users.
- Not applicable for guest user sessions or native mobile apps.
- Optional and based on the requirement that it can be configured.
- Executed only for the post-login requests.

Benefits of Session Validation

The Session Validation Context has the following benefits:

- Restricts access to ServiceNow® when hijackers copy a user's session cookies from one device to another to impersonate the session.
- Restricts the user's session access if they're using an insecure network.
- Configures the various rules and IP ranges by user group or role for user logins.

Session Validation context record

Policies in the session validation context execute post-login requests.

Use the fields in the session validation policy context record to define how your instance uses your policy.

Session Validation context form

Field	Description
Name	Name of the policy context. This field is static and can’t be changed.
Description	Description of the context.
Default Policy	Defines the default behavior of this context when evaluating the policy. Select from the following options. **Allow Policy** Denies access to all users by default, and only allows access when the conditions in the **Allow Policy** field evaluate to true. **Deny Policy** Allows access to all users by default, and only denies access when the conditions in the **Deny Policy** field evaluate to true.
Allow Policy	The policy used for this context. This field appears only when the **Default Policy** field is set to **Allow Policy**.

Session Validation context form (continued)

Field	Description
Deny Policy	The policy used for this context. This field appears only when the **Default Policy** field is set to **Deny Policy**.

You can choose the **Session Validation Policy** as Allow Policy or Deny Policy based on the policy input and policy conditions.

Note:

You can only use the IP, Role, and Group filter criteria for Session Validation policy.

Policy inputs and conditions

The **Policy Input** and **Policy Conditions** tabs display the inputs and conditions of the policy selected in the **Allow Policy** or **Deny Policy** field. These tabs serve as a reference; but they can't be used to change the policy inputs or conditions. To modify your policy, navigate to the policy using the reference icon next to the **Allow Policy** or **Deny Policy** field.

Activate Session Validation Context

Use Session Validation Context to restrict access to ServiceNow® when hijackers copy a user's session cookies from one device to another to impersonate the session or restricts the user's session access if they're using an insecure network.

Before you begin

Role required: admin

To use the session validation, you must perform the following steps:

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Authentication Policies** > **All Policies**.
2. Select the **Session Validation Policy** in the Policies (`sys_authentication_policy_list.do`) page.
3. Specify the **Policy Inputs** and **Policy Conditions**.
4. Set the conditions to **true** or **false** for the filters that you've added.
5. Select the **Active** check box to activate the policy after you set up the Session Validation Policy is set up with policy inputs and conditions.
6. Select a **Force AuthnRequest** check box for the default identity provider in the sso_properties table if the instance has SSO configured.
7. Navigate to **All** > **Adaptive Authentication** > **Authentication Policies** > **Properties**.

8. Set the system property `session.validation.enabled` to **Yes**.

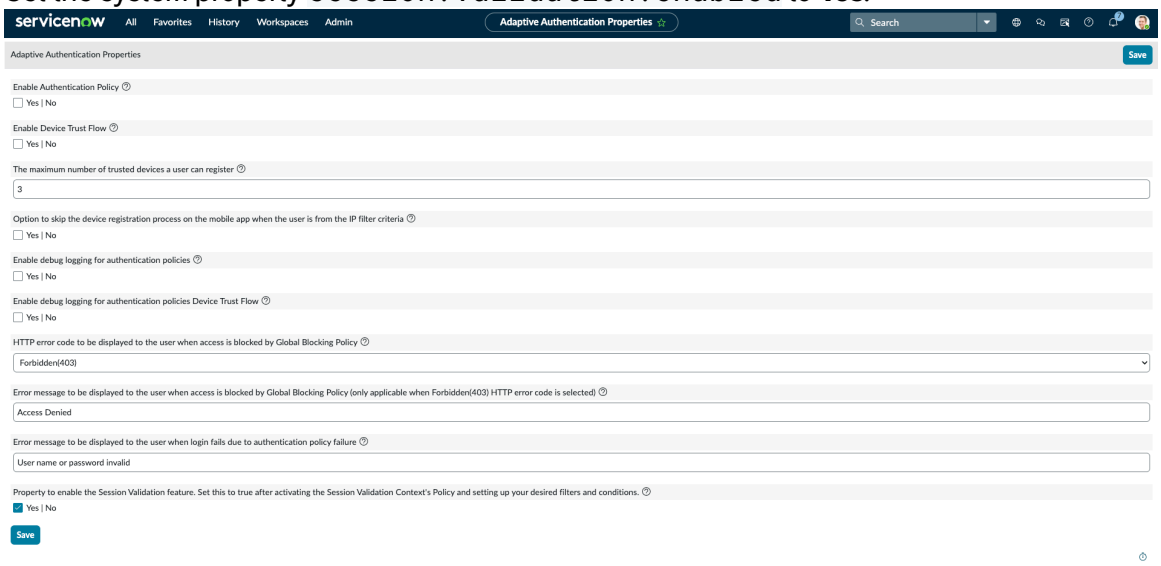

Result

The Session Validation feature is activated. You can configure the policy inputs and conditions for the policy to use the feature. To learn more, see Tutorial: Configuring session validation.

Tutorial: Configuring session validation

Configure session validation within the Adaptive Authentication framework to provide as an additional layer of protection for session or cookie hijacking.

Before you begin

Role required: admin

Plugin required: **Adaptive Authentication** (com.snc.adaptive_authentication)

To configure Session Validation, you must perform the following steps:

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Authentication Policies** > **All Policies**.
2. Select the **Session Validation Policy** in the Policies (`sys_authentication_policy_list.do`) page.
3. Select **Policy Inputs**.
   a. Select **New** or **Edit**.
   b. Choose the kind of Policy Input (Filter Criteria) that you want to create.

Available options are IP, Role, and Group Filter Criteria. Let's choose **IP Filter Criteria**.

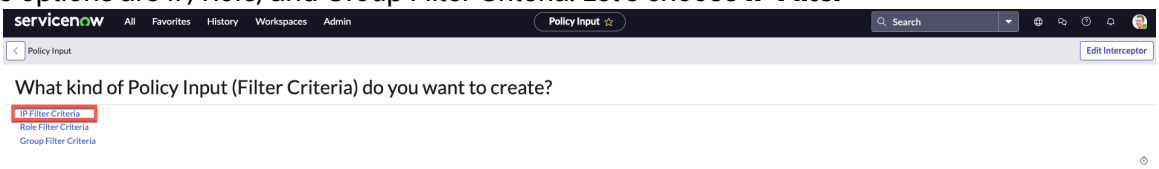

c. Fill the form with the filter details and provide the **IP Range**.

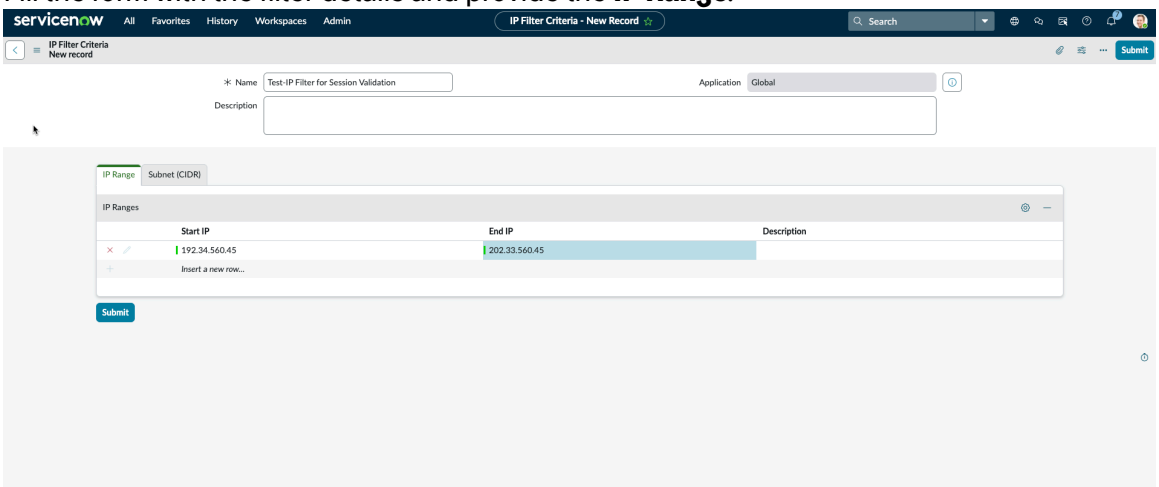

To learn more about how to create an IP Filter, see Create IP filter criteria.

d. Select **Submit**.

4. Select **Policy Conditions** on the Session Validation Policy page.

a. Select **New**.

b. Fill the form and set the Condition for the Policy Input.

Note: You can set the conditions to `true` or `false` based on the configuration of the policy input. In this example, it is set to `true`. Setting the condition to true in this case allows only the user with the configured IP address to log in.

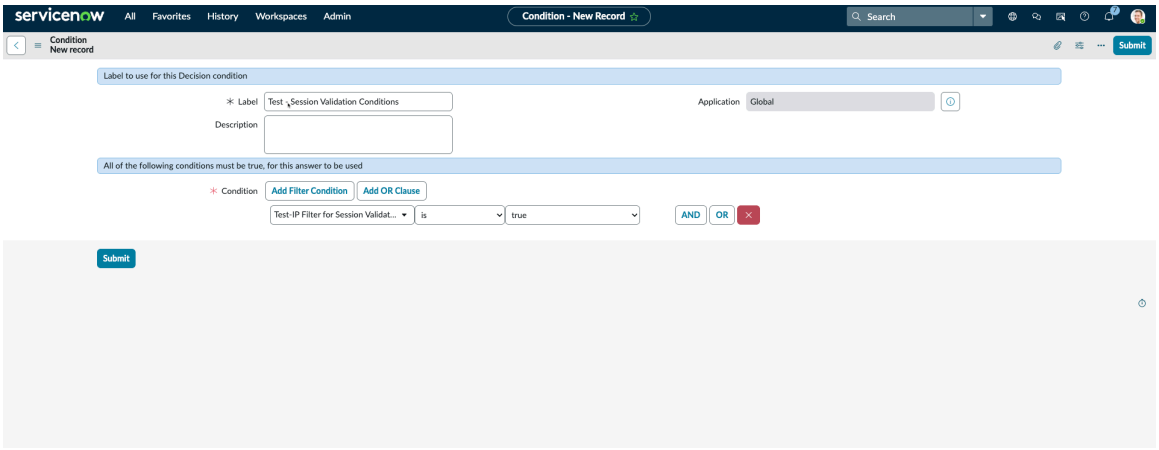

5. Select the `Active` check box to activate the policy after the Session Validation Policy is set up with policy inputs and conditions.

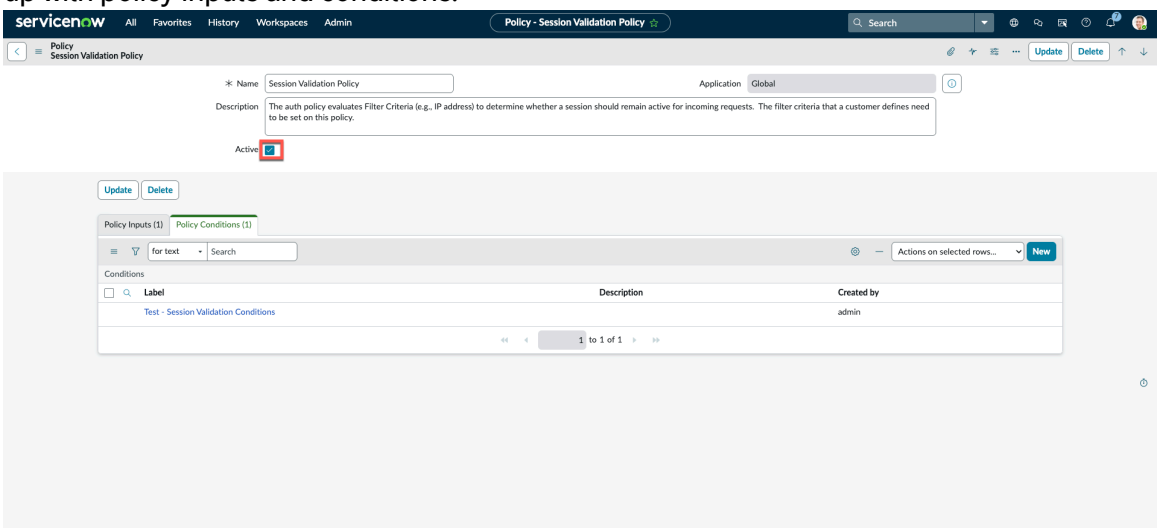

6. Navigate to **All** > **Adaptive Authentication** > **Authentication Policies** > **Properties** and enable the Session Validation property.

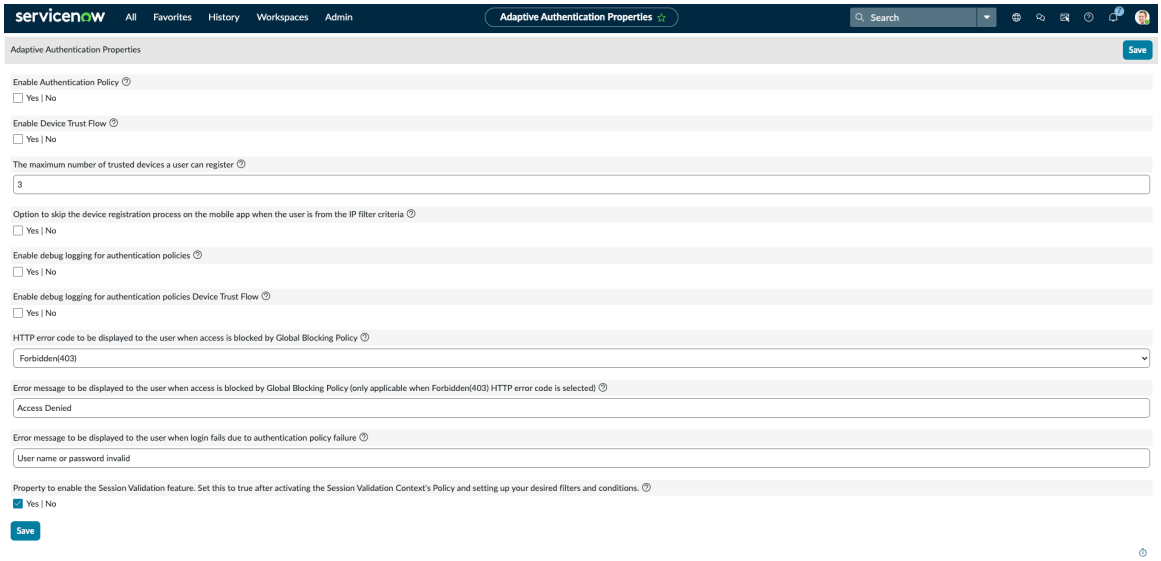

7. Navigate to **All** > **Adaptive Authentication** > **Auth Policy Contexts** > **Session Validation Context**.
8. Set the Default Policy to **Allow Policy** or **Deny Policy** to set the session validation context according to the policy input and policy conditions.

Note: By default:

- The Session Validation context is set to **Allow Policy**.
- Allow Policy is selected as **Session Validation Policy**.
- The Session Validation Context for an authentication policy can only be with **Allow Policy**.

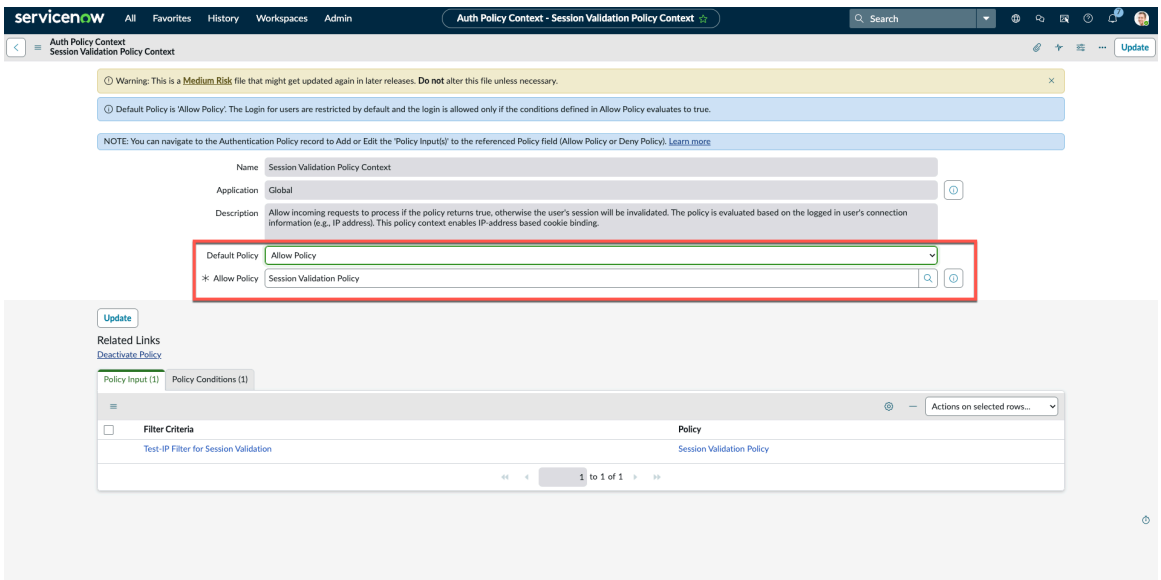

Result

The configuration evaluates the login session based on the following:

- Restricts access to the ServiceNow® instance when hijackers copy a user's session cookies from one device to another to impersonate a session.
- Restricts the user's session access if they're using an insecure network.

Authentication policies

Authentication policies evaluate authentication requests based on the specified policy conditions and either allow or deny access depending on the output of policy conditions evaluation. For example, access is allowed only if all the policy conditions specified in **Allow Access Policy** evaluate to true.

Use the built-in authentication policies or create an authentication policy according to your security requirements. You can find the policies on your instance by navigating to **Adaptive Authentication** > **Authentication Policies** > **All Policies**.

Note: At any point of time, either Allow Access Policy or Deny Access Policy can be executed but not both.

Policy	Description
Allow access policy	Denies all the authentication requests by default. Allows only the authentication requests that match the specified policy conditions.
Allow access pre-auth policy	Allow access pre-auth policy.
Allow non local login users	Choose this policy to allow non-local login users. Used in the context of SSO recovery flows.
Deny access policy	Allows the authentication requests by default. Denies only the authentication requests that match the specified policy conditions.
Global blocking policy	Denies the access requests of users and APIs before authentication. This policy can be used as an alternative to the IP Address Access Control.
Session Validation Policy	The authentication policy evaluates Filter Criteria (For example, IP address) to determine whether a session should remain active for incoming requests. You must set the filter criteria on this policy.
Local login deny policy	Choose this policy to block all local logins. Used in the context of SSO recovery flows.
Step-down MFA policy	Choose **Step-Down MFA Policy** as the default policy in situations where users do not require MFA authentication. When the policy conditions defined in the Step-Up MFA Policy evaluate to true, users are not required to login using MFA.
Step-up MFA policy	Choose **Step-Up MFA Policy** as the default policy to require MFA authentication when the policy conditions defined in the Step-Up MFA Policy evaluate to true.

Configure an authentication policy

Configure an authentication policy to define inputs and conditions to used to grant access to an instance or enforce multi-factor authentication.

Before you begin
Role required: admin

Procedure
1. Navigate to **All** > **Adaptive Authentication** > **Authentication Policies** > **All Policies**.

   Note: To see examples of completed policies, you can review these policies on your instance:
   - DEMO POLICY - Allow Local Login for Admins from the Trusted IP Range only
   - DEMO POLICY - Allow Local Login for Admins only
   - DEMO POLICY - Restrict Username and Password based Authentication for specific users

2. Click the **New** button to create a new policy record.
3. In the **Policy** form, fill in the fields.

Policy form

Field	Description
Name	Name of your policy.
Application	The scoped application for the policy. This field is filled automatically with the current application.
Description	Description of the policy
Active	Whether the policy is active.

4. From the **Policy Inputs** tab, click **Edit**.
5. Select one or more filter criteria from the **Collection** list and move them to **Policy Inputs List** of **Allow Access Policy** list.

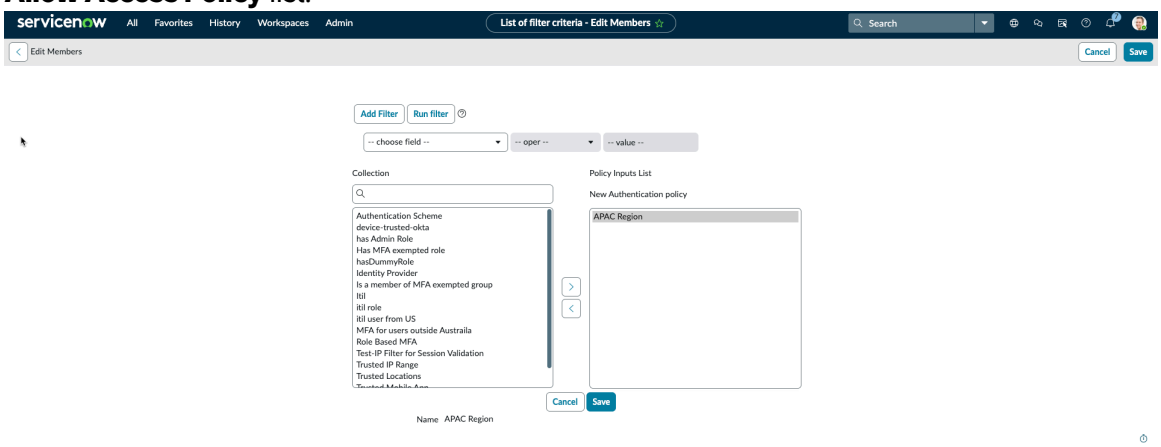

   Note: For information on creating your own filter criteria to use in this section, see Filter criteria.

6. From the **Policy Conditions** tab, click **New**.
7. On the form, fill in the fields:

Condition form

Field	Description
Label	Name to identify the condition.
Description	Description of the condition.
Condition	Logical combination of multiple policy inputs (filter criteria) that is used to evaluate authentication requests. For example, you can create conditions that allow only contractors from a list of trusted IP addresses.

8. **Optional:** Repeat step 7 to create additional policy conditions.

   Note: If you create multiple policy conditions, the final output of the access policy depends on the logical OR output of the all policy conditions. This means the policy will evaluate to true if any one of your policy conditions is met.

9. Click **Save**.

Add an authentication policy to an authentication policy context

Add an authentication policy to one of the authentication policy contexts. The authentication context uses the policy inputs and conditions to determine whether uses are granted access to the instance, or whether MFA is enforced for your users.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Auth Policy Contexts**, and select one of the entries depending on your needs.

   Pre Authentication Context

   Use the pre-authorization context to evaluate your policy before the user sees the login screen. Users are granted or denied access based on this evaluation.

   Post Authentication Context

   Use the post-authorization context to evaluate your policy after the user enters their login credentials. Users are granted or denied access based on this evaluation. Because this evaluation takes place after the instance identifies the user, policies in this context can make their evaluation based on user data, such as their role or group.

   MFA Context

   Use the MFA context to determine whether a user must use multi-factor authentication when logging in.

   Session Validation Context

   Use the session validation context to evaluate the IP-Address that are set based on the conditions defined as a filter and allows access to the instance within a session.

2. In the **Default Policy** field, select a value.

The value in this field determines how the context uses the outcome of your policy's conditions.

The available options in this field depend on the selected context. For details on these contexts, see Authentication policy contexts.

3. Assign a policy to the context.
The name of the field depends on which context you have selected.

   Pre-authentication and Post-authentication context

   These contexts have a **Deny Policy** or **Allow Policy** field, depending on the selection in the **Default Policy** field.

   MFA context

   This context has a **Step-Up MFA Policy** or **Step-Down MFA Policy** field, depending on the selection in the **Default Policy** field.

   Session validation context

   This context has a **Allow Policy** or **Deny Policy** field, depending on the selection in the **Default Policy** field.

4. Click **Update**.
After updating the record, you can see your policy's inputs and conditions in the **Policy Input** and **Policy Conditions** tabs.

Adaptive Authentication Events

You can use the adaptive authentication events table to know about the events.

Following are some of the logs that can be tracked in the adaptive authentication events table:

- `glide.adaptive.auth.log.success.event`: To know the `success` events for Post Auth, API logins.
- `glide.adaptive.auth.log.mfa.relax.event`: To know the `mfa_relaxation` events for MFA Context.

Note:

- For Pre-Auth Context, Rest API Access policies of type "`pre_auth`": By default `Failure` events are logged. `Success` events `can't` be logged.
- For Post Auth Context, Rest API Access policies of type other than "`pre_auth`": By default `Failure` events or logged. For `Success` events to log, you need to enable the `glide.adaptive.auth.log.success.event` property.
- For MFA: By default `MFA Enforcement` events are logged.
- For MFA Relaxation events to log, you need to enable the `glide.adaptive.auth.log.mfa.relax.event` property. `MFA Relaxation` would log the same event as `MFA Enforcement`, but with the `Result` Field populated with `False`.

Configure adaptive authentication properties

After activating Adaptive Authentication, configure adaptive authentication properties according to your security requirements.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Authentication Policies** > **Properties**.
2. Configure these properties:

Adaptive Authentication Properties

Property	Description	Value
Enable Authentication Policy (glide.authenticate.auth.policy.enabled)	Option to enable the authentication policy.	Yes | No
Enable debug logging for authentication policies (glide.authenticate.policy.debug)	Option to enable debug logging for the authentication policies.	Yes | No
HTTP error code to be displayed to the user when access is blocked by Global Blocking Policy (glide.authenticate.global.blocking_policy.error_code)	HTTP error code that displays during login when the Global Blocking Policy blocks a user login.	Select from: ○ Forbidden(403) ○ Not Found(404)
Error message to be displayed to the user when access is blocked by Global Blocking Policy (only applicable when Forbidden(403) HTTP error code is selected) (glide.authenticate.global.blocking_policy.error_message)	Error message that displays when the Global Blocking Policy blocks access.	Text field
Enable Device Trust Flow (glide.authenticate.preauth.allow.trusted.device)	Option to enable the trusted device flow.	Yes | No
Maximum number of trusted device a user can register (glide.trusted.device.max.count)	It is the maximum number of trusted device a user can register.	Text field
Skip the device registration for device trust flow if the user is from the trusted network (glide.authenticate.preauth.skip.user.registration)	Option to skip registration if the user is trying to register from the trusted network	Yes | No
Property to enable the Session Validation feature. Set this to true after activating the Session Validation Context's Policy and setting up your desired filters and conditions (session.validation.enabled)	Option to enable the Session Validation feature. Set this to true after activating the Session Validation Context's Policy and setting up your desired filters and conditions.	Yes | No

Tutorial: Configure adaptive authentication

Use these example steps to configure adaptive authentication on an instance.

To use this tutorial, you must have an instance with Adaptive Authentication activated. For details on this process, see Activate adaptive authentication.

The example guides you through creating a new policy and applying it to an instance. In this tutorial you will:

Create a filter criteria record

Create a group filter criteria record to use as an input for your policy. This record allows your policy to determine access based on a user's group. In these steps you define the group or groups the policy uses to determine access.

Create a policy

Create a policy that determines whether a user can access the instance. This policy uses the group filter criteria record you create as an input. In these steps you also define policy conditions that define how the policy uses the policy input to determine user access.

Configure a policy context

Configure the **Post Authentication Policy Context** to use your new policy. When configured, your instance denies access to users within the group defined in the filter criteria record.

Create a filter criteria record

Learn how to create a criteria record to use as a policy input for your adaptive authentication policy.

Before you begin

Role required: admin

To deny access to your instance based on user groups, you must create a group filter criteria record. This record defines a user group or a set of user groups that your policy can grant or deny access. In this example, you will create a group filter criteria record for a single user group.

For details on user groups and how they are used in your instance, see Exploring user administration.

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Filter Criteria** > **Group Criteria**.
2. Click **New** to create a new record.
3. In the **Name** field, enter a name for your record.
   For example, `Denied Groups`.
4. In the **Description** field, enter a brief description.
   For example, `User groups denied access to the instance`.
5. In the **Groups for criteria** list, double-click **Insert a new row...**.
6. Enter the name of a user group, or click the reference icon ( ) to select a group from a list.
   If you want to create a new user group for your filter criteria, click the reference icon ( ), and then click the **New** button. For more details on creating user groups, see Create a user group.
7. After you have added your user group, click **Submit** to save your criteria record.

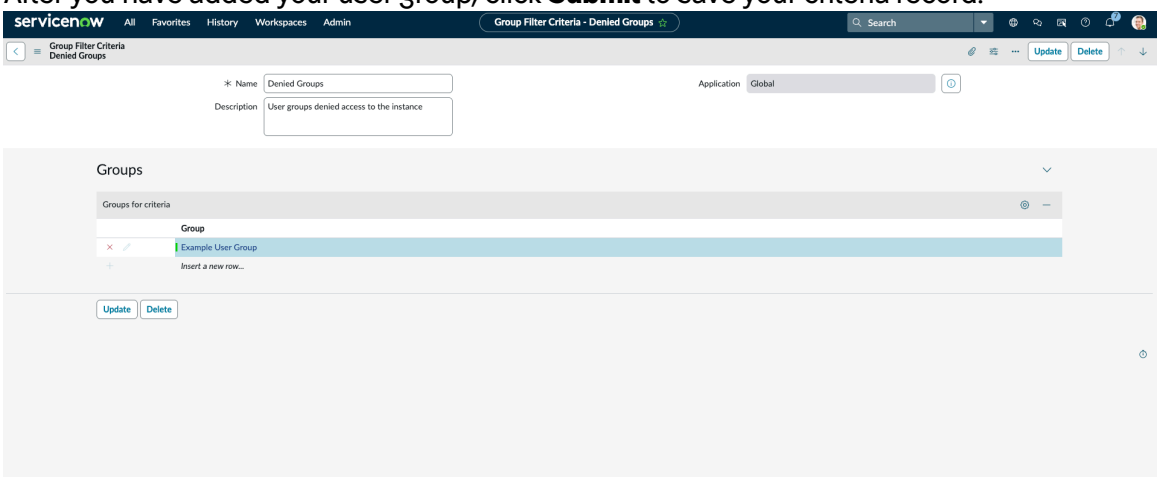

Create a policy

Learn how to create a policy to deny access to user groups defined in your group filter criteria.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Authentication Policies** > **All Policies**.
2. In the **Policies** list, click **New**.
3. In the **Name** field, enter a name for your record.
   For example, `Deny by Group`.
4. In the **Description** field, enter a brief description.
   For example, `User groups denied access to the instance`.
5. Right-click the form header and click **Save**.
   After saving, the **Policy Inputs** and **Policy Conditions** lists display on the form.
6. In the **Policy Inputs** list, click **Edit**.
7. In the list collector, add the filter criteria created in the previous steps to your policy.

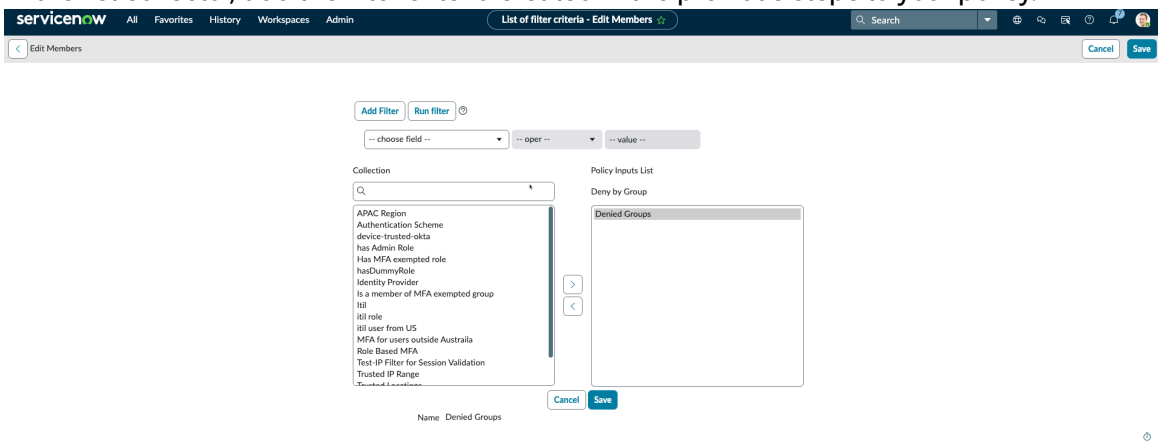

8. Click **Save**.
   The instance saves the policy input and displays the policy record displays again.
9. In the **Policy Conditions** list, click **New**.
10. In the **Label** field, create a label for this condition, for example `Deny by group`.
11. In the **Description** field, enter a brief description.
    For example, `Deny access to users in Denied Groups`.
12. In the **Condition** field, select the filter criteria you created in previous steps, then select **is**, and **true** to complete the condition.

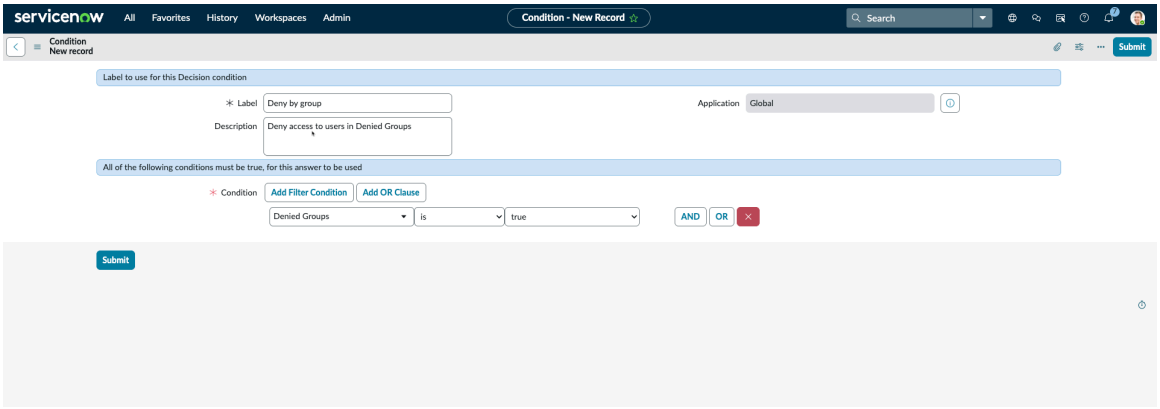

13. Click **Save**.

Configure a policy context

Configure the Post Authentication Policy Context to use your new policy. When configured, your instance denies access to users within the group defined in the filter criteria record.

Before you begin

Role required: admin

Procedure

1. **Adaptive Authentication** > **Auth Policy Contexts** > **Post Authentication Context**.
   You must use the **Post Authentication Context** because the policy used in this tutorial must evaluate what group a user is in. This information is only available after enters their credentials.
2. In the **Default Policy** field, select `Deny Policy`.
   This selection grants access to users by default, and only denies access when the policy conditions evaluate to true.
3. In the **Deny Policy** field, select the policy you created the previous steps.
4. Click **Update** to save the policy context record.

Adaptive Authentication for Trusted Mobile Apps

Access your ServiceNow from untrusted networks by using the Now Mobile app.

Adaptive Authentication for Trusted Mobile Apps enables users to access the ServiceNow instance using the Now Mobile app. The instance is protected behind a trusted IP network boundary.

The following are some of the scenarios where you need access to the instance while you are outside your network.

As an employee, you require access the Now Mobile app and Employee Service Center (portal) from outside the network. The Trusted mobile app filter enables you to identify the incoming request originating from a trusted Now Mobile app that is linked to a user account.

As an admin, you can configure the policy. The policy enables users to register their mobile devices and access the instance on a trusted network.

As a user, you can register your mobile device by scanning the QR displayed on the instance. After the registration, you can log in to your instance from the Now Mobile app, using your credentials.

Note: To register the mobile device, you must be on a trusted network. After the registration, you can log in to the instance from other networks as well.

Features

As an admin, you can use the following features:

- Configure policies by using trusted app filter criteria.
- Create policy conditions with a new filter criteria for trusted mobile apps.
- Support all login methods on the trusted mobile app for the users (Local login, SAML, OIDC, and MFA).
- Revoke a trusted device.
- Add security events for all device registration or revocation actions. The events can be used to notify the user.
- Support security events for signature failures, cookie validation failures, invalid tokens, invalid headers, or query parameters.
- Control the maximum number of registered devices.
- Prevent a new device registration, unless the user removes an existing-paired device.
- Capture the last time that a device was used.

Activate Trusted Mobile App

Activate Adaptive Authentication with Trusted Mobile app by using the authentication policy and filter conditions.

Before you begin

Make sure the Adaptive Authentication (*com.snc.adaptive_authentication*) plugin is installed.

Role required: admin

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Authentication Policies** > **Properties**.
2. On the Adaptive Authentication Properties page, enable the following properties:
   - Enable the Authentication Policy (*glide.authenticate.auth.policy.enabled*)
   - Enable Device Trust Flow (*glide.authenticate.preauth.allow.trusted.device*)

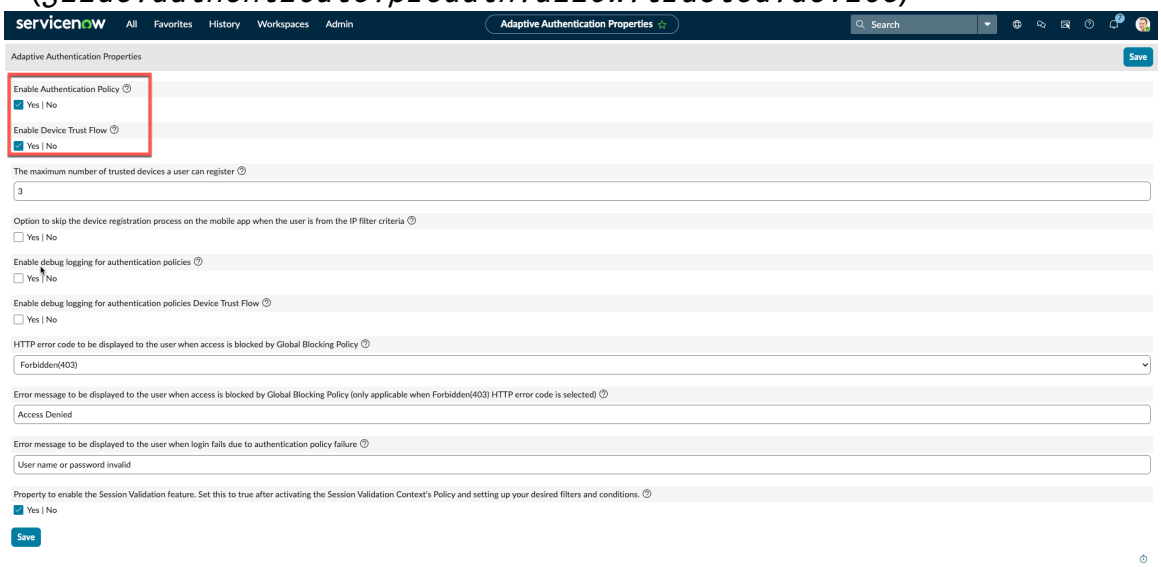

Note: To disable the device trust flow property, you must remove the conditions with trusted mobile filter. Otherwise, an error message is displayed to remove the conditions.

3. Navigate to **All** > **Adaptive Authentication** > **Auth Policy Contexts** > **Pre Authentication Context**.
4. Define the conditions in the Pre Authentication context.
   For more information, Pre authentication context.

   Note: By default, the policy condition is **Deny Policy**. You can change to **Allow Policy**. These policies are direct opposites.

   - With the **Allow Policy**, all users are denied access by default, and it only allows access when the allow access policy conditions are true.
   - With the **Deny Policy**, all users are allowed access by default, and it only denies access when the deny access policy conditions are true.

   In the Policy Input, the policy input **Trusted Mobile App** is a policy input for the trusted mobile app.

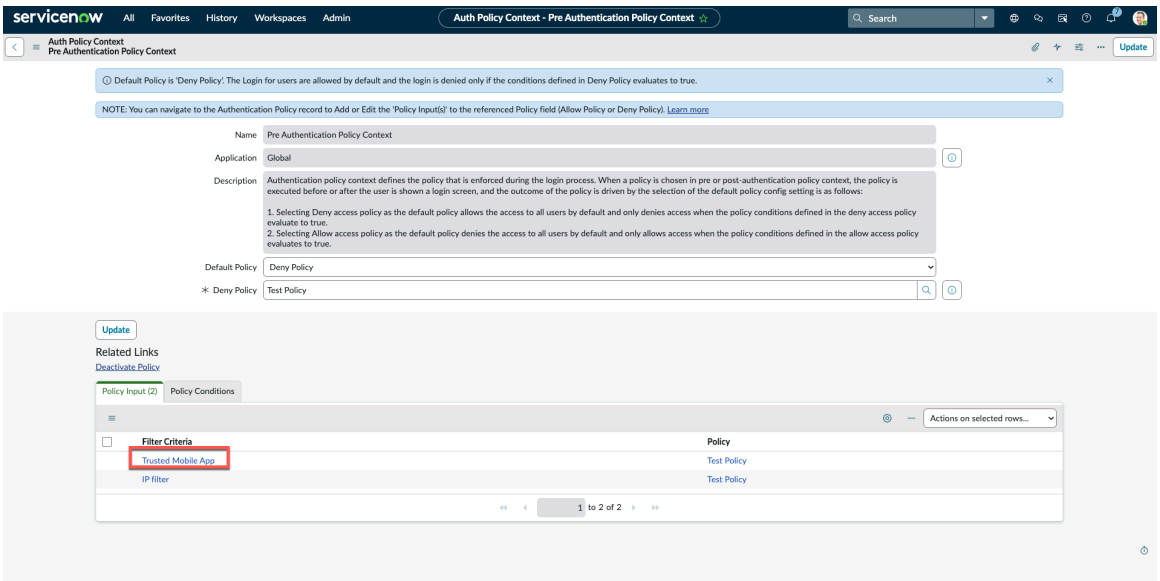

5. In the Policy Conditions, create the condition by clicking **New**.
6. On the form, fill the fields.

Filter Condition form

Field	Description
Label	Name of the condition
Description	Description of the condition
Application	Application scope for this record.
Condition	Conditions based on AND and OR. Because the Auth Policy is **Allow Policy**, the condition for Trusted Mobile App is set true in the example shown on the image.

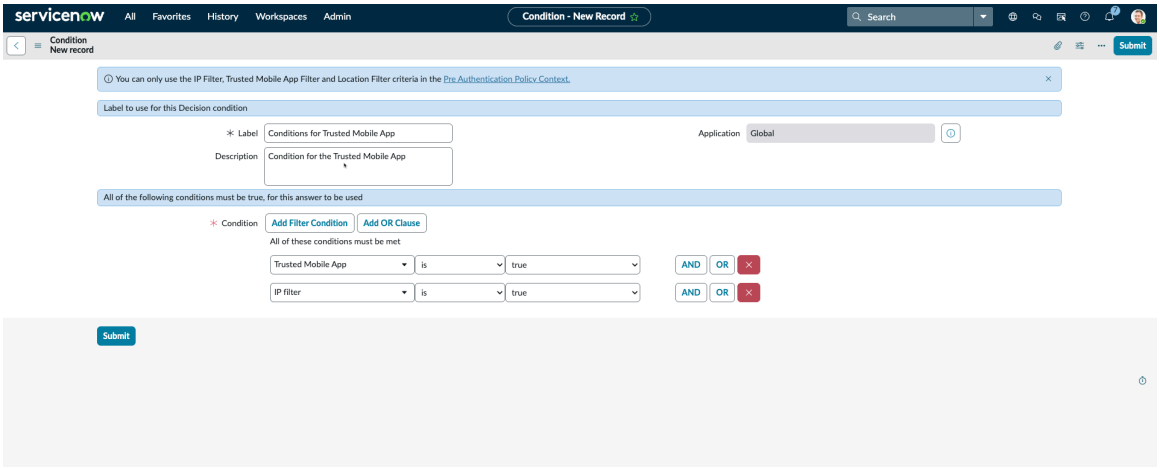

7. Click **Submit**.

Result

The policy input and filter conditions are created for the Trusted Device feature. Users can proceed with using the Trusted Device feature to access the ServiceNow instance from untrusted networks by using the Now Mobile app. For more information, see Register a trusted device.

Register a trusted device

Register a trusted device to access the ServiceNow instance outside the network.

Before you begin

You must be in the trusted network to perform the trusted device registration.

Role required: none

Procedure

1. Navigate to one of the following menu options:
   - On the ServiceNow AI Platform, navigate to **All** > **Self-Service** > **My Profile**.

     Note: You can also access your profile by clicking your user name in the instance header.

   - On the Now Support, click the profile and select **Trusted Device**.
2. In your user profile, click **Register a trusted mobile device** in the Related Links section. The Register Trusted Devices page is displayed.
3. Register a new device by clicking **Add a new trusted mobile device**.

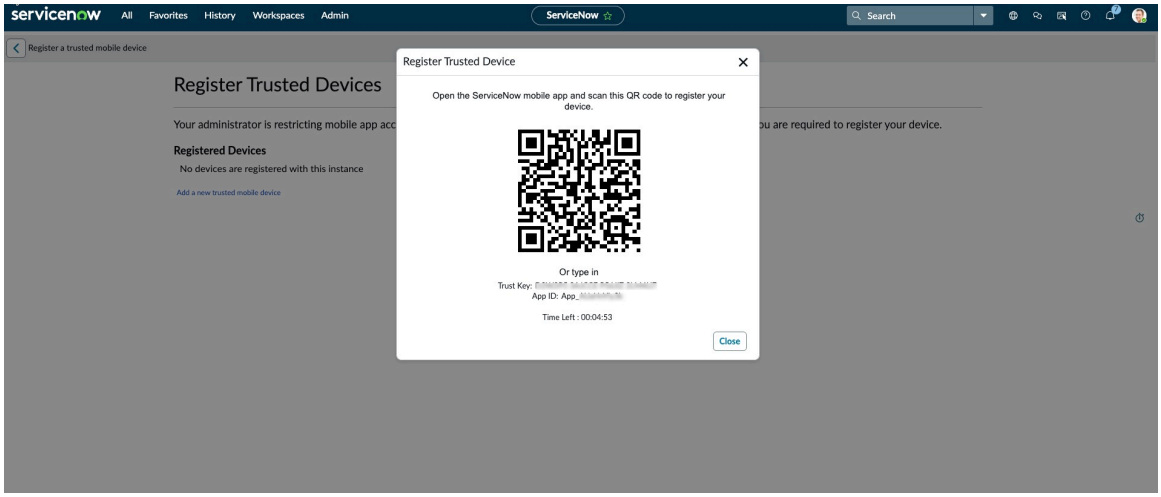

4. In the Device Registration screen on your ServiceNow mobile app, tap the **Scan QR code** button and scan the QR code that is displayed on your laptop or desktop.

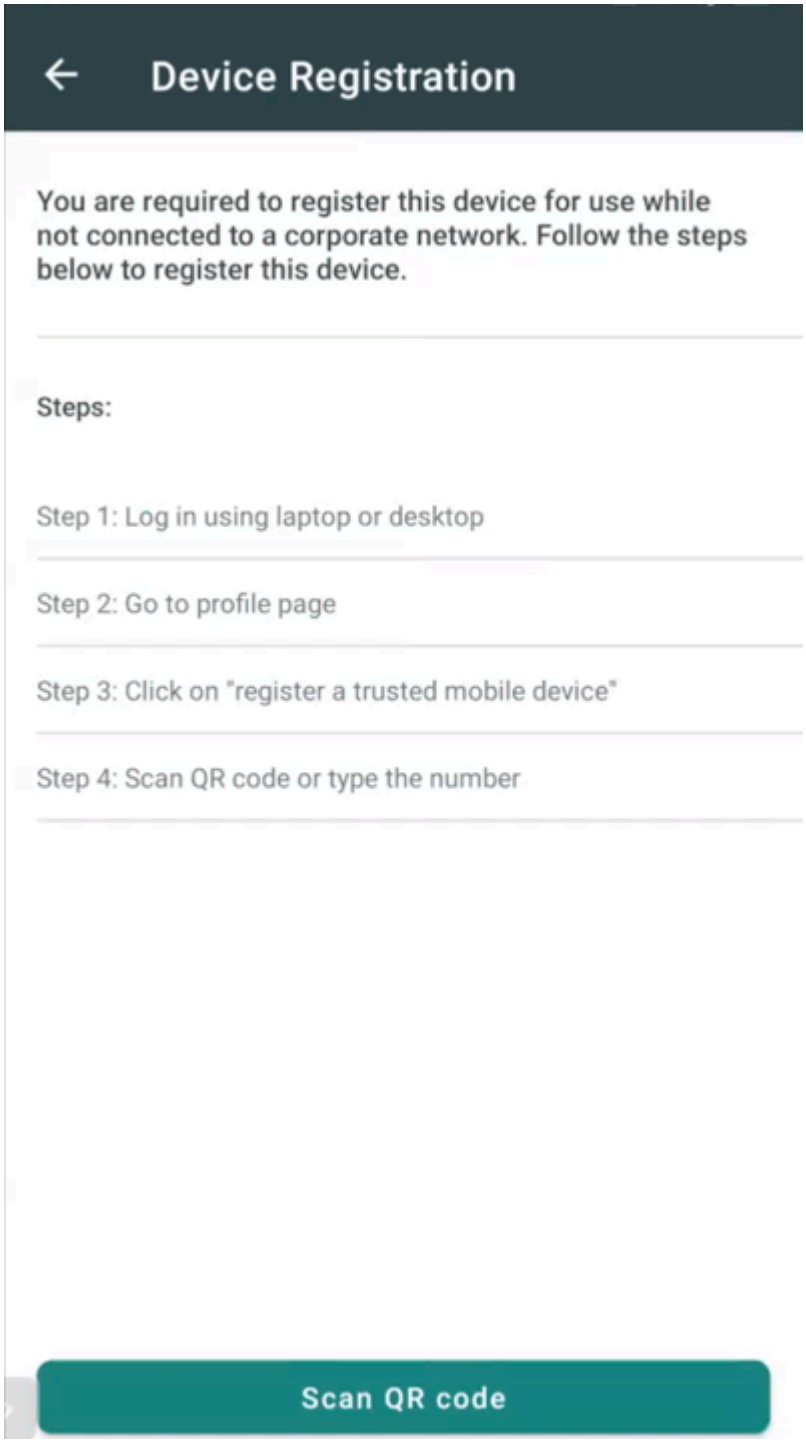

The Registration process will complete and you will be prompted to your login page to complete your authentication.

5. Specify your credentials and log in to the Now Mobile app.

The mobile home screen is displayed.

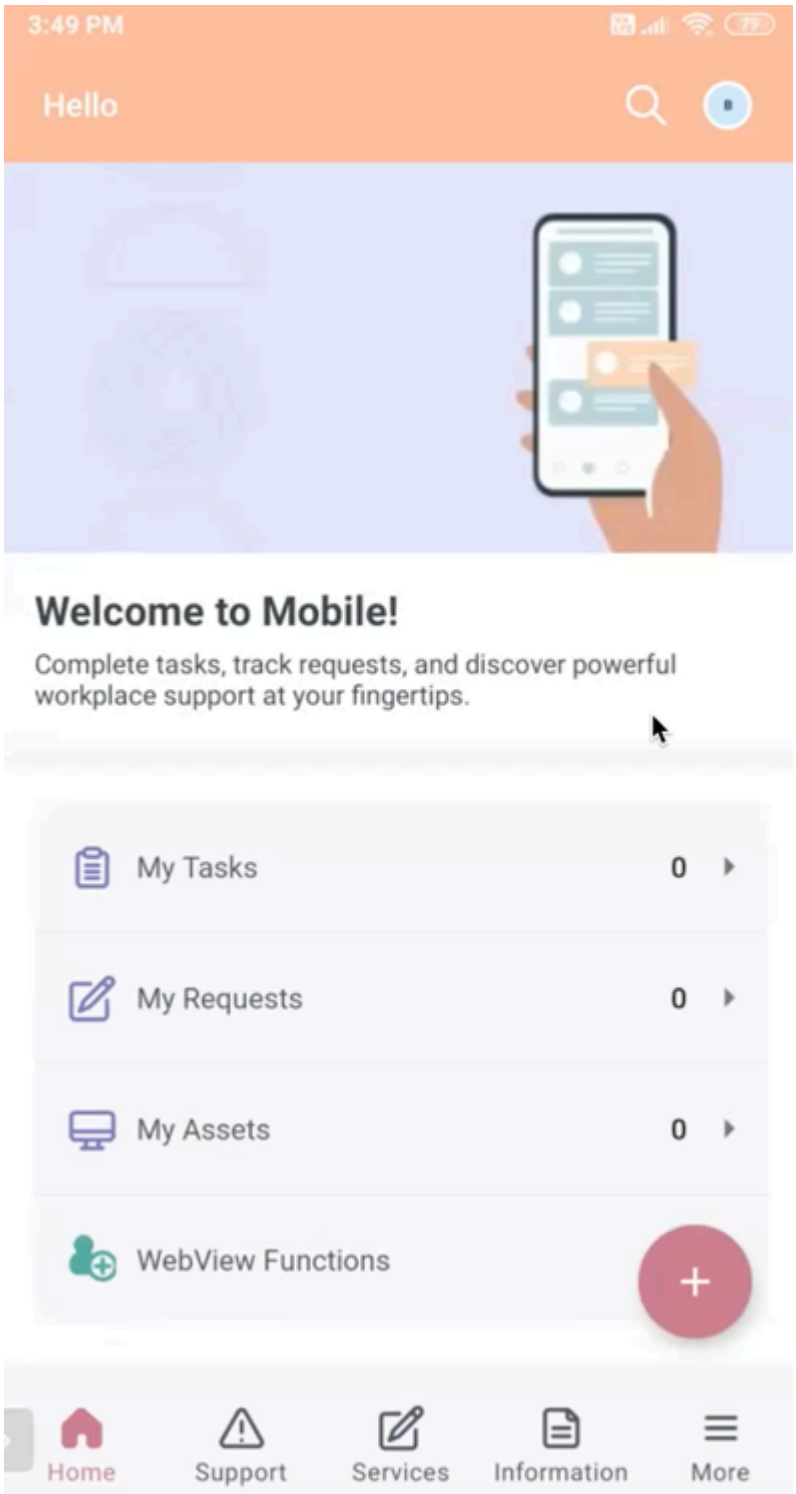

Result

In the Register Trusted Devices page on your laptop or desktop, the registered device is displayed. You can use the icon if you want to remove the registered device from the page.

What to do next

Navigate to **All** > **Adaptive Authentication** > **Device Trust** > **Device Registration** to view all the details of the registered device.

Manage your trusted device

Manage your trusted device from the Trusted Device registration page.

Before you begin

Role required: none

Procedure

1. Navigate to one of the following menu options:
   - On the ServiceNow AI Platform, navigate to **All** > **Self-Service** > **My Profile**.

     Note: You can also access your profile by clicking your user name in the instance header.

   - On the Now Support, click the profile and select **Trusted Device**.
2. In your user profile, click **Register Trusted Device** in the Related Links section.
   The Register Trusted Devices page is displayed.
3. On the Register Trusted Devices page, remove the device by clicking the icon.

   Note: The deleted mobile device must be registered again to access the instance.

Registration details of registered devices

View the details of devices that are registered with your ServiceNow instance.

To view all the details of the registered device with the instance, navigate to **All** > **Adaptive Authentication** > **Device Trust** > **Device Registration**.

Use the filter to identify the device. Click the **AppID** field to know more about the registered device.

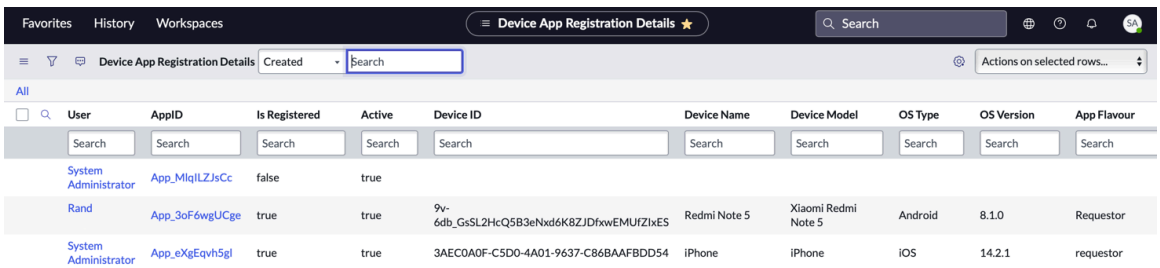

The following details are displayed for the device:

- AppID
- Device ID
- User
- Device Name
- Device Information such as:

- OS Type
- App Flavour
- App Version
- OS Version
- Device Model

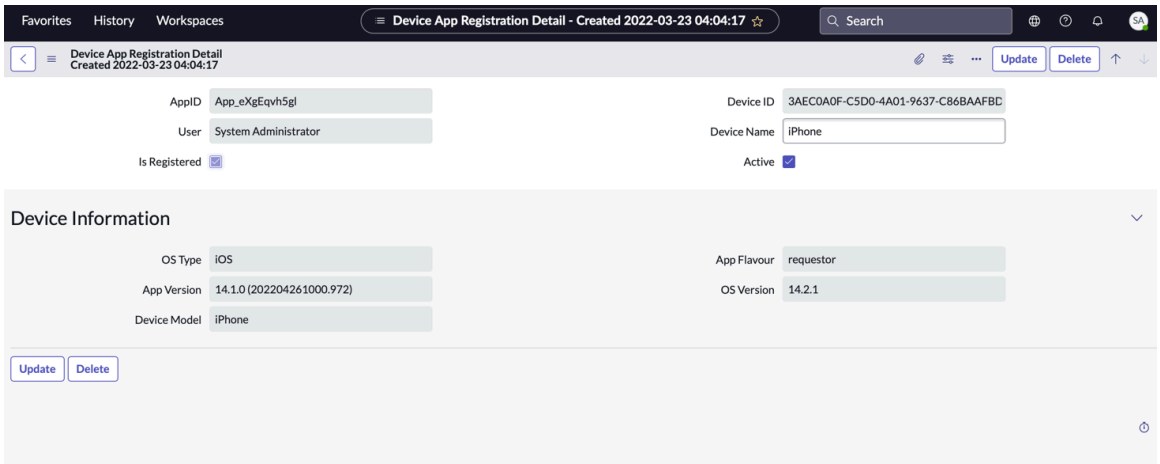

Trusted Mobile App troubleshooting

Review these troubleshooting scenarios to resolve issues with Trusted Mobile App.

The following sections describe some troubleshooting scenarios and ways to troubleshoot.

Time difference between instance and mobile (clock-skew)

If there is a clock-skew, you can adjust or reset the clock on your mobile device for using the Trusted device.

QR code expiry within five minutes

The QR code for Trusted device feature expires within five minutes. In such a case, you can click for a new QR code.

Note: The QR code expiry can be extended. To extend the code expiry, contact your admin.

System property has changed

If the system property has changed, verify the registration based on the system property for the valid credential that can be used for login.

Maximum device registered

If the maximum number of device registrations is reached, remove the existing device or contact your admin to change the maximum device count.

API authentication

Authentication configurations for API.

API based Authentication in ServiceNow®'s validates the identity of a user who accesses an instance, and then authorizes the user to features that match the user's role or job function while making an API call to the ServiceNow® instance.

Following are the types of API Authentication in ServiceNow®:

- Certificate based authentication
- OAuth
- Token-based authentication

Certificate based authentication

Certificate-based authentication lets you mutually authenticate inbound API requests using certificates from a trusted Certificate Authority (CA).

Certificate-based authentication for Inbound web services

Authenticate inbound requests to ServiceNow SOAP and REST APIs. To set up mutual authentication for inbound web services, see Set up Certificate-based authentication.

OAuth

OAuth based authentication validates the identity of the client that attempts to establish a trust on the system by using an authentication protocol.

OAuth 2.0 - Open Authorization is the industry-standard protocol for authorization, that ocuses on client developer simplicity while providing specific authorization flows for web applications, desktop applications, and mobile devices.

Inbound

Create an endpoint for external clients that want to access your instance. This creates an OAuth client application record and generates a client ID and client secret that the client needs to access the restricted resources on the instance. For more information see, OAuth inbound.

Token-based authentication

Token-based authentication for inbound REST APIs configuration using API Key or HMAC.

Support API tokens for REST API end points so that the ServiceNow® for the user authentication.

Token-based authentication for inbound REST APIs configuration can be performed on the ServiceNow® instance with the API Key or HMAC token.

API Key and HMAC Authentication for inbound REST APIs

Support API tokens for REST API endpoints so that the ServiceNow® user name and password isn't visible in the webhook URL.

Enable API key-based authentication to securely authenticate inbound webhook URL.

To use the API Key and HMAC Authentication, you must install the (Plugin: `com.glide.tokenbased_auth`) in the ServiceNow® instance.

Warning: Use **POST** request when submitting any sensitive information to the server.

Installing API Key and HMAC Authentication has dependencies on the following plugins:

- REST API Auth Scope Plugin (`com.glide.rest.auth.scope`)
- REST API Access Policy Plugin (`com.glide.rest.policy`)
- Authentication scope (`com.glide.auth.scope`)

Benefits

API Key and HMAC Authentication for inbound REST APIs enables:

- Ability to specify API key or HMAC token for REST API authentication.
- Ability to associate a user account with the API key or HMAC token.
- Ability to specify a token as query parameter or header within the REST API call.
- Ability to associate authentication scope with API key or HMAC token configurations so that API keys can only be used to invoke APIs associated with a particular scopes.
- Ability to associate an API key or HMAC token configuration with an authentication profile that can be used in API access policies.

Activate API Key and HMAC Authentication

You can activate the plugin API Key and HMAC Authentication (`com.glide.tokenbased_auth`) in your ServiceNow® instance.

Before you begin

Role required: admin.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the API Key and HMAC Authentication plugin (`com.glide.tokenbased_auth`) using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Configure API key - Token-based authentication

Configure an API key to support authentication for REST API endpoints.

Before you begin

Role required: admin

Plugin required: API Key and HMAC Authentication (`com.glide.tokenbased_auth`)

Procedure

1. Create an inbound authentication profile.

   a. Navigate to **All** > **System Web Services** > **API Access Policies** > **Inbound Authentication Profiles**.

   b. Select **New**.
   The system displays the message `What kind of authentication profile?`

   c. Select **Create API Key authentication profiles**.

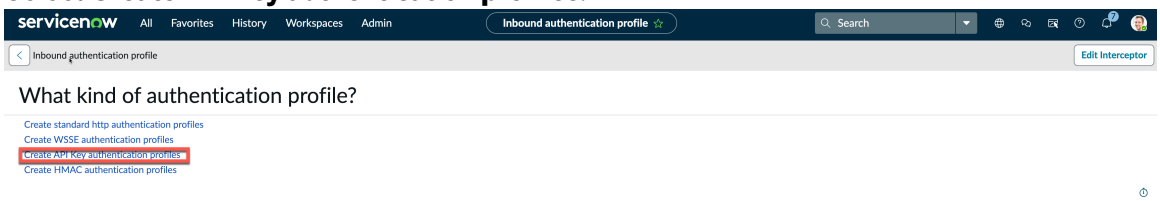

   d. On the form, fill in the fields.

API Key authentication profiles

Field	Description
Name	Name to identify the authentication policy.
Application	Scope of the authentication policy.
Auth Parameter	Select the auth parameter for the authentication request. You can select the default options or create a new auth parameter: ▪ x-sn-apikey: Auth Header ▪ x-sn-apikey: Query Parameter header

Note: The selected option has to be defined in the REST call in the Auth Header or Query Parameter.

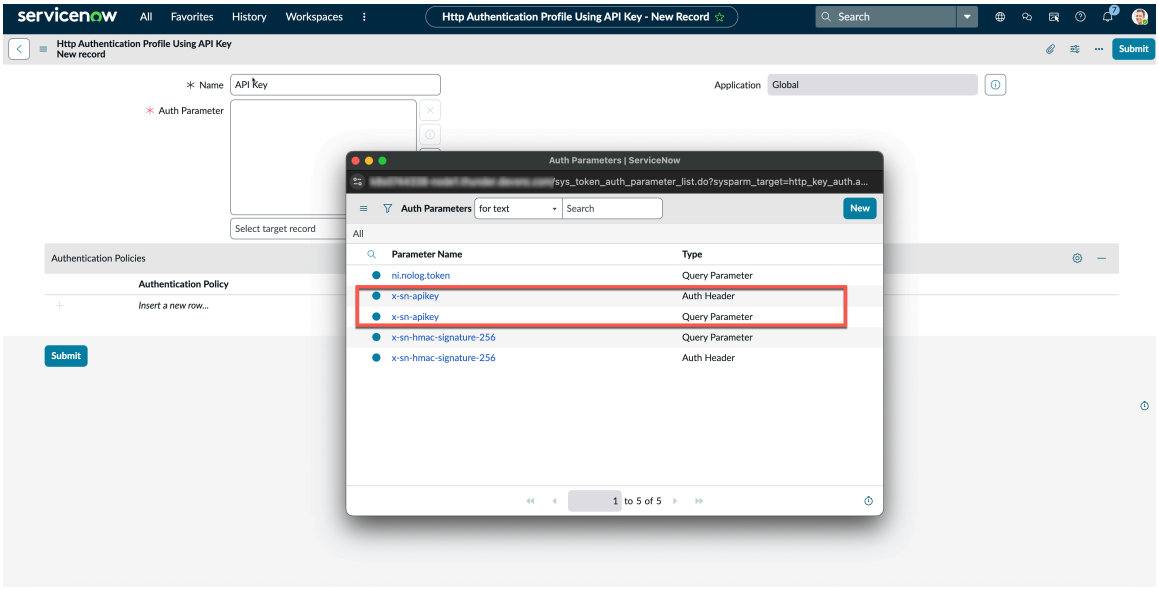

Note: If you wish to add a prefix for the API-Key, Open the Auth Parameter file that you have selected and specify the **Prefix** field.

e. Submit the form.

2. Create a REST API key.

a. Navigate to **All** > **System Web Services** > **API Access Policies** > **REST API Key**.

b. Select **New**.

c. On the form, fill the fields:

API Key

Field	Description
Name	Name to identify the REST API Key
Description	Description for the REST API Key.
Active	Status of the REST API Key.
User	User associated to the REST API Key. Use the look-up icon to select the user.
Auth Scope	Option to add auth scope to manage the authority of the REST API Key.
Token	The REST API key generated by the ServiceNow AI Platform. Copy the key to use as part of the REST API call within the Header or Query parameter.
Expiry	Time after which the credential is no longer valid. Empty value means no expiration.

Field	Description
	Note: For more information, about expiry of token, see Cleaning up token Expiry .

d. Submit the form.

e. Open the record that was created to view the token generated by the ServiceNow AI Platform for the user.

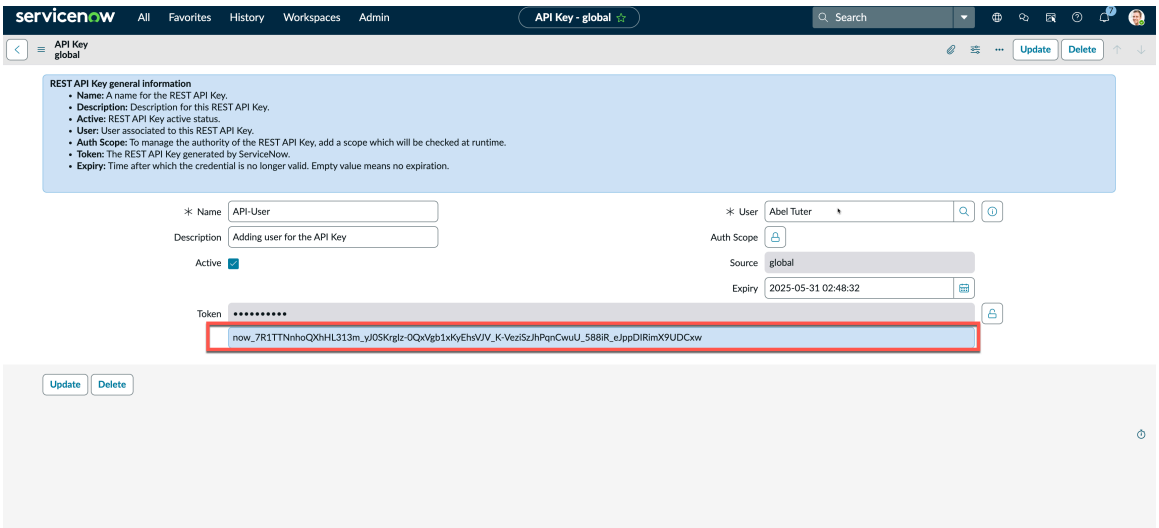

3. Create a REST API Access policy.

a. Navigate to **All** > **System Web Services** > **REST API Access Policies**.

b. Select **New**.

c. On the form, fill in the fields.

API Access Policies

Field	Description
Name	Unique name of the API access policy.
Active	Option to make the API access policy active.
REST API	The REST API to which the access policy is applied. For example, **Attachment API**.
REST API PATH	API path of the REST API. This field is auto-populated based on the selected REST API. For example, **now/attachment**.
HTTP Method	Method used for interacting with the API. This field is auto-populated based on the selected REST API.

Field	Description
Version	Version of the API. For example, **v1**. This field is auto-populated based on the selected REST API. **Note:** If you want to create an authentication policy for all versions of a REST API, you must create individual policies for each version.
Resource	Child resource of the REST API. This field is auto-populated based on the selected REST API. For example, **/now/attachment**
Application	Scope of the application.
Global	Enable this field to apply auth policy to all methods, versions, and resources for the API. **Note:** Token Based Auth isn't allowed in the Global REST API Policy.
Apply to all methods	Enable this field to apply the auth policy for the API to all the methods, versions, and resources for the API.
Apply to all resources	Enable this field to apply the auth policy for the API to all the versions.
Apply to all versions	Enable this field to apply the auth policy for the API to all the resources.

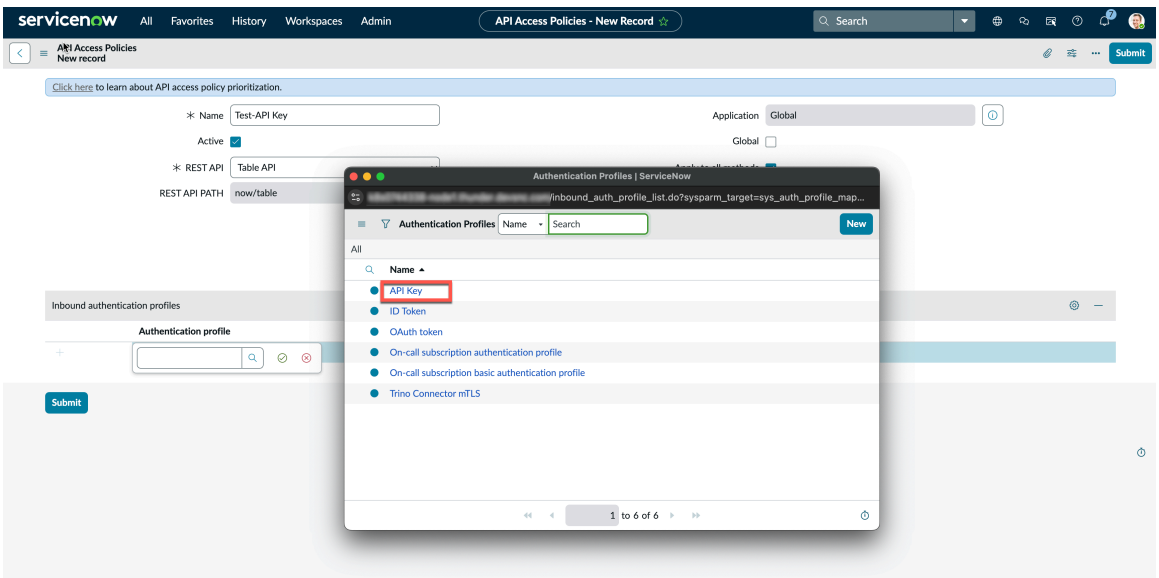

d. Add the API Authentication profile that was created.

e. Submit the form.

You can send the REST API call with the x-sn-apikey (token) that was generated by the ServiceNow AI Platform during the API Key creation within the Header or Query parameter based on the configuration for authentication.

Warning: Use **POST** request when submitting any sensitive information to the server.

Configure HMAC - Token-based authentication

Configure HMAC to support authentication for REST API endpoints.

Before you begin

Role required: admin

Plugin required: API Key and HMAC Authentication (`com.glide.tokenbased_auth`)

Note: While configuring the HMAC, ensure the name of the script that starts with HMAC to populated in the script list view.

Procedure

1. Create an HMAC configuration.

   a. Navigate to **All** > **System Web Services** > **API Access Policies** > **HMAC Configuration**.

   b. Select **New**.

   Note: You can also use the **Default HMAC SHA256 Base64 Encoding** which is created when installing the plugin.

   c. On the form, fill the fields:

HMAC Configuration

Field	Description
Name	Name for the HMAC configuration.
Application	Scope of the configuration.
Description	Detailed description about the configuration.
Hash Algorithm	Choose the Hash Algorithm. Options available: ▪ HmacSHA256 ▪ HmacSHA384 ▪ HmacSHA512
HMAC util script	Utility script for HMAC.

Field	Description
	Note: If you validate HMAC authentication using Request body, Time stamp and Secret with no Key Id, then do the following: ▪ Create a customized script include as HMAC util script ▪ Use that script include util to interpret the passed in timestamp, and request body. For every shared secret created in the ServiceNow® instance, there is a key ID. You need to configure the key ID into Default key id in HMAC auth profile.

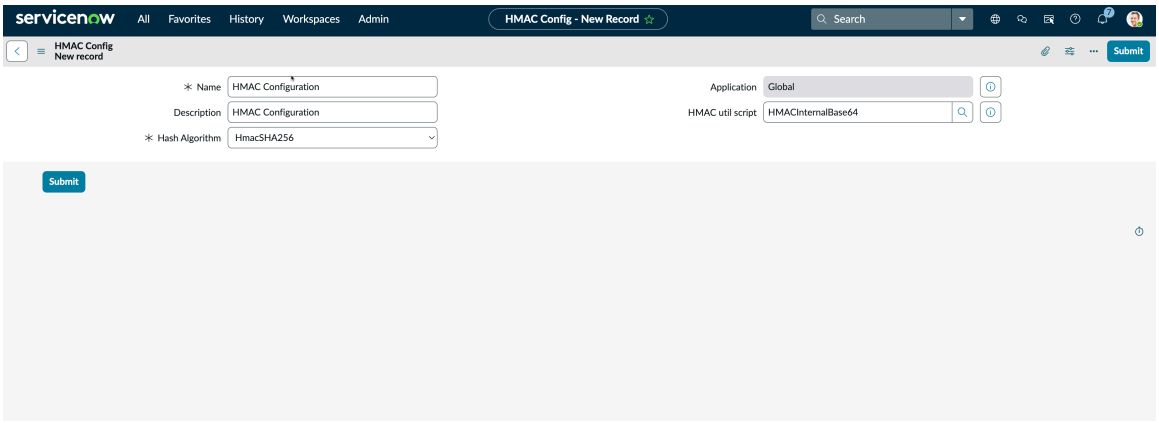

d. Submit the record.

2. Create an inbound authentication profile.

a. Navigate to **All** > **System Web Services** > **API Access Policies** > **Inbound Authentication Profiles**.

b. Select **New**.
The system displays the message `What kind of authentication profile?`

c. Select **Create HMAC authentication profiles**.

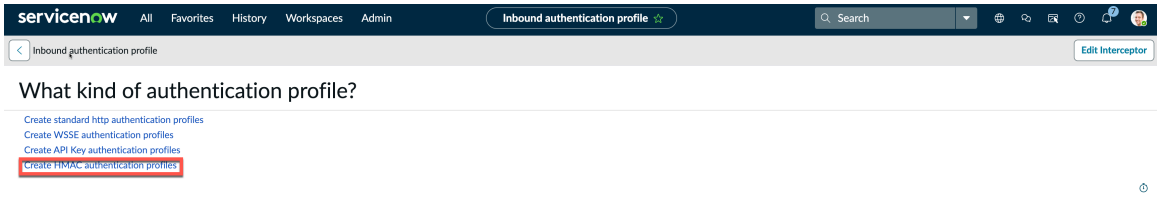

d. On the form, fill in the fields.

HMAC authentication profiles

Field	Description
Name	Name to identify the authentication policy.
Application	Scope of the authentication policy.
Configuration for HMAC	Select the HMAC configuration that was created.
Auth Parameter	Select the auth parameter for the authentication request. You can select the default options or create a new auth parameter: ▪ x-sn-hmac-signature-256: Auth Header ▪ x-sn-hmac-signature-256: Query Parameter
Default Key ID of Shared Secret	The token information that can be updated in this field for using HMAC.

Note: The selected option has to be defined in the REST call as part of the Auth Header or Query Parameter.

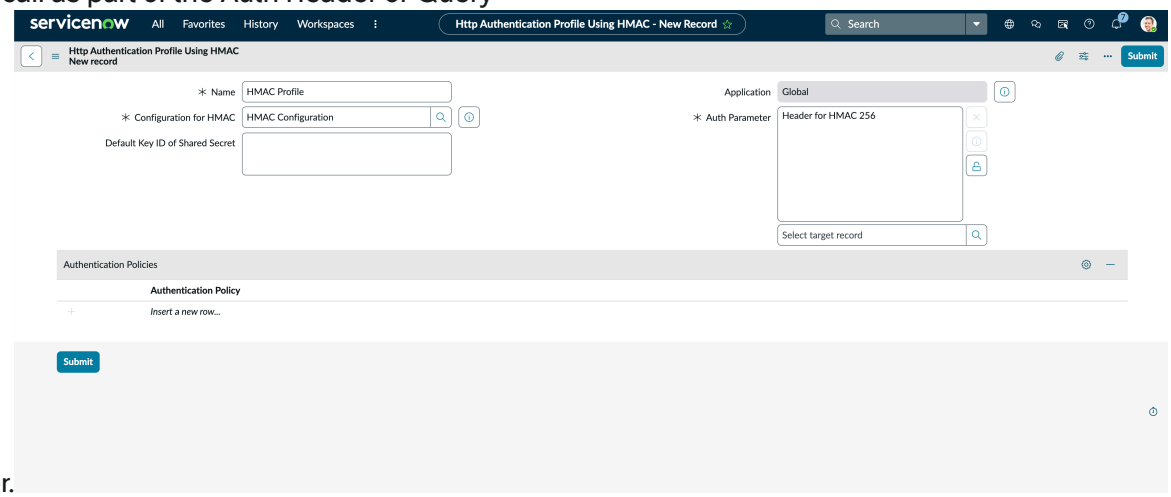

e. Submit the form.

3. Create an HMAC secret.

a. Navigate to **All** > **System Web Services** > **API Access Policies** > **REST API HMAC Secret**.

b. Select **New**.

c. On the form, fill the fields:

REST API HMAC Secret

Field	Description
Name	Name to identify the REST API HMAC secret.
Description	Description for the REST API HMAC secret.
Active	Status of the REST API HMAC secret.
User	User associated to the REST API HMAC secret. Use the look-up icon to select the user.
Key ID	Key ID that must be sent as part of the REST call. Key ID is generated after submitting the form.
Shared Secret	Shared secrets of the user. For example, the password.
Source	Source of the record.
Expiry	Time after which the credential is no longer valid. Empty value means no expiration. **Note:** For more information, about expiry of token, see Cleaning up token Expiry .

d. Submit the form.

e. Open the record that was created.

Find the Key ID generated by the ServiceNow AI Platform for the user.

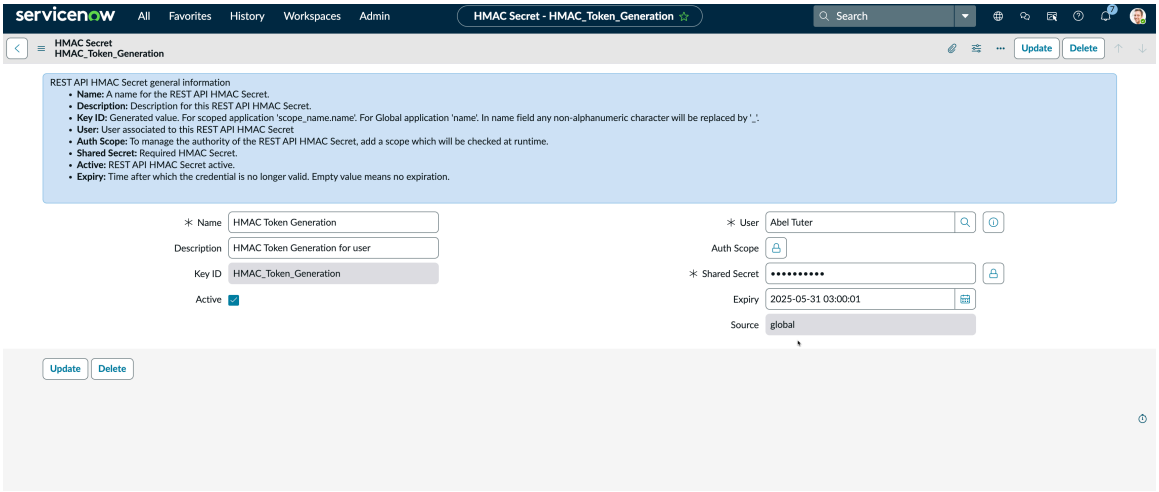

Note: You can add the Key ID that was generated during the Key ID in the Authentication Profile that was created for HMAC if you don't want to specify the Auth or Query parameter for the API call.

4. Create a REST API access policy.

a. Navigate to **All** > **System Web Services** > **REST API Access Policies**.

b. Select **New**.

c. On the form, fill in the fields.

API Access Policies

Field	Description
Name	Unique name of the API access policy.
Active	Option to make the API access policy active.
REST API	The REST API to which the access policy is applied. For example, **Attachment API**.
REST API PATH	API path of the REST API. This field is auto-populated based on the selected REST API. For example, **now/attachment**.
HTTP Method	Method used for interacting with the API. This field is auto-populated based on the selected REST API.
Version	Version of the API. For example, **v1**. This field is auto-populated based on the selected REST API.

Field	Description
	Note: If you want to create an authentication policy for all versions of a REST API, you must create individual policies for each version.
Resource	Child resource of the REST API. This field is auto-populated based on the selected REST API. For example, **/now/attachment**
Application	Scope of the application.
Global	Enable this field to apply auth policy to all methods, versions, and resources for the API. **Note:** Token Based Auth isn't allowed in the Global REST API Policy.
Apply to all methods	Enable this field to apply the auth policy for the API to all the methods, versions, and resources for the API.
Apply to all resources	Enable this field to apply the auth policy for the API to all the versions.
Apply to all versions	Enable this field to apply the auth policy for the API to all the resources.

d. Add the API authentication profile that was created.

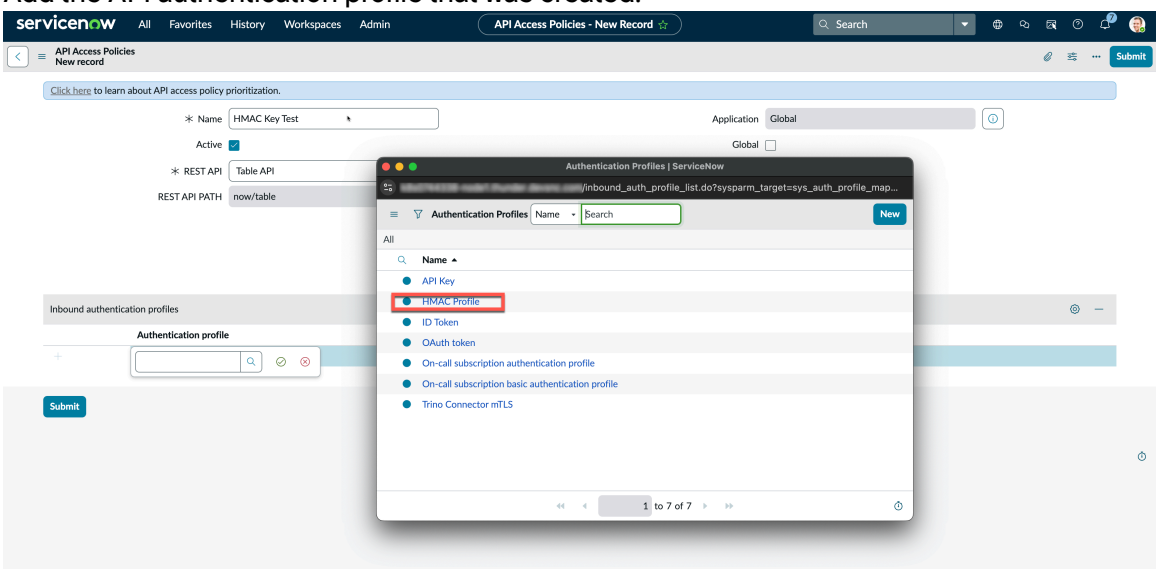

e. Submit the form.
You can send the REST API call:

- With the x-sn-hmac-signature-256 that was generated by ServiceNow® during the API Key creation within the Header or Query parameter based on the configuration for authentication.
- With pre-request Script with Shared Secret specified as part of the request.

> ⚠ **Warning:** Use **POST** request when submitting any sensitive information to the server.

Cleaning up token Expiry

Details about how to clean up token expiry by using different system properties.

Once the API Key or HMAC secret expires, the details is retained for 7 days. A scheduled job `Clean expired Token Auth Credentials` deletes the expired API Keys or HMAC Secrets after that. You can configure this behavior and control the deletion using the following properties:

- **com.snc.platform.security.token.auth.cleanup**: Use this property if you want to delete expired API Keys and HMAC tokens. By default `true`.
- **com.snc.platfrom.security.token.auth.days.expired.api_key.is.kept**: Set the value based on your requirement to determine the number of days you want the keep the expired `API` token in the system. By default 7.
- **com.snc.platfrom.security.token.auth.days.expired.hmac_key.is.kept**: Set the value based on your requirement to determine the number of days you want the keep the expired `HMAC` token in the system. By default 7.

To navigate to the token expiry property, enter `sys_properties.list` in the navigation bar and then search for the following properties. Change the expiry based on your requirement.

Basic authentication

Legacy API authentication method using username and password, with restricted usage and varying behaviour in zBoot and upgraded instances.

Basic authentication is a legacy method for authenticating API requests using a username and password combination. While it remains available for compatibility in certain scenarios, its use is strongly discouraged over token-based authentication methods.

Basic authentication should only be used in limited scenarios where alternative methods aren't viable.

Behaviour across instance scenarios

The behaviour of basic authentication varies depending on how an instance is provisioned or upgraded.

zBoot instances (newly provisioned instances)

In zBoot scenarios, basic authentication is restricted by default to improve security.

- Basic authentication may be blocked unless specific conditions are met
- Access may be limited to users with required roles or permissions
- Enforcement is stricter compared to upgraded instances

These restrictions minimises exposure to insecure authentication methods and encourage adoption of modern alternatives.

Upgraded instances

In upgraded instances, basic authentication may continue to function under certain conditions to support backward compatibility.

- Existing integrations using basic authentication may continue to work
- Restrictions may be applied gradually depending on patch or upgrade level
- Behaviour may differ based on system configuration and applied changes

This approach allows existing users to transition to more secure authentication mechanisms without immediate disruption.

Restrictions and limitations

Basic authentication is subject to several restrictions that vary based on system context and applied updates.

Restrictions	Details
Access restrictions	• Basic authentication may be inactive by default in some scenarios. • Access may only be permitted for specific users or roles. • Certain environments may enforce stricter controls.
Behavioural restrictions	• Behaviour differs between: ○ zBoot instances (stricter enforcement) ○ Upgraded instances (transitional behaviour) • Allowed usage may be time-bound or limited in scope depending on configuration.

API access policy

API access policy defines the permissions and duration of access to an API.

API access policy enables you to restrict access to inbound ServiceNow® APIs based on the authentication type and the specified filter criteria of the access policy.

Following are the three level auth policy support for the non-interactive sessions:

- **REST API**: Admin can define a global API auth policy or API specific auth policy.
- **SOAP API**: Admin can define a global API auth policy or API specific auth policy.
- **System/Export Processor**: Admin can define five base system auth profiles, which can be used to apply auth profile to processors API auth policy or processor specific auth policy.

Note: If you add a global auth policy, then the policy applies for all the REST APIs, SOAP APIs, or System/Export Processors.

Following are the plugins auto-installed for API policies and authentication scope:

- com.glide.rest.policy
- com.glide.soap.policy
- com.glide.processor.policy
- com.glide.rest.auth.scope

You can configure the default Global Blocking Policy or create a custom API access policy according to your security requirements. And apply filter criteria that contain filter conditions or queries that are used as policy inputs for an authentication policy.

The following API access policies are supported in ServiceNow®:

- REST API access policies
- SOAP API access policies

For information about policies related to export processors, see Access policy for System or Export Processors.

REST API access policies

REST API access policies allow you to restrict access to inbound REST APIs based on the authentication type and the specified filter criteria of the access policy.

A REST API, also known as RESTful API is a type of application programming interface (API) that adheres to the guidelines of REST architectural style. REST APIs provide a high degree of flexibility making it prevalent across the web.

Filter criteria contains filter conditions or queries that are used as policy inputs for an authentication policy.

You can configure the default Global Blocking Policy or create a custom API access policy according to your security requirements. For example, you can create a custom API access policy that allows only OAuth 2.0 authentication type from a specified range of IP addresses. Authentication requests of other authentication types and access requests from IP addresses other than the specified IP addresses are denied.

Activate REST API access policy

You can activate the REST API Access Policy plugin (com.glide.rest.policy) if you have the admin role. The application includes demo data and installs related ServiceNow® Store applications and plugins if they are not already installed.

Before you begin

Role required: admin.

About this task

The following items are installed with REST API Access Policy:

- Plugins: com.gilde.auth.profile, com.snc.adaptive_authentication, com.snc.platform.security.oauth
- Tables: sys_api_access_policy, sys_auth_profile_mapping, auth_policy_mapping, inbound_auth_profile, std_http_auth.

For more information, see Adaptive authentication.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the REST API Access Policy plugin (com.glide.rest.policy) using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Create an authentication profile

Create an authentication profile and add one or more authentication policies to the profile. You can also configure the **ID Token** and **OAuth Token** authentication profiles that are available by default.

Before you begin

Role required: admin

Note: You can apply authentication policies, IP range, role-based, user-based, and so on with mutual authentication and customized authentication.

Procedure

1. Navigate to **All** > **System Web Services** > **API Access Policies** > **Inbound Authentication Profiles**.
2. Select **New**.
   The system displays the message. `What kind of authentication profile?`
3. Select **Create standard http authentication profiles**.
4. On the form, fill in the fields.

Standard Authentication Profile form

Field	Description
Name	Name to identify the authentication policy.
Description	Description of the authentication policy.
Active	Option to make the authentication policy active.
Application	Scope of the authentication policy.
Type	Type of the authentication available. You can select **Basic Auth**, **ID Token**, **Certificate based Auth**, or **OAuth**.

Field	Description
OAuth Entity	OAuth Entity profile. This field appears only when **ID Token** or **OAuth** is selected from **Type**.

5. Double-click **Insert a new row**.
6. Select an authentication policy from the list and select the save icon .

Note: Don't select Allow Access Policy or Deny Access Policy. These policies are meant only for user logins.

You can add one or more authentication policies for an authentication profile.

Note:

Authenticate Header [WWW-Authenticate].

When REST API Access Policy is active we will return the most recently mapped authentication profile in the authenticate header. In the case you want the server to return all the authentication schemes, use the `glide.security.response.authenticate.header.auth_profile.first_scheme_on` property and set it to **false**. The response is returned with multiple header. For example:

```
< WWW-Authenticate: BEARER realm="Service-now"
< WWW-Authenticate: BASIC realm="Service-now"
```

Create REST API access policy

Create an API access policy and map an authentication profile to restrict the authentication type for a REST API. For example, you can create an API access policy that allows only ID token authentication for a REST API.

Before you begin

Make sure that an authentication profile is created. For more information, see Create an authentication profile.

Role required: api_service_admin, adaptive_auth_policy_admin

Procedure

1. Navigate to **All** > **System Web Services** > **REST API Access Policies**.
2. Select **New**.
3. On the form, fill in the mandatory fields and submit.

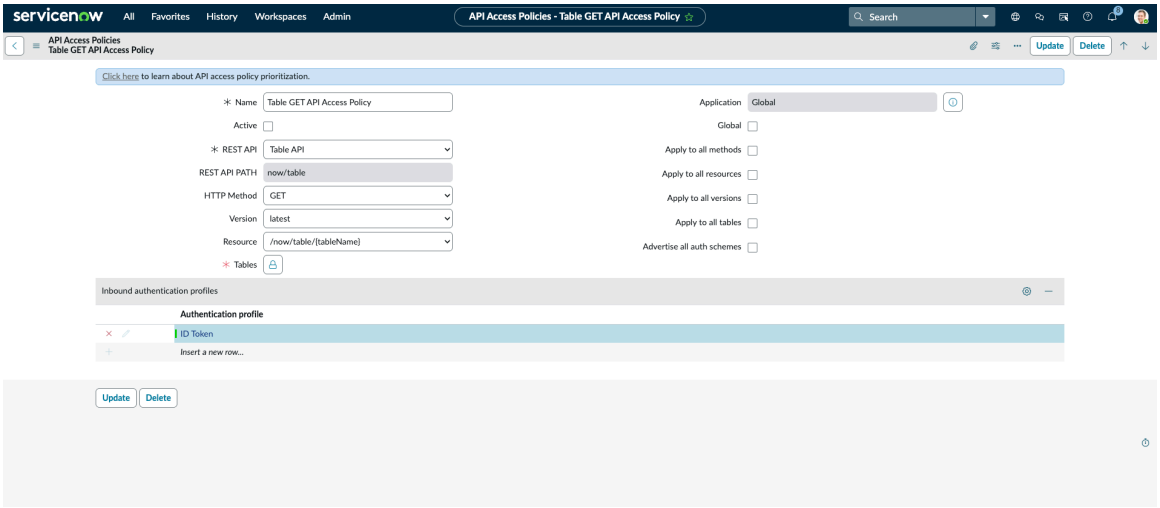

Note: You must reopen the submitted form to populate additional fields.

API Access Policies

Field	Description
Name	Unique name of the API access policy.
Active	Option to make the API access policy active.
REST API	The REST API to which the access policy is applied. For example, **Attachment API**.
REST API PATH	API path of the REST API. This field is auto-populated based on the selected REST API. For example, **now/attachment**.
HTTP Method	Method used for interacting with the API. This field is auto-populated based on the selected REST API.
Version	Version of the API. For example, **v1**. This field is auto-populated based on the selected REST API. **Note:** If you want to create an authentication policy for all versions of a REST API, you must create individual policies for each version.
Resource	Child resource of the REST API. This field is auto-populated based on the selected REST API. For example, **/now/attachment**
Table	The tables to which the API access policy applies. This option only applies to policies for the *Table* API.
Application	Scope of the application.
Global	Option to apply the policy to all methods, versions, and resources for the API.

Field	Description
Apply to all methods	Option to apply the policy to all the methods, versions, and resources for the API.
Apply to all resources	Option to apply the policy to all or the API versions.
Apply to all versions	Option to apply the policy to all or the API resources.
Apply to all tables	Option to apply the policy to all tables. This option only applies to policies for the *Table* API.
Advertise all auth schemes	Determines whether the `WWW-Authenticate` header includes all configured authentication schemes. When set to `false` (default), the header includes only the most recently configured authentication profile in the policy. When set to `true`, the header lists all configured authentication schemes.

Note: To understand more about the API access policy prioritization, see API access policy prioritization.

4. Double-click **Insert a new row**.
5. Select an inbound authentication profile from the list and select the save icon.
   For example, you can add **Basic Auth**, **ID Token**, **Certificate based Auth**, **OAuth** or **WSSE Auth**.

   a. To add one or more inbound authentication profiles, select **New** to create a new profile.

   b. Choose **What Kind of authentication profiles?**.
   - **Create standard http authentication profiles**
   - **Create WSSE authentication profiles**
   - **Create API Key authentication profiles**
   - **Create HMAC authentication profiles**

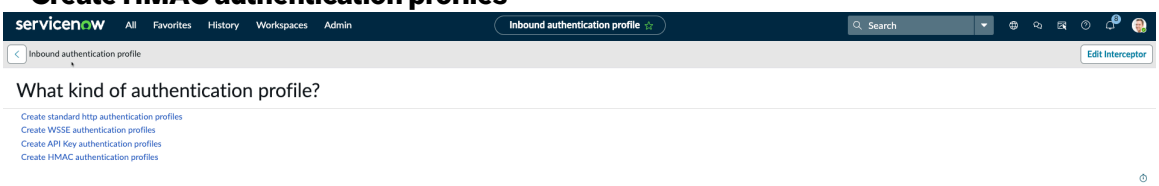

   c. After creating the authentication profile, save the record.

6. Select **Submit** to submit the REST API access policy.

API access policy prioritization

Learn about the policy prioritization logic if there are multiple API access policy configured for your ServiceNow® instance.

API access policies are prioritized based on the type of REST API policy set on your ServiceNow® instance.

The approach is defining different weights for each part of the API such as,method, resource, and version.

The API policy is prioritized for non-global first, then global. In other words, a non-global access policy will always override a global API access policy.

Prioritization logic are as follows:

Prioritization

Fields	Priority	Prioritization Logic
Method, resource, and version	1	If the 3 fields matches with the policy then that policy takes the 1st priority.
Method+ resource	2	If the 2 fields matches with the policy then that policy takes the 1st priority.
Resource + version	3	If the 2 fields along with the field **Apply to all methods** matches with the policy then that policy takes the 3rd priority.
Resource	4	If the field along with the field **Apply to all methods** matches with the policy then that policy takes the 4th priority.
Method + version	5	If the 2 fields along with the field **Apply to all resources** matches with the policy then that policy takes the 5th priority.
Method	6	If the field along with the field **Apply to all resources** matches with the policy then that policy takes the 5th priority.
Version	7	If the field along with the fields **Apply to all methods** and **Apply to all versions** matches with the policy then that policy takes the 7th priority.
Global and Apply to all methods	8	If the fields **Global** is `true` and **Apply to all methods** is `false` then that policy takes the 8th priority.

Prioritization (continued)

Fields	Priority	Prioritization Logic
Global and Apply to all methods	9	If the fields **Global** is `true` and **Apply to all methods** is `true` then that policy takes the 9th priority.

REST API Auth Scope

Use the REST API Auth Scope to provide access to a specific REST API

Earlier, the client without any scope have access to all APIs unless the API is protected by REST API Auth Scope or the client has `useraccount` scope.

After creating REST API Auth Scope record, to access this REST API you must associate the same Auth Scope to the OAuth Entity which should have access to this REST API. For a new OAuth Entity, the default Auth Scope is empty.

Note: Unless you have REST API Auth Scope record, the REST API can be accessed by any valid OAuth Entities.

You must manually link the Auth scope within the OAuth Entity. The `useraccount` is a special scope, if it's associated with an OAuth Entity it can access any API even if you have created a REST API Auth Scope record with a different auth scope.

To learn more about how to use the API Auth Scope in the new Inbound integration experience, see Inbound integrations.

Note:

- After the REST API Auth scope is enabled and added to the auth scope for the REST API, then all the existing OAuth token won't able to access this API anymore unless admin adds this auth scope to the corresponded OAuth entity
- The admin is responsible to making sure the oauth_entity has the right auth scope after to link the auth scope with the REST API.
- OAuth access tokens issued by ServiceNow supports the auth scope.
- OIDC token that is not issued by ServiceNow is validated by ServiceNow.
- OIDC token has its scope from IDP when you require an ID token. Here the auth scope is for ServiceNow instead of third party (IdP).

Configurations for REST API Scope

To configure the REST API Scope, perform the following tasks:

- Create an auth scope
- Link auth scope with the REST API
- Link auth scope with OAuth entity
- Perform OAuth flow to get OAuth access token
- Use the OAuth access token to make the API call

Activate REST API Auth Scope

You can activate the REST API Auth Scope plugin (com.glide.rest.auth.scope) to link the OAuth entity with authentication scopes.

Before you begin

Install the following plugins:

- OAuth 2.0
- REST API Provider
- Authentication scope
- REST API Scope

Note: The *REST API Scope* plugin is added as part of the Tokyo release.

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the REST API auth Scope (`com.glide.rest.auth.scope`) plugin using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

REST API Auth Scope properties and tables

The REST API Auth Scope plugin (com.glide.rest.auth.scope) includes the following system properties, tables, and scripts.

REST API Auth Scope Properties

REST API Auth Scope adds the following system properties.

Properties

Name	Description
com.glide.rest.api.auth.scope.check	This property is used to turn-off the platform level auth scope check. If set to false, during the runtime the auth scope check is skipped whether it is linked or not linked with the REST API. By default this property is set to true. This properties is used when you want to revert to previous release behavior.

Properties (continued)

Name	Description
glide.oauth.token.scope.useraccount	This property is only used when *useraccount* auth scope is deleted and is added back manually by end user. In this case, the sys ID for the *useraccount* is changed. It is required to update this property to the new sys_id. During the runtime, auth scope sys ID is used instead of auth scope name.

REST API Auth Scope Tables

REST API Auth Scope the following tables.

Tables

Name	Description
Authentication Scope (sys_auth_scope)	This table defines the auth scope that can be linked with the REST API and OAuth entity. The auth scope name must be unique and it is global.
REST API Auth Scope (sys_api_access_scope)	This table links REST API with auth scope.

Configure REST API Auth scope

Link the OAuth entity with an auth scope to manage the token to access the REST APIs that are linked with the auth scope.

Before you begin
Install the following plugins:

- OAuth 2.0
- REST API Provider
- Authentication scope
- REST API Auth Scope

Note: The *REST API Auth Scope* plugin is added as part of the Tokyo release.

Role required: admin

Procedure

1. Navigate to **All** > **API Auth Scopes** > **REST API Auth Scope**.
   The REST API Auth Scopes page is displayed.
2. To configure a new REST API Auth Scope, click **New**.
3. On the form, fill in the fields.

REST API Auth Scope

Name	A unique name that identifies the REST API Auth Scope.
Active	Select the check box to make the configuration active.
Application	Read-only application scope.
REST API	The REST API to which the auth scope is applied. For example, the Table API.
Auth Scope	Select the auth scope from the lookup icon.
REST API PATH	API path of the REST API. This field is auto-populated based on the selected REST API. For example, now/table.
HTTP Method	Method used for interacting with the API. Select the method from drop-down list. You can disable the **Apply auth scope to all http methods in this API** field on the form manually to select the method.
REST API Version	Version of the API. For example, v1. This field is auto-populated based on the selected REST API. You can disable the **Apply auth scope to all versions in this API** field on the form manually to select the version.
Resource	Child resource of the REST API. This field is auto-populated based on the selected REST API. For example, /now/table. You can disable the **Apply auth scope to all resources in this API** field on the form manually to select the resources.
Apply auth scope to all http methods in this API	When enabled, applies the auth scope to all the http methods in the API.
Apply auth scope to all versions in this API	hen enabled, applies the auth scope to all versions in the API.
Apply auth scope to all resources in this API	When enabled, applies the auth scope to all resources in the API

4. Click **Submit**.
   Based on the selected REST API and Auth Scope, the APIs retrieves information that is particular to the scope.

Example: Consider creating three REST API Auth Scope for Table API

The first auth scope is mapped to the **Table API** with all the http methods, versions, and resources enabled.

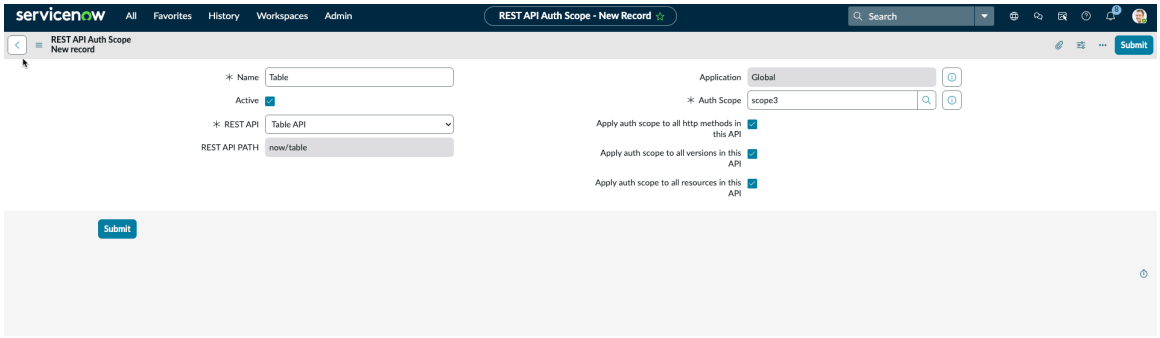

The second auth scope is mapped to the **Table API** with all the versions and resources enabled. But, you choose the HTTP Method, in this example, the **GET** method.

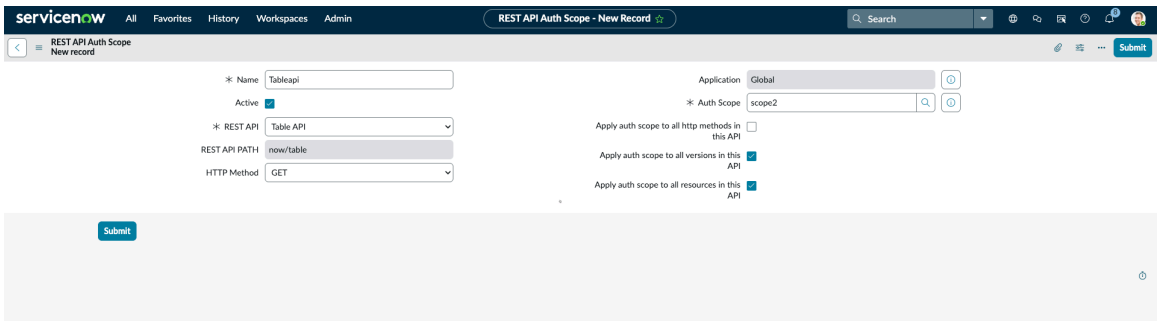

The third auth scope is mapped to the **Table API** without the http methods, versions, and resources enabled. But, you choose the HTTP Method, Version, and Resource manually. In this example, HTTP Method is **GET**, REST API Version is **latest**, and Resource is `/now/table/{tableName}`.

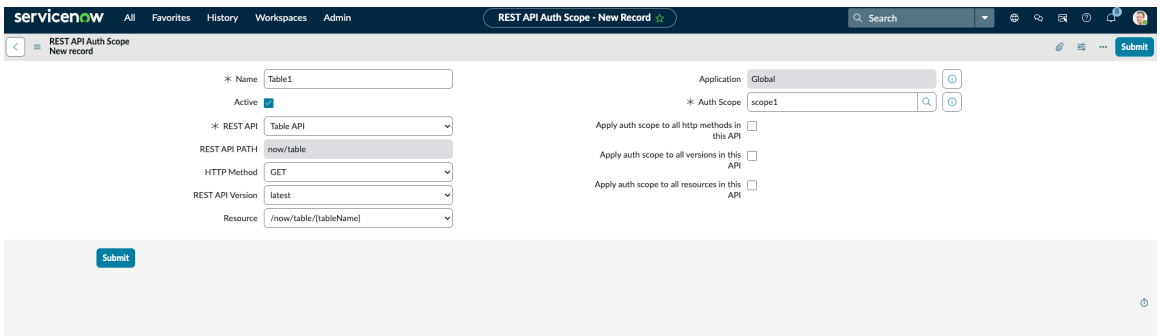

If all these auth scopes are created, you can use **GET** method with all the three scopes, but for **POST**, **PUT**, **DELETE**, or **PATCH** methods only **scope3** can be used.

REST API scope troubleshooting

Troubleshooting actions can help resolve common issues when setting up or running the REST API scope.

Troubleshooting

Issue	Action
REST API is linked with auth scope, however in runtime there is no auth scope check even using Bearer token authentication.	• Make sure the sys_api_access_policy record is active. Runtime ignores inactive records. • Check if property *com.glide.rest.api.auth.scope.check.enable* is set to false. • Check if the OAuth token has `useraccount` auth scope.
REST API is linked with *auth_scope1*, however the access token which has *auth_scope2* is also able to access it.	• Check if this record is active. • Check for this REST, check if any other records, which have the same APIs but different apply methods, versions, or resource.
REST API is linked with auth scope, however in runtime there is no auth scope check for *basicAuth* and *mutualAuth*.	It is expected since the REST API auth scope only applies to the OAuth access token or OIDC token. It doesn't apply BasicAuth, Session Cookie and Certificate based authentication.
REST API call return 403 when using the OAuth access token.	Check for the error message "Missing required api access scope". If found then the auth scope check fails for this REST API
Pre-defined `useraccount` is deleted and not sure to restore.	Export `useraccount` as `xml` from the other instance and import it or create an `useraccount` and modify system property *glide.oauth.token.scope.useraccount* to the newly created sys_id record.

Frequently asked questions

Following are some of the frequently asked question when using the REST API Auth scope:

Can one OAuth token be linked with several auth scopes?

Yes, one `oauth_entity` can be linked with multiple auth scopes, every OAuth token issued by this `oauth_entity` has the same auth scopes.

Can different OAuth tokens with different auth scopes access the same REST API?

Yes, for the same REST API, it may be accessed by different auth scopes. As long as one auth scope is matched, the auth scope returns the results.

Can OAuth access token with `useraccount` auth scope access any REST APIs?

Yes, the `useraccount` has full access to auth scope.

Can OAuth access token OAuth scope be changed dynamically?

Yes, the auth scoped is not hard-coded with the access token in the `oauth_credential` table. Instead auth scope is getting from linked `oauth_entity` during runtime.

Can OAuth token keep same auth scopes after refresh?

Yes, auth scope will not change after token refresh, unless `oauth_admin` modify auth scope linked with `oauth_entity`.

Pre-defined `useraccount` auth scope record is deleted, can a new auth scope with name `useraccount` be created?

Creating a new auth scope with the same `useraccount` doesn't work. In the runtime, it uses the `sys_id` instead of name to do the auth scope check, modify the system property *`glide.oauth.token.scope.useraccount`* to the newly created `sys_id` record.

If admin modify auth scoped linked with `oauth_entity`, are all the existing OAuth access token issued by this OAuth entity changed also?

Yes, the auth scope is not directly linked with the OAuth access token, it is getting from `oauth_entity` during runtime.

Can different OAuth access tokens issued by the same `oauth_entity` have different auth scopes?

No, all access to the token is issued by the same `oauth_entity` and always have the same auth scopes.

Can a user define different auth scopes for a particular endpoint?

No, there is a unique constrain check for a particular REST API endpoint. However for the same REST API endpoint, it may have more than one matched auth scopes.

Is the auth scope check used for BasicAuth also?

No, auth scope check is only OAuth access token and OIDC token, it is not applied for *`basicAuth`* and *`mutualAuth`*

SOAP API access policies

SOAP API access policies allow you to restrict access to inbound SOAP APIs based on the authentication type and the specified filter criteria of the access policy.

The SOAP API access policies enables you to have the ability to apply the inbound authentication profile and API access policies to inbound SOAP and scripted SOAP APIs.

You can leverage ServiceNow API access policies like IP range and role based restrictions to allow or disallow inbound SOAP API calls based on the authentication.

As an admin, you can perform the following actions to apply the policies.

- Activate SOAP API Access Policy and Authentication Profile plugins. For more information, see Activate SOAP API access policy.
- Create SOAP API Access Policies and associate those policies with an authentication profile. For more information, see Create SOAP API access policy and Create an authentication profile.

  Note: The policies are applicable to SOAP table API or scripted SOAP API. Besides the standard `http` and `WSSE` authentication profile.

- Create authentication policies such as IP range, role-based restrictions and associate this policy to the authentication profile. For more information, see Create an API authentication policy.

Activate SOAP API access policy

For SOAP API access policy, install the SOAP API Access Policy (`com.glide.soap.policy`) plugin.

Before you begin
Role required: admin

About this task
The following item is installed with SOAP API Access Policy Plugin: Processor Access Policy (`com.glide.processor.policy`)

Dependent Plugin: Authentication Profile (`com.glide.auth.profile`)

Procedure
1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the SOAP API Access Policy plugin (`com.glide.soap.policy`) using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Create an authentication profile

Create an authentication profile and add one or more authentication policies to the profile. You can also configure the **ID Token** and **OAuth Token** authentication profiles that are available by default.

Before you begin
Role required: admin

Note: You can apply authentication policies, IP range, role-based, user-based, and so on with mutual authentication and customized authentication.

Procedure
1. Navigate to **All** > **System Web Services** > **API Access Policies** > **Inbound Authentication Profiles**.
2. Select **New**.
   The system displays the message. `What kind of authentication profile?`
3. Choose **What Kind of authentication profiles?**.
   - **Create standard http authentication profiles**
   - **Create WSSE authentication profiles**

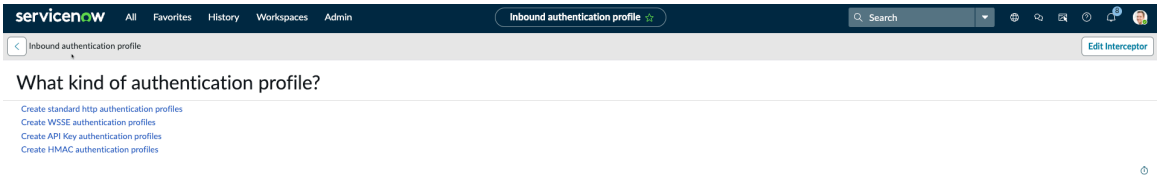

4. On the form, fill in the fields.

Standard Authentication Profile form

Field	Description
Name	Name to identify the authentication policy.
Description	Description of the authentication policy.
Active	Option to make the authentication policy active.
Application	Scope of the authentication policy.
Type	Type of the authentication available. You can select **Basic Auth**, **ID Token**, **Certificate based Auth**, **OAuth**, or **WSSE** (In case of WSSE Authentication profile).
OAuth Entity	OAuth Entity profile. This field appears only when **ID Token** or **OAuth** is selected from **Type**.

5. Double-click **Insert a new row**.
6. Select an authentication policy from the list and select the save icon .

Note: Don't select **Allow Access Policy** or **Deny Access Policy**. These policies are meant only for user logins.

You can add one or more authentication policies for an authentication profile.

When there's a change in the authentication profile, the Authorization header returns a value specific to the changes made at that time. To have the ability to get all the authentication schemes returned in the `WWW-Authenticate` header, you must activate `glide.security.response.authenticate.header.auth_profile.first_scheme_only` to **false**. The response is returned with multiple headers. For example:

```
< WWW-Authenticate: BEARER realm="Service-now"
< WWW-Authenticate: BASIC realm="Service-now"
```

Create SOAP API access policy

Create an API access policy and map an authentication profile to restrict the authentication type for a SOAP API. For example, you can create an API access policy that allows only ID token authentication for a SOAP API.

Before you begin

- Make sure that an authentication profile is created. For more information, see Create an authentication profile.
- Role required: admin

Procedure

1. Navigate to **All** > **System Web Services** > **API Access Policies** > **SOAP API Access Policies**.
2. Click **New**.
3. On the form, fill in the fields.

API Access Policies form

Field	Description
Name	Unique name of the API access policy.
SOAP API	SOAP API to which the access policy is applied. For example, **GetProperty API**.
Application	Scope of the application.
Active	Option to make the API access policy active.

4. In the Inbound Authentication section, double-click **Insert a new row**.
5. Select an inbound authentication profile from the list and click the save icon .
   For example, you can add **Basic Auth**, **ID Token**, **Certificate based Auth**, **OAuth** or **WSSE Auth**.

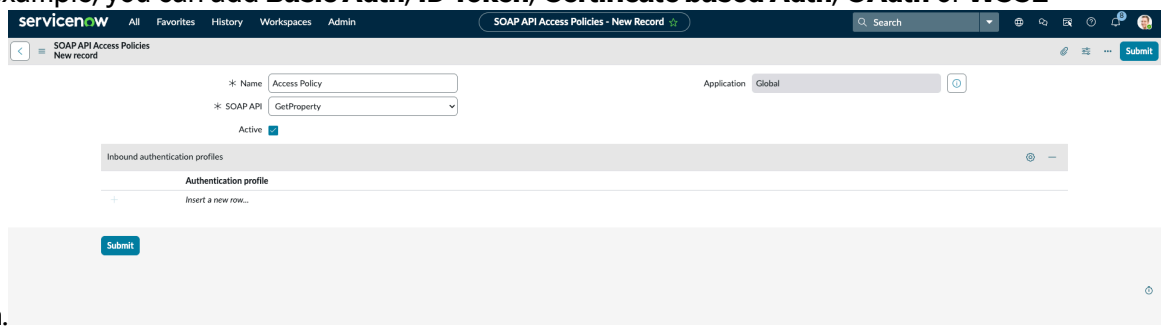

   a. To add one or more inbound authentication profiles, click **New** to create a new profile.

   b. Choose **What Kind of authentication profiles?**.
   - **Create standard http authentication profiles**
   - **Create WSSE authentication profiles**

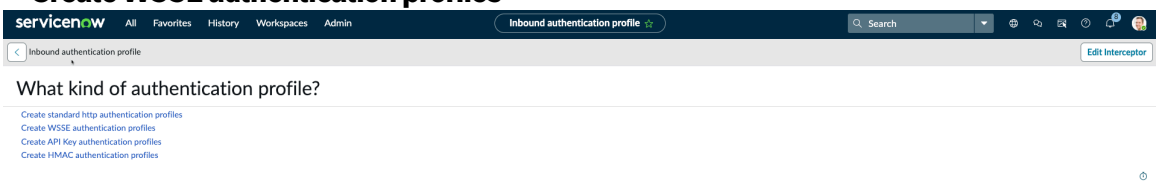

c. To create a WSSE authentication profile, on the form, fill in the fields.

WSSE Authentication Profile

Field	Description
Name	Unique name of the API access policy.
Description	Description of the authentication profile.
Application	Scope of the application.
Active	Option to make the API access policy active.
Type	WSSE Auth as authentication profile WSSE (Web Security).

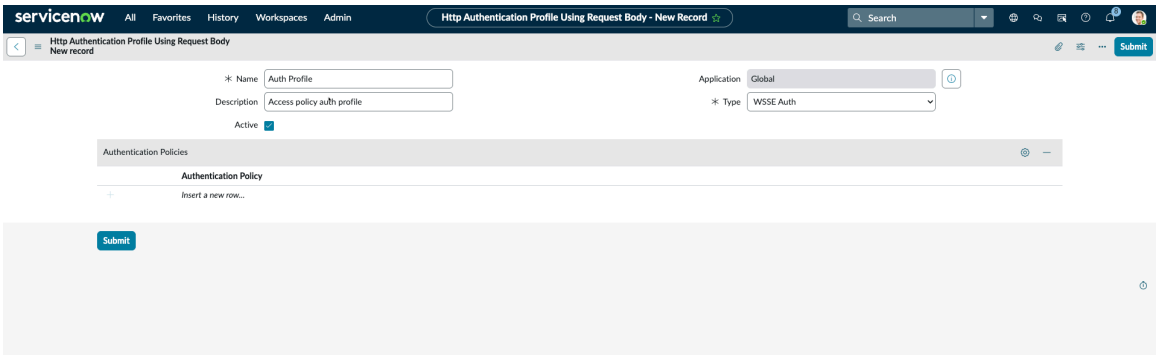

d. After creating the authentication profile, save the record.

6. Click **Submit** to submit the SOAP API access policy.

Create a global API access policy to protect SOAP APIs

Create a single global API access policy to protect all the SOAP APIs.

Before you begin

- Role required: admin
- Install the **Processor Access policy** (`com.glide.processor.policy`) plugin
- Make sure that an authentication profile is created. For more information, see Create an authentication profile.

The following steps describes how to create a single global access policy to protect all the SOAP APIs.

Note: The policies defined at the individual SOAP API levels overrides the **SOAPProcessor** level global access policy.

Procedure

1. Navigate to **All** > **System Security** > **Process Access Policies**.
2. On the form, fill in the fields.

Process Access Policy form

Field	Description
Processor	Name to identify the authentication policy. For example select, **SOAPProcessor** (Authentication profile).
Application	Scope of the authentication policy. Default: Global
Authentication Profile	Type of the authentications profile. Select, **Global SOAP Auth Profile**.

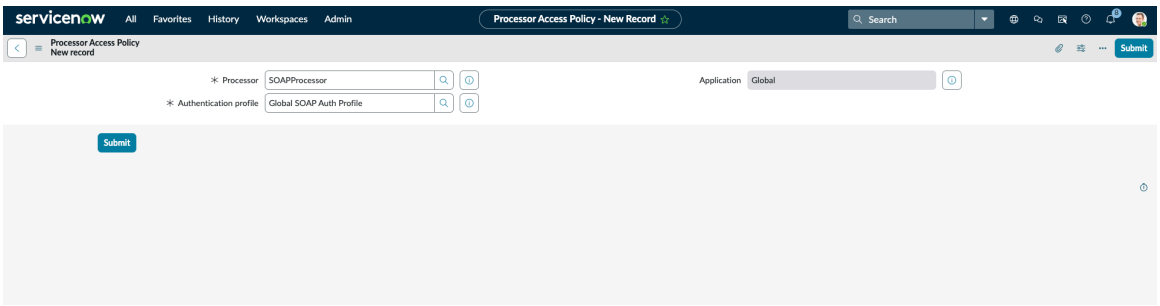

3. Select **Submit**.

Filter criteria for APIs

Filter criteria contains filter conditions or queries that are used as policy inputs for an authentication policy. Policy inputs are used to group one or more filter criteria and define the policy conditions of an authentication policy. For example, an IP Filter Criteria define an IP address in the Classless Inter-Domain Routing (CIDR) format or a range of IP addresses.

You can create filter criteria based on the user's IP address, role, and the user group to which they belong.

Note: You also use the filter criteria created from the **Adaptive Authentication** module. For more information, see Adaptive authentication.

You can create filter criteria for APIs using the same process as filter criteria for adaptive authentication. For details see Filter criteria.

Related topics

- Create IP filter criteria
- Create role filter criteria
- Create group filter criteria

API Authentication Policies

Authentication policies evaluate authentication requests based on the specified policy conditions and either allows or denies access depending on the matching criteria.

You can use the built-in Global Blocking Policy or create an authentication policy according to your security requirements. Global Blocking Policy denies the authentication requests of users and APIs based on the specified filter criteria.

Note: Do not use or modify Allow Access Policy and Deny Access Policy. These policies are meant only for user logins.

Create an API authentication policy

Authentication policies allow you to enforce access restrictions on the APIs based on the specified filter criteria.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Web Services** > **API Authentication Policy**.
2. Click **New**.
3. On the form, fill in these fields.

Policy form

Field	Description
Name	Name to identify the policy.
Description	Short description of the policy.
Application	Scope of the application.

4. Right-click the form header and select **Save**.
5. From the Policy Inputs tab, select **Edit** to add the existing Filter Criteria.
   You can also create a new Policy Input. For more information, see Create policy inputs.
6. Move one or more filter criteria from the Collections list to the Policy Inputs

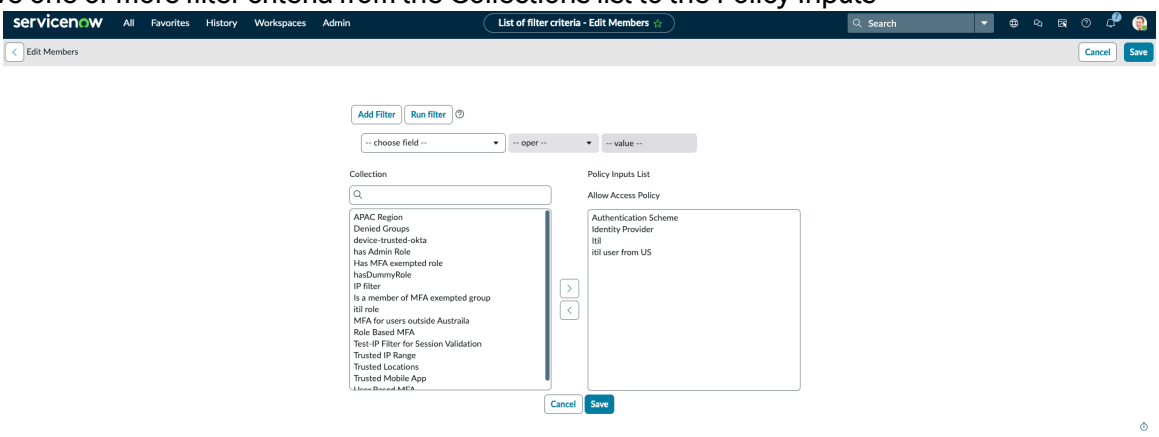

List.

7. Select **Save**.
8. From the Policy Conditions tab, select **New**.
9. On the form, fill in these fields.

Condition form

Field	Description
Label	Name of the policy condition.

Field	Description
Description	Short description of the policy condition.
Application	Scope of the application.
Condition	One or more conditions that are combined with **OR** filter.

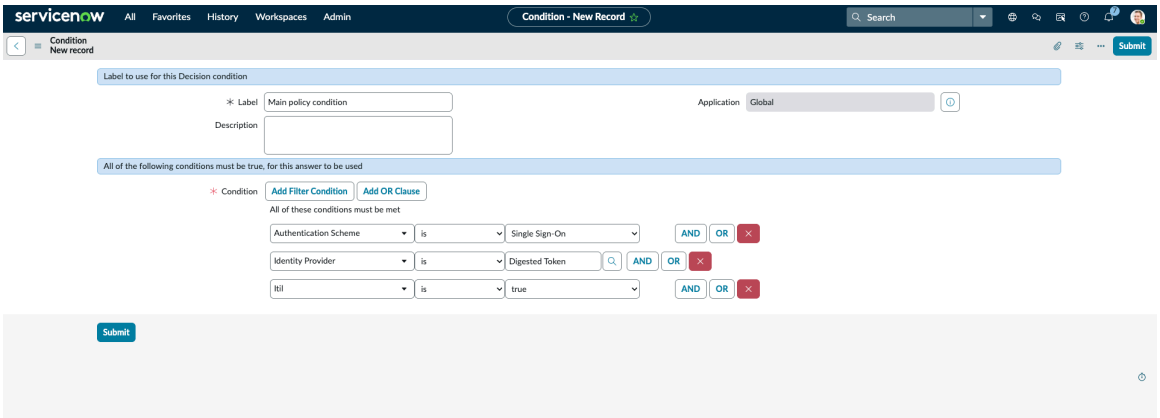

10. Select **Submit**.

Configure global blocking policy for APIs

Global blocking policy denies the authentication requests of users and APIs based on the specified policy conditions. This policy can be used as an alternative to the IP Address Access Control.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Web Services** > **API Access Policies** > **Global Blocking Policy**.
2. From the **Policy Inputs** tab, click **Edit**.
3. Select one or more filter criteria from the **Collection** list and move them to **Global Blocking Policy** list.
   You can also add additional filters.
4. From the **Policy Conditions** tab, click **New**.
5. On the form, fill these fields:

Condition Form

Field	Description
Label	Name to identify the condition.
Description	Description of the condition.
Condition	Logical combination of multiple policy inputs (filter criteria) that is used to evaluate authentication requests. For example, you can create conditions that allow only contractors from a list of trusted IP addresses.

Access policy for System or Export Processors

Ability for System or Export Processors to leverage processor access policy to secure all the export endpoints.

Ability to secure all the export endpoints and apply the inbound authentication profile and processor access policies at a global or instance level.

Note:

- Non-public processors, including the export processors like CSV, PDF, etc are supported for access policies.
- Script processors are also supported for access policies.

Sample use cases

- Admin can block the RSS processor if not intending to use it, by leveraging the API access policy.
- Admin can create an authentication profile with basic authentication and associate an authentication policy that always evaluates to false.

  Example, Create IP criteria with a range from `0.0.0.0` to `255.255.255.255` (add Ipv6 address space as well) and then add policy condition with the false operator. By this way, policy conditions will always evaluate as false, and the API access policy will block the access irrespective of where the request is originating.

- Allowing access to the processor only from a trusted network.

Activate Processor access policy

For Processor, install the Processor Access policy (`com.glide.processor.policy`) plugin.

Before you begin
Role required: admin

About this task
The following items are installed with Processor Access Policy:

- System Property: com.glide.auth.profile.supported.processor.list
- Module in the navigation: Processor Access Policies

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Processor Access Policy (`com.glide.processor.policy`) plugin using the filter criteria and search bar.

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

3. Select **Install** to start the installation process.

Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

What to do next

.
Configure authentication profile for the processors, for more information see, Configure Authentication profile for Processor.

Configure Authentication profile for Processor

Apply authentication profile for the export processors.

Before you begin

Role required: admin

Plugin required: Processor Access Policy (`com.glide.processor.policy`)

Procedure

1. Navigate to **All** > **System Security** > **Processor Access Policies**.
2. To add authentication profile to the processor, click **New**.
3. On the form, fill in the fields.

Process Access Policy form

Field	Description
Processor	Name to identify the authentication policy. **Note:** Public processors are not supported. If a non-supported processor is selected, then an error is displayed during submission of the processor.
Application	Scope of the authentication policy. Default: Global
Authentication Profile	Type of the authentications profile. You can select **ID Token** or **OAuth Token**.

4. Click **Submit**.
   The authentication profile is applied to the processor.

For example, OAuth authentication profile is configured for the CSV Processor. In this case, you have to use OAuth access token for the exporting using CSV as an export option.

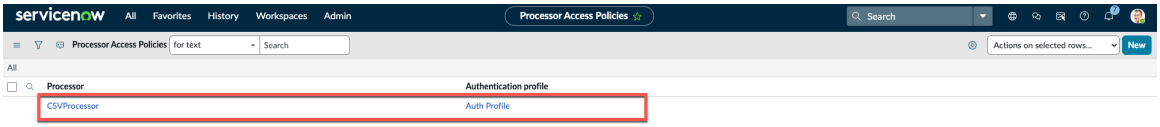

Authentication factors

Authentication factors help identify and verify callers, allowing only authorized users to access AI voice agents on the ServiceNow AI Platform.

Explore	Configure
Learn the features and business value of authentication factors for voice agents.	Understand how to configure authentication factors for voice agents.

Explore authentication factors for AI voice agents

Authentication factors are the elements used for caller identification and authentication. In secure voice agent environments, the process begins with identifying the caller, followed by authenticating their identity before granting access. A robust security strategy combines multiple factors to confirm that only authorized users interact with AI voice agents.

When configuring an AI voice service to support natural, conversational exchanges, it’s crucial to select authentication factors that reliably verify a user's identity. Caller access to specific voice agents is determined by the authentication types and methods configured by the administrator.

In this context, two categories of authentication mechanisms are supported:

Single-factor authentication

Single-factor authentication requires the caller to verify their identity through one method. Any of the six supported factors can be configured as a standalone factor.

Multi-factor authentication

Multi-factor authentication (MFA) requires callers to pass two verification methods in sequence. This raises the assurance level of the session and restricts access to sensitive data and actions.

- Primary factor: The initial verification method (for example, Soft PIN or TOTP).
- Secondary factor: An additional verification method that increases confidence in the caller's identity (for example, SMS OTP or Okta Verify push notification).

Note: MFA is enabled by default. To make single-factor authentication the default behavior, set the `glide.voice.authenticate.mfa_mandatory` system property to false.

Overview of the supported authentication factors

Time-based one-time password (TOTP) authentication

TOTP is a temporary numeric code generated by an authenticator app, such as Okta Verify, on the caller's registered device. Codes are generated locally and are resistant to interception, making TOTP well-suited for both single-factor and MFA configurations. Callers can enter the code via keypad or by speaking the digits.

Push notification - Okta Verify

Callers approve an authentication request via a push notification sent to their registered mobile device. This factor requires no code entry and is low-friction. It is effective as both a primary and secondary factor. An internet connection and a registered device with Okta Verify installed are required.

Soft PIN authentication

Soft PIN is a 6-digit numeric code the caller enrolls in advance. It is device-independent and quick to use across conversational AI channels, such as AI voice agents. Callers can enter the PIN through keypad or by speaking the digits. Because a PIN can be observed or shared, Soft PIN is best used alongside a second factor for sensitive actions.

SMS One-time passcode (OTP) authentication

SMS OTP delivers a temporary numeric code to the caller's registered mobile number. It is widely recognized and requires no app installation. Callers can enter the code via keypad or by speaking the digits. SMS OTP is susceptible to SIM-swapping and delivery delays and should not be the sole factor for critical operations.

Email One-time passwords (OTP) authentication

Email OTP delivers a temporary numeric code to the caller's registered email address. It is easy to deploy and familiar to most users. Callers can enter the code via keypad or by speaking the digits. Email OTP is susceptible to email account compromise and phishing, and should not be used as a standalone factor for sensitive operations.

Knowledge-based authentication (Security Questions)

KBA presents the caller with pre-configured questions, such as `"What are the last four digits of your employee ID?"`. The answers can be validated against ServiceNow AI Platform tables or external systems via custom scripts. KBA is used primarily for caller identification and low-risk authentication scenarios. Because answers can be social-engineered, KBA should not be used as a standalone factor for sensitive actions. Callers can respond via keypad or by speaking their answer.

For details on configuring voice input for authentication factors, see Configure voice input for authentication factors.

To learn more about voice service and how to create them, see Create an AI voice assistant.

Related topics

- Time-based one-time password (TOTP) authentication
- Push notification - Okta Verify
- Soft PIN authentication
- SMS One-time passcode (OTP) authentication
- Email One-time passwords (OTP) authentication
- Knowledge-based authentication (Security Questions)

Configure authentication factors for AI voice agents

To secure voice agent environments, configure authentication factors that first identify the caller, then authenticate them before granting access.

Configuration involves:

- User identification
- User authentication

User identification

User identification forms the foundational layer of security for AI voice agents. You can configure knowledge-based authentication (KBA) to establish the caller identity before any authentication steps.

User authentication

After identification, authentication confirms the caller's identity before enabling access to sensitive resources or actions. ServiceNow AI Platform supports both single-factor and multi-factor authentication (MFA), enabling administrators to tailor security configurations according to assurance levels and user roles. Supported factors include numeric PINs, authenticator app codes, one-time passwords over SMS and Email, push notifications, and knowledge-based questions.

Recommendations

- You're encouraged to use multi-factor authentication as the default approach for optimal security.
- You should regularly review and update authentication factor configurations to address evolving threats and maintain conformance.

The following matrix shows the authentication factor combinations based on assurance level and authentication strategy:

Authentication factor combinations

Factor	Security Assurance	Recommended as Single Factor	Recommended as First Factor (MFA)	Recommended as Second Factor (MFA)	Learn more about configuration
Time-based one-time password (TOTP) authentication	High	Y	Y	Y	Authenticator Applications

Authentication factor combinations (continued)

Factor	Security Assurance	Recommended as Single Factor	Recommended as First Factor (MFA)	Recommended as Second Factor (MFA)	Learn more about configuration
Push notification - Okta Verify	High	Y	Y	Y	Configure push notification (Okta Verify)
Soft PIN authentication	Medium	Sometimes	Y	Y	Configure Soft PIN
SMS One-time passcode (OTP) authentication	Medium	N	N	Y	Multi-factor authentication Providers
Email One-time passwords (OTP) authentication	Medium	N	Sometimes	Y	Configure Email OTP
Knowledge-based authentication (Security Questions)	Low	N	Y	N	Configure knowledge-based authentication

Important: Y = Recommended | Sometimes = Use with Caution | N = Not Recommended

Configure voice input for authentication factors

Configure how callers provide authentication responses by speaking or using the phone keypad.

Before you begin

Role required: auth_factors_admin

Note:

- For KBA, voice input is configured per question. For more information, see Create KBA questions
- For all other numeric factors — Authenticator App (TOTP), Email OTP, SMS OTP, and Soft PIN, use the following procedure.

Procedure

1. Navigate to **All** > **Authentication Factors** > **Services** > **Service Configurations**.
2. Select **New**.
3. Specify the following fields on the form.

Service Configuration form

Field	Description
Authentication Factor	Select the factor to configure. Options: **Authenticator App (TOTP)**, **Email OTP**, **SMS OTP**, **Soft PIN**.
Application	Global application scope is selected by default.
Service Profile Table	Select the table for the AI voice agent service. Leave empty to apply the configuration to all AI voice agent services.
Service Profile	Select the specific AI voice agent service the configuration applies to. Leave empty to apply the configuration to all AI voice agent services.
Input Type	Select how callers provide their response. Options: ○ **Text**: The caller enters the response using the phone keypad. ○ **Voice**: The caller speaks the response aloud. When **Voice** is selected, the **Input Format Description**, **Input Format Example**, and **Input Validation Pattern** fields are displayed with preset values. These fields are read-only.
Input Format Description	This field appears only when **Input Type** is set to **Voice**. Expected format of the spoken response. Read-only. Preset by the platform for each factor.
Input Format Example	This field appears only when **Input Type** is set to **Voice**. Example spoken responses to guide the voice agent. Read-only. Preset by the platform for each factor.
Input Validation Pattern	This field appears only when **Input Type** is set to **Voice**. Regular expression pattern used to validate the transcribed response. Read-only. Preset by the platform for each factor.

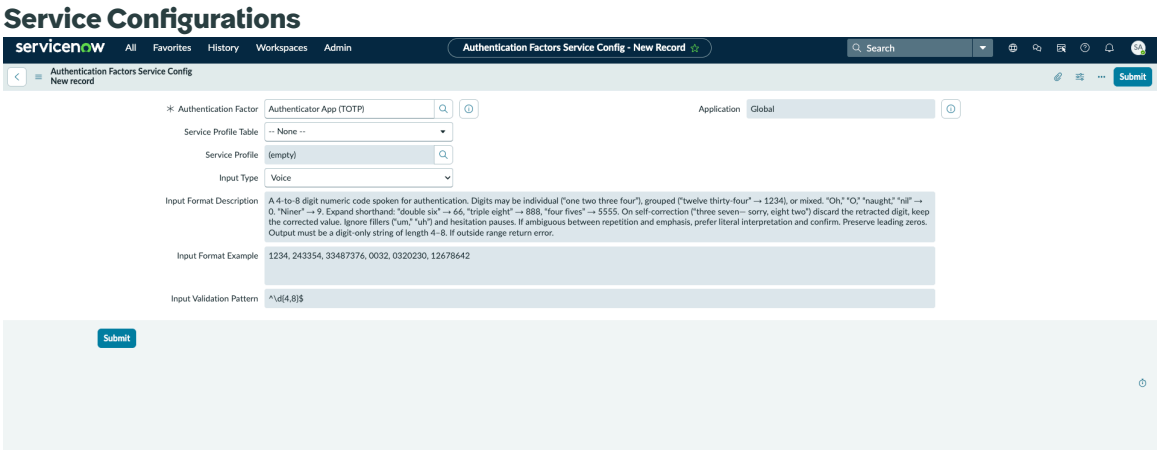

4. Select **Submit**.

Result

The input type configuration is saved. During an AI voice agent session, callers authenticate using the configured input method for the selected factor and service.

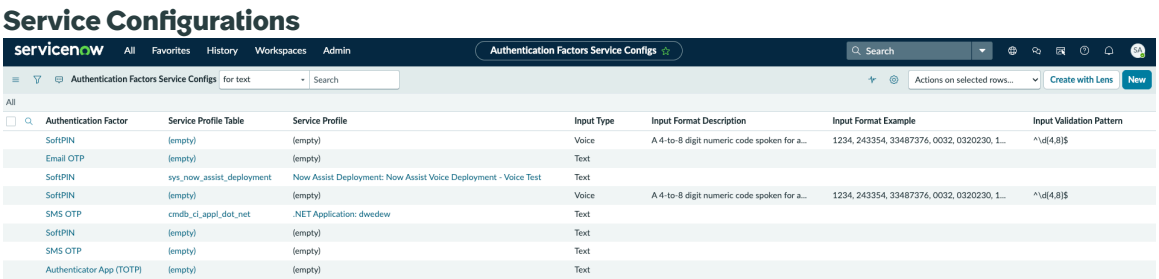

Time-based one-time password (TOTP) authentication

A time based one-time password (TOTP) is a secure authentication factor that verifies user identity by generating a unique, time-sensitive code.

TOTP authenticator apps like **Okta Verify** generate temporary numeric codes on a registered device, usually a mobile phone. These codes are only valid for a short time and can't be used more than once. This approach offers enhanced security compared to static passwords.

Use case

TOTP authenticator apps are suitable for users who require stronger protection than standard multi-factor authentication (MFA) methods.

Key strengths

The TOTP authenticator apps method offers the following advantages:

- Local generation: One-time passwords are generated directly on the user's device and don't require an internet or mobile network connection.
- Security: TOTP is resistant to interception and SIM-swapping attacks, providing robust protection for the authentication process.

Important considerations

While TOTP authenticator apps are a secure and convenient authentication method, there are few considerations to keep in mind:

- Initial setup: Users must set up the authenticator app on a registered device.
- Device management: Users must re-enroll when devices are replaced or reset.
- Phishing risk: One-time codes can be compromised if entered on untrusted or malicious sites.

TOTP authenticator apps are an effective method to strengthen your organization's security posture. For detailed configuration instructions, see Authenticator Applications.

Push notification - Okta Verify

The **Okta Verify** app push notification enables users to securely approve authentication requests directly on their enrolled mobile devices.

During authentication, users receive a push notification via the **Okta Verify** app, enabling them to review request details and approve access without manually entering a verification code.

Use case

Okta Verify push notification is recommended as a robust and convenient second factor for ServiceNow AI Platform authentication flows, providing enhanced security and a seamless user experience.

Key strengths

The **Okta Verify** push notification method offers the following advantages:

- A streamlined process enables quick and easy authentication.
- Resilient against various attack vectors, reinforcing the authentication process.

Important considerations

While **Okta Verify** push notification is a secure and convenient authentication method, there are few considerations to keep in mind:

- Connectivity: Requires a reliable internet connection and a compatible mobile device.
- User Vigilance: Users should carefully review each authentication prompt to help prevent unintended approvals.
- Device Security: Registered devices must be protected with screen locks or biometric safeguards to help prevent unauthorized access.
- Device Re-enrollment: If a device is replaced or reset, users must re-enroll to maintain uninterrupted authentication.

Okta Verify push notification enhances organizational security while providing a convenient authentication experience. For detailed configuration instructions, see Configure push notification (Okta Verify).

Configure push notification (Okta Verify)

Configure **Okta Verify** to receive push notifications for secure and convenient identity verification.

Before you begin

Role required: auth_factors_admin

The push notification authentication factor on **Okta** must be enabled. To learn more, see Okta Support Center.

Note:

- Currently, only **Okta Verify** is supported for push notifications with ServiceNow AI voice agents.
- The users must have the **Okta Verify** app installed and configured on their phones.

Procedure

1. Log in to your Okta admin account.
2. Navigate to **Security** > **API**.
3. Select **Tokens**.

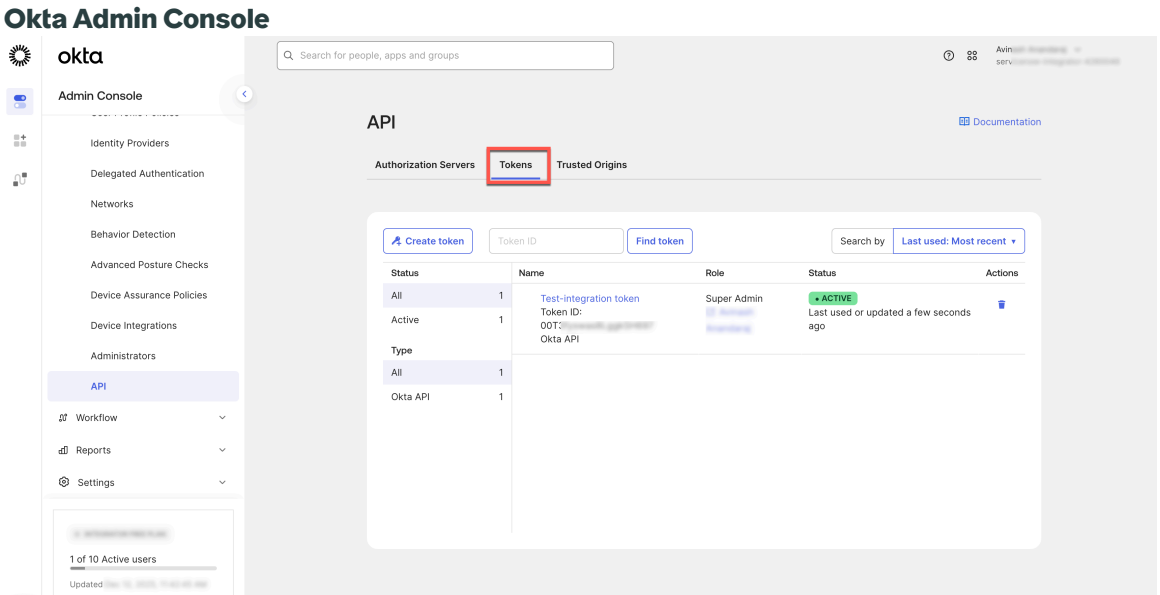

4. Select **Create** token to create token and copy its value.

   For example, `00EYJEm|b34aYVi0Yaav3KnqxsItV_mPI9_p2N46Cc`.

5. On the instance, open the **Connection & Credential Alias** record.

   For example, `<instance_url>/sys_alias.do?sys_id=692e16a0ffb72210d487fffffffff1b`.

Connection & Credential Alias new record

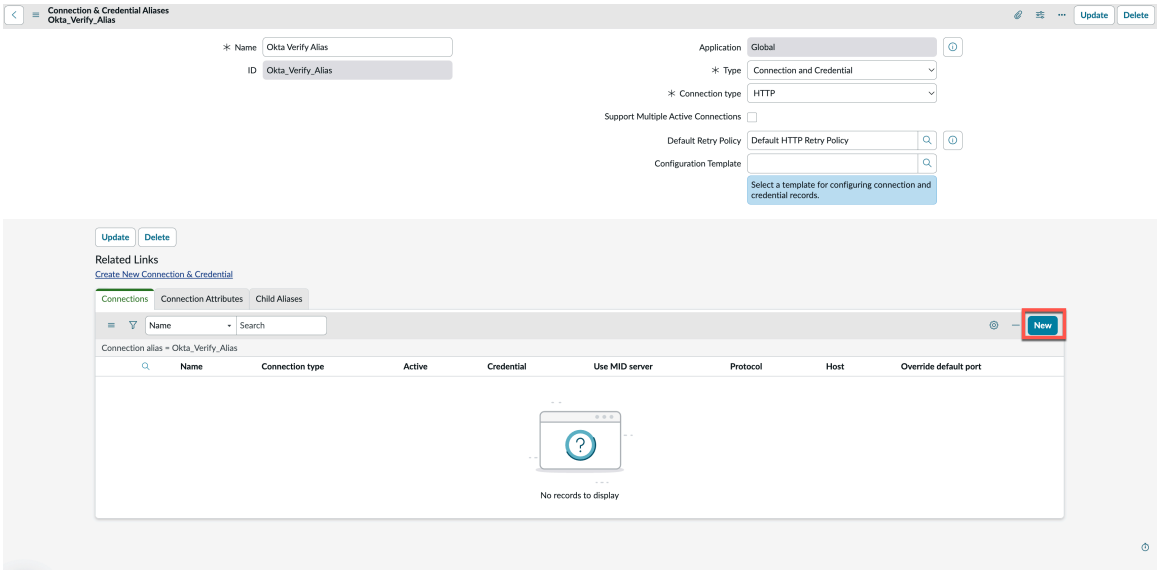

6. Select **New** under the Connections tab.
7. Specify the following field on the **HTTP(s) Connections** form.
   - Name
   - Connection alias
   - Connection URL

HTTP(s) Connections details

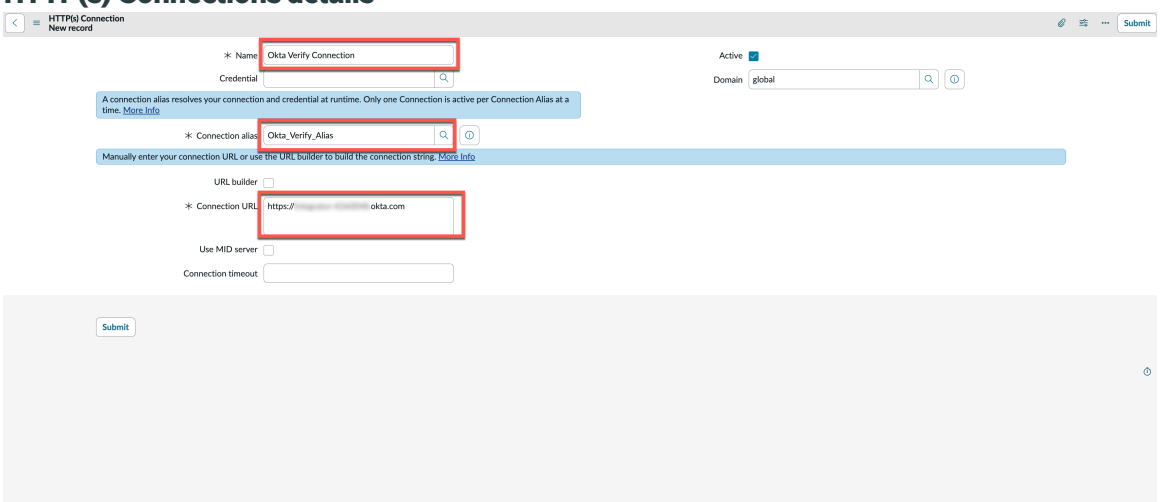

8. Select the **Search** icon next to Credentials to add new **API Key Credentials**.

API Key Credentials selection

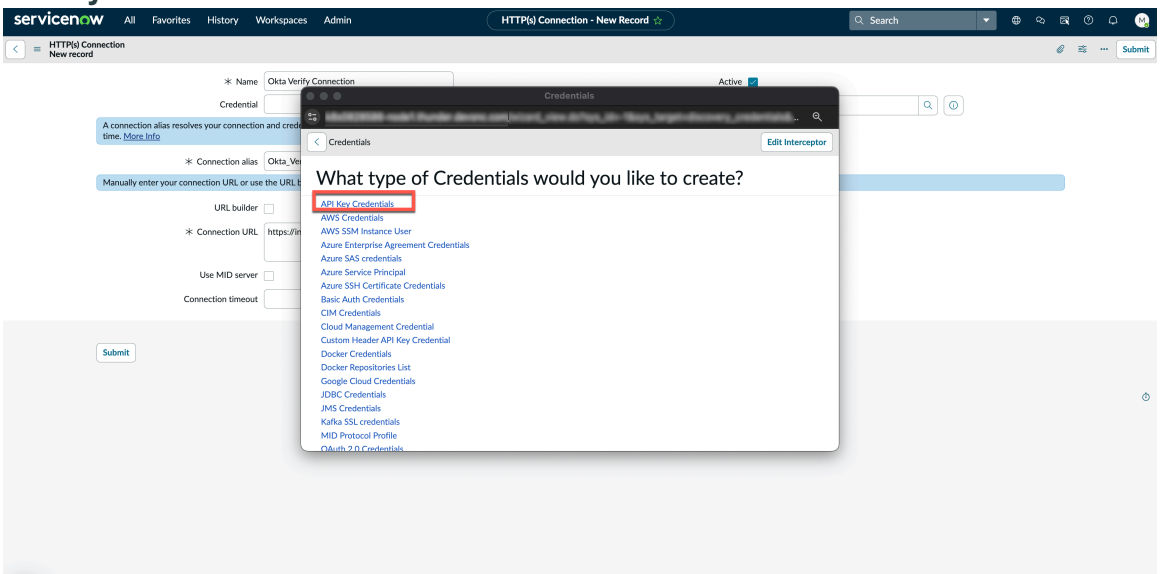

9. Specify the **Name** and **API Key** on the API Key Credentials form.

Note: The API key should be prefixed with 'SSWS'. If the copied token value is `00iJ88ydm9ZLnDNE8AeGsmjcUC-ZFyCt3pts2pNTwZ` you must provide the value as `SSWS 00iJ88ydm9ZLnDNE8AeGsmjcUC-ZFyCt3pts2pNTwZ`.

API Key Credential details

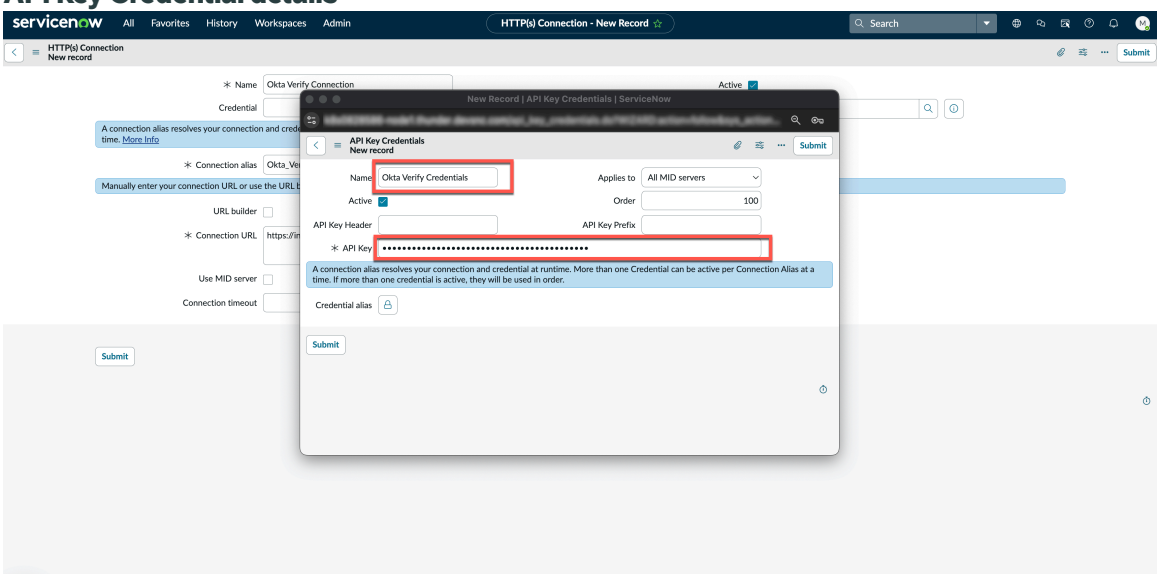

10. Select **Submit** to submit the API Key Credentials.

11. Select **Submit** on the Connection record.

HTTP(s) Connections record completion

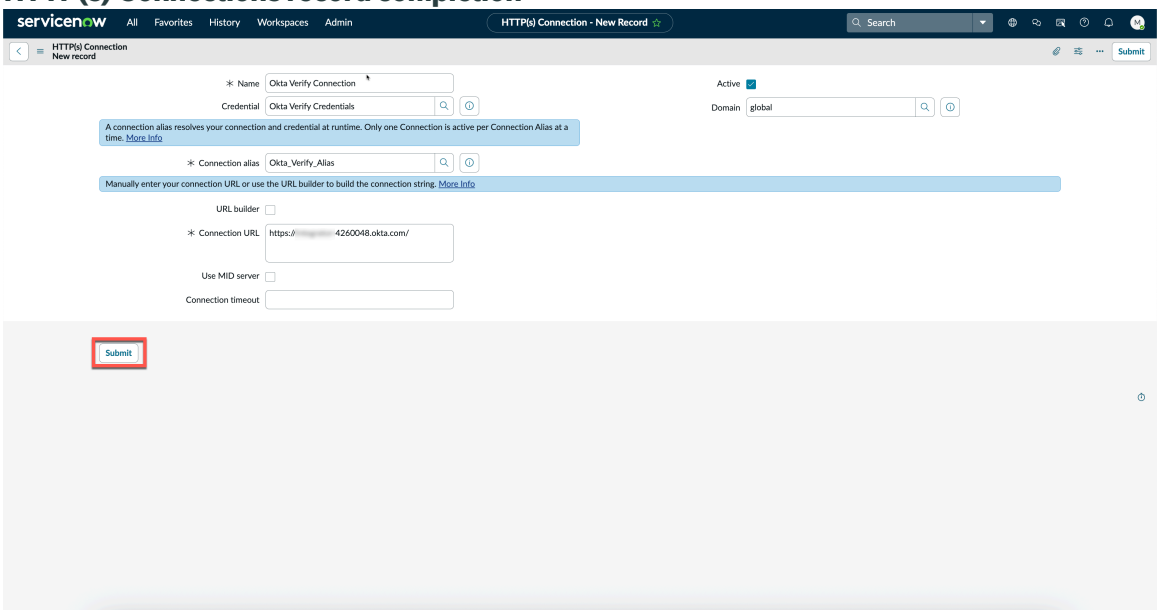

Soft PIN authentication

Soft PIN is a six-digit numeric PIN that verifies a caller's identity during an AI voice agent session.

When to use Soft PIN

Soft PIN is appropriate for low-risk caller verification, such as confirming a returning user before granting access to self-service tasks.

Soft PIN can be configured as a single factor, the first factor in a multi-factor authentication flow, or a second factor.

Soft PIN is a medium-assurance factor and is not suitable as the only authentication factor for sensitive operations. For those flows, combine Soft PIN with a higher-assurance factor such as Okta Verify push notification or a time-based one-time password (TOTP). For guidance on combining factors, see Explore authentication factors for AI voice agents.

How Soft PIN works

Each user enrolls a Soft PIN before it can be used for authentication. Users can change their PIN by re-enrolling at any time.

When Soft PIN is selected as an authentication factor for an AI voice agent service, the agent prompts the caller for the PIN during the session. The platform validates the response against the user's enrolled PIN and returns the result to the orchestrator.

Note: Soft PIN supports both Text and Voice input.

Enrollment rules

The system enforces the following rules on the chosen PIN:

Enrollment rules

Rule	Behavior
Length	Exactly six digits.
Repetition	No single digit can repeat more than twice consecutively. For example, `111234` is rejected.
Sequences	Ascending or descending numeric sequences longer than two digits aren't allowed. For example, `123456` and `987654` are rejected.
History	The new PIN can't match any of the user's previous five PINs.

Limitations

A six-digit numeric PIN provides lower assurance than time-based codes or push notifications. PINs are vulnerable to reuse, observation, and social engineering.

Availability

The administrator manages the following conditions on the instance. Soft PIN enrollment is available when both are met:

- Install Now Assist for Platform `sn_genai_platform` for activating AI voice agents.
- The system property `glide.auth_factors.Soft PIN.enrollment.enabled` is set to `true` (default).

When the plugin is not installed, no Soft PIN module exists on the instance and the enrollment URL is not available. When the plugin is installed but the property is set to false, the enrollment option is hidden from the user profile, the navigation menu, and the Service Portal. Users who navigate directly to the enrollment URL see the following message:

```
Soft PIN enrollment is not available at this time. Please contact your administrator for more details.
```

System property

Property	Description	Default state
`glide.auth_factors.Soft PIN.enrollment.enabled`	Controls whether the Soft PIN enrollment option appears in the user profile, the navigation menu, and the Service Portal. Requires the AI Voice Agents plugin.	true

Related topics

Configure Soft PIN

Authentication factors

Configure Soft PIN

Users are required to configure Soft PIN before it can be used for authentication with ServiceNow AI Platform.

Before you begin

Role required: none

Perform the following:

- Administrator must enable Soft PIN on the instance before you can enroll.
- Install Now Assist for Platform `sn_genai_platform` for activating AI voice agents and set `glide.auth_factors.softpin.enrollment.enabled` property to `true`. To know more, see Availability in Soft PIN authentication.
- The user record must exist in the `sys_user` table.

Procedure

1. Navigate to the Soft PIN enrollment page using any of the following:
   - **Service Portal**: Go to your **User Profile** and select **Enroll Soft PIN**.

Soft PIN on Service Portal

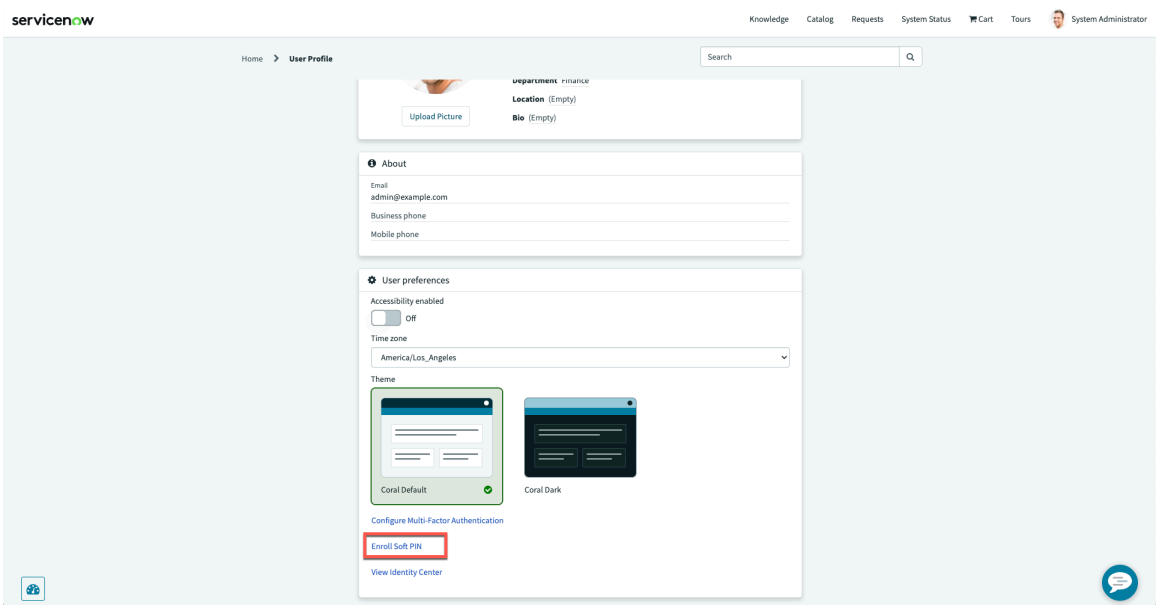

   - **Platform UI**: Go to your **User Profile** and select **Enroll Soft PIN** under Related Links.

Soft PIN on Platform UI

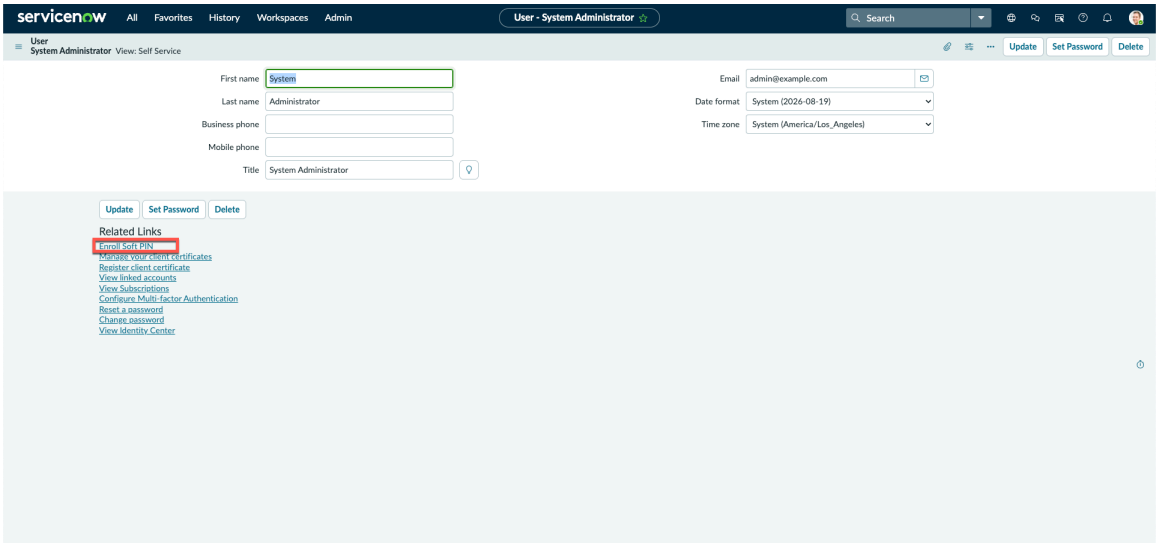

- **Navigation menu**: Select **All** > **Authentication Factors** > **Soft PIN** > **Enroll**.

Soft PIN on Navigation menu

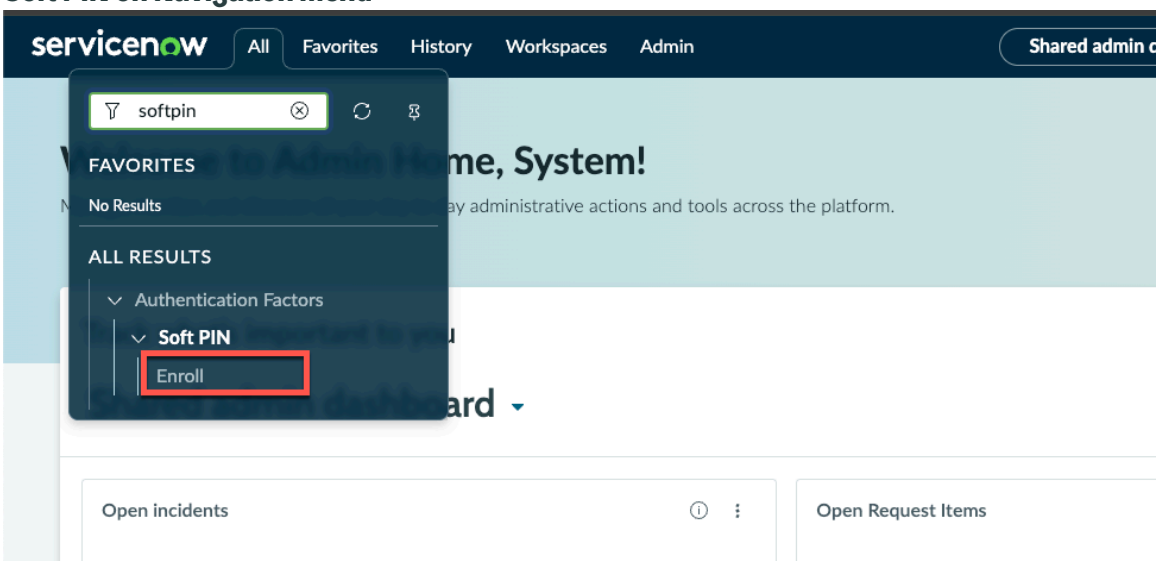

2. Specify a six-digit PIN that satisfies the enrollment rules:

   Rules for Soft PIN:

   - The PIN must be exactly six digits in length
   - No single digit must be repeated more than twice consecutively
   - Don’t use ascending or descending numeric sequences longer than two digits
   - Can’t reuse any of your previous five PINs

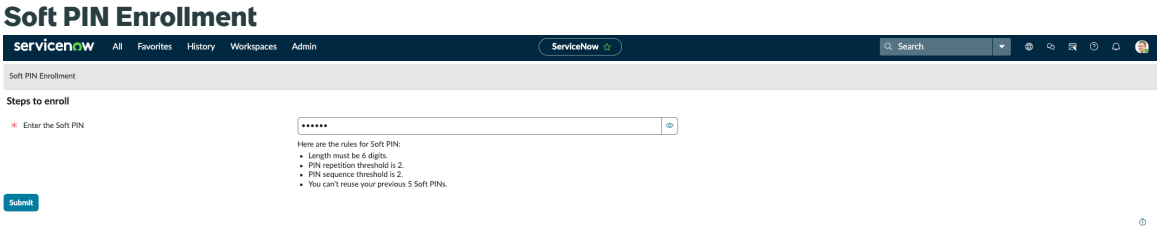

3. Select **Submit**.

Result

You can use the submitted Soft PIN to authenticate various ServiceNow system or service. For example, AI voice service.

Note: Users are eligible for Soft PIN enrollment only if their user ID is present in **sys_user** table. If its missing, an enrollment failure message is displayed.

Update your Soft PIN

To change your PIN, return to the enrollment page from any of the entry points before and enter a new PIN. The new PIN must satisfy the same rules and can't match any of your previous five PINs. Submitting a new PIN replaces the existing one immediately.

Related topics

Soft PIN authentication

Authentication factors

SMS One-time passcode (OTP) authentication

SMS one-time password (OTP) authentication is a method used to verify user identity by sending a temporary, numeric code to the user's registered mobile number. The user enters this code to complete authentication.

Use case

- Serves as a secondary authentication factor, particularly for users without authenticator apps.
- Recommended for medium-risk scenarios, such as verifying changes to profile information and approving login attempts.

Key strengths

The SMS OTP method offers the following advantages:

- Familiar and easy for users to adopt
- No additional app installation required
- Simple to deploy and integrate

Important considerations

While SMS method is a convenient authentication method, there are several considerations to keep in mind:

- Vulnerable to SIM-swapping, interception, and delivery delays.
- Reliant on mobile network availability.
- Not recommended as the sole authentication factor for high-risk or sensitive operations.

SMS OTP can enhance overall security when used appropriately. For detailed configuration instructions, see Multi-factor authentication Providers.

Email One-time passwords (OTP) authentication

Email OTP for AI voice agents sends a one-time numeric code to the caller's email address. The caller retrieves the code from their email and provides it to the agent to verify their identity.

When to use Email OTP

Email OTP is appropriate for caller verification when the caller has access to their email during the session and SMS delivery is not preferred or available. Email OTP can be configured as a single factor, the first factor in a multi-factor authentication flow, or a second factor.

Email OTP is a medium-assurance factor and is not suitable as the only authentication factor for sensitive operations. For those flows, combine Email OTP with a higher-assurance factor such as Okta Verify push notification or a time-based one-time password (TOTP). For guidance on combining factors, see Explore authentication factors for AI voice agents

How Email OTP works

and returns the result to the orchestrator. Email OTP supports both **Text** and **Voice** input.

Email source configuration

Email OTP determines which email address to send the code to by reading from a configuration record in the Email OTP Service Configurations table. Each record specifies a user record location: which table to look in, which column holds the email, and which column links that record back to the user's identity in sys_user. The table is domain-separated.

By default, the ServiceNow AI Platform ships an instance-level configuration base system that reads the Email field from the User (sys_user) table, keyed by Sys ID. Administrators can override this by pointing to a different table — for example, a custom contact table or a group table — as long as the table has an email field and a column that references sys_user. Service profile-specific configurations take precedence over the instance-level default.

Limitations

Email OTP requires the caller to have access to their email during the session. Latency depends on email delivery times.

Availability

Email OTP is available base system on ServiceNow AI Platform. No plugin installation or system property change is required to use Email OTP.

Related topics

- Configure Email OTP
- Explore authentication factors for AI voice agents

Configure Email OTP

Configure the Email one-time password (OTP) to enable OTP-based authentication for users in your instance.

Before you begin

Role required: `auth_factors_admin`

Install Now Assist for Platform `sn_genai_platform` for activating AI voice agents before Email OTP can be configured.

Procedure

1. Navigate to **All** > **Authentication Factors** > **Email OTP Factor** > **Email OTP Service Configuration**.
2. Select **New**.
3. Specify the following fields on the form:

Email OTP Service Configuration

Field	Description
Service Profile Table	Select the table associated with the AI voice agent service profile. For example, `sys_now_assist_d...` (Now Assist Deployment).
Service Profile	Select the specific AI voice agent service profile to associate with the Email OTP configuration. For example, `Now Assist Deployment: Auth Factors Voice`.
Table	Select the table that contains the user records. Defaults to `User [sys_user]`.
Active	Select to enable the Email OTP configuration. Enabled by default.
Email Column	Select the column from the user table that contains the email address used to send the OTP. Defaults to `Email`.
User Column	Select the column used to identify the user record. Defaults to `Sys ID`.

Email OTP Service Configuration

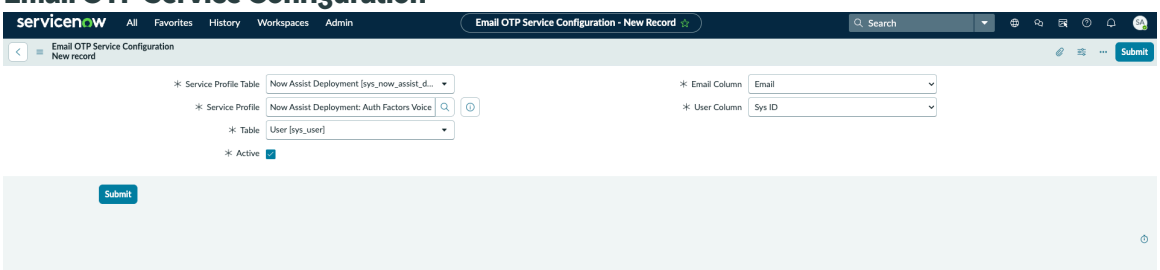

4. Select **Submit**.

Result

The Email OTP configuration is saved. During an AI voice agent session, a one-time passcode is sent to the user's registered email address for authentication.

Related topics

- Email One-time passwords (OTP) authentication
- Explore authentication factors for AI voice agents

Knowledge-based authentication (Security Questions)

Knowledge-based authentication (KBA) is an identification and authentication method that verifies callers by prompting them to answer preconfigured questions across conversational AI channels, such as AI voice agents. KBA can be used to identify a caller, authenticate a caller, or both within the same interaction.

KBA validates answers against records in ServiceNow AI Platform tables. For callers whose data resides outside ServiceNow, admins can configure scripts to validate answers against external systems in real-time. External data is never imported or stored in ServiceNow AI Platform.

How KBA works

Identification locates the caller by matching their answer to a record in ServiceNow AI Platform or an external system. For example, a caller provides their business phone number, and the system finds a matching record. Identification runs once per session and establishes who the caller is before any sensitive interaction begins.

Authentication verifies that the caller is who they claim to be. The caller answers one or more questions, and the system validates those answers against stored or externally sourced data.

KBA questions can be configured for identification, authentication, or both phases, depending on admin configuration.

External source validation

When caller data is not stored in ServiceNow AI Platform, admins can configure a custom script on an answer record to validate the caller's response against an external system, such as a CRM or order management platform. The script receives the caller's answer as input and returns a match result.

- For identification, the script returns the matched record.
- For authentication, the script returns a `true` or `false` result.

Note: Only `snc_external` users can be authenticated using external source.

Script execution is limited to 15 seconds by default. To learn more configuration properties, see System Properties.

Context persistence

Starting `"nowassist-aia-voice", version: "5.0.3"` release, answers collected during identification and authentication are persisted as session context and are available to subsequent authentication questions. This means a caller does not have to repeat information they already provided. For example, if a caller provides a booking reference during identification, that value is accessible to authentication scripts without prompting the caller again.

Note: Context persistence is available only for scripted answers, it doesn't capture non-scripted answer responses.

Key strengths

- No additional device or internet connectivity is required.
- Familiar to most users.

Limitations

- KBA relies on information the caller knows, which can be guessed, obtained from public records, or exposed through social engineering.
- KBA is not recommended as the sole verification method for sensitive operations.
- KBA is best suited for low-risk scenarios, such as general IT support or public documentation access.

For detailed configuration instructions, see:

- Create KBA questions
- Create KBA answers
- Map KBA questions to answers
- AI voice agent service mapping with KBA

Configure knowledge-based authentication

Configure knowledge-based authentication (KBA) to identify and authenticate callers by prompting them to answer preconfigured questions across conversational AI channels, such as AI voice agents.

Before you begin

Role required: auth_factors_admin

Procedure

1. Navigate to **All** > **Authentication Factors** > **Knowledge Based Factor**.
2. Perform the following:
   - Create KBA questions
   - Create KBA answers
   - Map KBA questions to answers
   - AI voice agent service mapping with KBA

   After creating the KBA configurations, you can use it as an identification and authentication method.

Create KBA questions

Create knowledge-based questions to use for caller identification and authentication in AI voice agent interactions.

Before you begin

Role required: auth_factors_admin

Default questions

A set of KBA questions default is available to all AI voice assistants in AI Voice Assistant Designer. You can use these questions as-is or edit them.

Default KBA questions

Question	Default answer source	Type	Category
Phone Number	sys_user.mobile_phone	Identification	Phone, triggers Automatic Number Identification (ANI)
Employee ID	sys_user.employee_number	Identification or Authentication	Others
Email	sys_user.email	Identification or Authentication	Others
Manager	sys_user.manager	Identification or Authentication	Others
Zip Code	sys_user.zip	Identification or Authentication	Others

Procedure

1. Navigate to **All** > **Authentication Factors** > **Knowledge Based Factor** > **Questions**.
2. Select **New** on the Knowledge Based Questions page.
3. Specify the following fields on the form:

Knowledge-Based Questions

Field	Description
Question	Define a security question to be used for user identity verification. Example: `What is your business phone number?`
Application	Global application scope is selected by default.
Keyword	Enter the primary keyword that best describes the question. Example: `business_phone`.
Category	Select the category. Options: ○ Phone Number: Enables partial phone number matching and makes the caller's ANI (automatic number identification) available by default for identification. ○ Others: Standard matching; ANI is not used. **Note:** Phone Number category can't be used for authentication.
Channel	Select the channel for which the question is configured. Currently, only **Voice** is supported.

Field	Description
Input Type	Select how the user provides their answer. Options: ○ **Text**: The caller enters a response via phone keypad. ○ **Voice**: The callers speak their response.
Type	Select when the question is available for use. Options: ○ **Identification**: The question is available for identification configuration only. ○ **Authentication**: The question is available for authentication configuration only. ○ **Identification or Authentication**: The question is available for both.

4. Select the Input Type:

   a. Input Type as **Text**: No additional fields are required.

   Knowledge Based Question - Text

   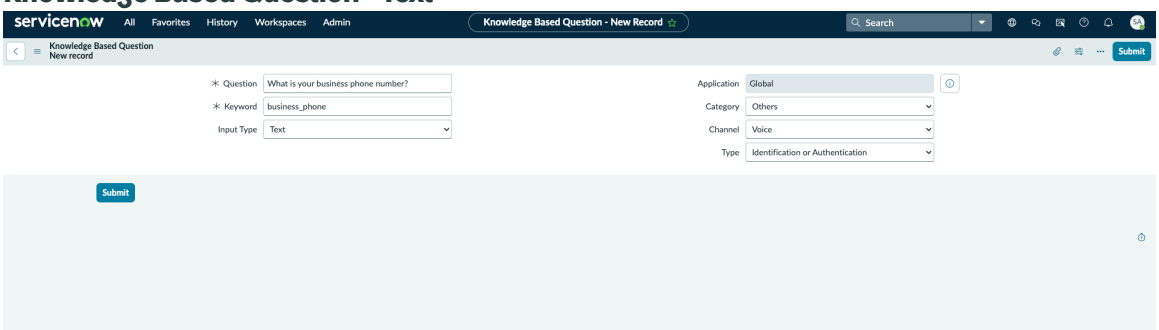

   Note: If Input Type is set to **Text**, no additional fields are required.

   b. Input Type as **Voice**: Specify the following additional fields:

   Voice input fields

Field	Description
Input Format Description	Describe the expected format and structure of the spoken response. For example: `5-character alphanumeric employee ID starting with the letter E.`
Input Format Example	Specify comma-separated examples of valid spoken responses. For example: `E372B, E481K, E529D.`
Input Validation Pattern	Specify a regular expression pattern to validate the spoken response against the

Field	Description
	expected format. For example: `^E[0-9]{3}[A-Z]$`.

Knowledge Based Question - Voice

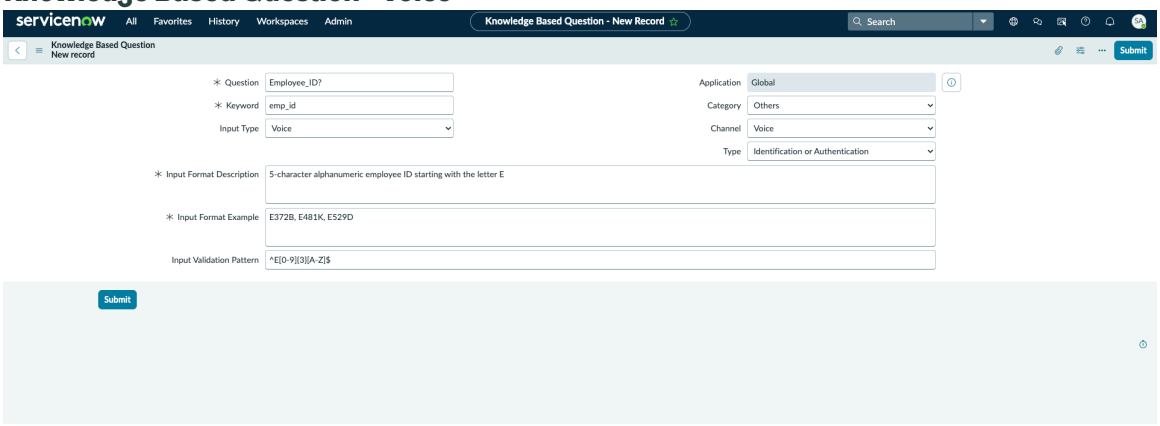

5. Select **Submit**.

Result

You're redirected to the Knowledge Based Questions list view. Verify if your question is successfully added.

Knowledge Based Questions - list

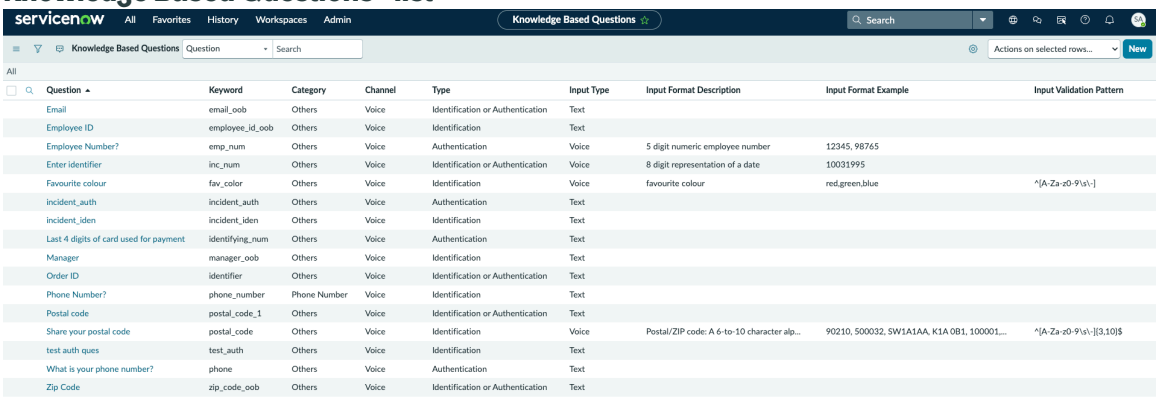

Question	Keyword	Category	Channel	Type	Input Type	Input Format Description	Input Format Example	Input Validation Pattern
Email	email_oob	Others	Voice	Identification or Authentication	Text			
Employee ID	employee_id_oob	Others	Voice	Identification	Text			
Employee Number?	emp_num	Others	Voice	Authentication	Voice	5 digit numeric employee number	12345, 98765	
Enter identifier	inc_num	Others	Voice	Identification or Authentication	Voice	8 digit representation of a date	10031995	
Favourite colour	fav_color	Others	Voice	Identification	Voice	favourite colour	red,green,blue	^[A-Za-z0-9\s\-]
incident_auth	incident_auth	Others	Voice	Authentication	Text			
incident_iden	incident_iden	Others	Voice	Identification	Text			
Last 4 digits of card used for payment	identifying_num	Others	Voice	Authentication	Text			
Manager	manager_oob	Others	Voice	Identification	Text			
Order ID	identifier	Others	Voice	Identification or Authentication	Text			
Phone Number?	phone_number	Phone Number	Voice	Identification	Text			
Postal code	postal_code_1	Others	Voice	Identification or Authentication	Text			
Share your postal code	postal_code	Others	Voice	Identification	Voice	Postal/ZIP code: A 6-to-10 character alp...	90210, 500032, SW1A1AA, K1A 0B1, 100001,...	^[A-Za-z0-9\s\-]{3,10}$
test auth ques	test_auth	Others	Voice	Identification	Text			
What is your phone number?	phone	Others	Voice	Authentication	Text			
Zip Code	zip_code_oob	Others	Voice	Identification or Authentication	Text			

Create KBA answers

Create knowledge-based answers for the preconfigured security questions to confirm the user's identity.

Before you begin

Role required: auth_factors_admin

KBA validates answers against records in ServiceNow AI Platform tables by default. To #validate against#an external source such as a CRM or order management platform, select Identification or Authentication in the Script Configuration field. External data is never imported or stored in ServiceNow AI Platform.

The User Column field determines whether an identified caller can proceed to authentication. When populated with a field that links to a system user account, the caller is eligible for

authentication. When left empty or mapped incorrectly, the caller is treated as a guest — the caller can access personalized, non-sensitive information but can't be authenticated.

Procedure

1. Navigate to **All** > **Authentication Factors** > **Knowledge Based Factor** > **Answers**.
2. Select **New** on the Knowledge Based Answers page.
3. Specify the following fields on the form:

Common fields

Field	Description
Script Configuration	Select the type of script-based validation. Options: ○ **None**: Validate the caller's answer against a ServiceNow AI Platform table. ○ **Identification**: Validate a caller's identity via a custom script against an external system. ○ **Authentication**: Validate a caller's authentication via a custom script against an external system. **Note:** When Script Configuration is set to **Identification** or **Authentication**, the Answer Table, Answer Column, and User Column fields are replaced by a script editor.
Description	Define a description for the answer. For example: Business Phone Number.
Application	Global application scope is selected by default.

4. Specify fields based on the Script Configuration selected:

   a. Script Configuration as **None**: Validate against a ServiceNow AI Platform table:

Internal source fields

Field	Description
Answer Table	Select the answer table. **Note:** The selected table should have a field referenced to the `sys_user` table.
Answer Column	Select the answer column from the **Available** list. Example: `Business phone`.

Field	Description
User Column	Select the field that links records in the Answer Table to a system user account. For example: sys_user. All fields of type reference from the Answer Table are displayed and can be selected. When populated with a valid `sys_user` reference, the answer supports identification and authentication. When left empty, the answer supports guest identification only and the caller can't be authenticated. Whether this field is required is controlled by the property. To know more, see System Properties.

Knowledge Based Answer - Script Configuration: None

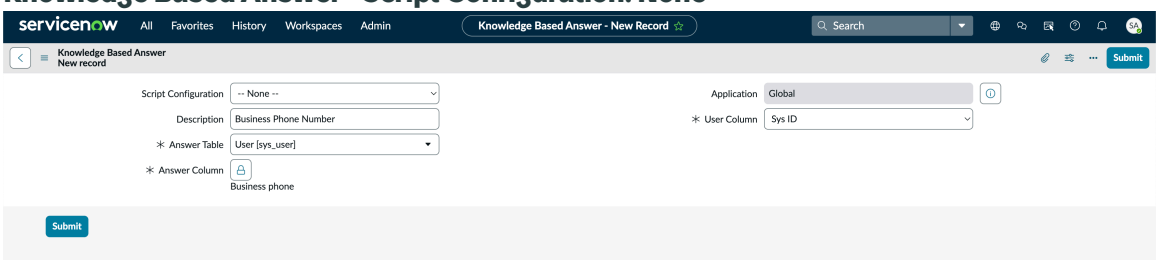

b. Script Configuration as **Identification**: Validate against an external system during identification:

Note: KBA Answers - **Script Configuration**: Works for only users with `snc_external` role users.

External source fields (Identification)

Field	Description
Script	Define the custom script to validate the caller's answer against an external system and return the location of the matched ServiceNow AI Platform record. Script input: `user_input` , the caller's answer. Script output: ▪ `table_name`: The ServiceNow AI Platform table where the matched record is located. ▪ `sys_id`: The `sys_id` of the matched record within the specified table.

Field	Description
	For script execution time limits, see System Properties.

Script Configuration as Identification

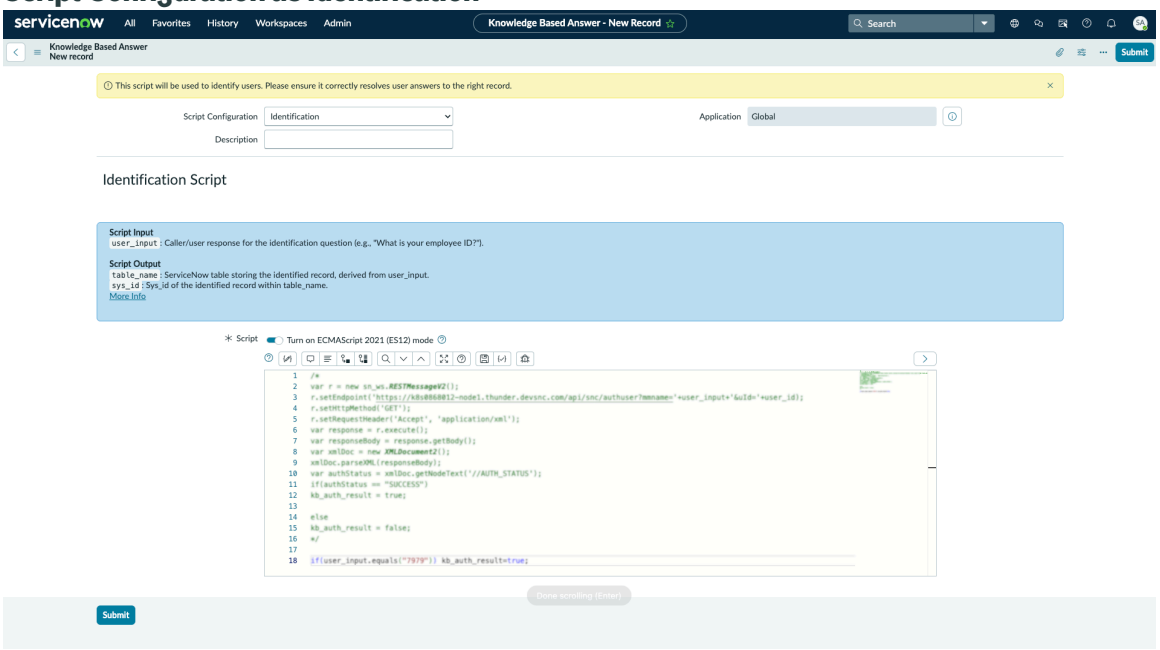

c. Script Configuration as **Authentication**: Validate against an external system during authentication:

> **Note:** KBA Answers - **Script Configuration**: Works for only users with `snc_external` role users.

External source fields (Authentication)

Field	Description
Script	Define the custom script to verify the caller's answer against an external system and return a pass or fail result. Script input: ▪ `user_id`: The sys_id of the user being authenticated. ▪ `user_input`: The caller's answer. ▪ `kba_session_context`: a JSON object containing questions presented to the caller and the answers they provided, captured across identification and authentication steps in the current session. Only questions configured with a scripted answer are included. This allows the script to reference prior responses when validating the current answer — for

Field	Description
	example, to cross-check answers or apply multi-field logic. Example: `{ "employee_id":"EMP12345", "date_of_birth":"1Jan2026" }` Script output: `kb_auth_result`, set to `true` if the answer matches, `false` if it does not. For script execution time limits, see System Properties.

Script Configuration as Authentication

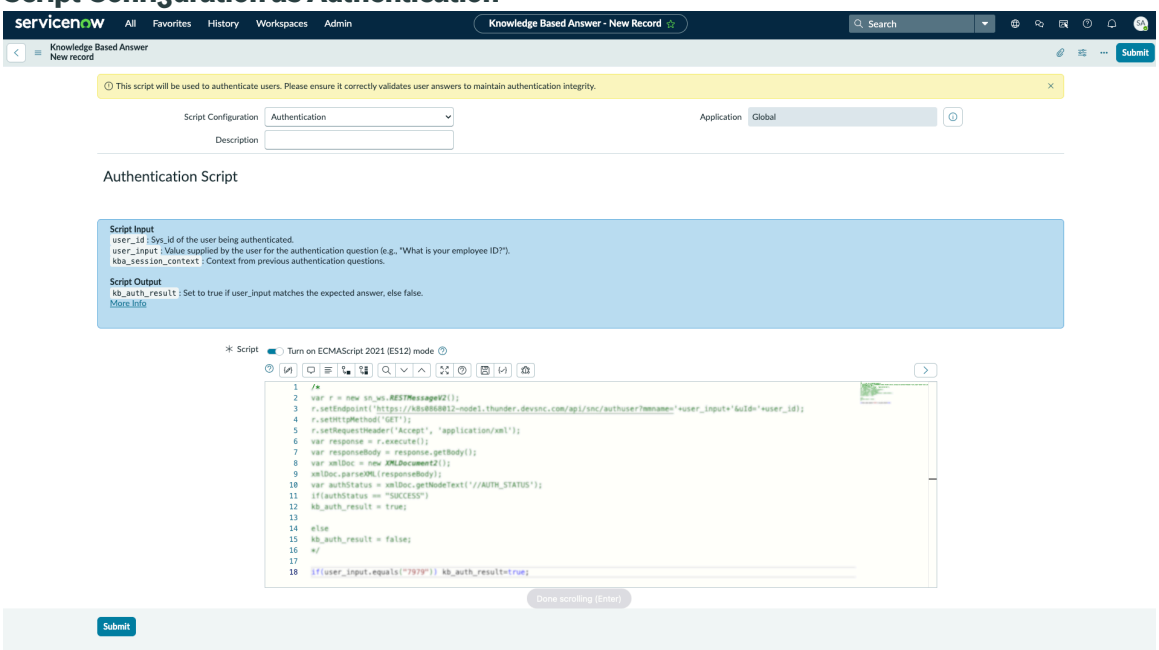

Note:

- If a question uses scripted answers, only one answer is permitted.
- If the answers are non-scripted, then multiple answers are allowed for that question.

However, for any given question, all associated answers must consistently follow the same approach: either all answers must support identification only, or all must support both identification and authentication. You can't mix answer types for the same question.

5. Select **Submit**.

Result
You're redirected to the Knowledge Based Answers list view. Verify if your answer is successfully added.

Knowledge Based Answers - list

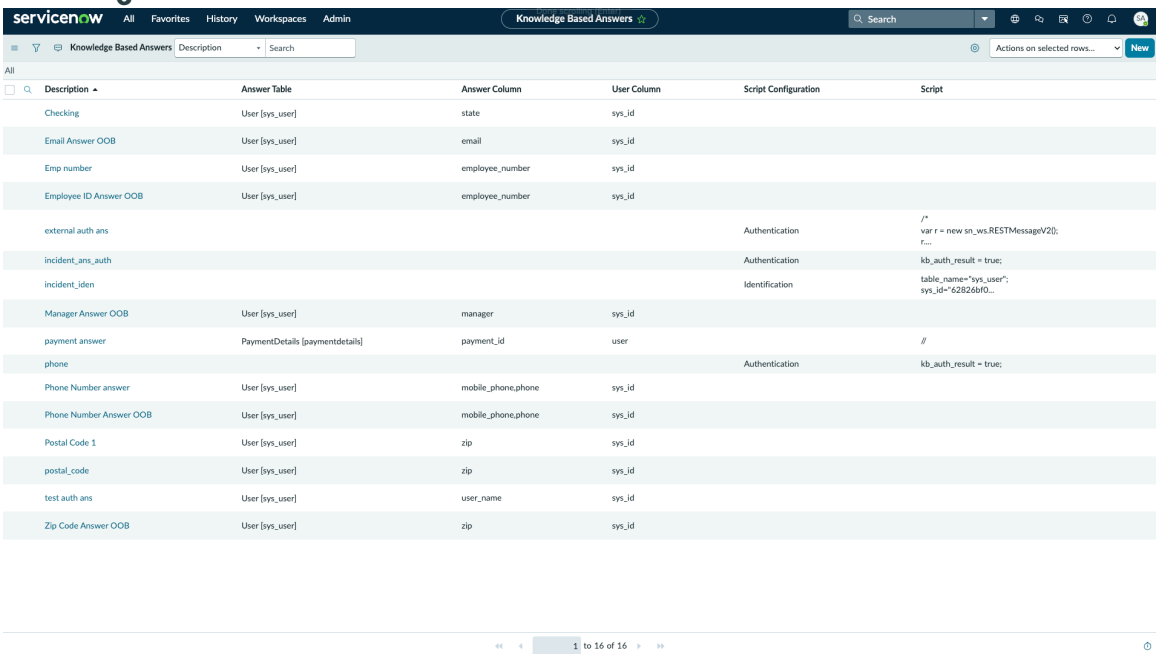

The following properties control the behavior of knowledge-based authentication (KBA), including security question validation, answer matching, and user identification settings.

System properties

System properties	Description	Default Value
`glide.auth_factors.kba.enable_user_column`	Controls whether the User Column field is mandatory when creating or updating an answer. **Note:** All answers mapped to the same question must follow the same pattern — either all supporting identification and authentication, or all supporting identification only.	`true`
`glide.auth_factors.kba.script_execution_timeout`	Controls the execution time limit in seconds for external source validation. Minimum: 1 second, Maximum: 30 seconds.	`15`

Map KBA questions to answers

Create knowledge-based questions and answer mapping to confirm the user's identity.

Before you begin

Role required: auth_factors_admin

Procedure

1. Navigate to **All** > **Authentication Factors** > **Knowledge Based Factor** > **Question Answer Mappings**.
2. Select **New** on the Knowledge Based Question Answer Mappings page.
3. Specify the following fields on the form:

Knowledge Based Question Answer Mappings

Field	Description
Question	Select the question that you would like to map with the answer. Example: `What is your business phone number?`
Application	Global application scope is selected by default.
Answer	Select the answer for mapping to the question. Example: `Business Phone Number`.

Knowledge Based Question Answer Mapping

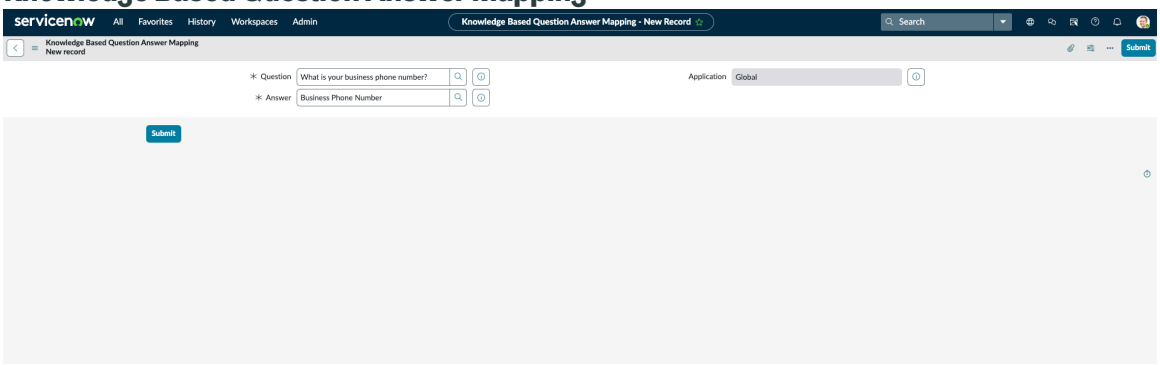

4. Select **Submit**.

Result

You're redirected to the Knowledge Based Question Answer Mappings list view. Verify if your mapping is successfully added.

Knowledge Based Question Answer Mappings - list

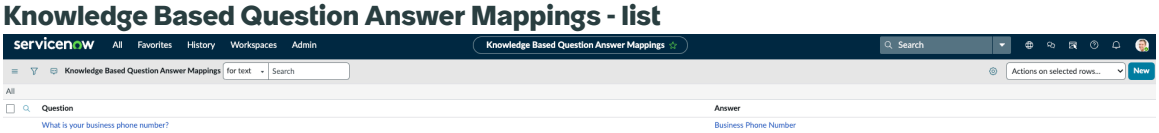

AI voice agent service mapping with KBA

Specify the questions used for caller identification and authentication with a specific AI voice agent service.

Before you begin

Role required: auth_factors_admin

Service mappings are created automatically when KBA questions are selected in AI Voice Assistant Designer. You can also create and manage mappings directly in this table. Changes made here are reflected on the Caller Verification screen in Assistant Designer instantly. To learn more about voice services and how to create and manage them, see Create an AI voice assistant.

Procedure

1. Navigate to **All** > **Authentication Factors** > **Knowledge Based Factor** > **Question Service Mappings**.
2. Select **New** on the Knowledge Based Question Service Mappings page.
3. Specify the following fields on the form:

Knowledge Based Question Service Mappings

Field	Description
Service Profile Table	Select the table for your service.
Application	Global application scope is selected by default.
Service Profile	Select the service profile. Example: Choosing a voice agent: **Now Assist Voice Deployment - 22**.
Question	Select the question to assign to this service mapping. **Note:** A validation error is returned for the following: ○ The selected question must have a corresponding question-answer mapping configured. If it does not, the question is cleared and a validation error is returned. ○ When a question with Type = **Authentication** is assigned a Service mapping with Type = **Identification**, the mapping action is aborted with an error message and a validation error is returned.
Type	Select the type for which you would use the service mapping. Options:

Field	Description
	○ **Authentication**: The question is used during the authentication phase. ○ **Identification**: The question is used during the identification phase.
Usage	Available when Type is set to **Identification**. Select the role of this question in the identification flow. Options: ○ **Primary**: Attempted first to identify the caller. ○ **Fallback**: Attempted if the primary identifier does not return a match.
Active	Select to set the configuration active.

Knowledge Based Question Service Mapping

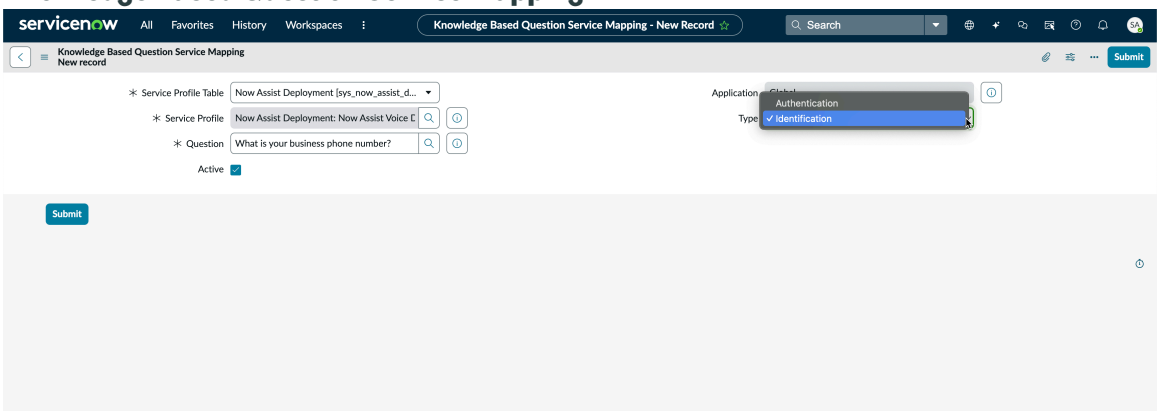

4. Select **Submit**.

Result

You're redirected to the Knowledge Based Question Service Mappings list view. Verify if your mapping is successfully added.

Knowledge Based Question Service Mappings - list

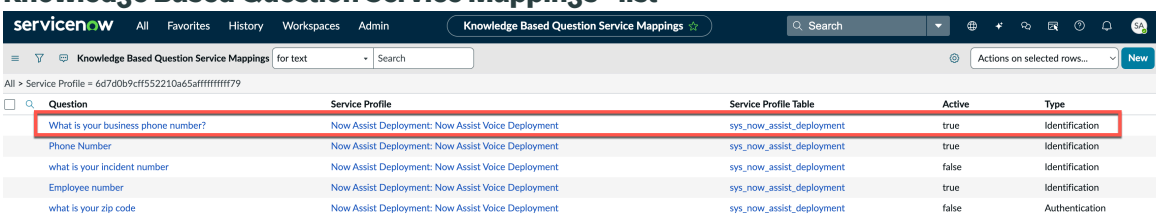

Certificate-based authentication

Certificate-based authentication lets you mutually authenticate user logins or inbound API requests using certificates from a trusted Certificate Authority (CA).

Note: Certificate Based Authentication is not supported on the On-Prem and edge encryption enabled instance.

Certificate-based authentication for user interface logins

Enable end users to use PIV (Personal Identity Verification) or CAC (Common Access Card) cards to log in to the ServiceNow AI Platform or Service Portal instead of using a user name and password. To set up mutual authentication for user interface logins, see Set up Certificate-based authentication.

After Certificate-based authentication is set up, end users can finalize their set up and log in. See Log in using Certificate-based authentication.

Certificate-based authentication for Inbound web services

Authenticate inbound requests to ServiceNow SOAP and REST APIs. To set up mutual authentication for inbound web services, see Set up Certificate-based authentication.

Set up Certificate-based authentication

Set up mutual authentication for either user interface-based logins or inbound web services.

Before you begin

Role required: admin

Check that your instance is using an ADCv2 load balancer. For more information, see theADCv2 Migration knowledge article. If your instance is not using the ADCv2 load balancer, contact Now Support.

Procedure

Set up Certificate-based authentication in order to:

- Allow end users to securely log in to the ServiceNow AI Platform or Service Portal using PIV or CAC cards. After certificate-based authentication is enabled, you can self-register the PEM certificate or an administrator can map the certificate for you. See Log in using Certificate-based authentication.
- Enable mutual authentication for inbound web services. Once Certificate-based authentication is set up, the system uses the provided certificates to mutually authenticate requests to access ServiceNow REST and SOAP APIs.

Activate Certificate-based authentication

You can activate the Certificate-based authentication plugin (com.glide.auth.mutual) for ServiceNow AI Platform if you have the admin role.

Before you begin

Role required: admin.

About this task

The following Tables are installed with Certificate-based authentication:

- sys_user_certificate
- sys_ca_certificate
- sys_ca_certificate_api_track

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Certificate-based authentication plugin (com.glide.auth.mutual) using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Register CA certificate

Register root certificates or intermediate certificates to make them available for authentication.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **Certificate Based Authentication** > **CA Certificate Chain**.
2. Click **New**.
3. On the form, fill in the fields:

Mutual Auth CA Certificate form

Field	Description
Name	Name to identify the certificate.
Expiration notification	Option to warn users when a certificate is about to expire.
Notify on expiration	List of users to be notified when the certificate expires.
Warn in days to expire	Number of days when a notification is sent to users before a certificate expires.
Active	Option to make the client certificate active.
Format	PEM
Type	Type of certificate. Options include: ○ **CA Cert**: The root CA certificate. Can also include intermediate certificates in the chain. CA certificates are automatically synced with the load balancer. Use this option

Field	Description
	when possible to avoid missing a required certificate in the chain. ○ **Intermediate Cert**: An intermediate certificate in the certificate chain. This certificate remains on the instance only and is not synced with the load balancer. Only use this option if you need to add an intermediate certificate to an existing chain.
Short description	Short description of the user client certificate.

Note: During the certificate upload, the read-only fields, **Valid from**, **Expires**, **Expires in days**, **Issuer**, and **Subject**, **Certificate Chain**, and **PEM Certificate** are extracted and auto-populated.

4. Click **Submit**.
5. **Optional:** Click **Validate Stores/Certificates** to validate the certificate.

Map PEM certificate to user

Map PEM certificates to users to enable them to log in using PIV or CAC cards or to authenticate inbound requests. You can map multiple PEM certificates to a user.

Before you begin

- Role required: admin
- Make sure that you have the Privacy Enhanced Mail (PEM) certificate of the user.

Note: After the Map PEM certificate to User configuration, the "verify certificate" will fail. This is because the PEM certificate is not stored.

Procedure

1. Navigate to **All** > **Certificate Based Authentication** > **User to Certificate Mapping** and click **New**.
2. On the form, fill in these fields:

User Client Certificate form

Field	Description
Name	Name of the user client certificate.
Expiration notification	Option to warn users when a certificate is about to expire.
Warn in days to expire	Number of days when a notification is sent to users before a certificate expires.
Notify on expiration	List of users to be notified when the certificate expires.
Active	Option to make the client certificate active.
User	User who is mapped to the client certificate. The system receives the client certificate from either the inbound request or certificate registration, and then uses the

Field	Description
	user designated in this field to initiate a session to execute the request.
Short description	Short description of the user client certificate.
Format	Privacy Enhanced Mail (PEM) format is a base-64 encoded Distinguished Encoding Rules (DER) certificate.
Type	Client cert. This field is read only.

Note: During the certificate upload, the read-only fields, **Valid from**, **Expires**, **Expires in days**, **Issuer**, and **Subject** are extracted and auto-populated.

3. Click the attachments icon and upload the certificate.
4. Click **Submit**.
   The certificate is validated and mapped to the specified user if the certificate is from a trusted Certificate Authority (CA).

Configure Certificate-based authentication properties

Use system properties to enable or disable certificate-based authentication features.

Before you begin

Role required: sso_config_admin

Procedure

1. Navigate to **All** > **Certificate Based Authentication** > **Properties**.
2. On the form, fill in the fields:

Certificate Based Authentication Properties form

Property	Description
Enable certificate based authentication	Option to enable to Certificate-based authentication for both user interface logins and inbound web services. Default: true **Note:** On the Portal pages, use the Form Layout to add the field to the form and then enable the property.
Show 'Log in with PIV/CAC' option in login screen	Displays the **Log in with PIV/CAC card** option on the login screen. Allows users to log in using Certificate-based authentication using the user interface. Default: false
Enable auto-redirect for certificate based login	Determines whether to require that the user click **Log in with PIV/CAC card** after selecting a registered certificate and entering their PIN. Activate to automatically log in the user after they select a registered client

Property	Description
	certificate and enter their PIN. Deactivate to require that the user click **Log in with PIV/ CAC card** after they select a registered client certificate and enter their PIN. Default: false

Log in using Certificate-based authentication

After your administrator sets up Certificate-based authentication, you can register the client certificate and log in using your PIV (Personal Identity Verification) or CAC (Common Access Card) card.

Register client certificate for your PIV or CAC card

Before you log in to ServiceNow AI Platform using your PIV or CAC card, you must register the client certificate of your PIV or CAC card. If you are not able to register the client certificate, contact your administrator. Your administrator can also register the client certificate of your PIV or CAC card.

Before you begin

- Make sure that Certificate-based authentication is enabled.
- Make sure that a card reader is connected to your computer and your PIV or CAC card is inserted.
- Role required: none

About this task

If you need an admin to register your client certificate, see Map PEM certificate to user.

Procedure

1. Log in to ServiceNow AI Platform using your user name and password.
2. Click your name from the **User Menu** and select **Profile**.
3. From the **Related Links**, click **Register client certificate**.

   If a valid is certificate is available, the following message displays:

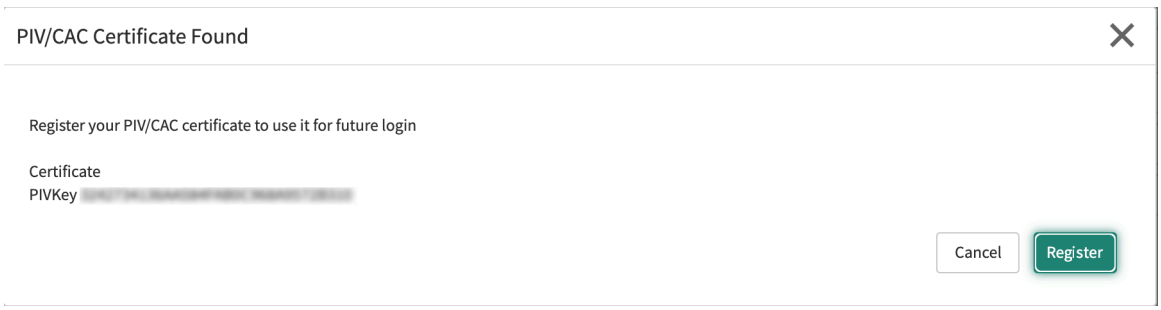

4. Click **Register**.

   After the registration is successful, the following message displays:

   ```
   The PIV/CAC certificate has been successfully registered and
   linked to the user account.
   ```

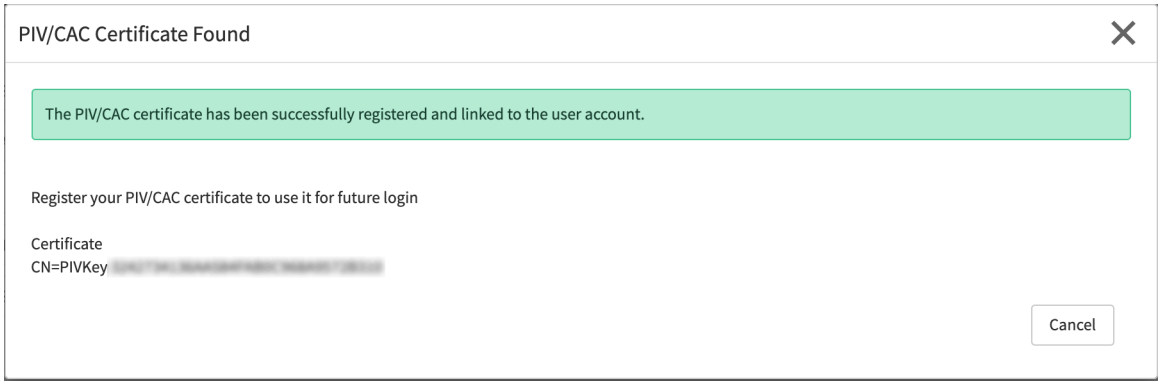

The next time you log in to your ServiceNow AI Platform, you can log in using your PIV or CAC card. For more information, see Log in to ServiceNow AI Platform using PIV or CAC card.

Log in to ServiceNow AI Platform using PIV or CAC card

You can log in with your PIV or CAC card instead of user name and password when Certificate-based authentication is enabled on ServiceNow AI Platform.

Before you begin

- Role required: none
- Make sure that Certificate-based authentication is enabled.
- Make sure that a PIV or CAC card reader is connected to your computer.
- Make sure that a client certificate of your PIV or CAC card is mapped to you. For more information, see Register CA certificate.

Procedure

1. Insert your PIV or CAC card into the card reader.
2. Navigate to your instance in a browser.
   Browser prompts for a PIN for your PIV or CAC card.
3. Enter the PIN of your PIV or CAC card in the browser prompt.

   Note: If you forget your PIN, contact your administrator.

4. If you enter a correct PIN, the browser displays a prompt to select a certificate.

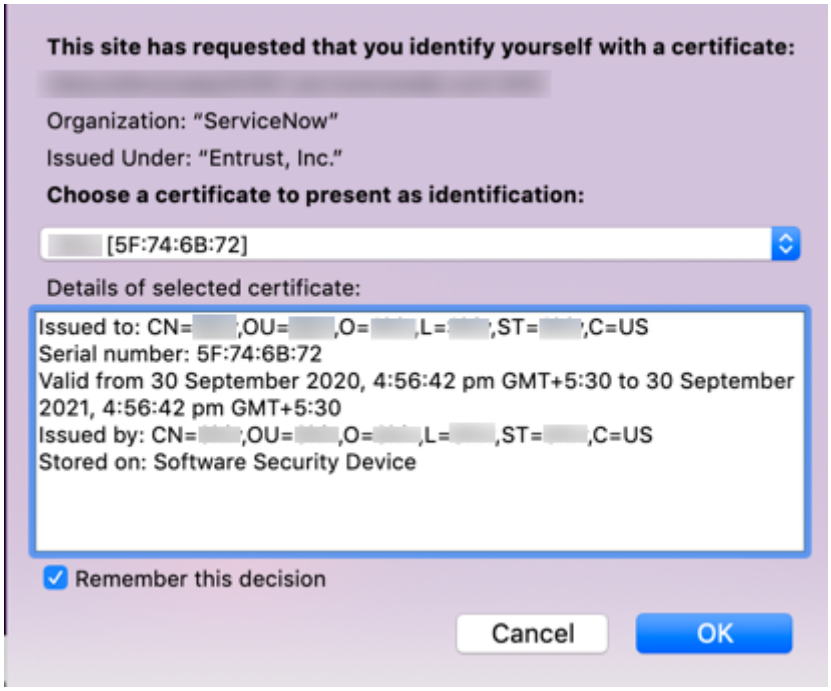

5. Select a certificate from the browser prompt.

   If the certificate is valid and mapped to you, you are redirected to the login page.

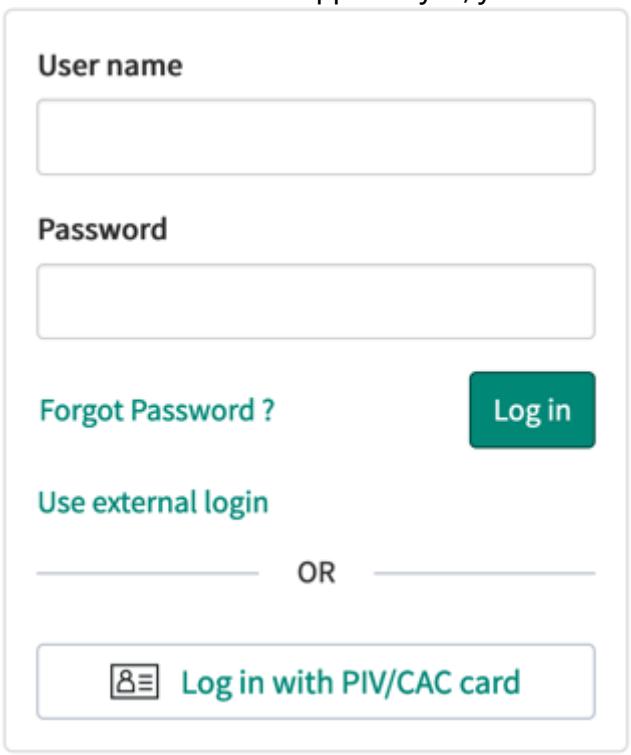

6. Click **Log in with PIV/CAC card** button.

   To log out of the ServiceNow AI Platform, you must remove the PIV or CAC from the card reader and then close the browser.

Manage your client certificates

View and delete client certificates associated with your account.

Before you begin

- Make sure that Certificate-based authentication is enabled.
- Role required: none

Procedure

1. Log in to ServiceNow AI Platform using your user name and password.
2. Click your name from the **User Menu** and select **Profile**.

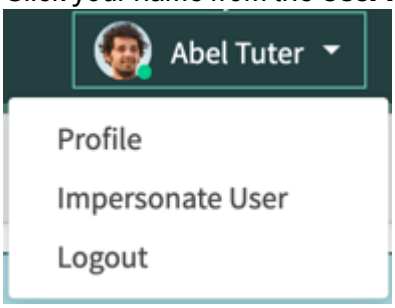

3. From the **Related Links**, click **Manage your client certificates**.
   The User Client Certificates [sys_user_certificate] table opens and displays certificates that are associated with your account.
4. View or delete certificates associated with your account.

Custom instance URLs

You can enable your ServiceNow instance to be accessible from a company-branded or custom URL.

Custom URL overview

You can use custom URL to enable ServiceNow instance to be accessible from one or more company-branded or custom URLs.

Custom URLs can be associated with specific portal. For example, admin can define `http://support.acme.com` for CSM portal and `http://hr.acme.com` for HR portal. In such situation there is a need to redirect users to different IdPs for authentication on the basis of the custom URL that users are accessing.

Important:

- Do not create a custom URL with more than 100 domains per instance.
- You must delete custom URL record from ServiceNow instance first and then delete any Domain Name Server (DNS) entries from the DNS server.
- Any deletion of DNS entry from DNS server prior to deletion of custom URL record from ServiceNow instance would result in blocking the deletion of other corresponding custom URL records from ServiceNow.

From Tokyo, you can allow users to auto-redirect to specified IdPs defined at the custom URL record.

Note: Custom URLs are not available for on-premise customers or developer instances. Also, the URL must be public-facing.

Only the owner of the top-level domain (TLD) or any subdomain can configure the custom URL to a DNS subdomain. For example, your instance might have the following designated URL and additional custom URLs:

Custom URL example

Example URLs	Usage
`https://acme.service-now.com`	The initial domain name for Acme that came with the ServiceNow instance.
`https://support.acme.com`	A custom URL that associates to your ServiceNow instance. This URL is referred to as an alias (CNAME) of the initial domain name.
`https://US-support.acme.com`	A secondary custom URL that associates to a service portal on your instance. Your instance can support multiple custom URLs to the same service portal.

Custom URL considerations outside of your instance

Before you can associate a custom URL, you must own (or purchase) a URL through a domain provider. There are also specific configurations necessary before you can create and associate a custom URL on your instance.

Custom URL configurations

Configuration items	Description
Set the CNAME with the provider	The CNAME record must be set as the ServiceNow instance URL.
Determine your dedicated VIP status	The status of VIP.

Note: When deleting or updating CNAME records which point to your instance, you must follow this sequence to avoid dangling records in the instance. First, delete the CNAME records from your instance and then remove or update the CNAME setting at your DNS provider.

Activate custom URLs

Enable custom URLs to be set up on your ServiceNow instance. You can activate the Custom URL plugin (com.snc.customurl) if you have the admin role.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Plugins**.
2. Find and click the **Custom URL** plugin.

   Note: Do not select the **Custom URL - Internal** plugin, which is an internal component for scripted custom URL APIs.

3. On the Custom URL record, click the **Activate/Repair** related link.

4. In the Plugin Activation window, click **Activate**.
   When the Plugin Activation window reopens with a message that the plugin is activated, click **Close & Reload Form** to stay on this form.
5. In the Plugin Files related list, find the following property, and change the setting value:

Option	Description
glide.customurl.enabled	To enable custom URLs, set the value to **True**. By default, this property is set to **False**, which means that you cannot associate a custom URL. **Note:** To disable this feature, set this property back to **False**.

6. Click **Update**.

Set a custom URL as the instance URL

Add a custom URL to your instance configuration to use instead of your ServiceNow URL.

Before you begin

Role required: admin
You must activate the custom URL plugin and have purchased or registered a URL before adding the custom URL to your instance.

Procedure

1. Navigate to **All** > **Custom URL** > **Custom URLs**.
2. You should either:
   - Click **New** to associate a new domain name that is new to your instance.
   - Select a custom URL already set up to set as your instance URL.
3. Fill in the appropriate fields:

Custom URL fields

Field	Description
Domain Name	Fully qualified domain name (FQDN) of the custom URL. The FQDN is a CNAME redirect created in the name server record for the custom domain. **Note:** For example, in the name server for acme.com, you might create an entry: `support.acme.com 300 IN CNAME acme.servicenow.com`
Is Instance URL	Check box to enable this custom URL for all outbound URLs. Only one active custom URL can be the instance URL.

Field	Description
	To enable this setting for a custom URL, click **Set as Instance URL** on the URL record. Any previous custom URLs are then removed.
Status	Status of the custom URL record. If the status is **Active**, the custom URL is provisioned and ready to use.
Service Portal	Service portal that you want to use when you redirect users to your instance using the custom URL.
Identity Provider	Identity Provider for the custom URL enables you to allow users to auto-redirect to specified IdPs defined at the custom URL record.

A custom URL should activate within six hours on your instance.

Note:

- You must delete custom URL record from ServiceNow instance first and then delete any Domain Name Server (DNS) entries from the DNS server.
- Any deletion of DNS entry from DNS server prior to deletion of custom URL record from ServiceNow instance would result in blocking the deletion of other corresponding custom URL records from ServiceNow.

Custom URL with Identity Provider

Set your custom URL with the Identity Provider to enable the user to login with their IdP's.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Custom URL** > **Custom URLs**.
2. Click **New**.
3. Provide the details of the IdP in the **Identity Provider** field.
   For information about other fields, see Set a custom URL as the instance URL.

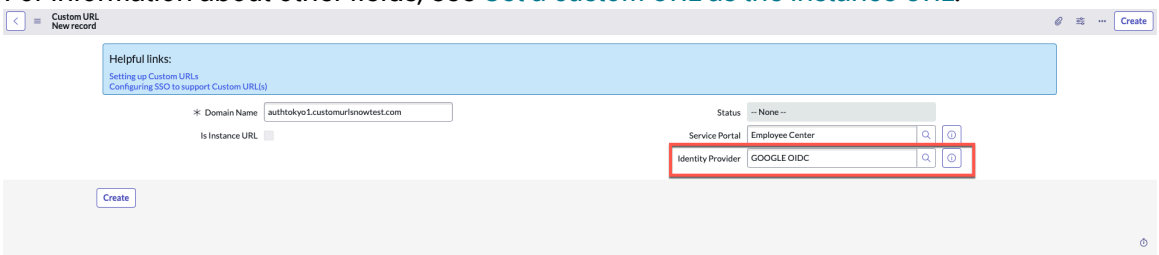

4. Click **Create**.

   The record is created and displayed on the Custom URL page.

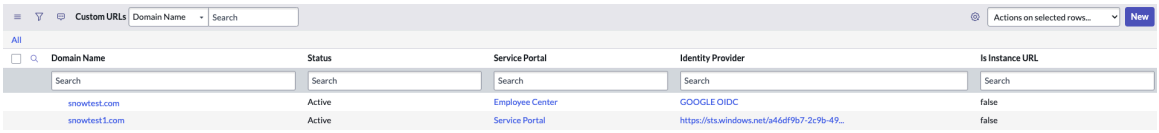

When the custom URL is accessed, the user is redirected to the Identity Provider that is configured. In this case, accessing the `snowtest.com`, the user is navigated to the **Employee Center**, and then redirected to the Google Identity Provider.

Note:

- If the Service Portal field is empty and the Identity Provider field is defined, when the user accesses the custom URL, the user is directly navigated to the configured Identity Provider.
- If both the Service Portal and Identity Provider field is defined, when the user accesses the custom URL with the Service Portal that is defined, the user is navigated to the configured Identity Provider.
- If both the Service Portal and Identity Provider field is defined, when the user accesses the custom URL with the different portal that is not defined, the user is navigated to the auto-redirect Identity Provider if configured on the instance.

5. Use the credentials to login to the application.

Custom URL datacenter job information

Every custom URL that is associated to your instance has a corresponding ServiceNow datacenter job which runs and shows URL information that is pertinent to your instance as described in the table.

Job field	Description
Job ID	Unique ID of the job that checks the domain of the custom URL.
Last Run At	Date and time when the job last ran.
Payload	List of domains or custom URLs that were sent to the datacenter for CERT provisioning.
Poll Count	Number of times that the results have polled for this job.
Result	Verifies and validates each domain or custom URL sent in a payload.
Status	Status of the datacenter job.

Generate SP metadata for SAML/SSO custom URL installations

A SAML or SSO installation needs the SP metadata generated for the IdP before the custom URL instance generates.

Before you begin

Role required: admin

The IdP needs SP metadata for the instance to authenticate and forward requests.

Note: Adding the Assertion Consumer Service URL (SP login URL) might be different for each IdP (Azure, ADFS, or Okta).

Procedure

1. Choose your installed SSO plugin:

Option	Description
Multi-Provider SSO	Navigate to **Multi-Provider SSO** > **Identity Providers**. Choose an IdP and click the **Generate Metadata** button. The integration automatically generates the instance's SP metadata from the system property settings.
SAML 2 SSO	Navigate to **SAML 2 Single Sign-on** > **Metadata**. The integration automatically generates the instance's SP metadata from the system property settings.

2. Copy the SP metadata in the text box.

   For example:

```
<EntityDescriptor xmlns="urn:oasis:names:tc:SAML:2.0:metadata"
 entityID="https://yourinstance.service-now.com">
  <SPSSODescriptor AuthnRequestsSigned="false"
 WantAssertionsSigned="true"
 protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protoc
ol">
  <SingleLogoutService
 Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect"
 Location="https://yourinstance.service-now.com/navpage.do" />

 <NameIDFormat>urn:oasis:names:tc:SAML:1.1:nameid-format:emailA
ddress</NameIDFormat>
  <AssertionConsumerService isDefault="true" index="0"
 Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"
 Location="https://yourinstance.service-now.com/navpage.do" />
  <AssertionConsumerService index="1"
 Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"
 Location="https://yourinstance.service-now.com/consumer.do"/>
 </SPSSODescriptor>
</EntityDescriptor>
```

3. Provide the instance SP metadata to the IdP.
   For example, SSOCircle allows a user to provide the SP metadata online.
4. **Optional:** To set up custom URLs in Azure:
   a. Go to **App Registrations**.
   b. Select **All apps** from the menu.
   c. Select the **ServiceNow App**.
   d. Click settings to configure the URL.
5. **Optional:** To set-up custom URLs in Okta:

a. Create two ServiceNow UD Okta Applications.

b. One Okta Application for the "service-now.com" instance URL.

c. One Okta Application for the custom URL.

Note:

- Disable the **Disable Force Authentication** within the Okta configuration for the **Test Connection** to run successfully.
- If you're testing the Identity Provider record associated with the base URL, ensure you've to login to the instance with the base URL.
- If you're testing the Identity Provider associated with the Custom URL, ensure you've to login to the instance with the Custom URL.

6. **Optional:** To use OAuth authentication, set up the redirect URL as all the registered custom URLs in the OAuth application endpoint configuration for the external client applications. The redirect URL is synonymous with the callback URL that the authorization server redirects to.
7. **Optional:** To use Google reCAPTCHA service, set up an API key pair.

Custom URL errors and fixes

A list of common errors and associated fixes for a custom URL setup and configuration.

Errors during setup

Error message	Fix
Unable to create a Custom URL. Try again later.	There might be an issue outside of your control interfering with the custom URL creation. Run at a different time before contacting support. **Note:** It usually takes 30 minutes to create a custom URL. If it is taking longer, contact .
Unable to submit your new Custom URL request because another Custom URL request for your instance is still in progress.	Check the status on your Custom URL Jobs before you submit a new request.
You must clear the following properties before the instance URL can be set: Glide Proxy, Glide Servlet.	
The provisioning of <custom_domain> is still in progress. This process can take up to X hours. Your instance administrator will receive a notification when this process completes.	After the provisioning on an instance starts for a custom URL, you must wait for the process to complete before the status changes.
<custom_domain> is now set as the new instance URL. <base.service-now.com> is	Only one URL can be designated as the instance URL. The other URLs that are associated with your instance can be

Error message	Fix
still in service, but all new URLs, such as in notifications, will use <custom_domain>.	active, but only the instance URL can use notifications.
The Custom URL <custom_domain> is set to be immediately removed from the instance configuration and revert back to <base.service-now.com>. <custom_domain> continues an association to this instance as long as the CNAME record in your DNS is set.	This message confirms that you intend to change your custom URL. By accepting this confirmation, you initiate the change of the URL to your instance. Any URL that is on your Domain Name list, as a custom URL, can be active unless you remove it from the CNAME record on your provider.
You cannot modify glide.servlet.uri. This property is set by Custom URL.	
You cannot modify glide.proxy.host. This property is set by Custom URL.	
The CNAME record for <custom_domain> does not point to <base.service-now.com>.	The configuration on the URL provider side is not correct. Check your configuration on the CNAME record.
Missing CNAME record for <custom_domain>.	The CNAME record must be configured from your URL provider before the custom URL can be set for your instance.

Installation exits

Installation exits are customizations that exit from Java to call a script before returning back to Java.

Note: Functionality described here requires the **Admin** role.

Available installation exits

Navigate to **System Definition** > **Installation Exits**. Some installation exit names (Login, Logout, ValidatePassword, ExternalAuthentication) are reserved and cannot be changed. Other installation exits can override these with custom script that replaces the script in the default installation exit.

The following installation exits are available in the base system:

Installation Exit	Description
Login	Takes a username and password pair and authenticates with the user object
Logout	Takes the user to the welcome page upon signing out; can be overridden by LogoutRedirect
LogoutRedirect	Takes the user to a specified URL upon signing out
ExternalAuthentication	Authenticates using header, parameter, or cookie; can be overridden by DigestSingleSignOn and PGPSingleSignOn
DigestSingleSignOn	Authenticates using header, parameter, or cookie and decrypts Digest encryption
PGPSingleSignOn	Authenticates using header, parameter, or cookie and decrypts PGP encryption

Installation Exit	Description
ValidatePassword	Active by default, starting with the Helsinki release; allows customers to define their own password validation; can be overridden by ValidatePasswordStronger
ValidatePasswordStronger	Requires passwords be at least 8 characters long and contain a digit, an uppercase letter, and a lowercase letter
GetIntegrationSessionTimeout	Implements the default integration session timeout behavior.

Login modifications

The following modification to the **Login** installation exit sets each user's session timeout value as the user is logging in. In this particular example, if the user name is *admin*, the session is set to timeout in 30 seconds.

```
gs.include("PrototypeServer");

var Login = Class.create();
Login.prototype = {
 initialize : function() {
 },

        process : function() {
          // the request is passed in as a global
          var userName = request.getParameter("user_name");
          var userPassword =
 request.getParameter("user_password");

          var authed = GlideUser.authenticate(userName,
 userPassword);
          if (authed) {
             //
 ************************************************************
             // customization - if the userName == admin, set
 the session
             // timeout to be 30 seconds. You can implement your
 own
             // session timeout algorithm here by checking to
 see if a user
             // belongs to a certain group or has a certain
 role.
             // Values of setMaxInactiveInterval exceeding 1440
 minutes are
             // treated as one day (1440 minutes).

           if (userName == "admin") {
              request.getSession().setMaxInactiveInterval(30);
             }
             //
 ************************************************************
             return GlideUser.getUser(userName);
          }

          this.loginFailed();

          return "login.failed";
```

```
        },

        loginFailed : function() {
          var message = GlideSysMessage.format("login_invalid");
          var gSession = GlideSession.get();
          gSession.addErrorMessage(message);

          var userName = request.getParameter("user_name");
          EventManager.queue("login.failed", "", userName, "");
        }

}
```

Session timeout can also be set according to IP address.

```
gs.include("PrototypeServer");

var Login = Class.create();
Login.prototype = {
 initialize : function() {
 },

        process : function() {
          // the request is passed in as a global
          var userName = request.getParameter("user_name");
          var userPassword =
 request.getParameter("user_password");

          var authed = GlideUser.authenticate(userName,
 userPassword);
          if (authed) {

          //
 ************************************************************
          // customization - if the user is logging in from a
 particular IP
          // range starting with XXX.XXX you can implement your
 own
          // session timeout algorithm here by checking the
 login IP
          //
          // Values of setMaxInactiveInterval exceeding 1440
 minutes are
          // treated as one day (1440 minutes).

          var clientIP =
 gs.getSession().getClientIP().toString();

          // if client IP starts with specified range
          if (clientIP.indexOf('XXX.XXX') == 0) {
             // set to 10 hours
             request.getSession().setMaxInactiveInterval(60 * 60
 * 10);
          }
          //
 ************************************************************
```

```
            return GlideUser.getUser(userName);
          }

          this.loginFailed();

          return "login.failed";
        },

        loginFailed : function() {
          var message = GlideSysMessage.format("login_invalid");
          var gSession = GlideSession.get();
          gSession.addErrorMessage(message);

          var userName = request.getParameter("user_name");
          EventManager.queue("login.failed", "", userName, "");
        }

}
```

Related topics

IP range based authentication

One way to secure a web-based application is to restrict access based on the IP address.

You can block access to a specific address or range of addresses that you suspect belong to malicious individuals. The instance allows you to control access by IP address.

Note: Use the Adaptive Authentication (AA) pre-authentication context policy to enforce IP based authentications and restrictions for additional capabilities. For more information, see Adaptive authentication.

Notes and Limitations:

- The system won't let you lock yourself out, so if you try to add a rule such that your current address would be locked out, the system warns you and refuses your insert.
- If you're inside of a corporate intranet, be very careful about setting up your IP rules. The IP address you see on your own computer (like 10.10.10.25) generally bears no relationship to the IP address you will actually appear as out on the internet. Your company likely proxies and/or NATs your address into a predictable set of outbound addresses which you will likely need to ask your network team about.
- A user whose access is restricted based on an access rule gets a 403 error on their browser.
- Restricted users do not use transactions, semaphores, or count towards any server resource counts.
- This feature does not supersede or override your existing access control rules if, for example, you're running a VPN to our data center. It's an additional check that must be met in addition to any access controls we may have set up on your PIX.
- Allow rules always supersede deny rules. So if an address is both allowed (by one rule) and denied (by a second rule) it is, in fact, allowed.
- Asterisks and CIDR blocks are not currently supported.
- Regarding forwarded proxy addresses, the allow rules are applied to each address in the chain and then the deny rules are applied to each address in the chain if none of the allow rules matched.

- IP range based authentication can effect the transfer of update sets. If IP address access control is enabled on the source instance, add the IP addresses of all application nodes supporting your instance as exceptions.

  Note: To find your instance IP information, Log in to ServiceNow - NOW Support, and Search for the **My IP Information** service catalog item.

Note: To learn more about the *com.snc.ipauthenticator* and *glide.ip.authenticate.strict* properties, which restrict instance access to specific IP ranges, see the following topics in Instance Security Hardening Settings:

Restrict access to specific IP ranges plugin [Updated in Security Center 1.3]

IP Address Access Control

Apply an IP access control to outbound traffic, inbound traffic, or bidirectional traffic. The system only blocks an IP address if a matching Deny rule exists and no matching Allow rule exists. By default, there are no restrictions on access to your instance.

Before you begin

Note: Use the Adaptive Authentication (AA) pre-authentication context policy to enforce IP based authentications and restrictions for additional capabilities. For more information, see Adaptive authentication.

Role required: admin

Important: IP Address Control access control rules can affect the MID Server outbound traffic. If you configure the denylist, add the MID Server IP addresses as allowlist exceptions to avoid connectivity disruptions. For MID Server IP requirements, see MID Server system requirements

Procedure

1. Navigate to **All** > **System Security** > **IP Address Access Control** to see a list of your IP access controls.
   You might have to activate the IP Range Based Authentication [com.snc.ipauthenticator] plugin.
2. Complete the form.

   Note: To find your instance IP information, Log in to ServiceNow - NOW Support, and Search for the **My IP Information** service catalog item.

Field	Description
MID Server traffic	IP access control rules apply to MID Server outbound connections. When using rules, add the MID Server host IP addresses to allowlist to avoid blocking MID Server traffic. For a list of MID Server IP requirements, see MID Server system requirements
Type	Type of access control rule to include.

Field	Description
	○ *Allow*: Any IP address in this range can interact with this instance. ○ *Deny*: Any IP address in this range cannot interact with this instance unless it is listed in an Allow rule. Also, when adding deny rules, you cannot deny your own public IP address or your instance does not update a deny rule. **Note:** To support maintenance, upgrades, and Customer Service and Support, some ServiceNow internal IPs cannot be blocked by Deny rules.
Direction	Direction of the IP access control rule. ○ **Inbound**: Choose Inbound to allow or deny inbound transactions. These are transactions initiated from outside of your instance. ○ **Outbound**: Choose Outbound to allow or deny outbound transactions. These are transactions initiated from within your instance. ○ **Bidirectional**: Choose Bidirectional for the configuration to apply for both Inbound and Outbound.
Active	When selected, the form is active.
Description	Description of the access control.
Range Start	Starting range of IP addresses to allow or deny. **Note:** These rules also affect transferring update sets. To ensure that IP address access control does not cause update sets to fail, add the target instance as an exception.
Range End	Ending range of IP addresses to allow or deny. **Note:** To limit access to specific VPN addresses only, enter a Deny range of **0.0.0.0** through **255.255.255.255** into the **Deny** field, and only enter the specific allowed VPN ranges.

3. Click **Submit**.

Find denied IP addresses

Find Denied IP addresses in the instance's node log files.

Before you begin

Role required: admin.

About this task

Log entries for blocked IP address appear as follows: `2015-10-21 18:37:43 (175) http-30 WARNING *** WARNING *** Security restricted: Access restricted (xx.xx.xxx.xxx not authorized)`.

Note: Denied IP addresses are the viewable instance's node log files, not viewable from the system logs.

Procedure

1. Navigate to **All** > **System Logs** > **Utilities** > **Node Log File Browser.**.
2. Browse the logs by criteria, such as time period and message.
3. You can also download log files when you know which log you are looking for, by navigating to **System Logs** > **Utilities** > **Node Log File Download**.

Lightweight Directory Access Protocol integration

An LDAP integration allows your instance to use your existing LDAP server as the primary source of user data.

Administrators integrate with a Lightweight Directory Access Protocol (LDAP) directory to streamline the user login process and to automate administrative tasks such as creating users and assigning them roles. An LDAP integration allows the system to use your existing LDAP server as the primary source of user data. Typically, an LDAP integration is also part of a single sign-on implementation.

The integration uses the LDAP service account credentials to retrieve the user distinguished name (DN) from the LDAP server. Given the DN value for the user, the integration then rebinds with LDAP with the user's DN and password. The password that the user enters is contained entirely in the HTTPS session. The integration never stores LDAP passwords.

The integration uses a read-only connection that never writes to the LDAP directory. The integration only queries for information, and then updates its internal database accordingly.

Note: For detailed information about setting up the integration, see LDAP integration setup.

Note: If your instance is using an LDAP integration and the Active Directory settings require users to reset their password upon login, your users will not be able to log in the instance. The instance cannot change any user's active directory password.

Features of LDAP integration

LDAP integration features include the following.

Scheduled LDAP refresh

A scheduled scan of your LDAP server is usually run once a night. It queries all applicable user records' attributes and compares them with the account on our servers. If there is a difference, we modify our user record with the changed attribute. The load placed upon the LDAP server during the refresh depends on how many records are queried, and the number of attributes being compared. We recommend scheduling the refresh during off-peak hours. A large refresh operation can affect other scheduled operations, such as running reports, and should be planned to minimize any conflicts.

LDAP listener

LDAP listener is our version of a persistent query (or persistent search). We issue a standing query for changes made to your LDAP server, and constantly listen for a response. Assuming your server supports a persistent search, any changes made to any of your applicable LDAP accounts are returned to the LDAP listener and sent to your instance within approximately 10 seconds. This is an extremely useful tool, allowing us to have a nearly real-time copy of your users' account details, without having to wait for the next scheduled refresh.

On-demand LDAP login

After an LDAP integration is established, the instance can allow new users to log in to the system even if they do not yet have an account on the instance. When a new user attempts to log in to the instance, the integration checks to see if this user has an account in the instance. If the integration does not find an existing user account, it automatically queries the LDAP server for the username that was entered. If a matching LDAP account is found, the integration tries to authenticate with the password the user entered. If the password is valid, the instance creates an account for the user, populates the account with all applicable LDAP information, and logs the user in to the instance.

On-demand login uses the LDAP User Import transform map. For more information on transform map requirements, see LDAP transform maps.

LDAP data population

Note: Functionality described in this integration is not available by default. This integration involves post-deployment customization performed by an experienced administrator or by ServiceNow professional services consultants.

An integration to the LDAP servers allows you to quickly and easily populate the instance's database with user records from the existing LDAP database. To prevent data inconsistencies, you can create, ignore, or skip incoming LDAP records.

You can also limit the data the integration imports by specifying LDAP attributes, thereby importing only the data that you want to expose to an instance. Typically, the LDAP attributes you specify become part of the integration transform map. If you do not specify any LDAP attributes, the integration imports all available object attributes from the LDAP server. The instance stores imported LDAP data in temporary import set tables, so the more attributes you import, the longer the import time. For more information, see Specify the LDAP attributes.

LDAP authentication

Use LDAP authentication to access using LDAP credentials.

When a user enters network credentials in the login page:

1. The instance passes the credentials to an LDAP server to find the instance.
2. With RDNs, it validates the user's DN string. It validates only if at least one of the LDAP OU configurations with `table=sys_user` has an RDN configured.
3. The LDAP server responds with an authorized or unauthorized message that the system uses to determine whether access should be granted.

By authenticating against your LDAP server, users access the platform with the same credentials that they use for other internal resources on your network domain. Also, you can reuse any existing password and security policies that are already in place. For example, the LDAP server may already have account lockout and password expiration policies.

When you enable LDAP, the system updates user records with these fields.

LDAP user record updates

Field	Description
Source	Identifies whether or not LDAP is used to validate a user. If the source starts with ldap, then the user is validated via LDAP. If the source does not start with ldap, then the password on the user record is used to validate the user upon login.
LDAP Server	Identifies which LDAP server authenticates the user when there are multiple LDAP servers.

Note: The system does not support LDAP password authentication through a MID Server. An instance must be able to directly connect with an LDAP server to support password authentication.

Understand LDAP integration

An LDAP integration allows your instance to use your existing LDAP server as the primary source of user data.

LDAP integration prerequisites

- The directory services server must be LDAP v3 compliant
- Inbound network access through the firewall must be allowed (to the LDAP server)
- External IP or Name of the LDAP server
- User credentials with read-only access
- For LDAPS, a PKI certificate

LDAP integration timing

LDAP integrations are usually done before the instance Go Live, but can be integrated at any time.

LDAP server data integrity

Some users are concerned about a third party (the instance in this case) making changes (writing) to your LDAP server. In an LDAP integration, your instance does not write to the internal LDAP directory. The instance queries for information, and updates its database accordingly.

No changes are made to the internal LDAP server by the instance. The service account is read only.

Most changes (including additions) to your LDAP server are available to the instance within seconds, depending on how many components of the full LDAP integration are in place.

To keep LDAP records synchronized, schedule a periodic scan of the LDAP server to pick up changes.

The instance does not synchronize department records. Users and group memberships are kept up-to-date by the LDAP Listener mechanism and a daily full LDAP Browse, but the instance does not delete any of these entries once they disappear from LDAP.

If an entry were to be deleted, the entire history would also get deleted, and any references to it would be cleared or deleted. Configuration Items (CIs), SLA Agreements, Software Licenses, Purchase Orders, and Service Catalog Entries all have a reference to Department, and if

Department is deleted, then those references get cleared. There are many references to Users, and so deleting a user would lose all history of what that user did. Currently, the decision to delete or not to delete is made by our customers.

Security

The connection is made from a single machine using a fixed IP address through a specific port on your firewall. Authentication is done with a read-only LDAP account of your choosing. You can use standard LDAP, or load the public side of an SSL certificate installed on your directory, in which case we can use LDAPS. To add another layer of security, we also offer the option of a point-to-point IPSEC VPN tunnel. Speak to your account manager for details and pricing.

Secure LDAP connections

Connection	Description
MID Server	To shield your LDAP server from external network traffic, install a MID Server on the local network and configure the system to communicate with the MID Server over a secure channel.
LDAPS	To establish an encrypted LDAPS connection, load the public side of your LDAP server's SSL certificate. The integration uses the certificate to encrypt all communication between the LDAP server and the instance.
VPN	To secure the LDAP server with an encrypted point-to-point IPSEC VPN tunnel, speak to your account manager for details and pricing.

Another security aspect to consider is the data shared in an LDAP integration. To limit the data exposed to your instance, specify attributes in your transform map. For more information, see LDAP transform maps.

Importing LDAP data to the instance

It is recommended that attributes are defined to import only required data. Defined attributes get mapped into the instance user database.

We cannot answer the question of which specific attributes are needed because this is determined by the scope of the project and business requirements.

Supported types of LDAP servers

The instance has successfully integrated with Microsoft Active Directory, Novell, Domino (Lotus Notes), and Open LDAP. We use JNDI to interface with the LDAP Server. As long as your LDAP server is LDAP v3 compliant, the integration is successful.

LDAP single-sign-on

Along with the data population functionality provided with the LDAP import, you can use the External Authentication functionality supported by the application to prevent your users from needing to sign on each time.

Multiple LDAP domains

The recommended method for handling multiple domains is to create a separate LDAP server record for each domain. Each LDAP server record must point to a domain controller for that domain. This means the local network must allow connections to each of the domain controllers.

After expanding to more than one network domain, it is critical that you identify unique LDAP attributes for the application usernames and import coalesce values. A common unique coalesce attribute for Active Directory is objectSid. Unique usernames may vary based on the LDAP data design. Common attributes are `email` or `userPrincipalName`.

Handling query limits

By default, Active Directory 2000/2003 has an LDAP query limit (maxPageSize) of 1000 objects to prevent excessive loads and denial of service attacks. We have two methods of dealing with this limit.

The default method is to break up the query to return less than 1000 objects at a time. For example, query only for object starting with the letter 'a', then query for 'b' objects. The more efficient method for large environments is to enable paging. Paging is supported by default on all Microsoft Active Directory servers. It automatically splits the results into multiple result sets, so we don't have to split up the query into multiple requests.

LDAP query type

If an LDAP password is supplied then a "Simple Bind" is performed. If no LDAP password is supplied then "none" is used, in which case the LDAP server must allow anonymous login.

LDAP authentication

We use provided service account credentials for LDAP to retrieve the user DN from the LDAP server. Given the DN value for the user, we then rebind with LDAP given the users DN and the provided password.

Password storage

The password that the user enters is contained entirely in their HTTPS session. We do not store that password anywhere.

Setting up LDAP authentication

These fields on the user record pertain to LDAP:

- **Source**: The Source field identifies whether or not a user is validated using LDAP. If the source field starts with "ldap", then the user is validated via LDAP. If the Source field does not start with "ldap", then the password on the user record is used to validate the user upon login.
- **LDAP Server**: The instance supports multiple LDAP servers, so the LDAP Server field determines which server should be used to authenticate the user.

LDAP integration requirements

Review the requirements for LDAP integration, which include a PKI certificate an LDAP compliant directory services server.

LDAP integration requires:

- An LDAP v3 compliant directory services server
  - Allows inbound network access through the firewall (to the LDAP server)
  - (Optional) Accepts anonymous login
  - (Optional) Supports paging for large LDAP queries

- The external IP address or fully-qualified domain name of the LDAP server. You can also use a MID server.
- A read-only LDAP account of your choosing
- For multiple domains, network access for each domain controller
- For LDAPS, a PKI certificate
- For LDAP listener, a Microsoft Active Directory server that supports persistent queries (ADNotify)

Supported LDAP servers

Using JNDI to interface with the LDAP server, the instance has successfully integrated with the following servers:

- Microsoft Active Directory
- Novell
- Domino (Lotus Notes)
- Open LDAP

LDAP query limits

By default, Active Directory 2000/2003 has an LDAP query limit (maxPageSize) of 1000 objects to prevent excessive loads and denial of service attacks. The system has two methods of dealing with this limit.

- The default method is to break up the query to return fewer than 1000 objects at a time. For example, query only for objects starting with the letter a, then query for b objects.
- The more efficient method for large environments is to enable paging, which is supported by default on all Microsoft Active Directory servers. Paging automatically splits the results into multiple result sets so the integration does not have to split up the query into multiple requests.

LDAP integration setup

Administrators can enable LDAP integration to allow sign-on of users from their company LDAP directory.

LDAP typically uses one of these types of communication channels.

LDAP communication channels

Connection	Description	LDAP import support?	LDAP authentication support?
MID Server connection	Communicates over HTTP on port 80 by default. This communication channel does not require a certificate. The connection between the MID Server and the instance is over HTTPS (port 443). You can use the MID Server to import data over LDAP, but you cannot use the MID Server for LDAP authentication. Proceed to Define the LDAP Server.	Yes	No

LDAP communication channels (continued)

Connection	Description	LDAP import support?	LDAP authentication support?
Standard LDAP integration	Communicates over TCP on port 389 by default. This communication channel does not require a certificate. Proceed to Define the LDAP Server.	Yes	Yes
SSL-encrypted LDAP integration (LDAPS)	Communicates over TCP on port 636 by default, This communication channel requires a certificate. Proceed to Install the LDAP X.509 SSL certificate to obtain and upload the certificate.	Yes	Yes
VPN connection	Communicates over an IPSEC tunnel. Purchase or create an IPSEC tunnel on your local network. Proceed to Define the LDAP Server.	Yes	Yes

If using a MID Server, the MID Server connects to the instance and the MID Server also connects to the LDAP server. In both cases, the MID Server initiates the connection:

1. First, the MID Server connects to the LDAP server via LDAP on Port 389.
2. Then, the MID Server initiates an HTTPS encrypted connection to the instance on Port 443 to push the data to the instance.

For more information about VPNs, Mid Servers, and LDAP, see You Don't Need A VPN Part II on the community.

Install the LDAP X.509 SSL certificate

You can install an X.509 certificate for your LDAP integration.

Before you begin

Role required: admin

Procedure

1. Purchase or generate an SSL certificate on your LDAP server.
2. Navigate to **LDAP** > **Certificate** and click **New**.
3. Fill in the form fields:

Field	Description
Name	The certificate name.
Expiration notification	Select this option to send a notification to the users selected in the **Notify on expiration** field. By default, this is enabled.
Notify on expiration	Select the users to revive the notification regarding certificate expiration. If no users are selected, the logged in user is added by default, along with the last two logged in users with the administrator role.
Warn in days to expire	The number of days before expiration that the instance send the notification. Enter a value of at least 20. Instances upgraded to Istanbul and later releases have this value set to 20 unless a greater value is specified.

Field	Description
Active	A check box to indicate that this certificate is active.
Format	The format of the certificate.
Type	The certificate container. The instance recognizes certificates from trust stores, Java keystore, and PKCS#12 keystores.
Valid from	The instance automatically adds the certificate valid from date to this field. Attach the certificate to the X.509 certificate record to populate this field.
Expires	The instance automatically adds the certificate expiration date to this field. Attach the certificate to the X.509 certificate record to populate this field.
Expires in days	The calculated number of days to expiration.
Short description	A description for the certificate.
Issuer	The instance automatically adds the certificate issuer to this field. Attach the certificate to the X.509 certificate record to populate this field.
Subject	The instance automatically adds the certificate subject to this field. Attach the certificate to the X.509 certificate record to populate this field.
PEM Certificate	Enter the value of the X509 certificate.

Note: The integration does not currently sign the certificate in communications between the instance and the IdP.

4. Click **Save**.

What to do next
Click **Validate Stores/Certificates** to test the trust store and certificate.

Define an LDAP server

Create a new LDAP server record in the instance.

Before you begin
Role required: admin.

Procedure
1. Navigate to **All** > **System LDAP** > **Create New Server**.
2. Fill in the form fields.

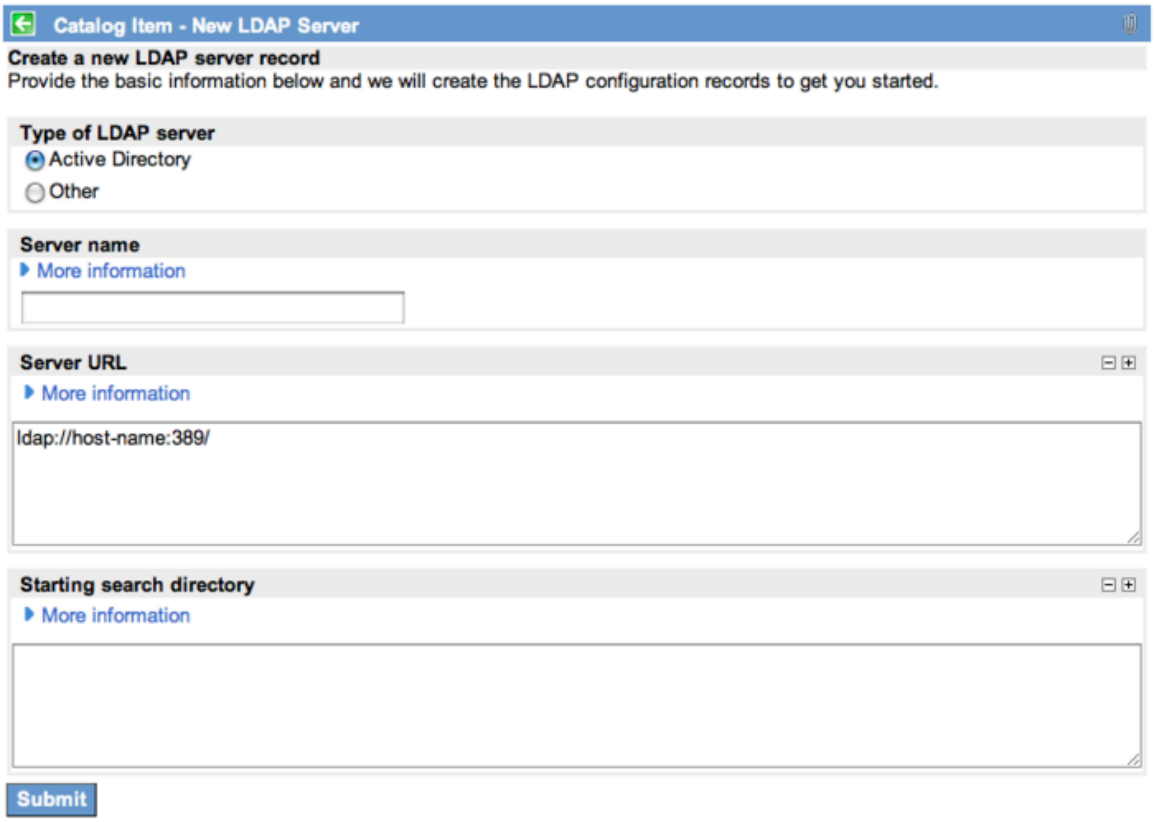

In the **Server URL** field, the valid URLs of all servers appear separated by a space. Servers are first ordered by operational status, with servers that are **Up** listed first, then ordered by the **Order** value that you specify. The first server listed is the primary LDAP server. The others are redundant servers.

Note: There is a slight delay between the change in the actual operational status and the display.

Alternatively, you can add a redundant LDAP server by navigating to an existing LDAP server record and inserting a row in the LDAP Server URLs embedded list.

3. Click **Submit**.

Note: You can also modify an existing LDAP server record by navigating to **System LDAP** > **LDAP Servers** and making the needed changes.

4. Make changes to the fields as necessary.

LDAP server form

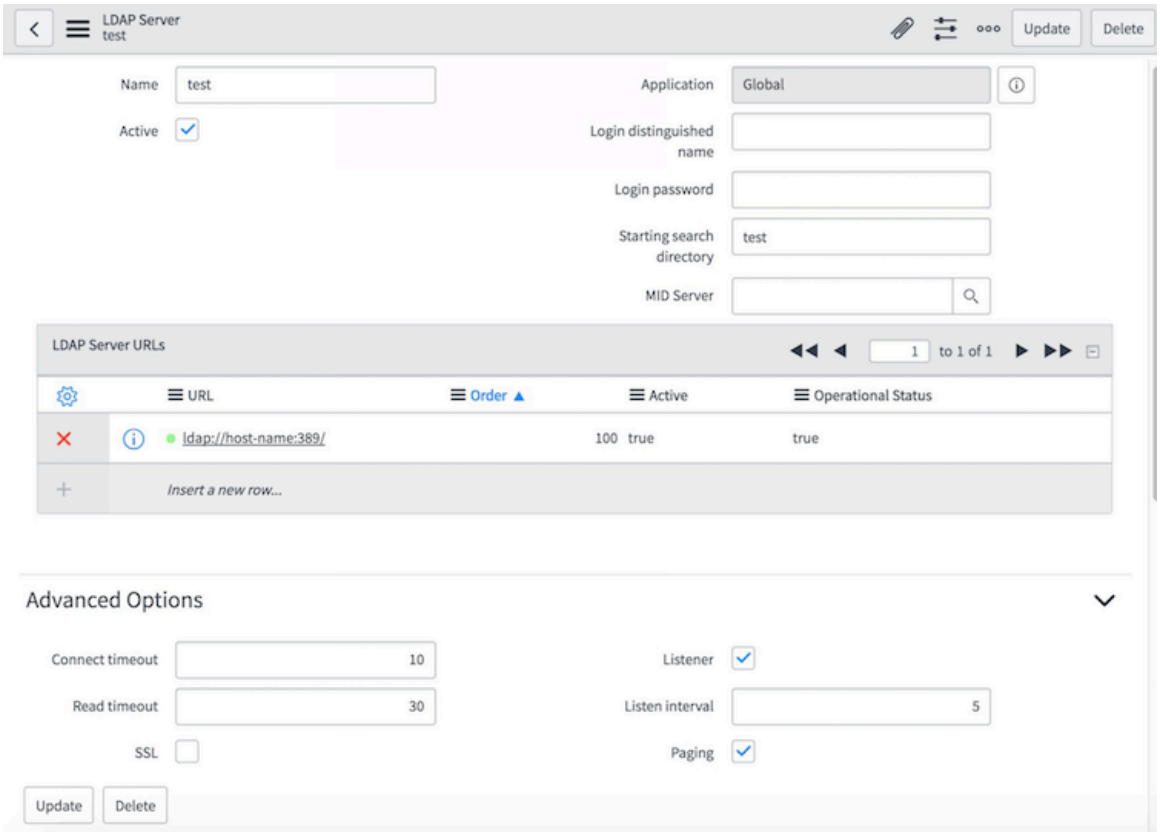

Field	Description
Name	Enter the name of the server.
Active	Select this check box if the server is active.
LDAP Server URLs	Enter the URLs of the primary and backup LDAP servers. Servers are first ordered by operational status, with servers that are Up listed first, then ordered by the Order value that you specify. The first server listed is the primary LDAP server. The others are redundant servers.
Server URL	Enter the URL of the server. Configure the form to add this field if necessary. It is a calculated read-only field that shows the list of LDAP servers that you can also see in the **LDAP Server URLs** field, separated by a space, and ordered by operational status and the order values of the URLs.
Login distinguished name	Enter the distinguished name (DN) of the user authenticating the LDAP connection. To access an LDAP directory server, the username must be in the full distinguished name format: `servicenow@service-now.com`
Login password	Enter the server's password.
Starting search directory	Enter the relative distinguished name (RDN) of the default search directory. All queries to this LDAP server will start from this RDN.

Field	Description
MID Server	Select the MID Server you want to use to connect to the LDAP server. Using a MID Server to establish an LDAP connection prevents you from having to expose the LDAP server to external network traffic. It also eliminates the need to establish a VPN tunnel between your LDAP server and ServiceNow data centers. **Note:** ○ The MID Server user must have the user_admin role in order to be able to read LDAP server configuration records. ○ The following are not available with the MID Server: ▪ LDAP authentication ▪ SSL connection
Connect timeout	If a MID Server is configured, the connection times out after 10 seconds, regardless of this setting. This setting is hard-coded and cannot be altered.
Read timeout	Specify the number of seconds the integration has to read LDAP data. The integration stops reading LDAP data after the connection exceeds the read timeout. If you enable an SSL connection, you can also set a read timeout value with the `com.glide.ssl.read.timeout` system property. If you enter timeout values for both this field and the system property, the lowest timeout value takes precedence.
SSL	Select this check box to require the LDAP server to make an SSL-encrypted connection. If you selected a MID Server, this field is not available. If you use an LDAPS integration and the default SSL port is 636, no further configuration is necessary; SSL is automatically enabled. If the LDAPS integration uses another SSL port, define the alternate SSL connection properties. **Note:** Be sure a network administrator configures the local firewall to allow the application server to access the LDAP server. If the LDAP server is located within an internal network, the firewall forwards (or NATs) the application server's IP address through the firewall on the correct port.
Listener	Select this check box to enable the integration to periodically poll Microsoft Active Directory servers or LDAP servers that support persistent search request control. Additionally, if you selected a MID Server, the listener functionality is available for that MID Server. See Enable an LDAP listener and set system properties for more information.
Listen interval (timeout value)	Specify the listener timeout value in the number of minutes that the integration listens for LDAP data with every connection. The integration stops listening for LDAP data after the connection exceeds the listen interval.

Field	Description
Paging	Select this check box to have the LDAP server split up LDAP attribute data into multiple result sets rather than submit multiple queries.

Note: If you provide an LDAP password, the integration performs a Simple Bind operation. If you do not provide an LDAP password, the LDAP server must allow anonymous login or the integration cannot bind to the LDAP server.

Result

When an LDAP Server record is set to active, the system automatically tests every connection to validate it.
Validations include:

- The LDAP server is accessible at the provided URL and port
- The LDAP server URL is properly formatted
- The login credentials are valid

Starting with the Fuji release, the system displays colored dots next to each server URL:

LDAP server connection icons

Color	Description
Green	The server if active and operational.
Gray	The server is neither active nor operational.
Red	The server is active but not operational.

LDAP server connection status

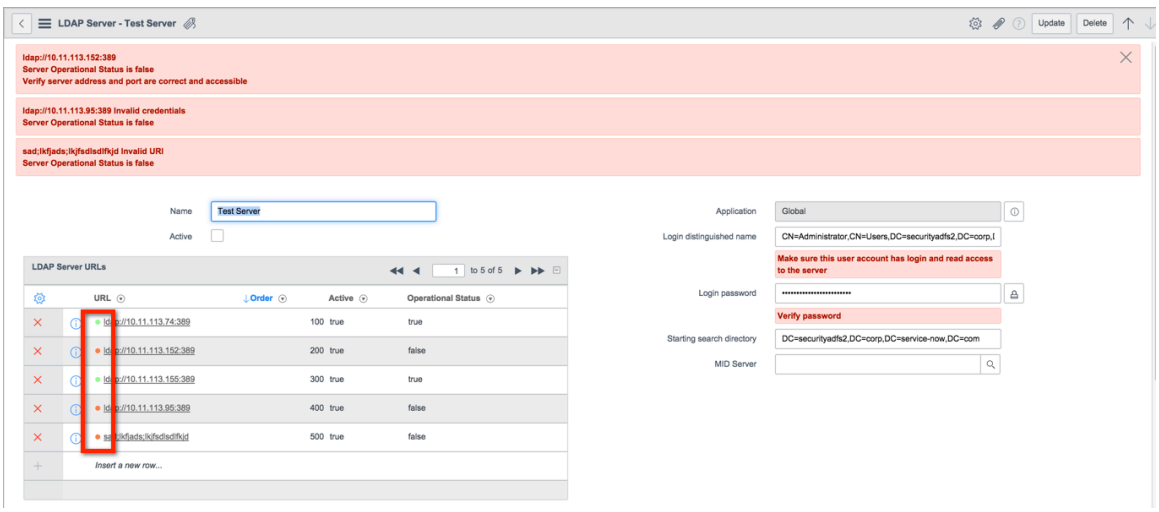

Enable an LDAP listener and set system properties

Enabling a listener is optional. If enabled, a listener notifies the system to process LDAP records soon after there is an update on the LDAP server.

Before you begin

Role required: admin.

About this task

A listener is a dedicated process that periodically searches for changes on the LDAP server.

The listener can be deployed on a Microsoft Active Directory server that supports persistent queries (ADNotify), or on an LDAP server that supports persistent search request control (with OID 2.16.840.1.113730.3.4.3).

If the LDAP server supports a persistent search, the LDAP listener recognizes any user and group changes made to any of the applicable LDAP accounts and forwards them to your instance within approximately 10 seconds. This allows the instance to have a nearly real-time copy of your users' account details without having to wait for the next scheduled refresh. The LDAP listener can only synchronize objects that map to the User [sys_users] and Group [sys_user_group] tables.

Note: If a user is added via the listener, but the user does not meet the requirements as defined by the OU filter, then the instance ignores the record on the LDAP server. If it meets the criteria, the user is added to the instance.

To enable a listener:

Procedure

1. Navigate to **All** > **System LDAP** > **LDAP Servers**.
2. Select the LDAP server to configure.
3. Select the **Listener** check box.
4. Click **Update**.

Note:

The system only imports user records that match the LDAP OU filter. Incoming user records that do not meet the filter requirements are flagged as invalid and ignored by the import. Administrators can enable verbose LDAP logging to determine if incoming records are not matching the LDAP OU filter.

5. **Optional:** Navigate to the System Properties [sys_properties] table and set LDAP listener system properties.

LDAP listener properties

Property	Description
glide.ldap.listener.use_background_transaction	When true, the LDAP listener is started as a background running the LDAP listener as a background transaction, t Listener Start/Stop Transaction can cancel the transactio duration is reached, 5 minutes by default. This behavior listener from waiting indefinitely. **Note:** This property applies only to LDAP connections that do not use a MID Server. Use *glide.ldap.listener.mid.use_backgro* to control the behavior of LDAP connections that go Server. ○ Type: true | false ○ Default value: false ○ Location: Add to the System Properties [sys_properties
glide.ldap.listener.mid.use_background_transaction	When true, the LDAP listener is started as a background running the LDAP listener as a background transaction, t

Property	Description
	Listener Start/Stop MID Transaction can cancel the trans maximum duration is reached, 5 minutes by default. This LDAP listener from waiting indefinitely. **Note:** This property applies only to LDAP connections that use a MID Server. Use *glide.ldap.listener.use_background_* control the behavior of LDAP connections that do no Server. ○ Type: true | false ○ Default value: false ○ Location: Add to the System Properties [sys_properties
glide.ldap.listener.mid.one_listener	When true, only a single ECC queue message is created the LDAP listener through a MID Server. When false, mult messages can be created, leading to the creation of mult stop the LDAP listener. ○ Type: true | false ○ Default value: true ○ Location: Add to the System Properties [sys_properties

Specify the LDAP attributes

Specify the attributes included in LDAP server queries by using the LDAP server **Attributes** field. This can enhance performance as well as security.

Before you begin

Role required: admin

About this task

By default, the system loads all of the attributes for each object that it has permission to read from your LDAP server. Using the **Attributes** field, you can specify and thereby limit the attributes the LDAP query returns. Using this approach for large LDAP imports can greatly improve the speed of those imports.

Procedure

Explicitly define attributes where possible.

If there is information that you do not want exposed to the instance, exclude the attribute. If you do not specify LDAP server attributes, user transactions may freeze for extended periods of time when new attributes are added to an LDAP server object because the system will be busy loading data from the new attributes.

Note: To use the manager lookup scripts described in Select or Create a Transform Map for LDAP Data, specify **manager** and **dn** (distinguished name) in the **Attributes** field. Neither attribute is required to be a part of a transform map.

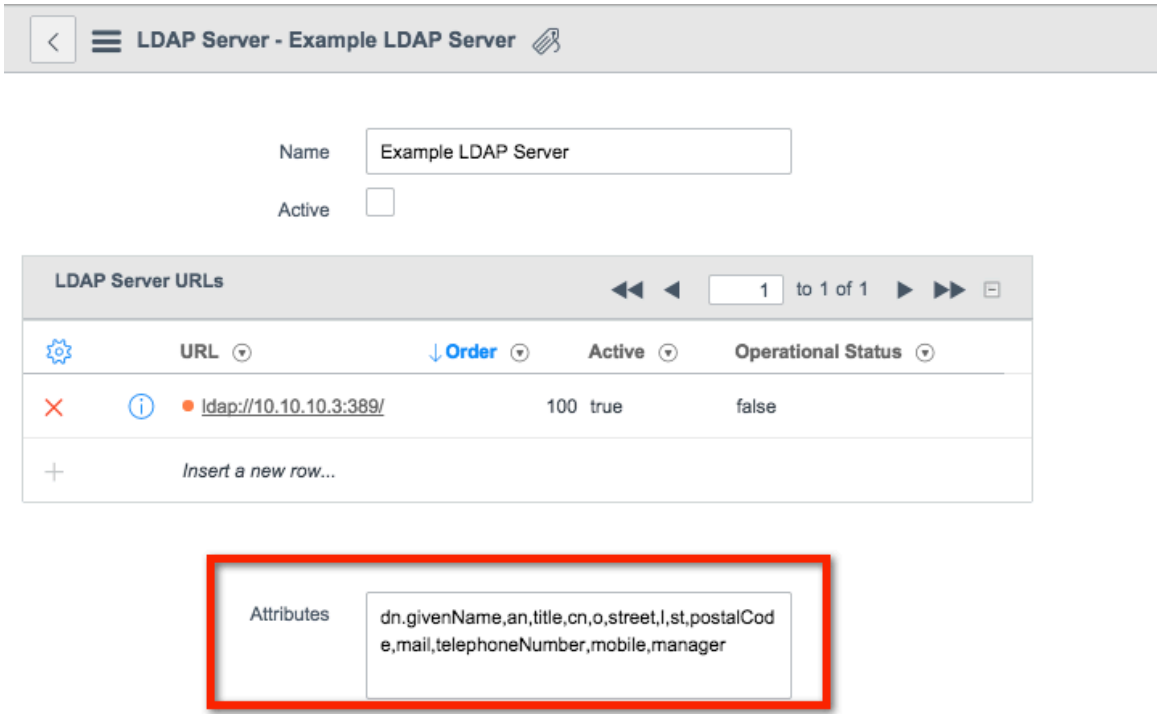

Test an LDAP connection

The instance tests the connection automatically every time a user opens the LDAP Server form. Alternatively, you can manually test the connection to the LDAP server from the LDAP server form.

Before you begin

Role required: admin

About this task

By default, error messages appear on the LDAP server form if there are any issues connecting to the LDAP server.

Note: Employees can also verify connectivity between the instance and the LDAP server. Contact Technical Support for assistance verifying LDAP connectivity.

To manually test a connection:

Procedure

1. Navigate to **All** > **System LDAP** > **LDAP Servers**.
2. Select the LDAP server to test.
3. Under **Related Links**, click **Test connection**.
4. Under **Related Links**, click **Browse** to verify that the appropriate LDAP directory structure is visible to the system.
5. **Optional:** If the connection was successful, click **Browse** to view the source LDAP directory structure that is visible to the instance.

   Note: The **Filter** and **RDN** fields on the left of the Browse window are ignored when you use the search field on the right.

Result

The instance changes the operational status of LDAP servers depending on the result of the connection test.

- If your instance establishes a connection to a server that has a **Operational Status** value of **down**, the **Operational Status** value is automatically changed to **up**. This functionality is supported for both automatic and manual connection tests.
- If a connection cannot be established to a server that has a **Operational Status** value of **up**, the **Operational Status** value is automatically changed to **down**. This functionality is supported for automatic connection tests only, not manual tests.

Define LDAP organizational units

An organizational unit (OU) definition specifies the LDAP source directories available to the integration.

Before you begin

Role required: admin.

About this task

OU definitions can contain locations, people, or user groups. Every LDAP server definition contains two sample OU definitions: one for importing groups into the system and the other for users.

Procedure

1. Navigate to **All** > **System LDAP** > **LDAP Servers**.
2. Select the LDAP server to configure.
3. In the **LDAP OU Definitions** related list, select either the **Groups** or **Users** sample OU definition.
4. Complete the LDAP OU Definition form (see table).
5. Click **Update**.
   The system automatically tests the connection to the LDAP server.
6. Under **Related Links**, click **Browse** to view the LDAP directory records that the OU definition returns.

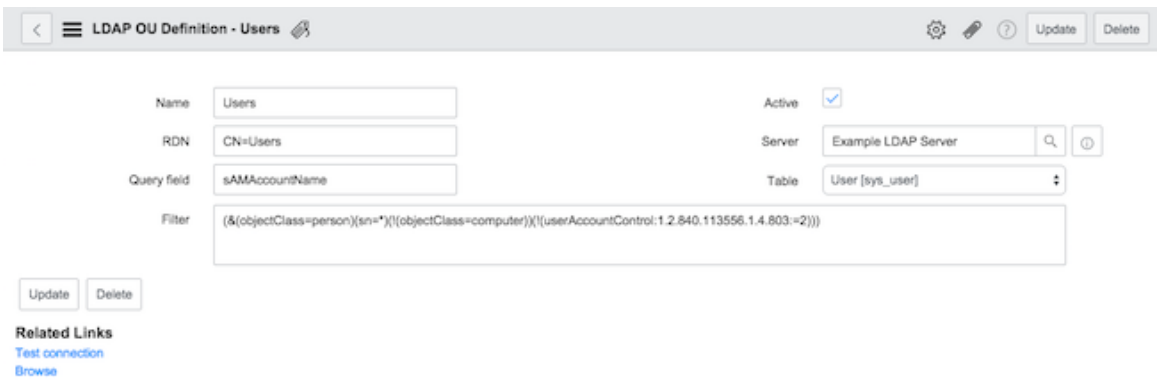

OU Definition form

Field	Description
Name	Specify the name the integration uses when referencing this OU. The name you enter here becomes an LDAP target in the data source record.
RDN	Specify the relative distinguished name of the subdirectory you want to search. This RDN is combined with the start-searching directory from the LDAP server definition to identify the subdirectory containing information for this organizational unit. For example, the sample OU definition uses the RDN value of `CN=Users` to search the

Field	Description
	LDAP directory `CN=Users,DC=service-now,DC=com` and any directory below this point. This field must match a subdirectory in your LDAP system.
Query field	Specify the name of the attribute within the LDAP server to query for records. The query field must be unique in both single and multiple domain instances. For best results, use email addresses or other credentials that uniquely identify the user in a multiple domain instance. Active Directory uses the **sAMAccountName** attribute. Other LDAP servers tend to use the **cn** attribute. **Note:** The **Query field** must map to the **User ID** field in the User [sys_user] table. For example, if an Active Directory user logs in as `joe.example`, there must be a user record with a **User ID** value of **joe.example** and an LDAP record with an **sAMAccountName** value of **joe.example**.
Active	Select this check box to activate the OU definition and to allow administrators to test importing data. However, the integration can only bring data into the system from active OU definitions.
Table	Specify the table that receives the mapped data from your LDAP server. For users, select **User (sys_user)**, and for groups, select **Group (sys_group)**.
Filter	Enter an LDAP filter string to select specific records to import from the OU. The more specific the LDAP filter query, the more efficient the query is. For example, the Users LDAP OU definition uses the following filter to select records that are classified as a person, have an **sn** attribute value, are not computers, and are not flagged as inactive: `(&(objectClass=person)(sn=*)(!(objectClass=computer)) (!(userAccountControl:1.2.840.113556.1.4.803:=2)))` You can find a description of LDAP filter syntax by searching the internet for `LDAP Filters RFC`.

Example: Example organizational unit definitions

Suppose you have an LDAP server with the following directory structure:

dc=my-domain,dc=com

- ou=Groups
  - cn=Development
  - cn=HR
  - cn=Sales
- ou=Users
  - ou=Development
  - ou=HR
  - ou=Sales

Further suppose that you want to exclude the HR group and HR users from the application. Do the following:

1. Create an LDAP server record with a starting search directory of dc=my-domain,dc=com.
2. Create an OU definition record for ou=Groups with a filter to exclude cn=HR.
3. Create an OU definition record for ou=Users with a filter to exclude ou=HR.

If you do not specify additional attributes or filters with an OU definition, the LDAP query returns the entire sub-tree from the starting directory and RDN.

In these examples, an OU definition with the RDN value of ou=Groups and no filter would have returned all groups. Likewise, an OU definition with the RDN value of ou=Users and no filter would have returned all users and child organizational units.

Create a data source for LDAP

Each LDAP organizational unit (OU) definition has its own related list of data sources.

Before you begin

Role required: admin

About this task

Note: Both the **LDAP Server** and **LDAP OU Definition** must be active for the test load action to function properly. When the test load is activated for the first time, the system samples up to 20 records to determine the length of the import set fields. If the sampled records do not contain values for the **User ID** field, the system sets the field length for all subsequent imports to the default length of 40. The import truncates any imported data that exceeds the import set table field length. Additionally, the **User ID** field is truncated to a maximum of 40 characters. Be aware that the 20 loaded records cannot be transformed and are for testing purposes only. If the test records contain values for the **User ID** field, the field length is set based on the field length of the longest user ID in the test records.

To create a new data source:

Procedure

1. Navigate to **All** > **System LDAP** > **LDAP Servers**.
2. Select the LDAP server to configure.
3. In the **LDAP OU Definitions** related list, select an item, such as **Groups** or **Users**.
4. In the **Data Sources** related list, click **New**.
5. Complete the Data Source form (see table).
6. Click **Submit**.
7. Under **Related Links**, click **Test Load 20 Records** to test whether the data source can bring LDAP data into the import table.

Data Sources form

Field	Description
Name	Specify the name the integration uses when referencing this data source.
Import set table name	Enter the name of the staging table where the system temporarily places the imported LDAP records and attributes. Review this table to view imported LDAP records. You can use the same import set table name for all LDAP data sources.

Field	Description
Type	Select **LDAP** to indicate the imported data is LDAP data. After you select the type **LDAP**, the form displays the **LDAP target** field.
LDAP target	Select the LDAP OU definition associated with this data source.

Auto provision LDAP users

You automatically provision users who are in the LDAP server but not yet in your instance.

Before you begin

Role required: admin

Procedure

Create the following properties in the System Properties [sys_properties] table:

LDAP properties

LDAP property	Description
glide.ldap.authentication	Enables LDAP authentication by using LDAP to authenticate users. Set this property to **true** (the default value).
glide.ldap.user.autoprovision	Enables LDAP the system to automatically create users in the User [sys_user] table when the user exists in LDAP but is not yet in the instance. Set this property to **true** (the default value).

Both of these properties must be set to **true** for auto-provisioning to work.

LDAP integration via MID Server

Administrators can integrate using an LDAP data source over a Management, Instrumentation, and Discovery (MID) Server.

The MID Server enables communication and movement of data between the ServiceNow AI Platform and external applications, data sources, and services. For details on installing a MID Server, see MID Server installation.

Using a MID Server to establish an LDAP connection prevents you from having to expose the LDAP server to external network traffic. It eliminates the need to establish a VPN tunnel between your LDAP server and datacenters. The MID Server user must have the user_admin role to be able to read LDAP server configuration records.

Note: The MID Server does not allow using the UI action `<instance>/sys_ui_action.do?sys_id=1b4f7ef30a0001060058e223c9a5744c` to refresh user and group records from LDAP.

A MID Server connection communicates over HTTP on port 80 by default. This communication channel does not require a certificate. The connection between the MID Server and the instance is over HTTPS (port 443). The instance connects to the LDAP server directly, using LDAP or LDAPS. This connection can either be over the internet or through a VPN tunnel.

Note: LDAP cannot communicate via the MID Server with password authentication.

For a secure communication over SSL, you must add an SSL certificate for the MID Server. Change the LDAP server URL from LDAP to LDAPS and change to port 636.

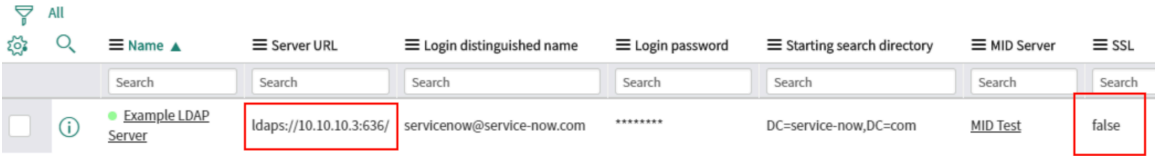

Note: If you create a new LDAP server, the MID Server SSL flag sets to false by default. You can ignore this behavior.

To set connection properties for a specific LDAP server, see Define an LDAP server.

Configure LDAP connection monitoring

Change or disable LDAP connection monitoring and notifications.

Before you begin
Role required: admin

About this task
The instance automatically sends an email to users configured in the LDAP Admins group when an LDAP server connection fails. This uses the email notification, which is launched by the **LDAP Connection Test**scheduled job. This email notification is enabled by default.

Note: The instance does not send the email notification unless there is at least one member in the LDAP Admins group. Make sure to populate this group with the users you want to receive the email.

By default, the scheduled job tests the connection every 15 minutes. To change this interval or disable monitoring:

Procedure

1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. Open **LDAP Connection Test**.
3. Do one of the following:
   - Change the interval in the **Repeat Interval** field.
   - Disable monitoring by clearing the **Active** check box.

Import binary data through a MID Server

As an administrator, you can import binary large object (BLOB) data with an LDAP integration through the MID Server.

Before you begin
Role required: admin

About this task

Procedure

1. Add the name of the LDAP column you want to import binary data from to the system property `glide.ldap.binary_attributes`.
2. Add a MID Server property with the Name `glide.ldap.binary_attributes` and the same value you set for the system property.

Troubleshooting LDAP integration via MID Server

You may encounter issues in the following areas while integrating LDAP via MID Server.

You can troubleshoot these issues by viewing the outputs found in the External Communication Channel (ECC) Queue (**Discovery** > **Output and Artifacts** > **ECC Queue**).

Test Connection Issues

When defining OUs within the server, there is a **Test connection** related list that is used to verify the LDAP connection. When you click this link, the ECC Queue should show a single output message with a topic name of **LDAPConnectionTesterProbe**. After the test has completed on the MID Server, the ECC Queue should show an input message with the same topic name. If the **Name** column for the input message shows **true**, the test was successful. Drill down into the record to view the payload and ensure it does not contain error messages.

Test Connection

Created	Agent	Topic	Name	Source	Queue	State	Processed
2013-07-29 13:24:17	mid.server.local_mid	LDAPConnectionTesterProbe	true	04a952038f21010036bf21ca47e79a30	input	processed	2013-07-29 13:24:19
2013-07-29 13:24:04	mid.server.local_mid	LDAPConnectionTesterProbe		04a952038f21010036bf21ca47e79a30	output	processed	2013-07-29 13:24:17

Browse Issues

When defining OUs within the server, there is a **Browse** related list that is used to view the LDAP directory records that the OU definition returns. When you click this link, the ECC Queue should show a single output message with a topic name of **LDAPBrowseProbe**. After data has been returned from the MID Server, the ECC Queue should show an input message with the same topic name. If the **Name** column for the input message shows **true**, the test was successful. Drill down into the record to view the payload and ensure it does not contain error messages.

Load Import Issues

When uploading data (for example, using the Test Load 20 Records feature), the ECC Queue should show a single output message with a topic name of **LDAPProbe**.

After data has been returned from the MID Server, the ECC Queue should show another input message called **LDAPProbeCompleted**. The **Name** column for this input message shows the total number of records returned.

An additional input messages, also named **LDAPProbe**, is displayed. The **Name** column for this input message displays the highest record number in the batch. If the total number of records returned is 258 and the batch size is 200 (the default), two LDAPProbe (200, 258) incoming messages will be received, and one LDAPProbeCompleted (258) incoming message will be received.

Drill down into the record to view the payload and ensure it does not contain error messages.

Import Load

Created	Agent	Topic	Name	Source	Queue	State	Processed
2013-07-29 13:09:48	mid.server.local_mid	LDAPProbeCompleted	11	ed0a0d7a8f32010036bf21ca47e79a56	input	processed	2013-07-29 13:09:51
2013-07-29 13:09:48	LDAPProbeResult	LDAPProbe	11	ed0a0d7a8f32010036bf21ca47e79a56	input	processed	2013-07-29 13:09:51
2013-07-29 13:09:36	mid.server.local_mid	LDAPProbe		ed0a0d7a8f32010036bf21ca47e79a56	output	processed	2013-07-29 13:09:46

Also keep an eye out for an output message called **LDAPProbeError**.

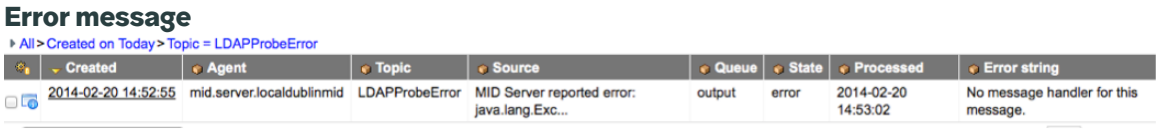

Click the link in the **Name** column to view the details of the error.

LDAP paging

LDAP paging does not work if the paging size on the LDAP server is less than 1000. Set the MID Server property *`glide.ldap.max_results`* to a value less than or equal to the LDAP server paging size.

LDAP fails to import binary data

To import binary data via LDAP, such as a user photo, you must include the binary attribute in the MID Server property *`glide.ldap.binary_attributes`*. For the user photo example, the attribute may be `jpegphoto`.

Import and map data

LDAP import maps match fields in your LDAP database to fields in your instance.

Note: LDAP mapping has a performance effect, so the recommended approach is to schedule it during off-peak hours, or process a few records at a time to maintain system availability.

Define a transform map that only imports the needed or required attributes. Depending on the version of the instance you are using, the method for specifying LDAP mapping relationships varies.

The easiest way to know whether or not you are running a version which uses the System LDAP application for LDAP integration is to find the application from the application navigator.

The **Run Business Rules** option is applied only for the target table. Only transform maps associated to the target table run the business rules associated with different tables. If you are updating a user group and have business rules running on a user group table, the group must have roles defined.

LDAP import mapping options

System LDAP application?	Map
Yes	Use a transform map to specify your mapping.
No	Use a LDAP legacy import map to specify your mapping, or the default LDAP transform that is included in baseline instances. Remember to adjust the **Coalesce** field to match against the correct fields.

Scheduled imports

A scheduled import allows administrators to import LDAP data on a regular schedule. By default, the LDAP integration includes two sample scheduled imports:

- Example LDAP User Import
- Example LDAP Group Import

Neither example is active by default. Change these scheduled imports to meet your company's business needs.

LDAP transform maps

The transform map moves data from the import set table to the target table (User or Group).

The LDAP integration uses standard import sets and transform maps. You can also create custom LDAP transform maps.

Important: Whether you select or create custom LDAP transform maps, there should be one active transform map for a set of source and target tables. Enabling multiple transform maps for the same source and target tables can produce duplicate entries in the target table unless you coalesce against the matching fields.

Default LDAP transform maps

By default, the system provides two transform maps for LDAP data.

Default LDAP transform maps

Transform Map	Source Table	Target Table	Description
LDAP User Import	[ldap_import]	[sys_user]	Default transform map for creating user records from LDAP credentials as part of LDAP on-demand login. Contains mappings for an Active Directory LDAP server.
LDAP Group Import	[ldap_group_import]	[sys_user_group]	Default transform map for creating group records from LDAP OUs. Contains mappings for an Active Directory LDAP server.

Note: By default, the system does not have a transform map for LDAP department records.

Requirements for custom LDAP transform maps

If you choose to create a custom transform map, the transform map must meet the following mapping requirements.

Requirements for custom LDAP transform maps

Source Table	Source Field	Target Table	Target Field	Coalesce	Description
`ldap_import`	`u_source`	`sys_user`	source	false	The **u_source** field identifies the LDAP DN of the imported user or group. The system uses this field to determine that a user requires LDAP authentication, to find a user's manager, and to put users into groups.

Requirements for custom LDAP transform maps (continued)

Source Table	Source Field	Target Table	Target Field	Coalesce	Description
`ldap_import`	Select one of the following fields: • `u_samaccountname` • `u_dn` • `u_cn`	`sys_user`	`user_name`	true	If LDAP integrates to Active Directory, select **u_samaccountname** as the source field. If other LDAP directories are used, select **u_dn** or **u_cn** as the source field.

Differences between LDAP transform maps and legacy import maps

When specifying LDAP mapping relationships using transform maps, there is a major difference in how reference fields are set for manager and department.

When using a transform map, it is necessary to use a transform script to create references. This is because the value associated with an LDAP attribute like "manager" is the distinguished name (DN) of the manager.

Without some extra logic in place, the result is the creation of a user record with a manager name that is the distinguished name of that user in LDAP. The integration includes a transform script to facilitate the creation of these references. The default transform map "LDAP User Import" includes transform scripts for these references.

Existing mapping relationships

When updating legacy import maps to transform maps, you can retain the LDAP mapping relationships that existed prior to the addition of the System LDAP application. The LDAP server has a **Map** field that is a reference to the legacy import map.

Note: By default this field is hidden, so you have to configure the form to display it.

If you want to transition to using a transform map, clear the reference to the legacy import map.

LDAP import map settings

Verify and use attributes to limit the fields the integration imports from the LDAP source. Additionally, it is important to map the user_name field to the LDAP attribute that contains the user's login ID. For Active Directory this is usually the sAMAccountName attribute. If you would like to import and coalesce on a binary attribute (such as objectSID or objectGUID), you have to create a custom transform script.

Note: Any value mapped to the user_name field must be unique.

If you do not specify a transform map (such as LDAP User Import), the integration uses the following default mappings:

LDAP import default mapping

User field or variable	LDAP attribute
user_name	sAMAccountName
email	mail
phone	telephoneNumber
home_phone	homePhone
mobile_phone	mobile
first_name	givenName
last_name	sn
title	title
department	department
manager	manager
middle_name	initials
u_memberof	groups
u_member	members
u_manager	manager

LDAP data transformation

If an LDAP attribute contains simple data, the transform map links an imported LDAP attribute to an appropriate field in the target table (User or Group). For example, sample data in the sAMAccountName attribute maps to the **User ID** field in the User table.

If the imported LDAP data maps to a reference field, the instance searches for an existing matching record. If no matching record exists, the instance creates a new record for the reference field unless the field mapping specifies otherwise.

For example, suppose the LDAP attribute l maps to the **Location** reference field in the User table. Whenever the import brings in an attribute value that does not match an existing location record value, the transform map creates a new location record. The new location record has the same value as the imported attribute, and the imported user record now has a link to the new location record.

However, there are times when LDAP attribute returns a distinguished name (DN), which is essentially a reference to another record within the LDAP directory. For example, the manager attribute typically contains the distinguished name for the manager of the current LDAP directory entry. An imported DN typically uses a long text string such as: cn=Beth Anglin,ou=Users,dc=my-domain,dc=com.

⚠ **Warning:** Make sure your target fields are long enough to contain a DN. Many text fields use the default length of 40, which may not be long enough for some DN values. The ServiceNow system truncates any value that exceeds the field length.

Administrators do not typically want the system to create new users from the DN value because the new user has no association with an existing user. Instead, administrators want the import to locate the manager's existing user record and associate it with the newly imported user. The `LDAPUtils` script include contains the `setManager` and `processManagers` functions that

can parse a DN and search for an existing user. For best results, use these functions to create a custom transform map.

For example, the `LDAP User Import` transform map script calls the `setManager` function:

```
//
// The manager coming in from LDAP is the DN value for the
 manager.
// The line of code below will locate the manager that matches
 the
// DN value and set it into the target record.  If you are not
// interested in getting the manager from LDAP then remove or
// comment out the line below
ldapUtils. setManager (source , target ) ;
```

In some cases, the integration imports a user's record before importing the associated manager's user record. To handle such cases, you may want to call the `processManagers` function after the transform completes. For example, the **LDAP User Import** transform map uses an `onComplete` transform script to call the `processManagers` function.

```
// It is possible that the manager for a user did not exist
 in the database when // the user was processed and therefore
 we could not locate and set the manager field. // The
 processManagers call below will find all those records for
 which a manager could  // not be found and attempt to locate
 the manager again. This happens at the end of the  // import
 and therefore all users should have been created and we should
 be able to  // locate the manager at this point
ldapUtils. processManagers ( ) ;
```

Remove or comment out the `setManager` and `processManagers` function calls if your LDAP integration does not use the manager attribute.

LDAP scripting

Create custom transform maps, scripts, and business rules to specify requirements when importing data.

Custom transform maps should include `onStart` and `onAfter` transform scripts.

The `onStart` script should call the `LDAPUtils` script include and start logging. For example, the **LDAP User Import** transform map has an `onStart` script that uses this code:

```
gs.include ( "LDAPUtils" ) ; var ldapUtils  = new LDAPUtils
 ( ) ;
ldapUtils. setLog (log ) ;
```

The `onAfter` script should call the `addMembers` function. For example:

```
ldapUtils.addMembers (source , target ) ;
```

Set disabled Active Directory users to inactive

Use the following script to automatically deactivate users when the associated AD user is disabled.

Before you begin

Role required: admin

About this task

You can identify disabled Active Directory users by checking the value of the `userAccountControl` attribute. This rule executes whenever the `userAccountControl` value changes and deactivates user accounts if the **User Account Control** signifies a disabled AD account.

Use the following script to automatically deactivate users when the associated AD user is disabled.

Procedure

1. Configure the User form and create a new integer field called **User Account Control**.
2. Add mapping for userAccountControl (external) to the new field.
3. Create a new business rule with the following properties:

Disable AD Users business rule

Business rule field	Value
Name	Disable AD Users
Table	User [sys_user]
When	Before
Condition	current.u_user_account_control.changes()

The Script field should contain the following:

```
var disabledFlag = 2;
//perform a bitwise comparison on userAccountControl to see if
 the 2 bit flag is enabled
if (current.u_user_account_control & disabledFlag) {
  gs.log('Disabling user: ' + current.user_name +
 'userAccountControl=' + current.u_user_account_control);
  current.active='false';
  current.locked_out='true';
}
```

Assign LDAP field values

You can use a script to assign a value to any field for which there is a field mapping.

For example, to assign a value to the sys_user.company field, create a field map for the company field and add a transform script of:

```
company = "Don's Sporting Goods";
```

Exclude particular LDAP users

If you cannot completely filter the LDAP user list using LDAP filter properties, you can exclude users with a map script.

After you have run the logic to identify a user that should not be imported, set the user_name field to an empty string and this user will not be imported.

```
user_name='';
```

One way to identify users to filter out is to look for a string in the `distinguishedName` attribute. For example, this script excludes accounts that are not in a Users OU. You might use this script if you have too many Users OU to include in the target OU LDAP Option.

```
//vdn is a variable mapped to distinguishedName
gs.include("LDAPUtils");
var vdn = source.getElement(this.distinguishedName);
if (vdn.indexOf('OU=Users')<0) {
  user_name='';
  gs.log('LDAP Import Skipping User: ' + vdn);
}
```

A more complex method of filtering is to use regular expressions.

```
//vcn is a variable mapped to cn
//vdn is a variable mapped to distinguishedName
//c is the regular expression string
gs.include("LDAPUtils");
var vdn = source.getElement(this.distinguishedName);
var vcn = source.getElement(this.cn);
var c = /^[a-z][a-z][a-z][0-9][0-9][0-9]$/;
var nvcn = vcn.toLowerCase();
//test to see if the cn is in the form of 3 letters followed by
 3 numbers, only import these
if (c.test(nvcn)) {
 user_name = nvcn;
} else {
 gs.log("LDAP import rejected username: " + vcn + " for DN: " +
 vdn);
 user_name = "";
}
```

Set choice action for reference field imports

The LDAP transform map determines how fields in the Import Set table map to fields in existing tables such as Incident or User.

Before you begin

Role required: admin

About this task

If the LDAP transform map updates a field in the import set table, the integration automatically creates a new record whenever there is a new record in the LDAP data. If the LDAP transform map updates a reference field storing data from another table, the administrator can choose to create, ignore, or reject new LDAP records.

For example, if the integration receives a new department record that does not match any existing department, you may want to update all of the other LDAP record fields without creating a new department record in the instance. The transform map allows you to set the record creation options for each reference field.

Procedure

1. Navigate to **All** > **System LDAP** > **Transform Maps**.
2. In the Field Maps related list, select one of the following actions from the **Choice action** field:

- **create** – creates a new reference field record if a matching record does not exist.
- **ignore** – ignores new records in the reference field and completes processing of all other fields in the transform map.
- **reject** – stops the transform for the entire record.

Note: The field map only displays the **Choice action** field for reference fields.

Verify LDAP mapping

After creating an LDAP transform map, refresh the LDAP data to verify the transform map works as expected.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System LDAP** > **Scheduled Loads**.
2. Click your LDAP import job.
3. Click **Execute Now**.

LDAP integration troubleshooting

If you are integrating your LDAP server and have questions, these items may help you troubleshoot the issue.

Preliminary checks

- If the LDAP is unavailable, users cannot log in to the instance. A good practice is to have local accounts for administrators so that if the LDAP is down, administrators can still access the instance.
- Check the service account to ensure that it is not expired or locked out.
- Check the format of the username. Instead of using just the username, try using the domain with the username, or username@domain.
- Verify that you have changed the `system_id` entry on the `ldap_server_config` record. If you modify the `system_id` unintentionally with an update set, `system_id` points to the wrong node for the target instance and does not work.

Error codes

The LDAP log file lists industry standard error codes for both LDAP and Active Directory (AD). The LDAP log file is contained in the wrapper file. The LDAP error codes are two-digit numbers, while the Active Directory error codes are three-digit numbers. For a list of the most-common error codes, see LDAP Error Codes.

Multiple domain integration

You can integrate multiple domains within the same forest or in completely non-trusted domains. It is recommended that you create a separate LDAP server record for each domain. Each LDAP server record must point to a domain controller for that given domain. This means you will have to allow connections to each of the domain controllers. Multiple AD forests through LDAP with one LDAP account is not supported.

When you expand to more than one domain, it is critical that you identify unique LDAP attributes for the application usernames and import coalesce values. A common unique coalesce attribute for Active Directory is `objectSid`. Unique usernames will vary based on your LDAP data design. Common unique attributes are `email` or `userPrincipalName`.

Incoming records

See LDAP transform maps to set how the integration processes incoming LDAP records that are missing matching values in reference fields.

Common authentication errors

- User Cannot Log In (Invalid DN)
- Invalid CN
- Invalid Connection

Automatic LDAP connection tests

You can manually test connections to LDAP servers or allow ServiceNow to automatically test the connections.

The system tests the connection automatically:

- Every time a user opens the LDAP Server form.
- Through the LDAP Connection Test scheduled job, which runs every 15 minutes by default.

  You can change how often this scheduled job runs. If this scheduled job is not able to establish a connection, a new one-time schedule job retries the connection test after either five minutes, or half the **Repeat Interval** value in the scheduled job, whichever occurs first.

Error messages appear on the form if there are any issues connecting to the LDAP server. Also supported are test connections for servers behind a MID server.

View the LDAP monitor

You can view current information about LDAP servers and listeners using LDAP monitor.

Before you begin
Role required: admin

About this task
The available states are:

- Active
- Inactive
- Error
- Active (Shutting down...)
- Error (Shutting down...)

In addition to its current state, the monitor also shows:

- The last message detected by the listener, such as waiting for LDAP changes, error connecting, and so forth.
- The last LDAP user change, such as new user, updated user, and so forth.
- The last error that occurred.

To view LDAP monitor:

Procedure

Navigate to **All** > **LDAP** > **System LDAP** > **LDAP Monitor**.

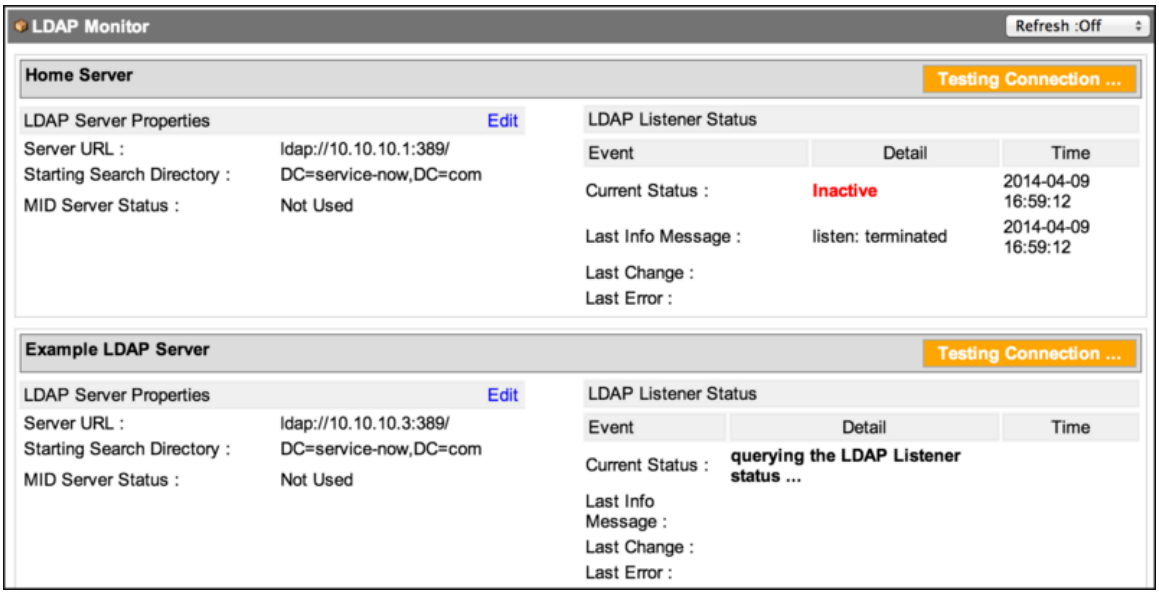

See the table for descriptions of the properties and fields in the screen.

LDAP monitor

Field	Description
Refresh	You can configure the refresh rate by clicking the **Refresh** field in the LDAP Server Monitor header bar, and selecting the number of seconds between each data refresh. You can also select **None** to suppress refreshing.
Connection Status	The server connection indicator is located on the right side, above the LDAP Listener Status fields. When the server is connected, the box is green and shows *Connected*. When the server is not connected, the box is red and shows *Not Connected*. When the server connection is being tested, the box is yellow and shows *Testing Connection*.
LDAP Server Properties	
Edit	As you monitor LDAP servers, you can make changes to the properties by clicking **Edit** in the LDAP Server Monitor screen.
Server URL	The combination of the server name and server port where the LDAP Server is listening. Frequently, the port is set to one of the following: • 389: the default port for connecting to LDAP in clear text • 636: the standard port for connecting to LDAP via an SSL connection Example value: ldap://10.10.10.3:389/ Your LDAP Server may have more than one URL address. This does NOT establish multiple directory structures from which you can import data, which is done by creating another LDAP Server entry, but does provide for redundancy when you have multiple LDAP Servers to avoid a single point of failure. The LDAP URL

Field	Description
	addresses are separated with a space character, and the system automatically tries each server address in turn until a valid connection can be made.
Starting search directory	The starting directory or RDN (Relative Distinguished Name) where the system begins searching for users or groups. Example value: DC=service-now,DC=com No data ABOVE this point is available for import. The instance has visibility into the specified directory and directories BELOW it in the LDAP hierarchy.
MID Server Status	The current connection status of the MID Server.
LDAP Listener Status	
Current Status	This indicates whether the listener is active.
Last Info Message	This shows the last message the LDAP server received relating to user and group changes, and the time the message was received.
Last Change	This shows the last change made to the LDAP server, and the time it was made.
Last Error	This shows the last error that occurred on to the LDAP server, and the time it occurred.

LDAP error codes

The LDAP Log file lists industry standard error codes for both LDAP and Active Directory (AD).

Standard error codes

Standard LDAP errors

Error / data code	Text	Description
0	LDAP_SUCCESS	Indicates the requested client operation completed successfully.
2	LDAP_PROTOCOL_ERROR	Indicates that the server has received an invalid or malformed request from the client.
3	LDAP_TIMELIMIT_EXCEEDED	Indicates that the operation's time limit specified by either the client or the server has been exceeded. On search operations, incomplete results are returned.
4	LDAP_SIZELIMIT_EXCEEDED	Indicates that in a search operation, the size limit specified by the client or the server has been exceeded. Incomplete results are returned.

Standard LDAP errors (continued)

Error / data code	Text	Description
5	LDAP_COMPARE_FALSE	Does not indicate an error condition. Indicates that the results of a compare operation are false.
6	LDAP_COMPARE_TRUE	Does not indicate an error condition. Indicates that the results of a compare operation are true.
7	LDAP_AUTH_METHOD_NOT_SUPPORTED	Indicates that during a bind operation the client requested an authentication method not supported by the LDAP server.
8	LDAP_STRONG_AUTH_REQUIRED	Indicates one of the following: In bind requests, the LDAP server accepts only strong authentication. In a client request, the client requested an operation such as delete that requires strong authentication. In an unsolicited notice of disconnection, the LDAP server discovers the security protecting the communication between the client and server has unexpectedly failed or been compromised.
9		Reserved.
10	LDAP_REFERRAL	Does not indicate an error condition. In LDAPv3, indicates that the server does not hold the target entry of the request, but that the servers in the referral field may.
11	LDAP_ADMINLIMIT_EXCEEDED	Indicates that an LDAP server limit set by an administrative authority has been exceeded.
12	LDAP_UNAVAILABLE_CRITICAL_EXTENSION	Indicates that the LDAP server was unable to satisfy a request because one or more critical extensions were not available. Either the server does not support the control or the control is not appropriate for the operation type.
13	LDAP_CONFIDENTIALITY_REQUIRED	Indicates that the session is not protected by a protocol such as Transport Layer Security (TLS), which provides session confidentiality.
14	LDAP_SASL_BIND_IN_PROGRESS	Does not indicate an error condition, but indicates that the server is ready for the next step in the process. The client must send the server the same SASL mechanism to continue the process.

Standard LDAP errors (continued)

Error / data code	Text	Description
15		Not used.
16	LDAP_NO_SUCH_ATTRIBUTE	Indicates that the attribute specified in the modify or compare operation does not exist in the entry.
17	LDAP_UNDEFINED_TYPE	Indicates that the attribute specified in the modify or add operation does not exist in the LDAP server's schema.
18	LDAP_INAPPROPRIATE_MATCHING	Indicates that the matching rule specified in the search filter does not match a rule defined for the attribute's syntax.
19	LDAP_CONSTRAINT_VIOLATION	Indicates that the attribute value specified in a modify, add, or modify DN operation violates constraints placed on the attribute. The constraint can be one of size or content (string only, no binary).
20	LDAP_TYPE_OR_VALUE_EXISTS	Indicates that the attribute value specified in a modify or add operation already exists as a value for that attribute.
21	LDAP_INVALID_SYNTAX	Indicates that the attribute value specified in an add, compare, or modify operation is an unrecognized or invalid syntax for the attribute.
22-31		Not used.
32	LDAP_NO_SUCH_OBJECT	Indicates the target object cannot be found. This code is not returned on following operations: Search operations that find the search base but cannot find any entries that match the search filter. Bind operations.
33	LDAP_ALIAS_PROBLEM	Indicates that an error occurred when an alias was dereferenced.
34	LDAP_INVALID_DN_SYNTAX	Indicates that the syntax of the DN is incorrect. (If the DN syntax is correct, but the LDAP server's structure rules do not permit the operation, the server returns code 53: LDAP_UNWILLING_TO_PERFORM.)
35	LDAP_IS_LEAF	Indicates that the specified operation cannot be performed on a leaf entry. (This code is not currently in the LDAP specifications, but is reserved for this constant.)

Standard LDAP errors (continued)

Error / data code	Text	Description
36	LDAP_ALIAS_DEREF_PROBLEM	Indicates that during a search operation, either the client does not have access rights to read the aliased object's name or dereferencing is not allowed.
37-47		Not used.
48	LDAP_INAPPROPRIATE_AUTH	Indicates that during a bind operation, the client is attempting to use an authentication method that the client cannot use correctly. For example, either of the following cause this error: The client returns simple credentials when strong credentials are required...OR...The client returns a DN and a password for a simple bind when the entry does not have a password defined.
49	LDAP_INVALID_CREDENTIALS	Indicates that during a bind operation one of the following occurred: The client passed either an incorrect DN or password, or the password is incorrect because it has expired, intruder detection has locked the account, or another similar reason. See the data code for more information.
49 / 52e	AD_INVALID CREDENTIALS	Indicates an Active Directory (AD) AcceptSecurityContext error, which is returned when the username is valid but the combination of password and user credential is invalid. This is the AD equivalent of LDAP error code 49.
49 / 525	USER NOT FOUND	Indicates an Active Directory (AD) AcceptSecurityContext data error that is returned when the username is invalid.
49 / 530	NOT_PERMITTED_TO_LOGON_AT_THIS_TIME	Indicates an Active Directory (AD) AcceptSecurityContext data error that is logon failure caused because the user is not permitted to log on at this time. Returns only when presented with a valid username and valid password credential.
49 / 531	RESTRICTED_TO_SPECIFIC_MACHINES	Indicates an Active Directory (AD) AcceptSecurityContext data error that is logon failure caused because the user is not permitted to log on from this computer. Returns only when

Standard LDAP errors (continued)

Error / data code	Text	Description
		presented with a valid username and valid password credential.
49 / 532	PASSWORD_EXPIRED	Indicates an Active Directory (AD) AcceptSecurityContext data error that is a logon failure. The specified account password has expired. Returns only when presented with valid username and password credential.
49 / 533	ACCOUNT_DISABLED	Indicates an Active Directory (AD) AcceptSecurityContext data error that is a logon failure. The account is currently disabled. Returns only when presented with valid username and password credential.
49 / 568	ERROR_TOO_MANY_CONTEXT_IDS	Indicates that during a log-on attempt, the user's security context accumulated too many security IDs. This is an issue with the specific LDAP user object/ account which should be investigated by the LDAP administrator.
49 / 701	ACCOUNT_EXPIRED	Indicates an Active Directory (AD) AcceptSecurityContext data error that is a logon failure. The user's account has expired. Returns only when presented with valid username and password credential.
49 / 773	USER MUST RESET PASSWORD	Indicates an Active Directory (AD) AcceptSecurityContext data error. The user's password must be changed before logging on the first time. Returns only when presented with valid user-name and password credential.
50	LDAP_INSUFFICIENT_ACCESS	Indicates that the caller does not have sufficient rights to perform the requested operation.
51	LDAP_BUSY	Indicates that the LDAP server is too busy to process the client request at this time but if the client waits and resubmits the request, the server may be able to process it then.
52	LDAP_UNAVAILABLE	Indicates that the LDAP server cannot process the client's bind request, usually because it is shutting down.
52e	AD_INVALID CREDENTIALS	Indicates an Active Directory (AD) AcceptSecurityContext error, which is returned when the username is valid

Standard LDAP errors (continued)

Error / data code	Text	Description
		but the combination of password and user credential is invalid. This is the AD equivalent of LDAP error code 49: LDAP_INVALID_CREDENTIALS.
53	LDAP_UNWILLING_TO_PERFORM	Indicates that the LDAP server cannot process the request because of server-defined restrictions. This error is returned for the following reasons: The add entry request violates the server's structure rules...OR...The modify attribute request specifies attributes that users cannot modify...OR...Password restrictions prevent the action...OR...Connection restrictions prevent the action.
54	LDAP_LOOP_DETECT	Indicates that the client discovered an alias or referral loop, and is thus unable to complete this request.
55-63		Not used.
64	LDAP_NAMING_VIOLATION	Indicates that the add or modify DN operation violates the schema's structure rules. For example, The request places the entry subordinate to an alias. The request places the entry subordinate to a container that is forbidden by the containment rules. The RDN for the entry uses a forbidden attribute type.
65	LDAP_OBJECT_CLASS_VIOLATION	Indicates that the add, modify, or modify DN operation violates the object class rules for the entry. For example, the following types of request return this error: The add or modify operation tries to add an entry without a value for a required attribute. The add or modify operation tries to add an entry with a value for an attribute which the class definition does not contain. The modify operation tries to remove a required attribute without removing the auxiliary class that defines the attribute as required.
66	LDAP_NOT_ALLOWED_ON_NONLEAF	Indicates that the requested operation is permitted only on leaf entries. For example, the following types of requests return this error: The client requests a delete operation on a parent entry. The

Standard LDAP errors (continued)

Error / data code	Text	Description
		client request a modify DN operation on a parent entry.
67	LDAP_NOT_ALLOWED_ON_RDN	Indicates that the modify operation attempted to remove an attribute value that forms the entry's relative distinguished name.
68	LDAP_ALREADY_EXISTS	Indicates that the add operation attempted to add an entry that already exists, or that the modify operation attempted to rename an entry to the name of an entry that already exists.
69	LDAP_NO_OBJECT_CLASS_MODS	Indicates that the modify operation attempted to modify the structure rules of an object class.
70	LDAP_RESULTS_TOO_LARGE	Reserved for CLDAP.
71	LDAP_AFFECTS_MULTIPLE_DSAS	Indicates that the modify DN operation moves the entry from one LDAP server to another and requires more than one LDAP server.
72-79		Not used.
80	LDAP_OTHER	Indicates an unknown error condition. This is the default value for NDS error codes which do not map to other LDAP error codes.
775	USER_ACCOUNT_LOCKED	Indicates users are unable to log in because the user account is locked.

Customized error codes

Customized LDAP error codes

Error / data code	Text
10000	LDAP_ERROR_GENEREL
10001	LDAP_ERROR_MAL_FORMED_URL
10002	LDAP_ERROR_UNAUTHENTICATED_BIND
10300	LDAP_ERROR_COMMUNICATION_EXCEPTION
10301	LDAP_ERROR_SOCKET_TIMEOUT
10302	LDAP_ERROR_CONNECTION_REFUSED
10303	LDAP_ERROR_CONNECTION_RESET
10304	LDAP_ERROR_NO_ROUTE

Customized LDAP error codes (continued)

Error / data code	Text
10305	LDAP_ERROR_UNKNOW_HOST
10400	LDAP_ERROR_SSL_EXCEPTION
10401	LDAP_ERROR_SSL_EMPTY_CERT_STORE
10402	LDAP_ERROR_SSL_CERT_NOT_FOUND
10403	LDAP_ERROR_SSL_CERT_EXPIRED
10500	LDAP_ERROR_INVALID_SEARCH_FILTER_EXCEPTION

Send a one-time password when the LDAP server is down

An LDAP property is available to send a one-time password to a user if the user is unable to log in because the LDAP server is down. You can also configure another property to control how long the password is valid.

Before you begin

Role required: admin

To receive a one-time password, the user must have notifications enabled on their user profile. The notification is an email message only. SMS messages are not supported.

About this task

Both properties are enabled by default. The default value for property that controls password validity is 10 minutes.

Procedure

1. Open the list of system properties by entering `sys_properties.list` in the filter of the application navigator.
2. Find the *`glide.ldap.onetime.password.enabled`* property.
3. Set the property to `true`.
4. To change the password validity time for a user, set the following property to an integer number of minutes: *`glide.authenticate.onetime.password.validity`*.

LDAP record synchronization

Administrators can synchronize inactive, disabled, or deleted LDAP records with their LDAP records.

LDAP record synchronization is the process of detecting inactive records on the LDAP server and updating the corresponding LDAP records. Detecting inactive LDAP records involves defining consistent data indicators for each user object, importing LDAP data, and evaluating the data indicators.

A data indicator can be:

- A date field
- Membership in a specific OU (identify by parsing the *dn* attribute), using the *`useraccountcontrol`* attribute
- A combination of these indicators

Imported data comes into the instance through import set tables where the data can be evaluated and processed.

The import process can use LDAP refresh filterson multiple import jobs to divide different types of user records and segregarte records for separate processing.

LDAP refresh filters

Filters on the LDAP refresh process can be used to specify processing that ignores inserts of disabled users.

You can loosen the LDAP OU filter to bring all of the data in to your import set table (including inactive users) and then specify processing that ignores inserts of disabled users. The sample 'Users' OU definition that the instance provides in its out-of-box LDAP sample contains a filter.

This filter is important because it defines which user records are brought into the import set table to be evaluated. While achieving a smaller data load, a limitation of this filter is that it filters out inactive users, so the inactive user records are not imported into the import set temporary tables. Since there is not visibility of the inactive user records, there is no ability to evaluate the record indicators.

LDAP OU filter

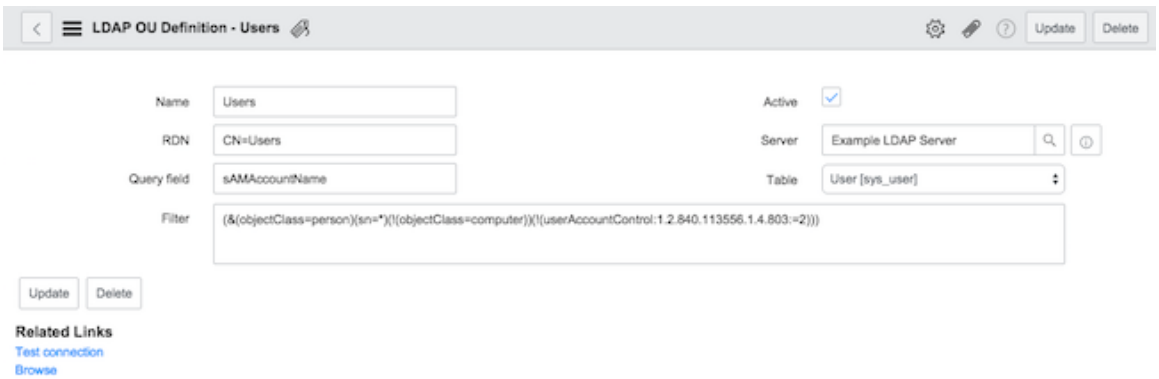

To use filtering within the main LDAP refresh process, change the filter to bring in all of the user records. The result is that all the records will be loaded into the import set temporary table where they can be evaluated and transformed.

Note: There is a precaution here: because the filtering brings in all the records, you may end up with a vast amount of older inactive LDAP accounts that should not be inserted into the instance. A user record should never be created for a disabled user.

LDAP extraction

Implement an LDAP extraction process to detect inactive users.

To detect inactive users using LDAP extraction, create a separate LDAP data source scoped specifically to inactive user accounts. For example, target a inactive users organizational unit (OU) or apply a query filter that matches inactive account flags. In the Table Transform Map for that data source, add a transform script that sets `target.active = false` for each record. Because the data source returns only inactive users, the script deactivates only those accounts in ServiceNow.

Benefits

Benefits to this method include:

- Simple scripting
- Existing user records aren't involved in processing
- Inactive users aren't loaded into a temporary import table
- No performance impact

Drawbacks

Drawbacks to this method include:

- An additional process is created
- The extract set must be placed in a location where your data source can access it

Alternative method

LDAP refresh filters use multiple import jobs to divide different types of user records, segregating records for separate processing.

Inactive LDAP user accounts

Detect that an existing, current, user account is inactive or has been disabled or deleted from an Active Directory (AD) LDAP.

A common LDAP integration issue is how to detect disabled or deleted users in an Active Directory (AD) and then deactivate them in the instance. In an Active Directory LDAP, a filter is usually set to exclude inactive users when refreshing, so the instance is not aware of users that are disabled or deleted in AD. The issue is how to detect that an existing, current user is inactive or has been deleted from AD.

For more information on locating inactive accounts, see Find inactive LDAP accounts by using the userAccountControl field.

Note: The recommended approach is to deactivate user records and all other types of records, not delete them. Each record is linked to other records, and deleting a record destroys all the relationships to those other records. Deactivating records keeps those relationships in place.

Find inactive LDAP accounts by using the userAccountControl field

Identify when an Active Directory (AD) user is deleted (or made inactive).

Before you begin

Role required: admin

About this task

One method is to track the active status of AD users and create a business rule to update corresponding accounts when an AD account is inactive.

Procedure

1. Create a new string field on the User [sys_user] table to track the value of the AD **userAccountControl** field.
   For example: `u_ad_user_account`.
2. Create an LDAP transform script to set the field value.

Example

```
target.u_ad_user_account = source.userAccountControl
```

3. Update the LDAP filter to show disabled AD accounts.

   Example
   Here is an example of a filter.

   ```
   (&(objectClass=person)(sn=*)(!(objectClass=computer))(!
   (userAccountControl:1.2.840.113556.1.4.803:=2)))
   ```

   Here is an example of a replacement filter you can use.

   ```
   (&(objectClass=person)(sn=*)(!(objectClass=computer)))
   ```

4. Create an onChange business rule to set the active field to false whenever the u_ad_user_account field has the value 514.
   '514' indicates an inactive account.

LDAP script examples

The following script examples assume you use an Active Directory (AD) for your LDAP server.

userAccountControl attribute values script

This example tests the source for the userAccountControl attribute values associated with a disabled user (514 or 546).

```
//Deactivate LDAP-disabled users during transform based on
 'userAccountControl' attribute
if(source.u_useraccountcontrol == '514' ||
 source.u_useraccountcontrol == '546'){
   target.active=false;
   target.locked_out=true;
}
```

Here is an example using a bitwise check:

```
if(source.u_useraccountcontrol & 2){
   active = false;
}
```

userAccountControl attribute script

This example examines the userAccountControl attribute but does not test for specific values. It also contains the option of reactivating LDAP user accounts.

```
/*
* Deactivate LDAP-disabled users during transform based on
 'userAccountControl' attribute
* Convert the userAccountControl attribute back to a hex value
*/
var ctrl = parseInt(source.u_useraccountcontrol, 10);
ctrl = ctrl.toString(16);

/*
* The only digit we care about is the final one
* A final hex digit value of '2' in 'ctrl' means disabled
*/
if(ctrl.substr(-1) == "2"){
```

```
   //Deactivate and lock the user account
   target.active = false;
   target.locked_out = true;

   //Ignore any insert of a disabled record
   if(action == 'insert'){
      ignore = true;
   }
}
/* Optional: Uncomment else block to reactivate and unlock the
 user account
else {
   target.active = true;
   target.locked_out = ctrl.substr(-2, 1) == "1";
}
*/
```

onBefore transform map script

Here is an example of a onBefore transform map script. The script identifies disabled records and records being inserted. If an insert of a disabled user is occurring, then the operation transform ignores the record.

```
//Ignore any insert of a disabled record as defined by the
 'userAccountControl' attribute
var uc = source.u_useraccountcontrol;
if((uc == '514' || uc == '546') && action == 'insert'){
   ignore = true;
}
```

DN member script

This script example introduces flexibility by not relying on the 546 and 514 userAccountControl values, but instead checking whether the user is a member of a particular Distinguished Name (DN). You can use this script either in the **Script** field of the 'Table Transform Map' record or in an onBefore transform map script.

```
//Deactivate LDAP-disabled users during transform based on OU
 membership in 'dn'
if(source.u_dn.indexOf('OU=Disabled Accounts') > -1){
   target.active = false;
   target.locked_out = true;
}
```

Active Directory Application Mode (ADAM)

Active Directory Application Mode (ADAM) is an Lightweight Directory Access Protocol (LDAP)-compliant directory service.

Note: A basic level of understanding with Microsoft Windows Server and Active Directory is needed for understanding this topic. You must also have administrator permissions on the server you are configuring for ADAM.

These are sample procedures. Due to installation and environment variations, we cannot offer direct support. We recommend working with a Microsoft consultant.

ADAM has a simple install and runs as a service on Windows operating systems. It can be fully customized and distributed as an application component or used as a stand-alone LDAP

directory. ADAM uses the same technologies found on Active Directory Domain Controllers (including replication and delegation features) and has its own administration and customization features. It can be run as a Windows service. ADAM can be installed on Windows XP, 2000, 2003, and 2008 operating systems. ADAM is included as part of Windows Server 2003 R2 and Windows Server 2008. A download is available at http://www.microsoft.com/downloads http://www.microsoft.com/downloads for earlier operating systems.

Security

Some company security policies prohibit external vendors and partners from connecting directly to an Active Directory (AD) Domain Controller. If exposing certain AD objects or attributes to an external vendor or partner is prohibited, access to objects and attributes can be blocked using AD Security Access Control Entries (ACE or ACL). Depending on security requirements, this method can introduce complexity in the integration. Consolidating multiple domains and forests is recommended. If all LDAP imports and authentications need to be channeled through a single source, ADAM can be used as a consolidated source. With the release of Windows 2008 this functionality has been renamed to Light-Weight-Directory Service, LDS. Installation and configuration is similar to Windows Server 2003 R2.

Recommended Knowledge

For this task, you must understand AD, object classes and attributes. To have a successful integration, you need to be knowledgeable of the current AD object structure, familiar with Active Directory delegations, and have a strategy on how to use ADAM and for what purposes. If you are not familiar with AD or ADAM, work with your AD administrator to configure a new ADAM environment.

Trusts

If *`userProxy`* objects is used, the computer hosting ADAM needs to be a member of the domain that has the AD accounts, or a member of a trusted domain.

Internal Connectivity

If *`userProxy`* objects is used, the ADAM computer must be able to connect to the related Domain Controllers to perform proxy authentication.

Related topics

http://www.microsoft.com/downloads/en/details.aspx?familyid=9688f8b9-1034-4ef6-a3e5-2a2a57b5c8e4&displaylang=en%7C

Configuring an instance with ADAM

The first install copies the ADAM files to your computer, registers requires components, and creates the application shortcuts.

Before you begin

Role required: admin

About this task

By default, all of the application files are installed to `%systemroot%\ADAM`.

- Windows Server 2003 R2 - ADAM can be installed using the **Control Panel > Add and Remove Programs > Optional Component Manager**.
- Windows Server 2000 & Windows XP - Downloaded http://www.microsoft.com/downloads from Microsoft.

Create the first instance service which functions as the first directory service hosted by ADAM. Do one of the following:

Procedure

- Run `adaminstall.exe` from the ADAM folder.
- Use the **Create an ADAM instance** shortcut from the **Start Menu > Programs > ADAM** folder.

1. Select the **A unique instance** install option.

   Note: You can use this option to install an instance replica on a second server to provide a fault tolerant system.

2. Complete the fields.

ADAM Instance

Field	Description
Instance Name	used primarily to identify the Windows Service name and display name
Ports	sets the port numbers to be used for LDAP and LDAPS Listeners. The default LDAP port is 389, LDAPS is 636. If these ports are in use on the server, the setup wizard selects new ports. Work with your network administrator to determine the best ports to use
Application Directory Partition	creates an application directory partition. Not needed at this step, we recommend creating the new partition now. A good practice is to use the same distinguished name as your forest or domain, but replace the highest level domain with adam instead of com or local. For example, if your forest partition is `dc=myCompany,dc=com`, you could create the ADAM partition as `dc=myCompany,dc=adam`
File Locations	selects the location(s) for the ADAM partition data.
Service Account Selection	selects a service account that the instance runs as. For stand-alone services, you can use the default network service account. If you plan on using replicas, you need to use an account that has access to all ADAM instances.
ADAM Administrators	the delegation on the ADAM directory that leverages Windows integrated authentication. This is how the initial access is granted for administration. Once the initial account is granted rights, this user or group delegates rights to other Windows users or ADAM users. You can select the default to only grant admin access to the current user, or grant access to a different user or group based on your needs.
Import LDIF Files	the files to import. MS-UserProxy is the most important file to import, but it's worth adding all available files since there is little overhead to the schema and you won't have to worry about extending it later if your needs expand. Confirm the details and the wizard complete the configuration.

Related topics

http://www.microsoft.com/downloads/en/details.aspx?familyid=9688f8b9-1034-4ef6-a3e5-2a2a57b5c8e4&displaylang=en%7C

Set up the ADAM console

Set up the ADAM console. Even though there are many similarities between ADAM and Active Directory, the administration can be very different since there is no **Users and Computers** management console.

Before you begin
Role required: admin

About this task
Most of the general administration is performed using the ADAM ADSI MMC console available from the **ADAM** start menu. The first time you run the ADAM ADSI console, you must connect to the partition you created.

Procedure

1. Right-click the **ADAM ADSI Edit** item in the left frame.
2. Give the new connection a name and update the server name and port fields with the information used when you created the instance.
3. Select **distinguished name** or **naming context** and specify the distinguished name of the application partition you created earlier.
   You can connect to the `Configuration and Schema` partitions for advanced configuration options.
   You should now be able to see into the partition and the default containers for LostAndFound, NTDS Quotas, and Roles. The Roles container has not been configured yet.

Related topics

http://www.microsoft.com/downloads/en/details.aspx?familyid=9688f8b9-1034-4ef6-a3e5-2a2a57b5c8e4&displaylang=en%7C

Create containers and organizational units for ADAM

Logically group objects stored in ADAM into containers and organizational units (OU) just as they would be in Active Directory.

Before you begin
Role required: admin

Procedure

1. Right-click the root partition and navigate to **New** > **Object** > **organizationalUnit**.

   Note: You can also view the list of other objects that are available. This list varies based on the schema extensions installed when you imported the LDF files.

2. When prompted for a value, enter the name of the OU such as `Users`.
   The screen displays a **More Attributes** button.
3. Use the button to assign values to additional attributes.
   For the OUs and containers, no additional values are needed.
   After creating the OUs, the new OUs are listed as a child of the root object.

Related topics

http://www.microsoft.com/downloads/en/details.aspx?familyid=9688f8b9-1034-4ef6-a3e5-2a2a57b5c8e4&displaylang=en%7C

Delegation with ADAM

Once the OU structure is created, define the permission delegations to properly secure the objects to limited users.

As with Active Directory, there are two general ways to grant permissions:

- Add users to a group that already has the appropriate permissions assigned.
- Define new permissions on the ADAM objects.

For this task, we discuss object level permissions. Refer to the Group Administration section for information on group memberships.

Since we don't have a Users and Computers console for ADAM, all object level permissions are defined using the Active Directory utility `DSACLS.exe`. This file is found in the ADAM program directory. When running ADAM utilities it is best to launch the ADAM Tools Command Prompt. This ensures the proper versions of the tools. DSALCS is used to view and set object access rights.

Example: "`dsacls \\localhost:50010\dc=myCompany,dc=adam`" displays the permissions assigned to the root of partition `dc=myCompany,dc=adam` running on the localhost, port 50010. DSACLS is a complex tool used to create complex delegation. Run `"DSACLS /?"` for usage notes.

Related topics

Create containers and organizational units for ADAM

Use ADAMSync to populate ADAM

http://www.microsoft.com/downloads/en/details.aspx?familyid=9688f8b9-1034-4ef6-a3e5-2a2a57b5c8e4&displaylang=en%7C

Populating ADAM Objects

ADAM Objects include User Objects, UserProxy Object, and Group Objects.

User Objects

Users can be created using the ADAM ADSI Edit console just as we did for OU creation. Users can also be administered using AD command line tools, which is beyond the scope of this document. The only mandatory attribute for new user objects is the `cn`, which is a short name or the user's full name. There are also a wide range of optional attributes similar to Active Directory user attributes. You can access the full list of attributes by selecting properties from the user object.

UserProxy Objects

For ServiceNow LDAP integration we recommend you use `UserProxy` objects in ADAM which creates a proxy account that links to the related AD user account. This allows you to have ADAM authenticate logon credentials using AD usernames and passwords from the domain without ServiceNow directly connecting to the Domain Controller. `UserProxy` objects are very similar to AD and ADAM User objects except that do not store passwords and has an `objectSID` attribute that contains the SID from the linked AD User object. This is how the proxy works. `UserProxy` objects are created using the `ADSIEdit` console or command line tools, but this can be tedious. It is recommended that you use an automated process as defined below.

Group Objects

Groups are created using the ADSIEdit console and AD command-line tools. Group concepts are similar to AD and are used to integrate groups and members to ServiceNow. The biggest difference is ADAM groups can contain members from ADAM or from trusted AD Domains.

Automating ADAM Object Creation

If you are interested in synchronizing Active Directory accounts to ADAM, we recommend you use Microsoft *ADAMSync* tool. This is the most common use of ADAM for ServiceNow LDAP integration.

About Permission Delegation

ADAM contains some built-in groups with default permissions. These groups are found in the container *cn=roles,dc=myCompany,dc=adam*. These are similar to domain level groups and have rights to objects in the current partition. Similar to AD Forests you can also set a higher level of permissions using the default groups in *cn=roles,cn=configuration,dc=myCompany,dc=adam*. You must connect to the configuration partition in *ADSIEdit*. The Administrators group by default includes the account specified during the setup. This member is not always visible since it's inherited through the configuration groups. Administrators have full control of all partition objects. The Readers group does not contain any members by default and has read access to all objects in the partition. The Users group is a dynamic group just as it is in Active Directory. Transitively it includes all ADAM users created in the partition.

Related topics

http://www.microsoft.com/downloads/en/details.aspx?familyid=9688f8b9-1034-4ef6-a3e5-2a2a57b5c8e4&displaylang=en%7C

Testing and troubleshooting ADAM setup

The primary tool used for testing is LDP. This allows you to fully test user authentication.

Most of the object management can be completed using the ADAM ADSI Edit console which will provide access to the entire collection of objects and attributes. The highest level of control and troubleshooting ADAM services is using the Windows service created during the instance setup. The service name will vary and depends on the name of the instance created. This service must be running in order for the ADAM service to run. If you are experiencing connection problems, you should review the network configurations to ensure you have the appropriate network access to connect to the server and ADAM port. For each ADAM instance installed, a Windows Event Log is created. This is also a great tool for troubleshooting ADAM services.

The Windows Security Event Log is also helpful when troubleshooting *userProxy* authentications. All *userProxy* logon attempts are logged in the Security Log and reference the remote client device address, the distinguished name of the user trying to log on, and the result or status code.

Related topics

http://www.microsoft.com/downloads/en/details.aspx?familyid=9688f8b9-1034-4ef6-a3e5-2a2a57b5c8e4&displaylang=en%7C

Backup and recovery with ADAM

All ADAM data can be backed up using standard file system backup methods.

Redundancy

ADAM has built-in replication utilities based on the same technology as AD. A full read and write replica of an ADAM partition can exist on the same or different computer. You can use this replica in a variety of ways to provide a fault-tolerant LDAP integration with the instance. One option is to expose both partitions to the instance through the firewall and define both servers in the LDAP Properties server field.

Related topics

Active Directory Application Mode (ADAM)

http://www.microsoft.com/downloads/en/details.aspx?familyid=9688f8b9-1034-4ef6-a3e5-2a2a57b5c8e4&displaylang=en%7C

Use LDAPS with ADAM

The default configuration for *userProxy* object authentication is to enforce LDAPS (secure LDAP) communications. LDAPS requires SSL certificates to secure the network traffic.

To remove this requirement make the following change using the *ADSIEdit* console connected to the configuration partition.

```
Object: CN=Directory Service, CN=Windows NT, CN=Services,
 CN=Configuration
Attribute: msDS-Other-Setings
Value: change RequiresSecureProxyBind from 1 (enforced) to 0
 (disabled)
```

Restart the ADAM service to use the new setting.

To support secure binds and encrypt the user and password information being transmitted, a SSL certificate must be installed on the server and any LDAP client. Since there is limited and controlled uses to the ADAM service, it is feasible to use a self-signed certificate which would meet the needs without incurring certificate costs or building a Certificate Authority (CA) infrastructure. If you already have a CA, you can issue a certificate. Otherwise, create a self-signed certificate.

Creating a Self-Signed Certificate

To use the *selfssl* utility, Internet Information Services (IIS) must be installed. This service can be removed after you generate the certificate. You can get the selfssl.exe utility from the IIS Resource Kit. If IIS is already installed, create a new website so that the current sites will not be impacted during the certificate generation. Selfssl needs to temporarily attach the new self-issued certificate to a valid web site.

Selfssl is a command-line tool and has the following common parameters.

Selfssl Parameter Descriptions

Parameter	Description
/T	Adds the cert to 'Trusted Certificates' on the local machine
/N:cn	Set the common name of the certificate. This must match the fully qualified domain name of the server running the web service using the certificate
/K	Sets the strength of the key size in bits
/V	Number of days the cert is valid

Selfssl Parameter Descriptions (continued)

Parameter	Description
/S	Web site ID to attach the certificate to
/P	IP port of the web service

The common name attribute should match the external name or address that the instance will use to connect to your ADAM computer. You will need to get the IIS Website site id unless you are using the default website which is 1 and does not need to be defined in the selfssl command. A sample command to generate a certificate for myCompany would be:

```
selfssl /N:CN=myCompany.externaldomain.com /K:1024 /V:3650 /
S:12345 /P:50001 /T
```

This statement creates a certificate that is valid for 10 years. Set the value to any duration, but be aware the new certificate must be generated and submitted to the instance before the old one expires. We recommend making a note of the expiration date on the certificate.

Once the certificate is generated you can remove it from the website, or delete the entire web site if you created a temporary site.

Related topics

http://www.microsoft.com/downloads/en/details.aspx?familyid=9688f8b9-1034-4ef6-a3e5-2a2a57b5c8e4&displaylang=en%7C

Assign the certificate to ADAM

Install an SSL certificate on the server and any LDAP client to support secure binds and encrypt the user and password information being transmitted.

Before you begin

Role required: admin

About this task

Because there is limited and controlled uses to the ADAM service, it is feasible to use a self-signed certificate to meet your needs without incurring certificate costs or building a Certificate Authority (CA) infrastructure.

Procedure

1. Open the Certificates MMC console and create two console connections, one for Local Computer Certificates, and the other for Local Computer Services Certificates on the new ADAM service.
   The new certificate can be found under `Certificates (Local Computer)\Personal\Certificates`.
2. Copy the certificate to the container for the ADAM service `Certificates - Service (ADAM Service Name)\ADAM_ADAM Service Name\Trusted Root Certificates\Certificates` and copy the certificate to `Certificates - Service (ADAM Service Name)\ADAM_ADAM Service Name\Personal \Certificates`.
3. Open the details tab on the certificate that you copied, note the Valid from date stamp, and assign read access to the certificate key file.
   Go to `C:\Documents and Settings\All Users\Application Data \Microsoft\Crypto\RSA\MachineKeys` and identify the certificate with the matching

time stamp. Assign Read & Execute rights to the service account running ADAM. By default, this is **Network Service**.

4. Restart the ADAM service to activate the new certificate.

Related topics

http://www.microsoft.com/downloads/en/details.aspx?familyid=9688f8b9-1034-4ef6-a3e5-2a2a57b5c8e4&displaylang=en%7C

Export the public key certificate

LDAPS clients, including the instance need the public key certificate in order to make a secure connection to ADAM.

Before you begin

Role required: admin

About this task

From the server certificate consoles you used above, export a public key to be used by the clients.

Procedure

1. Select the certificate, right-click, and select **all tasks/export**.
   Do not export the private key. Select the default DER encoded binary X.509 format and specify the export file name.
2. Install the public certificate on the LDAP clients that connect to the server using LDAPS. When prompted, add the certificate to the *Trusted Root Certificate Authorities* store.

Related topics

http://www.microsoft.com/downloads/en/details.aspx?familyid=9688f8b9-1034-4ef6-a3e5-2a2a57b5c8e4&displaylang=en%7C

Active Directory Application Mode (ADAM) Access Account

The system requires a user account to read the Active Directory Application Mode (ADAM) object information that is imported into the application instance.

Create the account by using one of the following methods:

- Create a local ADAM user account and assign it a password and assign permissions.
- Assign permission to a Windows domain account on the ADAM partition.
- Use a *userProxy* account.

When using ADAM as an LDAP source, you must specify the fully qualified distinguished name (FQDN) of the ADAM account in the instance's LDAP server's **Login distinguished name** field.

Related topics

Active Directory Application Mode (ADAM)

http://www.microsoft.com/downloads/en/details.aspx?familyid=9688f8b9-1034-4ef6-a3e5-2a2a57b5c8e4&displaylang=en%7C

Test the LDAPS connections

Test the LDAPS connections. There are two console connections, one for Local Computer Certificates, and the other for Local Computer Services Certificates on the new ADAM service.

Before you begin

Role required: admin

Procedure

1. Run `LDP.exe` from the ADAM install folder `c:\windows\adam`.
   Verify that the ADAM version is selected because this is not the standard Windows LDP client.
2. Open a new connection by using the **Connection/Connect** menu.
   The server name must match the CN that is assigned to the certificate.
3. Enter the **LDAPS port** and select the **SSL** check box.
   The results of a successful connection are some general server information and no errors.
4. Bind (log in) to the service.
   To replicate typical LDAP client connections, select the Simple bind option. Enter a valid ADAM user or `userProxy` distinguished name in the user field and the associated password.
   If you see a return message stating 'Authenticated as:....' then you have successfully connected using LDAPS.

Related topics

http://www.microsoft.com/downloads/en/details.aspx?familyid=9688f8b9-1034-4ef6-a3e5-2a2a57b5c8e4&displaylang=en%7C

Use ADAMSync to populate ADAM

Administrators use MS ADAMSync to populate LDAP directories that use MS ADAM.

Note:

This document assumes you have at least a basic level of understanding with Microsoft Windows Server, Active Directory, and ADAM and that you already have a functional Active Directory Application Mode (ADAM) instance with a partition.

These are sample procedures. Due to the complexity and the fact that it is running in your environment, we cannot offer direct support. We recommend you work with Microsoft or a Microsoft consultant if you run into any trouble.

Once ADAM has been installed and the first partition has been created, you can populate it with objects.

The following options are available:

- Manual object creation using GUI or scripts. This option is inefficient and slow.
- Integrate with Active Directory using Microsoft Integration Information Server. This option ultimately provides the most flexibility and functionality but does require some advanced configurations. There is a free version of MIIS available that is compatible with Active Directory, ADAM, andMicrosoft Global Address Lists from Exchange. Unless you already have experience with MIIS we advise that you don't attempt to implement a new environment for LDAP integration only.
- Use ADAMSync, a synchronization tool that Microsoft provides with ADAM. This is the option that is explained here.

Define ADAM user accounts

Define user accounts in ADAM. One user account is used for the instance to connect with and the other user account is for ADAMSync.

Before you begin

Role required: admin

About this task

These accounts can be local ADAM User objects, UserProxy objects, or a Windows account from a trusted domain.

The *ADAM User* account requires read-only access to the directory structure you are importing to your instance. The best way to accomplish this is to add the account to the member attribute on the Readers group found in cn=roles,dc=myCompany,dc=adam.

New ADAM User accounts are disabled by default. You will need to enable the new accounts and set a password.

Procedure

1. Enable users by changing the attribute msDS-UserAccountDisabled to FALSE.
2. Right-click the user object and reset the password.
3. Test the new accounts by using LDP as defined in Active Directory Application Mode (ADAM) to make sure they can connect.
   Use the **LDAP** > **View/Tree** option, leaving the Base DN blank to make sure that you can view the objects in the directory by using the new accounts. The Configuration, Schema, and the domain partition should be visible in the left pane. Traverse the domain partition. If you are using a new local ADAM account, it will show 'No Children' which means that you don't have read access to the objects. Verify the Setup group memberships and re-test.

   ADAMSync uses the *ADAMSync User* account to manage objects in the ADAM partition. This account requires admin level rights since it will create, update, and delete ADAM objects.

   ADAMSync uses the *ADAMSync AD* account to read the AD objects that will be synchronized to ADAM.

Set up ADAMSync

ADAMSync is included with Windows Server 2003 R2. Download and install ADAMSync if you are using a different OS.

Extending the schema

The ADAM schema needs to be extended to support ADAMSync.

1. Run the following command from c:\windows\adam to import the ADAMSync schema extensions. You may have to change the server:port and add credentials if the current user doesn't have access. See the AdamSyncMetadata.ldf file for details.

```
ldifde -i -f MS-AdamSyncMetadata.LDF -s localhost:50000 -j . -c
 "cn=Configuration,dc=X" #configurationNamingContext
```

2. Do the same with MS-AdamSchemaW2k3.ldf to support Windows 2003 attributes.

```
ldifde -i -u -f MS-AdamSchemaW2K3.LDF -s localhost:50000 -j .
 -c "cn=Configuration,dc=X" #configurationNamingContext
```

Recommended schema changes

Here are some additional schema changes we recommend.

1. Open a new MMC console and add the ADAM Schema Snap-in.
2. Connect to the ADAM instance.
3. Expand the Classes folder and locate the userProxy class, open **Properties**.

4. Verify the following optional attributes on the Attributes tab, add any that do not already exist.
   - company
   - department
   - givenNane
   - mail
   - physicalDeliveryOfficeName
   - sAMAccountName
   - sn
   - telephoneNumber
   - title
   - userAccountControl
   - userPrincipalName
5. Restart the ADAM Service to enable the new settings.

Install the ADAM configuration file

Install the ADAM configuration file through the Windows command line.

Before you begin
Role required: admin

Procedure

1. Install the configuration file.

```
C:\WINDOWS\adam>adamsync /install localhost:50000
MS-AdamSyncConf-SNC.XML
```

2. Run the synchronization file to log to the console.

```
C:\WINDOWS\adam>adamsync /sync localhost:50000
"ou=users,dc=service-now,dc=adam" /log -
```

3. Review the results by using the ADSIEdit console.
   You should see the new objects and attributes that were created by ADAMSync.
4. Run ldap to test the UserProxy authentication.
   Automating the sync process

   Set up the sync process as a Windows Scheduled Task. You must either provide the credentials in the config file, command line, or run the Scheduled Task with an account that has access.

   Special notes
   - You can create multiple configuration files and scheduled jobs to sync ADAM from multiple sources.

     This example imports the sAMAccountName attribute which can be used as the application logon. If you are going to sync source you need to make sure you have a unique attribute

value that can be used for the logon credentials. sAMAccountName is guaranteed to be unique within a domain, but not across multiple domains.

- If you are using Microsoft Exchange, we recommend excluding cn=SystemMailbox* objects as part of the object-filter configuration.

Example ADAM configuration files

All of the configurations for ADAMSync are stored in xml files.

Default configuration file with comments

There is a default configuration file called MS-AdamSyncConf.xml included with the ADAMSync install. Make a copy of this file so you have a base example to refer to in the future. This example is the default configuration file with comments added.

```
<?xml version="1.0"?>
<doc>
 <configuration>
<!-- Sync File Description -->
<description>MyCompany ADAMSync Configuration</description>
  <security-mode>object</security-mode>;
<!-- source-ad-name = fqdn of the domain controller -->;

 <source-ad-
name>;fully.qualified.domain.name.of.domain.controller</sour
ce-ad-name>;
<!-- source-ad-partition = root AD domain partition -->;

 <source-ad-partition>;dc=myCompany,dc=com</source-ad-
partition>;
<!-- source-ad-account = use this to specify an account to
 connect to AD -->;
<!-- if not used, the current user will be used  -->;
  <source-ad-account>;</source-ad-account>;
  <account-domain>;</account-domain>;
<!-- target-dn = target ADAM OU -->;
  <target-dn>;ou=servicenow
 users,dc=myCompany,dc=adam</target-dn>;
  <query>;
<!-- base-dn = should be the root AD partition if you want all
 users -->;
   <base-dn>;dc=myCompany,dc=com</base-dn>;
<!-- object-filter = standard ldap query format, this will grab
 all users -->;
<!-- need to review results to see if you should modify this
 filter -->;
   <object-filter>;(objectCategory=person)</object-filter>;
   <attributes>;
<!-- include=userproxy requires objectSID to link back to the AD
 account -->;
    <include>;objectSID</include>;
    <include>;givenName</include>;
    <include>;sn</include>;
    <include>;description</include>;
    <include>;title</include>;
    <include>;company</include>;
    <include>;department</include>;
```

```
    <include>;mail</include>;
    <include>;physicalDeliveryOfficeName</include>;
    <include>;telephoneNumber</include>;
    <include>;sAMAccountName</include>;
   </attributes>;
  </query>;
<!-- map for user-to-userproxy object types -->;
  <user-proxy>;
    <source-object-class>;user</source-object-class>;
    <target-object-class>;userProxy</target-object-class>;
  </user-proxy>;
  <schedule>;
   <aging>;
    <frequency>;0</frequency>;
    <num-objects>;0</num-objects>;
   </aging>;
   <schtasks-cmd>;</schtasks-cmd>;
  </schedule>;
 </configuration>;
 <synchronizer-state>;
  <dirsync-cookie>;</dirsync-cookie>;
  <status>;</status>;
  <authoritative-adam-instance>;</authoritative-adam-instance>;
  <configuration-file-guid>;</configuration-file-guid>;
  <last-sync-attempt-time>;</last-sync-attempt-time>;
  <last-sync-success-time>;</last-sync-success-time>;
  <last-sync-error-time>;</last-sync-error-time>;
  <last-sync-error-string>;</last-sync-error-string>;
  <consecutive-sync-failures>;</consecutive-sync-failures>;
  <user-credentials>;</user-credentials>;

 <runs-since-last-object-update>;</runs-since-last-object-updat
e>;
  <runs-since-last-full-sync>;</runs-since-last-full-sync>;
 </synchronizer-state>;
</doc>;
```

LDAP filters configuration file

You can provide any level of filtering in the object-filter value in the configuration file. Use standard LDAP query syntax with the following xml escape characters in place of the standard operators.

- AND = "&" replace with &
- OR = "|" (vertical line) replace with |
- NOT = "!" replace with !

Reference configuration file

Here's an actual configuration file that can be referenced as a sample.

```
<?xml version="1.0"?>;
<doc>;
 <configuration>;
<description>;SNCTest ADAMSync Configuration</description>;
  <security-mode>;object</security-mode>;
```

```
 <source-ad-name>;domaincontroller.service-now.com</source-ad-
name>;

 <source-ad-partition>;dc=service-now,dc=com</source-ad-
partition>;
  <source-ad-account>;</source-ad-account>;
  <account-domain>;</account-domain>;
  <target-dn>;ou=servicenow
 users,dc=service-now,dc=adam</target-dn>;
  <query>;
   <base-dn>;dc=service-now,dc=com</base-dn>;
   <object-filter>;(objectCategory=person)</object-filter>;
   <attributes>;
    <include>;objectSID</include>;
    <include>;givenName</include>;
    <include>;sn</include>;
    <include>;description</include>;
    <include>;title</include>;
    <include>;company</include>;
    <include>;department</include>;
    <include>;mail</include>;
    <include>;physicalDeliveryOfficeName</include>;
    <include>;telephoneNumber</include>;
    <include>;userAccountControl</include>;
   </attributes>;
  </query>;
  <user-proxy>;
    <source-object-class>;user</source-object-class>;
    <target-object-class>;userProxy</target-object-class>;
  </user-proxy>;
  <schedule>;
   <aging>;
    <frequency>;0</frequency>;
    <num-objects>;0</num-objects>;
   </aging>;
   <schtasks-cmd>;</schtasks-cmd>;
  </schedule>;
 </configuration>;
 <synchronizer-state>;
  <dirsync-cookie>;</dirsync-cookie>;
  <status>;</status>;
  <authoritative-adam-instance>;</authoritative-adam-instance>;
  <configuration-file-guid>;</configuration-file-guid>;
  <last-sync-attempt-time>;</last-sync-attempt-time>;
  <last-sync-success-time>;</last-sync-success-time>;
  <last-sync-error-time>;</last-sync-error-time>;
  <last-sync-error-string>;</last-sync-error-string>;
  <consecutive-sync-failures>;</consecutive-sync-failures>;
  <user-credentials>;</user-credentials>;

 <runs-since-last-object-update>;</runs-since-last-object-updat
e>;
  <runs-since-last-full-sync>;</runs-since-last-full-sync>;
 </synchronizer-state>;
</doc>;
```

Configure Microsoft Active Directory for secure LDAPS communication

Use certificate pairs to enable Microsoft Active Directory (AD) LDAPS communications.

Note: These procedures were designed and tested using Windows 2003 R2 Standard Edition and work with all versions of Windows 2003.

Secure LDAP (LDAPS) communication is similar to SSL (HTTPS) communication in that both encrypt the data between servers and clients. To accomplish this, the server and clients share common information by using certificate pairs. The server holds the private key certificate and the clients hold the public key certificate. These certificates are required to enable Microsoft Active Directory (AD) LDAPS communications.

To configure LDAPS for Active Directory you must:

- Ensure that the Active Directory domain is set up and that the instance is able to connect to the Active Directory server through the firewall.
- Verify that there is a Certificate Authority (CA) that can issue a certificate for the domain controller (DC). If you don't already have a CA infrastructure there are two options.
  - Setup a stand-alone CA to issue the certificate
  - Request a third party certificate
- If you already have a CA in place, you can generate a certificate from an internal CA.

All certificates have a defined expiration date which can be viewed in the certificate properties. If the certificate expires, all LDAPS traffic fails, and your users can no longer log into the instance. To resolve this, a new certificate must be issued and installed on your instance.

The default expiration for Microsoft CA certificates is one year. External CA certificates are usually purchased in one year increments. Note when your certificate expires, or use the application's Expiration Notification function (located in **System LDAP** > **Certificates**). Ensure that you have a new certificate ready before the old one expires. This gives you time to install and test the new certificate before the old one expires.

Set up a stand-alone certificate authority for active directory

The first step to configure Microsoft Active Directory for SSL access is to set up a stand-alone Certificate Authority (CA).

Before you begin

Role required: admin

About this task

Do not worry about addition resource utilization because both of the required services (IIS & CA) can be disabled after issuing the certificate(s).

Procedure

1. Install Internet Information Server (IIS).
2. Install Certificate Authority Services in stand-alone mode.
3. Verify the Certificate Services web application is installed and active.

What to do next

Using the IIS Manager console, expand the local computer and select *Web Sites*. The state of **Default Web Site** should be *Running*. You should also see a *CertSrv* application listed under the **Default Web Site**. If the site is not running or the application is missing, you must resolve the issue before you proceed.

Generate a certificate from an internal certificate authority

When you configure Microsoft Active Directory for SSL access, you must generate an internal certificate and request the external certificate.

Before you begin

Role required: admin

About this task

These steps apply to Microsoft CA services. If you have a different internal CA platform, see your local CA administrator for assistance.

Procedure

1. From the domain controller (DC) you want to create a certificate for, browse to `http://localhost/certsrv` or specify the CA server name if it is on a remote server.
2. From the Welcome page, click **Request a certificate** and select *advanced certificate request*.
3. On the Advanced Certificate Request page, select **Create** and submit a request to this CA.
4. Complete the Advanced Certificate Request as follows:

Advanced Certificate Request fields

Field	Entry
Name	The fully qualified domain name (FQDN) of the DC that is requesting the certificate.
E-Mail	The email address of the person responsible for the certificate.
Company	Your company name.
Key Options settings	
Create new key set	Select it.
CSR	Microsoft RSA SChannel Cryptographic Provider.
Key Usage	Exchange.
Key Size	1024 is recommended. The instance supports up to 2048.
Automatic key container name	Select it.
Store certificate in the local computer certificate store	Select it.

5. Click **Submit**.
   You are directed to a page that provides your **Request ID**, make note of this ID.
6. To process the pending request, complete the following:
   a. Open the Certificate Authority management console.
   b. Expand the server node and select **Pending Requests**.
   c. Locate the Request ID for the request you just submitted, right-click, and select *All Tasks/ Issue to approve the request and issue the certificate*.
7. To retrieve the issued certificate, complete the following:

a. From the DC you made the request from, browse to `http://localhost/certsrv`, or specify the CA server name if it is on a remote server.

b. Select `View the status of a pending certificate request`.

c. Select the link to the new certificate.

d. Select the link to **Install this certificate**.

What to do next

You need to request a third party certificate. Certificates from external CAs can be purchased for as little as $30 per year. For detailed procedures on requesting a certificate from an external CA, see Microsoft article 321051. After it is received, installed, and tested, follow the export procedure.

Test the LDAPS connectivity locally

Test the LDAPS connectivity after installing the internal and third party certificates when you configure Microsoft Active Directory for SSL access.

Before you begin

Role required: admin

Procedure

1. Ensure that Windows Support Tools are installed on the domain controller (DC).
   The Support Tools setup (suptools.msi) can be found in the `\Support\Tools` directory on your Windows Server CD.
2. Navigate to **Start** > **All Programs** > **Windows Support Tools** > **Command Prompt**.
   On the command line, enter `ldp` to start the tool.
3. From the *ldp* window, select **Connection** > **Connect** and supply the local FQDN and port number (636).
   Also select the *SSL*.
   If successful, a window displays and lists information related to the Active Directory SSL connection. If the connection is unsuccessful, try restarting your system and repeat this procedure.

Export the public key certificate to trust the LDAP certificate

Export the public key certificate and import it into the application when you configure Microsoft Active Directory for SSL access.

Before you begin

Role required: admin

About this task

If your Certificate Authority is not a trusted third party vendor, you must export the certificate for the issuing CA so we can trust it, and, by association, trust the LDAP server certificate. For MS Certificate Services users, you can view the certificate path by viewing the certificate in the console used to export; select the **Certificate Path** tab. You must export all certificates in the chain. You can find the CA certificate in the same folder as the LDAP certificate by looking for the name in the Certificate Path. Submit all certificates for importing to your instance.

Procedure

1. From a current or new MMC console, add the Certificate (Local Computer) snap-in.
2. Open the `Personal/Certificates` folder.
3. Locate the new certificate.

The Issued to column shows the FQDN of the domain controller.

4. Right-click the certificate and select *All Tasks/Export*.
5. Export to DER or Base-64 format.
   Name the file using the format `MyCompany.cer`. This is the public key certificate the needs to be used on the instance to communicate securely with your domain controller.
6. Test LDAPS locally before you submit the certificate to the instance.

What to do next
After completing this procedure, import the public key certificate into the application.
See Install the LDAP X.509 SSL certificate to upload the certificate into the application.

LDAP global catalog usage

A DC can be granted the Global Catalog (GC) role. Global Catalog (GC) role is an LDAP-compliant directory consisting of a partial representation of every object from every domain within a forest.

Administrators configure Active Directory to host Lightweight Directory Access Protocol (LDAP) directory information using one of the following hosting methods.

- The common method of hosting LDAP directory information is to use the default LDAP or LDAPS (secure LDAP) on ports 389 or 636. These standard LDAP ports always exist on a Domain Controller (DC) and are rarely changed. Accessing this directory partition provides access to all of the objects within the domain that is hosted on the DC. There is no way to access objects from other domains using this method.
- A DC can also be granted the Global Catalog (GC) role. Global Catalog (GC) role is an LDAP-compliant directory consisting of a partial representation of every object from every domain within the forest. This LDAP directory can be accessed on port 3268, with LDAPS on port 3269. LDAPS and the default LDAP ports' certificate requirements are the same.

Global Catalog LDAP dependencies

- The domain controller that your instance connects to must have the Global Catalog role enabled.
- Firewall rules must allow inbound traffic to the domain controller on port 3268 (LDAP) or 3269 (LDAPS).

Special notes

- Not all attributes are replicated to the GC partition. Common attributes such as first name, last name, email, phone number, description, and address are included. Additional attributes can be added to the GC but should be limited to minimize the impact to forest replication traffic.
- Standard LDAP integrations usually use sAMAccountName as the instance's UserID and as the coalesce key in the LDAP import map since this is guaranteed to be unique within a domain. This attribute is no longer unique when viewing an entire forest of domains. A new unique attribute needs to be identified and as the UserID and the coalesce key. These do not need to be the same attribute and may vary based on your forest design. Consult your Active Directory administrator. Typically, the userPrinicpalName is a unique attribute across domains but this may not be a user-friendly name to login with, but it could be used for the unique identifier on imports. A common attribute that is used for the UserID is email address. These decisions impact the LDAP Properties and LDAP Mapping.
- The value used for the coalesce key on the LDAP import map must be unique and exist on every object being imported. If it is not unique or does not exist, incorrect records are updated with changes.

- If you already have an LDAP integration and wish to change it to a GC, change the import coalesce key. The new key values must be imported before you can change the coalesce key.
- If you make any changes to your LDAP integration that break your integration, your first step should be to revert those changes. After that, contact Customer Service and Support with complete information about what you're attempting.

OpenLDAP minor schema modification

In OpenLDAP 2.3 systems that use the back-bdb (Berkley backend), administrators make a minor modification to their schema to facilitate the integration.

Warning: The customization described here was developed for use in specific instances, and is not supported by Now Support. This method is provided as-is and should be tested thoroughly before implementation. Post all questions and comments regarding this customization to our community forum.

In OpenLDAP 2.3, back-bdb has limited support for inequality indexing (ordering). It is implemented only for generalizedTime and ChangeSequenceNumber syntax. It cannot be supported on syntax that support substrings. Search filters containing inequalities are processed using the presence index.

We recommend creating a custom attribute for this purpose, instead of changing what is already indexed or present in the schema (for example, *servnowid*).

Modify the OpenLDAP schema

Modify the OpenLDAP schema. These steps detail a schema modification to OpenLDAP 2.3 provided by one of our customers that helped them integrate with their instance.

Before you begin

Role required: admin

About this task

Warning: The customization described here was developed for use in specific instances, and is not supported by Now Support. This method is provided as-is and should be tested thoroughly before implementation. Post all questions and comments regarding this customization to our community forum.

To modify the OpenLDAP schema for integration with the instance:

Procedure

1. Create a custom attribute.

   Example

```
attribute ( 1.3.6.1.4.1.3403000.2.1.8

    NAME 'servnowid'
  ORDERING caseIgnoreOrderingMatch
  EQUALITY caseIgnoreMatch
  SYNTAX '1.3.6.1.4.1.1466.115.121.1.15' )
```

2. Include the attribute in the selected objectclass OID.

Example

```
objectclass ( 1.3.6.1.4.1.3403000.2.2.1
     NAME 'BcfUserIdentifiers' SUP top AUXILIARY
  MAY ( uniqid $ unixid $ servnowid ) )
```

In OpenLDAP 2.3, you can dynamically change the server configurations, but you can only extend the schema. You cannot modify or delete the existing schema. Instead of creating another objectclass for this attribute in the dynamic configuration, use the static configuration file, slapd.conf.

3. In slapd.conf, include indexing for the new attribute in the bdb section of your main database backend.

Example

```
database bdb (configs here) ....

index servnowid pres

(other indexes here) .....
```

4. As root, run slapindex to index this attribute to make it available in search filters.
Make sure that the OpenLDAP daemon is not running or is in read-only mode before starting slapindex.

Record LDAP deletions

By default, the instance does not delete any entries after they disappear from LDAP.

Deleting an entry, also referred to as a record, also deletes the entire history and references to the deleted entry.

For example, configuration items (CIs), SLA agreements, software licenses, purchase orders, and service catalog entries all have a reference to Department, and if a department is deleted, then the integration clears all references to the department. Also, deleting a user results in losing all history of what that user did.

Decide whether to retain or Delete all records from a table according to your organization's needs.

Limit concurrent sessions

You can limit the number of concurrent interactive sessions for a user or role on an instance across all nodes.

Explore

Activate

Learn the features and business value of limit concurrent sessions.	Understand how to Activate limit concurrent sessions.
Set  Set the limit concurrent sessions for a user or role.	Disable  Know about the how to disable limit concurrent sessions.

Explore limit concurrent sessions

You can limit the number of concurrent interactive sessions for a user or role on an instance across all nodes.

Concurrent interactive sessions refer to the number of sessions a user can have active per ServiceNow instance. An active instance session occurs with every new login to a specific ServiceNow instance. By default, there are no limitations on the number of active instance sessions a user can have.

With the Jakarta release, you can limit the number of active concurrent sessions per user. When the user logs in after hitting the maximum number of sessions active, the oldest active session terminates and a new interactive session becomes active. If a user tries to access a closed session through a browser, the user is redirected to the login page.

Note: The **Limit concurrent sessions** plugin must be active to enable a maximum session limit. Limits are set through the glide.authenticate.max.concurrent.interactive.sessions property. A maximum limit value applies to any user or role that has the limit property active. A user or a role connected to the user must have the limit_concurrent_sessions value set to true for the limit on sessions to initiate. For the Jakarta release, this feature does not support sessions created through the native mobile app or non-interactive mechanisms.

A typical use case if a maximum concurrent session of 1 is set:

1. The user accesses the initial ServiceNow instance through Chrome.
2. After the user successfully logs in, ServiceNow creates session 1 (S1) for the user.
3. The user decides to initiate another access to the ServiceNow instance through Firefox.
4. After the user successfully logs in, ServiceNow creates session 2 (S2) for the user.
5. Since the user has a maximum concurrent session limit of 1, the S1 session invalidates when S2 is created.
6. When the user goes back through Chrome to access the S1 ServiceNow instance, the user is redirected to the login page as S1 is invalid.

Concurrent session limits work with all the ServiceNow authentication mechanisms: SAML, LDAP, and local database authentication. It also works with Multi-factor authentication and all interactive ServiceNow authentication mechanisms. The source of the session is viewable through the **sys_user_session** table, under the column **Type**. The values can be:

- Web Browser
- Mobile Browser
- ServiceNow Mobile App
- Non-interactive (SOAP, WSDL, OAuth)

Activate and configure limit concurrent sessions plugin

You can activate the Limit Concurrent Sessions plugin (com.glide.limit.concurrent.sessions) if you have the admin role.

Before you begin

Role required: admin

About this task

Procedure

1. Navigate to **All** > **System Definition** > **Plugins**.
2. Find and click the Limit Concurrent Sessions plugin.
3. On the System Plugin form, review the plugin details and then click the **Activate/Upgrade** related link.
4. Click **Activate**.
5. To enable this feature and set a maximum limit of concurrent sessions, go to the **Plugin Files** tab, find the following properties, and change the setting values.

Option	Description
glide.authenticate.limit.concurrent.interactive.sessions	You can enable the ability to limit concur rent sessions by setting the value to **True**. By default, this property is set to **False**, which means there is no limit on the number of inter active sessions a user can have active. **Note:** To disable this feature, set this property back to **False**.
glide.authenticate.max.concurrent.interactive.sessions	You can set the maximum number of concur rent active interactive sessions a user can have on the instance across all nodes.

6. **Optional:** You can also amend the following properties, if necessary.

Option	Description
glide.authenticate.session.types.to.limit.concurrency	This property limits session types. By default, only the web browser sessions have a limit. Session types include: ○ Web Browser (1) ○ Mobile Browser (2) ○ ServiceNow Mobile App (3) ○ Non-interactive (10)

Option	Description
	You can configure and set the value to '1' for web browser, '2' for mobile browser, or '1,2' for both. **Note:** Only web and mobile browser sessions can have a limit. There are no limits for sessions that originate from the ServiceNow mobile app or non-interac tive sessions.
glide.authenticate.limit.concurrent.sessions.across.all.nodes	This property restricts the limit of concurrent sessions per node instead of restricting them across all nodes of a ServiceNow instance. By default, the value is set to true, which limits user sessions across all nodes. If the property is set to false, only the sessions on that node and not the ones on the other nodes are sub ject to the limit.

7. Click **Update** to have the settings take effect.

What to do next
Set a concurrent session limit by user or role.

Related topics
List of plugins (Zurich)

Set a concurrent session limit by user or role

You can set a concurrent session limit on a specific user or on a particular role.

Before you begin
Role required: admin

Procedure
1. Navigate to **All** > **User Administration** > **Users** or **User Administration** > **Roles**.
2. Select a user or role that you want to set a concurrent session limit, check the **Limit Concurrent Sessions** check box, and click **Update**.
   The user or role has a limit of how many concurrent sessions can be open at one time.

Disable a concurrent session limit by user or role

You can disable a concurrent session limit on a specific user or on a particular role.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **User Administration** > **Users** or **User Administration** > **Roles**.
2. Select a user or role that you want to disable a concurrent session limit, uncheck the **Limit Concurrent Sessions** check box, and click **Update**.
   The user or role is not subject to a limit of how many concurrent sessions can be open at one time.

Local authentication

Use ServiceNow® local authentication to secure the users login on a local device.

Local authentication is a method of verifying a user's identity on a local device. It uses unique characteristics or credentials possessed by the user to validate the user's login.

Use ServiceNow®'s local authentication features to securely access the ServiceNow® instance on your local devices.

Login and authentication security

Configure login security options to control access to your instance.

Configure your login security Understand how to configure login security.	Define Login Scenarios Define login scenarios.
Password Reset Learn about login password resets.	

Explore Login and authentication security

Configure login security options to control access to your instance.

Security options

You can control several aspects of user login and authentication security:

Feature	Description	Related topics
Log in and log out controls	Control several dimensions of the log in and log out process for users, such as specifying a landing page that the user sees upon login and control how users log out.	• Define login scenarios • Configure the logout confirmation prompt • Installation exits • Specify lockout for failed login attempts
Authentication security	Control the password reset process and features like the Remember Me option. You can also use IP address-based controls for access to the instance and implement a nonce to be used with single sign-on digest authentication.	• Configure your password policy • Password Reset • Remember me • IP range based authentication • Implement a nonce

Define login scenarios

You can direct all users to the same page after login.

Before you begin

Role required: admin

About this task

When users log on to an instance directly, such as going to http://{instance_name}.service-now.com/, the system does the following:

1. Accesses the value in the property *glide.entry.page.script*. The default value of the property is derived from a script include named *CMSEntryPage*.
2. Directs the user to the instance login page if the entry page requires a login.
3. Applies login rules, if any, to the user.

To force the system to direct all users to the same page after login:

Procedure

1. Navigate to **All** > **Content Management** > **Configuration** > **Configuration Page**.
2. Select a value for the *Login page* field, or create a new page as desired.

   If this page is not the site default page, it always redirects here. If it is a site default page, it applies login rules. If this value is null, the system uses navpage.do as the entry page. Do not enter a login page here; otherwise, users need to log in twice.

   Logging Into an Instance to access a record:

   When users log into an instance to access a record by its globally unique identifier (sys_id), such as http://{instance}.service-now.com/incident.do?sys_id={sys_id}, then the system does the following:

   a. Directs the user to a login page if not already logged in.

   b. Directs the user to the appropriate record if they are allowed to access it. If the user does not have access rights to the record, a denial of access message appears.

   Logging Into the Service Portal or a CMS site:

   When users log on the Service Portal or a CMS site, such as http://<instance>.service-now.com/site-name/page.do, the system does the following:

   - If there is a value in the *Login page* field on the Service Portal or the CMS site form, it directs the user to that login page and applies login rules, if any, to the user.
   - If there is no login page specified, it directs the user to the value in the *Home page* field on the Service Portal or the CMS site form.

   Logging Into the Service Portal or a CMS Site to access a record:

   When users log on to the Service Portal or a CMS site to access a record, such as http://{instance}.service-now.com/ess/incident_detail.do?sysparm_document_key=incident,{sys_id}, the system follows the same procedure and finally takes the user to the record. If the user does not have access rights to the record, a denial of access message appears.

Logins and the employee self-service portal

The system keeps track of the first starting page that a user is trying to access even if the user wants to log in to the Employee Self-Service Portal.

Consider the following scenarios.

Example 1:

1. A user is not logged in, and then tries to access a record using a specific SYS ID in the URL.
2. The system redirects the user to the login page.
3. Rather than logging in, the user tries to access another site, such as the Employee Self-Service (/ess) Portal.
4. The system redirects the user to the login page again.
5. The user logs in and is redirected to the record that the user was first trying to access rather than the Employee Self-Service Portal.

Example 2:

1. A user is not logged in, and then tries to access a record using a specific SYS ID in the URL through the Employee Self-Service (/ess) Portal.
2. The system redirects the user to the login page.
3. Rather than logging in, the user tries to access another record through the Employee Self-Service Portal.
4. The system redirects the user to the login page again.
5. The user logs in and is redirected to the first record rather than the second.

Specify a login landing page

By default, users see their homepage upon login. You can specify a different login landing page by using a system property or the content management system.

Before you begin
Role required: admin

About this task
To specify a login landing page for all users, change the property value on the sys_properties table.

Procedure
1. Type `sys_properties.list` in the navigation filter.
2. Locate the **glide.login.home** system property.
3. In the **Value** field, enter the name of the page that all users see upon login.

   Use `<page name>.do`; you may omit the `http://''instance''.service-now.com/` portion of the URL. To determine the page name or the URL of a page in the system, you can point to a link. Some possible pages are `welcome.do` and `incident.do`.

   To specify a Core UI dashboard as the login landing page, set the property to `$dashboards.do?dashboard=<SYS_ID>`. Replace <SYS_ID> with the sys_id of the dashboard.

   To specify a Platform Analytics dashboard as the login landing page, set the property to `$pa_dashboards.do?id=<SYS_ID>`. Replace <SYS_ID> with the sys_id of the Platform Analytics dashboard.

   To direct users to service portal, set the property to `/sp`

   Note: This property is system-wide, so setting it affects all users. To set a login specifically for users with no roles, you can apply these same steps and use the **glide.entry.loggedin.page_ess** property.

   You can also specify a login landing page with the content management system.

Specify lockout for failed login attempts

The system provides inactive script actions that enable you to specify the number of failed login attempts before a user account is locked and to reset the count after a successful login.

Before you begin
Role required: password_reset_admin

Procedure

Navigate to **All** > **System Policy** > **Script Actions** to view or activate the scripts.

Note: Starting with the Kingston release, following a zBoot, the script actions **SNC User Lockout Check with Auto Unlock** and **SNC User Clear** are activated.

To learn more about properties that affect failed login attempts, see Managing failed login attempts (instance security hardening) in the Instance Security Hardening Settings.

Script action	Description
SNC User Lockout Check with Auto Unlock	• Uses the value of the `glide.user.max_unlock_attempts` property to set the limit for failed login attempts. • Unlocks the user account after the time period that is specified for the `glide.user.unlock_timeout_in_mins` property. If no value is specified, then the system unlocks the user account after the default period of 15 minutes.
SNC User Lockout Check	Tracks the number of failed login attempts and locks the user account after a specified number of failed login attempts (default: 5).
SNC User Clear	Updates the user record after a successful login: Resets the number of failed login attempts and updates the date of the last login.

What to do next

Each time a user attempts to log in, the action is recorded in an event log. You can view a log of failed login attempts.

1. Navigate to **System Policy** > **Event Logs**.
2. Filter for **login.failed** in the **Name** field. You can view the attempted login name, date, and IP address logged for the attempt.

Make UI pages public or private

You can make pages public if you want your users to see the pages without logging in.

Before you begin

Role required: admin

About this task

Most pages are only viewable by logged in users. A limited number of pages are public so that users do not have to log in to view them, such as the welcome page, the front page, and the login and logout pages.

Warning: Several base system public pages are required for the functionality of many features. Do not disable base system public pages.

Procedure

1. In the application navigator filter, type `sys_public.list`.
2. Click **New**.
3. In the sys_public table, create a record with the following values.

Field	Description
Page	The name of the page. For example: $sp
Active	When selected, the page is publicly accessible. Deselect the Active option when you want the page to be private.

4. Click **Save**.
   By setting active to true, the page is public, so anyone visiting `<instance_name>/sp` or `<instance_name>/$sp.do` can access the page.

Password complexity requirements

Passwords in your ServiceNow® instance must meet complexity requirements.

Explore

Learn the features and business value of Password complexity requirements.

Enable

Understand how to enable Password complexity requirements.

Configure

Configure the Password complexity requirements.

Reference - Unsupported password characters

Know about the unsupported password characters.

Explore Password complexity requirements

Passwords in your ServiceNow® instance must meet complexity requirements.

The Passwords complexity requirements policy setting determines whether passwords must meet a series of strong-password guidelines.

Password complexity requirements are enforced when a password is changed or created.

Password complexity requirements adhere and work based on the following:

- If the *glide.apply.password_policy.on_login* property is enabled, the password policy check is enforced against the user during login. During login, the user must adhere to the password policy and change the password for the instance.
- The password policy requirements are based on the Basic Multilingual Plane (BMP) that contains characters for all modern languages. ServiceNow instances are shipped with BMPs of around 10,000 characters.
- Passwords within the permissible BMPs can be set for your instance. Passwords that do not adhere within these BMPs are not allowed.

If needed, you can require that passwords are changed regularly, at least every 90 days.

Requirements and banned characters

To enable a secure network environment, it is necessary for users to use strong passwords that include combination of letters, numbers, and symbols. These combinations help to prevent unauthorized users who usually use manual or automated methods to guess weak passwords.

The password to your instance must meet the following requirements:

- Minimum 8 characters.
- Maximum 100 characters.
- Contains lower case and upper case characters
- Contains special characters.
- Contains digits.

You can ban a list of common bad passwords, such as the following:

- Predictable and repeating sequences such as "123456", "qwerty", "!@#$%^", "aaaaa", and so on.
- Employee name or user names.
- Relevant brand or product names.
- Locations, such as company headquarters, city, country, and so on.
- Company-specific internal terms or abbreviations.
- Emojis.

Note: User, brand, or company-specific characters that cannot be used in the password can be configured in the Password Policy or Exclude Password page.

Enable password policies on your instance

Implement password policy controls at login. Force users to change their password if the password does not meet the password policy criteria.

Before you begin

Role required: admin

About this task

The Password Policy plugin (com.glide.password_policy) is enabled by default. The policy goes into effect when a user changes or resets the password.

The **Password Strength Preset** field is automatically set to **Default Strong**. If you want to add new criteria, you can perform the following procedure.

If you customized your instance through the *ValidatePasswordStronger* installation exit or your Password Reset credential store *pwd_cred_store* property, then see password policy properties to know how to implement a password policy for your instance.

Note: The active password policy is highlighted for the Instance as shown.

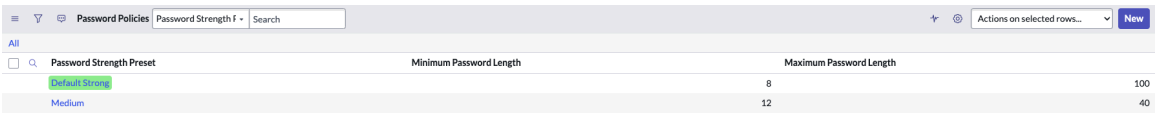

To change the password policy navigate to **All** > **Password Reset** > **Credentials Stores**, select the credentials and change the **Password policy** field to the required policy input.

Procedure

1. Navigate to **All** > **Password Policy** > **Password Policies**.
2. Click **New**.
   The Password Policy form appears.
3. Specify the **Name** for your password policy.
4. In the Password Policy Criteria section, select one of the following presets from the **Password Strength Preset** field.

Password Strength Preset	Description
Default	Auto-populates the fields for required pass word characters as follows: ○ Sets **Minimum Uppercase Character(s)** to 1. ○ Sets **Minimum Lowercase Character(s)** to 1. ○ Sets **Minimum Numeric Character(s)** to 1. ○ Sets **Minimum Special Character(s)** to 0. The minimum password length is 8 charac ters, and the maximum is 100.
Medium	Auto-populates the fields for required pass word characters as follows: ○ Sets **Minimum Uppercase Character(s)** to 1. ○ Sets **Minimum Lowercase Character(s)** to 1. ○ Sets **Minimum Numeric Character(s)** to 1. ○ Sets **Minimum Special Character(s)** to 1. The minimum password length is 12 charac ters, and the maximum is 40.

Password Strength Preset	Description
High	Auto-populates the fields for required pass word characters as follows: ○ Sets **Minimum Uppercase Character(s)** to 1. ○ Sets **Minimum Lowercase Character(s)** to 2. ○ Sets **Minimum Numeric Character(s)** to 1. ○ Sets **Minimum Special Character(s)** to 3. The minimum password length is 8 charac ters, and the maximum is 100.
Default Strong	Auto-populates the fields for required pass word characters as follows: ○ Sets **Minimum Uppercase Character(s)** to 1. ○ Sets **Minimum Lowercase Character(s)** to 1. ○ Sets **Minimum Numeric Character(s)** to 1. ○ Sets **Minimum Special Character(s)** to 1. The minimum password length is 8 charac ters, and the maximum is 100.
Custom	Auto-populates the fields for required pass word characters as follows: ○ Sets **Minimum Uppercase Character(s)** to 1. ○ Sets **Minimum Lowercase Character(s)** to 1. ○ Sets **Minimum Numeric Character(s)** to 1. ○ Sets **Minimum Special Character(s)** to 1. The minimum password length is 8 charac ters, and the maximum is 100. You can also customize the **Password Policy Script** that is displayed.
Advanced	Selecting **Advanced** displays **Password Rule Script** and **Password Strength Script**. Based on your requirement you can customize these scripts.

Note: Password policy is applied based on the selected preset.

5. On the form, fill in the remaining fields.

Password Policy form

Field	Description
Minimum Password Length	Minimum length of the password. This option is displayed for all the presets except for **Advanced**. Set this field to a minimum of 8–10 characters.
Maximum Password Length	Maximum length of the password. This option is displayed for all the presets except **Advanced**. Set this field to a maximum of 100 characters.
Minimum Uppercase Character(s)	Minimum number of uppercase characters in the password, from 0 to 10.
Minimum Lowercase Character(s)	Minimum lowercase characters in the password, from 0 to 10.
Minimum Numeric Character(s)	Minimum numeric of characters in the password, from 0 to 10.
Minimum Special Character(s)	Minimum number of special characters in the password, from 0 to 10.
Included Special Characters	Allow a restricted set of special characters without any delimiter. For example, if you enter `$,!`, then users can only use "$" and "!" as special characters in the password. No other special characters can be used. A password with other special characters is not allowed.
Excluded Special Characters	Allow a restricted set of special characters without any delimiter. For example, if you enter `@$!`, then users cannot use '@', '$' and '!' as special characters in their passwords. **Note:** This option is available if the *glide.password_policy.use_excluded_spec* property is enabled.
Disallow User Data	Option to disallow user data that is authentication-related.
Sequence Length Threshold	The sequence length of your password.
Repetition Length Threshold	The repetition length of your password.

Field	Description
	Note: ○ Both the sequence length threshold and repetition length threshold can have a maximum of eight characters. These fields enable you to restrict weak combinations of passwords that have predictable and repeating sequences such as "123456", "qwerty", "!@#$%^", "aaaaa", and so on. ○ If **Password Strength Preset** is set to **Default Strong**, then the length for both sequence length threshold and repetition length threshold is set to four characters.
Test Your Password	Specify your actual password in this field.

6. Click **Test Your Password**.
7. After the password is tested as valid, click **Submit** to submit the password.

Note: Always test your password before submitting.

Password policy properties

The password policy properties enable you to administrate password policies, exclude list passwords, and apply a password policy during login.

Navigate to **Password Policy** > **Properties** to view and edit the password policy properties.

Property	Description
glide.enable.password_policy	Enables a password policy for your instance. The policy goes into effect when a user changes or resets a password. This property is automatically set to **true**.

Property	Description
	Note: • If your instance is customized, via the *ValidatePasswordStronger* installation exit or your Password Reset credential store [pwd_cred_store], then you must create this property and add it to your system properties. • Prior to the Orlando release, if your instance was customized with the *ValidatePasswordStronger* installation exit, then you had to create the Password Policy property to make the Password Policy work. • Starting with the Orlando release, there is no installation exit customization. The Password Policy properties work by default. These properties can be manually turned off.
glide.enable.blacklist_password	Prohibits using specific passwords. The administrator can insert passwords into the Excluded Password table. This property is automatically set to **true**.
glide.apply.password_policy.on_login	Forces users to change the passwords during their next login if the existing passwords are not in compliance with the current password policy. This property is automatically set to **false**. Setting the value to **true** enforces a password policy during login. **Note:** Enabling this property might force a significant number of users who are not in compliance with the new password policy to change their passwords.
glide.password_policy.user_excluded_special_char	Enables users to use the excluded special character option on the Password policy form.
glide.validate.sys_user.password.field	Enables validation on the user password against the password policy when an admin is editing the sys_user form or list view.
glide.password.policy.generate.password.field.disable	Disables the **Password** field on the set password pop-up on the sys_user form.
glide.user.show.password.field	Enables the **Password** field on the sys_user form.
glide.password_policy.debug	Enables debug logging for the password policy.

Configure your password policy

Password policy criteria enables you to secure your password and adhere to the minimum password complexity requirements.

Before you begin

Role required: admin

About this task

The Password Policy [com.glide.password_policy] plugin is enabled by default. It goes into effect when a user changes or resets the password. If you customized your instance, through the ValidatePasswordStronger installation exit or your Password Reset credential store [pwd_cred_store], see password policy properties.

Procedure

1. Navigate to **All** > **Password Policy** > **Password Policies**.

   Note: **Default Strong** preset in enabled as a default password acceptance criteria. In case, if you want to add a new criteria, you can perform the following steps.

2. Click **New**.

   The Password Policy New record page, has the following sections that you must specify for setting your password:

   - **Password Policy Criteria**
   - **Sequence Matching**
   - **Test Your Password**

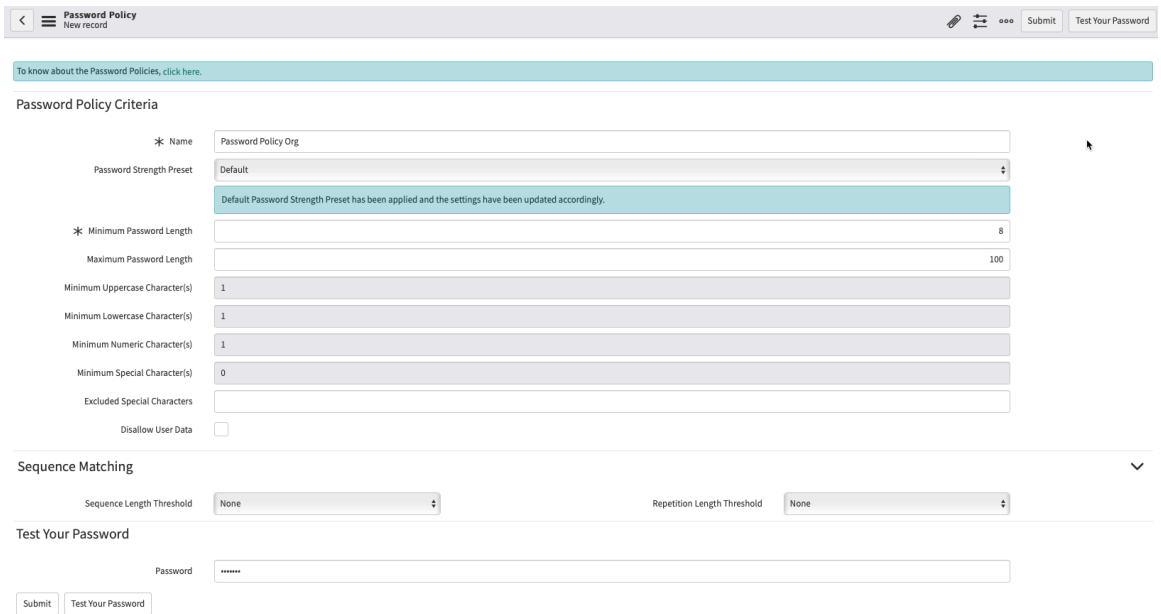

3. Specify the **Name** for your password policy.
4. In the **Password Policy Criteria** section, select the preset from the **Password Strength Preset**. The presets available are as follows:

Password Strength Preset and its description

Password Strength Preset	Description
Default	Selecting **Default** auto-populates the password characters required as follows: ○ One Minimum Uppercase Character ○ One Minimum Lowercase Character ○ One Minimum Numeric Character
Medium	Selecting **Medium** auto-populates the password characters required based on characters as follows: ○ One Minimum Uppercase Character ○ One Minimum Lowercase Character ○ One Minimum Numeric Character ○ One Minimum Special Character
High	Selecting **High** auto-populates the password characters as follows: ○ One Minimum Uppercase Character ○ Two Minimum Lowercase Character ○ One Minimum Numeric Character ○ Three Minimum Special Character
Default Strong	Selecting **Default Strong** auto-populates the password characters required based on characters as follows: ○ One Minimum Uppercase Character ○ One Minimum Lowercase Character ○ One Minimum Numeric Character ○ One Minimum Special Character
Custom	Selecting **Custom** auto-populates the password characters required based on characters as follows: ○ One Minimum Uppercase Character ○ One Minimum Lowercase Character ○ One Minimum Numeric Character ○ One Minimum Special Character You can also customize the **Password Policy Script** that is displayed.
Advanced	Selecting **Advanced** displays **Password Rule Script** and **Password Strength Script**. Based on your requirement you can customize these scripts.

Note: Password policy is applied based on the selected preset.

5. Specify the fields described in the table:

Password Policy form

Field	Description
Minimum Password Length	Minimum length of the password. This option is displayed for all the presets except **Advanced**. **Note:** **Minimum Password Length** is a required field and recommended setting it to a minimum of 8–10 characters.
Maximum Password Length	Maximum length of the password. This option is displayed for all the presets except **Advanced**. **Note:** **Maximum Password Length** is an optional field and recommended to set it to a maximum 100 characters.
Minimum Uppercase Character(s)	Minimum number of uppercase characters in the password, from 0 to 10.
Minimum Lowercase Character(s)	Minimum lowercase characters in the password, from 0 to 10.
Minimum Numeric Character(s)	Minimum numeric of characters in the password, from 0 to 10.
Minimum Special Character(s)	Minimum number of special characters in the password, from 0 to 10.
Included Special Characters	Allow a restricted set of special characters without any delimiter. For example, if you enter "$,!" users can only use "$" and "!" as special characters in the password. No other special characters can be used, and a password with other special characters is not allowed.
Excluded Special Characters	Allow a restricted set of special characters without any delimiter. For example, when '@ $!' is entered, users should not be able to use '@', '$' and '!' as special characters in their passwords. **Note:** This option is available if the glide.password_policy.use_excluded_special_char property is enabled.
Disallow User Data	It is enabled to disallow the user data.

6. In the Sequence Matching section, specify the **Sequence Length Threshold** and **Repetition Length Threshold**.

   Note:

   - Both the sequence length threshold and repetition length threshold can have a maximum of eight characters. These fields enable you to restrict weak combinations of passwords that have predictable and repeating sequences such as "123456", "qwerty", "!@#$%^", "aaaaa", and so on.
   - For **Default Strong**, the length for both sequence length threshold and repetition length threshold is selected as four characters.

7. In the **Test Your Password** section, specify your password.
8. Click **Test Your Password**.
9. Once the password is valid, click **Submit** to submit the password.

   Note: Always test your password before submitting.

Configure password for a user

Set your user's password for the instance based on the password policy that is configured.

Before you begin

Users created for setting the password for their first login. For more information, see Create a user.

To populate the **Password** field directly on the user form, enable the **Enable to show the password field on the sys_user Form** (glide.user.show.password.field). To know more about the properties, see Password policy properties.

Role required: admin

Procedure

1. Navigate to **All** > **User Administration** > **Users**.
2. Select the user from the list in the Users page.
3. To set the password based on the password policy, click **Set Password**.

   The Set Password pop-up is displayed.

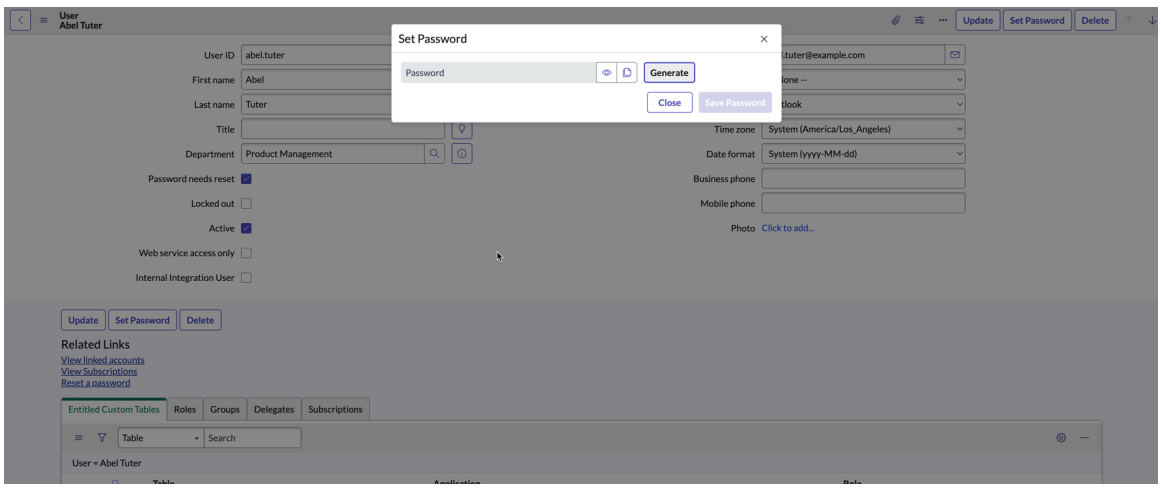

4. In the Set Password, perform the following.

a. Click **Generate** to generate the password.

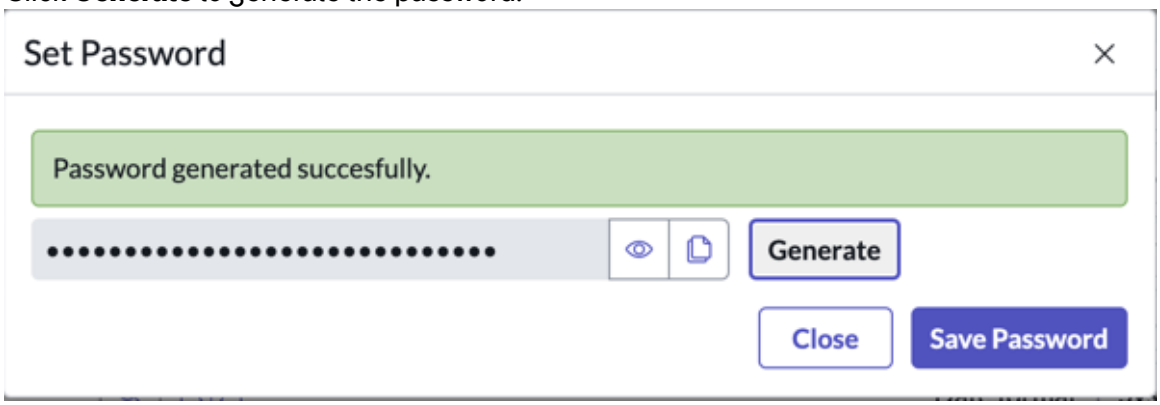

b. Click the **View** icon to view the password.

c. Click the **Copy** icon to copy the password to share to the user.

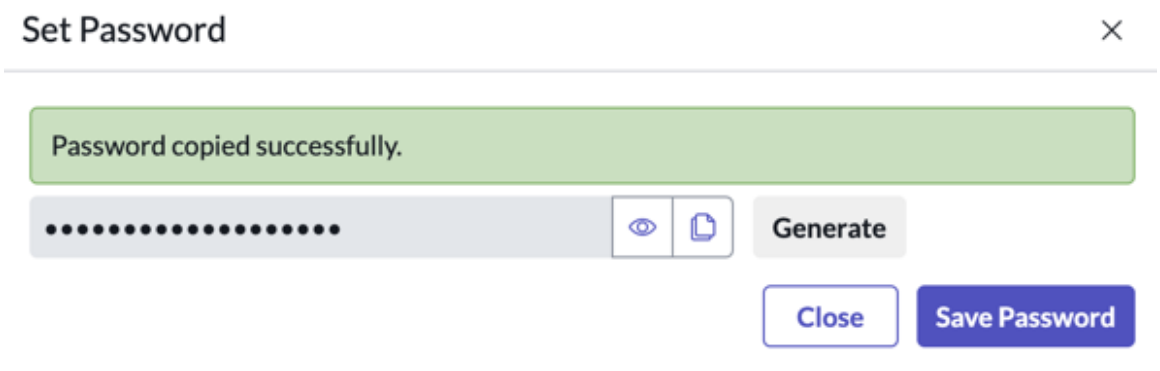

5. Click **Save Password**.

The password is set for the user. Also the Password needs rest check-box is automatically enabled.

The user during the first login must use the same password for login and change the password upon login adhering to the password policy configured by the administrators.

Exclude passwords through password policies on your instance

Add passwords to the Excluded Password table to prohibit specific passwords from being used by users on your instance.

Before you begin
Role required: admin

About this task

You can ban as many passwords as necessary. Some of the examples are as follows:

- Predictable and repeating sequences, such as "123456", "qwerty", "!@#$%^", "aaaaa", and so on.
- Employee name or user names.
- Relevant brand or product names.
- Locations, such as a company headquarters, city, country, and so on.
- Company-specific internal terms or abbreviations.
- Emojis.

Procedure

1. Navigate to **All** > **Password Policy** > **Excluded Passwords**.
2. To add a password to your exclude list, click **New** and enter the password.
3. Navigate to **Password Policy** > **Exclusion List Management** to govern your users' password and prevent users from using bad password for the instance.

   The **Exclusion List Management** has about 5000 commonly used passwords.

Unsupported password characters

There are password characters that are not supported. Users cannot use these characters, based on ServiceNow password complexity requirements.

To enable a secure network environment, it is necessary for users to use strong passwords that include combination of letters, numbers, and symbols. These combinations help to prevent unauthorized users who usually use manual or automated methods to guess weak passwords.

- The password policy requirements are based on the Basic Multilingual Plane (BMP) that contains characters for all modern languages. ServiceNow instance is shipped with BMPs of around 10000 characters.
- Password characters within this permissible BMPs can be set for your instance, the password characters that do not adhere within this BMPs are not allowed.

Note: User or company-specific characters that cannot be used in the password can be configured in the exclude password list.

For more information about the password policy properties, see password policy properties.

Password Reset

The default self-service Password Reset process enables a user to reset the password without assistance from service desk agents.

Example: The default self-service password reset flow

1. If a user does not remember the password, the user can click the **Forgot Password?** link on the login screen.
2. The Password Reset application starts. On the **Identity** page, the user identifies himself or herself by entering a **Username**.
3. On the **Verify** page, the user proves that they are the person who is associated with the username. In this example, the user enters the email address that is associated with the user

profile. The admin can configure a different verification method or can require additional verifications, for example, a personal question that only the user can answer.

4. The **Reset** page tells the user to check email for instructions.
5. The user opens the email and clicks the **here** link to reset the password. The link is valid for a period that you specify (use the **password_reset.request.expiry** property).
6. The **Reset Password** page guides the user to reset the password.

The default self-service Password Reset process (`com.glideapp.password_reset`) defines:

- The URI that specifies where users are redirected when they click **Forgot Password?**. By default this value is **/$pwd_reset.do?sysparm_url=ss_default**, which is the same value used in the *glide.security.password_reset.uri* property. In previous releases, this value was set to **/reset_password.do**.
- The **Enable Password Reset URL** option, which specifies that the user should receive an email with a link to reset their password after they click **Forgot Password?**.
- The Personal Data - Enter Email Address verification flow that specifies the three-step password reset flow.

See Configure your Password Reset process for instructions on accessing this form and configuring the fields.

Note:

- This feature works for locally authenticated users who enter the username and password specified in their user record. Users logging in to the instance via an SSO solution or an LDAP integration cannot reset passwords using the example self-service Password Reset process.
- The end user must enable and configure notification preferences. See Subscription-based notifications. Administrators can xmodify the email that is sent to the end user.

Modify the Password Reset notification email text

Users of the self-service Password Reset process receive an email notification when they request password reset. You can modify the text of the email and other aspects of the notification.

Before you begin

Role required: password_reset_admin

About this task

This process is relevant only if users are Subscription-based notifications.

Procedure

1. Navigate to **All** > **System Notification** > **Notifications**.
2. Select the **Password Reset - Reset Link** notification.
3. Modify the text of the email in the **What it will contain** section.
   For information on configuring other aspects of the notification, see Create an email notification.

Configure Password Reset properties

You can specify properties that configure the Password Reset experience for end users.

Before you begin
Role required: password_reset_admin

About this task
While there are no range limits for the values you can enter for properties, consider using only positive integer values starting at 1. When you determine the limit for the upper range of a property, consider the task that the user is performing.

For example, you would not want to allow 100 attempts for users to verify their identity. A more common value is 3 attempts. Similarly, you may not want to force users who are completing the enrollment process to spend time selecting and answering 30 security questions. The more commonly used number of security questions is between 5 and 7.

Procedure
1. Navigate to **All** > **Password Reset** > **Properties**.
   For information about the Password Reset properties, see Password Reset global properties.
2. Update settings as needed and then click **Save**.

Remember me

When the **Remember me** check box is selected at login, a cookie is stored on the user's computer. This cookie automatically authenticates the user upon subsequent visits.

If the user logs out, the cookie is destroyed. The default value of the **Remember me** check box is controlled by one property, and whether or not the check box appears on the login page is controlled by a different property.

Two properties, *glide.ui.user_cookie.life_span_in_days* and *glide.ui.user_cookie.max_life_span_in_days* control the glide_user system generated cookie expiration value. When a user accesses an instance with 'remember me' enabled, the access resets the cookie expiration period until the maximum (glide.ui.user_cookie.max_life_span_in_days) life span limit is reached.

Note: To learn more about these properties, see the following topics in Instance Security Hardening Settings:

- Minimize absolute session timeout duration [Updated in Security Center 1.3]
- Minimize session window timeout duration [Updated in Security Center 1.3]

Change the default value of the Remember me check box

You can change the default value of the **Remember me** check box.

Before you begin
Role required: admin

Procedure
1. Navigate to **All** > **System Properties** > **UI Properties**.
2. Locate the *Default value of "Remember me" checkbox on login page* property (*glide.ui.remember.me.default*).
3. To set the default value of the **Remember me** check box to **No**, clear the property check box.
4. To restore the default value of the **Remember me** check box to **Yes**, select the property check box.

Remove the Remember me check box

You can remove the **Remember me** check box so users do not have access to this feature.

Before you begin

Role required: security_admin

Note: To learn more about this property, see Remove remember me in Instance Security Hardening Settings.

Procedure

1. Elevate your role to security_admin.
2. Navigate to **System Properties** > **UI Properties**.
3. Locate the *Remove "Remember me" checkbox from login page* property (*glide.ui.forgetme*).
4. Select the property check box.
   This setting removes the **Remember me** check box, invalidates existing cookies, and disables Remember me functionality entirely.
5. To restore the **Remember me** check box to the login page, clear the property check box.

Configure the logout confirmation prompt

You can enable a logout confirmation prompt to prevent users from inadvertently logging themselves out.

Before you begin

Role required: admin

About this task

Note: The following procedure only works in UI versions earlier than Core UI, which is the most recent and most commonly used one.

Procedure

1. Navigate to **All** > **System Properties** > **System**.
2. Locate the **Prompt user to confirm a logout request** property and select the check box.
3. When the user clicks the **Logout** button, a confirmation dialog box displays.

Implement a nonce

You can implement a nonce to be used with single sign-on digest authentication.

To use a nonce with the unencrypted token or encrypted token methods of single sign on, these steps apply with only a few minor changes.

Note: The nonce is used only for login requests, not for any other type of request. If the system receives a nonce value after login, the nonce is not consumed.

Benefits

The usage of a nonce prohibits a malicious user from performing a replay attack in order to log into your system.

Nonce process flow

When a customer has implemented the digested token Single Sign-on and wishes to add the security of a nonce, they follow a certain process flow.

1. A user logs into the customer's portal.
2. The customer generates the required SSO parameters and appends a random nonce to the end. For example, if the customer were forwarding the authentication response via the query string, it may look something like this:

```
SM_USER=itil&DE_USER=V1QuWMmxSfBgfRS099XOcAjKo5Q=&NONCE=1407743018
```

The instance receives this request and retrieves the authentication variables. Before attempting to verify the integrity of the authentication response, the instance checks the nonce against an internal table (u_authentication_nonce) to verify that it does not yet exist. If the nonce does not exist within that table, the nonce is then added to the table and the authentication process is allowed to continue. However, if that nonce value already exists within the table, the authentication attempt is cancelled and an error code of `failed_missing_requirement` is returned, which typically takes the user back to the login page.

Implement a nonce

Add a cryptographic nonce to the authentication header to ensure that it can only be used once.

- Create a system property called *`glide.authenticate.header.nonce_key`* and set its value to whatever variable name you're using for the nonce, such as NONCE or NCE.
- Create a new table called `u_authentication_nonce`. Add a field to the table called `u_nonce`.
- Go to **System Properties** > **Installation Exits** and create an item called `DigestSingleSignOnNonce` which overrides ExternalAuthentication (*`glide.authenticate.external_property`*).
- Add the following code to the script portion of the newly created DigestSingleSignOnNonce.

```
gs.include("PrototypeServer");

var DigestSingleSignOnNonce = Class.create();
DigestSingleSignOnNonce.prototype = {

 process : function() {

  var headerKey =
 GlideProperties.get("glide.authenticate.header.key",
 "SM_USER");
  var headerDigestKey =
 GlideProperties.get("glide.authenticate.header.encrypted_key",
 "DIGEST");
  var headerNonceKey =
 GlideProperties.get("glide.authenticate.header.nonce_key",
 "NCE");
  var fieldName =
 GlideProperties.get("glide.authenticate.header.value",
 "user_name");
  var fkey =
 GlideProperties.get("glide.authenticate.secret_key");
```

```
 // Look in the Headers
 var data = request.getHeader(headerKey);
 var encdata = request.getHeader(headerDigestKey);
 var nonce = request.getHeader(headerNonceKey);

 // If not, then check the URL Parameters
 if (data == null || encdata == null || nonce == null) {
  data = request.getParameter(headerKey);
  encdata = request.getParameter(headerDigestKey);
  nonce = request.getParameter(headerNonceKey);
 }

 // then maybe its a cookie
 if (data == null || encdata == null || nonce == null) {
  var cookies = request.getCookies();
  data = GlideCookieMan.getCookieValue(cookies, headerKey);
  encdata = GlideCookieMan.getCookieValue(cookies,
headerDigestKey);
  nonce = GlideCookieMan.getCookieValue(cookies,
headerNonceKey);
 }

 // if found run encryption
 if (data != null && encdata != null && nonce != null) {
  try {

   // Replace all spaces with plus(+)'s, converted in url
   encdata = encdata.replaceAll(' ', '+');

   // ----- Encrypt the username|nonce
   var key = this.getDigest( data + "|" + nonce, fkey);

   // Check for match of received encoded data
   // and your encoding of user name
   if (encdata == key) {
    var ugr = new GlideRecord("sys_user");
    ugr.initialize();
    if (!ugr.isValidField(fieldName)) {
     GlideLog.warn("External authorization is set to use
field: '"+ fieldName + "' which doesn't exist");
     return "failed_missing_requirement";
    }
    ugr.addQuery(fieldName, data);
    ugr.query();
    if (!ugr.next()) {
     var userLoad = GlideUser.getUser(data);
     if (userLoad == null)
      return "failed_authentication";

     ugr.initialize();
     ugr.addQuery(fieldName, data);
     ugr.query();
     if (!ugr.next())
      return "failed_authentication";

    }
```

```
    if (this.processNonce(nonce)){
     var userName = ugr.getValue("user_name");
     return userName;
    }
    else return "failed_missing_requirement";
    }
   else {

    return "failed_authentication";
   }
  } catch(e) {
   gs.log(e);
   return "failed_authentication";
  }
  // Encoded data didn't match recieved Encoded data
 } else {

  return "failed_missing_requirement";
 }
},

getDigest : function( data, fkey ) {
 try {
  // default to something JDK 1.4 has
  var MAC_ALG = "HmacSHA1";
  return SncAuthentication.encode(data, fkey, MAC_ALG);

 } catch (e) {
  gs.log(e.toString());
  throw 'failed_missing_requirement';
 }
} ,

processNonce : function( sentNonce ) {
 var ngr = new GlideRecord("u_authentication_nonce");

 ngr.addQuery("u_nonce", sentNonce);
 ngr.query();
 if (ngr.next()) {
  gs.log("This SSO entry has already been processed! (Nonce: "
+ sentNonce + ")");
  return false;
 }
 var ngrNew = new GlideRecord("u_authentication_nonce");
 ngrNew.initialize();
 ngrNew.u_nonce = sentNonce;
 ngrNew.insert();
 gs.log("Inserted new nonce: " + sentNonce);
 return true;
}
};
```

- Once you've saved your new installation exit, go to the DigestSingleSignOn installation exit and make sure that it is set Active=false.

Your instance should now be configured to implement a nonce.

Multi-factor authentication

Learn how to activate, use, and configure Multi-factor authentication (MFA).

Explore  Learn the features and business value of MFA.	Configure  Understand how to Activate MFA.
Use  Use MFA.	MFA with Single Sign-on (SSO)  Know about how to configure MFA with Single Sign-on (SSO).

Multi-factor Authentication enforcement

Enforcement of MFA for non-SSO logins to ServiceNow from the Yokohama release.

MFA enhances security by requiring users to provide two or more evidence to prove their identity during login. MFA protects accounts from threats like phishing and account takeovers.

MFA being a critical security tool against various identity takeover-related attacks. ServiceNow enforces MFA by default post-Yokohama upgrade and making it mandatory for non-SSO log ins (users performing login with only username and password or LDAP based authentication) to ensure a better security posture and reduce the risk of breaches.

MFA enforcement is carried though a MFA policy that is activated by default from Yokohama or upgrade to Yokohama. The MFA is enforced the **Enforce MFA for non-SSO logins** policy is `Active` and honored through the MFA Context.

Note:

- The policy is enabled for all **non-snc_external** users performing local (username and password) or LDAP based authentication.
- The instance admin can modify the enforcement scope by changing the MFA context policy, policy criteria, or policy conditions.

Here's more details about MFA enforcement:

- Enforced to all production and non-production instances.
- Enforced to all the **non-snc_external** users and **non-SSO** login.
- Integration with Basic auth and OAuth resource owner password credential grant does not require MFA from Yokohama.

To know more about the changes due to enforcement, see Changes due to the Multi-factor Authentication enforcement.

To know more about MFA, see Exploring Multi-factor Authentication.

Related topics

Configure Multi-factor authentication

Using Multi-factor authentication

Multi-factor Authentication enforcement properties

Troubleshooting Multi-factor Authentication enforcement

Frequently asked questions - Multi-factor Authentication enforcement

Changes due to the Multi-factor Authentication enforcement

Information about the changes that are expected due to the MFA enforcement.

MFA enforcement changes that are expected as follows:

- No impact on integration.
- Users who have already enrolled MFA continue to be challenged with MFA during local login.
- If a user has not set up MFA yet, after the Yokohama upgrade, they'll be initially exempted from undergoing MFA for a certain period during which they can enroll MFA.
- If the user fails to self-enroll, then they have MFA enforced on the completion of the exemption period (default 30 days) and must set up MFA.
- After 90 Days, MFA will be enforced by default. There will be no self-enrollment period for first-time logins after 90. Administrators can configure this period.

As an administrator, you can prepare for the MFA enforcement based on the following:

- Review the MFA Context Policy configuration and adjust the policy conditions according to your business requirements.
- Add the exempted group provided by default, if there are any other users who must be exempted.
- Review the MFA Enforcement Properties and adjust based on your requirements. To learn more, see Multi-factor Authentication enforcement properties.

Example: MFA enforcement scenario after 30 days of the MFA enforcement

- **Scenario 1 for Acme Corp**: In the instance, if the user doesn't have an active MFA policy. Imagine Sarah uses local authentication to access an instance. On upgrading to the Yokohama release during login, a message to enroll in MFA is displayed. There's 30 days to complete this setup. If Sarah doesn't complete the setup, after 30 days, the account requires MFA to log in, and won't be able to access it until MFA is set up.
- **Scenario 2 for Acme Corp**: In the instance, Anita was already MFA along with local authentication. Anita continues to require MFA without the 30-day self-enrollment window.

- **Scenario 3 for Acme Corp**: In the instance, Olivia uses Single Sign-On (SSO) for authentication. There's no impact on the login experience, and Olivia isn't enforced for MFA.
- **Scenario 4 for Globex Corp**: In the instance, if a user already had an MFA policy requiring - MFA for all local login attempts outside the company's trusted network. On upgrading to the Yokohama or a later release, MFA enforcement behavior for user logins remain the same.

Multi-factor Authentication enforcement properties

Configure the properties for MFA enforcement.

To set the MFA enforcement properties, navigate to **sys_properties.list** to validate or change the MFA enforcement properties.

MFA enforcement properties

Field	Description
glide.authenticate.multifactor.setup.bypass.count	Number of times a user can bypass setting up multi-factor authentication (Max possible bypass count is 3, beyond that it will treated as 3).
glide.authenticate.multifactor.self_enroll_days	Indicates number of days a user will be given an option to self enroll for MFA, post which user will be automatically challenged with MFA. Any value exceeding 90 will be treated as 90 days only.
glide.authenticate.multifactor.enforce.max_enroll_days	Indicates the maximum number of days a new user will be given an option to self-enroll for MFA, post which all new user's performing non-SSO logins are automatically challenged with MFA. Any value exceeding 270 is treated as 270 days only.
glide.authenticate.multifactor.enforce.debug	MFA Enforcement debug logger. Helps in debugging the flows for MFA logins. **Note:** This property doesn't exist on the instance, you have to create and enable the property, if needed.
glide.authenticate.hybrid_user_tracking.enabled	Property for enabling of hybrid users. The user accounts which are not marked as 'Web service access only' in corresponding sys_user record, but still performs integrations (For example, API logins) using the username and password, will be tracked in the 'User Login Info' table when this property is enabled. **Note:** This property doesn't exist on the instance, you have to create and enable the property, if needed.
glide.authenticate.hybrid_user_tracking.debug	Adds debug logging for the tracking of hybrid user API logins.Side

MFA enforcement properties (continued)

Field	Description
	Note: This property doesn't exist on the instance, you have to create and enable the property, if needed.

Troubleshooting Multi-factor Authentication enforcement

Troubleshooting information due to the MFA enforcement.

ServiceNow enforces MFA by default post-Yokohama upgrade and making it mandatory for non-SSO logins (users performing login with only username and password or LDAP based authentication) to ensure a better security posture and reduce the risk of breaches.

MFA enforcement is carried though a MFA policy that is activated by default from Yokohama or upgrade to Yokohama. Following are some of the troubleshooting task that you can perform if there's any change to the MFA behavior:

- Debug using the troubleshooting tools
- Navigate to the Log location and Debug properties
- Understand the MFA scenarios based on your users experience while using MFA
- Understand the MFA issue due to upgrade from a previous release

Debug MFA

Use the either of the following tools or a combination to understand the debug information:

- **Splunk** - to see the debug logs.
- System logs or Node logs.
- **HAR** logs to analyze the debug logs for the MFA.

Log location and Debug properties

Navigate to the following location to know more about logs:

- For system logs, navigate to **All** > **System Log** > **System Logs**.
- For node logs, navigate to **All** > **System Logs** > **Utilities** > **Node Log File Browser**.

The system debug logs and instance node logs are required for the debug purpose. Following are the debug properties that are required to be enabled:

- `glide.webauthn.debug.enabled`
- `glide.log.default_log_debug`
- `glide.authenticate.policy.debug`
- `glide.authenticate.hybrid_user_tracking.debug`

MFA issue based on scenarios

Scenario 1: User is not able to login using any of their second factor

Reset the MFA for the your users and delete the old user records from the following tables:

- `user_multifactor_auth`
- `sys_user_public_credential`
- `sys_user_multi_factor_setup`

Scenario 2: Admin is not able to login using any of their second factor

Another user with admin access can reset the MFA for any blocked admin user. If still the issue exist, reach out to ServiceNow Support.

Scenario3: Error observed during the MFA Setup or Validation

Check the warning "Associated Error Codes/Warning: Your 6-digit verification code is incorrect. Try again with the correct code".

Perform the following steps:

- In case of TOTP Authenticator App, if the date and time of the Authenticator MFA device and instance are not in sync (±30 sec), then the TOTP code is not accepted. Verify the device and instance date and time.
- In case of email, configure the user level notification, outbound email configuration, and user correctly in the `sys_user` table.
- In case of SMS, configure the Twillio or other SMS service provider integration correctly and set to active. Verify if the user's mobile number is configured correctly in the `sys_user` table.

Frequently asked questions - Multi-factor Authentication enforcement

Details about some of the FAQs due the MFA enforcement.

MFA being a critical security tool against various identity takeover-related attacks. ServiceNow enforces MFA by default post-Yokohama upgrade and making it mandatory for non-SSO logins (users performing login with only user name and password or LDAP based authentication).

MFA enforcement is carried though a MFA policy that is activated by default from Yokohama or upgrade to Yokohama. You can check the following links to learn more about MFA enforcement related questions:

- MFA enforcement exception
- MFA enforcement requirements – What and Why
- MFA enforcement scope
- MFA enforcement timeline
- MFA metrics
- MFA types
- MFA reset

MFA enforcement requirements – What and Why

FAQ related to MFA enforcement and why it's important.

1. What is the MFA?

Multi-factor Authentication (MFA) is a security process that requires you to provide two or more forms of verification before they can access an account or system. To learn more, see Exploring Multi-factor Authentication.

2. Why is the MFA enforcement mandate?

   MFA is mandated to protect your account and data security. Cyberthreats are ever-changing, and passwords alone no longer provide sufficient protection against unauthorized access.

   - With MFA enabled, even if attackers have your password, the attackers still need a second form of verification. This additional layer significantly blocks most unauthorized attempts, keeping your information more secure.
   - Setting MFA as the default, minimize the risk of security breaches and safeguarding your account automatically. This means you get enhanced peace of mind without having to make any extra security decisions.

3. Why is it important to enable MFA?

   Enabling MFA boosts your account security. Passwords alone aren't enough because passwords can be exposed in data breaches. With MFA, even if someone knows your password, they can't access your account without a second verification step.

4. Why does ServiceNow require MFA?

   ServiceNow is mandating MFA to protect you from these threats. It's a simple yet effective way to reduce unauthorized access. By requiring MFA, there's a strong layer of protection to every account, reducing security risks for you and all users.

5. What is the MFA requirement for existing customers?

   For existing customers upgrading their instance to the Yokohama or a later release:

   - If the instance doesn't already have the **Adaptive Authentication** – Multi-factor Authentication context turned on, automatically it's enabled as a default MFA policy.
   - All the internal users (users who don't have snc_external role) logging in with local or LDAP authentication must set up MFA within 30 days of their first successful login. During this time, you can log in normally but see a message at the time of login to enroll in MFA.
   - After 30 days, MFA will be required by default, and users won't be able to log in without completing the MFA setup.

6. What is the MFA requirement for new customers?

   For any instance using the Yokohama release or later, MFA is enabled by default for all internal users. It also applies to users who don't have the `snc_external` role and are logging in with local or LDAP authentication. From the first login, the users are required to set up and use MFA.

MFA enforcement scope

FAQ related to MFA enforcement scope and why it's important.

1. Which user, login, and environment types require MFA?

   From the Yokohama release onwards, with the new default secure MFA policy MFA enforced for the following scenarios:

   - All the users except the users having **snc_external** role.
   - All the users performing user name and password based local or Lightweight Directory Access Protocol (LDAP) authentication.
   - All customer instances, including production, subprod, and test instances, that didn't already have an active MFA policy before the upgrade.

The instance admin can modify the enforcement scope by changing the MFA context policy, policy criteria, or policy conditions.

2. Is MFA required for Single-Sign-On (SSO) logins?

   No. With the default secure MFA policy, MFA isn't required for SSO (SAML, OIDC, Certificate Based Authentication) login.

   Customers can collaborate with their Single Sign-On (SSO) provider (Identity Provider, or IdP) to enforce multi-factor Authentication (MFA) on the IdP side. If enforcing MFA on the IdP side isn't feasible, customers also have the option to enable the ServiceNow platform's MFA for SSO logins by following the instructions provided in Multi-factor Authentication with Single Sign-On.

3. Is MFA required for external users?

   No. With the default secure MFA policy, MFA isn't required for users having the snc_external role.

   - Admins can modify this behavior and enforce MFA for external users by updating the MFA policy conditions.
   - External users who were already undergoing MFA before the upgrade to Yokohama or later release continues to have MFA.
   - External users can visit their profile and self-enroll for MFA.

4. Is MFA required for Mobile App login?

   Yes. The MFA policy is applied to both web and mobile app log in with user name and password based non-SSO login.

5. Is MFA required for non-production and test environments?

   Yes. MFA is enforced for all customer instances, including production, non-production, dev, and test environments, if there's no active MFA policy existed on the instance before upgrading to Yokohama or later versions.

6. Is MFA required for developer instances?

   Yes. MFA is enforced for all developer instances that are on Yokohama or later release versions.

7. Is MFA required for API authentication?

   No. From Yokohama or a later release, MFA is only required for the user name and password-based interactive user logins. This means API authentication with basic auth works without requiring MFA. It's recommended to customers use alternative secure API authentication methods such as OAuth or mTLS. More details here.

   a. Clone

   b. Update set retrieval

   c. RPA

   Note: To enforce MFA for API authentication, set the `glide.authenticate.multifactor.for_integrations` system property to `true`. MFA is enforced only for users who have already enrolled in MFA. Users who have not enrolled are not affected.

8. Is there any impact on the clone setup process due to MFA?

   No, the clone setup process continues to work with user name and password and doesn't require MFA.

9. Is there any impact on the update set retrieval due to MFA?

   No, the update set retrieval continues to work with user name and password and doesn't require MFA.

10. Is there been any impact on RPA bots accessing ServiceNow instances?

    Yes, if the RPA bot uses the interactive user name and password login to access the ServiceNow instance, it must perform MFA. Admins can add RPA bot accounts to the **MFA Exempted User Group** if they want to relax MFA for RPA bot accounts.

11. Is MFA required for the OAuth-based integrations?

    The OAuth Resource owner password credential (ROPC) works with user name and password without requiring MFA. For Authorization code grant type MFA is required as part of the user login flow before giving the OAuth consent.

MFA enforcement timeline

FAQ related to MFA enforcement timelines and why it's important.

1. When is MFA enforced?

   According to the MFA policy, eligible users who haven't completed the MFA setup has a 30-day self-enrollment period. The behavior is controlled using the system property `glide.authenticate.multifactor.self_enrolment_period`. The property's default value is 30 days. It can be updated to a maximum of 90 days.

   All internal users (users who don't have a `snc_external` role) logging in with local or LDAP authentication must set up MFA within 30 days of their first successful login. During this time, you can log in normally but see a message at the time of login to enroll in MFA.

   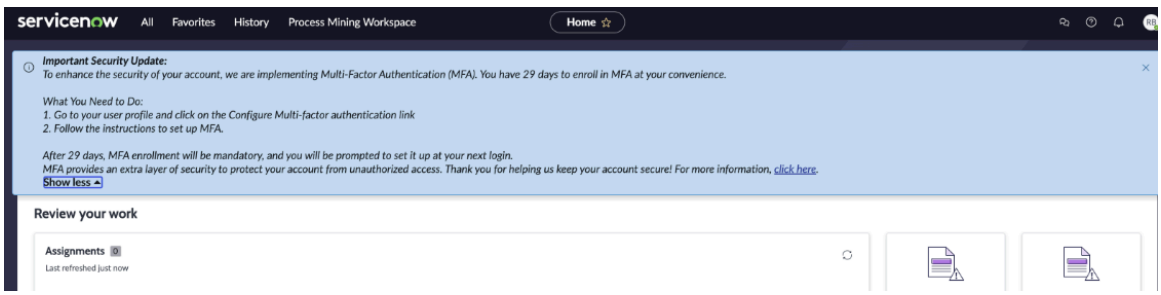

   After 90 days of upgrading to Yokohama or a later release, if an internal user (user without the `snc_external` role) logs in with local or LDAP authentication for the first time, they'll be required to use MFA immediately. You don't have the 30-day MFA self-enrollment window. This period is governed by a system property: `glide.authenticate.multifactor.enforcement.max_relaxation_period`. The maximum value for this property is 270 days.

2. How can the MFA enforcement timeline adjusted?

   - By updating the value of the property `glide.authenticate.multifactor.self_enrolment_period`, admins can provide a smaller or larger self-enrollment window. Set the property value to 0. The users are required to complete the MFA setup after their first login attempt with local or LDAP login

after upgrading to Yokohama or a later release. The maximum duration of the self-enrollment window can be 90 days. Property value set higher than 90 will be treated as 90.

- By updating the value of the property `glide.authenticate.multifactor.enforcement.max_relaxation_period` admin can decide how many days post upgrade to the Yokohama or a later release you get the MFA self-enrollment window.

3. How are end users informed about this upcoming change?

   End users performing local or LDAP authentication who will be enforced with MFA will see an information message after logging in. The same message is available when users visit their profile.

   Enforcement Message

On the User Profile	On Employee Service Center
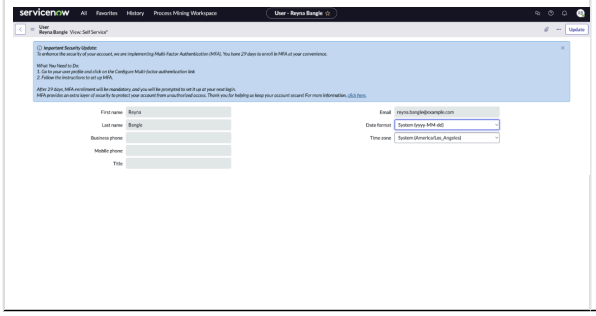	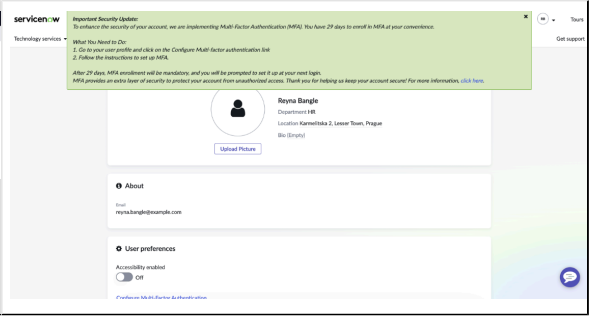

   This message won't appear for non-admin users performing SSO logins. The admin role will see a different information message after a successful login irrespective of the authentication method used for logging in.

   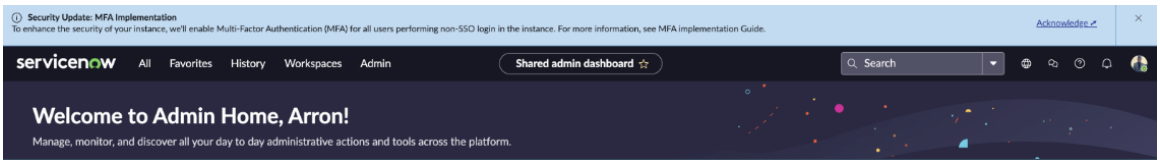

   This message continues to be displayed until one of the admins acknowledges the update by setting the `glide.authenticate.multifactor.enforcement.acknowledged` property value to true.

   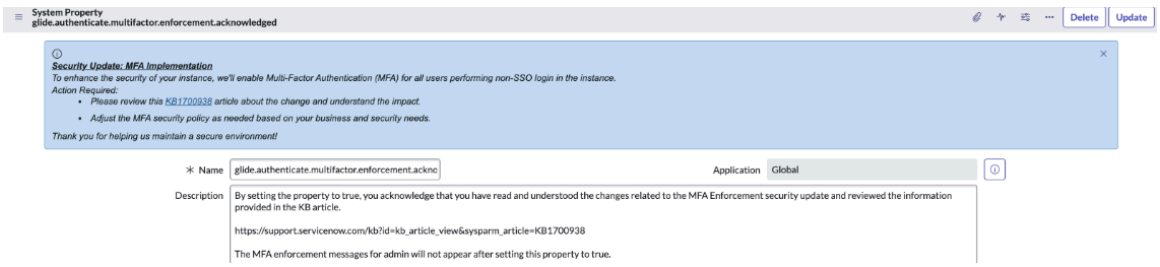

4. How to turn off the message displayed to end users about completing the MFA setup when they log in?

   Admins can update the value of the `glide.authenticate.multifactor.enforcement.show_user_info_message` system property to false to turn off the MFA enrollment information message shown to end users after login.

5. How to turn off the message displayed to administrators about the MFA enforcement?

The information message regarding MFA enforcement shown to admin users after login, stops appearing when one of the admins acknowledges it by updating the value of the `glide.authenticate.multifactor.enforcement.acknowledged` system property to true.

6. There's already an MFA policy defined using adaptive authentication based on the security needs of my organization in the instance. Is the policy impacted by the mandate?

   No, if the instance already has an active Adaptive authentication—MFA context policy, the new default secure MFA policy isn't enforced. If the instance had MFA property enabled (`glide.authenticate.multifactor`) but the MFA policy wasn't active, then the default secure MFA policy for enforcing MFA for all internal users (users who don't have `snc_external` role) performing user name and password-based local or LDAP login is enabled.

MFA enforcement exception

FAQ related to MFA enforcement exception and why it's important.

1. How can the MFA mandate be relaxed for specific users?

   In the Yokohama release, a new user group, MFA Exempted User Group record is added. Based on the default condition, there's an MFA policy added, any user who is a member of this group is enforced with MFA.

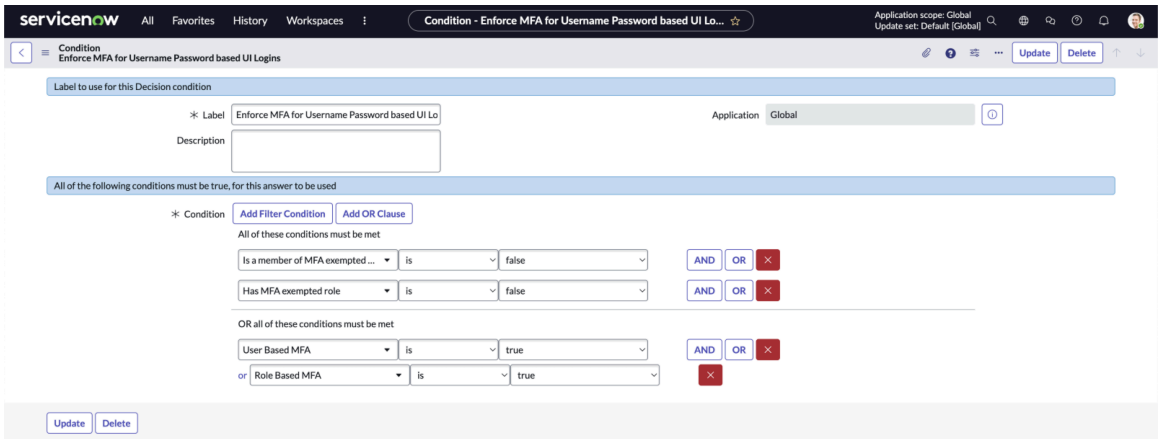

   To relax MFA for specific users, follow the procedure:

- Navigate to **MFA context**. The Step-Up MFA Policy associated with the MFA context record should be “Enforce MFA for non-SSO logins.

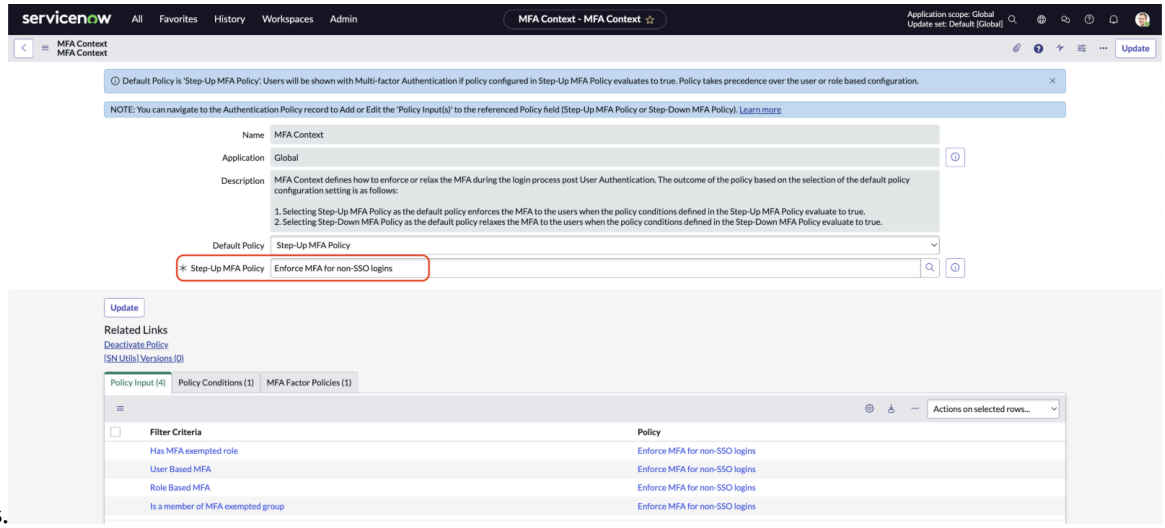

- Under the **Policy Input** related list, select the **Is a member of MFA exempted group** filter criteria record.
- Select **MFA Exempted User Group**.

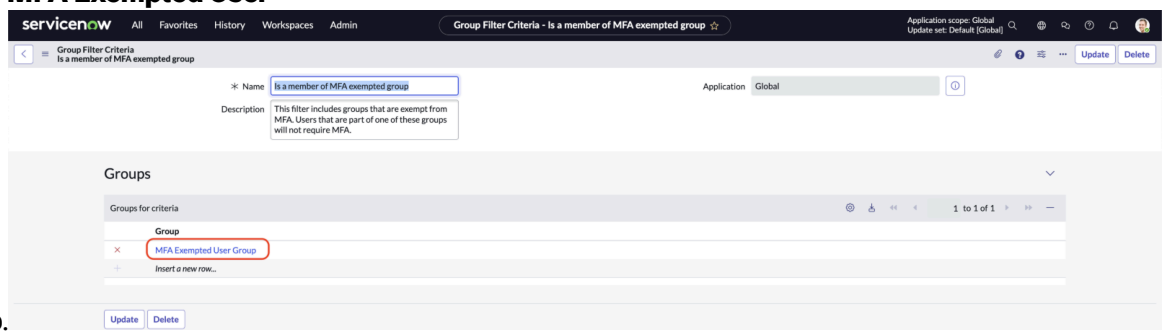

- Add users to this group as a member to exempt them from MFA enforcement.

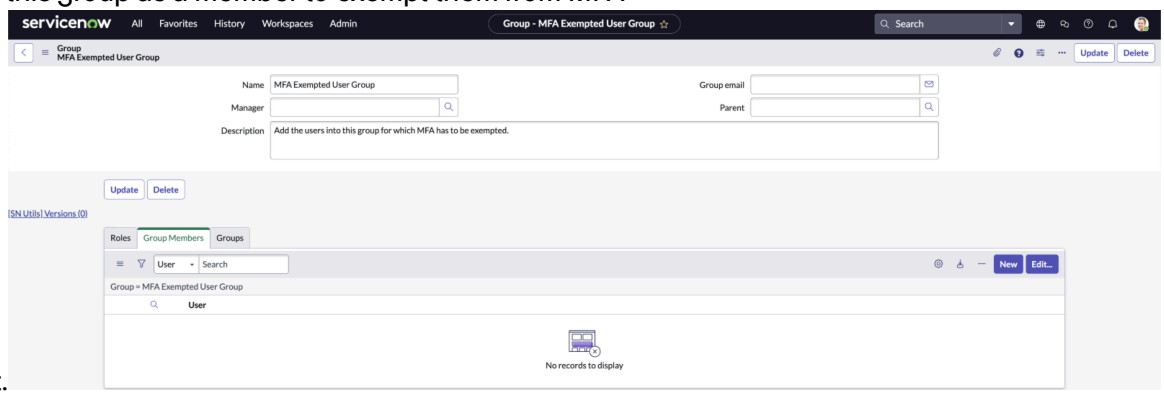

Note: If you have a different policy associated with the MFA context, you can add “Is a member of MFA exempted group” filter criteria to your policy and modify the policy conditions to exempt users of this group from MFA enforcement.

2. How can the MFAs mandate be relaxed for certain roles?

In the Yokohama release, an empty new role **Has MFA exempted role** filter criterion is added. There are conditions added to the MFA policy to exempt users who have the roles part of exempted role criteria from the MFA enforcement.

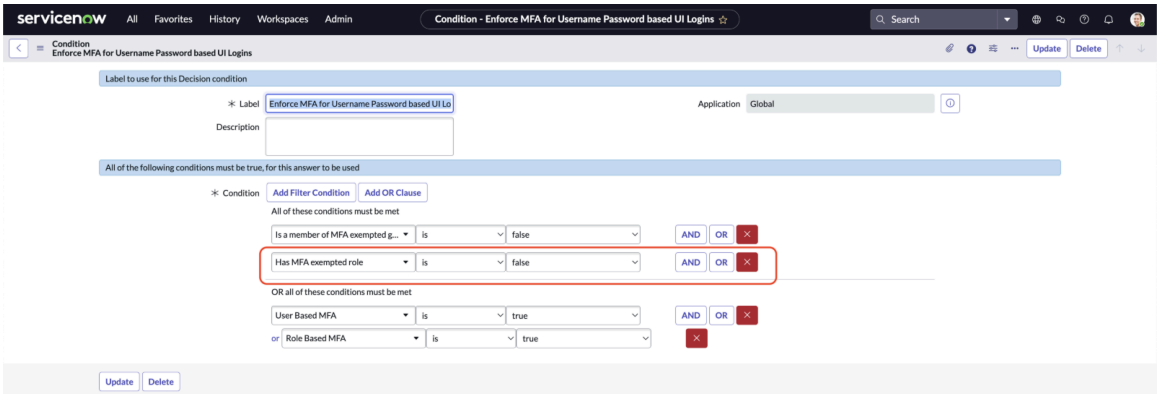

To relax MFA for specific roles, follow the procedure:

- Navigate to **MFA context**. The Step-Up MFA Policy associated with the MFA context record should be **Enforce MFA for non-SSO logins**.

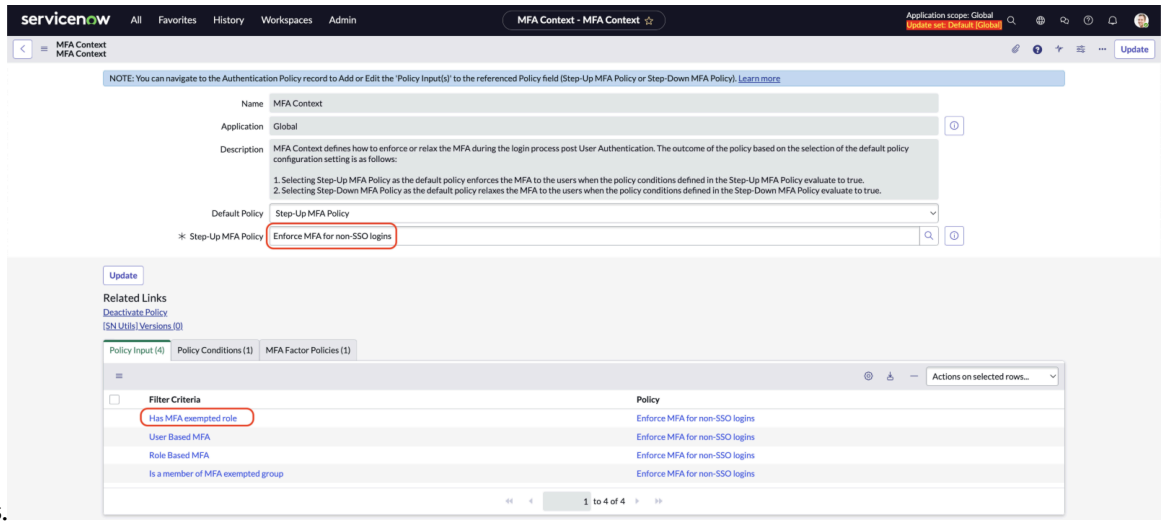

- Under the **Policy Input** related list, select **Has MFA exempted role** filter criteria record.

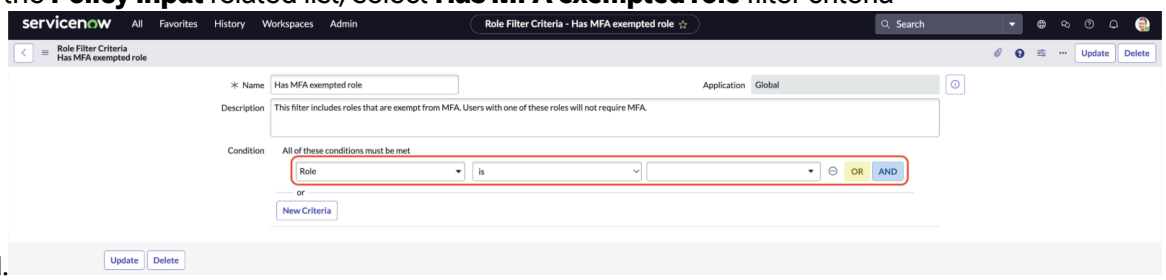

- Add the roles that you want to add to the condition. You can add multiple roles using the OR operator.

> **Note:** If you have a different policy associated with the MFA context, you can add **Has MFA exempted role** filter criteria to your policy. Modify the policy conditions to exempt users with exempted roles from the MFA enforcement.

3. How can the MFAs mandate be relaxed for certain groups?

In the Yokohama release, a user group **MFA Exempted User Group** is added. Based on the default, condition added to the MFA policy, the user or group who is a member of this group isn't enforced with MFA.

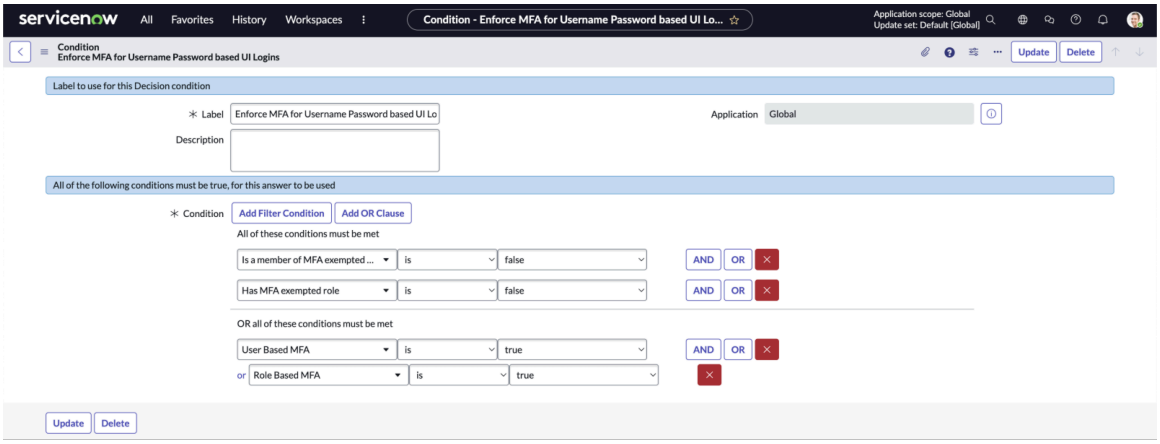

To relax MFA for specific groups, follow the procedure:

- Navigate to **MFA context**. The Step-Up MFA Policy associated with the MFA context record should be **Enforce MFA for non-SSO logins**.

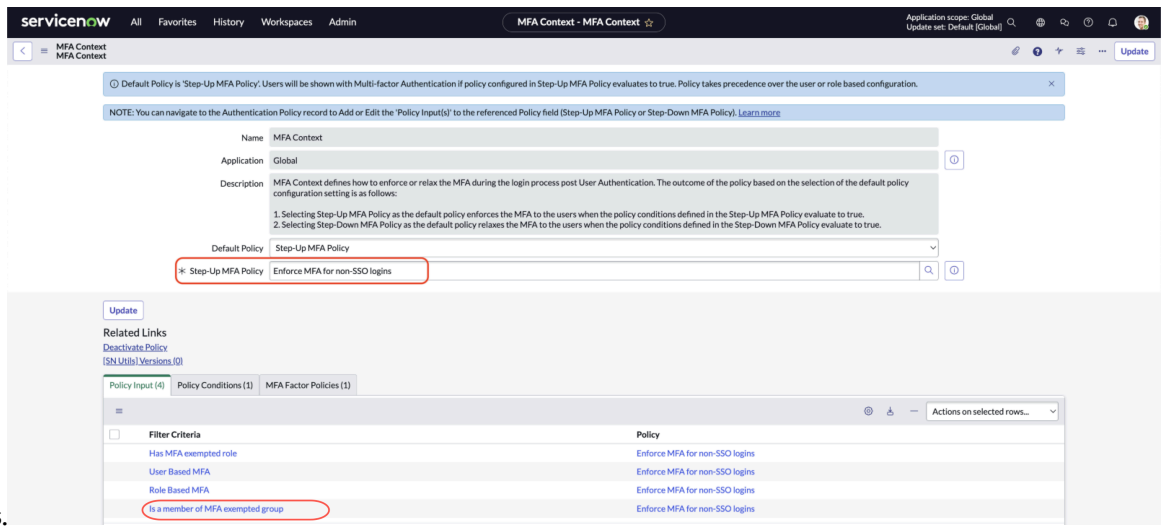

- Under the **Policy Input** related list, select the **Is a member of MFA exempted group** filter criteria record.

- Select **MFA Exempted User Group**.

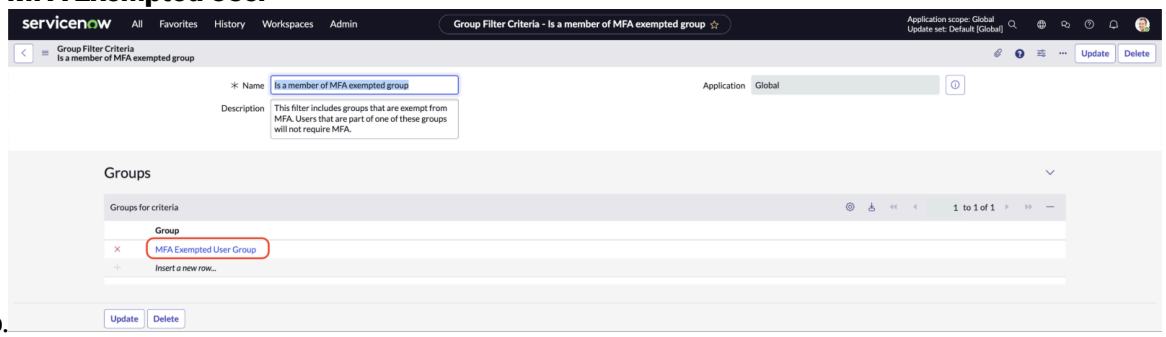

- Add the groups that you want to exempt from the MFA enforcement to this group.

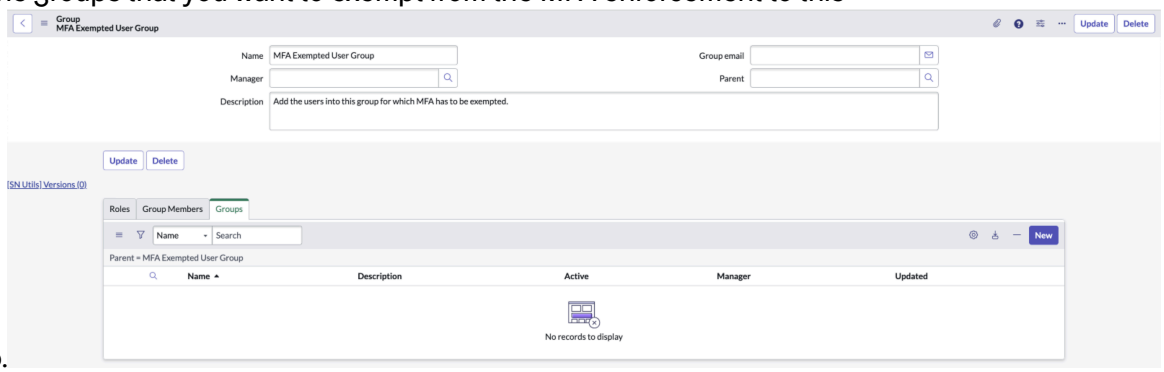

4. How can the MFAs mandate be relaxed for trusted networks?

- Navigate to **Adaptive Authentication** > **Filter Criteria** > **IP Filter Criteria**.
- Create a criterion to specify a trusted network. You can specify a list of IP ranges or subnets as part of the trusted network.

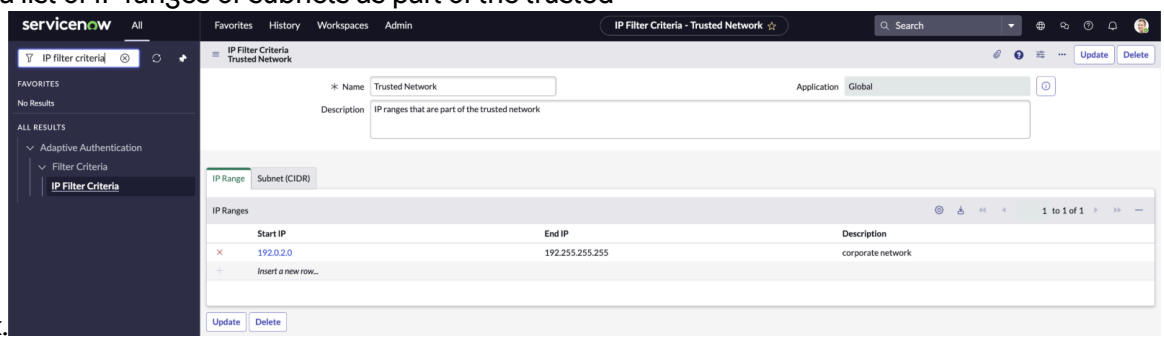

- Navigate to **Adaptive Authentication** > **Auth Policy Contexts** > **MFA context**.
- Open the policy associated with the context.

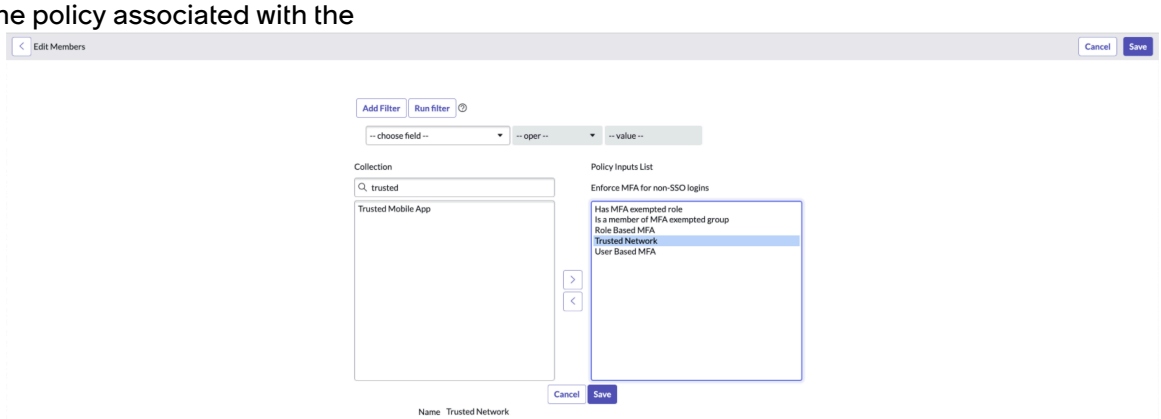

- Select the edit to add the **IP Filter Criteria** that you created to the **Policy inputs-related** list.

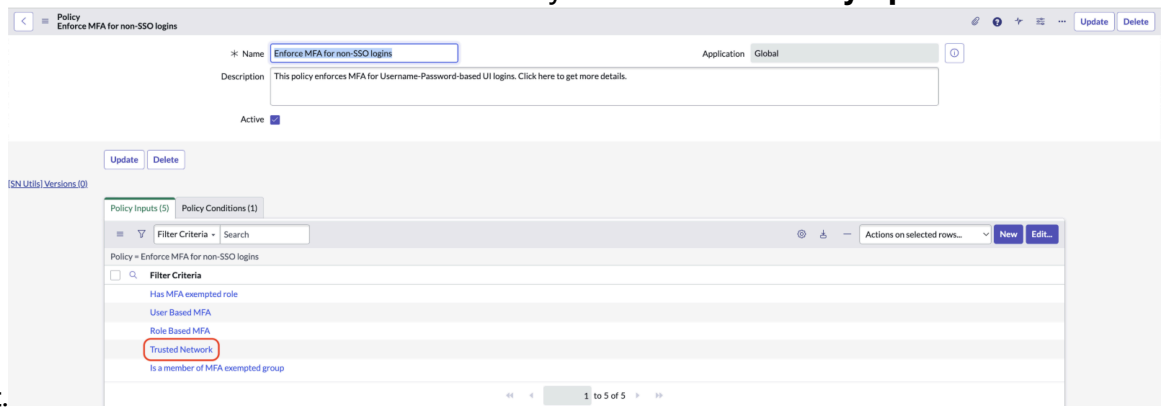

- Modify the policy condition to confirm it evaluates to false when users are part of the trusted network.

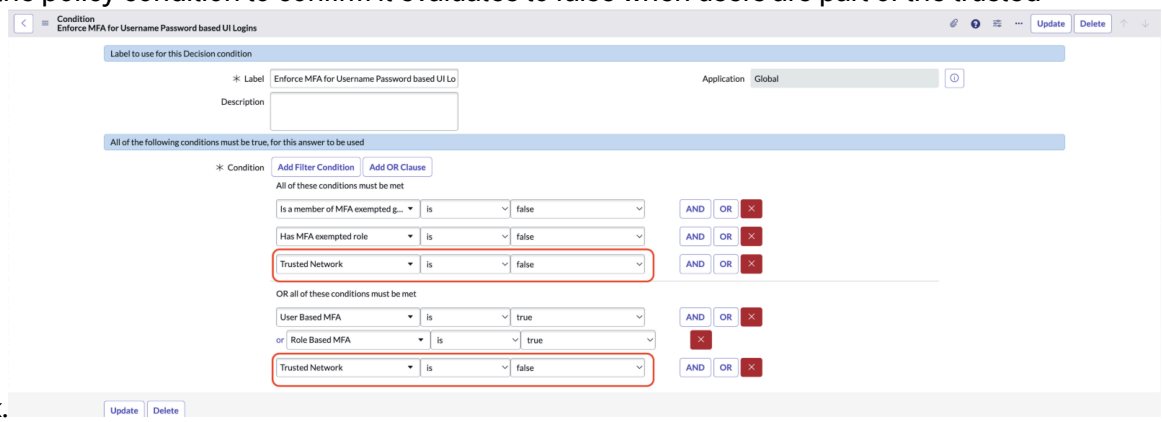

> **Note:** If you have a different policy associated with the MFA context, you can add the IP filter criteria created as part of Step 1 to your policy and modify the policy conditions to exempt MFA enforcement on the trusted network.

5. How can the MFAs mandate be relaxed for trusted locations?

You can use Location Filter Criteria which is available with the **Zero Trust – Location Based Access** (requires an additional subscription) plugin.

6. How to control the frequent MFA enforcement?

Use the Location Filter Criteria which is available with the **Zero Trust – Location-Based Access** (requires an additional subscription) plugin.

On the MFA validation page, there's a check box to remember a browser. MFA isn't enforced on the remembered browser:

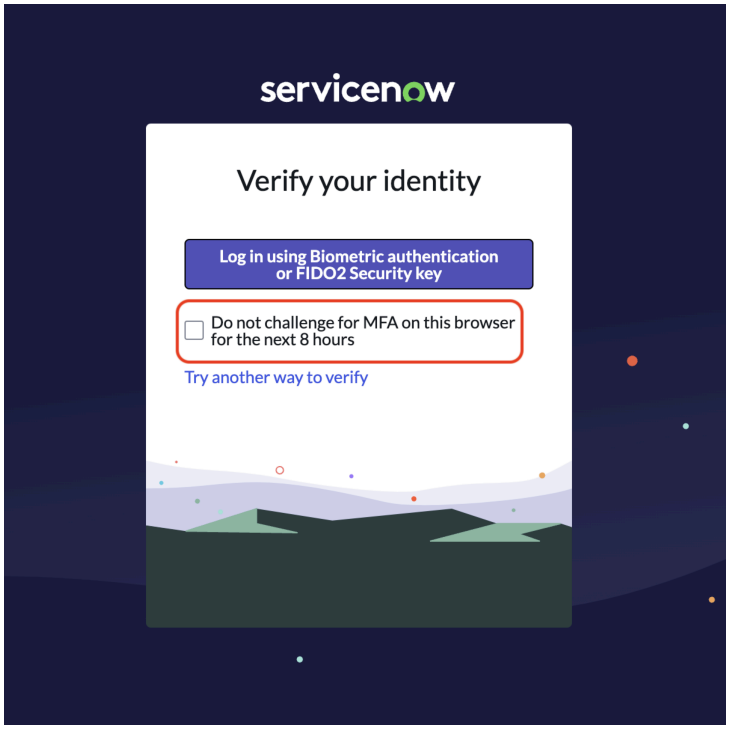

- The duration specified by this system property. `glide.authenticate.multifactor.browser.fingerprint.validity`. The default value of the property is 8 hours. This duration can be increased by up to 24 hours. Similarly using the `glide.authenticate.multifactor.remember.browser.default` system property the default value of the check box can be set to true.
- Navigate to **Multi-factor Authentication** > **Properties** and adjust these four properties to control the remembered browser feature.

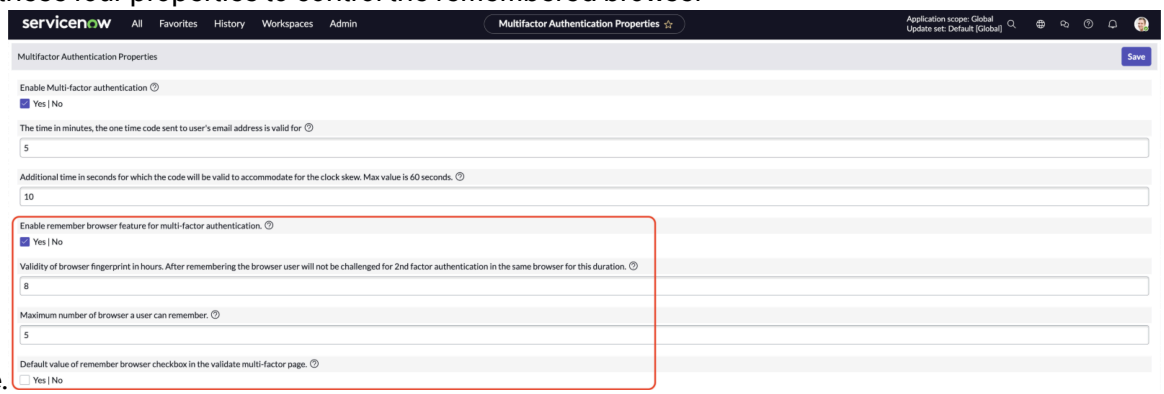

7. How does MFA work for accounts shared by users?

Single accounts shared by multiple users are a security risk. It isn't recommended to share an account with multiple users.

MFA metrics

FAQ related to understanding the MFA metrics.

1. How can I check the count of local logins which are not undergoing MFA?

You can navigate to **Security Center** > **Security Console** > **Security Metrics**. Under metrics for users click on local login not protected by MFA.

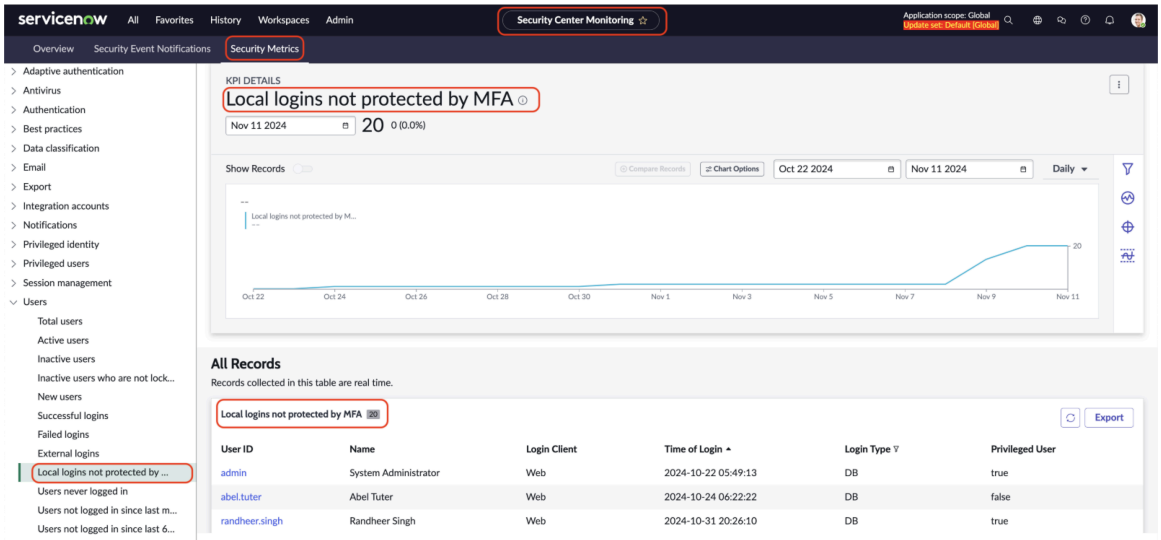

2. How can I check the number of users who still need to complete the MFA setup and are performing local login?

You can navigate to **Security Center** > **Security Console** > **Security Metrics**. Under metrics for users click on local login not protected by MFA.

MFA types

FAQ related to MFA types and why it's important.

1. What are the types of verification methods that are available for MFA with ServiceNow?

ServiceNow base system supports these verification methods.

a. Passkey

b. TOTP Authenticator apps such as Google Authenticator, Okta verify, Microsoft Authenticator, Authy, DUO

c. Biometric Authenticator (FIDO2) such as Windows Hello, Apple Touch ID, Face ID, android fingerprint sensor.

d. Hardware Security Keys (FIDO2) such as YubiKey, Thetis

e. Email One-time password (OTP)

f. SMS OTP - Multi-factor authentication with SMS com.snc.authentication.sms_mfaplugin installation and factor configuration are required to enable SMS OTP-based MFA.

2. Can a user configure multiple MFA factors or verification methods?

Yes, you can enroll for multiple MFA factors by visiting their user profile. For example, you can enroll a laptop with biometric authenticator, use the mobile phone with a passkey, and have an authenticator app setup.

3. What steps do users must perform to complete the MFA setup?

User can perform either of the following MFA options.

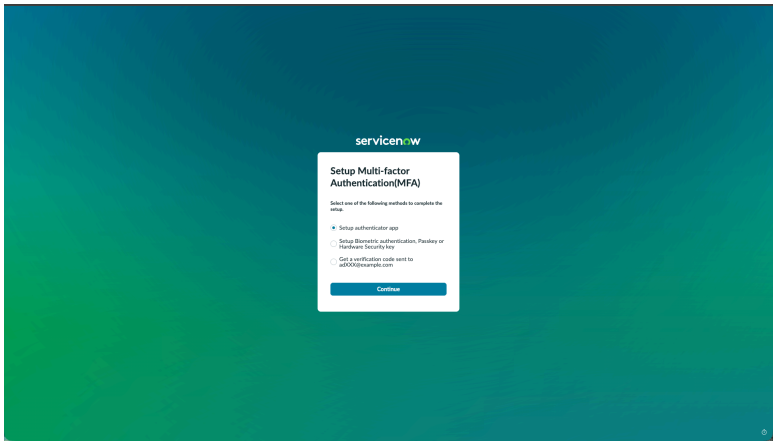

Refer the Multi-factor authentication documentation for more information about MFA setup.

4. Can the SMS and Email OTP-based MFA limited to certain users?

   Admin can set up MFA factor policies for email and SMS OTP-based MFA factors to limit these factors to certain user groups or roles.

5. The users don't have a mobile phone where they can set up an authenticator app. How can these users enable MFA?

   From the Xanadu release onwards, you can use a Biometric authenticator, passkeys, FIDO2 hardware security keys, and email OTP-based MFA without requiring an authenticator app setup on the mobile phone.

6. As an end user how to set up MFA?

   Refer the Using Multi-factor authentication documentation for more information about MFA setup.

MFA reset

FAQ related to MFA reset and why it's important.

1. Lost the authenticator app setup. How can I reset?

   It's better if the users enroll in multiple factors so that they aren't locked out due to MFA. For example, you can have the TOTP authenticator app, passkey, and email OTP-based MFA.

   In case users are unable to use any of the enrolled MFA factors the administrator can reset the MFA by following these steps.

   a. Clearing the Authenticator app setup:
      - Navigate to **All** > **Multi-factor Authentication** > **User Multi-factor Setup**.
      - Search for the user in the table.
      - Delete the record associated with the user

   b. Clearing the FIDO2 authenticators and passkeys:
      - Navigate to **All** > **Multi-factor Authentication** > **Web Authentication** > **User Public Credentials**.
      - Search for the user in the table
      - Delete the records associated with the user

   c. Clearing other multi-factor setups associated with the user

- Navigate to **All** > **Multi-factor Authentication** > **User Multi-factor Setup**.
- Search for the user in the table
- Delete the records associated with the user

In case no one including the admin can access the instance. The admin can perform a self-service MFA reset using a catalog item available at Now support.

2. How to avoid administrators from lockout due to MFA?

   The administrators enroll for multiple MFA factors so that they aren't blocked from accessing the instance.

   In case no one including the admin can access the instance. The admin can perform a self-service MFA reset using a catalog item available at Now support.

Exploring Multi-factor Authentication

Multi-factor Authentication (MFA) is an authentication method that requires users to provide information other than their basic credentials.

MFA is a security process that requires a user to provide two or more different verification factors to access a service or account. It adds an extra security layer of protection to your service beyond just a password, which makes it harder for unauthorized individuals to gain access.

By requiring multiple factors, MFA significantly enhances security and helps protect against various cyberthreats, including phishing and identity theft. Here's some insight about how MFA works:

- **First factor**: The user using their user name and password for login.
- **Second factor**: The user is prompted for a second factor that's with the user (An identity verification method such as an authenticator app or security key).

Further, these factors can be typically categorized into secured and less-secured based on their level of protection against common security threats.

- Secured Factors:
  - FIDO (Fast Identity Online): This factor uses hardware tokens or biometric authentication methods, providing a high level of security by confirming that the user has a physical device or unique biometric trait to verify their identity.
  - **TOTP (Time-Based One-time Password):** This factor generates a one-time password that is valid for a short period, usually 30 seconds. It's typically delivered through a mobile app, adding an extra layer of security by requiring the user to have access to a specific device and app.
- Less-Secured Factors:
  - **EMAIL:** This factor sends a verification code or link to the user's email address. While convenient, it's less secure because email accounts can be compromised.
  - **SMS:** This factor sends a verification code via text message to the user's phone number. It's also less secure due to the potential for SIM swapping and other mobile phone vulnerabilities.

To enhance security, it's recommended to prioritize the use of secured factors like FIDO and TOTP over less-secured factors like EMAIL and SMS.

Note:

- MFA is activated by default on ServiceNow.
- MFA is enabled using *glide.authenticate.multifactor* property. If you want to disable this property, you must provide a business justification about why you want to disable MFA.

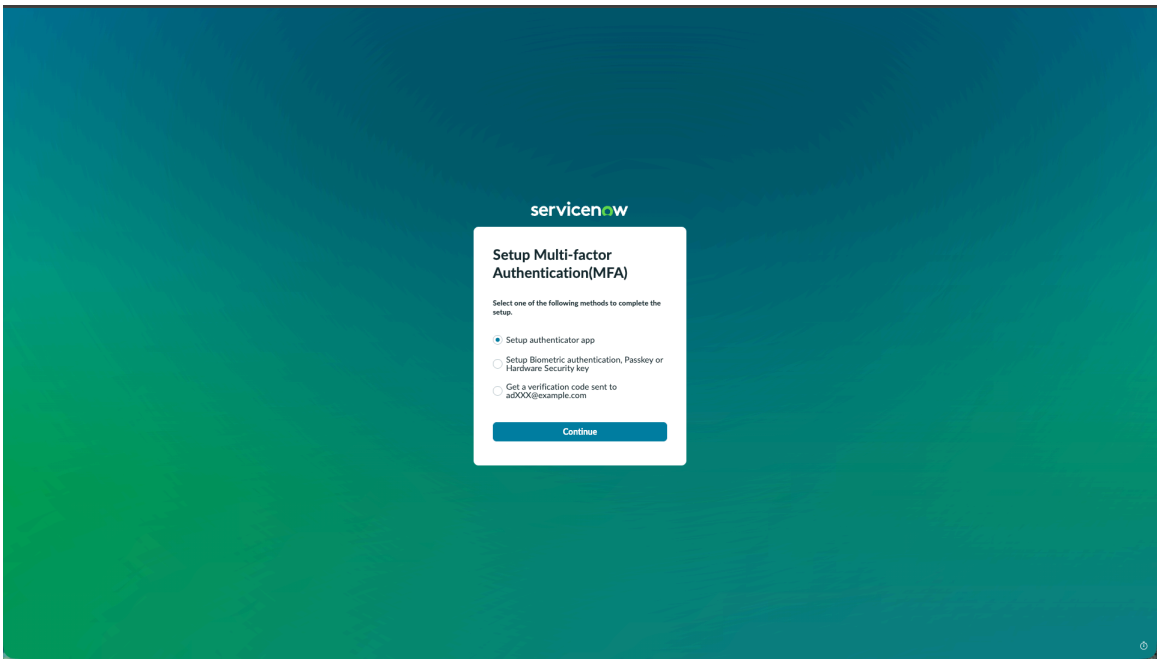

ServiceNow's MFA supports verification methods such as Authenticator App, Fast IDentity Online 2 (FIDO2), Passkey, and time-based One-time Password (OTP). Following are the details of available verification methods:

- **Authenticator App**: Apps that generate unique, temporary verification codes. For example: Okta, Google Authenticator, Microsoft Authenticator, and so on
- **FIDO2**: Physical devices that use public-key cryptography to validate user identities. For example: Hardware Keys (YubiKey), Biometric scanners (Apple's Touch ID).
- **Passkey**: Log in with a passkey by unlocking the device with a biometric sensor, PIN, or pattern.
- **OTP**: The secret key and the current time to generate a unique password that is only valid for a short period. For example: SMS (OTP) and Email (OTP).

You can use MFA along with the following:

- Local Database Authentication (native ServiceNow authentication) or Lightweight Directory Access Protocol integration
- SSO SAML or SSO OIDC. For more information, see Multi-Provider single sign-on (SSO).

Related topics

Multi-factor Authentication verification methods

Multi-factor Authentication system properties

Configure Multi-factor authentication

Configure multi-factor authentication (MFA) to improve your users security posture when using ServiceNow.

Firstly, to implement MFA, determine the MFA verification methods based on business and user needs.

After identifying the needs, with ServiceNow's MFA, you can select the right method or the combination of it for your users. Further, you can also use the various multi-factor authentication properties to enable, disable, and configure MFA to further enhance the user's MFA experience based on the authentication criteria.

Here's some of the important topics related to configuring MFA:

- Multi-factor Authentication verification methods
- Multi-factor Authentication system properties
- Multi-factor Authentication criteria

Multi-factor Authentication context

The Multi-factor Authentication (MFA) policy context uses a policy to define how and when MFA is enforced during the login process.

MFA context record

The MFA policy context defines whether your users must provide a second form of authentication when logging in. This context does not deny access to your instance as the post-authentication and pre-authentication policies. The policy you select in this context takes precedence over user or role-based configurations for multi-factor authentication.

To access the MFA context, navigate to **All** > **Multi-factor Authentication** > **MFA Context**.

Use the fields in the Post-authentication policy context record to define how your instance uses your policy.

Note:

- If the default policy is **Step-Up MFA Policy**, users will be shown with Multi-factor Authentication if policy configured in **Step-Up MFA Policy** evaluates to true. Policy takes precedence over the user or role based configuration.
- MFA with SSO login will only be available if `glide.authenticate.mfa.with.multisso.enabled` Property is set to true.
- You can navigate to the Authentication Policy record to Add or Edit the 'Policy Input(s)' to the referenced Policy field (**Step-Up MFA Policy** or **Step-Down MFA Policy**).
- MFA context policy applies only for user log ins. It does not apply for API authentication, basic auth, and OAuth resource owner password credential grant.

MFA context form

Field	Description
Name	Name of the policy context. This field is static and cannot be changed.
Description	Description of the context
Default Policy	Defines the default behavior of this context when evaluating the policy. Select from the following options. **Step-Up MFA Policy** Enforces MFA to users when the policy conditions defined in the **Step-Up MFA Policy** field evaluate to true.

MFA context form (continued)

Field	Description
	Step-Down MFA Policy Enforces MFA by default. MFA is not enforced only when the policy conditions defined in the **Step-Down MFA Policy** field evaluate to true.
Step-Up MFA Policy	The policy used for this context uses. This field appears only when the **Default Policy** field is set to **Step-Up MFA Policy**.
Step-Down MFA Policy	The policy used for this context uses. This field appears only when the **Default Policy** field is set to **Step-Down MFA Policy**.

Policy inputs and conditions

The **Policy Input** and **Policy Conditions** tabs display the inputs and conditions of the policy selected in the **Step-Up MFA Policy** or **Step-Down MFA Policy** field. These tabs serve as a reference, but cannot be used to change the policy inputs or conditions. To modify your policy settings, navigate to the policy using the reference icon ( ) next to the **Step-Up MFA Policy** or **Step-Down MFA Policy** field.

Note: Policy conditions can be created from here, but as a good practise it is recommended to add new policy conditions from policy page.

This example shows an MFA context record configured using a step-up MFA policy. This default policy means that MFA is enforced only when the conditions defined in the policy evaluate to true. The context uses a policy called **Step-Up MFA policy**. That policy has a set of inputs and conditions that are displayed in the **Policy Input** and **Policy Condition** tabs.

MFA policy context form

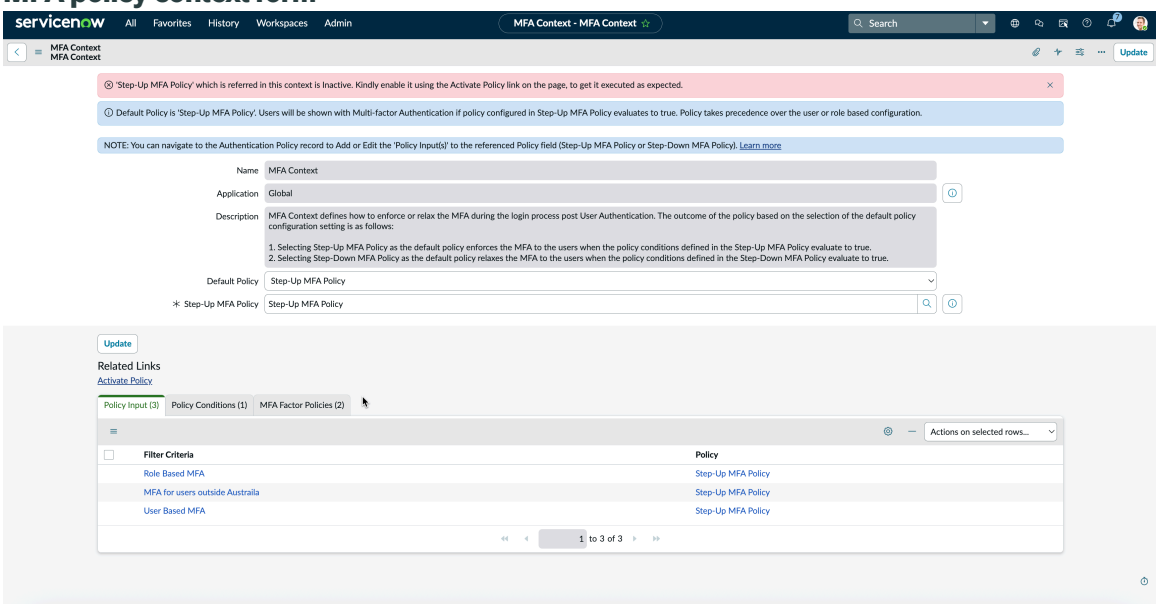

MFA factor policies

MFA factor policies are a critical component of an organization's security posture, enabling you to enforce additional verification steps beyond passwords. These policies define the authentication methods that users must employ to access providing a flexible and customizable approach to authentication. For more information, see Multi-Factor Authentication factor policies.

Multi-factor Authentication verification methods

ServiceNow's MFA supports verification methods such as Authenticator App, Fast IDentity Online 2 (FIDO2) and Time-based One-Time Password (TOTP).

The users can use the following options in addition to their user name and password to fulfill multi-factor authentication requirements. The users can setup MFA factors such as Authenticator applications, Biometric scanners, Hardware Keys, and SMS independently.

Authenticator Applications

An authenticator application is third-party software that generates temporary passcodes. Users can use these passcodes along with their password to log in into an instance that requires multi-factor authentication (MFA). For more detail on these applications, see Authenticator Applications.

Hardware Keys and Biometric scanners (Web Authentication)

The *Integration - Web Authentication* (com.snc.integration.webauthn) plugin allows hardware key or biometric reader authentication on your instance.

Biometric authenticators use fingerprint or facial recognition to identify users. Your users can use these authenticators on their devices as part of the multi-factor login process. For details on registering biometric authenticators, see Register a biometric authenticator.

Hardware keys are physical hardware that you can use to authenticate. Hardware keys are inserted into a port on your device to provide authentication. For details on registering hardware keys, see Register a hardware security key.

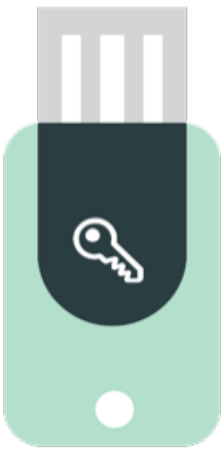

Passkey

Passkey authentication is a secure, password less method of log in, where the users use a passkey by unlocking their device with a biometric sensor, PIN, or pattern.

SMS

Admin can configure ServiceNow instance to require users who attempt to log in to the instance using SMS based OTP.

When users attempt to log in to ServiceNow, SMS OTP is sent to the mobile number associated with the sys_user record. Users can enter the six-digit verification code that it sent to the mobile device and verify their identity. For more information, see SMS as an MFA factor.

Email

Admin can configure ServiceNow instance to require users who attempt to log in to the instance using Email based OTP.

When users attempt to log in to ServiceNow, Email OTP is sent to the email address of the user. Users can enter the six-digit verification code that it sent to the email address and verify their identity. For more information, see Email as an MFA factor.

Related topics

- Web Authentication - MFA
- Multi-factor Authentication verification methods
- Email as an MFA factor

Web Authentication - MFA

Use the **Integration - Web Authentication (com.snc.integration.webauthn)** to allow hardware key or biometric reader authentication on your instance.

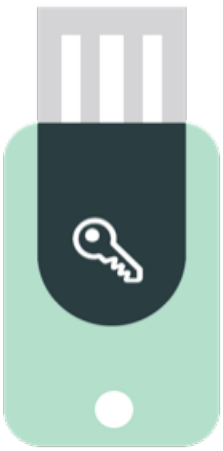

Hardware keys are physical hardware that you can use to authenticate. Hardware keys are inserted into a port on your device to provide authentication. For details on registering hardware keys, see Register a hardware security key.

Biometric authenticators use fingerprint or facial recognition to identify users. Your users can use these authenticators on their devices as part of the multi-factor login process. For details on registering biometric authenticators, see Register a biometric authenticator.

Configuring Multi-factor Authentication with Biometrics

Administrators can use the User Public Credentials list to view and manager user created credentials.

When a user registers an authenticator application, biometric authentication, or hardware key, you instance creates a record on the **User Public Credentials**[sys_user_public_credential] table. Use this table to see which users have registered an authenticator, as well as what types, and when they were registered and used. You can also mark these records as inactive to prevent the credentials to prevent users from using these credentials.

The **Integration - Web Authentication (com.snc.integration.webauthn)** plugin must be activated for Web Authentication (FIDO2).

Note: The **Integration - Web Authentication (com.snc.integration.webauthn)** plugin is installed by default.

User public credential form

Field	Description
Credential Nickname	Nickname for the credential. This nickname is chosen by the user when they register an authenticator.
User	User associated with the credential
Active	Whether the credential is active. Administrators can set a record to inactive to prevent a user from authenticating with this credential.
Authenticator	The type of authenticator registered by the user.
Registration Time	The date and time the user-created this credential
Last Used Time	The last date and time the user logged in with this credential.

Restricted authenticator types

If you have restricted an authentication method, such as biometric authenticators, users will not be able to create new credentials of that type. However, any credentials created before you made

this restriction will continue to work. You can disable records on the **User Public Credentials** table to prevent these credentials from being used after they have been created.

Authenticator configuration options

Use the Authenticator Configuration page to manage authenticator options on your instance.

Navigate to **Multi-factor Authentication** > **Web Authentication** > **Authenticator Configuration** to view and edit the default configuration options.

Authenticator Configuration form

Field	Description
Allowed authenticator type	Type of authenticators allowed to be registered. Select from: • **Platform** authenticators are attached or integrated into a device. Fingerprint readers or facial recognition available on mobile devices (such as Apple FaceID or TouchID) fall under this category. • **Roaming** authenticators can be removed from a computer or other client device and used elsewhere. Hardware keys fall under this category.
Attestation Type	Setting the value to direct or indirect will require importing authenticator metadata to attest to the provenance of an authenticator during registration. • None • Direct • Indirect
Platform self-attestation	Whether self-attestation is enabled for platform authenticators.
Cross platform self-attestation	Whether self-attestation is enabled for roaming authenticators.
User verification	Select from **Preferred** or **Required**. If required, web authentication flow prompts users for verification using PIN or Biometrics.
Verify user presence	Whether web authentication flow requires the user presence verification.
Resident key	Select from **Preferred** or **Required**. If required, the authenticator persists the public key credentials within the authenticator storage.
Timeout (In ms)	Maximum time limit for completing web authentication registration and authentication. Time is in milliseconds.

Multi-Factor Authentication factor policies

Use the MFA factor policies to specify the types of authentication factors that you would like to permit for your instance.

MFA factor policies are a critical component of an organization's security posture, enabling you to enforce additional verification steps beyond passwords. These policies define the authentication methods that users must employ to access your organization's resources, providing a flexible and customizable approach to authentication.

Implementing MFA factor policies is essential for enhancing the security of your organization's systems and data. These policies provide an additional layer of protection against cyberthreats, making it more difficult for attackers to gain unauthorized access.

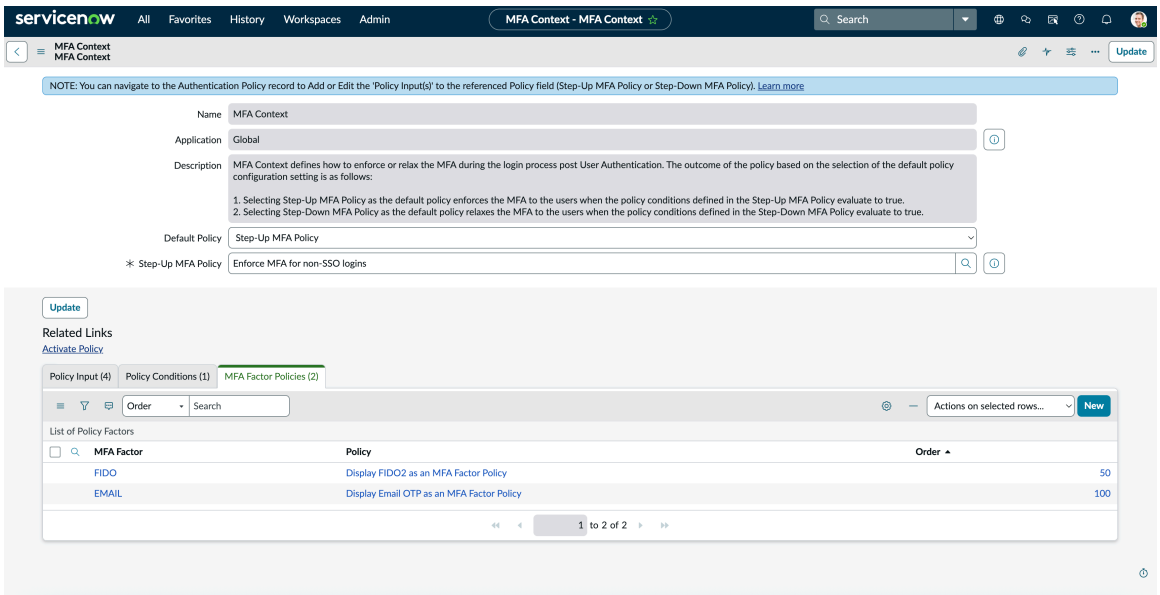

To use the MFA factor policies you have to configure the policy inputs and policy conditions along with an MFA Context. To know more, see Multi-factor Authentication context.

Following are the MFA factor policies available in ServiceNow that enables you to specify the types of authentication factors that are permitted or required:

- FIDO2
- SMS
- Email

To get the most out of MFA factor policies, you need to understand how to configure and manage them effectively. This includes defining the authentication methods, specifying policy inputs and conditions, and configuring policy enforcement (MFA Context). By understanding and implementing MFA factor policies, you can significantly improve the security and integrity of your organization's systems and data.

Related topics

Multi-factor Authentication context

FIDO2 as an MFA factor

SMS as an MFA factor

Email as an MFA factor

FIDO2 as an MFA factor

You can configure FIDO2 as an MFA factor policy to enforce MFA for yours.

FIDO2 is a password-less authentication standard that enables users to authenticate using a physical security key or biometric authentication. It provides a more secure alternative to traditional MFA methods, reducing the risk of phishing and other cyberattack.

The FIDO2 factor policy enhancement provides a secure authentication method to your multi-factor authentication (MFA) policies. You can configure FIDO2 as an MFA factor policy option, providing a higher level of security compared to traditional methods like Email and SMS.

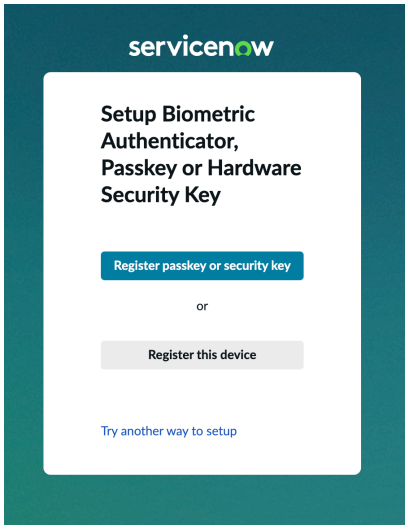

You can configure FIDO2 factor policy and when the users satisfies the factor policy condition, the during log in to ServiceNow, FIDO2 setup is displayed for the users who haven't already added registered Hardware key or Biometric on their profile.

If the registration is completed, then second factor validation screen is displayed to log in.

Note: FIDO2 can also be self-enrolled by the users. To know more about how to self-enroll, see Set up Multi-factor authentication on your user profile.

Key Benefits

The following are some of the key benefits of using FIDO2 as an MFA factor:

Exclusive FIDO2 Authentication

Ensure high-privilege accounts can only authenticate using FIDO2's strong authentication capabilities (biometrics, passkeys, or hardware security keys).

Exclusion of less secure methods

Suppress other authentication methods when FIDO2 is the only matching policy.

Forced enrollment

Require users to register a FIDO2 key if not already enrolled.

Granular control

Apply strict enforcement to specific roles or groups using policy-based targeting.

As a higher-security factor, FIDO2 has exclusive enforcement capabilities. When it's the only matching policy for a user:

- Overrides other enroll-enrolled factors.
- Forces FIDO2 registration, if the user isn't enrolled.
- Becomes the exclusive authentication option.

Example Configurations and User Behaviors

The following table illustrates how different user scenarios are handled based on their roles and enrolled factors.

Example Factor Policy Conditions:

- **FIDO2 Factory Policy**: Condition is "ITIL role should be true".
- **EMAIL Factor Policy**: Condition is "ASSET role should be true".

Example user	Has roles	Already enrolled factors	Matching policies	Behavior
andrew.och	**ITIL**	None	FIDO2	User is redirected to MFA setup with FIDO2 only. After registration, FIDO2 is the only authentication option.
abel.tuter	**ITIL**	Authenticator	FIDO2	User is redirected to MFA setup with FIDO2 only, even if the user has Authenticator as self-enrolled factor. **Note:** If the user hasn't registered to any MFA factor, then the user is redirected to MFA setup with FIDO2.
aileen.motterm	**ASSET**	Authenticator	Email	Sees Email and Authenticator options during log in. The user can choose either factor or optionally register FIDO2.
abraham.lincoln	**ASSET**, **ITIL**	Authenticator	Email and FIDO2	Sees Email and Authenticator option during log in. The user can register FIDO2 during validation. After registration, the user can see all the 3 factors.

By configuring FIDO2 as an MFA factor policy, you can significantly enhance the security of your authentication processes.

Configure FIDO2 as an MFA factor

Configure policy input and condition to display FIDO2 as an MFA factor policy for authentication.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Multi-factor Authentication** > **MFA Context**.
2. Select the **MFA Factor Policies** tab.
3. Select the **Display FIDO2 as an MFA Factor Policy**.

4. Select **New** to add **Policy Inputs**.
5. Select the filter criteria that you want to create.

   Following are the types of filter criteria:

   - IP Filter Criteria
   - Role Filter Criteria
   - Group Filter Criteria

   For example, Role Filter Criteria.

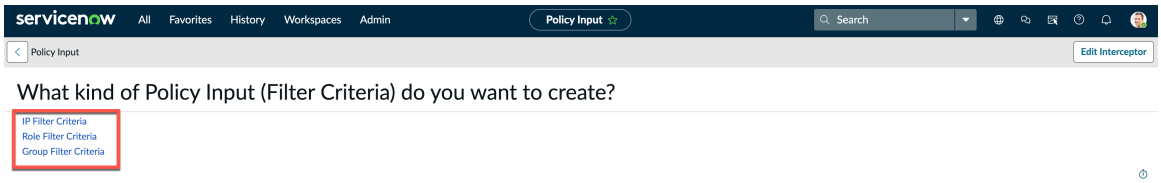

6. Select **Role Filter Criteria**, fill the fields for the role filter criteria and submit the record.

   The new policy is created. For more information, see Role Filter Criteria.

   Let's take an example of using **ITIL** role for the user (**andrew.och**) as the policy input and submit.

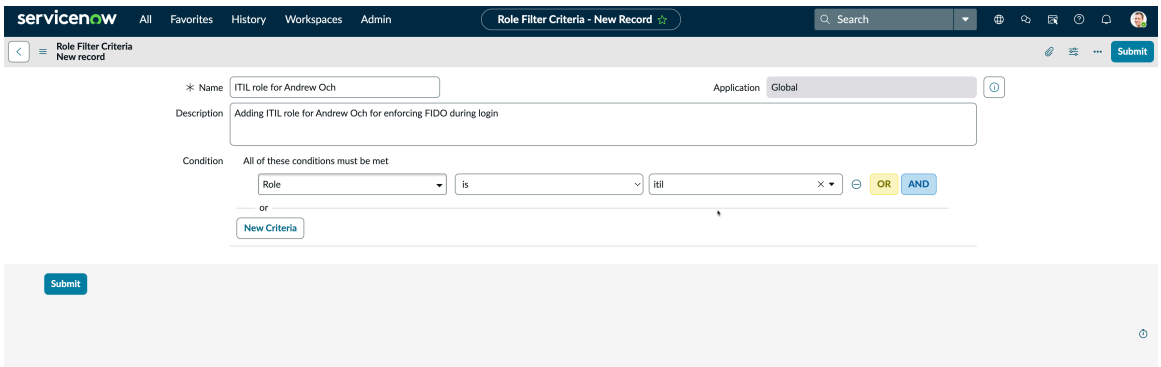

7. On the Policy - Display FIDO2 as an MFA Factor Policy page, select **Policy Conditions**.
8. Select **New** to add policy conditions.
9. On the form, fill in the fields.

Condition form

Field	Description
Label	Name to identify the condition.
Description	Description of the condition.

Field	Description
Condition	Logical combination of multiple policy inputs (filter criteria) that is used to evaluate authentication requests. Select the role-based filter criteria policy that was created for the condition.

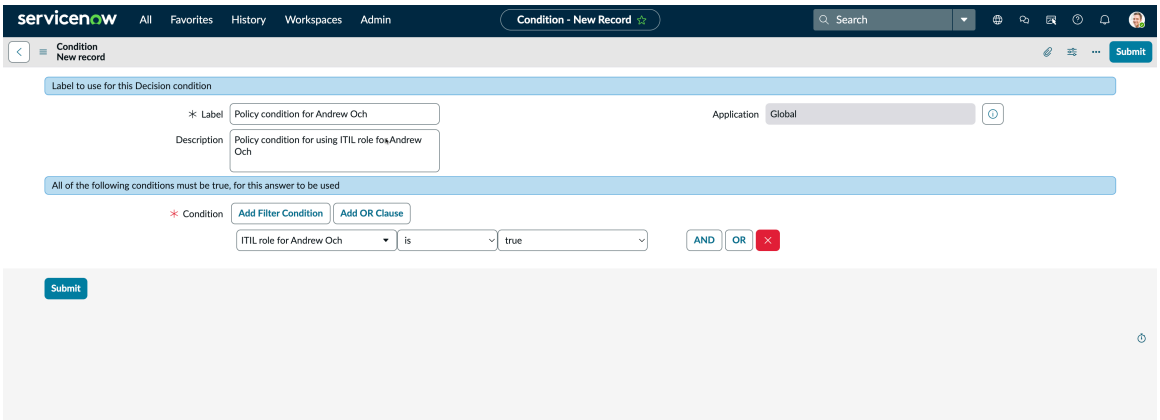

10. Select **Submit**.

    Based on the policy input and condition, if the user (**andrew.och**) tries to log in to the instance, the user is shown as the FIDO screen to either enroll and register.

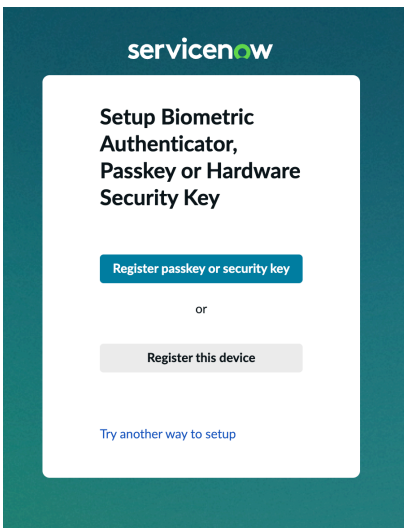

    To know more about different configuration example and user behaviors, see Example Configurations and User Behaviors.

11. **Optional:** Repeat step 8 to create additional policy conditions.

    Note: If you create multiple policy conditions, the final output of the access policy depends on the logical OR output of the all policy conditions. Based on the conditions the policy is evaluated.

SMS as an MFA factor

Multi-factor authentication (MFA) with SMS as a factor for your authentication.

Admin can configure ServiceNow instance to require users who attempt to login the instance using SMS based OTP.

When users attempt to login to ServiceNow, SMS OTP is sent to the mobile number associated with the sys_user record. User's can enter the six-digit verification code that it sent to the mobile device and verify their identity.

You can configure MFA with SMS using the out of the box Twilio as well. For more information, see Multi-factor authentication Providers.

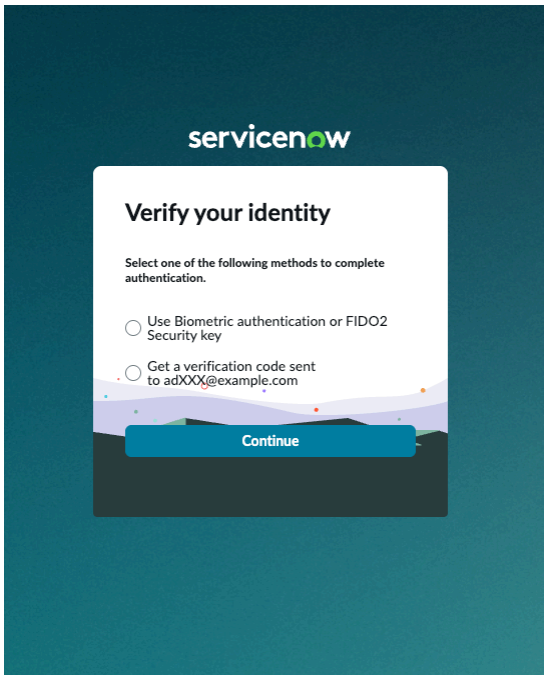

Further the MFA with SMS can be controlled based on the policy input and conditions using filter criteria. Following are the types of filter criteria:

- IP Filter Criteria
- Role Filter Criteria
- Group Filter Criteria

Activate the MFA with SMS plugin

For MFA with SMS, install the Multi-factor authentication with SMS (`com.snc.authentication.sms_mfa`) plugin.

Before you begin

Role required: admin

About this task

The following items are installed with Multi-factor authentication with SMS:

- Adaptive Authentication (`com.snc.adaptive_authentication`
- Notify - Twilio Direct Driver (`com.snc.notify.twilio_direct`)

Dependent Plugin: Integration - Multifactor Authentication (`com.snc.integration.multifactor.authentication`)

Note: You can load the demo data when installing the plugin if you are configuring a custom provider for generating the SMS.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Multi-factor authentication with SMS plugin (`com.snc.authentication.sms_mfa`) using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Configure SMS as an MFA factor

Configure policy input and condition to display SMS OTP as an MFA factor policy for authentication.

Before you begin

Plugin required: Multi-factor authentication with SMS (`com.snc.authentication.sms_mfa`).

Role required: admin

Note: The MFA context policy must be evaluated as `true` to apply the SMS factor policy.

Procedure

1. Navigate to **All** > **Multi-factor Authentication** > **MFA Context**.
2. Click the **MFA Factor Policies** tab.
3. Select the **Display SMS OTP as an MFA Factor Policy**.

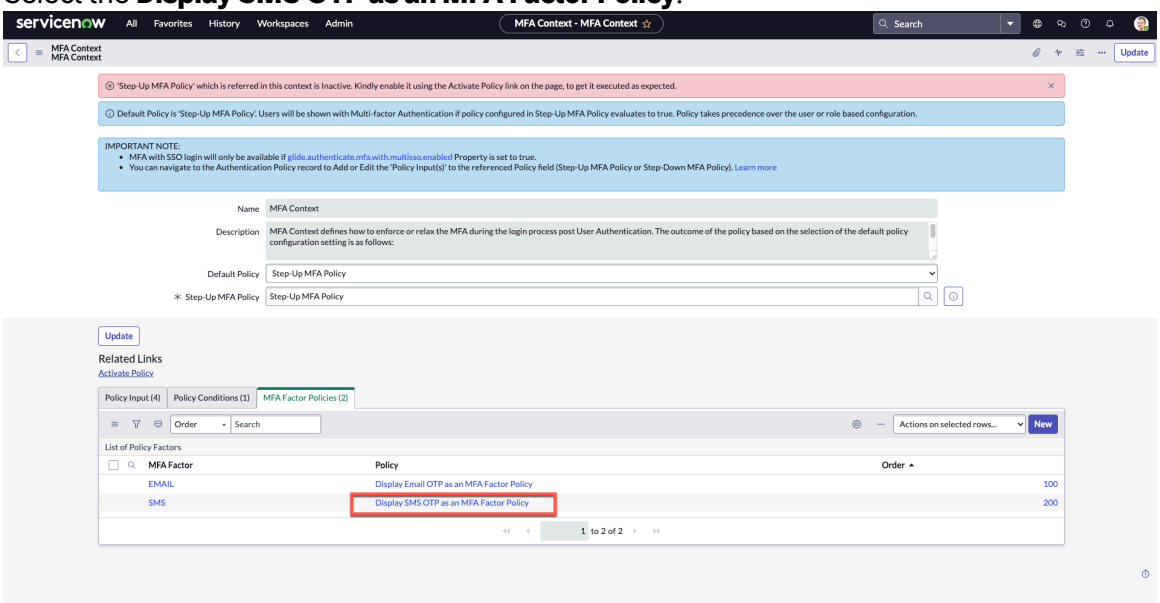

4. Click **New** to add **Policy Inputs**.

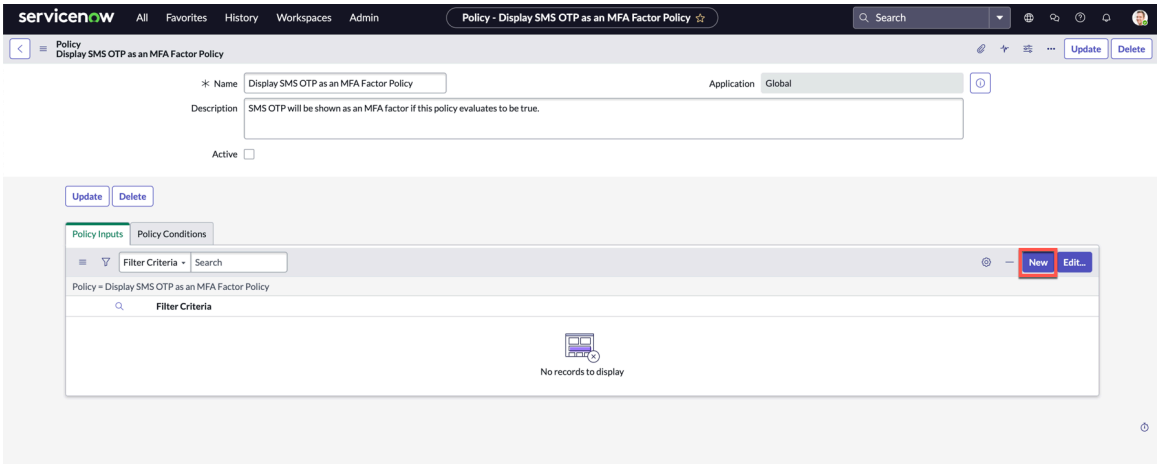

5. Select the filter criteria that you want to create.

   Following are the types of filter criteria:

   - IP Filter Criteria
   - Role Filter Criteria
   - Group Filter Criteria

   For example, Role Filter Criteria.

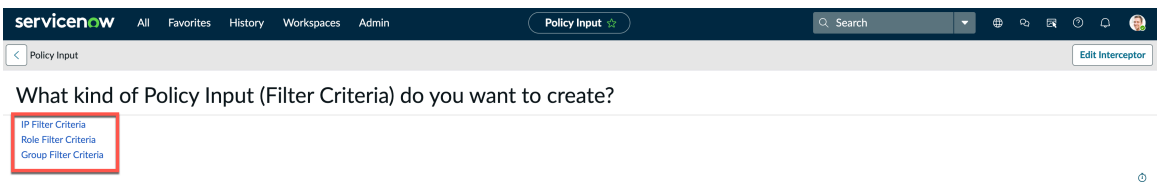

6. Click **Role Filter Criteria**, fill the fields for the role filter criteria and submit the record. The new policy is created. For more information, see Role Filter Criteria.
7. On the Policy - Display SMS OTP as an MFA Factor Policy page, click Policy conditions.

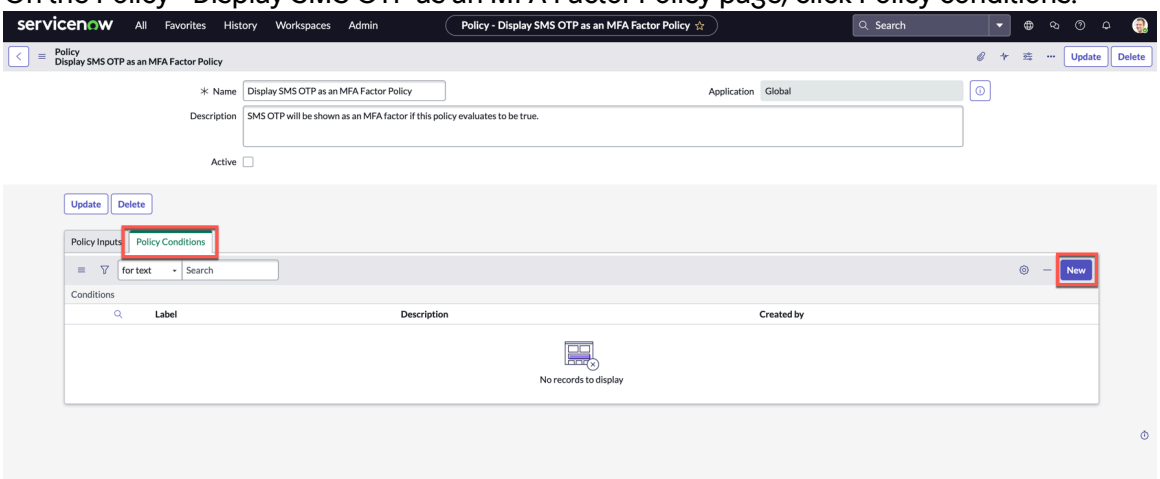

8. Click **New** to add **Policy Conditions**.

9. On the form, fill in the fields.

Condition form

Field	Description
Label	Name to identify the condition.
Description	Description of the condition.
Condition	Logical combination of multiple policy inputs (filter criteria) that is used to evaluate authentication requests. Select the role-based filter criteria policy that was created for the condition.

10. Click **Submit**.
11. **Optional:** Repeat step 8 to create additional policy conditions.

> **Note:** If you create multiple policy conditions, the final output of the access policy depends on the logical OR output of the all policy conditions. This means that the policy evaluates to true if any one of your policy conditions is met.

Based on the role filter (users) policy and the conditions specified for the role is matched, the SMS as MFA factor is shown as an option for authentication for the users.

Multi-factor authentication Providers

Use MFA providers to configure SMS and Email based authentication to ensure every user can login securely.

On the ServiceNow AI Platform, you can configure MFA providers with the mechanism of using multi-factor authentication such as Email and SMS.

You can use the following provider configuration available for MFA within the ServiceNow AI Platform:

- Email Provider Configuration
- Twilio Provider Configuration
- Infobip Provider Configuration.

> **Note:** Twilio and Infobip Provider Configuration are auto-populated by enabling the Load Demo Data when installing the Multi-factor authentication with SMS (`com.snc.authentication.sms_mfa`) and Notify - Twilio Direct Driver (`com.snc.notify.twilio_direct`) plugins.

You can also create your own provider configuration for enabling the multi-factor authentication with SMS and Email.

> **Note:** The Infobip Provider configuration is provided as a part of demo data, you can edit the fields based on your requirements for configuration your own provider.

Configure MFA Provider

Configure SMS and Email with the Provider to ensure every user can login securely.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Multi-factor Authentication** > **Providers**.

   Following provider configuration are available for MFA within the ServiceNow AI Platform:

   - Email Provider Configuration
   - Twilio Provider Configuration
   - Infobip Provider Configuration.

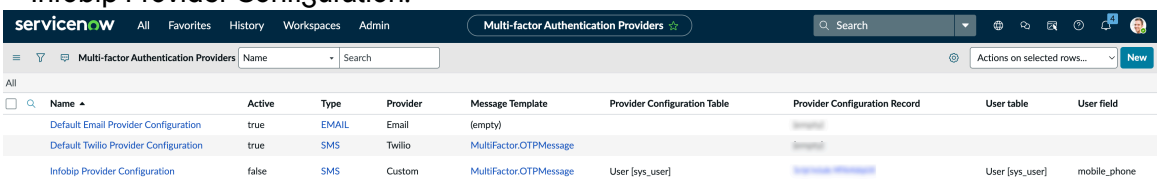

2. To create a new provider, click **New**.
3. On the form, fill the fields.

Condition form

Field	Description
Name	Name of the record.
Type	Description of the record.
Provider	Choose **Twilio** or **Custom**. **Note:** To configure Twilio, see Configure Notify with Twilio. When choosing **Custom**, you need to specify the following fields: ○ Provider Configuration Table ○ Provider Configuration Record ○ Script ○ User table ○ User field
Message Template	The message template for the record.
Active	Option to make the provider configuration active.

4. Click **Submit**.

   Based on the message template and provider configurations, the SMS or Email is sent to users as a factor for authentication during the login process.

Vonage Provider custom configuration (Tutorial)

Configure a SMS with Vonage Provider to ensure every user can login securely.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Web Services** > **Outbound** > **REST Message** and perform the REST message configuration based on the information from Vonage API Dashboard.

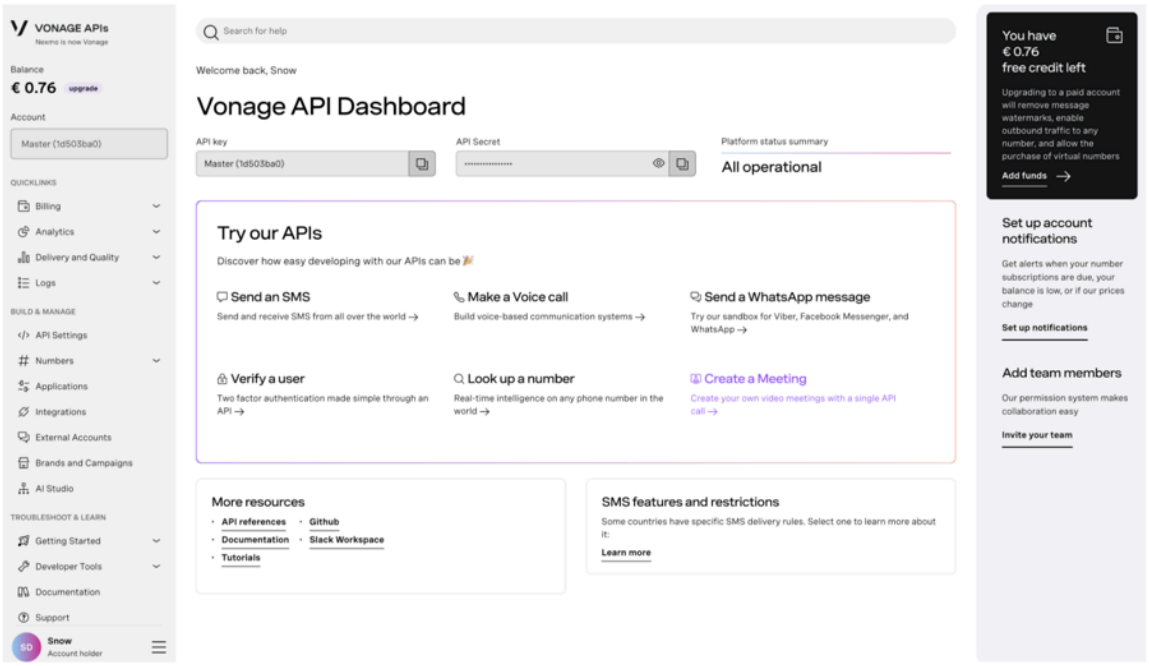

2. Click **New** to create a new **REST Message**.
3. Provide a **Name** and **Endpoint**.

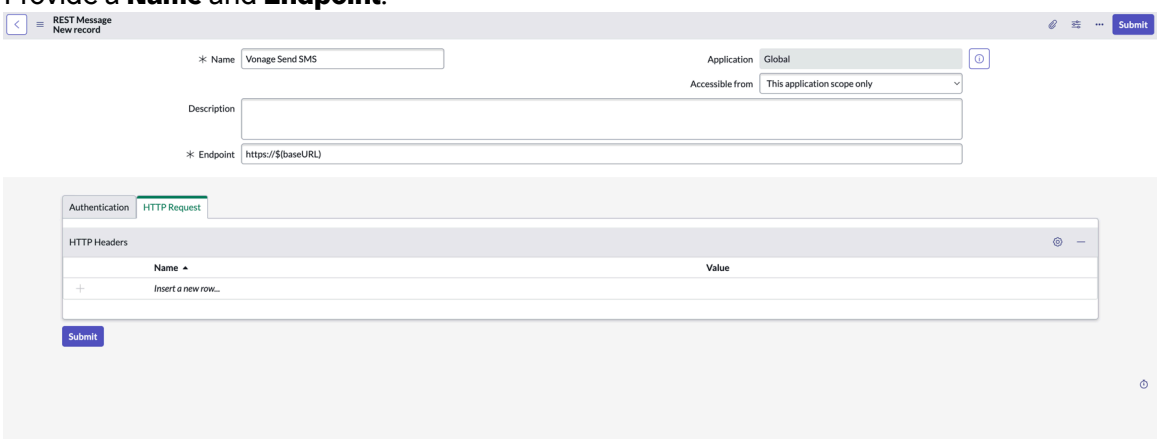

4. Click **Submit**.
5. Open the created record, in the **HTTP Methods** related list, click **New** and select the HTTP method as **POST**.

6. Fill the following fields:
   - Endpoint: `https://${baseURL}/sms/json`
   - Content:`api_key=${apiKey}&api_secret=${apiSecret}&text=${text}&from=${from}&to=${to}`
   - Content-Type: `application/x-www-form-urlencoded`

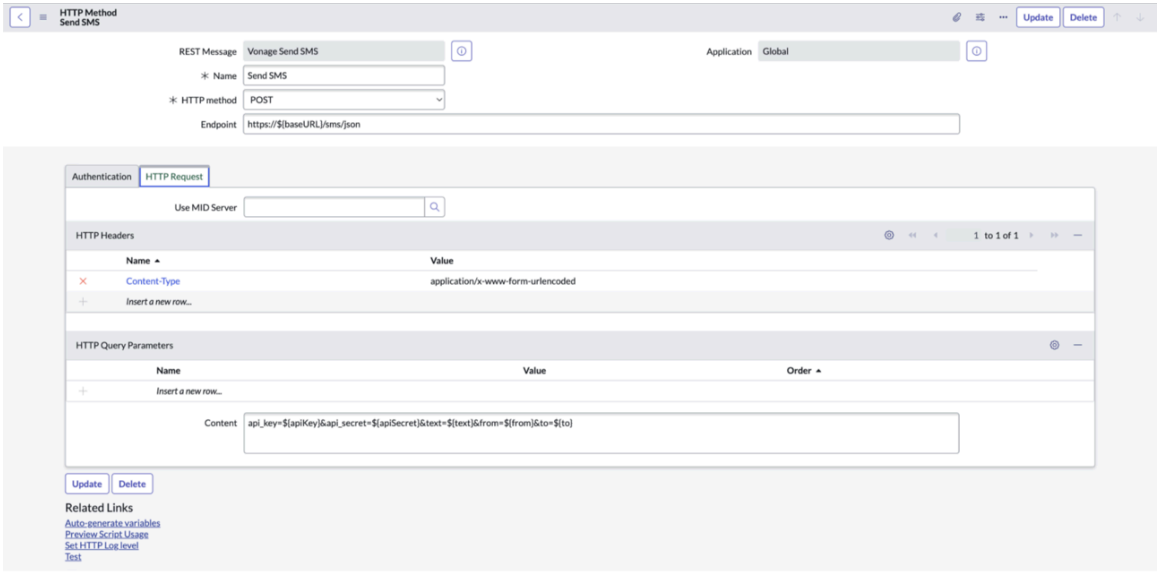

7. Update the record.
8. Click **Auto-generate variables** from the **Related Links** section.
9. Import SI shared in the folder and copy the apiKey and secret from the Vonage API dashboard.

> **Note:**
>
> The apiKey and secret are located in your Vonage account settings in the Vonage Dashboard.

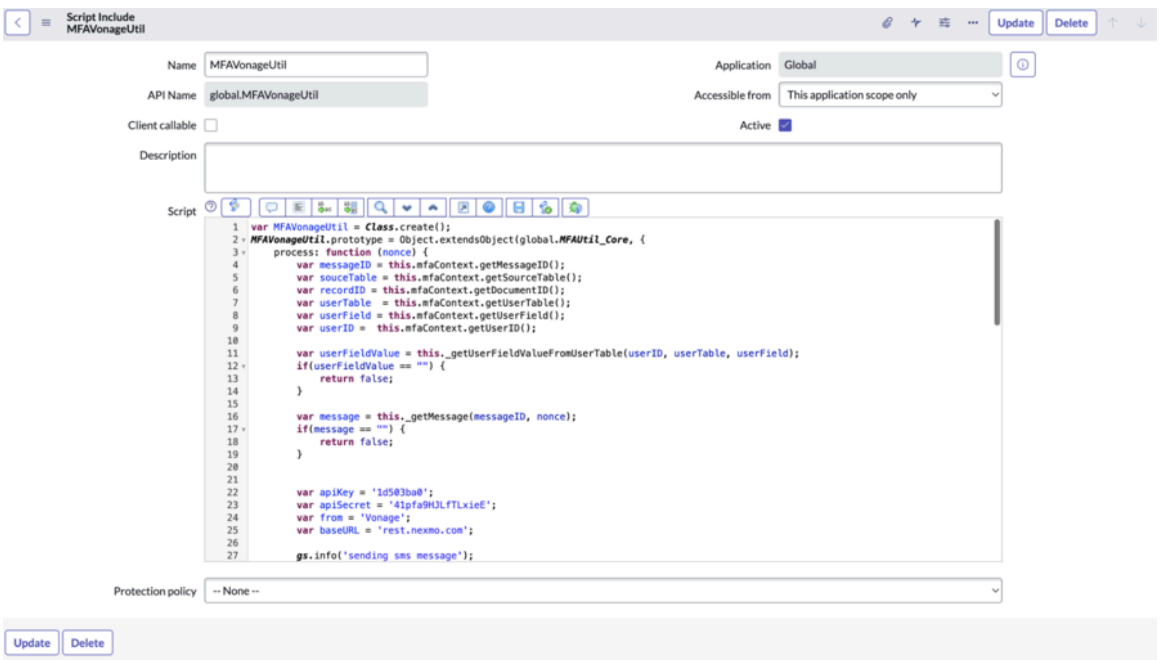

10. Create a custom Table Phone Numbers with two columns, For example:

- **Column Label**: User, **Type**: Reference, **Reference**: User (sys_user).
- **Column Label**: Phone Number, **Type**: Phone Number(E164).

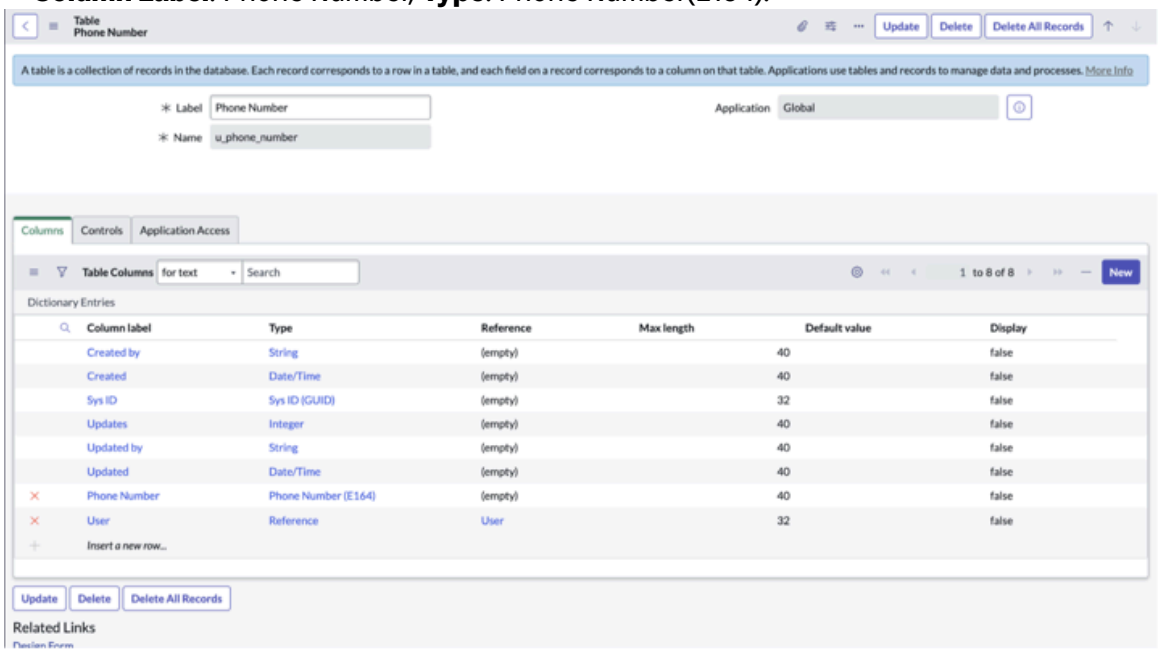

11. Create a custom provider in Multi-Factor Provider table.

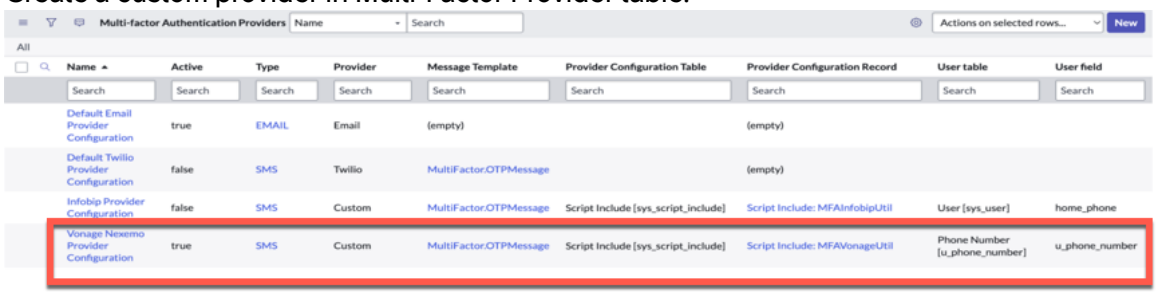

To know more about provider configuration, see Configure MFA Provider.

Email as an MFA factor

Multi-factor authentication (MFA) with Email as a factor for your authentication.

Admin can configure ServiceNow instance to require users who attempt to login to the instance using Email based OTP.

Note: MFA with Email is activated with the Integration - Multifactor Authentication (`com.snc.integration.multifactor.authentication`) plugin by default. You need to configure the policy inputs and conditions.

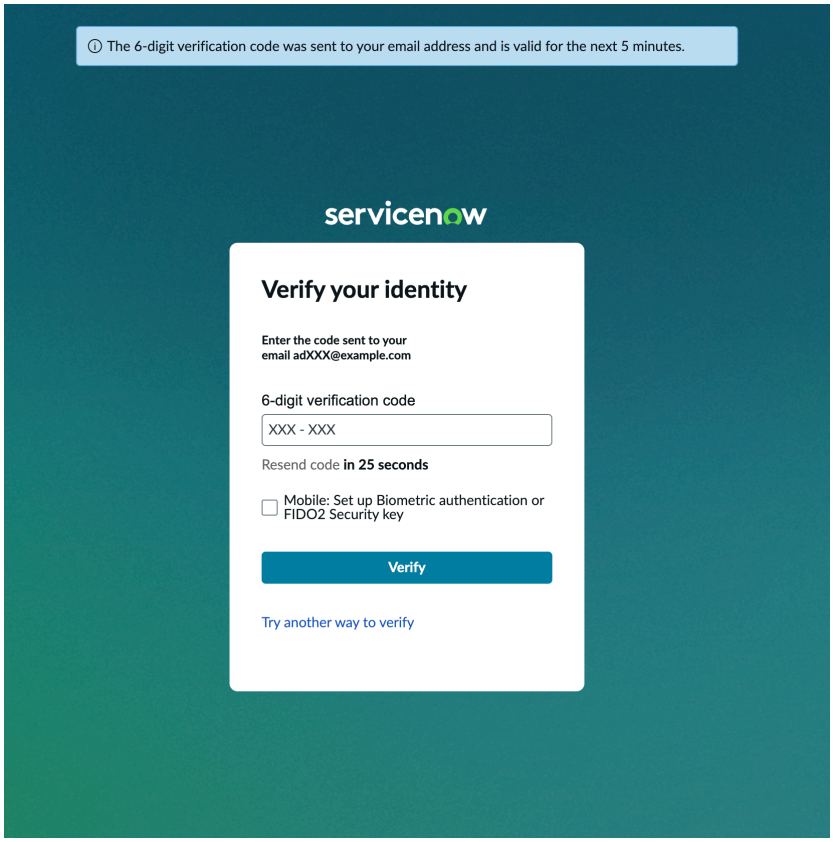

When users attempt to login to ServiceNow, Email OTP is sent to the Email address associated. User's can enter the six-digit verification code that it sent to the email address and verify their identity.

Configure Email as an MFA factor

Configure policy input and condition to display Email OTP as an MFA factor policy for authentication.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Multi-factor Authentication** > **MFA Context**.
2. Click **MFA Factor Policies** tab.

   Note: The **EMAIL** as **MFA Factor with a Policy** is available by default. You can edit the policy and specify the policy inputs and conditions.

3. Select the **Display Email OTP as an MFA Factor Policy**.

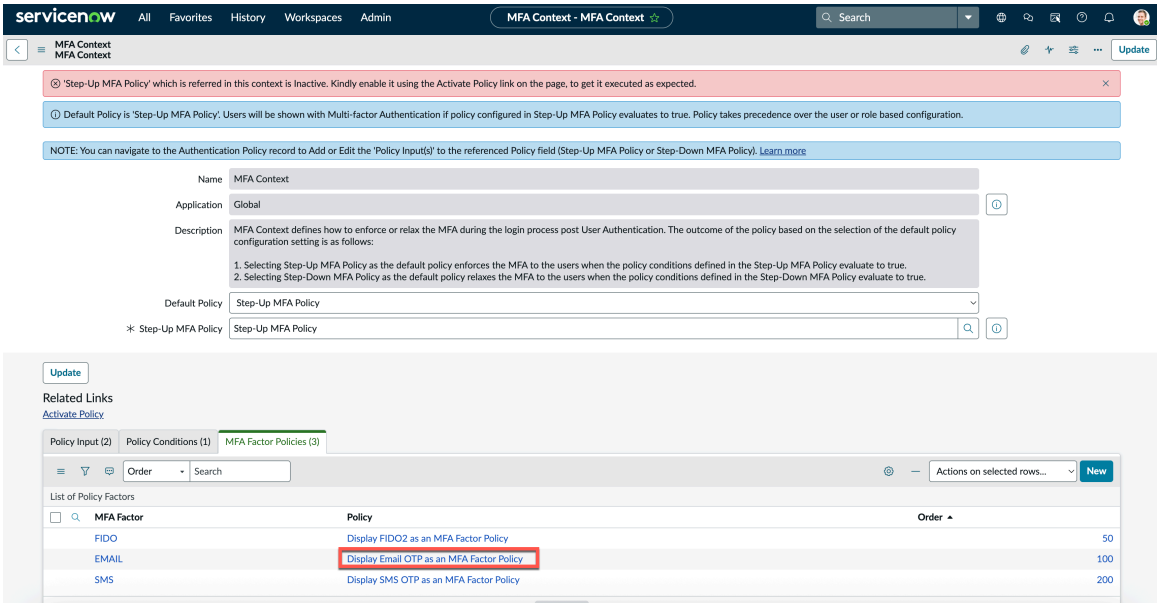

4. Click **New** to add **Policy Inputs**.
5. Select the filter criteria that you want to create.

   Following are the types of filter criteria:

   - IP Filter Criteria
   - Role Filter Criteria
   - Group Filter Criteria

   For example, Role Filter Criteria.

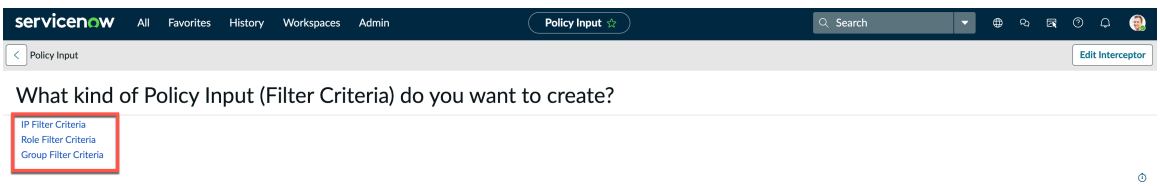

6. Click **Role Filter Criteria**, fill the fields for the role filter criteria and submit the record. The new policy is created. For more information, see Role Filter Criteria.
7. On the Policy - Display Email OTP as an MFA Factor Policy page, click **Policy conditions**.
8. Click **New** to add **Policy Conditions**.
9. On the form, fill in the fields:

Condition form

Field	Description
Label	Name to identify the condition.
Description	Description of the condition.

Field	Description
Condition	Logical combination of multiple policy inputs (filter criteria) that is used to evaluate authentication requests. Select the role based filter criteria policy that was created for the condition.

10. Click **Submit**.
11. **Optional:** Repeat step 8 to create additional policy conditions.

> **Note:** If you create multiple policy conditions, the final output of the access policy depends on the logical OR output of the all policy conditions. This means the policy will evaluate to true if any one of your policy conditions is met.

Based on the role filter (users) policy and the conditions specified for the role is matched, the Email MFA factor is shown as an options for authentication for the users.

Multi-factor Authentication system properties

Use system properties to enable and customize MFA to meet your security requirements.

Multi-factor Authentication properties

Property	Description
glide.authenticate.multifactor	Enable Multi-factor Authentication. **Note:** To enforce MFA for API authentication, set the `glide.authenticate.multifactor.for_integ` system property to `true`. MFA is enforced only for users who have already enrolled in MFA. Users who have not enrolled are not affected.
glide.authenticate.multifactor.setup.bypass.count	Number of times that a user can choose to skip the setup of MFA. The default is 0.
glide.multifactor.onetime.code.validity	Number of minutes that the reset code is valid. See Log in with Multi-factor Authentication. The default is 5. **Note:** This property is for email OTP validation time.
glide.authenticate.multifactor.clock_skew	Number of additional seconds that the reset code is valid. The maximum is 60. The default is `10`. The instance validates the code entered by the user against the single app-generated code that is generated at the `current time`. You can skew the time window with this property and allow one or more codes

Multi-factor Authentication properties (continued)

Property	Description
	to be generated during a time window to be considered valid. The property's value is used in the following calculation: `current time - X/2` and `current time + X/2`, where X is the value of this property. If you use the value of 10, for example, the instance considers any codes that the app generates within the time range `[the current time - 5 seconds]` and `[current time + 5 seconds]` to be valid. Use this property to prevent login issues where the user is unable to enter the correct code in the default time allotted.
glide.authenticate.multifactor.remember.browser.prompt	Set to true to prompt a user for MFA when they log in from a new device or browser. The default is yes.
glide.authenticate.multifactor.browser.duration	After MFA is completed in the browser, the user is not challenged for MFA in the same browser for this duration. The default is 8 hours.
glide.authenticate.multifactor.remembered.browsers.max	The number of browsers MFA remembers for this user.
glide.authenticate.multifactor.remember.browser.default	Default value of the remember browser check box in the validate multi-factor page.
glide.webauthn.enabled)	Option to enable passwordless authentication (FIDO2 based MFA) methods such as hardware key and biometric authentication.
glide.authenticate.multifactor.email.otp	Option to enable email based OTP as a factor for MFA.
glide.auth.mfa.ui.v2.enabled	Option to enable MFA factor independently for the users without setting up an authenticator app.

Multi-factor Authentication criteria

Use MFA criteria to determine which users and roles must use two-step verification.

Multi-factor criteria

Use MFA criteria to determine which users and roles must use two-step verification. You can use one of these criteria or a combination of them to suit your business needs. You can use one of these criteria or a combination of them to suit your business needs.

Note: It is recommended to use Adaptive Authentication policy based MFA.

User-based multi-factor criteria

Use user-based multi-factor criteria to select individual users who are required to log in using MFA. Administrators update the **Enable Multifactor Authentication** field on

a user record to enable or disable MFA requirements for a user. For details on this process, see Configure user-based multi-factor criteria.

Role-based multi-factor criteria

Use role-based multi-factor criteria to require MFA login for all users assigned to a specific role. The **Role-based multi-factor authentication** record on the **Multi-factor Criteria** [multi_factor_criteria] table contains the list of roles that require an MFA login. For details on maintaining this list, see Configure role-based multi-factor criteria.

Adaptive authentication policy-based multi-factor criteria

Use adaptive authentication to determine when your instance requires MFA. Adaptive authentication uses authentication policies to evaluate criteria like a user's IP address or user groups. For details on the adaptive authentication feature, see Adaptive authentication.

Configure user-based multi-factor criteria

Use user based multi-factor criteria to enable MFA for a user.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **User Administration** > **Users**.
2. Configure the list to show the **Enable Multi-factor Authentication** column.
3. Change the values of the **Enable Multi-factor Authentication** column for the selected users to **true**.

   When the user logs in with their user name and password, they are prompted to set up multi-factor authentication.

Configure role-based multi-factor criteria

Use role based multi-factor criteria to enforce Multi-factor authentication for all users assigned to specific roles.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Multi-factor Authentication** > **Multi-factor Criteria**.
2. In the **Multi-factor Criteria** list, open the **Role-based multi-factor authentication** record.
3. Use the **Multi-factor Roles** list to add or remove roles.

Option	Description
Add a role	Double-click **Insert a new row...** and enter or select a role name. Click the **Save Icon** ( ) to save the entry.
Remove a role	Click the delete icon ( ) to remove a role from the list.

4. Click **Update**.

Result

Your instance enforces multi-factor authentication for all users who are members of the roles listed in the **Multi-factor Roles** list.

Important: The record must be active to enforce role-based multi-factor authentication.

Configure adaptive authentication policy-based multi-factor criteria

Use adaptive policies to determine which users must use two-step multi-factor (MFA) verification.

Before you begin

Role required: admin

Note:

- If the default policy is **Step-Up MFA Policy**, users will be shown with Multi-factor Authentication if policy configured in **Step-Up MFA Policy** evaluates to true. Policy takes precedence over the user or role based configuration.
- MFA with SSO login will only be available if `glide.authenticate.mfa.with.multisso.enabled` Property is set to true.
- You can navigate to the Authentication Policy record to Add or Edit the 'Policy Input(s)' to the referenced Policy field (**Step-Up MFA Policy** or **Step-Down MFA Policy**).

Procedure

1. Navigate to **All** > **Adaptive Authentication** > **Auth Policy Contexts** > **MFA Context**.
   The **MFA Context** policy context record opens.
2. Select the default policy in the **Default Policy** field.
   This selection determines how your instance uses the policy conditions to determine whether to require MFA.

Default policies

Default Policy	Definition
Step-Up MFA Policy	Select to enforce MFA when the policy conditions defined in **Step-Up MFA Policy** evaluate to true.
Step-Down MFA Policy	Select to bypass MFA when the policy conditions defined in the **Step-Down MFA Policy** evaluate to true.

3. In the **Step-Up MFA Policy** field, select a policy to use with this context.
4. Click **Update**.
   After saving the record, the **Policy Input** and **Policy Conditions** lists update to display the policy inputs and conditions associated with the policy selected in the **Step-Up MFA Policy** field.

Multi-factor Authentication with Single Sign-On

You can use MFA with an SSO provider for your ServiceNow instance.

Since many users in your organization are logging in to your corporate network and accessing your ServiceNow instance, there is a stronger need for a secure authentication.

MFA with the combination of SSO provides an enhanced security for your instance. This capability provides flexibility to conditionally enable MFA for users.

Note:

- ServiceNow MFA is enforced after the user is redirected after the successful authentication at the Identity Provider.
- MFA was not enforced with SSO prior to the San Diego release.

For example, if you would like to give your external users an extra security protocol, then you can enforce MFA to only those users. In this way, you can add additional authentication abilities and control your user access.

For SSO-based login such as SAML, OpenID Connect, and Digest, you can enforce MFA for authentication.

MFA with SSO can be configured on-demand based on your requirement. With the help of authentication schemes and identity providers, you can enforce MFA for specific users with a specific login mechanism.

You can enforce MFA for the following conditions:

- Authentication Scheme
- Identity Provider

MFA with SSO is offered as a part of the Adaptive Authentication plugin (com.snc.adaptive_authentication). To know more on how to set up Adaptive Authentication, see Adaptive authentication.

Configuring MFA with SSO

Enforce MFA with SSO for your users within or outside your organization.

Before you begin

The MFA with SSO feature is offered as a part of the Adaptive Authentication plugin (com.snc.adaptive_authentication). You must enable the Adaptive authentication property for using MFA with SSO feature. To know more on how to set up Adaptive authentication, see Adaptive authentication.

Note: MFA with SSO login is available if the *glide.authenticate.mfa.with.multisso.enabled* property is set to **true**.

Role required: admin

Procedure

1. Navigate to **All** > **Multi-factor Authentication** > **Properties**.
2. Enable the **Enable Multi-factor authentication** and **Enable Multi-factor Authentication with SSO** check boxes.

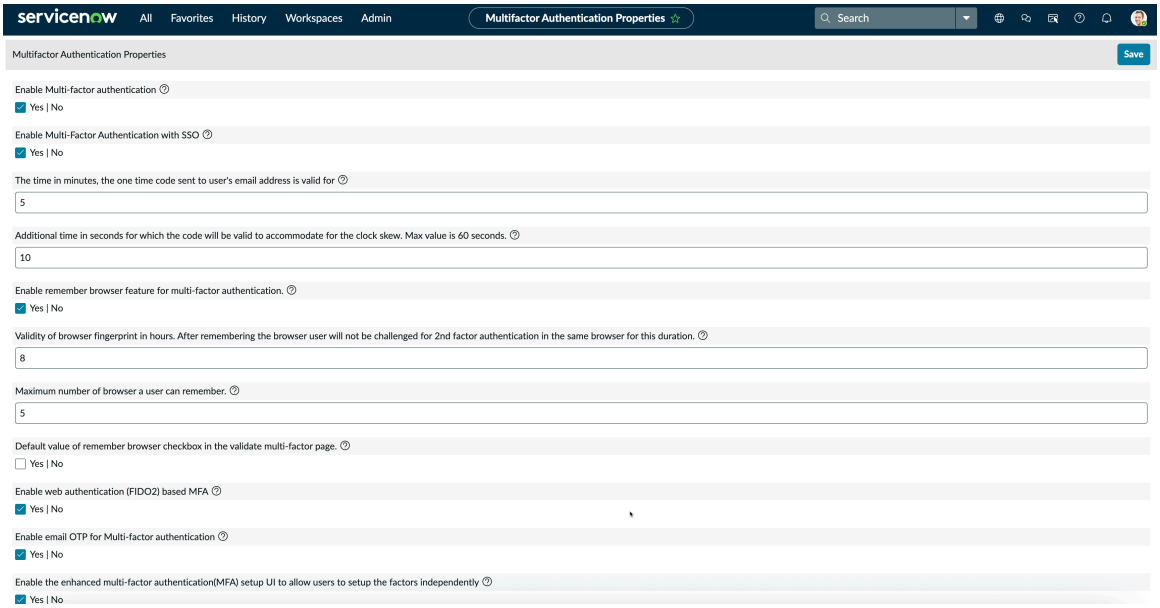

3. Click **Save**.
4. Navigate to **Multi-factor Authentication** > **MFA Context**.
   The MFA Context form appears.

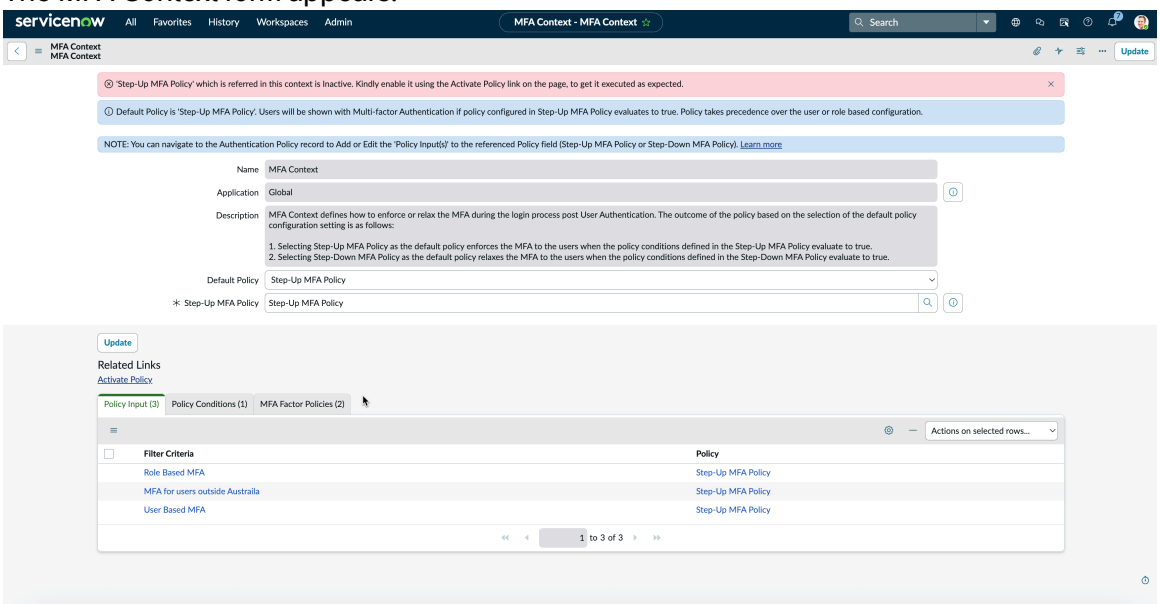

> **Note:** By default, the policy is **Step-Up MFA Policy**. Users are shown with Multi-factor Authentication if a condition configured in **Step-Up MFA Policy** evaluates to **true**. Policy takes precedence over the user or role-based configuration.

5. **Optional:** To edit a policy, go back to the Authentication Policy record, change the conditions, and come back.
   When you navigate to the Authentication Policy record, you can add or edit the 'Policy Inputs' to the referenced Policy fields, **Step-Up MFA Policy** or **Step-Down MFA Policy**.
6. To set a new policy condition, click **Policy Conditions**.
7. Click **New**.
8. On the form, fill in the fields.

Conditions form

Fields	Description
Label	Unique name for the condition that you want to create for the label.
Description	The policy condition description.
Conditions	The type of condition that you want to apply for the policy. You can add filter conditions and 'OR' clauses. **Note:** Adding various filter conditions and clauses enables you to challenge MFA for specific users.

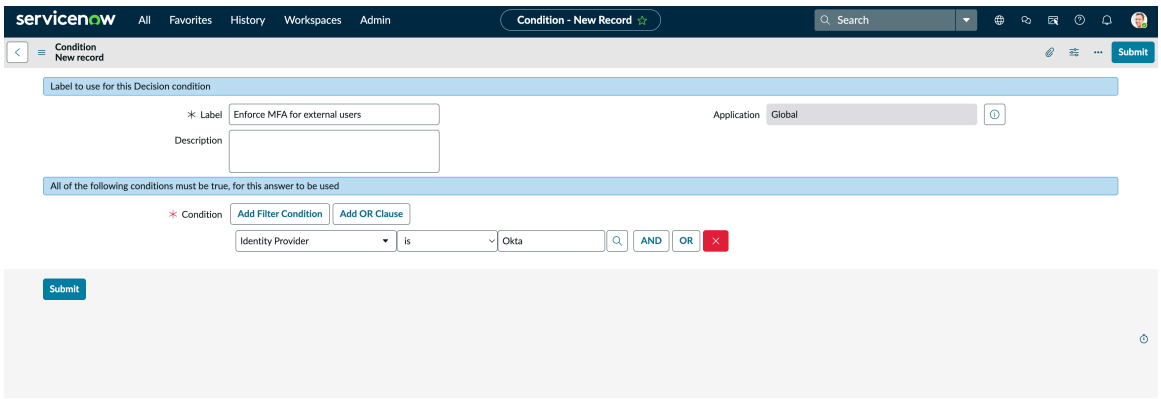

Example

To set up your conditions, consider the following example. Assume that you want the authentication scheme with the identity provider as your condition for external users. You can set the conditions as follows:

a. Select **Authentication Scheme** and set the criteria as **Single Sign-On** or **Username and Password**.

   Based on this selection, the user is provided with an SSO-based login or login form to specify the user name and password.

b. Select **Identity Provider** and specify the provider as the IDP record for which you want to enable MFA. For example, **Okta**.

Based on this selection, if the user logs in with SSO, the user is not challenged with MFA. In contrast, if the user logs in using Okta, then the user is challenged with MFA.

9. Click **Submit**.

Related topics

Adaptive authentication

Multi-factor Authentication with Single Sign-On

Reset Multi-factor Authentication (MFA) for users

Administrators can reset MFA for users who deleted the app, lost access to the device, or have no alternative MFA associated with their device.

Before you begin

Role required: admin

The following procedure describes how a ServiceNow® administrator can reset the MFA validation to unblock users and enable them to re-register MFA.

Procedure

1. Navigate to **All** > **Multi-factor Authentication** > **User Multi-factor Setup**.
2. Search for the user that you want to unblock.
3. Set the Validate to **false**.
4. Navigate to **All** > **Multi-factor Authentication** > **Web Authentication** > **User Public Credentials (sys_user_public_credential**.
5. Search for the user that you want to unblock.
6. Delete all the records for this user.
7. Navigate to **All** > **Multi-factor Authentication** > **User Recent Used Factors**.
8. Search for the same user.
9. Delete all the records for this user.

Result

When the unblocked user enters the credentials and logs in, the **Enable multi-factor authentication (MFA)** page is displayed. The user can follow the steps on the page to re-register MFA.

Reference topic - Multi-factor Authentication

Reference topic related to the configuration of MFA.

User Multi-factor Authentications

To access the details about the user's MFA, navigate to **All** > **Multi-factor Authentication** > **User Multi-factor Setup**.

User Multi-factor Authentications

Field	Description
User	The username of the user
Bypasses remaining	The total number of bypasses that is remaining for the user.
Multi-factor secret	The details of the multi-factor secret.
Validated	Is the multi-factor validated. Changing the value to **false** invalidates the existing authenticator app setup.

Note: The recently used MFA factors by the users is available in the User Recent Used Factors module. Navigate to **All** > **Multi-factor Authentication** > **User Recent User Factors**. For more information, see User Recent Used Factors.

Multi-factor Browser Fingerprints

To access the details about the user's browser fingerprints, navigate to **All** > **Multi-factor Authentication** > **Multi-factor Browser Fingerprints**.

Multi-factor Browser Fingerprints

Field	Description
User	The username of the user
Browser	The browser that the user used.
Browser Fingerprint	The browser fingerprint details.
Browser Fingerprint Cookie	The browser fingerprint cookie details.
Browser Fingerprint Cookie Expiry	The browser fingerprint cookie expiry details.

User Recent Used Factors

To access the details about the user's recently used factors, navigate to **All** > **Multi-factor Authentication** > **User Recent User Factors**.

User Multi Factor Setups

Field	Description
User	The username of the user
Multi Factor Type	The multi-factor type the user used for log in.
Is Recent Factor	The user's recent multi-factor detials.

MFA Metrics

View the different MFA metrics to understand the MFA adoption and usage.

MFA related metrics are available in Security Center. Security Center is a free application that you can download from the ServiceNow Store. For more information about Security Center, see Security Center landing page.

On Security Center, use the bar at the top of the page to navigate between the security monitoring console sections and select **Security Metrics** tab.

Following are the MFA metrics:

- Users enrolled for MFA: Total count of users enrolled to use MFA.
- Users using MFA bypass: Total count of users who are circumventing multi factor authentication.
- High privileged non-mfa user: Total count of high-privileged users that are not using MFA.
- Active MFA users: Total count of MFA users that are active on your instance.
- Locked out MFA users: Total count of MFA users who are locked out on your instance.
- Local logins not protected by MFA: Users that logged in without MFA.

For more information about Security Metrics, see Security metrics.

Using Multi-factor authentication

Learn how to use multi-factor authentication tools to securely access your instance.

Login with MFA

ServiceNow requires authenticator applications that support Time-based One-time Passwords (TOTP). ServiceNow tests MFA with the following authenticators:

- Google Authenticator
- Microsoft Authenticator
- LastPass Authenticator
- Authy
- FreeOTP
- Duo
- Okta Verify

Note:

- Other authenticators not listed might also be compatible, but are not tested by ServiceNow.
- For information related to browser specific behavior change, see this KB article.

Register an authentication app

Authenticator App You can use authenticator apps to use a second factor of authentication.	If your administrator has enabled multi-factor authentication (MFA) on your instance, you are prompted for a second authentication after entering your user name and password. For details on the MFA login process, see Log in with Multi-factor Authentication.
Validation with Authenticator app Enter the code displayed on your authenticator app to login.	If you haven't configured a second form of authentication, you will see a configuration page after logging in to guide you through the process of setting up an authentication app. For details on this setup, see Set up Multi-factor authentication for the first time.

Register an authentication device

After you've configured an authentication app, you can register other methods for authentication.

Biometric authenticators

You can use biometric authenticators like fingerprint or facial recognition as your second MFA authentication. If your administrator allows this option, you can configure biometric authenticators using the steps in Register a biometric authenticator.

Hardware key authenticators

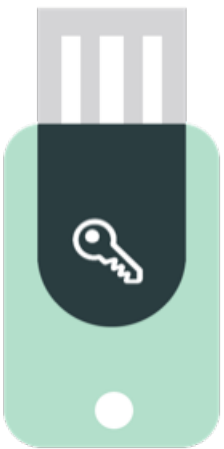

Hardware keys are physical security devices you can use for authentication. You can register a hardware device for use with your instance using the steps in Register a hardware security key.

Validation with Biometric or Hardware Key Use the biometric or hardware key to login.	Use the Biometric or Security Key to login. 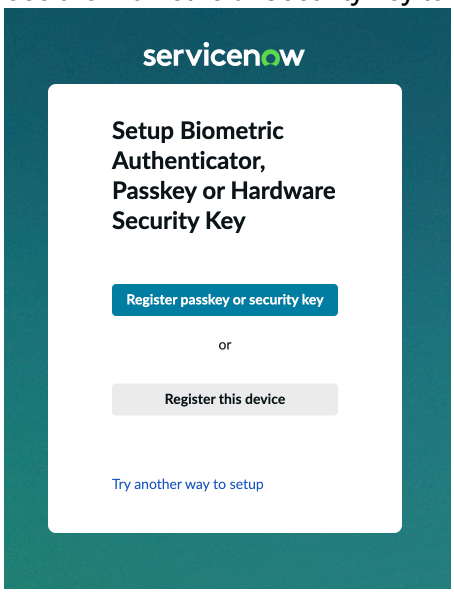

Register a phone number for OTP

SMS Use SMS based OTP to require users who attempt to login.	When users attempt to login to ServiceNow, SMS OTP is sent to the mobile number associated with the sys_user record. Users can enter the six-digit verification code that it sent to the mobile device and verify their identity. 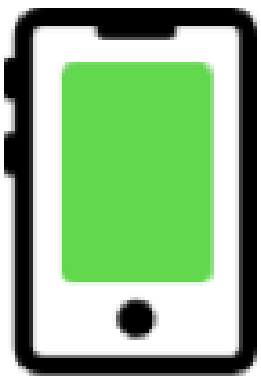
Validation with SMS Use the validation with SMS to login based on the OTP generated.	You need to enter the 6-digit code sent to the mobile number to login. The code sent is valid for the next 5 minutes. You can use resend code to again send the code.

Register an Email address for OTP

Email address Use Email based OTP to require users who attempt to login.	When users attempt to login to ServiceNow, Email OTP is sent to the email address associated to the user. User's can enter the six-

digit verification code that it sent to the mobile device and verify their identity.

Validation with Email Use the validation with Email to login based on the OTP generated.	You need to enter the 6-digit code sent to the email address to login. The code sent is valid for the next 5 minutes. You can use resend code to again send the code. 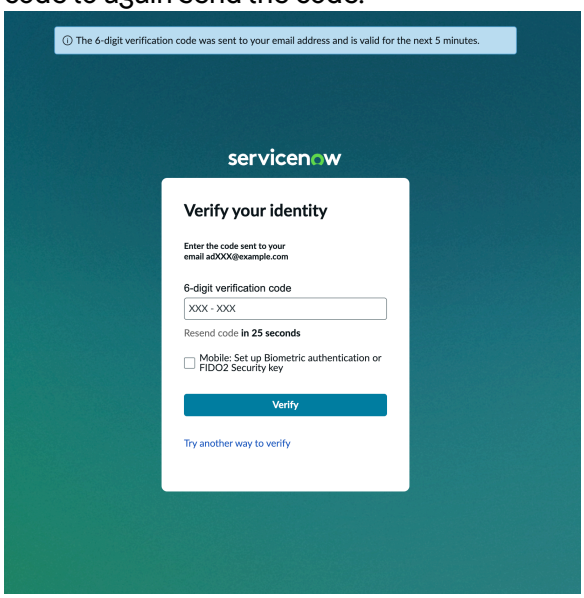

Set up Multi-factor authentication for the first time

If your administrator enabled MFA on your profile but you have not yet set up the application, you can set it up upon login.

Before you begin

Role required: none

Procedure

1. Log in to your instance using your user name and password.

   The multi-factor authentication set up screen intercepts your login.

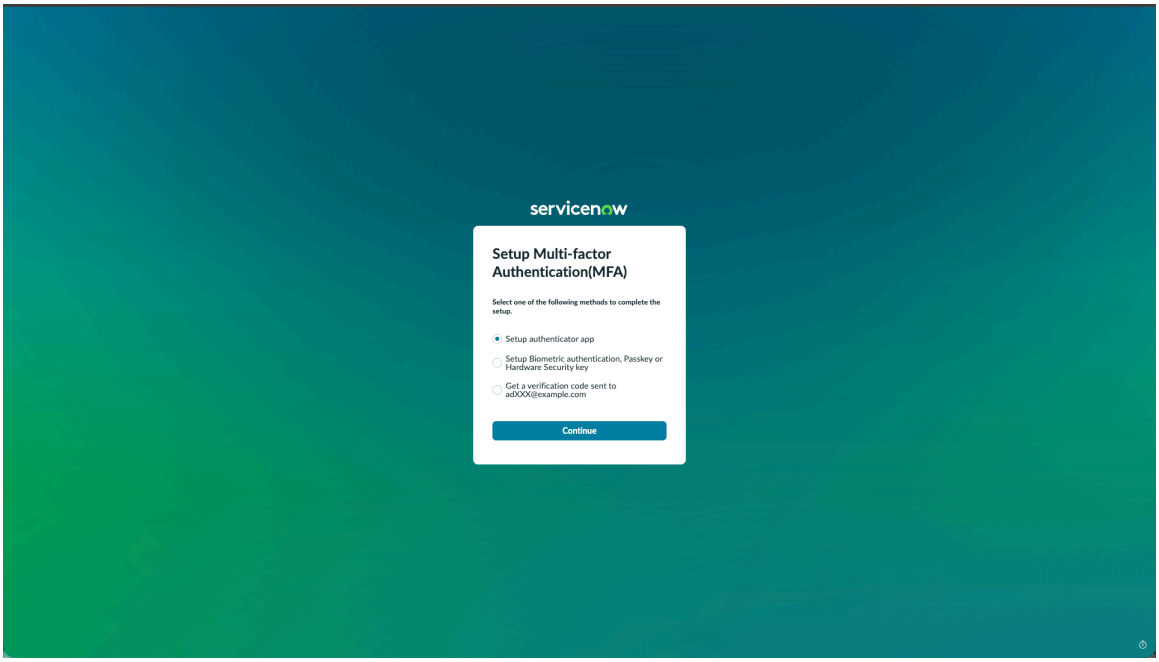

Note: If you want to skip the authentication set up now, click **Bypass Setup**. You can bypass multi-factor authentication for a limited number of times that your administrator allows. Eventually, you must configure multi-factor authentication.

2. Select one of the following methods to complete the mfa setup.

   a. **Setup authenticator app**

      Follow the instructions on the screen to pair device and login.

b. Setup Biometric authentication, Passkey or Hardware Security key

Select either of the option to complete the setup.

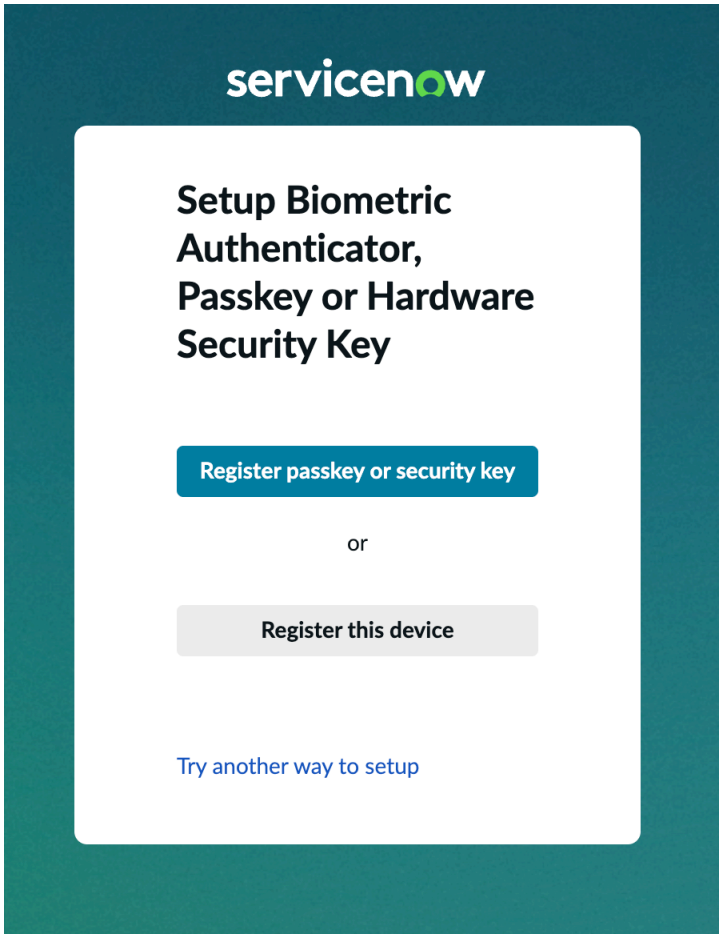

c. Get a verification code sent to your email

Enter the verification code that is sent to your email.

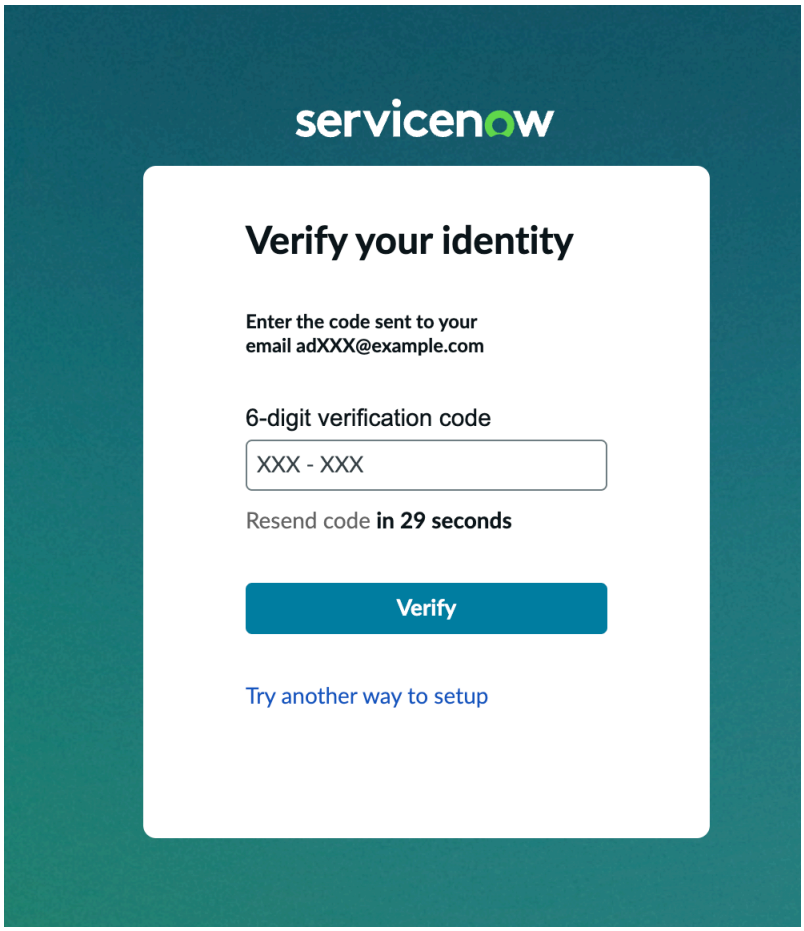

After the successful completion of the either of the setup, you are logged in to the instance.

Set up Multi-factor authentication on your user profile

Enable multi-factor authentication for your account in your user profile settings.

Before you begin
Role required: none

Multi-factor authentication must be enabled on your instance.

Note: Your administrator may require that you use multi-factor authentication. In this case, you are automatically prompted when you log in. See Set up Multi-factor authentication for the first time. Use the process below if your administrator allows you to opt-in to multi-factor authentication.

Procedure

1. Navigate to **All** > **Self-Service** > **My Profile**.

   Note: You can also access your profile by clicking on your user name in the instance header.

2. In your user profile, click **Configure Multi-factor Authentication** in the **Related Links** section.
3. Choose the MFA Authenticator type that you would like to use.

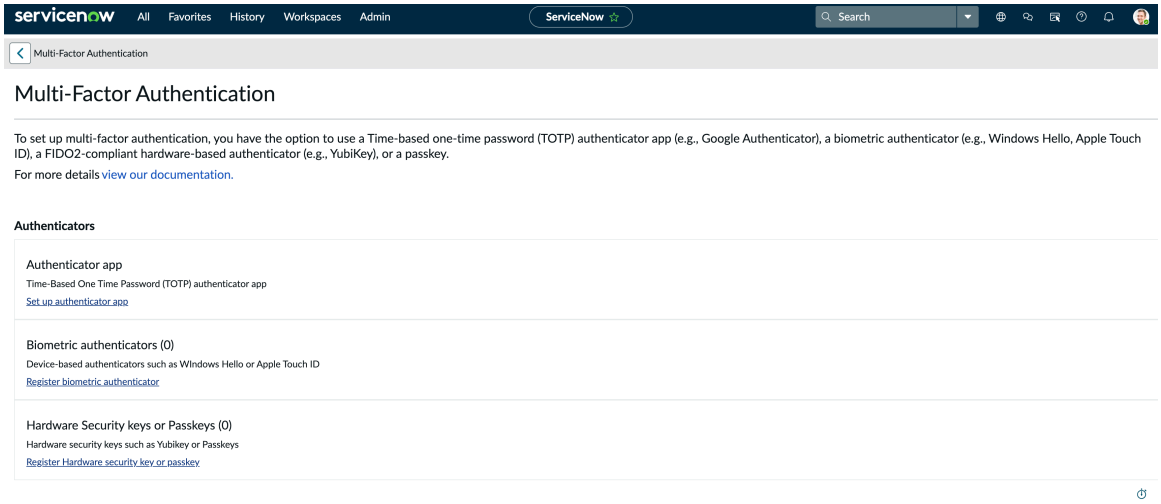

MFA Authenticators

MFA Authenticators	Description
Authenticator App	Time-Based One Time Password (TOTP) authenticator app Select **Set up authenticator app** and follow the instructions on the screen to register for an authenticator app as the second factor for authentication. 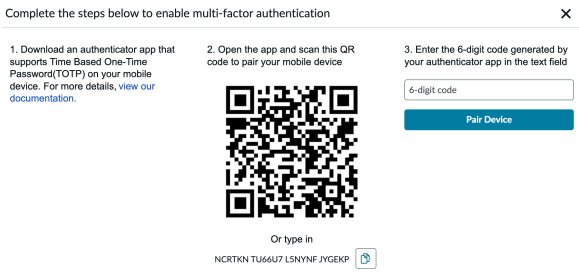
Biometric authenticators	Device-based authenticators such as Windows Hello or Apple Touch ID Select **Register biometric authentication** and follow the instructions on the screen to register for an authenticator app as the second factor for authentication. 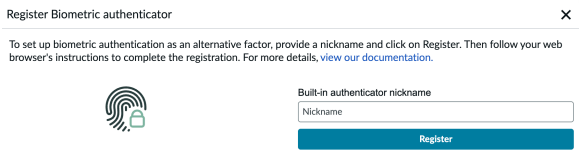
Hardware Security Keys	Hardware security keys such as YubiKey Select **Register hardware security keys** and follow the instructions on the screen

MFA Authenticators	Description
	to register for an authenticator app as the second factor for authentication. 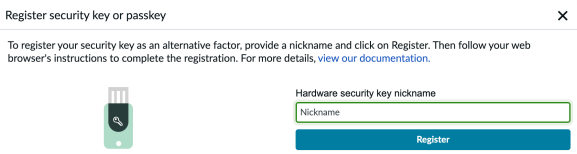

Result

Multi-factor authentication is enabled for your user account. You will be prompted to use multi-factor authentication the next time you log in.

Note: You cannot disable multi-factor authentication on your account after you have enabled it. If you need to disable multi-factor authentication, contact your administrator.

Log in with Multi-factor Authentication

Login with MFA when it is enabled by your administrator on your instance.

Before you begin

You must have MFA enabled for your profile. You can enable it yourself on your user profile or your administrator can enable it for you.
Role required: none

Note: The last used MFA verification factor is automatically shown for the users in the subsequent log in.

Procedure

1. Go to the URL of your instance to open the login screen.
2. Enter your user name and password.
3. Click **Log in**.
   The multi-factor authentication screen appears.

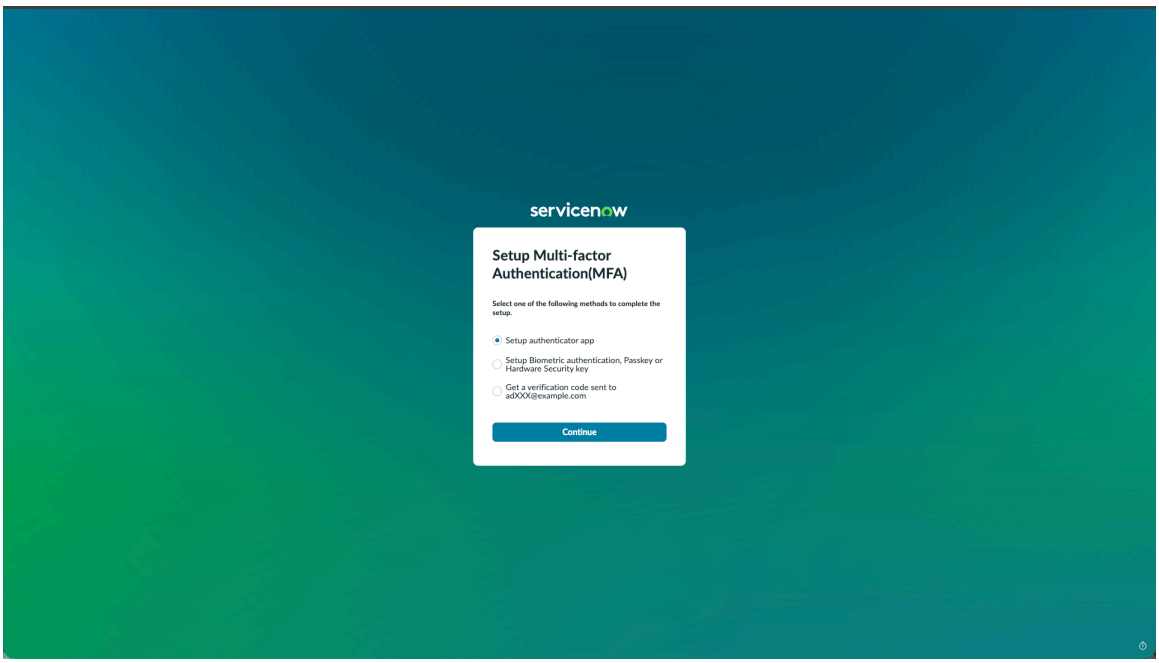

4. Select one of the methods to complete the MFA setup.
   For more information about each setup, see Set up Multi-factor authentication for the first time.

   If you wish to postpone the setup, select **Postpone setup**. The maximum number of time you can postpone the setup is displayed on the screen and this value is configured by your administrator.

Result
If you have already setup MFA, the recently used factor is displayed when you login your instance after entering your username and password.

Example: Recently used MFA factor

If the recent used factor was Biometric login, then upon login using your username and password, you're directly displayed with the Biometric MFA setup screen.

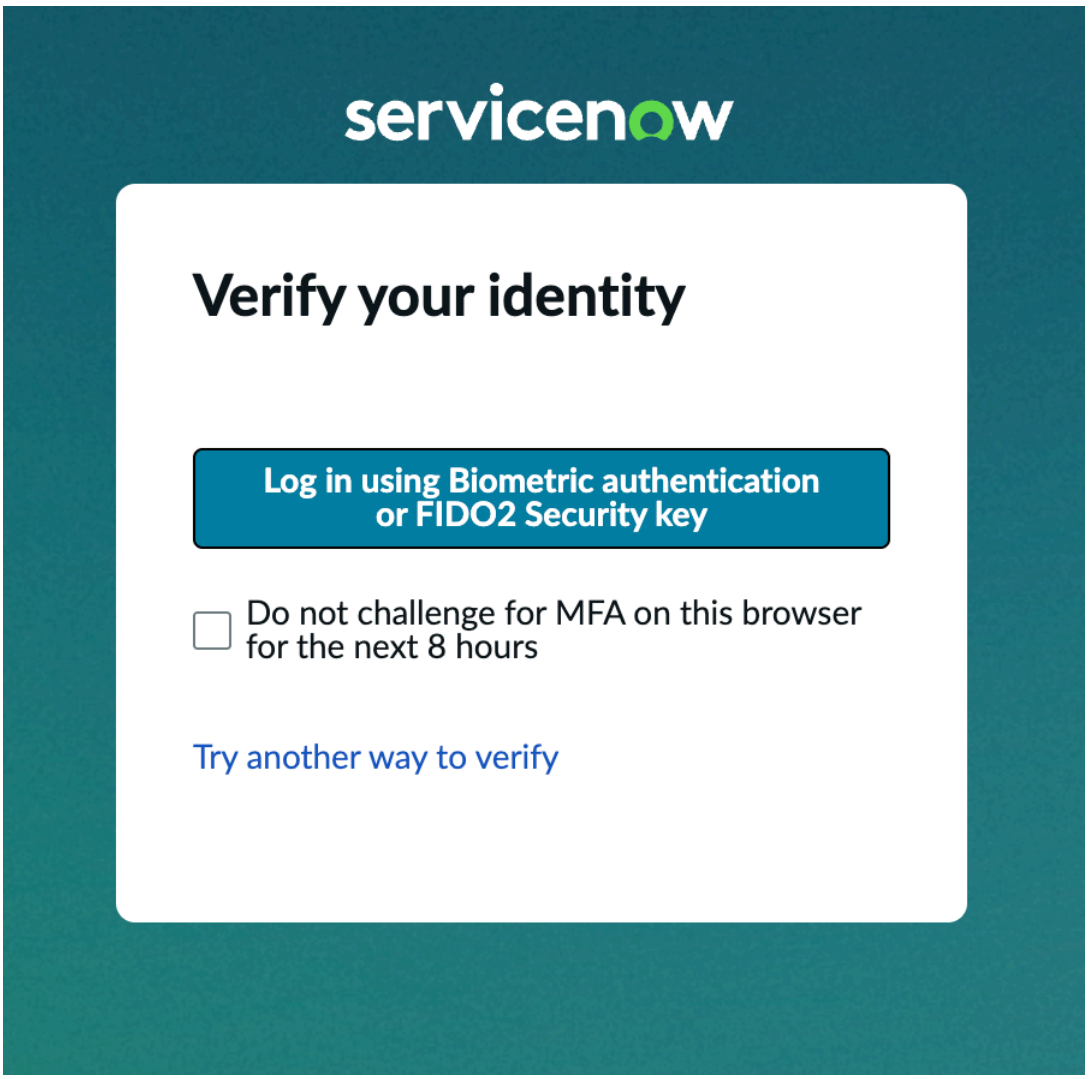

Authenticator Applications

Use third party authenticator applications to generate temporary MFA pass codes.

An authenticator application is third-party software that generates temporary pass-codes. You can use these pass-codes along with your password to login into an instance that requires multi-factor authentication (MFA).

If your administrator has enabled MFA on your instance, you see a prompt for a pass-code after entering your user and password during login.

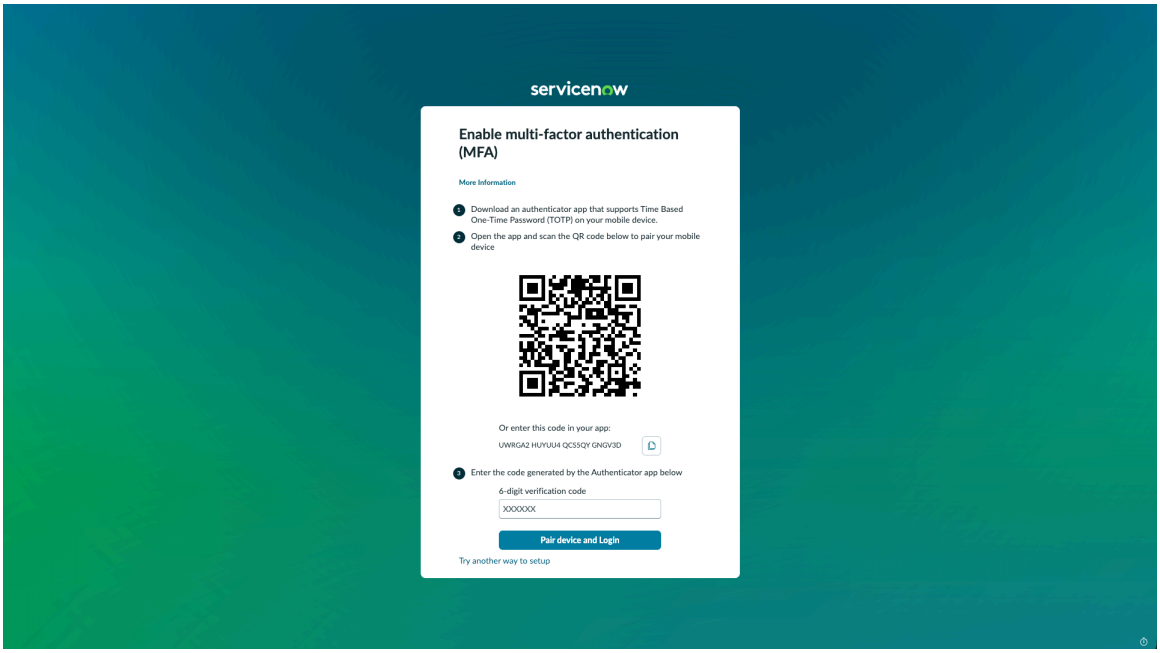

ServiceNow requires authenticator applications that support Time-based One-time Passwords (TOTP).

ServiceNow tests MFA with the following authenticators:

- Google Authenticator
- Microsoft Authenticator
- LastPass Authenticator
- Authy
- FreeOTP
- Duo
- Okta Verify

Note: Other authenticators not listed might also be compatible, but are not tested by ServiceNow.

Change an Authenticator app

Generate a new code to change an Authenticator app on your device.

Before you begin

You must have MFA enabled for your profile. You can enable it yourself on your user profile or your administrator can enable it for you.
Role required: none

Procedure

1. Navigate to **All** > **Self-Service** > **My Profile**.

Note: You can also access your profile by clicking on your user name in the instance header.

2. In your user profile, click **Muti-factor Authentication** in the **Related Links** section.
3. Under **Authenticator app**, select the edit icon.

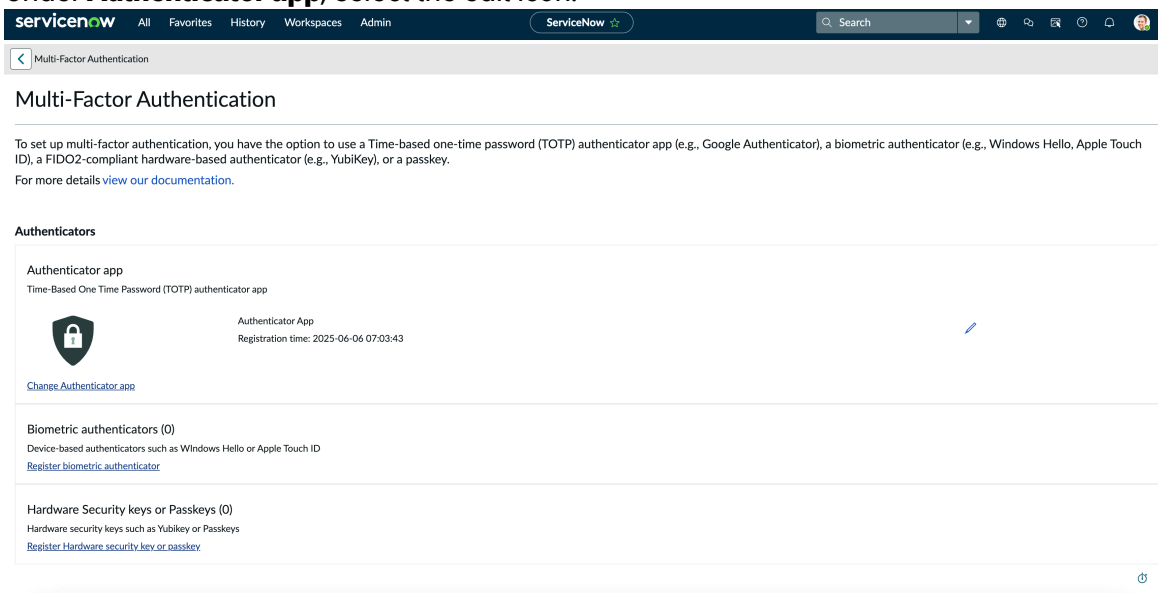

4. In the Change authenticator code window, click **Change**.
5. Follow the directions on the screen to register the authenticator app with the device.

Web Authentication

Your users can use hardware keys or their device's biometric readers (FIDO2) to authenticate to an instance.

Hardware keys

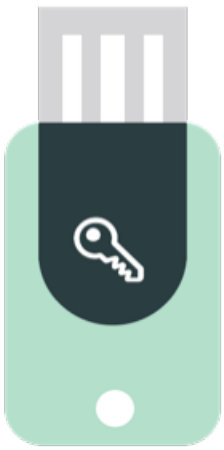

Hardware keys are physical hardware that you can use to authenticate. Hardware keys are inserted into a port on your device to provide authentication. For details on registering hardware keys, see Register a hardware security key.

Biometrics

Biometric authenticators use fingerprint or facial recognition to identify users. Your users can use these authenticators on their devices as part of the multi-factor login process. For details on registering biometric authenticators, see Register a biometric authenticator.

Choose the second factor that you wish to authentication and authenticate to your instance.

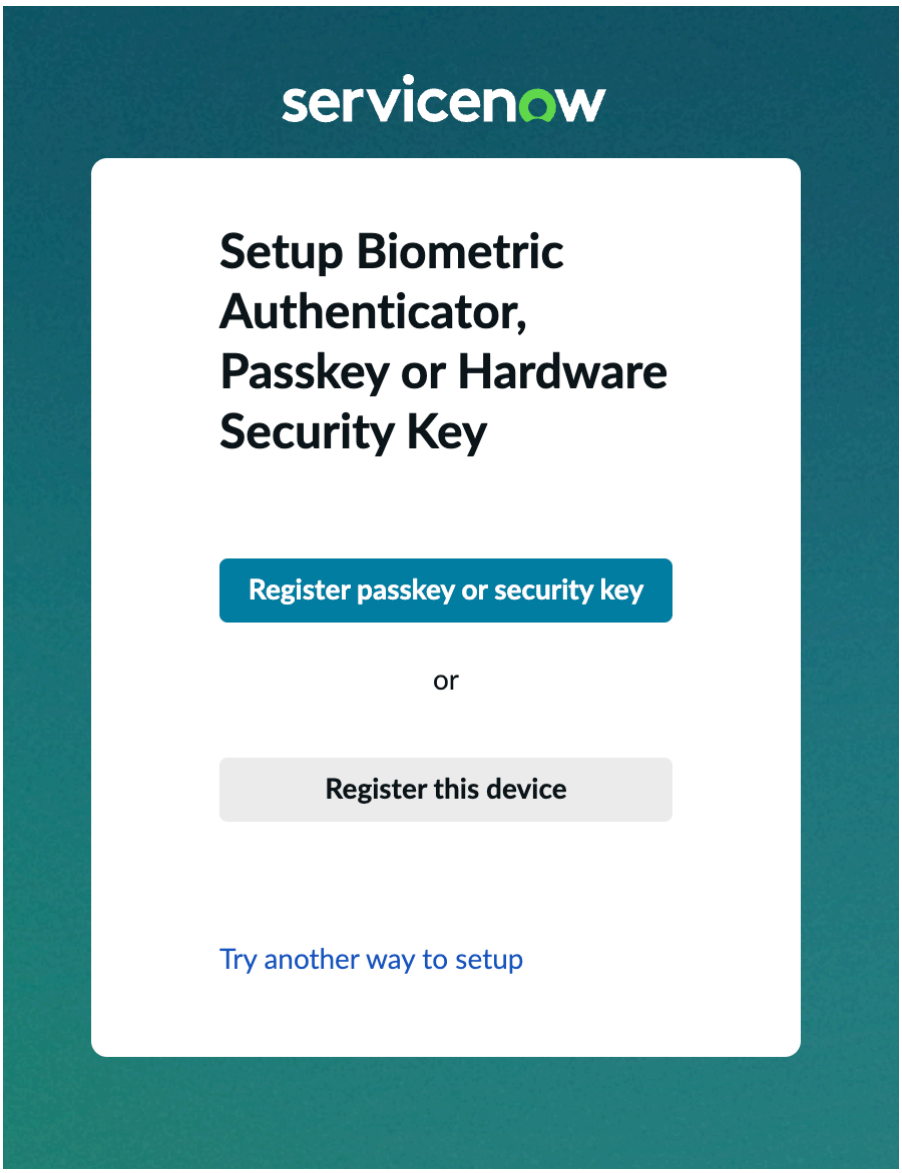

To configure the web authentication plugin, see .

Register a biometric authenticator

Register a biometric authenticator to use as part of your MFA login.

Before you begin

Role required: none

Procedure

1. Navigate to **All** > **Self-Service** > **My Profile**.
2. Under **Related Links**, click **Multi-factor Authentication**.
   The multi-factor authentication page opens.
3. Click **Register biometric authentication**.

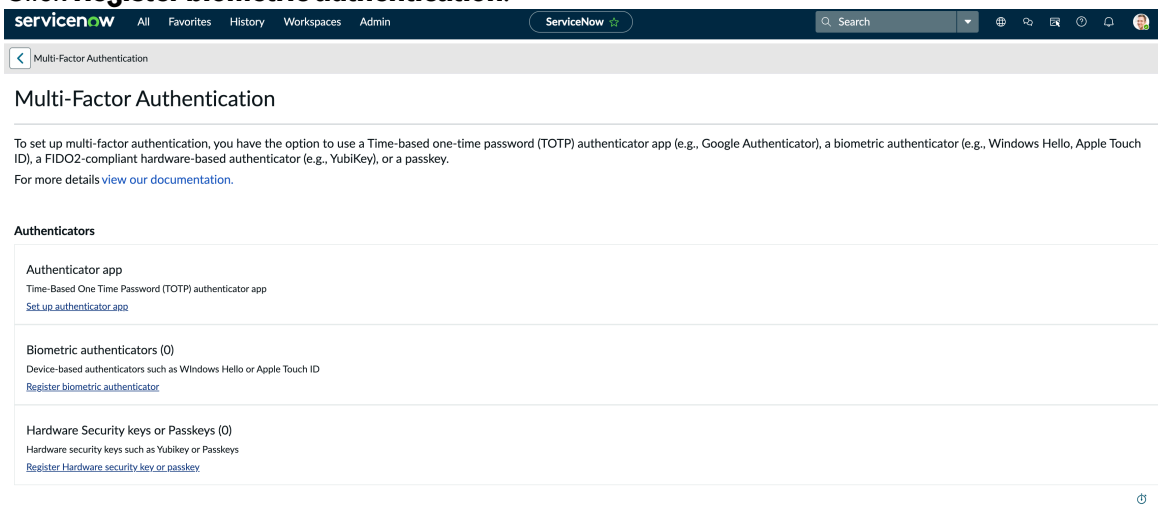

4. Enter a nickname for your authenticator and click **Register**.

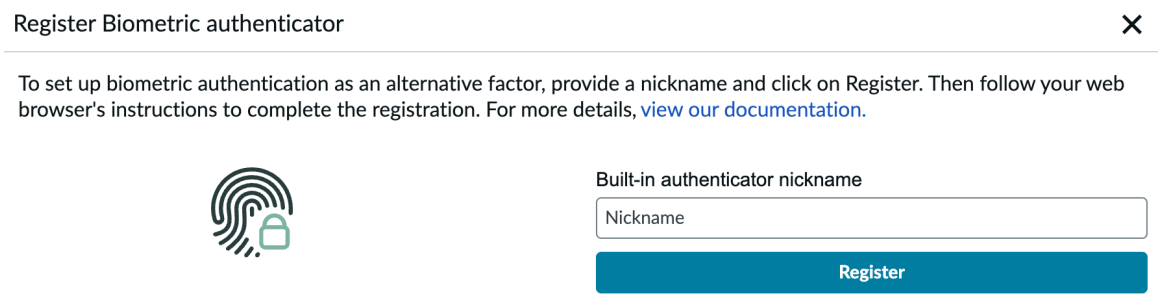

5. When prompted, follow the instructions on the screen to authenticate with your biometric authenticator.
   This message will vary based on your specific authenticator.

Example of a biometric authentication request

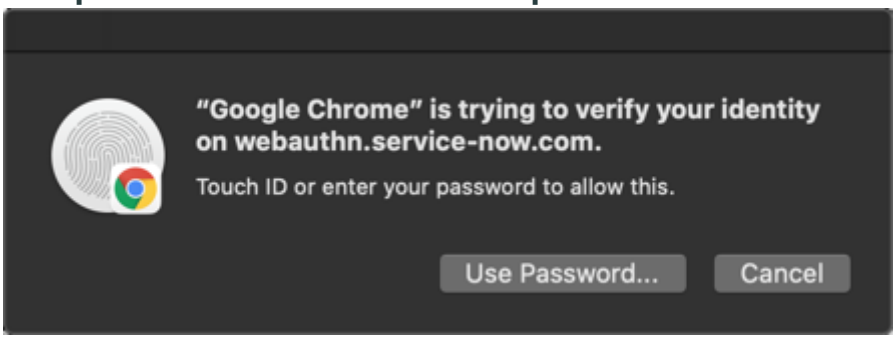

After authenticating successfully, you see a confirmation window. Click X to close the confirmation.

Result

Your biometric authenticator is registered. You can see the biometric authenticator listed on the Multi-factor authentication page.

Register a hardware security key

Register a hardware key to use as part of your MFA login.

Before you begin

Role required: none

Procedure

1. Navigate to **All** > **Self-Service** > **My Profile**.
2. Under **Related Links**, click **Multi-factor Authentication**.
   The Multi-factor authentication page opens.
3. Click **Register hardware security Key**.

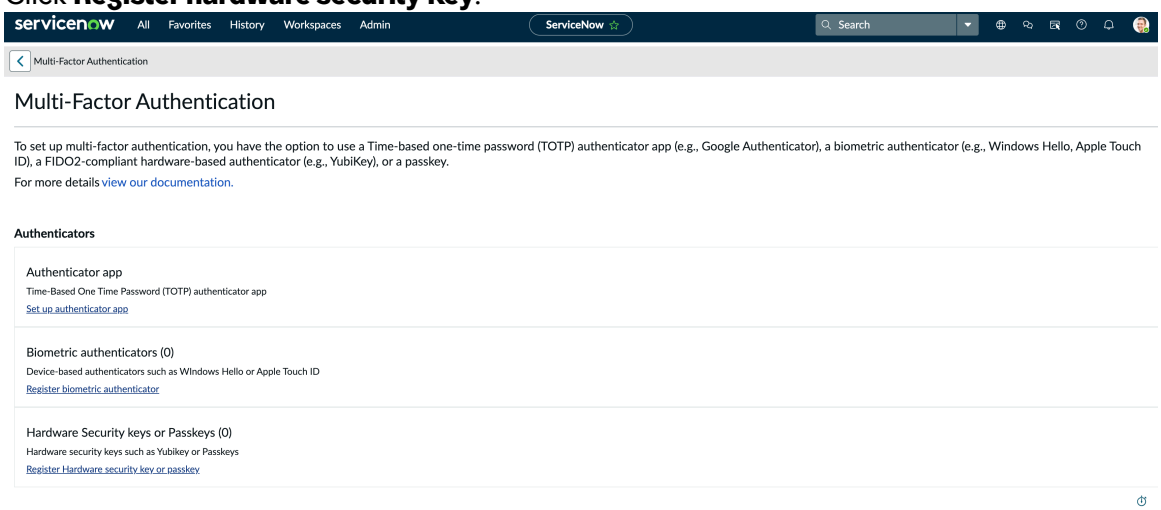

4. Enter a nickname for your hardware key and click **Register**.

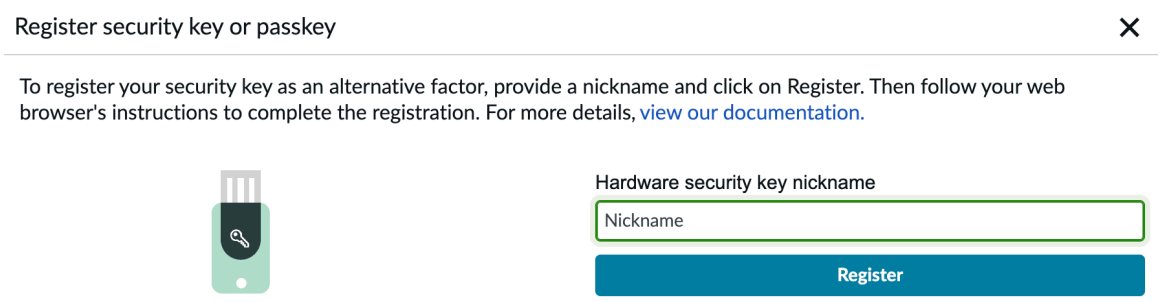

5. When prompted, insert your hardware security key and activate it.
   After authenticating successfully, you see a confirmation window. Click X to close the confirmation.

Result
Your hardware key is registered. You can see the hardware key listed in the Multi-factor authentication page.

MFA Dashboard

View the different MFA metrics to understand the MFA adoption and usage.

The MFA enhances security by requiring an additional verification step during the login process. The MFA Dashboard helps you to monitor and manage MFA settings for your organization. This dashboard provides an overview of MFA user enrollment, privileged admins who haven't opted MFA, and compliance. You can use it to ensure that all users have MFA enabled for enhanced security.

To access the MFA Dashboard, navigate **All** > **Multi-factor Authentication** > **MFA Dashboard**.

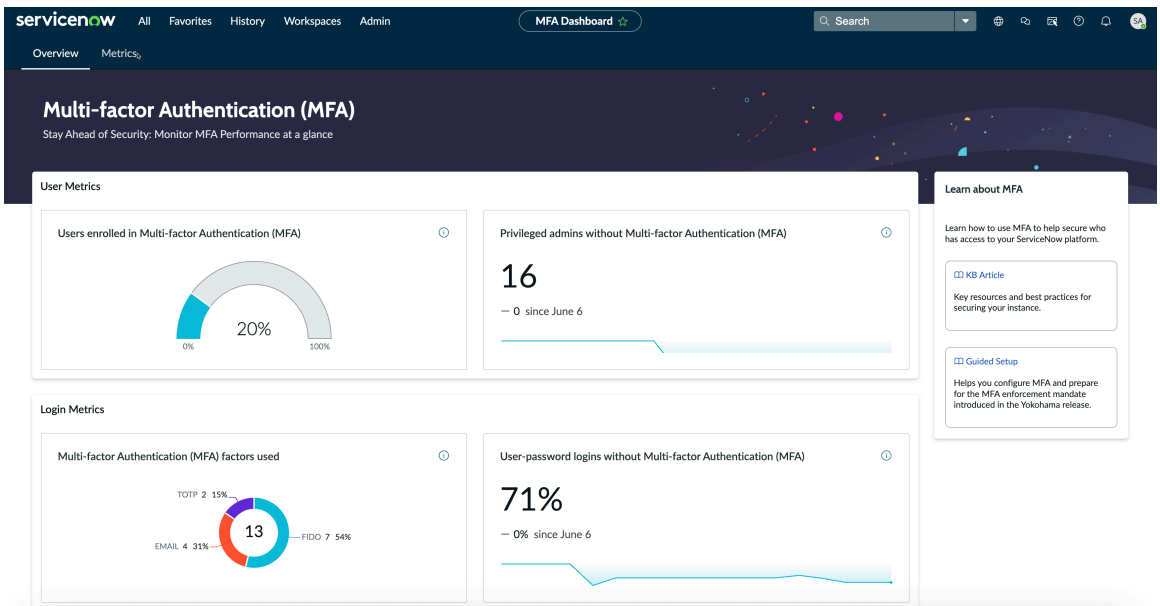

Note: You must enable MFA to view the metrics. For more information, see Multi-factor Authentication system properties.

Following are the types of metrics displayed on the MFA Dashboard:

- User Metrics
- Login Metrics

User Metrics

The following table provides details of the User Metrics in the MFA Dashboard.

User Metrics

Metrics	Description
Users enrolled in Multi-factor Authentication (MFA)	The percentage of users who can perform username-password based login and enrolled in MFA. This metric provides an insight on the adoption of MFA by the users over a period of time. **Note:** Ideally, the score should gradually increase and should be 100% over a period of time (Refreshed once a day to collect records for a day before).
Privileged admins without Multi-factor Authentication (MFA)	Privileged admins not using MFA is a significant risk to platform security. It's recommended that you get these people using MFA. **Note:** Privileged admins are the users who have at least one role from the sys_icenter_role_config table. (Refreshed once a day to collect records for a day before).

Login Metrics

The following table provides details of the Login Metrics in the MFA Dashboard.

Login Metrics

Multi-factor Authentication (MFA) factors used	Classification of MFA factors used during the username-password based login.
User-password logins without Multi-factor Authentication (MFA)	The percentage of username-password based logins without MFA. This metric provides an insight on the adoption of MFA over a period of time. **Note:** Ideally, the score should gradually decrease and should be zero over a period of time (Refreshed once a day to collect records for a day before).

User Metrics

User Metrics displays the user MFA enrollment trends on ServiceNow.

User Metrics on the MFA Dashboard are as follows:

- Users enrolled in Multi-factor Authentication (MFA)
- Privileged admins without Multi-factor Authentication (MFA)

Users enrolled in Multi-factor Authentication (MFA)

The percentage of users who can perform username-password based login and enrolled in MFA. This metric provides an insight on the adoption of MFA by the users over a period of time.

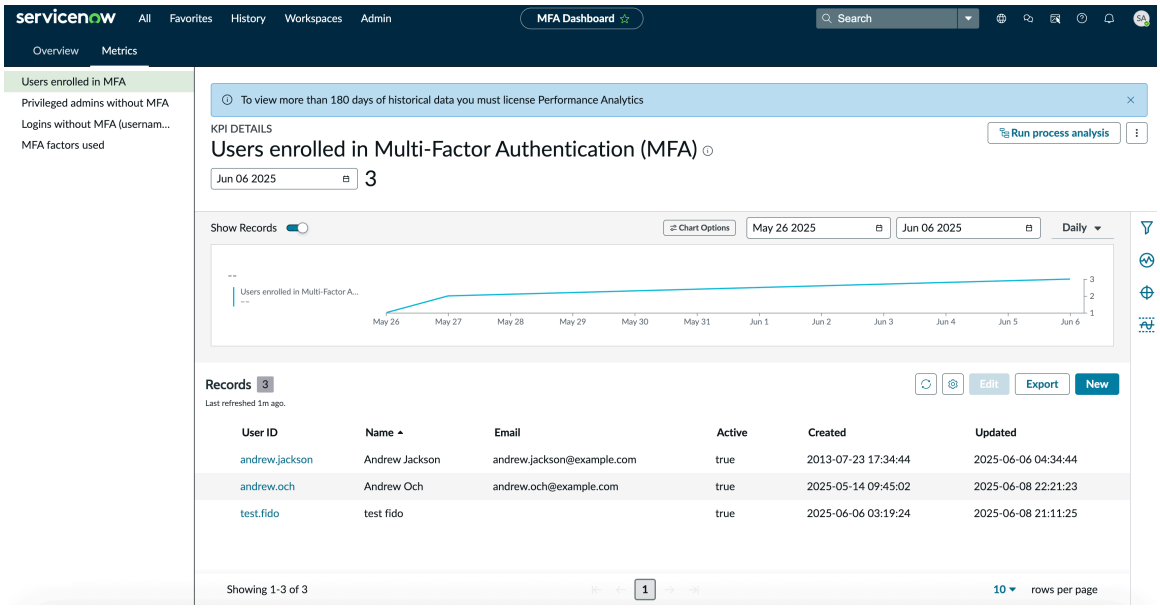

Note: Ideally, the score should gradually increase and should be 100% over a period of time (Refreshed once a day to collect records for a day before).

Privileged admins without Multi-factor Authentication (MFA)

Privileged admins not using MFA is a significant risk to platform security. It's recommended that you get these people using MFA.

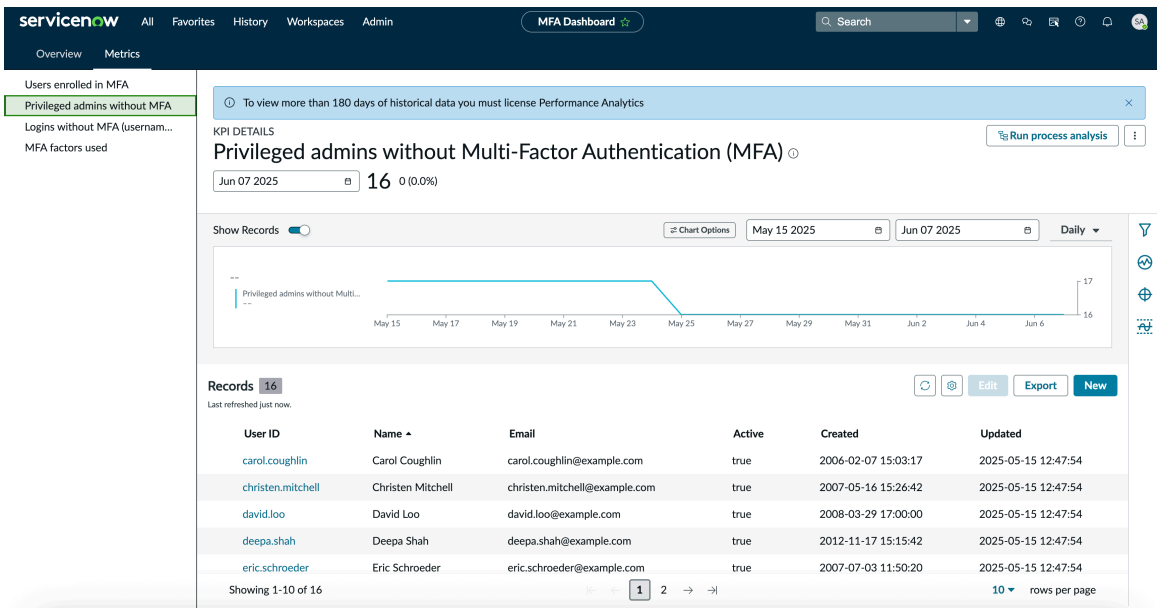

Note: Privileged admins are the users who have at least one role from the **sys_icenter_role_config** table. (Refreshed once a day to collect records for a day before).

Log in Metrics

Log in Metrics displays the log in trends on the ServiceNow.

Log in Metrics on the MFA Dashboard are as follows:

- Multi-factor Authentication (MFA) factors used
- User-password logins without Multi-factor Authentication (MFA)

Multi-factor Authentication (MFA) factors used

Classification of MFA factors used during the username-password based login.

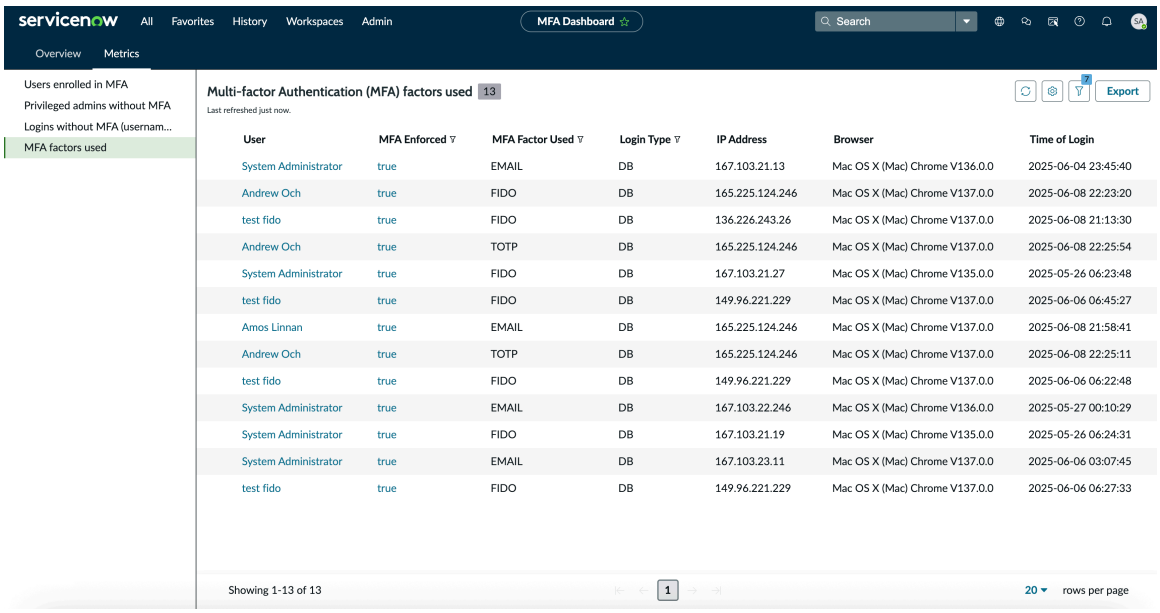

User-password logins without Multi-factor Authentication (MFA)

The percentage of username-password based log ins without Multi-factor authentication. This metric provides an insight on the adoption of MFA over a period of time.

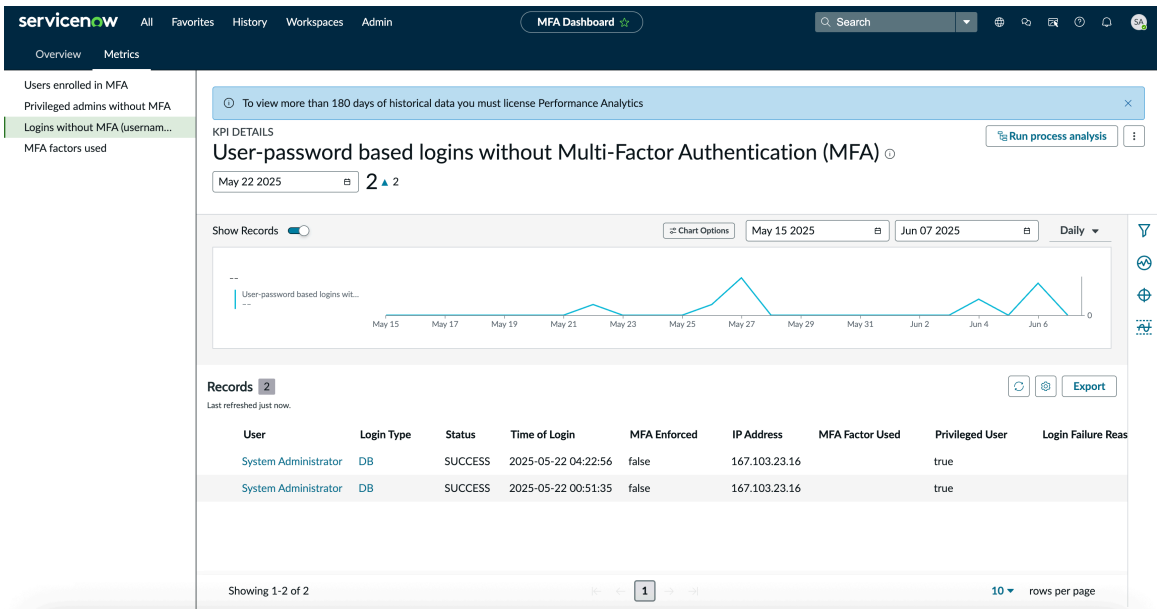

Note: Ideally, the score should gradually decrease and should be zero over a period of time (Refreshed once a day to collect records for a day before).

MFA Guided Setup

Use the MFA Guided Setup to step through the initial configuration of the MFA module and understand the requirements for MFA enforcements.

The MFA Guided Setup assists you to configure MFA for internal users who log in using a username and password.

To access the MFA Guided Setup, navigate to **All** > **Adoption Services** > **Guided Setup**. You can choose **All applications** and choose **Multi-factor Authentication**.

The Multi-factor Authentication guided setup opens in a new tab. Expand any category to view detailed status and related tasks on the page displayed.

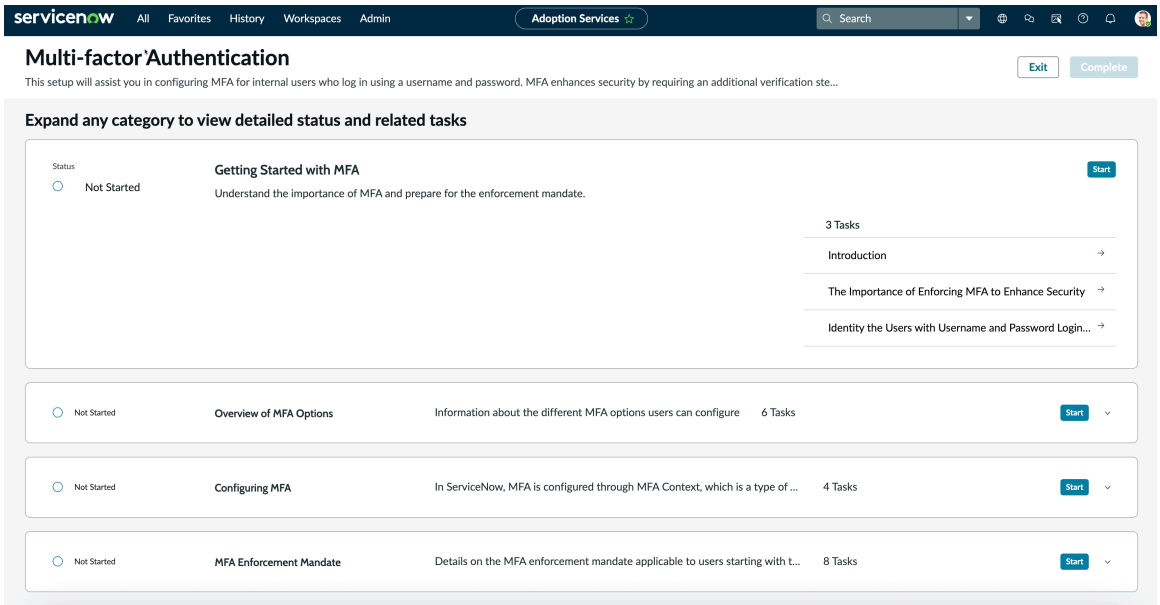

Complete the following task to become familiar with the MFA configurations:

- Getting Started with MFA
- Overview of MFA Options
- Configuring MFA
- MFA Enforcement Mandate
- User Communication and Onboarding
- Monitoring

To know more about how to use the Guided Setup, see Guided Setup.

Multi-Provider single sign-on (SSO)

External SSO allows organizations to use several SSO identity providers (IdPs) to manage authentication as well as retain local database (basic) authentication.

Multi-Provider Single sign-on (SSO) is an authentication method that enables users to access multiple applications with one login and one set of credentials.

For using SSO, you must understand the following:

- *Service Providers*: When users trying to access the ServiceNow instance are redirected to an Identity Providers (IdP) to validate their credentials after successful validation users are allowed to access the instance. Here, ServiceNow acts as a service provider and relies on an Identity Provider (IdP) for handling user authentication and granting access to the instance.
- Identity Providers: IdPs are external systems that validates the users identity and credentials to access a system.

To establish an SSO with to access ServiceNow, you must activate Multi-Provider Single sing-on (SSO) you must install the *Integration - Multiple Provider Single Sign-On Installer (com.snc.integration.sso.multi.installer)* plugin. For more information, see Activate Multi-Provider SSO plugin.

After successful installation of the plugin, you can customize the SSO properties, access tables and scripts that are shipped along with the plugin. For more information, see Multi-Provider SSO properties, tables, and scripts.

ServiceNow supports the following SSO methods:

- OpenID Connect
- SAML 2.0
- Digest Authentication (Token based authentication).

Choose the SSO method based on your requirement and learn more about how you need to prepare for configuring SSO. You must perform several steps to set up Multi-Provider SSO, including configuring properties, creating identity providers (IdPs), and configuring users to use SSO. For more information, see Multi-Provider SSO configurations.

After a successful configuration, the active IdPs in the instance are listed on the ServiceNow. You can list various SAML or OIDC Identity Providers (IdPs).

Note: A maximum of 10 IdPs can be listed on the login page. The IdP options won't be visible if the instance has Domain Support - Domain Extensions Installer (*com.glide.domain.msp_extensions.installer*) plugin installed and enabled.

The Zurich release of ServiceNow include the following enhancements on SSO:

- **List SAML IdPs on login page**: Log in using SAML and OIDC IdPs that are listed on the login experience on both the platform and portal login pages, making it easier for users to select their preferred IdP. Earlier only OIDC IdPs were listed.
- **Select group for Auto-Provisioning**: Select specific groups during the auto-provisioning configuration for SAML and OIDC, ensuring users are assigned to the correct groups automatically.
- **Configure multiple OIDC record using the same well-known URL**: Simplify OIDC setup by allowing the creation of OIDC records using the same well-known URL, streamlining the configuration process.
- **Enhanced External logout complete page**: Display of login failure reason to the user. Provision to log in again to ServiceNow on the external logout complete page in case of successful logout.
- **Enhanced error message**: Display of generic error message when Single Logout (SLO) is unsuccessful, ensuring consistent and secure communication.
- **Notification enhancements for SAML Certificate and Encryption Keystore**: Receive timely notifications to the admins for SAML certificate and Encryption Keystore updates expiry, ensuring that your SSO configurations remain secure and up-to-date.

Example: Why organization needs SSO

A globally dispersed corporation might require one SSO provider for their employees, a different one for their vendors, and local database authentication for their administrators. Alternatively, a company might implement SAML 2.0 and a digest token authentication solutions on the same instance.

Activate Multi-Provider SSO plugin

This integration requires the Integration - Multiple Provider Single Sign-On Installer (com.snc.integration.sso.multi.installer) plugin.

Before you begin

Role required: admin.

The com.snc.integration.sso.multi.installer plugin can also be used for OIDC, SAML, and Digest.

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Integration - Multiple Provider Single Sign-On Installer (com.snc.integration.sso.multi.installer) plugin using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Multi-Provider SSO properties, tables, and scripts

The Integration - Multiple Provider Single Sign-On Installer plugin includes the following system properties, tables, and scripts.

Properties

Multi-Provider SSO adds the following system properties.

Multi-Provider SSO properties

Name	Description
glide.authenticate.multisso.debug	Enables (true) or disables (false) debug logging for the multi-provider SSO integratio • Type: true | false • Default value: false

Multi-Provider SSO properties (continued)

Name	Description
glide.authenticate.multisso.enabled	Enables (true) or disables (false) multi-provider SSO. • Type: true | false • Default value: false **Note:** Setting this property to false will not disable multi-provider SSO if Account Recovery (ACR) is also enabled on the instance. To log in with a username and password ACR must also be disabled using the *glide.sso.acr.enabled* property. For details on this property see Account recovery properties.
glide.authenticate.multissov2_feature.enabled	This property determines if the MultiSSOv2 version is enabled in the instance.
glide.authenticate.show.max.sso.login.option	This property determines the maximum number of SSO options displayed on the log screen. **Note:** The default value is 5. The maximum value of the property is 10.
glide.authenticate.show.max.sso.login.option	This property determines the maximum number of IdPs displayed on the login scree "glide.authenticate.show.max.sso.login.optio **Note:** The default value is 10.

Tables

Multi-Provider SSO adds the following tables.

Multi-Provider SSO tables

Name	Description
SSO Properties [sso_properties]	Stores data for each IdP, schema, common SSO data, and so on.
SAML 2 Update 1 Properties [saml2_update1_properties]	Stores data for SAML 2.0 Update 1 configurations such as SAML certificates.
Digest Properties [digest_properties]	Stores data for digest token authentication configurations.
SSO Federation [sso_federation]	Stores data for each SSO federation.
OIDC Identity Provider [oidc_identity_provider]	Stores data for Open ID connect based identity providers.

Scripts

Multi-Provider SSO adds the following scripts.

Multi-Provider SSO scripts

Name	Description
MultiSSO	Allows a customer to have an SSO type defined on a company basis.
MultiSSOLogin	Allows each domain to have their own login script.
MultiSSOLogout	Allows each domain to have their own logout script.
MultiSSO_OIDC_custom	Allows a user to define a custom Single Sign-on script for OIDC connection.
MultiSSO_OIDC_logout_custom	Allows a user to define a custom logout script for OIDC connection.
MultiSSO_Abstract_Core	Provides a base class for all multi-provider SSO classes.
MultiSSO_ClientHelper	Provides a client callable utility functions for multi-provider SSO.
MultiSSO_DigestedToken	Provides a base system logic for digested token authentication.
MultiSSO_SAML2_Update1	Provides logic to process SAML 2.0 Update 1 authentication for a multi-tenant single sign-on.

Multi-Provider SSO configurations

You must perform several steps to set up Multi-Provider SSO, including configuring properties, creating identity providers (IdPs), and configuring users to use SSO.

Refer the following topics to know more about each of the configurations.

- Multi-Provider SSO (SAML) IdP authentication flow
- Configure Multi-Provider SSO properties
- Create an external identity provider
- Configure users for Multi-Provider SSO

Configure Multi-Provider SSO properties

Configure SSO properties and also add a property to the System Properties table to configure an IdP inclusion list.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Administration** > **Properties**.
2. Select the **Enable Multi-Provider SSO** check box.
   This selection adds the link Use external login to the login page.
3. To update the user table with the users in the IdP, select the **Enable Auto Importing** option.

4. To enable the debug messages to appear at the bottom of the content frame, select the **Enable debug logging for the Multi-Provider SSO integration** check box.
   If enabled, the debug logging feature slows down performance and uses up disk space to generate logs.
5. In the property **The field on the user table that identifies a user accessing the User identification login page**, enter the field on the User table that contains the value the IdP uses to identify the user. The default value is **user_name**.

   The default value is **user_name**.

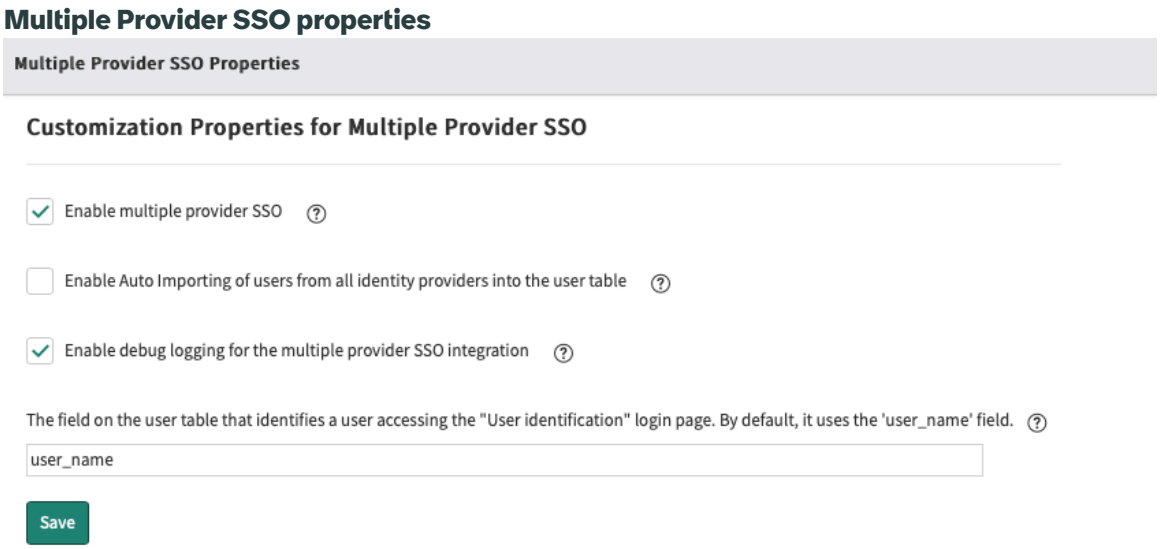

6. Click **Save**.
7. Instruct your users to click the **Use external login** link when they log in to the instance.

Related topics

SAML user provisioning

Create an external identity provider

After you have configured the multi-provider SSO properties, you can update or create new SAML 2.0 or digest token identity provider.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Identity Providers**.
2. To edit an identity provider record, click the record.
   - For digest token configurations, manually update the properties.
   - For SAML2 Update 1 configurations, automatically update the identity provider metadata with the **Import Identity Provider Metadata** related link or update the properties manually.
   - For OpenID Connect configurations, manually update the properties.
3. To create a new identity provider, click **New**.

- For digest token configurations: Click **Digest SSO** and enter the digest properties for multi-provider single sign-on.
- For SAML2 configurations: Click **MultiSSOV2_SAML2_custom** and import the identity provider metadata from a URL, as XML, or manually enter the identity provider information.

- For OpenID Connect: Click **OpenID Connect** and enter client ID, client secret, and well known configuration URL.

4. To make the IdP the failover IdP that is used when the default IdP is not available, select the **Default** check box.

   If you have SAML 2 Update 1 active and you upgrade to the Fuji release, the SAML 2 Update 1 IdP is selected as the default failover. No default failover IdP is selected for new instances or if you are upgrading from a release on which SAML 2 Update 1 is not active.

   Note: The metadata import process automatically creates a certificate record for the identity provider. Navigate to the **x509 Certificate** module to see the certificate.

   Note: Certificates for single-sign on should always be in PEM format to work with SAML certificates.

5. If E-Signature is active, configure the Identity Provider form and add the **Assertion Consumer URL for eSignature authentication** field.
   In most cases, this URL is: `https://YOURINSTANCE.service-now.com/consumer.do`. However, if you employ a customized method of handling the SAML authentication for E-Signature, you can set up your own consumer URL. If you are only using SAML 2.0 Update 1 and not using Multi-Provider Single Sign-on, configure the assertion consumer URL with E-signature SAML properties.

Generate instance service provider (SP) metadata for SAML

As part of your SSO configuration, you can generate the instance SP metadata to provide to the IdP.

Before you begin
Role required: admin

About this task
The IdP needs the instance SP metadata to authenticate and forward requests.

Procedure

1. Choose your installed SSO plugin:

Option	Description
Multi-Provider SSO	Navigate to **Multi-Provider SSO** > **Identity Providers**. Choose an IdP and click the **Generate Metadata** button. The integration automatically generates the instance's SP metadata from the system property settings.
SAML 2 SSO	Navigate to **SAML 2 Single Sign-on** > **Metadata**. The integration automatically generates the instance's SP metadata from the system property settings.

2. Copy the SP metadata in the text box.

   For example:

```
<EntityDescriptor xmlns="urn:oasis:names:tc:SAML:2.0:metadata"
 entityID="https://yourinstance.service-now.com">
  <SPSSODescriptor AuthnRequestsSigned="false"
 WantAssertionsSigned="true"
 protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protoc
ol">
  <SingleLogoutService
 Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect"
 Location="https://yourinstance.service-now.com/navpage.do" />

 <NameIDFormat>urn:oasis:names:tc:SAML:1.1:nameid-format:emailA
ddress</NameIDFormat>
  <AssertionConsumerService isDefault="true" index="0"
 Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"
 Location="https://yourinstance.service-now.com/navpage.do" />
  <AssertionConsumerService index="1"
 Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"
 Location="https://yourinstance.service-now.com/consumer.do"/>
 </SPSSODescriptor>
</EntityDescriptor>
```

3. Provide the instance SP metadata to the IdP.
   For example, SSOCircle allows a user to provide the SP metadata online.

Configure users for Multi-Provider SSO

Administrators can configure Multi-Provider SSO for individual users or for all users who belong to a company. You cannot configure Multi-Provider SSO for groups.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Identity Providers**.
2. Right-click an identity provider record and select **Copy sys_id**.
3. Copy the data to your clipboard.

4. Navigate to a user record or a company record.
5. Configure the form and add the **SSO Source** field.
6. In the **SSO Source** field, enter one of the following:
   - **SAML users**: enter **sso:** followed by the sys_id of the identity provider's record.
   - **SSO Federation users**: enter **federation:** followed by the sys_id of the federation record.
7. Click **Update**.

Test IdP connections

Testing the connection to an IdP validates the settings before enabling external authentication.

Before you begin

Role required: admin

About this task

The Jakarta release supports the test connection within a pop-up window. If your IdP does not work correctly with this option, you can turn off this default setting.
Refresh MultiSSO IDP Metadata job fetches and updates the certificate for the IdP while create, update, or testing the connection.

Note: For some idp cases, If test connection is failed, they need to create glide.authenticate.multisso.test.connection.mandatory with the value as false and you can activate the idp without the test connection.

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Identity Providers**.
2. Select a defined IdP or click **New** to define a new IdP.
3. **Optional:** Configure an identity provider if setting up a new IdP.
4. Click **Test Connection**, enter login credentials for the IdP to check login.
   You cannot activate the IdP until you have a successful test connection. If the test fails, you can update to save your configuration information, but you cannot activate this configuration.
5. Verify results using the **SSO Test Results/Summary** or the **Testing SSO Logs** section to see log messages.
   If there are any errors, refer to the Multi-SSO (SAML 2.0) errors and fixes

6. Click **Cancel** when testing is complete.

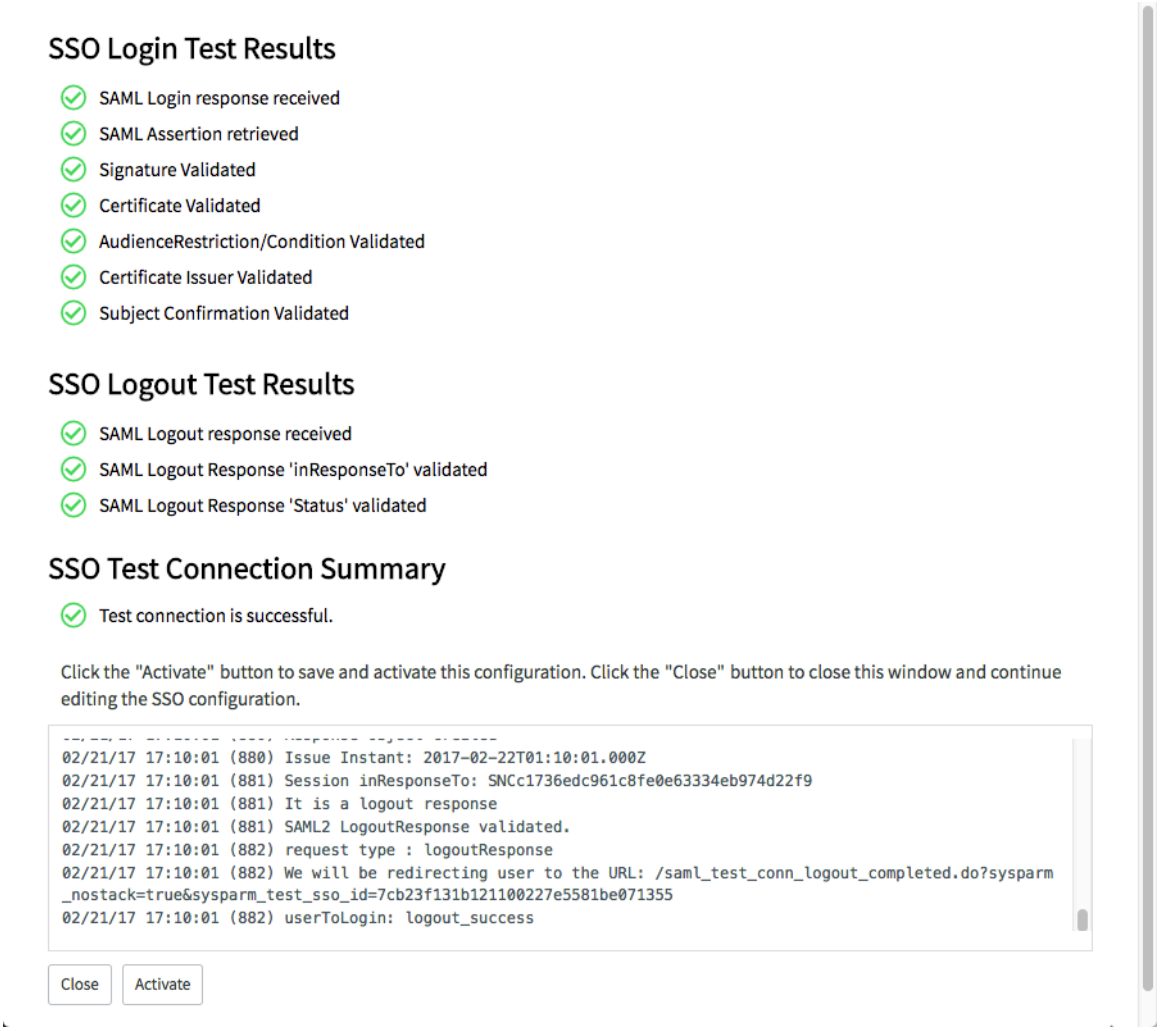

Common IdP connection errors

The following table describes some of the common IdP connection errors and their solutions.

Troubleshooting IdP test connections

Error messages	Solution
User Field validation failed. Invalid User Field '<field name>' is not a field on sys_user table.	Verify the contents of the User table field you selected matches the SAML NameID token.
Assertion issuer is invalid.	Verify **Identity Provider URL** contains a valid URL to your IdP. Each IdP URL must be unique.
AudienceRestriction validation failed.	Verify the **Audience URI** contains a valid URL to your instance.
Cannot logout of IdP's session.	Verify the **SingleLogoutRequest** URL contains a valid URL to your IdP's logout service.
Signature did not validate against the credential's key.	Verify the IdP has a valid certificate installed.

Troubleshoot script issues with SAML

Troubleshoot script issues with SAML. You might encounter script issues if SAML is already active at the time that you activate Multiple Single Sign-On and if you already customized the installation exits.

Before you begin
Role required: admin

Procedure

1. Back up the modified installation exit SAML2SingleSignon_update1 and script include SAML2_update1.
2. Revert both the installation exit and script include to the version that is available with the baseline system.
3. Activate or upgrade the **Integration - Multiple Provider Single Sign-On Installer** plugin. The system upgrades SAML and all necessary files to SAML 2 Update 1.
4. Open the Multiple SSO properties page and select the **Enable Multi-Provider SSO** check box to enable it.
5. Put the SAML2SingleSignon_update1 installation exit changes into the baseline script include **MultiSSO_SAML2_Update1 and the SAML2_update1** script include changes into the baseline SAML2_update1 script include.

Log in using Multi-Provider SSO

The recommended and most efficient method for users to log in using Multi-Provider SSO is to use a specifically configured URL.

Before you begin
Role required: admin

About this task
After multi-provider SSO is configured, you can send a URL to your users with the correct IdP in the parameter string. For example:

```
/login_with_sso.do?glide_sso_id=<sys_id of the sso configuration>
```

After a user successfully logs in to the IdP page, a cookie containing the IdP sys_id is added to the browser. The next time the user attempts to log in, the system redirects the user to log in to the IdP server, which automatically logs in to the instance.

If a URL parameter is not set or the browser cache has been cleared, users can also do the following:

Procedure

1. Click the **Use external login** link on the login page.

   The external login page appears. Users can click **Use local login** to return to the standard login page.

2. Enter the value for the specified field on the user table that you configured in Multi-Provider SSO properties.

   The user is redirected to the IdP server, where they log in.

Result

After users successfully log in to an IdP, they are automatically redirected to that IdP whenever they attempt to access the instance. To have a user access a different IdP, send the user a URL with the new IdP information in the parameter. The new IdP overwrites the old IdP in the cookie if the user successfully logs in. If the user does not log in successfully, the old IdP information is retained in the cookie.

Enable users to choose the identity provider for login

SSO federation support enables users to choose which IdP to log in to.

Before you begin

Role required: admin

About this task

SSO federations aggregate metadata from multiple IdPs and service providers, including your instance. Federations then publish the metadata as an XML file, which includes information like IdP names and IdP certificates. Administrators can then instruct the instance to read the XML file and automatically populate the SSO Properties table with all the necessary IdP information.

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Federation**.
2. Click **New**.
3. Fill in the fields, as appropriate (see table).
4. Click **Submit**.
5. After you configure a federation, enable the Refresh SSO Metadata scheduled job, and then configure the users who you want to access the federation IdPs. Use the sys_ID of the federation record you just created.

   Use the sys_ID of the federation record that you just created.

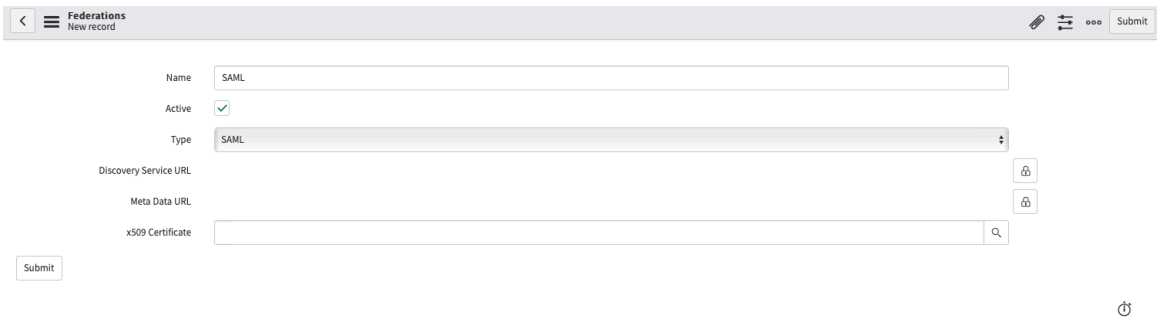

   The instance populates the SSO properties table with the IdP information. When users who are configured to use the federation log in, they are redirected to the discovery service URL you configured. Then they select the IdP and provide the necessary credentials. Alternatively, you can send users a URL with the IdP in the parameter.

Allowing users to choose the identity provider for login

Field	Description
Name	Enter a descriptive name for the federation.
Active	Select the check box to enable the instance to pull the XML file from the federation.
Type	Select the type of authentication this federation supports.

Field	Description
Discovery Service URL	Enter the URL of the discovery service for this federation. This is the site where users are directed to select an IdP and log in.
Meta Data URL	Enter the URL of the XML file that holds the federation metadata.
x509 Certificate	Select the federation certificate.
Domain	Select the domain that the data will belong to.

Note: The InCommon federated identity management IdP is preconfigured.

Use Service Portal with Multi-Provider SSO to redirect a URL

Service Portal uses a combination of system properties and script includes to determine how the system handles URL redirects for users logging in to the portal.

Before you begin
Role required: admin

About this task
If you are using the system property to automatically redirect to your primary Identity Provider (IdP), then Service Portal automatically redirects to that IdP. If you have multiple IdPs, Service Portal shows a link on the login page to **Use external login**.

Procedure

1. Configure the Service Portal login page
2. Redirect to Service Portal after login

Account recovery (ACR)

Administrators can configure account recovery (ACR) to perform recovery activities such as addressing SSO misconfiguration or expired certificates.

Note: Enabling ACR disables the local interactive log-ins (username or password based) when SSO is enabled to your instances.

ACR provides the following capabilities:

- Bypass your single sign-on (SSO) login to address issues with SSO configuration as an administrator.
- Log in with using SSO to perform tasks with an administrator account configured as an account recovery.
- ACR flows enable the administrators to use self-service capabilities to address account recovery when there's a need for recovery, for example, SSO miss-configuration, expired certificates.
- Reduce unauthorized access to the instance and provide a strong foundation to use ACR outside SSO use cases.

Fresh Instance

For a fresh instance to use ACR, you must do the following:

- Activate Mutli-SSO plugin (`com.snc.integration.sso.multi.installer`)
- Enable ACR (`glide.sso.acr.enabled`) - This is enabled by default in case of a fresh instance.
- Before enabling SSO property (`glide.authenticate.multisso.enabled`), the administrator must enroll as an ACR user.

  Note: Setting this property to false will not disable multi-provider SSO if Account Recovery (ACR) is also enabled on the instance. To log in with a username and password ACR must also be disabled using the *`glide.sso.acr.enabled`* property. For details on this property see Account recovery properties.

- Administrator must set a password for local login and register MFA before enrolling as an ACR user.

Upgraded Instance

For an upgraded instance to use ACR, you must do the following:

- Activate Mutli-SSO plugin (`com.snc.integration.sso.multi.installer`)
- Enable ACR (`glide.sso.acr.enabled`)

  Note: In case of upgraded instance, the administrator must enable ACR.

- Before enabling SSO property (`glide.authenticate.multisso.enabled`), the administrator must enroll as an ACR user.
- Administrator must set a password for local login and register MFA before enrolling as an ACR user.

Configure account recovery users

To use account recovery, you must register at least one admin account as an account recovery user. Single sign-on can't be activated on your instance until there is at least one account configured. For details on this process, see Configure an account recovery user from the Account Recovery Properties page.

Note: If you're upgrading an instance already using single sign-on to Rome or a later release, single-sign on will continue to function without a recovery user configured.

Account recovery configuration

The account recovery feature is included with the **Integration - Multiple Provider Single Sign-On Installer (com.snc.integration.sso.multi.installer) plugins**. The feature is enabled by default. You can change this and other account recovery settings using system properties. For details on these properties, see Account recovery properties.

Account recovery policy context

After you've registered an account recovery user and enabled single sign-on (SSO), your instance restricts all local logins. This restriction is defined in the **SSO - ACR Context** auth policy context. For more information about the context, see Account recovery context.

For details on how authentication policies and policy contexts, and how they work on your instance, see Adaptive authentication.

Configure an account recovery user

Configure an account recovery user to perform account recovery activities on your instance.

An account recovery user is a user account that administrators can use to perform account recovery tasks, such as addressing an SSO misconfiguration or addressing expired certificates.

Note: If you are using account recovery on your instance, you must configure an account recovery user. This step is necessary before enabling multiple-provider single sign-on on an instance.

Configure an account recovery user from the Account Recovery Properties page

Configure an account recovery from the Account Recovery properties page.

Before you begin

Role required: admin

For a fresh instance to configure ACR, you must do the following:

- Activate Mutli-SSO plugin (`com.snc.integration.sso.multi.installer`)
- Enable ACR (`glide.sso.acr.enabled`) - This is enabled by default in case of a fresh instance.
- Before enabling SSO property (`glide.authenticate.multisso.enabled`), the administrator must enroll as an ACR user.
- Administrator must set a password for local login and register MFA before enrolling as an ACR user.

For an upgraded instance to use ACR, you must do the following:

- Activate Mutli-SSO plugin (`com.snc.integration.sso.multi.installer`)
- Enable ACR (`glide.sso.acr.enabled`)

  Note: In case of upgraded instance, the administrator must enable ACR.

- Before enabling SSO property (`glide.authenticate.multisso.enabled`), the administrator must enroll as an ACR user.

  Note: Setting this property to false will not disable multi-provider SSO if Account Recovery (ACR) is also enabled on the instance. To log in with a username and password ACR must also be disabled using the *`glide.sso.acr.enabled`* property. For details on this property see Account recovery properties.

- Administrator must set a password for local login and register MFA before enrolling as an ACR user.

Procedure

1. Navigate to **All** > **Account Recovery** > **Properties**.
2. Select **Enable account recovery**.

   Note: You need to enable account recovery when SSO is enable. Account recovery users will be limited to SSO configuration and troubleshooting-related tasks.

3. Select the **here** text in the Step 2.

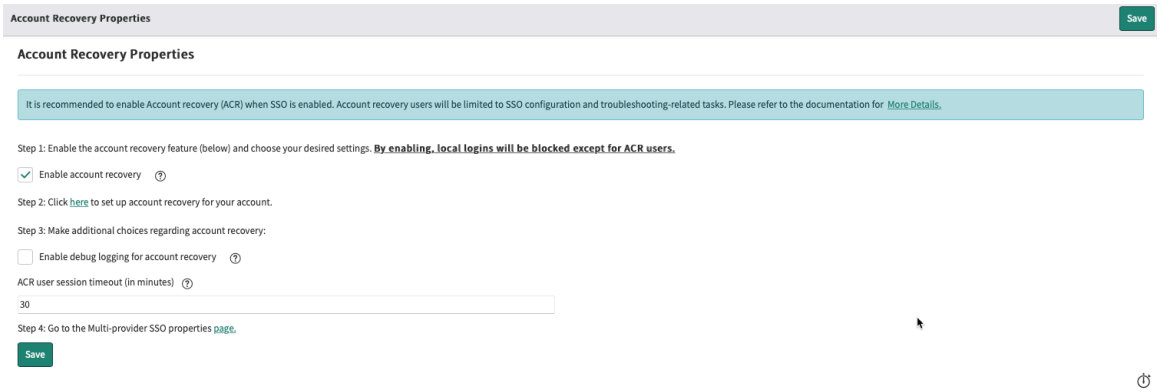

4. Following the on-screen directions in the **Configure account recovery for Multi-SSO** modal.

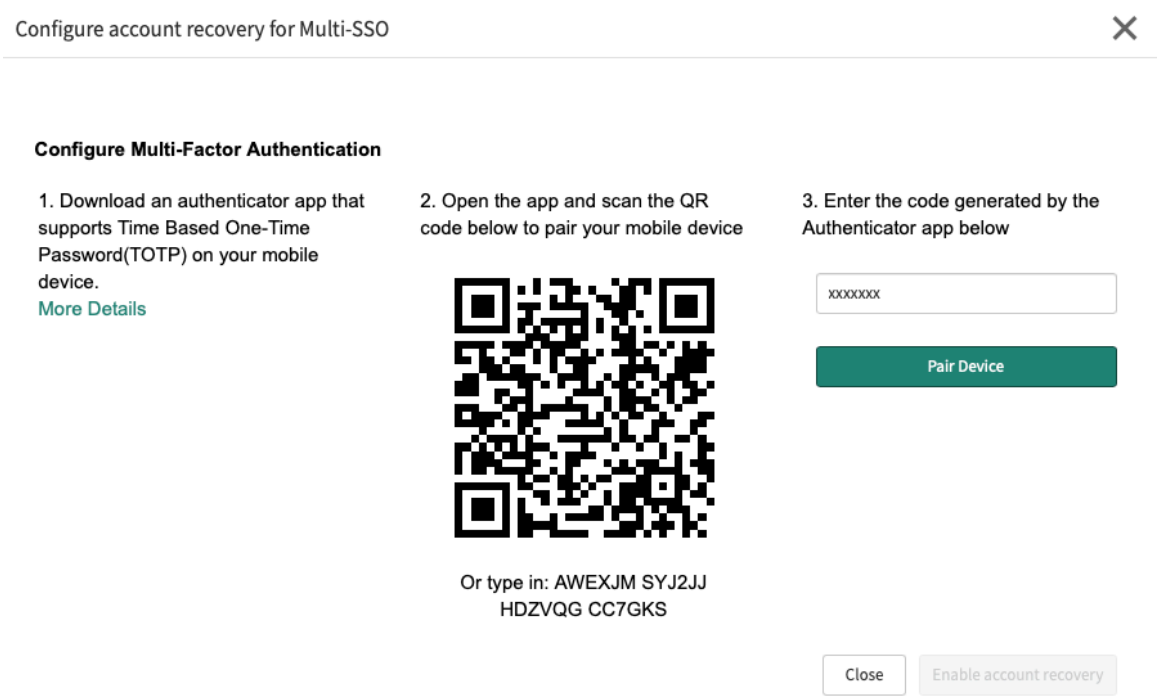

After completing the on-screen steps, the **Enable account recovery** is enabled.

5. Select **Enable account recovery**.
6. Select **Save**.

Result

You have configured your user account as an account recovery user. You can verify this account, and see any other configured account recovery users by navigating to **Multi-Provider SSO** > **Account Recovery** > **Users**.

Configure an account recovery user from an admin user profile

Configure an administrator as an account recovery user from the admin user profile.

Before you begin

Role required: admin

Procedure

1. Log into your instance using an administrator account.
2. Click on the user name in the instance header, and select **Profile**.

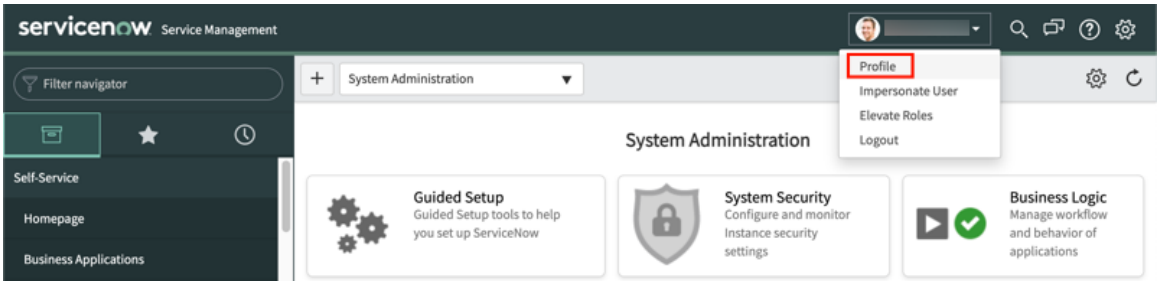

3. In the **User** form, click **Enable Account Recovery** in the **Related Links** section.

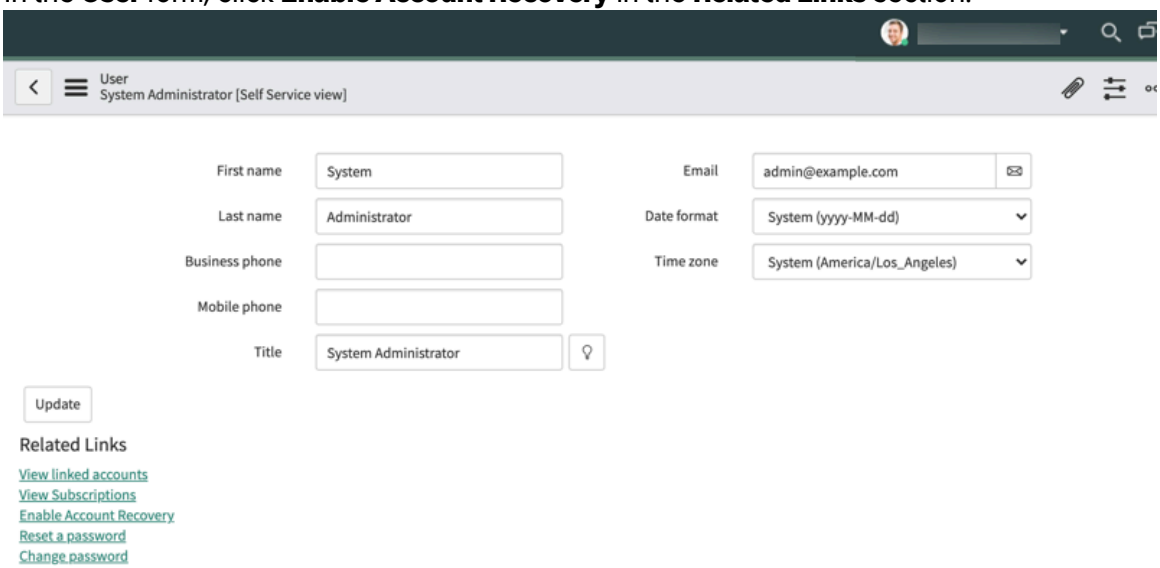

Note:

If the selected user already has account recovery enabled, **Enable Account Recovery** does not appear in the related links. There will be a **Disable Account Recovery** option instead.

4. Following the on-screen directions in the **Configure account recovery for Multi-SSO** modal.

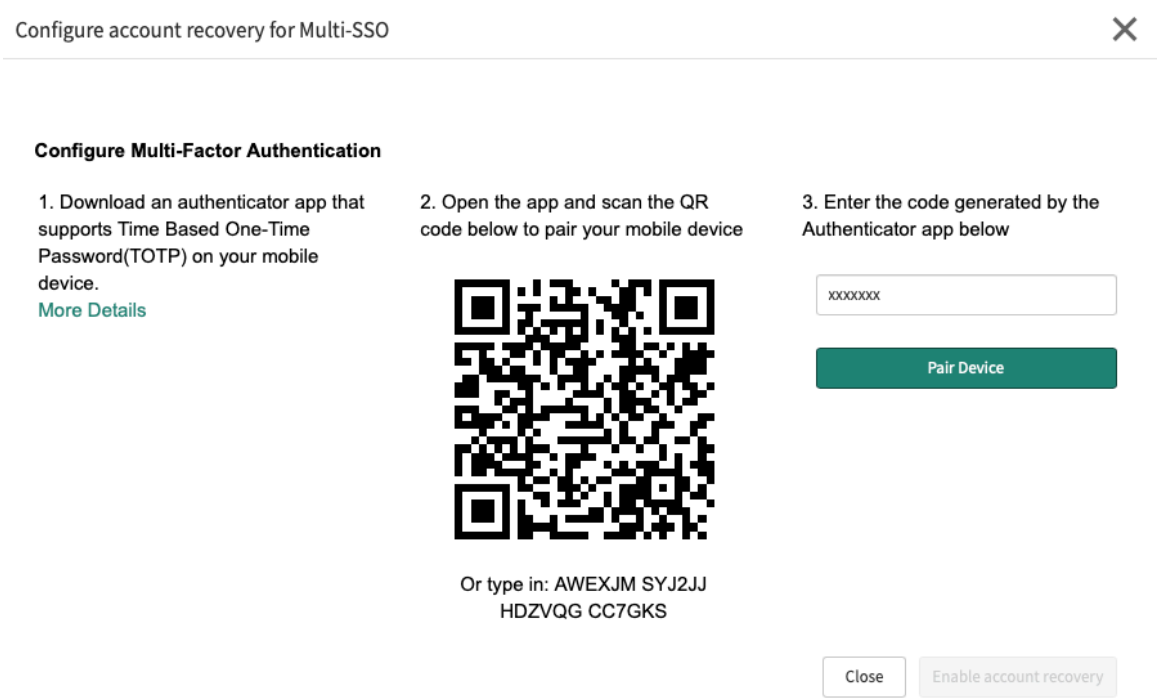

After completing the on-screen steps, the **Enable account recovery** is enabled.

5. Click **Enable account recovery**.

Result

You have configured your user account as an account recovery user. You can verify this account, and see any other configured account recovery users by navigating to **Multi-Provider SSO** > **Account Recovery** > **Users**.

Account recovery properties

Use system properties to configure Account Recovery (ACR) on your instance.

Access the account recovery properties on your instance by navigating to **Multi-Provider SSO** > **Account Recovery** > **Properties**.

Account recovery system properties

Property	Description
Enable account recovery feature [glide.sso.acr.enabled]	Whether the account recovery feature is enabled on your instance. This property is enabled by default.
Enable debug logging for account recovery [glide.sso.acr.debug.log.enabled]	Whether your instance includes account recovery information in debug logging. This property is disabled by default.
ACR user session timeout (in minutes) [glide.sso.acr.ui.session.timeout]	Minutes of inactivity before your instance terminates an account recovery user session. This property has a default value of 30.

E-signature for Multi-Provider SSO

E-signature with Multi-Provider SSO enables you to use the e-signature properties instead the SAML or OIDC properties for authentication.

For single sign-on (SSO) verification during authentication and to require users to provide credentials before submitting their electronic signature, you can configure the authentication to prompt for user credentials before submitting a signature.

You can install the Approvals with e-signature (`com.glide.e_signature_approvals`) plugin to have the E-signature configuration for SSO logins.

Note: You must install the Code Signing Signatures (`com.glide.code_signing.signatures`) to install the E-signature plugin.

Activate Approval with e-Signature plugin

The Approval with e-Signature plugin (com.glide.e_signature_approvals) allows users to approve requests by re-entering their login credentials.

Before you begin

Role required: admin

Note: You must install the Code Signing Signatures (`com.glide.code_signing.signatures`) to install the E-signature plugin.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the plugin using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Use Multi-Provider SSO to set up an SSO approval for a SAML 2.0 authentication

An SSO approval with e-signature requires configuration on the SAML IdP and the ServiceNow instance.

Before you begin

Role required: admin

About this task

The SAML IdP must support and honor the forceAuthn attribute in SAML assertion requests. E-signature doesn't function without this IdP setting. Set up an approval with e-signature using credentials from a SAML 2.0 authentication.

Procedure

1. Activate or upgrade to SAML 2.0 with the Activate Multi-Provider SSO plugin.
2. Activate the Approval with E-Signature plugin.
3. Navigate to **Multi-Provider SSO** > **Identity Providers** and verify your 2.0 SAML IdP configuration Advanced tab shows the **Force AuthnRequest** attribute checked.

Your SAML 2.0 IdP must support the **Force AuthnRequest** attribute, or e-signature isn't supported.

4. On the eSignature Approval tab, enter the following e-signature SAML properties.

Option	Description
Assertion Consumer URL for eSignature authentication	This property defaults to the appropriate URL. To configure this property, select the lock icon to make this field editable. After edits, select the icon to lock the field.
Assertion Consumer Index for eSignature authentication	If your Service Provider has more than one URL set for the AssertionConsumerURL, you can set the index to use for eSignature, starting with index 1 or more.
AuthnRequest URL for eSignature Authentication	You can enter the URL that points to the SAML 2.0 IdP AuthnRequest URL for eSignature authentication. If the URL is the same as the Assertion Consumer URL, you can leave this setting empty.
Authentication pop-up Dialog Width	When a user approves a request using eSignature, a dialog opens and a user can enter credentials. This setting controls the width of that dialog box. The default is 500.
Authentication pop-up Dialog Height	When a user approves a request using eSignature, a dialog opens and a user can enter credentials. This setting controls the height of that dialog box. The default is 300.

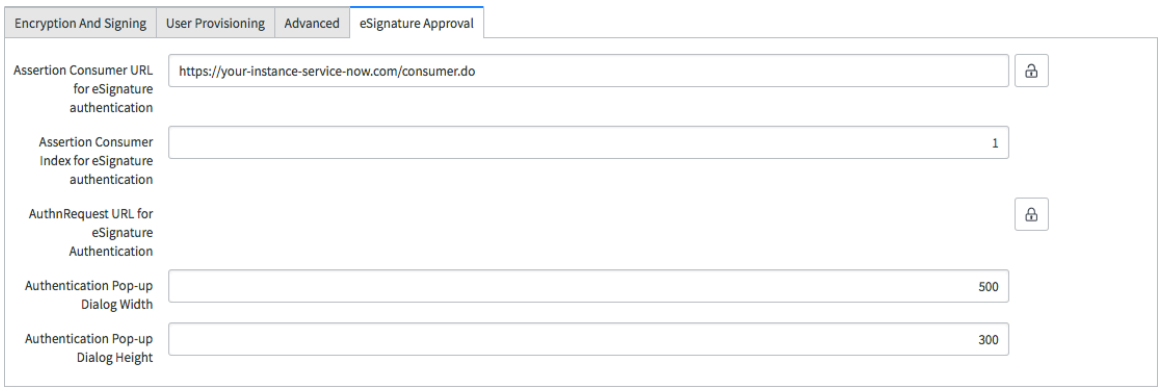

5. Select the **Generate Metadata** button underneath the tabs to regenerate the service provider metadata.
6. Copy the service provider metadata, and update it on the SAML IdP.

Use Multi-Provider SSO to set up an SSO approval for an OIDC authentication

An SSO approval with e-signature requires configuration on the SAML IdP and the ServiceNow instance.

Before you begin

Role required: admin

About this task

The SAML IdP must support and honor the forceAuthn attribute in SAML assertion requests. E-signature doesn't function without this IdP setting. Set up an approval with e-signature using credentials from a SAML 2.0 authentication.

Procedure

1. Activate the Approval with E-Signature plugin.
2. Navigate to **Multi-Provider SSO** > **Identity Providers** and verify your OIDC provider configurations
3. On the eSignature Approval tab, enter the following e-signature SAML properties.

Option	Description
Assertion Consumer URL for eSignature au thentication	This property defaults to the appropriate URL. To configure this property, select the lock icon to make this field editable. After edits, se lect the icon to lock the field.
Authentication pop-up Dialog Width	When a user approves a request using eSig nature, a dialog opens and a user can enter credentials. This setting controls the width of that dialog box. The default is 800.
Authentication pop-up Dialog Height	When a user approves a request using eSig nature, a dialog opens and a user can enter credentials. This setting controls the height of that dialog box. The default is 900.

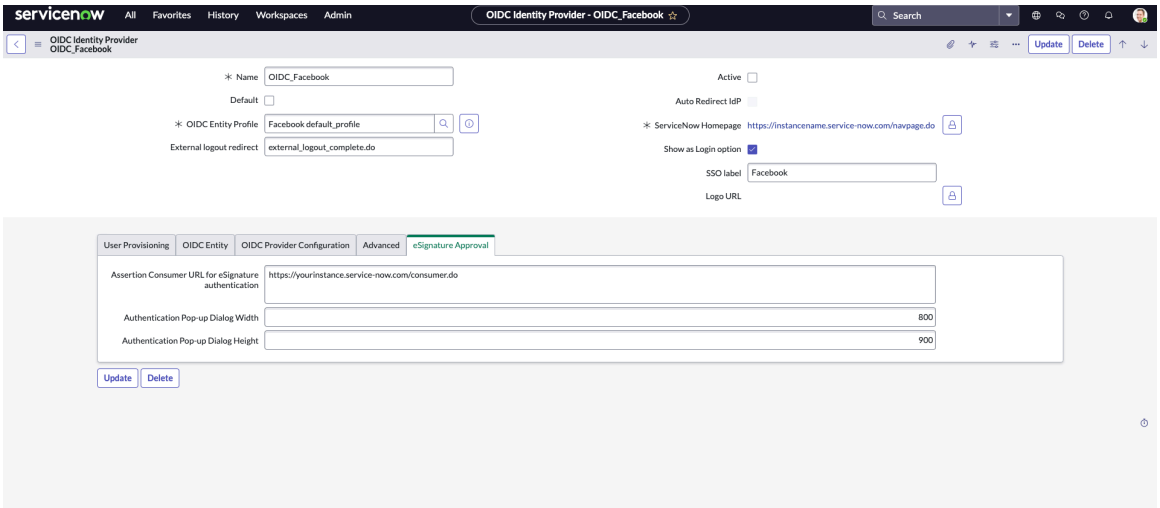

4. Select **Submit** if you are configuring the E-signature during the initial OIDC setup or **Update** if you want to update the details in the E-signature.

OpenID Connect (OIDC) as a Single Sign-On (SSO) identity provider (IdP)

OpenID Connect (OIDC) is an identity layer built on top of the OAuth protocol, which provides a modern and intuitive Single Sign-on (SSO) experience to you and your end users.

OIDC improves the log in experience for mobile applications by enabling users to log in to ServiceNow applications using their social identity provider. For example, administrators can configure Single Sign-on with a third-party identity provider that supports OpenID Connect. Users then have the option to log in to your custom ServiceNow application using their identity provider credentials.

Note: ServiceNow support the JSON Web Encryption (JWE) as part of the OIDC SSO flow. To know more, see the KB Article.

You can choose to use social identity providers like Google for your business-to-customer (B2C) users and enterprise identity providers like Okta, Microsoft Entra ID for your business-to-business users.

Create an OpenID Connect (OIDC) configuration for Single Sign-On (SSO)

Create or update an OpenID Connect (OIDC) configuration by using the Multi-Provider SSO plugin.

Before you begin

- Register an OIDC application with your Identity Provider (IdP) and note the Client ID, Client Secret, and Well-known configuration URL.
- Activate Multi-Provider SSO plugin. Multi-Provider SSO feature supports domain separation, you can specify different IdP's to different domains.
- Configure Multi-provider SSO properties
- Activate Approval with e-Signature plugin to enable eSignature for the OIDC IdP.
- Role required: admin

If you have a client ID, client secret, and well-known configuration URL of the identity provider, you can directly import the OIDC configuration for SSO.

Note:

- The **Login with OIDC** button is not shown on the login page for OIDC based IDPs if the domain separation plugin is installed.
- Admin can use OIDC based IdP to enable SSO for a user of specific company or domain.
- ServiceNow support the JSON Web Encryption (JWE) as part of the OIDC SSO flow. To know more, see the KB Article.

If you do not have the required information about the identity provider, you can manually configure OIDC for SSO. After completing the configuration, users can log in to ServiceNow applications using third-party social identity providers like Google Okta.

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Identity Providers**.
2. Choose one of the following options.
   - To update an existing configuration, click an OIDC Identity Provider record.
   - To create a new configuration, click **New** and select **OpenID Connect**.
3. For a new configuration, enter the OIDC configuration information in one of the following methods.

Option	Description
Import OpenID Connect Well-Known Configuration	If you have the well-known configuration URL along with your associated client credentials, you can directly import an OIDC configuration.

Option	Description
	Note: If you import the OIDC well-known configuration, all related fields are auto-populated.
Manually configure the OIDC Identity Provider form	If you do not have an existing OAuth OIDC En tity, close the Import OpenID Connect Well-Known Configuration pop-up and manually fill the fields in the OIDC Identity Provider form.

Import OpenID Connect Well-Known Configuration fields

Property	Description
Name	Unique name for the OIDC identity provider configuration.
Client ID	Client ID of the application registered in the third-party OIDC identity provider.
Client Secret	Client secret of the application registered in the third-party OIDC identity provider.
Well-known Configuration URL	URL that contains metadata about the third-party OIDC identity provider.

All required fields must be filled in on the OIDC Identity Provider form.

Before you manually fill the OIDC Identity Provider form, ensure that you already have an OAuth Entity Profile for the OIDC IdP.

If you do not have a OAuth Entity Profile, you can create it using the default External OIDC Provider templates, like Okta, Azure and others.

The grant type of the OAuth Entity Profile must be with an authorization code. For more information, see Configure an OAuth OIDC provider on the ServiceNow AI Platform.

Note: You can use the templates of third-party identity providers, Auth0, Azure AD, Google, and Okta are available in the demo data of the Multiple Provider Single Sign-On Installer plugin.

OIDC Identity Provider fields

Property	Description
Name	Name of the OIDC identity provider record.
Active	Option to make the OIDC IdP configuration active. **Note:** This option can only be set to active after a successful test connection.
Default	Option to set the OIDC IdP configuration as default when there are more than one OIDC configurations.
Auto Redirect IdP	Option to enable auto redirection of the users to the login page of the identity provider. This field shows when the **Set as Auto Redirect IdP** option is set under the Related Links section.

Property	Description
	Note: If you make a new Auto Redirect IdP configuration active, the *glide_sso_id* cookie automatically updates with the new Auto Redirect IdP. The *glide.authenticate.sso.update.idp.cookie* system property controls this feature.
OIDC Entity Profile	OAuth Entity Profile for the OIDC configuration.
ServiceNow Homepage	The URL of the login page used for authentication. This field is automatically set to your instance URL. The format of the URL is: `https://yourinstance.service-now.com/navpage.do`
External logout redirect	The URL where the integration redirects users after they log out. Typically, the portal, which is used for SSO. This field is automatically set to `external_logout_complete.do` For example, `https://yourinstance.service-now.com/external_logout_complete.do`
Show as login option	Option to display the OIDC IdP as a login option on the login page. The login option appears as the **login with Identity provider** button.
SSO label	Label of the OIDC IdP displayed on the login page. This field appears only when **Show as login option** is enabled.
Logo URL	Publicly available URL that contains logo of the OIDC IdP provider. This field appears only when **Show as login option** is enabled.

4. **Optional:** Enable automatic user provisioning in the User Provisioning tab>User Provisioning tab.

   (Optional) You can choose to enable automatic user provisioning during user login. When automatic user provisioning is enabled, a user record is automatically created in the ServiceNow instance if that user record does not exist.

   User Provisioning fields

Property	Description
Automatically provision users	Option to enable automatic user provisioning. This property creates a user in the instance User (sys_user) table when the user exits on IdP but does not exist in the User table.
Provision using	Data source to use to transform, an `ID Token`, `User Info` endpoint, or `Both ID Token and User Info` to a ServiceNow user. Use the Lookup list to select the pre-defined data source template, then open the record to configure the Transforms table mapping.
Provision data source	ID token data source used for user provisioning.
User Info Datasource	The user info endpoint datasource used for user provisioning. This field is displayed when `User Info` or `Both ID Token and User Info` are selected for the **Provision using** field.
Update User on next login	Option to enable user update during the next login.

Property	Description
Update User Interval Time (Seconds)	Minimum time interval in seconds to update a user record between subsequent logins. This field is automatically set to 3,600 seconds. For example, after a user logs in, the user record will be updated after 3,600 seconds until the next login. This field is available only when the **Update User on next login** field is enabled.
User roles applied to provisioned users	List of roles applied to the newly provisioned users.

5. OIDC Entity tab
   You can view and modify the OIDC client configuration and OIDC connect flow using the entity record.
6. OIDC Provider Configuration tab
   You can view and modify the well-know configuration URL of the OIDC IdP or ID token claim validation.
7. **Optional:** Advanced tab

   (Optional) Scripts that are run during single sign-on and logout.

   Advanced fields

Property	Description
Single Sign-On Script	Script that executes during Single Sign-On. This field is automatically set to `MultiSSO_OIDC_custom`.
Logout Script	Script that executes after the user logs out. This field is automatically set to `MultiSSO_OIDC_logout_custom`.

8. **Optional:** On the Continuous Authentication tab, configure the following fields:

   Note:

   - The Continuous Authentication tab appears only when you install the **Zero Trust - Continuous Authentication** (`com.snc.zero_trust_continuous_authentication`) plugin that requires license.
   - If you're using continuous authentication policy to protect access to table or data class, see Continuous Authentication (CA).

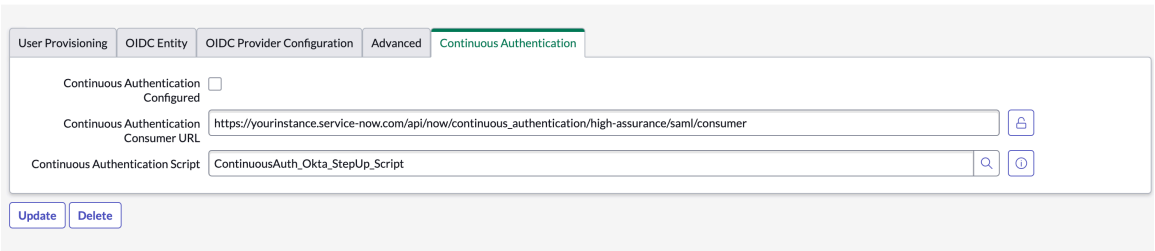

Continuous Authentication

Field	Description
Continuous Authentication Configured	Select the check-box to set the configuration active.
Continuous Authentication Consumer URL	Provide the Consumer URL from the Identity Provider.
Continuous Authentication Script	(Optional) Select the look-up icon to choose the script provided from the platform. In this configuration, for OIDC Okta: **ContinuousAuth_Okta_StepUp_Script**

9. **Optional:** On the eSignature Approval tab, configure the eSignature for the OIDC Idp.

   Note: The eSignature Approval tab appears only when you install the **Approval with e-Signature** plugin (com.glide.e_signature_approvals).

eSignature Approval fields

Property	Description
Assertion Consumer URL for eSignature authentication	If you employ a customized method of handling the OIDC authentication for eSignature, you can set up your own consumer URL. For example, if you are using Multi-Provider SSO, you do not need to use this property. This format of the URL is `https://yourinstance.service-now.com/consumer.do`
Authentication pop-up Dialog Width	Width of the authentication pop-up dialog. This field is automatically set to 800.
Authentication pop-up Dialog Height	Height of the authentication pop-up dialog. This field is automatically set to 900.

10. **Optional:** Navigate to the login page of the instance to verify that IdP appears as a login option.
    The URL should be in the following format: `https://yourinstance/login_with_sso.do?glide_sso_id=sysId_IdP`

    Note: If you have enabled **Selected as login Option**, you can go to the login URL of the instance.

Use Facebook-based Single Sign-On (SSO)

Log in to your ServiceNow instance by using your Facebook credentials on the Facebook-based SSO.

Before you begin

The Facebook-based SSO is shipped along with your ServiceNow instance.

You can define the Identity Provider (Idp) configurations to the **OIDC_Facebook** IdP as your **Identity Providers**. For more information about Idp configurations, see Configure a Facebook-based Single Sign-On (SSO).

Role required: admin

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Identity Providers**.
2. Select **OIDC_Facebook**.
3. On the OIDC_Facebook page, specify the following fields:

Note:

- Most of the fields are auto-populated when using the default IdP.
- Details of the ServiceNow Homepage must be provided.
- User-related details such as Client ID and Client Secret from Facebook must be provided within the IdP.

OIDC_Facebook Identity Provider Details

Field	Description
Name	Name of the OIDC IdP record. Enter `OIDC_Facebook`.
Default	Option to set the OIDC IdP configuration as default.
OIDC Entity Profile	OAuth Entity Profile for the OIDC configuration. Enter `Facebook default_profile`.
External logout redirect	The URL where the integration redirects users after they log out. Typically, this URL is the portal that is used for SSO. This field is automatically set to **external_logout_complete.do**. For example, `https://<yourinstance>.service-now.com/external_logout_complete.do`.
Active	Option to make the OIDC IdP configuration active. **Note:** This option can only be set to active after a successful test connection.
Auto Redirect IdP	Option to enable automatic redirection of the users to the login page of the identity provider.
ServiceNow Homepage	The URL of the login page that is used for authentication. This field is automatically set to your instance URL. The format of the URL is `https://<yourinstance>.service-now.com/navpage.do`

Field	Description
Show as Login option	Option to display the OIDC IdP as a login option on the login page. In this case, the login option appears as the **Login with Facebook** button.
SSO Label	Label of the OIDC IdP displayed on the login page. This field appears only when **Show as Login option** is enabled.
Logo URL	Publicly available URL that contains the logo of the OIDC IdP provider. This field appears only when **Show as Login option** is enabled.

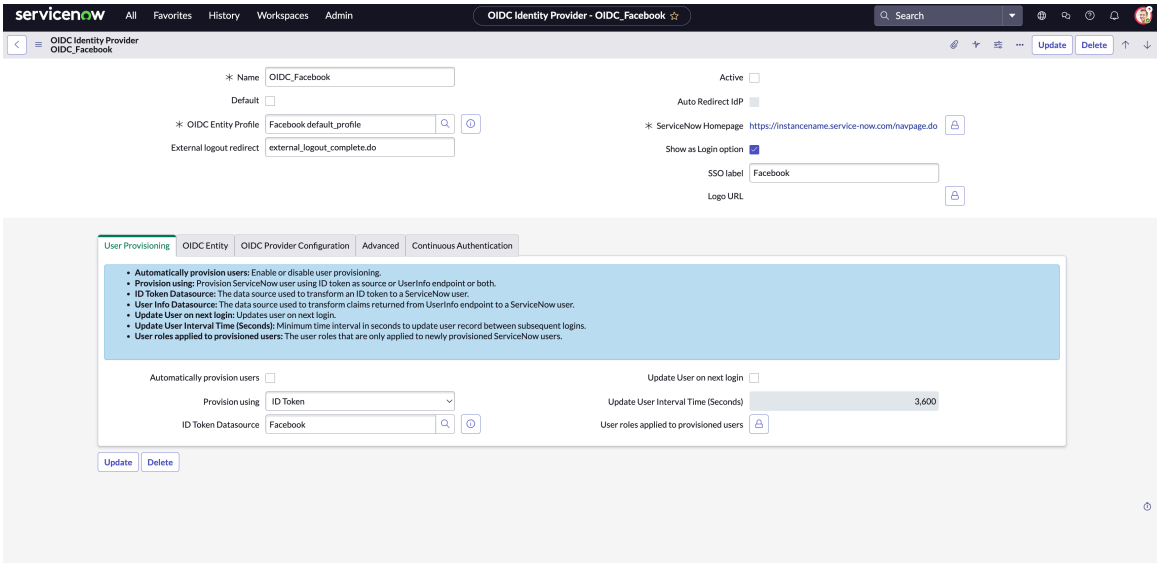

4. **Optional:** Open the **User Provisioning** tab, and fill in the fields.

> **Note:** You must configure the OIDC-related information such as Client ID and Client Secret of your users from Facebook.

User Provisioning tab

Field	Description
Automatically provision users	Option to enable automatic user provisioning. This property creates a user in the instance User [sys_user] table when the user exists on the IdP but doesn't exist in the User table. **Note:** You can choose to enable automatic user provisioning during user login. When automatic user provisioning is enabled, a user record is automatically created in the ServiceNow instance if that user record doesn't exist.
Provision using	The data source to use to transform to a ServiceNow user. Choices are as follows: ○ **An ID Token** ○ **User Info endpoint** ○ **Both ID Token and User Info**

Field	Description
	Use the Lookup list to select the pre-defined data source template, then open the record to configure the Transforms table mapping.
Provision data source	The ID token data source that is used for user provisioning.
User Info Datasource	The user info endpoint data source used for user provisioning. This field appears only when **User Info** or **Both ID Token and User Info** is selected from the **Provision using** field.
Update User on next login	Option to enable user updates during the next login.
Update User Interval Time (Seconds)	Minimum time interval (in seconds) to update a user record between subsequent logins. This field is automatically set to 3600 seconds. For example, after a user logs in, the user record will be updated after 3,600 seconds until the next login. This field is available only when the **Update User on next login** field is enabled.
User roles applied to provisioned users	List of roles applied to the newly provisioned users.

5. In the **OIDC Entity** tab, view and modify the OIDC client configuration and OIDC connect flow using the entity record.
   For more information related to OIDC-based configuration, see Configure an OAuth OIDC provider for accepting third-party token
6. In the **OIDC Provider Configuration** tab, view and modify the well-know configuration URL of the OIDC IdP.
7. **Optional:** Open the **Advanced** tab, and fill in the fields.

Advanced tab

Property	Description
Single Sign-On Script	The script that executes during Single Sign-On.
Logout Script	The script that executes after the user logs out.

Note: Scripts are run during single sign-on and logout.

8. To enable and test the configuration, click **Active**.
9. To update the record, click **Update**.
   The Facebook-based login option is displayed on the login form.
10. When logging in on the login form, do the following:
    a. Select the Facebook option.
    b. To log in to the ServiceNow instance, specify your Facebook credentials.

Configure a Facebook-based Single Sign-On (SSO)

Configure a Facebook-based SSO to your ServiceNow instance.

Before you begin

Have a valid Client ID that is configured as an IdP from Facebook.

Enable the following properties:

- Enable multiple provider SSO.
- Enable debug logging for the multiple provider SSO integrations.

Role required: admin

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Identity Providers**.
2. To create a new Facebook identity provider, click **New**.
3. Click **OpenID Connect**.
4. On the form, fill in the fields.

Import OpenID Connect Well Known Configuration form

Fields	Description
Name	Unique name for the OIDC identity provider configuration.
Client ID	The client ID of the application registered in the third-party OIDC identity provider.
Client Secret	The client secret of the application registered in the third-party OIDC identity provider.
Well known Configuration URL	The URL that contains metadata about the third-party OIDC identity provider.

5. Click **Import**.
   The Facebook-based IdP is created.
6. Select the Facebook IdP.
7. In the Facebook idP, do the following:
   a. Validate all the fields such as **Name**, **OIDC Entity Profile**, **External logout redirect**, and **ServiceNow Homepage**.
   b. Provide your **SSO label**.
8. In the **User Provisioning** tab, specify the fields that you need to configure users to specific user provisioning and roles.

   Only the mandatory fields are required. You can specify the remaining fields depending on what you need.

9. In the **OIDC Entity** tab, do the following:
   a. Click the entity.
   b. Set the **Redirect URL** field to your Facebook redirect URL.
10. In the **OAuth Entity Profiles** tab, do the following:

a. In the profile details, click a profile.

b. Select a scope and verify the details.
For example, select **scope-1**

11. In the **OAuth Entity Scopes** tab, click the **scope-1** link and add the scope as `email`.
12. To save the configuration, right-click the header and click **Save**.
13. To set the configuration as active, select **Active**.

Result
Users are displayed with the Facebook SSO option on the login form.

SAML

The Security Assertion Markup Language (SAML) is an XML-based standard for exchanging authentication and authorization data between security domains.

SAML exchanges security information between an identity provider (a producer of assertions) and a service provider (a consumer of assertions). SAML is a product of the OASIS Security Services Technical Committee. When implemented correctly, SAML is one of the most secure methods of single sign-on (SSO) available.

The SAML 2.0 integration enables SSO by exchanging XML tokens with an external Identity Provider (IdP). The IdP authenticates the user and passes a NameID token to the system. If the system finds a user with a matching NameID token (for example, the email address), the instance logs that user in.

If you are using the SAML 2.0 plugin for SSO authentication, you must set the glide.ui.rotate_sessions property to false. Otherwise, it interferes with the session information sharing that takes place between the instance and the Identity Provider. Users with the security_admin elevated privilege can access this property.

Note: It is recommended that customers using an existing SAML 2.0 integration upgrade to the Multi-Provider SSO plugin.

Multi-Provider SSO (SAML) IdP authentication flow

Describes the different entities that can authenticate a user through the SAML multi-SSO.

You can follow the authentication flow to understand when an entity authenticates a user using Multi-SSO.

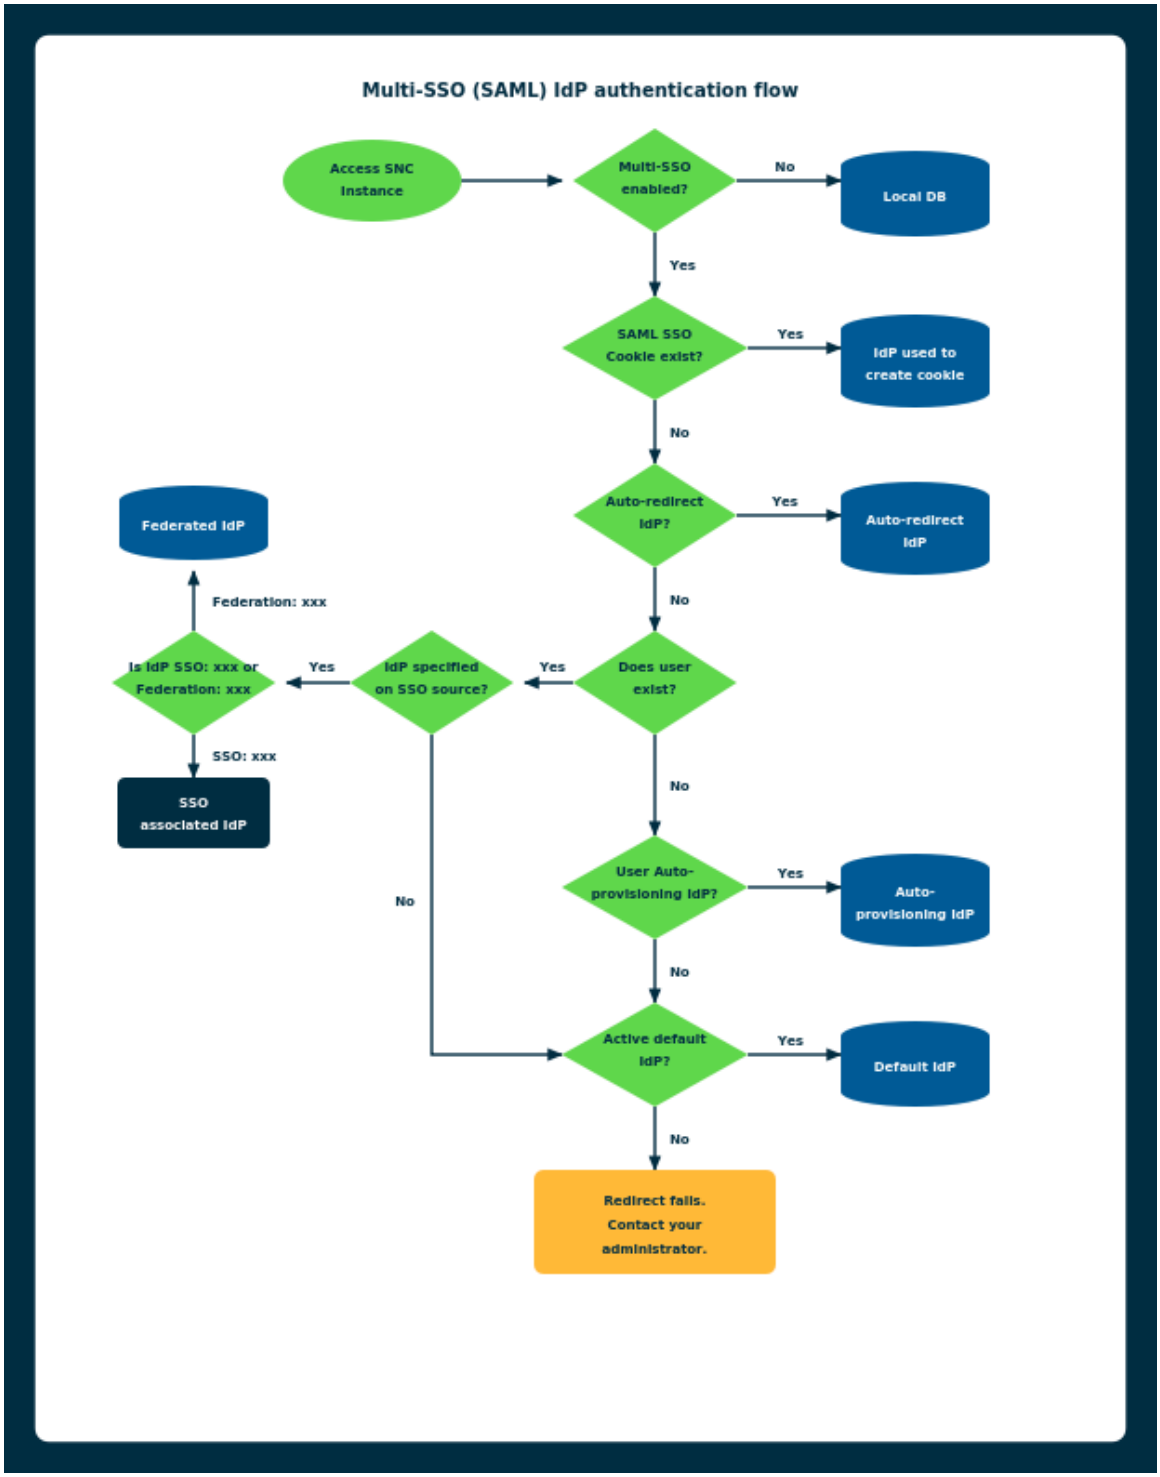

Local DB

If Multi-SSO is not enabled, authentication directs to a local DB.

SAML SSO Cookie IdP

If a SAML SSO cookie exists, the IdP which is specified with this cookie authenticates the user.

Auto-redirect IdP

If the auto-redirect IdP is enabled, this IdP authenticates the user.

Federated IdP

If the user browser is redirected to the external authorization (login_locate_sso.do) login screen, and the user exists in the user table with the IdP set in the **SSO Source** field as federation: xxx, then the federated IdP authenticates the user.

Associated IdP

If the user browser is redirected to the external authorization (login_locate_sso.do) login screen, and the user exists in the user table with the IdP set in the **SSO Source** field as sso: xxx, then the associated IdP authenticates the user.

Auto-provisioning IdP

If the user browser is redirected to the external authorization (login_locate_sso.do) login screen, and the user does not exist in the user table, but auto-provisioning is enabled, then the auto-provisioning IdP authenticates the user.

Note: If there is more than one auto-provisioning IdP enabled, the user can choose the auto-provisioning IdP they can use.

Default IdP

If the user browser is redirected to the external authorization (login_locate_sso.do) login screen, and the user either:

- Does not exist in the user table, auto-provisioning is not enabled, and there is an active default IdP
- Exists in the user table, an IdP is not specified on the SSO source user or company record, and there is an active default IdP

then the default IdP authenticates the user.

Identity Provider (IdP) system properties

An IdP generally offers an XML document containing their authentication and logout metadata.

For example, SSOCircle publishes their metadata online.

Browse the IdP metadata to find these entries:

- The `SingleSignOnService` element with a `Binding` attribute that contains a value of `HTTP-Redirect`. The `Location` attribute lists the URL the integration requires for the AuthnRequest service.
- The `SingleLogoutService` element with a `Binding` attribute that contains a value of `HTTP-Redirect`.The `Location` attribute lists the URL the integration requires for the SingleLogoutRequest service.

Note: The SAML 2.0 integration only supports binding to IdP services by HTTP-Redirect.

For example:

```
<SingleSignOnServiceBinding="urn:oasis:names:tc:SAML:2.0:binding
s:HTTP-Redirect"Location="https://
idp.ssocircle.com:443/sso/SSORedirect/metaAlias/ssocircle"/>
```

```
<SingleLogoutServiceBinding="urn:oasis:names:tc:SAML:2.0:binding
s:HTTP-Redirect"Location="https://
idp.ssocircle.com:443/sso/IDPSloRedirect/metaAlias/ssocircle"Res
ponseLocation="https://
idp.ssocircle.com:443/sso/IDPSloRedirect/metaAlias/ssocircle"/>
```

Set the IdP issuer URL

Provide the URL to the IdPs who will issue the security token.

Before you begin
Role required: admin

About this task
The integration verifies that each SAML response contains the same URL listed in this system property as the URL listed in the *Issuer* element. For example:

```
<samlp:Responsexmlns:samlp="urn:oasis:names:tc:SAML:2.0:protoc
ol"  Destination="https://demoi2.service-now.com/navpage.do"
  ID="s28da6774c88ae1eab292bf25fe625db81919d8e1e"
 InResponseTo="SNC841720c227c81948cfd68cadcad235c6"
 IssueInstant="2012-01-30T20:07:10Z"  Version="2.0"><saml:Issuer
 xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion">http://
idp.ssocircle.com</saml:Issuer>
   ...
   <saml:Assertion
 xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion"
     ID="s2f347f973c063836cf70ea38302d94976f9c5b851"
     IssueInstant="2012-01-30T20:07:10Z"
 Version="2.0"><saml:Issuer>http://
idp.ssocircle.com</saml:Issuer>
      ...
   </saml:Assertion></samlp:Response>
```

Procedure

1. Navigate to **All** > **SAML 2 Single Sign-on** > **Properties**.
2. In the property *The Identity Provider URL which will issue the SAML2 security token with user info.*, enter the URL to your IdP.

   Each IdP URL must be unique. By default, the integration contains the URL to SSOCircle. For more information, see http://idp.ssocircle.com.

Set the AuthnRequest service URL

Using the IdP's metadata, set the request service URLs for the integration's IdP.

Before you begin
Role required: admin

Procedure

1. In the property *The base URL to the Identity Provider's AuthnRequest service. The AuthnRequest will be posted to this URL as the SAMLRequest parameter*, enter the URL to the HTTP-Redirect binding obtained from the `SingleSignOnService` element.
2. Select the check box next to *Sign AuthnRequest* to enable the Identity Provider's single-sign on service to receive a signed AuthnRequest.
3. In the property *When SAML 2.0 single sign-on fails because the session is not authenticated, or this is the first login, redirect to this URL. This is the base URL where the initial SAML 2.0 AuthnRequest is sent using the SAMLRequest parameter*, enter the URL to the HTTP-Redirect binding obtained from the `SingleSignOnService` element.

   By default, the integration contains the URL to the SSOCircle service.

Set the SingleLogoutRequest service URL

Set the request service URLs for the integration's IdP by using the IdP's metadata.

Before you begin
Role required: admin

Procedure

1. In the *The base URL to the Identity Provider's SingleLogoutRequest service. The LogoutRequest will be posted to this URL as the SAMLRequest parameter* property, enter the URL obtained from the `SingleLogoutService` element.
   The LogoutRequest is posted to this URL as the SAMLRequest parameter. By default, the integration contains the URL to the SSOCircle service.
2. In the *URL to redirect users after logout, typically back to the portal that enabled the SSO (e.g. http://portal.companya.com/logout)* property, enter the URL where you want to redirect users after they successfully logout.

   If your IdP uses form-based authentication, enter the URL to your IdP's login form. If your IdP uses a non-form-based authentication method such as Kerberos, you should set the URL to a static logout page. This way, users who log out do not get immediately get redirected to the IdP and login again. By default, the integration contains the URL to the static UI page `external_logout_complete.do`.

(Optional) Enable signed logout requests

Some IdPs require the Service Provider to sign logout requests with a certificate.

Before you begin
Role required: admin

About this task
If your IdP requires signed logout requests, use the IdP's metdata to set the following system properties.

Procedure

1. In the **Advanced** tab, from the property *Sign LogoutRequest. Set this property to true if the Identity Provider's SingleLogoutRequest service requires signed LogoutRequest*, select **Yes** to specify that your IdP requires a signed logout request, or select **No** to use unsigned logout requests.
2. If you selected **Yes** to *Sign LogoutRequest*, then in *The protocol binding for the Identity Provider's SingleLogoutRequest service. (Value can be either "urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect" or "urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST".)* property, enter the one of the supported values listed in `Binding` attribute from the `SingleLogoutService` element.

   By default, the integration uses an HTTP-Redirect binding.
3. Click **Update**.
4. Install a Service Provider (SP) key store.

Service Provider (SP) system properties

These system properties define how the instance interacts with the IdP as a Service Provider.

Follow the sequential process to define IdP as Service Provider.

- Set the instance URL for SAML
- Set the audience URL for SAML
- Set up a NameID policy for SAML
- (Optional) Enable providing an authentication context class for SAML
- (Optional) Set keystore properties for signing logout requests for SAML

Set the instance URL for SAML

Set the instance-specific URLs so that the IdP can authenticate users.

Before you begin
Role required: admin

Procedure

1. In the property *The URL to the Service-now instance (usually this instance)*, enter the URL (including login page) of the instance for which the IdP authenticates.
   For example: https://yourinstance.service-now.com/navpage.do
2. In the property *The entity identification, or the issuer*, enter the base URL (excluding login page) of the instance for which the IdP authenticates.
   For example: https://yourinstance.service-now.com/

Set the audience URL for SAML

Enable your instance to verify that it is the intended recipient of a SAML response by using the Audience property.

Before you begin
Role required: admin

About this task
The integration verifies that each SAML response contains the same URL listed in this system property as the URL listed in the *Audience* element. For example:

```
<samlp:Responsexmlns:samlp="urn:oasis:names:tc:SAML:2.0:protoc
ol"  ID="s2cdc74f37f923e26fe1aeec42b70a93d24230334f"
  InResponseTo="90AA6073F01567BFB0DF194F596314E2"
 Version="2.0"  IssueInstant="2010-04-29T23:21:51Z"
 Destination="https://dloomac.service-now.com/navpage.do">
...
<saml:Conditions NotBefore="2012-01-30T19:57:10Z"
 NotOnOrAfter="2012-01-30T20:17:10Z"><saml:AudienceRestriction><
saml:Audience>https://
demoi2.service-now.com</saml:Audience></saml:AudienceRestriction
></saml:Conditions>
...
</samlp:Response>
```

Procedure

1. Navigate to **All** > **SAML 2 Single Sign-on** > **Properties**.
2. In the property *The audience uri that accepts SAML2 token. (Normally, it is your instance URI. For example: https://<instance name>.service-now.com.)*, enter the URL of your instance.

For example, https://demoi2.service-now.com. This URL must match the value of the *Audience* element in the SAML Response.

3. Click **Update**.

Set up a NameID policy for SAML

Set up a NameID policy for SAML. SAML 2.0 requires the IdP to exchange a NameID token with the service provider.

Before you begin
Role required: admin

About this task
For the SAML 2.0 integration the NameID token must map to a particular field in the User table. The integration uses the NameID token's value to determine what user the IdP authenticates.

Procedure

1. Browse the IdP metadata to find the `NameIDFormat` element that contains a value of `emailAddress`.
   The value of this element is the default format that the integration uses.
2. Review other `NameIDFormat` elements to determine if there are formats that match other fields in the User table.

Determine what User table field matches the NameID token

Identity providers specify what format the `NameID` token has.

Before you begin
Role required: admin

About this task
Setting up SAML 2.0 requires selecting a field from the User table that matches the format of the `NameID` token. Typically, IdPs offer the option to use an email address as the `NameID` token. Since the User table contains an email field, this field is a logical choice for use as a `NameID` token. To use another field from the User table as the `NameID` token, first verify that the IdP offers a `NameID` format that matches the value of a User table field. This may require adding the field to the User table.

Procedure

1. Compare the available formats in the IdP's `NameIDFormat` element to fields in the User table.
2. Select a `NameID` format where there is a matching value in the User table.
3. In the *The User table field to match with the Subject's NameID element in the SAMLResponse* field, enter the name of the User table field to search for matching values in the `NameID` token.

   By default, the integration uses the email field.

Set the IdP NameID policy

Specify what format the IdP uses for the `NameID` token.

Before you begin
Role required: admin

About this task
This format is listed as part of the IdP's metadata.

Procedure

1. In the property *The NameID policy to use for returning the Subject's NameID in the SAMLResponse. Your SAML identity provider will have to support this by declaring the policy in its metadata. The NameID value is used to match with the specified field in the User table to lookup the user.*, enter the value of the `NameIDFormat` element the integration uses.

   By default, the integration uses the SSOCircle `NameIDFormat` for email addresses.

2. Click **Save**.

Values in the User table field for SAML

Ensure that the integration's User table field contains appropriate matching values.

For example, if the integration uses the email field as the `NameID` token, ensure that the instance lists the same email address as the IdP. The integration fails to authenticate any user who does not have a matching value for the `NameID` token.

(Optional) Enable providing an authentication context class for SAML

You can enable the instance to send an authentication context class request to the IdP containing your instance's preferred authentication request format.

Before you begin
Role required: admin

About this task
If you enable creating an AuthContextClass message, then you must also specify an authentication context class reference format.

Note: Some IdP's do not allow the Service Provider to set the authentication context class. Disabling this setting allows the IdP to choose the authentication context class.

Procedure

1. From the property *Create an AuthnContextClass request in the AuthnRequest statement*, select **Yes** to specify a particular context class such as Password Protected Transport, or select **No** to have the IdP select the most appropriate context class.
2. If you selected **Yes** to *Create an AuthnContextClass request in the AuthnRequest statement*, then in *The AuthnContextClassRef method that we will request in our SAML 2.0 AuthnRequest to the Identity Provider* property, enter the URN of the context class you want to use for authentication (see table).

   #### AuthnContextClass URN options

Authentication type	Authentication context class URN
Forms-based authentication	`urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTranspor`
Kerberos-based authentication	`urn:federation:authentication:windows`

   By default, the integration uses a Password Protected Transport authentication method.

3. Click **Update**.

(Optional) Set keystore properties for signing logout requests for SAML

Set the keystore properties to enable the integration to sign logout requests by using your signed server and signed CA certificates.

Before you begin

Role required: admin

Procedure

1. In the property *The alias of key entry stored in SAML 2.0 SP Keystore used to sign SAML 2 requests*, enter the alias name that you created for the SAML 2.0 Keystore.
   By default, the integration looks for the alias *saml2sp*.
2. In the property *The password of key entry stored in SAML 2.0 SP Keystore used to sign SAML 2 requests*, enter the password for your SAML 2.0 Keystore.
   By default, the password is the same as the default alias name.
3. Click **Update**.
4. Regenerate your SP metadata.
   For more information, see SP metadata.

Create a service provider key store for SAML

Create a Java key store containing the following items for your instance to sign logout requests.

Before you begin

Role Required: admin.

About this task

- Signed server certificate for the instance
- Signed CA certificate
- Public and private key pair

You may create your own signed certificate with a private certificate authority or purchase one from a public certificate authority.

The following steps illustrate how to generate a new Java Keytool keystore file, create a certificate signing request (CSR), and import certificates. Any root or intermediate certificates need to be imported before importing the primary certificate for your domain. Type these commands in a command line interface.

Note: These instructions are not specific to the platform and require technical knowledge of security certificates to complete. Technical Support cannot assist in creating the certificates.

Procedure

1. Generate a Java keystore and key pair.

```
keytool -genkey -alias mydomain -keyalg RSA -keystore
 my.keystore
```

2. Generate a CSR for an existing Java keystore.

```
keytool -certreq -alias mydomain -keystore my.keystore -file
 mydomain.csr
```

3. Import a root or intermediate certificate authority CA certificate to an existing Java keystore.

```
keytool -import -trustcacerts -alias root -file Thawte.crt
 -keystore my.keystore
```

4. Import a signed primary certificate to an existing Java keystore.

```
keytool -import -trustcacerts -alias mydomain -file
 mydomain.crt -keystore my.keystore
```

Install a service provider keystore for signing SAML requests

Use the following steps to remove the existing example key store and install your own Service Provider key store containing your public and private key pair.

Before you begin

Role required: admin

Procedure

1. Create a Service Provider key store.
2. Navigate to **SAML 2 Single Sign-on** > **Certificate** or **Multi-provider** > **Administrator** > **x509 Certificate**.
3. Click **SAML 2.0 Keystore_Key2048_SHA256**.
4. Click the **Manage Attachments** link.
5. Select the Delete checkbox next to *saml2sp_key2048withsha256.jks*.
6. Click **Remove**.
7. Click **Choose Files** and select the Keystore containing your signed certificates.
8. Click **Attach**.
9. Close the Attachments popup.

   Note: It is recommended to provide different name for the certificate that is attached newly.

10. In *Key store password*, enter the password to access the SAML 2 alias.
11. Click **Update**.

Create self-signed BCFKS keystore for SAML

Generate a FIPS 140-2 compliant self-signed BCFKS keystore for use in SAML signing and encryption operations within the Multi-Provider SSO plugin.

Before you begin

Role required: admin

Do the following:

- Install Java on your machine and the key tool command-line tool accessible in your terminal (or "command prompt" if you are running it on a windows machine).
- Perform the following steps to create a keystore using FIPS-approved cryptographic algorithms (such as RSA 2048-bit or higher paired with SHA-256) that meets federal security requirements for identity federation and single sign-on implementations.

Procedure

1. Download the FIPS Provider Library.

Note:

Use the latest version `bc-fips-2.1.0.jar`. Make sure you use the most recent version.

2. Generate the FIPS-compliant keystore and certificate.

   a. Run the following key tool command to generate a self-signed certificate and keystore.

Key tool command

Running on Linux/macOS	Running on Windows:
`keytool -genkeypair \ -providername BCFIPS \ -providerclass org.bouncycastle.jcajce.provider.BouncyCastleFipsProvider \ -providerpath <path_to_bc-fips-<version>>.jar \ -alias <key_alias> \ -keyalg RSA \ -keysize <key_size> \ -keystore <keystore_name>.bcfks \ -validity <validity> \ -storetype BCFKS \ -storepass <keystore_password>`	`keytool -genkeypair ^ -providername BCFIPS ^ -providerclass org.bouncycastle.jcajce.provider.BouncyCastleFipsProvider ^ -providerpath <path_to_bc-fips-<version>>.jar ^ -alias <key_alias> ^ -keyalg RSA ^ -keysize <key_size> ^ -keystore <keystore_name>.bcfks ^ -validity <validity> ^ -storetype BCFKS ^ -storepass <keystore_password>`

   b. Replace placeholders (`<...>`) with appropriate values:
   - `<path_to_bc-fips-<version>>.jar`: Path to `bc-fips-<version>.jar`
   - `<key_alias>`: Alias for the key pair
   - `<key_size>`: 2048 or 4096
   - `<keystore_name>.bcfks`: Desired file name for the keystore
   - `<validity>`: Expiry in days
   - `<keystore_password>`: Password for the keystore

   c. Follow the prompts to enter additional DN (Distinguished Name) details for the certificate.

Note: When you are prompted for a password for the key (alias), press the Enter or Return key to use the same password you used for the keystore. Do not give a different password.

   d. Securely store the key alias and keystore password.
   Provide these credentials while:

- Creating the `sys_certificate` record for this keystore.
- Configuring the SAML Identity Provider to provide the signing key or encryption key alias and password.

Note: The key password is same as the keystore password specified during creation. Use the same password when configuring signing or encryption for the SAML Identity Provider.

3. Extract the Certificate Chain.

Key tool command

Running on Linux/macOS	Running on Windows:
`keytool -exportcert \ -provider org.bouncycastle.jcajce.provider.BouncyCastleFipsProvider \ -providerpath <path_to_bc-fips-<version>>.jar \ -storetype BCFKS \ -keystore <keystore_name>.bcfks \ -storepass <keystore_password> \ -alias <key_alias> \ -rfc \ -file <file_name>.cer`	`keytool -exportcert ^ -provider org.bouncycastle.jcajce.provider.BouncyCastleFipsProvider ^ -providerpath <path_to_bc-fips-<version>>.jar ^ -storetype BCFKS ^ -keystore <keystore_name>.bcfks ^ -storepass <keystore_password> ^ -alias <key_alias> ^ -rfc ^ -file <file_name>.cer`

Replace placeholders `(<...>)` with appropriate values:

- `<path_to_bc-fips-<version>>.jar`: Path to `bc-fips-<version>.jar`
- `<keystore_name>.bcfks`: keystore file name as given in previous step
- `<keystore_password>`: keystore password as given in previous step
- `<key_alias>`: Key alias as given in previous step
- `<file_name>.cer`: Desired file name for the extracted certificate in PEM format

4. Create a record on `sys_certificate` table.

   a. Log in to ServiceNow AI Platform.

   b. Navigate to **All** > **Multi-Provider SSO** > **Administration** > **x509 Certificate**.

   c. Click **New** to create a record.

   d. Select **BCFKS keystore** as Type.

   e. Attach the generated BCFKS keystore file `(<keystore_name>.bcfks)`.

Note: To configure certificate expiry notification, use **Notify on expiration** and **Groups to notify on expiration**, and set the notification timing using **Warn in days to expire** and **Frequency**.

f. Fill in other required fields, including the keystore password provided during keystore creation.

g. Click **Validate Stores/Certificates** related link to ensure the keystore is valid.

h. Copy the `sys_id` of this record.

(Optional) Advanced SAML properties

The following advanced settings allow you to further increase security and debug the integration.

Advanced settings

Navigate to **All** > **SAML 2 Single Sign-on** > **Properties**.

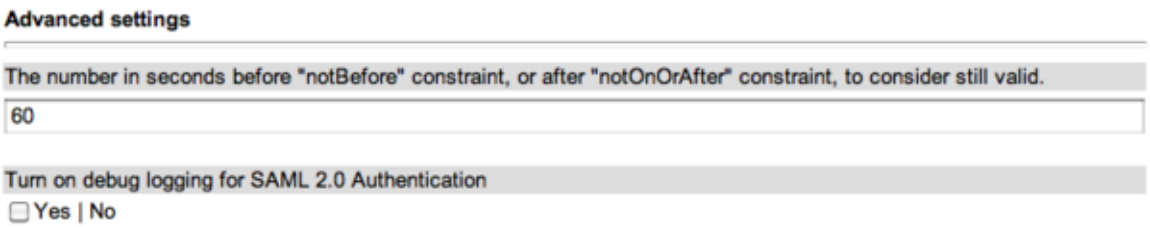

Advanced system properties table

Property	Description
The number of seconds "notBefore" constraint, or after "notOnOrAfter" constraint, to consider still valid	Enter the number of seconds to add to the *NotBefore* and *NotOnOrAfter* constraints to account for time differences between the IdP clock and SP clock. These constraints prevent against replay attacks by denying requests that are not made within the specified time frame. If the IdP clock and SP clock are significantly different, network latency may result in the SAML request being unauthorized. This property adds a grace period during which SAML requests and responses are still considered valid.
Turn on debug logging for SAML 2.0 Authentication	Select **Yes** to enable additional logging information for SAML 2.0 events.

Install the identity provider certificate

You can paste a PEM certificate into a X.509 Certificate form so the identify provider can verify communications with the service provider.

Before you begin

Role required: admin

About this task

The IdP's certificate is located within the IdP's metadata. The IdP developer determines where the certificate metadata resides when creating the local IdP.

Note: Certificates for single-sign on should always be in PEM format to work with SAML certificates.

Procedure

1. Navigate to **All** > **SAML Single Sign-on** > **Certificate**.
2. Fill in the form fields (see table).
3. Click **Save**.

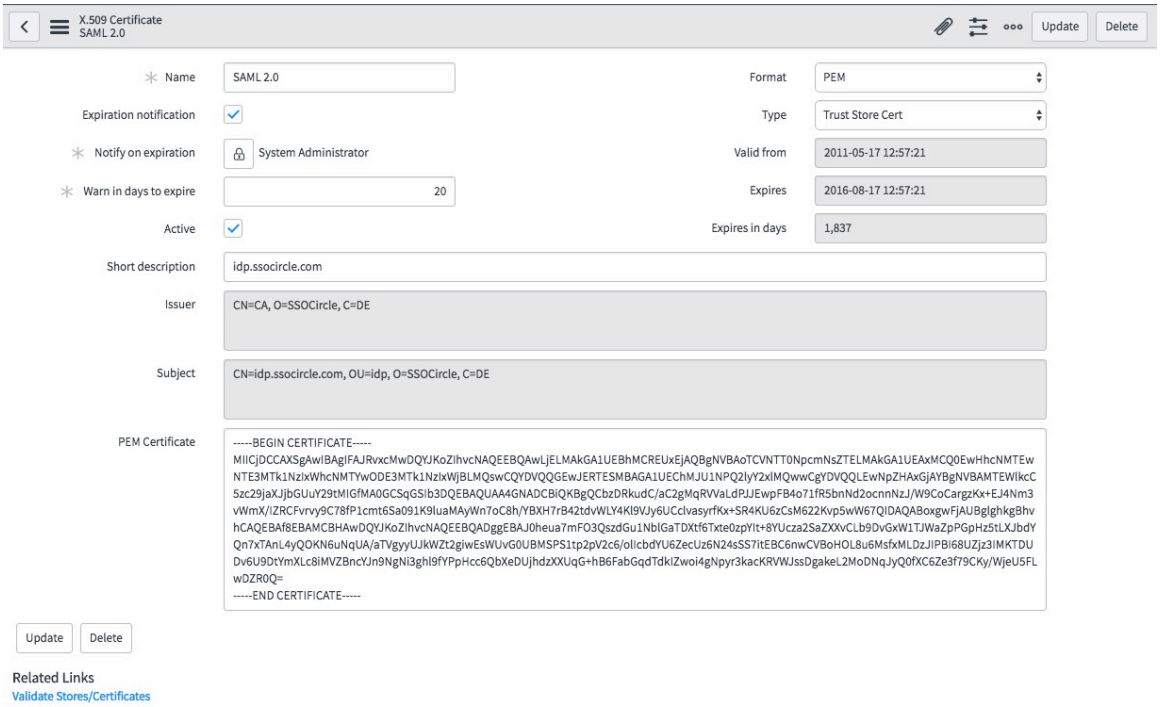

Note: The integration does not currently sign the certificate in communications between the instance and the IdP.

Field	Description
Name	The certificate name. Do not change the *Name* entry. The name of the X.509 certificate must be *SAML 2.0* in order for the integration to use it. This requirement is only true if you are not using Multi-Provider single sign-on (SSO).
Expiration notification	Select this option to send a notification to the users selected in the **Notify on expiration** field. By default, this is enabled.
Notify on expiration	Select the users to revive the notification regarding certificate expiration. If no users are selected, the logged in user is added by default, along with the last two logged in users with the administrator role.
Warn in days to expire	The number of days before expiration that the instance send the notification. Enter a value of at least 20. Instances upgraded to Istanbul and later releases have this value set to 20 unless a greater value is specified.
Active	A check box to indicate that this certificate is active.
Format	The format of the certificate. SAML uses the PEM format.
Type	The certificate container. The instance recognizes certificates from trust stores, Java keystore, and PKCS#12 keystores.

Field	Description
Valid from	The instance automatically adds the certificate valid from date to this field. Attach the certificate to the X.509 certificate record to populate this field.
Expires	The instance automatically adds the certificate expiration date to this field. Attach the certificate to the X.509 certificate record to populate this field.
Expires in days	The calculated number of days to expiration.
Short description	A description for the certificate.
Issue	The instance automatically adds the certificate issuer to this field. Attach the certificate to the X.509 certificate record to populate this field.
Subject	The instance automatically adds the certificate subject to this field. Attach the certificate to the X.509 certificate record to populate this field.
PEM Certificate	Enter the value of the X509 certificate.

What to do next
Click **Validate Stores/Certificates** to test the trust store and certificate.

Replace a missing certificate for SAML

If the **Certificate** module displays a blank page, the SAML 2.0 certificate record has been deleted. You can replace the missing certificate by manually creating a certificate record.

Before you begin
Role required: admin

Procedure
1. Navigate to **All** > **System Definition** > **Certificates**.
2. Create a new record called `SAML 2.0`.

> **Important:** You MUST use this name. This requirement is only true if you are not using Multi-Provider single sign-on (SSO).

3. Click **SAML 2 Single Sign-on** > **Certificate**.
4. In the *PEM Certificate* field, enter the value of the `ds:X509Certificate` element from your IdP's metadata.
5. Click **Save**.

Test the SAML integration

Test the SAML integration after you complete all the other setup tasks.

Before you begin
Role required: admin

Procedure
1. Log in to the instance as a user with the admin role.
2. Navigate to **SAML 2 Single Sign-on** > **Properties**.

3. In the property *Enable external authentication*, select **Yes**.

Note:

Enabling external authentication requires all users to use SAML 2.0 single sign-on. If anyone tries to access the application in an unauthenticated state, the instance automatically sends an authentication request to the (IdP) and redirects the user to the SAML IdP Authentication page.

4. Click **Save**.
5. Log out of the instance.
6. Browse to the instance URL.
   If the integration is functioning properly, the IdP should ask for the user's credentials.

Related topics

Multi-SSO (SAML 2.0) errors and fixes

Multi-SSO (SAML 2.0) errors and fixes

A list of common errors and associated fixes for a Multi-SSO (SAML 2.0) setup and configuration.

Errors during Multi-SSO (SAML 2.0) setup

Error in instance logs	Test Connection Message	SAML property	Diagnosis	Fix
NotAfter: <Thu Jun 05 22:57:44 PDT 2014>.	Ensure that the IDP x509 certificate is present, valid, and active.	N/A	The current certificate or the SAML assertion has expired.	• Sync the SNC clock with the SAML IdP server clock. • Update the SAML 2.0 certificate record.
• Unable to locate SAML 2.0 certificate. • Could not find a digital signature stored in the ServiceNow instance.	Ensure that the IDP x509 certificate is present, valid, and active	The PEM-formatted string should be entered into the PEM Certificate field.	The SAML certificate does not exist. It might be inactive.	Ensure that the correct PEM-formatted certificate is uploaded to the instance.
Certificates do not match. Expect: <certStr>, actual: <inboundCert>.	Ensure that the IDP x509 certificate is present, valid, and active.	N/A	The available certificate in SNC does not match the certificate in assertion. Causes include: • The certificate is updated on the IdP but not in	Confirm that the PEM-formatted string in the SAML 2.0 certificate record matches the X509 Certificate in the SAMLResponse for the user IdP.

Errors during Multi-SSO (SAML 2.0) setup (continued)

Error in instance logs	Test Connection Message	SAML property	Diagnosis	Fix
			the ServiceNow instance. • The certificate is in the wrong format.	
Failure to check the validity of the certificate.	Ensure that the IDP x509 certificate is present, valid, and active	N/A	The current certificate might have expired.	Update the SAML 2.0 certificate record.
Failure to validate signature profile.	Ensure that the IDP x509 certificate is present, valid, and active.	N/A	The assertion might be signed with a different certificate.	Check if the IdP has the same certificate as the SNC instance.
InResponseTo attribute in SubjectConfirmationData mismatch. Expect: <inResponseTo>, actual: <inResponseTo>.	Subject confirmation validation failed.	N/A	This error appears if either of the following situations occurs: • The IdP returns a SAMLResponse for a different SAMLRequest • A user bookmarks the URL with the SAMLRequest instead of just the instance URL • If a null value is expected, the response might be sent to a different node when the instance has multiple nodes.	The IdP admin should confirm that the expected SAMLReponse is being returned. This situation can be a load balancer or infrastructure issue.
SessionIndex value not found: <message>...	SessionIndex not valid.	N/A	The SessionIndex is required in the SNC instance. The IdP returns it in the SAML response to authenticate successfully.	The IdP admin should confirm that the SessionIndex is defined in the SAMLResponse.

Errors during Multi-SSO (SAML 2.0) setup (continued)

Error in instance logs	Test Connection Message	SAML property	Diagnosis	Fix
No valid SubjectConfirmation found.	Subject confirmation validation failed.	N/A	Conditions could be missing due to an error on the IdP. The StatusCode in the response would contain Responder instead of the expected Success.	Review SAMLResponse to determine if Conditions are included in the SAMLResponse. The valid subject confirmation data could be expired or not for the right audience.
Assertion audience mismatch. Expect: <propAudience>, actual: <audienceUri>. or AudienceRestriction validation failed. No matching audience found.	Ensure that the 'Audience URI' field is set correctly	The audience URI that accepts the SAML2 token. (Normally, it is your instance URI. For example: `https://demo.service-now.com`.)	The SNC instance configured audience URI must match the value in the IdP.	Locate <saml2:Audience> in the SAMLResponse in the logs and verify that the value matches the one on the instance.
Assertion issuer is invalid. Expect: <value on instance>, actual: <value returned by IdP>	Assertion issuer is invalid.	The Identity Provider URL that issues the SAML2 security token with user info.	The IdP entity id (issuer) does not match the value defined in the SNC instance.	• Check if IdP or SP is not configured properly. • Confirm that the SAML property (the Identity Provider URL that issues the SAML2 security token with user info) is set correctly.
Subject is valid in the future. Now: <now>, NotBefore: <notBefore> or Subject is expired. Now: <now>,	Subject validation confirmation failed.	The number in seconds before notBefore constraint, or after notOnOrAfter constraint, to	The IdP clock is not synced with SP clock.	Update the SAML property glide.authenticate.sso.saml2.clockskew to a larger value. The default is 180 seconds. Some cases require a setting of 300 or higher. You may also need to

Errors during Multi-SSO (SAML 2.0) setup (continued)

Error in instance logs	Test Connection Message	SAML property	Diagnosis	Fix
NotOnOrAfter: <notOnOrAfter>		consider still valid.		check the time on your IdP server.
Assertion is valid in the future, now: <now>, notBefore: <notBefore> or Assertion is expired, now: <now>, notOnOrAfter: <notOnOrAfter>	Assertion is invalid.	The number in seconds before notBefore constraint, or after notOnOrAfter constraint, to consider still valid.	IdP clock is not synced with SP clock	Update the SAML property to a larger value. Default of 60 seconds. Some cases require a setting of 300 or higher. You may also need to check the time on your IdP server.

Common login and IdP errors

Error or Symptom	Diagnosis	Fix
Login requests generate an infinite loop between the system and the IdP when High Security is active.	• Typically the URL endpoint is an error page or logout page. • The logout_redirect.do might create this loop when you define *glide.security.url.whitelist* without adding the IdP host name to the property value. **Note:** To learn more about this property, see Enforce URL allowlist check [Updated in Security Center 1.3, 1.5, and 2.0] in Instance Security Hardening Settings.	Set (or create) the system property glide.authenticate.failed_redirect to redirect failed authentication requests to this URL.
The token used to authenticate the user or the request is signed with the signature algorithm http://www.w3.org/2001/04/xmldsig-more#rsa-sha256 which is not the expected signature algorithm http://www.w3.org/2000/09/xmldsig#rsa-sha1.	Check the Alert Context tab for event details.	Navigate to the Advanced tab of the Relying Party Trust configuration dialog and verify that the algorithm is set to SHA-1 and not SHA-256.

Common login and IdP errors (continued)

Error or Symptom	Diagnosis	Fix
The error message `urn:oasis:names:tc:SAML:2.0:status:Requester` appears in your system log (syslog) table.	When your IdP (e.g. ADFS) responds with a status of `urn:oasis:names:tc:SAML:2.0:status:Requester`, it means the IdP rejected the login because of an issue with the request sent to it. Unfortunately, the SAML response received from the IdP in most cases does not provide further details for the error.	Review the SAML request sent to the IDP, and work with your IDP administrator, update your instance SAML settings to avoid the error. You may need to contact your IDP provider understand the reason for the login failure.

Redirect single sign-on (SSO) logins

When SSO is enabled, you can redirect users to specific pages or direct users to login locally.

For example, if a user attempts to go to `https://customerX.service-now.com`, an internal company portal can display instead of the default login page. Or, when a user logs out of an application, the browser can redirect them to a specific internal page. You can set redirection properties within the instance to ensure that users see an SSO login page rather than the default login page.

Note: The following properties do not force SSO. The login.do page is still accessible and users can login to the system if they have a local password set.

Redirection properties

When a user logs out, or if there is a failed attempt to sign on using SSO, you can define where the user is taken next, such as a main portal page or a knowledge base article with SSO login information. Use the following properties to specify the URLs. If one of these properties does not exist in your instance, you can create the property.

glide.authenticate.honor.sso_record.failed_requirement_redirect

URL to redirect users when they attempt to access a page that is private (for example, to view an incident) and do not provide SSO credentials. The property is typically set to a customer's login portal (for example, `http://portal.companya.com/`).

glide.authenticate.failed_redirect

URL to redirect users after a failed SSO attempt. You can redirect to a public knowledge article that describes the error and has helpful links (for example, `http://portal.companya.com/error`).

glide.authenticate.external.logout_redirect

URL to redirect users after logging out, typically back to the portal that enabled the single sign-on log in (for example, `http://portal.companya.com/logout`).

glide.authentication.external.disable_local_login

When set to true, requires SSO credentials for the main login page. Defaults to false. This property needs to be used in conjunction with the *`glide.authenticate.honor.sso_record.failed_requirement_redirect`* property.

The following table shows the relationship between the Installation Exit return values, the properties, and the expected behavior.

Forcing login using SSO only

Return value	Property
`failed_missing_requirement`	*glide.authenticate.honor.sso_record.failed_requir*
`failed_authentication`	*glide.authenticate.failed_redirect*
`<user_id>`	N/A

Restricting local login

As a security precaution, you should do more than rely on redirection properties to prohibit logging in locally. If a user should never log in locally and will always be authenticated by your internal single sign-on system, then a random password should be assigned to each user that is imported into the instance. The random password is most easily set at the time of the user import. If the user data is imported into your system through an import set, you can create an onBefore transform script using the following code .

```
var r  = new Packages. java. util. Random ( ) ;

 var str1  = Packages. java. lang. Long. toString (Packages.
 java. lang.
 Math. abs (r. nextLong ( ) ) , 36 ) ; var str2  = Packages.
 java. lang.
 Long. toString (Packages. java. lang. Math. abs (r. nextLong
 ( ) ) , 36
 ) ;

 var newPass  = str1  + str2 ;

target. user_password = newPass ;

 //password now set to a random string like this:
 //qvm81zdrn7cwwylpvw94eebk
```

Clone an instance with a SAML integration

Clone an instance with a SAML integration. Before you clone an instance that uses SAML 2.0, preserve the SAML SSO-related settings on the target instance or you might make the target instance inaccessible.

Before you begin

Role required: admin

Procedure

1. On the source instance, navigate to **System Clone** > **Preserve Data** > **Core Instance Properties**.
2. Make sure that the following SAML SSO-related properties are preserved using conditions.
   - glide.authenticate
   - glide.security
   - glide.entry
   - glide.script
   - glide.session
   - glide.saml2
   - com.glide.communications
   - com.snc.integration.saml_esig

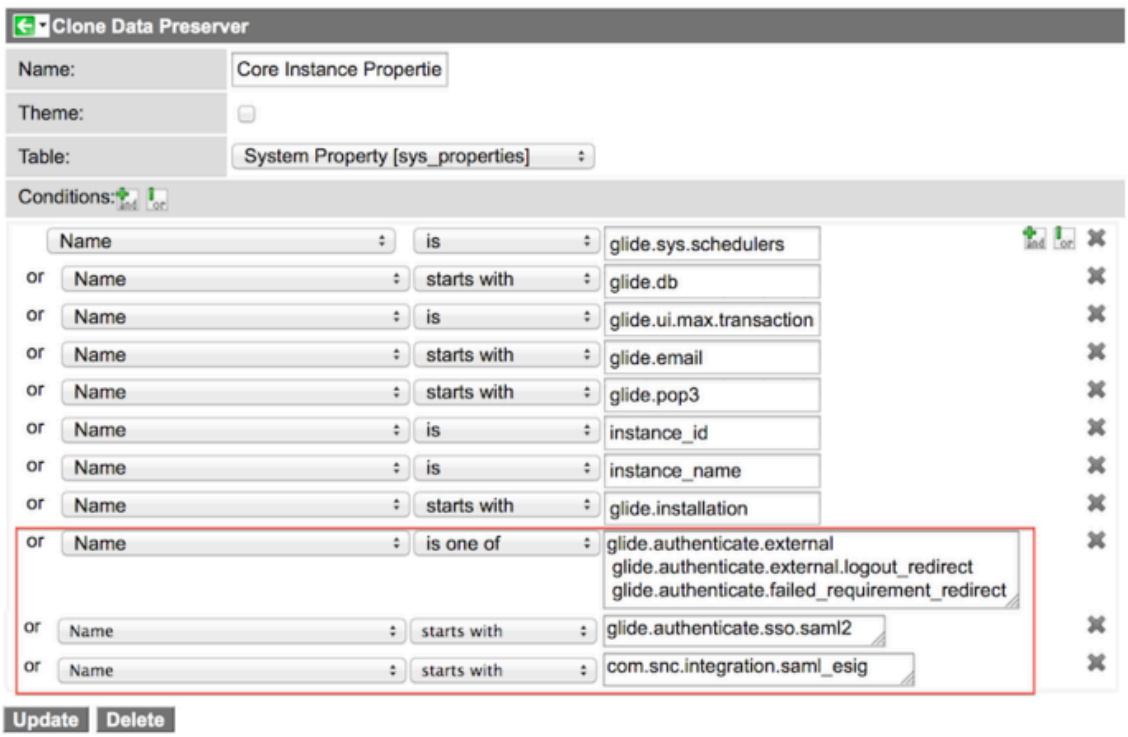

Note: When you create the clone, include attachments so that certificates carry over to the target instance. Also, make sure the **Theme** check box is cleared so these properties are preserved regardless of whether you preserve the instance theme.

3. On the source instance, navigate to **System Clone** > **Preserve Data** to preserve the SAML certificates on sys_certificate and SAML users on sys_user related to SAML/SSO/Multi SSO.

If you need them, export them into XML, then manually import them on the target.

> ⚠ **Warning:** Do not try to clone the SAML/SSO/Multi SSO setup from one system to another. Most transfers of SAML/SSO or Multi SSO settings do NOT work because they must be configured on the identity provider. If you overwrite a working setup, the target instance will fail to authenticate so your target instance will become inaccessible. Also, do not change the sys_id of your Multi SSO provider record; doing that will force your users to flush their cookies. For more information about cloning precautions, see Checklist before cloning an instance.

4. Exclude the Multi SSO tables sso_properties, digest_properties and saml2_update1_properties.
5. Manually create the SAML/SSO/Multi SSO records on each instance independently and set up the records on your identity provider as well.
6. Make sure that you manually create a LOCAL admin account on sys_user (not in LDAP or SAML) record on the target instance and with a sys_id that does not exist on the source instance.
7. Click **Update**.

SAML 2.0 concepts

Familiarize yourself with these SAML concepts.

Typical SAML process flow (diagram)

A typical SSO logic flow involves looking for an active session, checking user credentials, and creating the necessary token.

SSO Standard

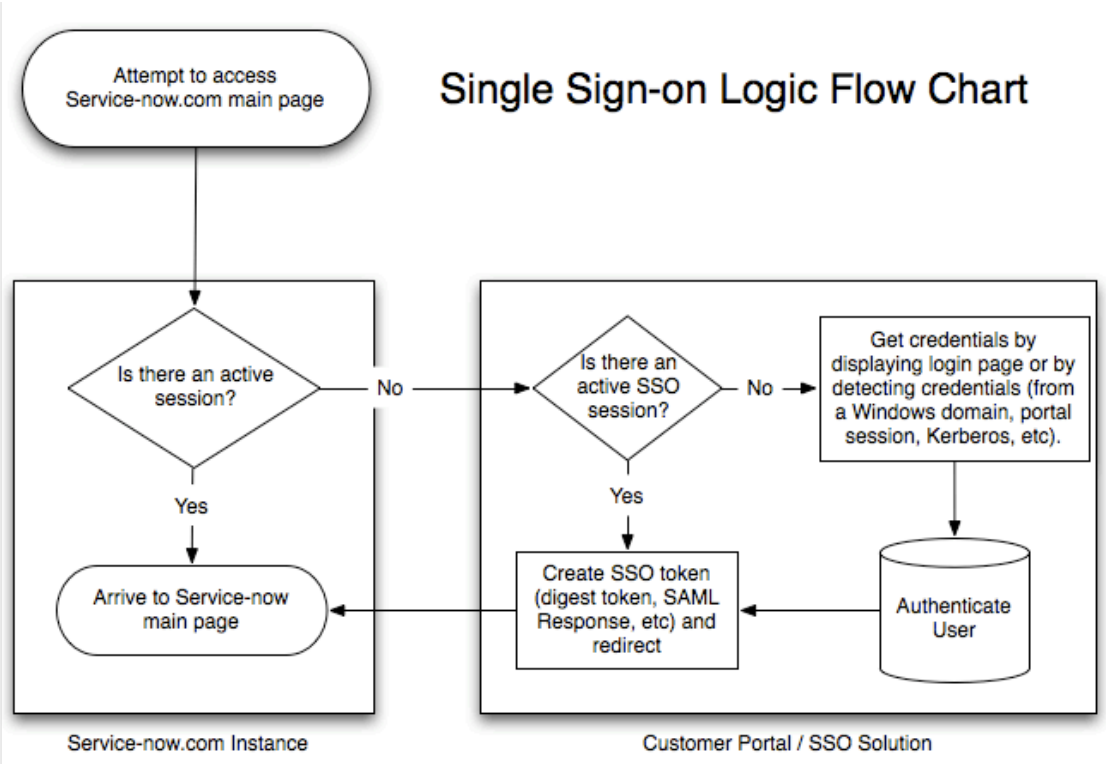

Login (AuthnRequest) process flow

SAML 2.0 specifies a Web Browser SSO Profile that involves exchanging information among an identity provider (IdP), a service provider (SP), and a principal (user) on a web browser.

The identity provider can be any SSO service offering SAML authentication services (for example SSOCircle). The service provider is always an instance. The message flow begins with a request for a secured resource at the service provider.

Request the target resource at the SP

The principal requests a target resource at the service provider:

https://instance.service-now.com/

The instance checks the request to see if the SAMLRequest and RelayState URL parameters are present. If they exist, the user has already validated with the IdP and can skip steps 2–6.

Issue AuthnRequest to Identity Provider

The instance constructs an `AuthnRequest` to be sent to the IdP using the `SAMLRequest` value. The instance also constructs and sends a `RelayState` URL parameter value.

The `RelayState` token is an opaque reference to state information maintained at the service provider. The value of the `SAMLRequest` parameter is the deflated and base64 encoded value of the `<samlp:AuthnRequest>` element:

```
<samlp:AuthnRequest
 xmlns:samlp="urn:oasis:names:tc:SAML:2.0:protocol"
   xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion"
    ID="identifier_1"    Version="2.0"
  IssueInstant="2004-12-05T09:21:59Z"
 AssertionConsumerServiceIndex="0"><saml:Issuer>https://
sp.example.com/SAML2</saml:Issuer><samlp:NameIDPolicy
 AllowCreate="true"
 Format="urn:oasis:names:tc:SAML:2.0:nameid-format:transient"/><
/samlp:AuthnRequest>
```

The integration then URL-encodes the `<samlp:AuthnRequest>` element and sends it as the `SAMLRequest` URL parameter.

The SSO service processes the `<samlp:AuthnRequest>` element by URL-decoding, base64-decoding and inflating the request, in that order. It then performs a security check. If the user does not have a valid security context, the IdP identifies the user by prompting for login credentials. If the user is already logged in, the IdP simply responds with the `SAMLResponse<tt> and <tt>RelayState` URL parameters (see step 3).

Respond with an SAMLResponse and RelayState

After collecting the required login credentials, the SSO service validates the request and responds with a document containing an XHTML form:

```
<formmethod="post"action="https://
instance.service-now.com/navpage.do" ...><input type="hidden"
 name="SAMLResponse" value="response ..." /><input type="hidden"
 name="RelayState" value="token ..." />
    ...
    <input type="submit" value="Submit" /></form>
```

The value of the `RelayState` parameter comes from this step. The value of the `SAMLResponse` parameter is the base64 encoding of the following `<samlp:Response>` element:

```
<samlp:Responsexmlns:samlp="urn:oasis:names:tc:SAML:2.0:protoc
ol"  ID="s2cdc74f37f923e26fe1aeec42b70a93d24230334f"
  InResponseTo="90AA6073F01567BFB0DF194F596314E2"
 Version="2.0"  IssueInstant="2010-04-29T23:21:51Z"
 Destination="https://
dloomac.service-now.com/navpage.do"><saml:Issuer
 xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion">http://
idp.ssocircle.com</saml:Issuer><samlp:Status
 xmlns:samlp="urn:oasis:names:tc:SAML:2.0:protocol"><samlp:Statu
sCode xmlns:samlp="urn:oasis:names:tc:SAML:2.0:protocol"
 Value="urn:oasis:names:tc:SAML:2.0:status:Success"></samlp:Stat
usCode></samlp:Status><saml:Assertion
 xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion"
    ID="s23e536bfc51b8487d4d3299dec162d9c2e338823b"
    IssueInstant="2010-04-29T23:21:51Z"
 Version="2.0"><saml:Issuer>http://
idp.ssocircle.com</saml:Issuer><Signature
 xmlns="http://www.w3.org/2000/09/xmldsig#">

...
      </Signature><saml:Subject><saml:NameID
 Format="urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress"
           NameQualifier="http://idp.ssocircle.com"
 SPNameQualifier="https://
dloomac.service-now.com/navpage.do">david.loo@service-now.com</s
aml:NameID><saml:SubjectConfirmation
 Method="urn:oasis:names:tc:SAML:2.0:cm:bearer"><saml:SubjectCon
firmationData
 InResponseTo="90AA6073F01567BFB0DF194F596314E2"
      NotOnOrAfter="2010-04-29T23:31:51Z"
 Recipient="https://dloomac.service-now.com/navpage.do" /
></saml:SubjectConfirmation></saml:Subject><saml:Conditions
 NotBefore="2010-04-29T23:11:51Z"
 NotOnOrAfter="2010-04-29T23:31:51Z"><saml:AudienceRestriction><
saml:Audience>https://
dloomac.service-now.com</saml:Audience></saml:AudienceRestrictio
n></saml:Conditions><saml:AuthnStatement
 AuthnInstant="2010-04-29T23:21:51Z"
 SessionIndex="s2dbf89ab99001e0e8cdaed67266d9d4b21b968a04"><saml
:AuthnContext><saml:AuthnContextClassRef>urn:oasis:names:tc:SAML
:2.0:ac:classes:PasswordProtectedTransport</saml:AuthnContextCla
ssRef></saml:AuthnContext></saml:AuthnStatement></saml:Assertion
></samlp:Response>
```

Validate SAMLResponse

The `SAMLResponse` value is base64 decoded and inflated to reveal the XML document in step 3. The login script extracts the XML value from the `//Subject/NameID` element and uses it to look up an existing user in the User table.

The login script also extracts the session ID from the `//AuthnStatement/@SessionIndex` element and stores it for the `LogoutRequest`.

Logout (LogoutRequest) process flow

During logout, the instance issues the SAML 2.0 *LogoutRequest* service call to the IdP.

This service logs the user out and then redirects back to the specified logout URL.

SAML 2 Logout

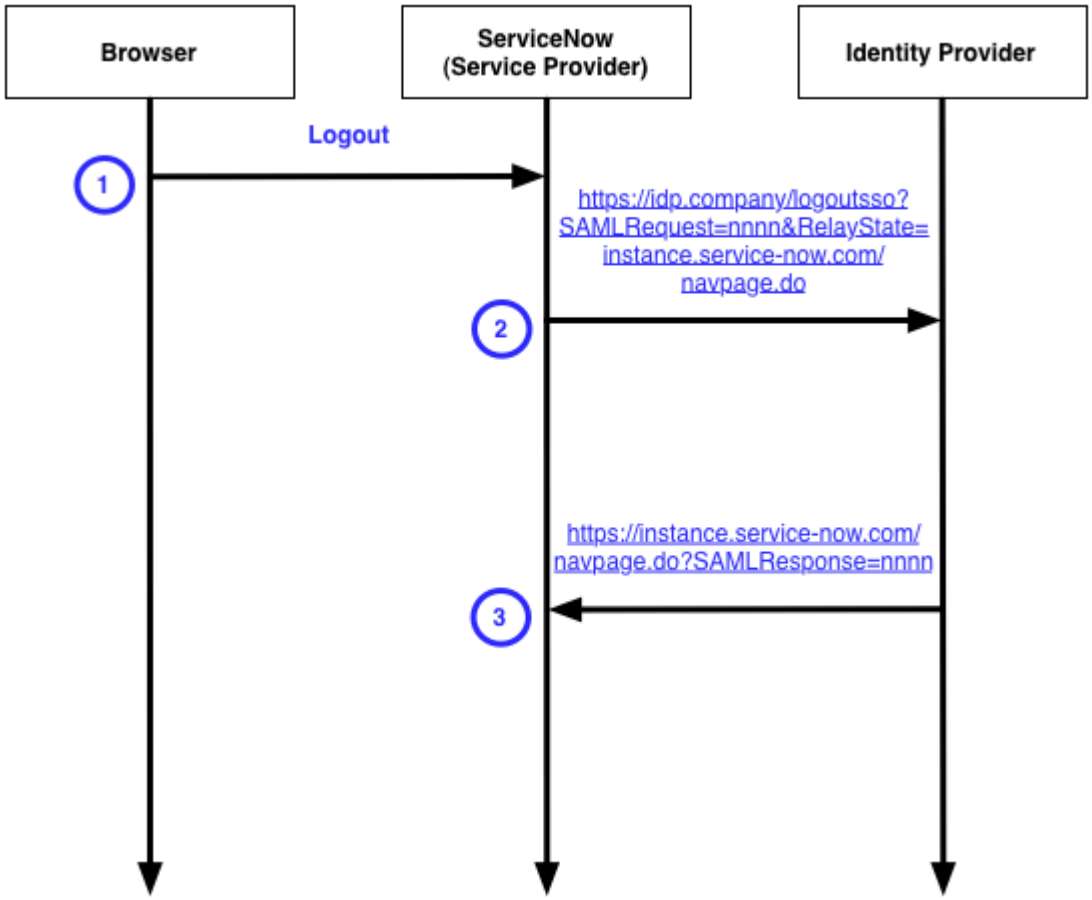

User Clicks the Logout Button

The user clicks the **Logout** button and the instance executes the logout script.

LogoutRequest issued

The logout script constructs a SAML 2.0 `LogoutRequest` and posts it to the preconfigured SingleLogoutRequest SAML 2.0 service at the IdP. The IdP deflates the request and then base64 encodes it. An example `LogoutRequest` looks like this:

```
<saml2p:LogoutRequestxmlns:saml2p="urn:oasis:names:tc:SAML:2.0:p
rotocol"  ID="21B78E9C6C8ECF16F01E4A0F15AB2D46"
  IssueInstant="2010-04-28T21:36:11.230Z"
 Version="2.0"><saml2:Issuer
 xmlns:saml2="urn:oasis:names:tc:SAML:2.0:assertion">https://
dloomac.service-now.com
 </saml2:Issuer><saml2:NameID
 xmlns:saml2="urn:oasis:names:tc:SAML:2.0:assertion"
 Format="urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress"
          NameQualifier="http://idp.ssocircle.com"
 SPNameQualifier="https://
dloomac.service-now.com/navpage.do">david.loo@service-now.com</s
```

```
aml2:NameID><saml2p:SessionIndex>s211b2f811485b2a1d2cc4db2b27193
3c286771104
 </saml2p:SessionIndex></saml2p:LogoutRequest>
```

User Logs Out

The user logs out of the IdP. The IdP redirects back to the instance, which in turns redirects back to the IdP since the user is not logged in.

URL information for an SSO provider

During a login challenge resulting from a URL link into the instance that requires an SSO session, the referring URL might need to be supplied to the SSO provider so that after authentication, the URL can be passed back to the instance and linked to the correct resource.

SSO Target Redirect

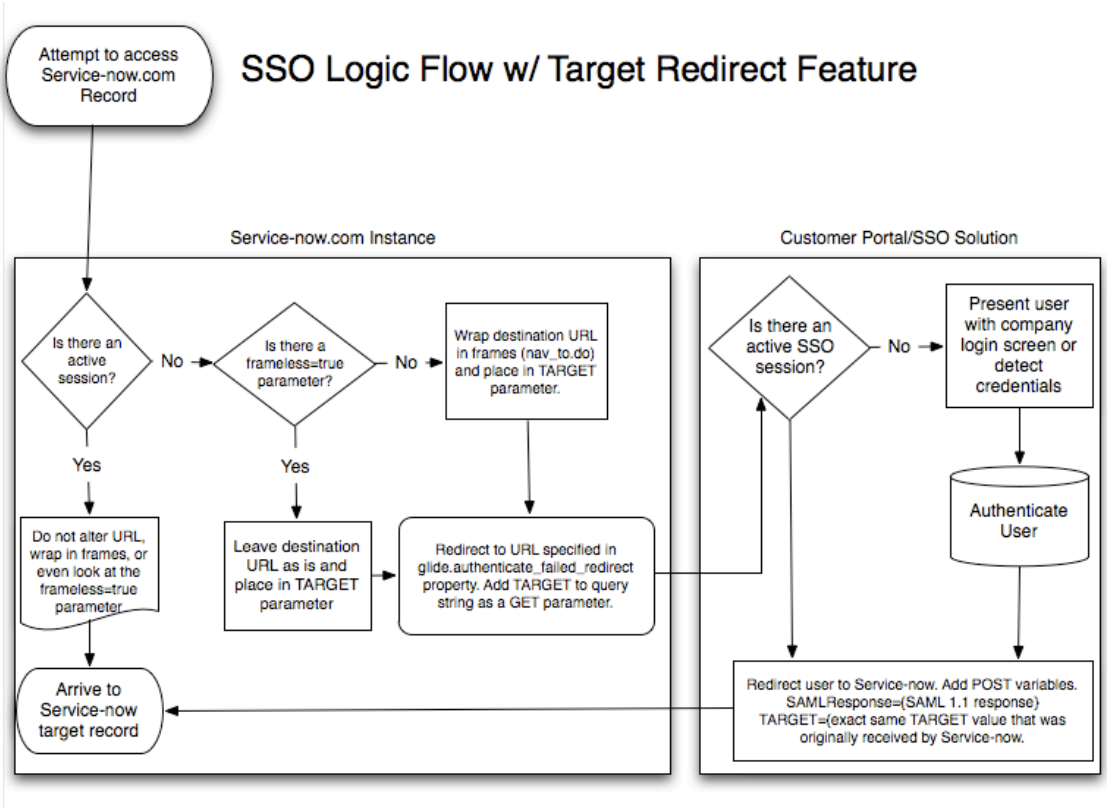

Installation exit return values have been enhanced to pass a URL instead of, or in addition to the URL defined by the properties. Usually, you would return a username or a predefined string value to control authorize or challenge the SSO session. The following examples show the extended behavior of passing a URL.

```
return
 "failed_missing_requirement:%26amp;TARGET=https://
instance.service-now.com/nav_to.do?uri=incident.do?sys_id=1234
5";
```

The example above passes the URL `https://instance.service-now.com/nav_to.do?uri=incident.do?sys_id=12345` to the SSO provider in the form of a URL parameter named TARGET.

Note: It is assumed that the SSO provider will use that information in the TARGET parameter to redirect back to the instance when the user credentials have been collected and authentication passed.

A colon : demarcates the two return values and an encoded & (%26amp;) concatenates the URL defined in the property *glide.authenticate.failed_missing_requirement* and the TARGET parameter.

SAML 2.0 configuration using Multi-Provider SSO

You can create or update a SAML 2.0 SSO configuration from the Multi-Provider SSO feature.

Before you begin

Role required: admin

About this task

Note: New to the Jakarta release, you must validate your configuration by using the Test Connection functionality before you can activate your IdP configuration. You can still use the Update functionality to save your configuration data, but it is not an active configuration without a successful test connection.

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Identity Providers**.
2. Do one of the following options.
   - To update a configuration, click an SSO configuration record.
   - To create a new configuration, click **New** > **SAML**.
3. Enter the IdP information by one of the following methods.

Option	Description
Using a metadata descriptor URL	Click the URL check box and enter the URL of the IdP that you are using.
Using metadata descriptor XML file	Click the XML check box and paste in the XML data generated from the IdP you are us ing.
Entering metadata manually	Close the popup window and manually enter the data in the property fields.

Note: All required fields must be filled-in on the Identity Provider form.

Multi-provider single sign-on fields

Property	Required	Description
Name	Yes	Enter the name for the IdP. This IdP is the auto redirect sys id.
Active	Yes	Active should be set to true for the IdP to be used for authentication. **Note:** The option to set this property only comes after a successful test connection.
Default	No	The Auto Redirect IdP, formerly known as the Primary IdP, automatically redirects users to access the base instance URL. This property sets this IdP configuration as the default.
Auto Redirect IdP	No	Sets this IdP configuration as the Auto Redirect IdP.

Property	Required	Description
		Note: If you make a new Auto Redirect IdP configuration active, the `glide_sso_id` cookie updates with the new Auto Redirect IdP. The `glide.authenticate.sso.update.idp.cookie` system property, automatically enabled, controls this feature.
Identity Provider URL	Yes	Enter the URL to your IdP. Each IdP URL must be unique.
Identity Provider's AuthnRequest	Yes	Enter the URL to the HTTP-Redirect binding obtained from the SingleSignOnService element.
Identity Provider's SingleLogoutRequest	No	Enter the URL obtained from the SingleLogoutService element.
ServiceNow Homepage	Yes	Enter the URL, including login page, of the instance for which the IdP authenticates. For example: https://yourinstance.service-now.com/navpage.do
Entity ID/Issuer	Yes	Enter the base URL, excluding login page. of the instance for which the IdP authenticates. For example: https://yourinstance.service-now.com/
Audience URI	Yes	Enter the base URL, excluding login page. of the instance for which the IdP authenticates. For example: https://yourinstance.service-now.com/
NameID Policy	Yes	Enter the value of the NameIDFormat element the integration uses.
External logout redirect	No	Enter the URL where the integration redirects users after they log out.
Failed Requirement Redirect	No	Enter the URL for redirecting failed authentication requests. By default, this is the URL endpoint of an error page or logout page configured in the IdP. You can populate this value in the glide.authenticate.failed_requirement_redirect field.
Client Type	No	Choose the client type, based on the type of your client. Options:**Iframe Embedded**. **Note:** If client type field is required for your configuration, you can edit the form and add the field. To know more, see Configure client type for OAuth and SSO records.

4. **Optional:** Encryption And Signing tab

Note:

- You should use your own self-signed or CA-signed certificate. The following types of certificates are supported:
  - BCFKS (FIPS-compliant) keystore (Recommended)
  - Java keystore
- The **FIPS approved mode** requires different certificates for **Encryption** and **Signing**.
- Use different certificates for **Encryption** and **Signing** for a better security posture.
- While using the certificates, make sure to update the following system properties with the sys_id of the certificates (x.509 Certificates):
  - **Signing** (`glide.authenticate.sso.saml2.keystore`)
  - **Encryption** (`glide.authenticate.sso.saml2.encryption.keystore`)
- Make sure to update the key alias and key password of the **Signing** and **Encryption** keystores in the Identity Provider record and generate the metadata (Select **Generate Metadata**).
- Upload the signing and encryption certificates present in the generated metadata (XML) to the Identity Provider.
- To configure certificate expiry notification, use **Notify on expiration** and **Groups to notify on expiration**, and set the notification timing using **Warn in days to expire** and **Frequency**.

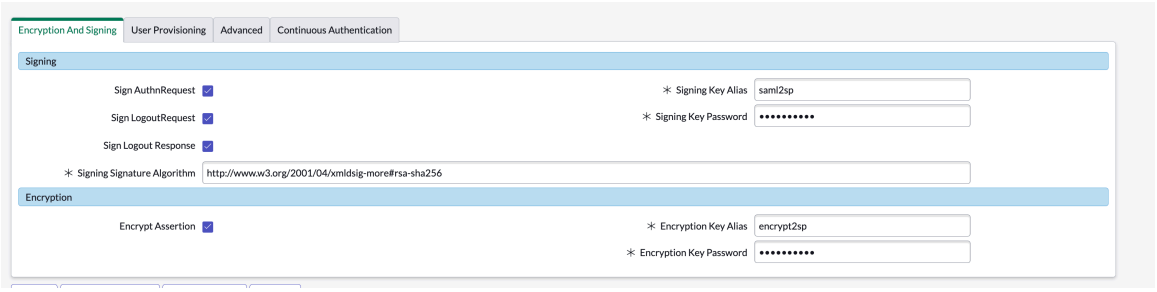

Encryption And Signing fields

Property	Description
Signing Key Alias	Enter the Signing alias of the key entry stored in SAML 2.0 SP Keystore.
Signing Key Password	Enter the Signing password of the key entry stored in SAML 2.0 SP Keystore.
Encryption Key Alias	Enter the Encryption alias of the key entry stored in SAML 2.0 SP Keystore.
Encryption Key Password	Enter the Encryption password of the key entry stored in SAML 2.0 SP Keystore.
Encrypt Assertion	Select the check box to encrypt the assertion in the SAML response. The metadata generated for the IDP embeds the x509 certificate, which the IDP uses to encrypt the assertion in the SAML response that it generates.
Signing Signature Algorithm	Enter the URL that points to the SAML 2.0 Identity Provider AuthnRequest Consumer for eSignature Authentication.

Property	Description
Sign AuthnRequest	Select the check box to enable the IdP single-sign on service to receive a signed AuthnRequest.
Sign LogoutRequest	Select the check box to enable the IdP single-sign on service to receive a signed LogoutRequest.
Sign Logout Response	Select the check box to enable the IdP single-sign on service to receive a signed Logout Response.

5. **Optional:** User Provisioning tab

User Provisioning fields

Property	Description
Auto Provisioning User	Enable automatic user provisioning, creates the users when user doesn't exists in the instance User Table based on the information provided by the IdP.
Update User Record Upon Each Login	Updates user information in the instance User table with the information in the IdP each time the user logs in using SAML.

6. **Optional:** Advanced tab

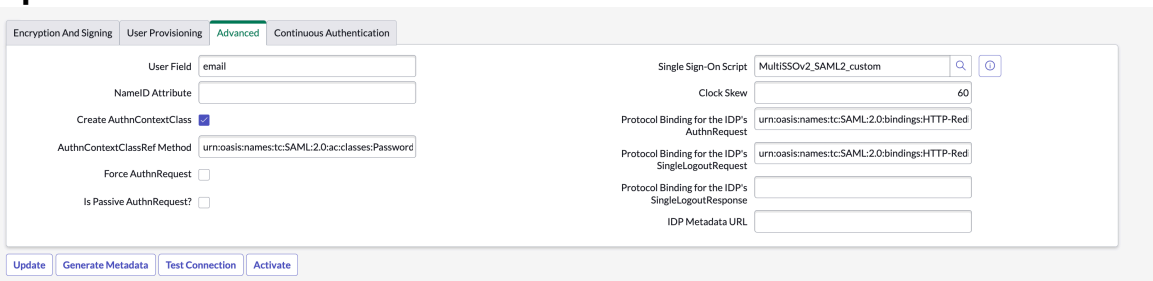

Advanced fields

Property	Description
User Field	Enter the field on the User table that contains the value the IdP requires to identify the user. This is a unique id as part of the response. For example, user name, employee id, and so on. In the sys user table, this unique id is matched with the user details.
NameID Attribute	Leave this field blank unless you configure a new NameID policy. If you configure a new policy, the system requires the User table it must use to identify the user logging in. The system matches the NameID token to the name of that User table field here.
Create AuthnContextClass	Select the check box to specify a particular context class such as Password Protected Transport. If the check box is cleared, the IdP selects the most appropriate context class.
AuthnContextClassRef Method	Enter the URN of the login mechanism you want the IdP to use to authenticate users.
Force AuthnRequest	Select the check box to force AuthnRequests to occur.

Property	Description
Is Passive AuthnRequest	Select the check box if the AuthnRequest is passive.
Single Sign-On Script	Select the Single Sign-On script. The default is *MultiSSOV2_SAML2_custom*.
Sign Logout Response	Enter the logout response details in this field.
Clock Skew	Enter the number of seconds between the two attributes that make up the SAMLResponse nonce. The default is 60. A valid SAMLResponse must fall between the *notBefore* and *notOnOrAfter* date-time values. See Sample SAML 2 Response with the SubjectConfirmation and SubjectConfirmationData Elements and Sample SAML 2 Response with the AudienceRestrictions and Audience Elements for a sample SAMLResponse message.
Protocol Binding for the IDP's SingleLogoutReuqest	Enter one of the supported values listed in the Binding attribute from the SingleLogoutService element.
Metadata URL from which IDP properties are imported	The IdP properties import from this URL. If set, it enables the automatic import of SAML certificate from the IdP if the previous certificate has expired. **Note:** If you upgrade from SAML2 Update 1 to Multi-Provider SSO or if you manually set up your SSO connection, the IdP Metadata URL does not automatically populate.
Request	An unique id as part of request, the id can be user name, employee id, and so on. **Note:** Both redirect and post binding is supported for request. The option to set this field only appears after a successful test connection. For more information see, Test IdP connections.
Response	An unique id as part of response, the id can be user name, employee id, and so on. **Note:** Both redirect and post binding is supported for response. The option to set this field only appears after a successful test connection. For more information see, Test IdP connections.

7. **Optional:** On the Continuous Authentication tab, configure the following fields:

Note:

- The Continuous Authentication tab appears only when you install the **Zero Trust - Continuous Authentication** (`com.snc.zero_trust_continuous_authentication`) plugin that requires license.
- If you're using continuous authentication policy to protect access to table or data class, see Continuous Authentication (CA).

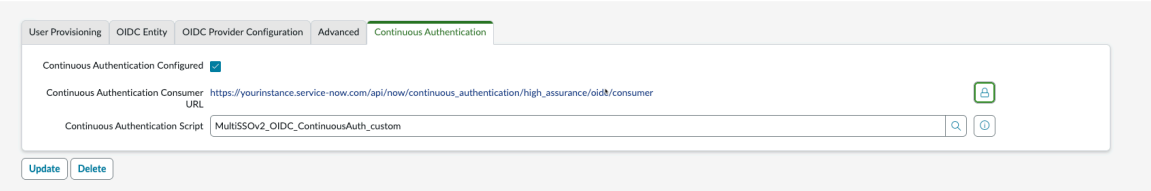

Continuous Authentication

Field	Description
Continuous Authentication Configured	Select the check-box to set the configuration active.
Continuous Authentication Consumer URL	Provide the Consumer URL from the Identity Provider.
Continuous Authentication Script	(Optional) Select the look-up icon to choose the script provided from the platform. In this configuration, for SAML: **MultiSSOv2_SAML2_ContinuousAuth_custom**

X.509 certificates for SAML

Store and activate the necessary IdP certificates for your SAML configuration.

The X.509 certificates are the IdP certificates that a SAML configuration uses. After you install a certificate, you can add as many certificates as necessary. When there are multiple certificates, the system uses the first active certificate that is found. If you set the URL for the **Metadata URL from which IDP properties are imported field**, the system automatically polls the IdP for a current, valid certificate when your certificate is no longer valid. It appends this certificate to your instance, and uses it for your active SAML configuration.

Note: Polling occurs if the IdP is accessible outside of your network.

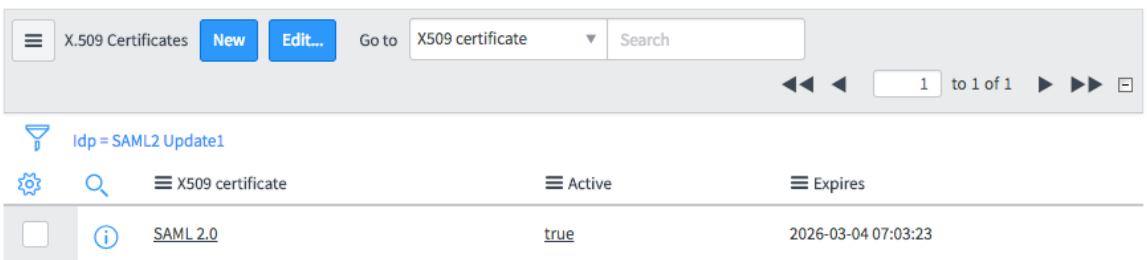

SAML Guided Tour

Use the SAML Guided Tour to configure SAML for single sign-on.

SAML Guided tour helps you to get trained and configure the SSO for your ServiceNow® instance. Admins can select the guided tour to quickly know the required actions when configuring SSO for the instances.

Before using the guided tour, you must configure a SAML application in your preferred Identity Provider, such as Okta, Microsoft Azure, ADFS, and so on. To learn how to configure a SAML application, refer the following documentation:

- Okta: For Okta, you must perform the following:
  - Add an Okta SAML application
  - Configure SAML 2.0 for ServiceNow®
- Microsoft Azure: Single sign-on (SSO) integration with ServiceNow®
- Ping: Configuring SAML SSO with ServiceNow®
- ADFS: Integration with SAML 2.0

To use the SAML Guided tour, activate the Integration - Multiple Provider Single Sign-On Installer plugin. For more information see, Activate Multi-Provider SSO plugin.

To use the SAML Guided Tour, follow these steps:

1. Navigate to **All** > **Multi-Provider SSO** > **Identity Provider**.
2. Click **New**.
3. Select **SAML**.
4. Click the **Help** icon.
5. Click **Take a Tour**.

SAML Guided tour use a series of steps that span across multiple pages to complete the configuration, following the steps and instructions provided as part of the tour and complete the tour. To know more about SAML configuration, see SAML 2.0 configuration using Multi-Provider SSO.

Integrating SAML 2.0 with other features

You can integration your SAML 2.0 solution with other features like E-Signature, deep linking, and ADFS.

Add deep linking support for SAML

Deep linking allows instances to support direct email links to a particular record in the system.

With the SAML 2.0 integration enabled, deep-linking URLs must pass an authentication check before the IdP redirects the user to the originally requested URL. For example, consider an email that contains this URL: `https://<instance name>.service-now.com/nav_to.do?uri=incident.do?sys_id=46c88ac1a9fe1981014de1c831fbcf6d`

The instance sends an authentication request to the IdP and uses the `RelayState` URL parameter to preserve the originally requested resource (in this case, `uri=incident.do?sys_id=46c88ac1a9fe1981014de1c831fbcf6d`). After the IdP authenticates the user, the instance reads the value of the RelayState URL parameter and redirects the user to the requested resource (if it exists in the instance).

To add support for deep linking verify that the identity provider supports the `RelayState` URL parameter.

ADFS integration with SAML 2.0

The ServiceNow Multi-Provider SSO plugin supports a SAML 2. single sign-on (SSO) integration with Microsoft ADFS.

For information about installing and configuring ADFS, see Active Directory Federation Services Overview. The Multi-Provider SSO plugin ha configured and tested with a SAML 2.0 SSO integration with ADFS 2.0, 3.0, Azure AD.ributes are not support SAML Response in

Set up ADFS for SAML

Set up ADFS for SAML. This procedure uses ADFS 2.0 and shows samportal.example.com as the ADFS website. Replace this with your ADFS website address.

Before you begin
Role required: admin

Procedure

1. Log into the ADFS 3.0 server and open the management console.
2. Right-click **Service** and select **Edit Federation Service Properties**.

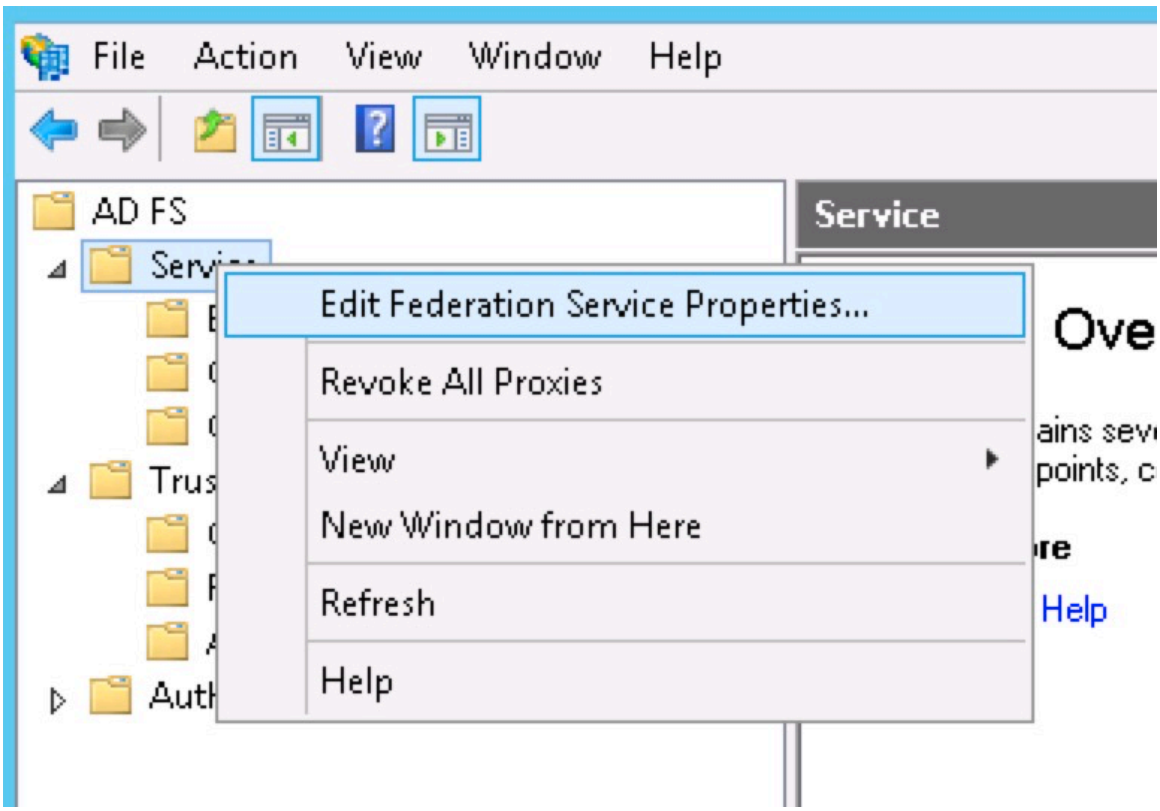

3. Confirm that the General settings match your DNS entries and certificate names.

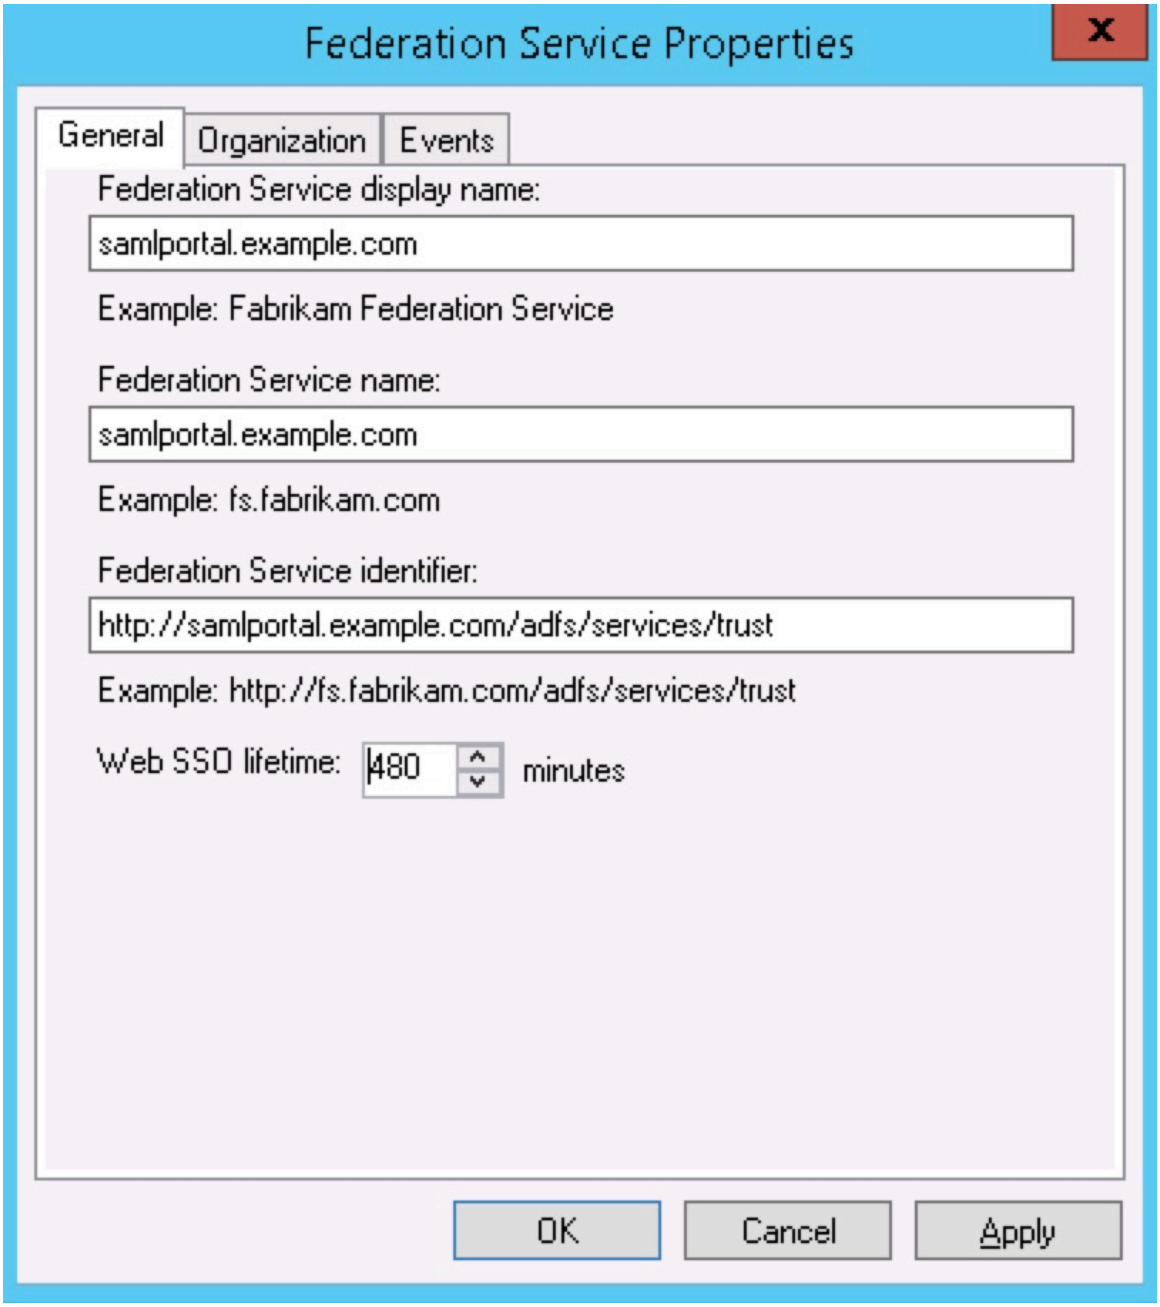

4. Browse to the certificates and export the Token-Signing certificate.

   a. Right-click the certificate and select **View Certificate**.

   b. Select the **Details** tab.

   c. Click **Copy to File**.
   The Certificate Export Wizard opens.

   d. Select **Next**.

   e. Ensure that the **No, do not export the private key** option is selected, and then click **Next**.

   f. Select **DER encoded binary X.509 (.cer)**, and then click **Next**.

g. Select where you want to save the file and give it a name and click **Next**.

h. Select **Finish**.
The instance requires that this certificate be in PEM format. You can convert this certificate using client tools or online tools such as SSL Shopper.

5. Use the DER/Binary certificate that you just created, and export it in Standard PEM format.

Set up the instance for ADFS

Configure your instance and SAML 2.0 settings to work with ADFS.

Before you begin

Perform these steps only after you have set up ADFS for SAML. For details on that process, see set up ADFS for SAML.

Role required: admin

Procedure

1. If not already active, Activate Multi-Provider SSO plugin.
2. Configure SAML, but when you install the IdP certificate, attach the PEM certificate you created when you Set up ADFS for SAML.
3. Click **Save**.
4. Verify that the **Issue** and **Subject** fields have values and that there are no errors.
If an error occurs, open the saved PEM formatted certificate in Notepad and copy and paste the certificate into the PEM Certificate field.
5. Verify that the SAML2SingleSignon_update1 installation exit is active.
6. Continue the SAML 2.0 configuration.

Note: When a certificate is updated on the ADFS server, you also need to upload an updated certificate to the instance.

Configure an ADFS relying party

Take the instance metadata and import it into your ADFS server. However, manual configuration of the relying party appears to be easier to implement.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Identity Providers** > **SAML2 Update1** > **Encryption And Signing** and verify that the SAML property Sign AuthnRequest (*glide.authenticate.sso.saml2.require_signed_authnrequest*) is not active.
Only keep this property active if your ADFS administrator can verify that you require signed requests.
2. Copy the metadata that you generated through the SAML 2 metadata link and save it to a file.
3. Log into the ADFS server and open the management console.
4. Select **Relying Party Trusts**.
5. Select **Add Relying Party Trust** from the top right corner of the window.

The add wizard appears.

6. Click **Start** to begin.
7. Use the **Import File** option to import the metadata file.
8. Give it a display name such as `ServiceNow` and enter any notes you want.
9. Select **ADFS 3.0 Profile**.
10. Do not select a token encryption certificate.
    It will use the certificate that is defined on the service that has already been exported. Defining a certificate prevents proper communication with the instance.
11. Do not enable any settings on the **Configure URL**.
12. Enter the instance site to which you connected as the Relying Party trust identifier.
    In this case, use `https://company.service-now.com` and click **Add**.
13. Permit all users to access this relying party.
14. Click **Next** and clear the **Open the Claims when this finishes** check box.
15. Close this page.
    The new relying party trust appears in the window.
16. Right-click on the relying party trust and select **Properties**.
17. Browse to the **Advanced** tab and set the **Secure hash algorithm** as either SHA-256 or SHA-1.
18. Browse to the Endpoints tab and add a **SAML Assertion Consumer** with a **Post** binding and a URL of `https://company.service-now.com/navpage.do`.

Configure the ADFS relying party claim rules

Edit the claim rules to enable proper communication with the instance.

Before you begin

Role required: admin

Procedure

1. Log into the ADFS server and open the management console.
2. Right-click the relying party trust and select **Edit Claim Rules**.
3. Click the **Issuance Transform Rules** tab.
4. Select **Add Rules**.
5. Select **Send LDAP Attribute as Claims** as the claim rule template to use.
6. Give the claim a name such as `Get LDAP Attributes`.
7. Set the **Attribute store** to `Active Directory`, the **LDAP Attribute** to `E-Mail-Addresses`, and the **Outgoing Claim Type** to `E-mail Address`.

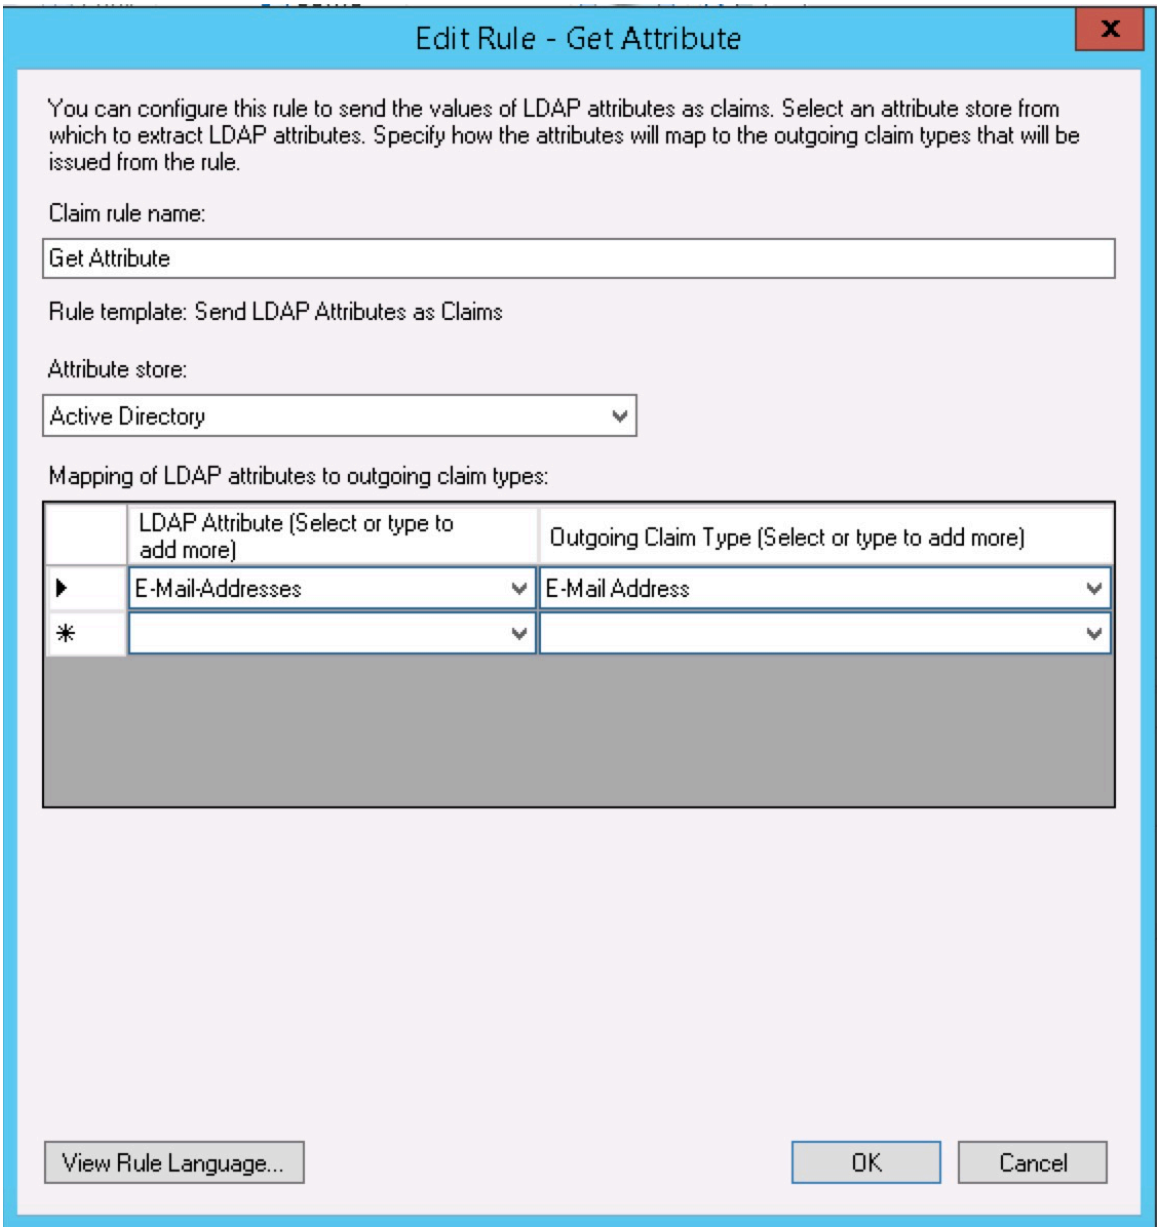


```
c:[Type ==
 "http://schemas.microsoft.com/ws/2008/06/identity/claims/windo
wsaccountname", Issuer == "AD AUTHORITY"]
=> issue(store = "Active Directory",
types =
 ("http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emaila
ddress"),
query = ";mail;{0}", param = c.Value);
```

8. Select **Finish**.
9. Select **Add Rules**.
10. Select **Transform an Incoming Claim** as the claim rule template to use.
11. Give the Claim a name such as `Email to Name ID`.
12. Set the **Incoming claim type** to the **Outgoing Claim Type** in the previous rule.
    For example, `E-Mail Address`.
13. Set the **Outgoing claim type** to `Name ID` and the **Outgoing name ID format** to `Email`.

Note: These values must match the Name ID policy you define during SAML 2.0 configuration.

14. Select **Pass through all claim values**.

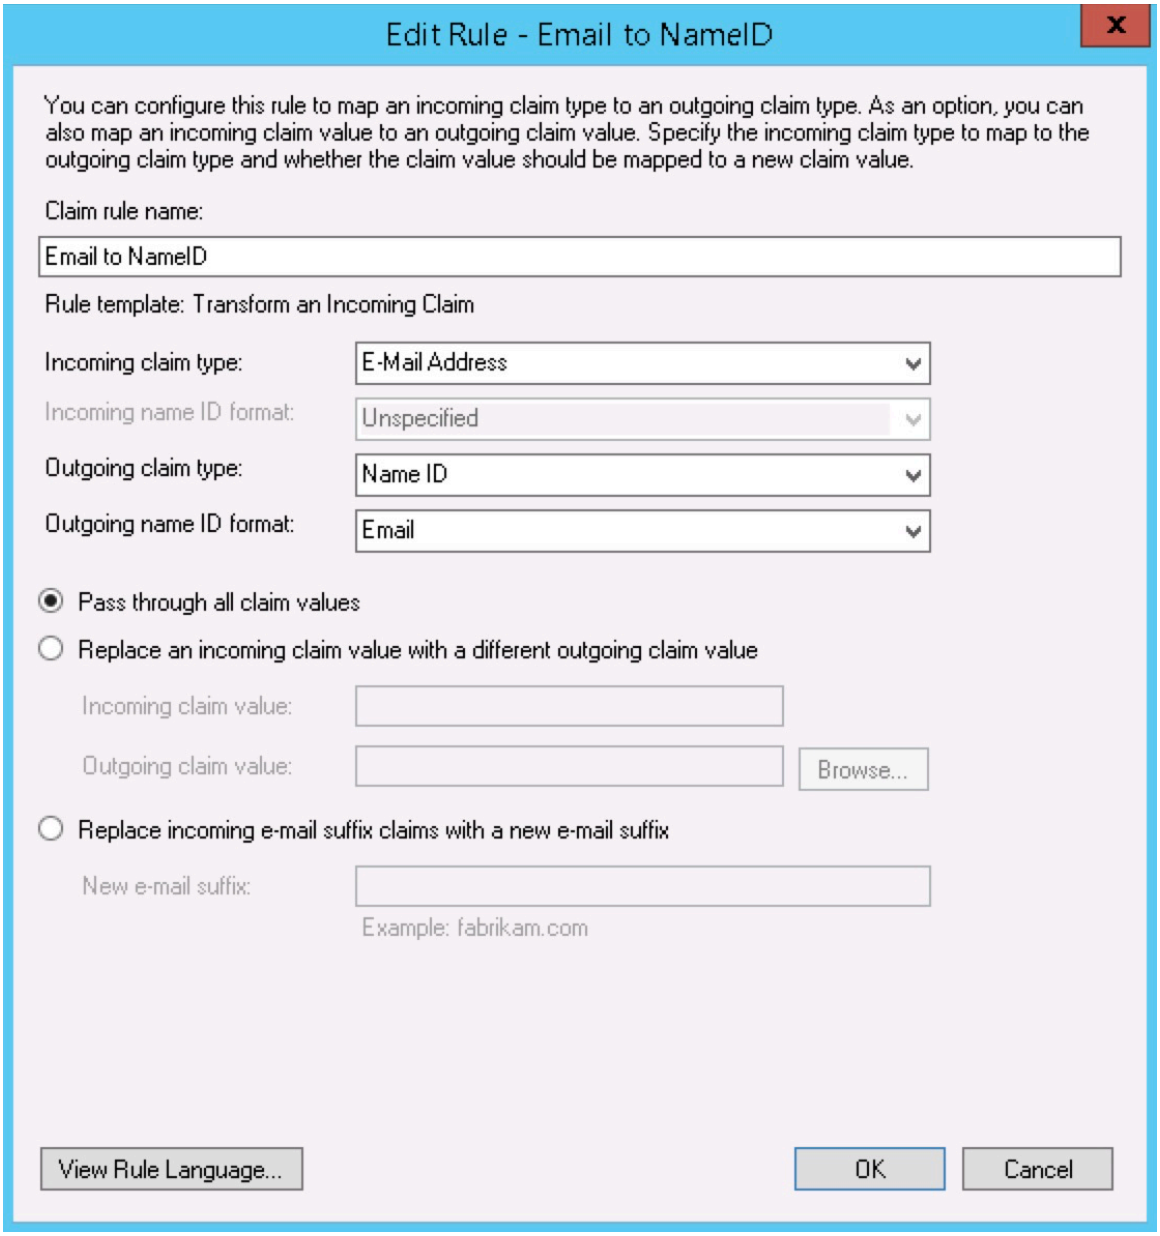

This claim rule should look similar to the following rule language.

```
c:[Type ==
 "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/emailad
dress"]
 => issue(Type =
 "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/nameide
ntifier",
Issuer = c.Issuer, OriginalIssuer = c.OriginalIssuer, Value =
 c.Value, ValueType = c.ValueType,
Properties["http://
schemas.xmlsoap.org/ws/2005/05/identity/claimproperties/forma
t"]
```

```
= "urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress");
```

15. Click **Finish**.

Create a SAML logout endpoint

Create a SAML logout endpoint to allow single logout.

Before you begin
Role required: admin

About this task
See this article on ADFS signout for more information.

Procedure

1. Go to **ADFS manager** > **Trust Relationships** > **Relying Party Trusts** > **properties**.
2. Under the Endpoints tab, click **Add SAML**.
3. Configure the settings:
   - **Endpoint Type**: SAML Logout
   - **Binding**: POST
   - **Trusted URL**: The URL of your ADFS server. For example,

     ```
     https://myadfsserver.domain.net/adfs/ls/?wa=wsignout1.0
     ```

   - **Response URL**: The application logout URL. For example,

     ```
     https://{instancename}.service-now.com/external_logout_complete.do
     ```

Test the ADFS configuration

Test your ADFS configuration to verify that it is properly functioning as an identity provider.

Before you begin
Role required: admin

Procedure

1. Open an Internet Explorer browser.
2. Navigate to your ADFS portal.

   For example, `https://samportal.example.com/adfs/ls/idpinitiatedsignon.aspx`. This page contains a drop-down list of all configured Relying Party Trusts.

3. Select the relying party that is associated with your instance.
4. Click **Continue to Sign In**.

   If you have configured the SAML 2.0 external authentication properly, you should be automatically logged into the instance.

5. Test a direct login URL by navigating to `https://samportal.example.com/adfs/ls/idpinitiatedsignon.aspx?logintoRP=https://company.service-now.com`.

(Workaround) Enable service provider-initiated authentication

Use this workaround if authentication fails because you do not have SAML 2.0 Update 1. This issue can happen if users attempt to skip IdP authentication and navigate directly to the instance.

Before you begin
Role required: admin

About this task
This error occurs when the instance doesn't provide ADFS with the needed definition and semantics for the SPNameQualifier attribute in the SAMLResponse.

Enable service provider-initiated authentication by doing one of the following actions:

Procedure

- Upgrade to SAML 2.0 Update 1 and clear the option to create an AuthnContextClass request.
- Modify the **SAML2** script include to comment out the definitions of the `SPNameQualifier` attribute when you have SAML 2.0 active (not SAML 2.0 Update 1).

  Comment out these lines in the createNameID and createNameIDPolicy functions:

```
//nid.setSPNameQualifier (serviceURL ) ;

 //nameIdPolicy. setSPNameQualifier (serviceURLStr ) ;
```

What to do next
If you do not want the login prompt from your ADFS server to appear when you access the instance, set the following SAML 2.0 Update 1 property to false: **Create an AuthnContextClass request in the AuthnRequest** statement (*`glide.authenticate.sso.saml2.createrequestedauthncontext`*).

(Workaround) Support Kerberos authentication

A workaround is available for the SAML 2.0 integration that changes the authentication context from forms-based authentication to Windows-based authentication.

Before you begin
Role required: admin

About this task
Currently, the SAML 2 integration uses a PasswordProtectedTransport or "forms-based authentication" authentication context. This authentication context requires the IdP to present users with a form for authentication credentials. With Kerberos, a SAML session is already active through an established Windows login, so the user does not need to authenticate with the IdP.

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Identity Providers**.
2. Open the **SAML2 Update1** IdP record.
3. Set the **The AuthnContextClassRef method that we will be included in our SAML 2.0 AuthnRequest to the Identity Provider** to one of the following:

AuthnContextClassRef method values
`urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport` (Default)

```
urn:federation:authentication:windows
```

4. Click **Update**.

Azure AD Integration with SAML 2.0

Integrate ServiceNow with Azure Active Directory (Azure AD).

Integrating ServiceNow with Azure AD enables you to:

- Control in Azure AD who has access to ServiceNow.
- Enable users to automatically sign-in to ServiceNow with their Azure AD accounts.
- Manage your accounts in the Azure portal.

Prerequisities

To get started, you need the perform the following:

- An Azure AD subscription. If you don't have a subscription, you can get a free account.
- A ServiceNow single sign-on (SSO) enabled.
- A ServiceNow instance or tenant of ServiceNow that supports San Diego versions or later.
- A ServiceNow tenant must have the Multiple Provider Single Sign On Plugin enabled.
- For automatic configuration, enable the multi-provider plugin for ServiceNow.

Set of Actions for Configuration

Following are the set of actions that needs to performed to configure Azure AD:

- Add ServiceNow from the gallery to Azure AD.
- Configure Azure AD SSO
- Create an Azure AD test user
- Assign the Azure AD test user
- Configure ServiceNow

Add ServiceNow from the gallery

Add ServiceNow from the gallery to your list of managed SaaS apps on Azure AD.

Before you begin

Role required: Azure admin

Procedure

1. Sign in to the Azure portal by using a Microsoft account.
2. Select the **Azure Active Directory** service from the left pane.
3. Navigate to **Enterprise Applications** > **All Applications**.
4. To add new application, select **New** application.
5. In the Add from the gallery section, enter ServiceNow in the search box.
6. Select ServiceNow from results panel, and then add the app.
7. Wait a few seconds while the app is added to your tenant.

Configure Azure AD SSO

Configure Azure AD SSO in the Azure portal.

Before you begin

Role required: Azure admin

Procedure

1. In the Azure portal, on the ServiceNow application integration page, find the Manage section.
2. Select single sign-on.
   On the Select a single sign-on method page, select SAML.
3. On the Set up single sign-on with SAML page, select the pen icon for Basic SAML Configuration to edit the settings.

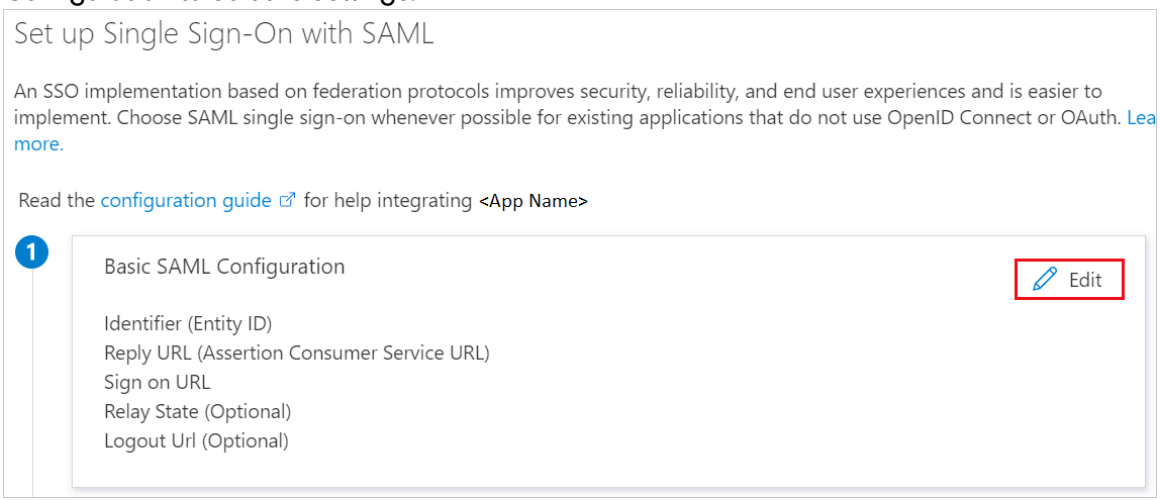

4. In the Basic SAML Configuration section, perform the following:

   a. In **Sign on URL**, enter one of the following URL patterns:

   ```
   https://<instancename>.service-now.com/navpage.do
   https://<instance-name>.service-now.com/login_with_sso.do?glide_sso_id=<sys_id of the sso configuration>
   ```

   Note: You need to provide the sys_id within this URL.

   b. In Identifier (Entity ID), enter a URL with the pattern: `https://<instance-name>.service-now.com`.

   c. For **Reply URL**, enter one of the following URL patterns:

   ```
   https://<instancename>.service-now.com/navpage.do
   https://<instancename>.service-now.com/customer.do
   ```

   d. In Logout URL, enter a URL with the pattern: `https://<instancename>.service-now.com/navpage.do`

   Note: You must update the actual sign-on URL, Reply URL, Logout URL and identifier. the values shown in these URLs are for demo purpose.

5. On the Set up single sign-on with SAML page, in the SAML Signing Certificate section, find **Certificate (Base64)**.

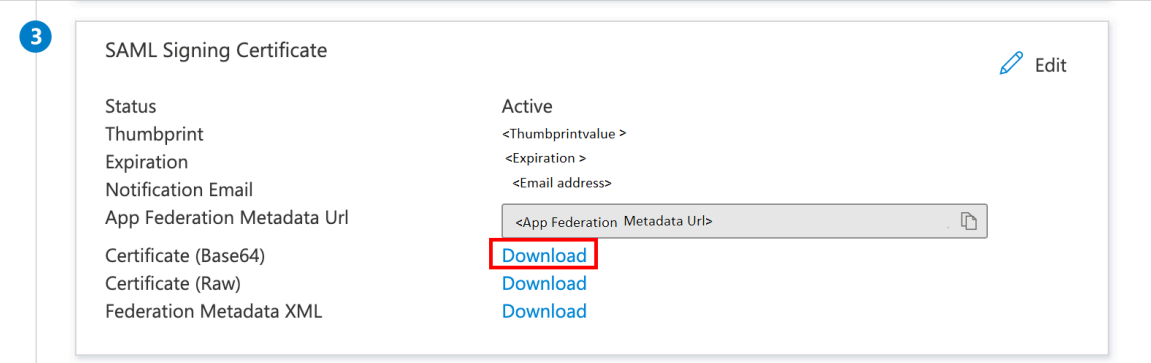

a. Select the copy button to copy **App Federation Metadata Url**, and paste it into Notepad. This URL is required for further configuration.

b. Select **Download** to download Certificate(Base64).

6. In the Set up ServiceNow section, copy the appropriate URLs, based on your requirement.

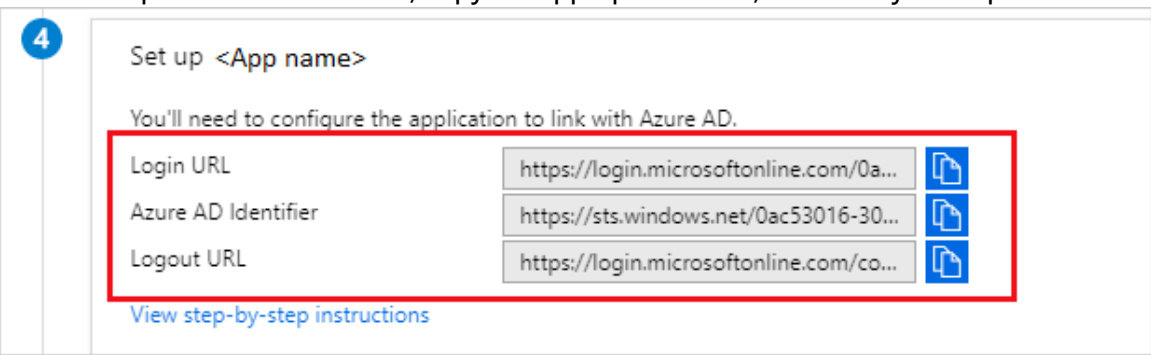

Create an Azure AD test user

Create a test user on Azure AD.

Before you begin
Role required: Azure admin

Procedure

1. On the Azure portal, navigate to **All** > **Azure Active Directory** > **Users** > **All users**.
2. Select **New user**.
3. In the User properties, provide the following information:
   - Name
   - User name
   - Password
4. Click **Create**.

Assign the Azure AD test user

Assign the Azure AD test user that is created to use Azure single sign-on by granting access to ServiceNow.

Before you begin
Role required: Azure admin

Procedure

1. On the Azure portal, navigate to **All** > **Enterprise Applications** > **All applications**.
2. From the applications list, select ServiceNow.
3. On the app's page,in the Manage section, select **Users and groups**.
4. Select **Add user**.
5. In the **Add Assignment** window, select **Users and groups**.
6. In the **Users and groups** window, select the `test user` that was created from the users list.
7. Choose the role for the assigned user from the **Select a role** drop-down if required.
8. In the **Add Assignment** window, select **Assign**.

Configure ServiceNow

Configure ServiceNow with Azure AD details to use SSO.

Before you begin

Plugin: `Integration - Multiple Provider single sign-on Installer`

Enable the Multi Provider SSO properties:

- Select **Enable multiple provider SSO**.
- Select **Enable Auto Importing of users from all identity providers into the user table**.
- Select **Enable debug logging for the multiple provider SSO integration**.
- Enter **email**, The field on the user table that.....

Role required: admin

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Identity Providers**.
2. On the Identity Providers page, select **New**.
3. On the Identity Providers windows, select **SAML**.
4. On the Import Identity Provider Metadata, you can either:
   - **URL**: App Federation Metadata URL to auto-populate the details on the Identity Provider configuration page.
   - **Import**: Import the XML to import the details on the Identity Provider configuration page.
5. Right click on the top of the screen, click **Copy sys_id**, and use this value to the **Sign on URL** in **Basic SAML Configuration** section.
6. On the form, fill in the fields.

Multi-provider single sign-on fields

Property	Required	Description
Name	Yes	Name for the IdP. This IdP is the auto redirect sys id.
Active	Yes	Active should be set to true for the IdP to be used for authentication.

Property	Required	Description
		Note: The option to set this property only comes after a successful test connection.
Default	No	Auto Redirect IdP, formerly known as the Primary IdP, automatically redirects users to access the base instance URL. This property sets this IdP configuration as the default.
Auto Redirect IdP	No	IdP configuration that you can set as the Auto Redirect IdP. **Note:** If you make a new Auto Redirect IdP configuration active, the *glide_sso_id* cookie updates with the new Auto Redirect IdP. The *glide.authenticate.sso.update.idp.cookie* system property, automatically enabled, controls this feature.
Identity Provider URL	Yes	URL to your IdP. Each IdP URL must be unique.
Identity Provider's AuthnRequest	Yes	URL to the HTTP-Redirect binding obtained from the SingleSignOnService element.
Identity Provider's SingleLogoutRequest	No	URL obtained from the SingleLogoutService element.
ServiceNow Homepage	Yes	URL, including the login page, of the instance for which the IdP authenticates. For example: https://yourinstance.service-now.com/navpage.do
Entity ID/Issuer	Yes	Base URL, excluding the login page of the instance for which the IdP authenticates. For example: https://yourinstance.service-now.com/
Audience URI	Yes	Base URL, excluding the login page of the instance for which the IdP authenticates. For example: https://yourinstance.service-now.com/
NameID Policy	Yes	Value of the NameIDFormat element the integration uses.
External logout redirect	No	URL where the integration redirects users after they log out.
Failed Requirement Redirect	No	URL for redirecting failed authentication requests. By default, this is the URL endpoint of an error page or logout page configured in the IdP. You can populate this value in the glide.authenticate.failed_requirement_redirect field.

7. **Optional:** Encryption And Signing tab

Note: Use your own certificates for the encryption and signing.

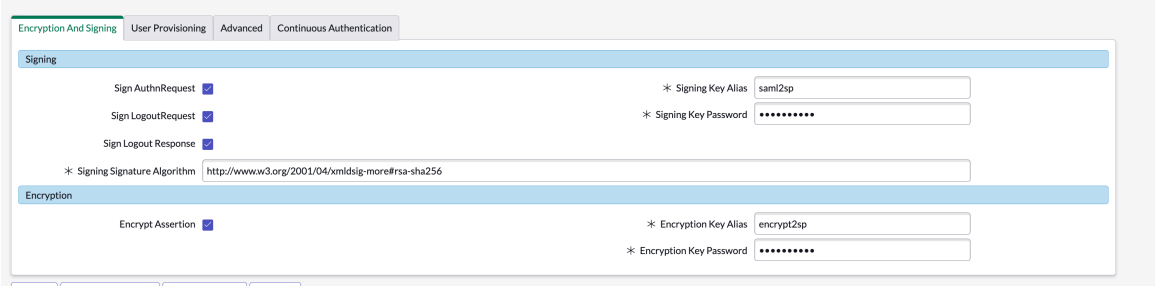

Encryption and Signing fields

Property	Description
Signing/ Encryption Key Alias	Alias of the key entry stored in SAML 2.0 SP Keystore.
Signing Key Password	Password of the key entry stored in SAML 2.0 SP Keystore.
Encrypt Assertion	Check box to encrypt the assertion in the SAML response. The metadata generated for the IDP embeds the x509 certificate, which the IDP uses to encrypt the assertion in the SAML response that it generates.
Signing Signature Algorithm	URL that points to the SAML 2.0 Identity Provider AuthnRequest Consumer for eSignature Authentication.
Sign AuthnRequest	Check box to enable the IdP single-sign on service to receive a signed AuthnRequest.
Sign LogoutRequest	Check box to enable the IdP single-sign on service to receive a signed LogoutRequest.

8. **Optional:** User Provisioning tab

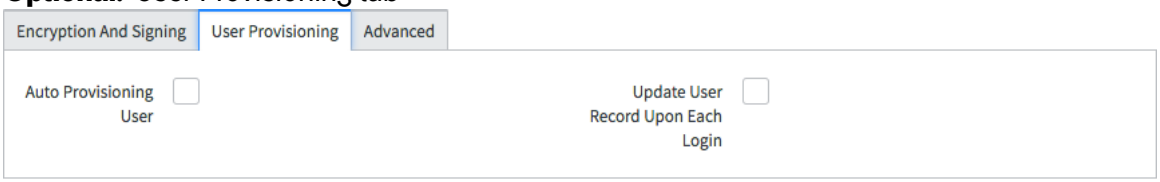

User Provisioning fields

Property	Description
Auto Provisioning User	Automatic user provisioning, creates the users when user doesn't exists in the instance User Table based on the information provided by the IdP.
Update User Record Upon Each Login	Update to user information in the instance User table with the information in the IdP each time that the user logs in using SAML.

9. **Optional:** Advanced tab

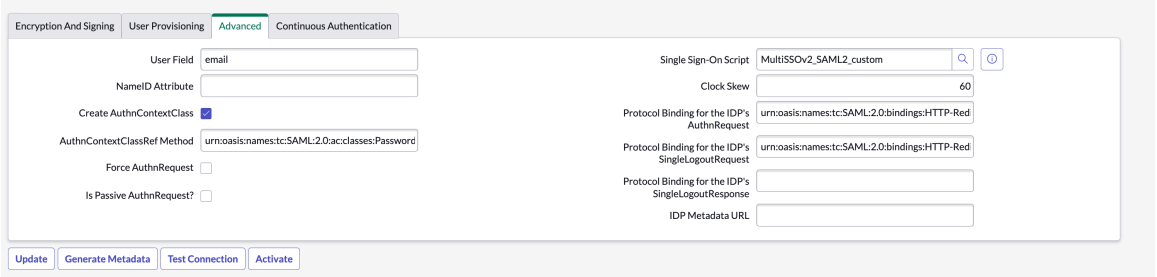

Advanced fields

Property	Description
User Field	Field on the User table that contains the value that the IdP requires to identify the user. This unique id is part of the response. For example, a user name, employee id, and so on. In the sys user table, this unique id is matched with the user details.
NameID Attribute	Field that you leave empty unless you configure a new NameID policy. If you configure a new policy, the system requires the User table it must use to identify the user logging in. The system matches the NameID token to the name of that User table field.
Create AuthnContextClass	Check box to specify a particular context class such as Password Protected Transport. If the check box is cleared, the IdP selects the most appropriate context class.
AuthnContextClassRef Method	URN of the login mechanism that you want the IdP to use to authenticate users.
Force AuthnRequest	Check box to force AuthnRequests to occur.
Is Passive AuthnRequest	Check box if the AuthnRequest is passive.
Single Sign-On Script	Single Sign-On script. The default is *MultiSSOV2_SAML2_custom*.
Sign Logout Response	Logout response details in this field.
Clock Skew	Nnumber of seconds between the two attributes that make up the SAMLResponse nonce. The default is 60. A valid SAMLResponse must fall between the *notBefore* and *notOnOrAfter* date-time values. See Sample SAML 2 Response with the SubjectConfirmation and SubjectConfirmationData Elements and Sample SAML 2 Response with the AudienceRestrictions and Audience Elements for a sample SAMLResponse message.
Protocol Binding for the IDP's SingleLogoutReuqest	One of the supported values listed in the Binding attribute from the SingleLogoutService element.
Metadata URL from which IDP properties are imported	IdP properties import from this URL. If set, it enables the automatic import of SAML certificate from the IdP if the previous certificate has expired. **Note:** If you upgrade from SAML2 Update 1 to Multi-Provider SSO or if you manually set up your SSO connection, the IdP Metadata URL does not automatically populate.
Request	Unique id as part of request. The id can be a user name, employee id, and so on.

Property	Description
	Note: Both redirect and post binding is supported for request. The option to set this field only appears after a successful test connection. For more information, see Test IdP connections.
Response	Unique id as part of response. The id can be a user name, employee id, and so on. **Note:** Both redirect and post binding is supported for response. The option to set this field only appears after a successful test connection. For more information, see Test IdP connections.

10. Select **Test Connection** at the upper-right corner of the page.
11. Enter your credentials.
    The SSO Logout Test Results are displayed.
12. Select **Activate** to activate the configuration.

Email links with external authentication

You can use email links when using the digestive token external authentication, however, you must establish how to handle links in email notifications.

The default links contain a URL that directs you to a specific location in the instance, like an Incident or Change Request, without incorporating SSO credentials. Below are examples for directing the user to the location in the instance without logging in on the instance login page.

- Unencrypted HTTP technique to connect to the /demo instance (it does not navigate to specific record):

```
https://<instance
 name>.service-now.com/?SM_USER=user_name&DE_USER=lQjIVp7aRJtyP
x5+2O/vgU24tbE=
```

- Link (in an email notification) to a specific record, so that the user first goes to the company's own login portal:

```
https://login.company_portal_page.com/nav_to.do?uri=incident.
do?sys_id=009f8eda0a0a0b2b01ab4eb094223466%26sysparm_stack=inci
dent_list.do%3Fsysparm_query=active=true
```

  You must set the *glide.email.override.url* property in your instance to contain the URL of the company portal page. If this property does not exist, you can create it.

- The company portal must then take that URL and construct the redirect URL to the instance preserving the segment necessary to access the specific record, and adding the SSO credentials to the end of the URL:

```
https://<instance
 name>.service-now.com/nav_to.do?uri=incident.do?sys_id=009f8ed
a0a0a0b2b01ab4eb094223466%26sysparm_stack=incident_list.do%3Fsy
sparm_query=active=true&SM_USER=user_name&DE_USER=lQjIVp7aRJtyP
x5+2O/vgU24tbE=
```

Add E-Signature support for SAML

Configure the following properties for E-Signature with Security Assertion Markup Language (SAML) 2.0 update 1.

When E-signature is active with Multi-SSO, SAML properties aren't used. The system adds E-signature properties to the SAML2 Update1 Properties [saml2_update1_properties] table:

Property	Description	Default
Assertion Consumer Index for eSignature authentication	An index number that identifies the endpoint.	1
Assertion Consumer URL for eSignature authentication	The URL that identifies the consumer.	`https://yourinstance.service-now.com/consumer.do`
AuthnRequest URL for eSignature Authentication	The URL for authentication	none

If you're using E-Signature with SAML 1.0 or SAML 2.0 (not including update 1), see the special configuration instructions: E-signature for Multi-Provider SSO.

Note: If you're a Life Science Customer using E-Signature, deactivate the User self-lockout prevention business rule.

Migrating an existing SAML 1.1 integration to SAML 2.0

To migrate from a SAML 1.1 integration to a SAML 2.0 integration, contact customer support.

Update your existing SAML 2.0 integration

Update your existing SAML 2.0 integration.

Before you begin

Role required: admin

About this task

Request the SAML 2.0 Update 1 plugin. Contact Customer Service and Support to request the *SAML 2.0 Single Sign-On - Update 1: security enhancements* plugin. The plugin applies updated versions of the *SAML2SingleSignon* installation exit (login script), *SAML2Logout* installation exit (logout script), and *SAML2* script include (script object).

Merge customizations from existing installation exit scripts into new scripts. The update saves an inactive copy of the integration's original installation exit scripts. You can use these copies to merge any customizations that you made to the login and logout scripts to the new versions of these installation exits.

Merge Customizations from Existing Installation Exit Scripts into New Scripts

Original Installation Exit Script Name	Original Script Status	New Installation Exit Script Name	New Script Status
SAML2SingleSignon	Inactive	SAML2SingleSignon_update1	Active
SAML2	Inactive	SAML2_update1	Active
SAML2Logout	Inactive	SAML2Logout_update1	Active

You can navigate to the SAML 2.0 login and logout installation exit scripts using these paths:

- **SAML 2 Single Sign-on** > **Login script**.
- **SAML 2 Single Sign-on** > **Logout script**.
- **System Definition** > **Installation Exits**.

You can navigate to the SAML 2.0 update 1 script include using these paths:

- **SAML 2 Single Sign-on** > **Script object**.
- **System Definition** > **Script Includes**.

Test the Update.

Procedure

1. Add a system property called `glide.authenticate.sso.saml2.debug` with a value of *true*.
2. Attempt SAML 2.0 login.
3. Review the system log.
   SAML2 validation errors begin with the text *SAML2ValidationError*.
4. Identify and fix typical login errors.
   For more information, see Multi-SSO (SAML 2.0) errors and fixes.

Sample SAML 2 responses after the update

The following sections illustrate the new required elements and attributes that the IdP should provide in the SAML Response.

Sample SAML 2 Response with Issuer Element

The following SAML 2 response uses the *Issuer* element.

```
<samlp:Responsexmlns:samlp="urn:oasis:names:tc:SAML:2.0:protoc
ol"  Destination="https://demoi2.service-now.com/navpage.do"
  ID="s28da6774c88ae1eab292bf25fe625db81919d8e1e"
 InResponseTo="SNC84172Oc227c81948cfd68cadcad235c6"
 IssueInstant="2012-01-30T20:07:10Z"  Version="2.0"><saml:Issuer
 xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion">http://
idp.ssocircle.com</saml:Issuer>
   ...
   <saml:Assertion
 xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion"
     ID="s2f347f973c063836cf70ea38302d94976f9c5b851"
     IssueInstant="2012-01-30T20:07:10Z"
 Version="2.0"><saml:Issuer>http://
idp.ssocircle.com</saml:Issuer>
      ...
   </saml:Assertion></samlp:Response>
```

Sample SAML 2 Response with the SubjectConfirmation and SubjectConfirmationData Elements

The following SAML 2 response uses the *SubjectConfirmation* and *SubjectConfirmationData* elements with the *NotOnOrAfter* and *Recipient* attributes.

```
<saml:SubjectConfirmationMethod="urn:oasis:names:tc:SAML:2.0:cm:
bearer"><saml:SubjectConfirmationData
```

```
 InResponseTo="SNC841720c227c81948cfd68cadcad235c6"
     NotOnOrAfter="2012-01-30T20:17:10Z"
 Recipient="https://
demoi2.service-now.com/navpage.do"/></saml:SubjectConfirmation>
```

Sample SAML 2 Response with the AudienceRestrictions and Audience Elements

The following SAML 2 response uses the *AudienceRestrictions* and *Audience* elements with the *NotBefore* and *NotOnOrAfter* attributes.

```
<saml:ConditionsNotBefore="2012-01-30T19:57:10Z"
 NotOnOrAfter="2012-01-30T20:17:10Z"><saml:AudienceRestriction><
saml:Audience>https://
demoi2.service-now.com</saml:Audience></saml:AudienceRestriction
></saml:Conditions>
```

SAML user provisioning

If users exist in your IdP but are not in your instance, SAML user provisioning can automatically create the users in your instance's User [sys_user] table.

SAML user provisioning is supported for SAML 2.0 Update 1 when Multi-SSO is enabled.

How SAML user provisioning works

When SAML user provisioning is enabled and the system encounters a new user that is not in the instance, the instance automatically creates a record in a temporary table with the name `u_imp_saml_user_<suffix>`, where `<suffix>` is an automatically generated text identifier. The system also creates transform map that specifies the data relationships between the import table and the User table. Each IdP in identified in the system has its own transform map. The transform map is created once for each IdP. Administrators can update it as necessary.

When the user logs in, they access an IdP to log in.

- The system presents a list of all IdPs that are able to use SAML user provisioning. If there is only one IdP that can use SAML user provisioning, that one is used automatically.
- If none of the above conditions are true, the system uses the Auto Redirect IdP.

Administer SAML user provisioning

Update the User table with the users in your IdP by first setting up field mapping and then enabling user provisioning through Multi-SSO IdP settings.

Before you begin

Set up your IdP mapping to identify which fields in the IdP are mapped to the correct fields in the User table.

Role required: admin

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Properties**.
2. Select **Enable Auto Importing of users from all identity providers into the user table** (*`glide.authenticate.multisso.user.autoprovision`*) to activate this feature.
3. Click **Save**.
4. Navigate to **Multi-Provider SSO** > **Identity Providers**.

5. Open the Identity Provider record that you want to use.
6. Create a record in the User table when the user does not already exist by selecting **Auto-provision Users**.
   If you upgraded to this release, you need to configure the form and add this field.
7. Enable user records to be updated when users log in to the IdP and the information on the IdP is out of date with the information on the User table by selecting **Update User Record Upon Each Login**.
   If you upgraded to this release, you need to configure the form and add this field.
8. Click **User Provisioning Transform Map** to see the map that the system automatically creates.
9. Make changes to the map as needed.

Result
When the first unknown users try to log in, the system creates the fields in the import set table from the metadata.xml file.

Note: You cannot map the fields from the IdP table until this first user logs in.

SAML 2.0 troubleshooting

Before contacting support, try the troubleshooting solutions available in the knowledge base on Hi.

Note: The instance does not support solutions provided by external sites.

To know more, see the KB0759251.

Other Common Issues

Error or Symptom	Solution
Error message: "is not a function." This issue might occur in a multi-node environment. If the plugin does not get activated on all nodes, an error like the following appears: org.mozilla.javascript.EcmaError: *[JavaPackage org.opensaml.saml2.core.impl.AuthnRequestBuilder]* is not a function.	This error occurs because the plugin was not active and did not load the .jar file. Therefore, the code appears to be missing. Contact Technical Support to restart nodes that are missing the plugin.
SAML does not authenticate users accessing CMS pages.	By default, CMS pages are public and therefore do not require authentication. If you want SAML to authenticate CMS pages, change the *view_content.do* public page from *active=true* to *active=false*.
Cannot redirect a user back to a CMS page after SAML authentication.	By default, the SSO integration uses a URL parameter called

Other Common Issues (continued)

Error or Symptom	Solution
	URI to control where the user is directed after authentication at the IdP. SSO ignores relative URLs. For example, SSO cannot redirect users to a */ess* relative URL. Instead, the user has to navigate to a URL such as */nav_to.do?uri=/ess*, which uses deep linking syntax. However, this puts the ESS portal inside the main navigation content IFrame. In other words, the site does not take up the full page, but rather loads as a page in your instance. For more information, see CMS Sites and Single Sign-On. If you change the CMS entry page to make it private by setting *view_content.do* to *active=false*, deep linking behavior then requires a customization to the Installation Exit login script. Create a script that looks for the *URI* portion of the URL and constructs a *RelayState* URL parameter containing the relative URL path to redirect users after authenticating at the IdP.
SAML does not redirect users to the appropriate page after authentication.	Determine if the relay state is passed out to the IdP and then passed back during

Other Common Issues (continued)

Error or Symptom	Solution
	authentication. You can do this with a browser capable of saving HTTP request headers and POST info, such as Chrome with its built-in developer tools, or Firefox with the add-on called HTTPfox. For Internet Explorer, use a third-party application such as Fiddler. The goal is to watch the requests pass from the client (browser) to the instance, and from the client to the IdP.

Monitor the event queue for login activities

Every single sign-on integration creates events for login activities.

You can use these events to monitor for login failures and determine if there are any security concerns to address.

Monitoring the event queue for login failures

Event Name	Description	Record	Parameter 1	Parameter 2
external.authentication.succeeded	External authentication succeeded and the user accessed the instance URL.	Session ID	User ID of user who successfully logged in	The URL the user accessed (which may be a deep link)
external.authentication.failed	The single sign-on requirements are not present or are missing.		Session ID	The missing authentication requirements
external.authentication.failed	The user does not exist in the User [sys_user] table		User ID	The string, "User does not exist"
external.authentication.failed	The user is locked out.		User ID	The string, "User locked out."

Event queue login events

The SAML 2.0 integration creates events for login activities.

You can use these events to monitor for login failures and determine if there are any security concerns to address.

Login activities events

Event name	Description	ID used with event	Event string
saml2.logout.validation.failed	The logout response from the IdP failed validation against your logout request. The event validates the <inResponseTo> element against the session ID (ID attribute of the <saml2p:LogoutRequest> element). For example, see the workflow for logout request issued.	Session ID	The string, "SAML2 LogoutResponse validation failed.'
external.authentication.succeeded	External authentication successed.		The string, 'Authentication Successed'
external.authentication.succeeded	External authentication succeeded and the user accessed the instance URL.	Session ID & User ID of user who successfully logged in	The URL the user accessed (which may be a deep link)
external.authentication.failed	The single sign-on requirements are not present or are missing.	Session ID	The missing authentication requirements
external.authentication.failed	The user does not exist in the User [sys_user] table.	User ID	The string, 'User does not exist'
external.authentication.failed	The user is locked out.	User ID	The string, 'User locked out'

OAuth Inbound and Outbound authentication

OAuth based authentication validates the identity of the client that attempts to establish a trust on the system by using an authentication protocol.

OAuth 2.0 - Open Authorization is the industry-standard protocol for authorization, that ocuses on client developer simplicity while providing specific authorization flows for web applications, desktop applications, and mobile devices.

It is a standard that is designed to allow a website or application to access resources hosted by other web apps on behalf of a user.

Instead of using the resource user's credentials to access protected resources, the client obtains an access token. Access tokens are issued to third-party clients with the user's approval, the client then uses the access token to access the protected resources.

From Zurich, you can configure OAuth integration with the following enhancements:

- Increase client secret length to 2048 characters to meet security requirements of third-party systems like Azure DevOps (ADO).
- Provide a JSON Web Key Set (JWKS) URL to automatically manage and update the public key for JSON Web Tokens (JWT) signature validation.
- Request OAuth tokens using the JWT grant type signed with Enhanced Security (ES) algorithms.
- Configure a unique identifier for JWT tokens.

Inbound

Create an endpoint for external clients that want to access your instance. This creates an OAuth client application record and generates a client ID and client secret that the client needs to access the restricted resources on the instance. For more information see, OAuth inbound.

Outbound

Use a third-party OAuth provider that provides the authorization for access to your instance. Specify an OAuth profile and OAuth scope when you are connecting to another OAuth provider. For more information see, OAuth outbound.

OAuth 2.0

OAuth 2.0 lets users access instance resources through external clients by obtaining a token rather than by entering login credentials with each resource request.

You must have the security_admin role to manage the OAuth integration. Configure OAuth 2.0 for the following scenarios:

- **OAuth external client scenario** (Inbound): Your instance provides an endpoint for third-party clients to pull data from the instance.
- **OAuth provider scenario** (Outbound): Your instance pulls data from a third-party provider.

Note: You must user authenticate for the first time to fetch the token post which, you don't need to authenticate using a user account before the token expiry.

Both, simple security and high security frameworks support OAuth 2.0. High Security is recommended. See for information about which versions have high security already active and how to activate high security.

Key concepts of the OAuth 2.0 implementation

Concept	Description
Resource Owner	An entity capable of granting access to a protected resource. A resource owner who is a person is called an end user. The resource owner is always a user account.
Client	An application that, with the authorization of the resource owner, makes requests for protected resources on behalf of the resource owner.
Resource Server	The server that hosts the protected resources, capable of accepting and responding to requests for protected resources.

Concept	Description
Authorization Server	The server that issues access tokens to the client after successfully authenticating the resource owner and obtaining authorization.
Authorization Request	The permission that a client requires to access a protected resource. The authorization request is always an HTTP POST message that contains the ID of the client that is acting on behalf of the resource owner and credentials that authorize the request.
Authorization Grant	A credential that represents the authorization from the resource owner to access a resource. The authorization grant is either user login credentials or a refresh token.
Access Token	A secure string that a client uses to access protected resources. An instance issues access tokens to clients that have a valid authorization grant. Each access token has a specific scope, lifespan, and other attributes. By default, an instance issues access tokens with a 30-minute lifespan in the scenario where the instance is the OAuth provider. For third-party tokens, 30 days.
Refresh Token	A credential that a client uses to obtain new access tokens without requiring additional user authorization. An instance issues a refresh token to a client when it is first authorized to have an access token. By default, an instance issues refresh tokens with a 100-day lifespan in the scenario where the instance is the OAuth provider. For third-party tokens, 365 days.
Self-signed certificates	The ServiceNow AI Platform does not support self-signed certificates. The OAuth client does not utilize the certificate trust store or allow connection to OAuth endpoints that incorporate a self-signed certificate.
User agent	The user who delegates access rights to a client application, which is often a website. The access rights permit the client application or website to access data in the instance that the user has access rights to. The user agent is used in the scenario.

OAuth grant types

A grant type is the way that the client obtains the access token. The following grant types are supported:

- **Authorization code**: The consumer first gets an authorization code and then uses it to get an access token. You can Specify an OAuth profile and specify this grant type. The process that uses the authorization code is also referred to as auth code flow or authorization code flow.
- **Resource owner password credentials**: The consumer of the resource already has the user credentials to get the access token. This process is also referred to as password flow.
- **Client credentials**: The consumer of the resource uses the client ID and client secret that is already configured in the application registry.

Storage of authentication credentials

The OAuth client secret is stored as a *password2* type field, which is encrypted with KMF. User passwords, which are used to check incoming endpoint requests, are stored as a hash value in the User table in a *password* type field (SHA 256). For details on this encryption, see Password2 encryption with KMF

Set up OAuth

Set up and activate OAuth, enable the OAuth system property, create an OAuth application endpoint for external client applications to access the instance, and set OAuth parameters.

Before you begin

Role required: admin

Procedure

1. Make sure the OAuth plugin is active and the OAuth activation property is set to true.
2. Create an OAuth application registry using one of the following methods:
   - Create an endpoint for external clients that want to access your instance. This creates an **OAuth client application** record and generates a client ID and client secret that the client needs to access the restricted resources on the instance.
   - Use a third-party OAuth provider that provides the authorization for access to your instance.

     Specify an OAuth profile and Specify an OAuth scope when you are connecting to another OAuth provider.
3. Configure your client applications to create an HTTP POST that requests an OAuth token. The application must also be able to parse the JSON response to use the returned access token and refresh token.

Activate OAuth

By default, the **OAuth 2.0 (com.snc.platform.security.oauth)** plugin is active on new and upgraded instances. If the plugin is not active on your instance, you can activate it.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the plugin using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Set the OAuth property

To generate OAuth 2.0 tokens to registered applications, the `com.snc.platform.security.oauth.is.active` property must be active for the instance.

Before you begin

Role required: admin

Procedure

1. To use OAuth 2.0, enter **sys_properties.list** in the navigator and select **New**.
   You can also open the system properties list by navigating to **All** > **System Properties** > **All Properties** > .
2. Fill out the form with the following settings:
   - **Name**: com.snc.platform.security.oauth.is.active
   - **Type**: true | false
   - **Value**: true

Change OAuth password parameter

Use this property to ensure only POST body parameters are accepted as input for all supported grant types.

Sending sensitive information over URI query parameters might lead to sensitive information disclosure by clients, the server, or any host between the requests. Starting with the Madrid release, this new property ensures only the POST body parameters are accepted as input for all supported grant types. Supported grant types include:

- authorization code
- password
- client credential
- refresh token

OAuth password parameter property

Property	Description
`glide.oauth.allow.parameters.in.post.body.only`	This property is set to true for zBoots only, as part of the OAuth 2.0 plugin. If you need this setting for your instance, create and set the property to true.

OAuth inbound

OAuth Inbound authentication allows trusted external applications to securely access ServiceNow APIs, ensuring controlled and authorized connections.

You must have one of the following roles to configure or manage OAuth integrations in the system:

- `oauth_admin`
- `mi_admin`
- `admin`

Inbound authentication enables external applications such as third-party systems or other ServiceNow® instances to connect securely to ServiceNow APIs. Inbound authentication confirms that only trusted clients can access your ServiceNow instance in a controlled and

secure manner. ServiceNow supports several OAuth 2.0 grant types, each designed for specific integration scenarios. Use the following information to choose the grant type that best fits your use case:

Authorization Code Grant

Ideal usage scenario	Functionality
Applications that must access user data on behalf of user with the user's consent. Example: Web, mobile, or desktop applications acting on behalf of a user.	The user initiates the login process from the client application, which redirects them to a ServiceNow login page. After the user logs in and grants consent, the client application receives an authorization code. The client application exchanges the authorization code with the ServiceNow instance for an access token. Authorization code grant is the most secure and widely used workflow for user-facing integrations. It supports both confidential clients (with a client secret) and public clients using Proof Key for Code Exchange (PKCE).

For more information about Authorization code grant workflow and configuration, see

Client Credentials Grant

Ideal Usage Scenario	Functionality
Client applications such as back-end services or automated system integrations that must access ServiceNow APIs without user involvement.	The client application authenticates directly with the ServiceNow instance using its own credentials (client ID and secret). Once authenticated, the application receives an access token to access the ServiceNow APIs.

For more information about Client credentials grant workflow and configuration, see .

Third party ID Token Flow

Ideal Usage Scenario	Functionality
Federated authentication scenarios where ServiceNow trusts identity or access tokens issued by external identity providers such as Azure AD or Okta.	The client application obtains an ID or access token from a trusted third-party identity provider, and includes it in the authorization header when making API requests to the ServiceNow instance. ServiceNow validates the token and, if trusted, grants access based on the identity it asserts. This enables seamless single sign-on (SSO) and federated authentication across systems.

For more information about Third party token flow and configuration, see.

JWT Bearer Grant

Ideal Usage Scenario	Functionality
Client applications that need secure access to ServiceNow	**When acting on behalf of a user:**

Ideal Usage Scenario	Functionality
resources, either on behalf of a user or as themselves, without requiring user interaction or storing a shared secret. The client application creates a signed JSON Web Token (JWT) that includes identity-related claims, such as the user or system it represents. It then presents it to the ServiceNow instance to request for access token.	The token represents a previously authenticated user. This enables secure, seamless access without prompting the user for credentials or consent again. The signed token asserts the user's identity, enabling ServiceNow to trust the request without requiring real-time user interaction. **When acting as itself:** The token identifies and authenticates the client application directly. Instead of using a shared secret, the client signs the token with a private key, making it a more secure alternative to the client credentials grant.

For more information about JWT bearer grant workflow and configuration, see .

Resource Owner Password Credentials Grant

Ideal Usage Scenario	Functionality
Highly trusted internal client applications in controlled environments where the app collects the user's credentials directly.	The client application collects the user's user name and password, and redirects them to the ServiceNow instance to obtain an access token. The workflow bypasses redirection and consent screens, but exposes user credentials to the client application. ServiceNow recommends that you implement the Resource owner password credentials grant only in legacy or controlled environments.

For more information about Resource owner password credentials grant workflow and configuration, see .

Implicit Grants

Ideal Usage Scenario	Functionality
Legacy Single-page applications (SPAs) or browser-based apps that can't securely store a client secret, and require a lightweight, client-side authentication flow.	The user logs in through a browser. The client application receives the access token directly in the URL after login, bypassing the intermediate authorization code step. This flow was originally designed for public clients that run entirely in the browser, where securely storing secrets isn't possible. While it simplifies implementation, it exposes tokens in the browser, increasing the risk of interception. For stronger security, ServiceNow recommends implementing an authorization code grant with PKCE.

For more information about Implicit grant workflow and configuration, see OAuth implicit grants.

OAuth Scopes

You can scope the OAuth authentication scope support for REST API. OAuth Scope provides access to only the particular REST APIs. For more information, see REST API Auth Scope.

New Inbound integrations experience

The new inbound integration workflow in the ServiceNow Machine Identity Console provides enhanced experience for managing inbound integrations.

Inbound integrations methods

Inbound Integrations allow you to configure and manage external applications to access ServiceNow APIs. To view detailed usage statistics for each OAuth client, see View Inbound API Integration Usage dashboard.

On the instance, navigate to **All** > **Machine Identity Console** > **Inbound integrations** > **New integration**. Choose the application connection type (grant type).

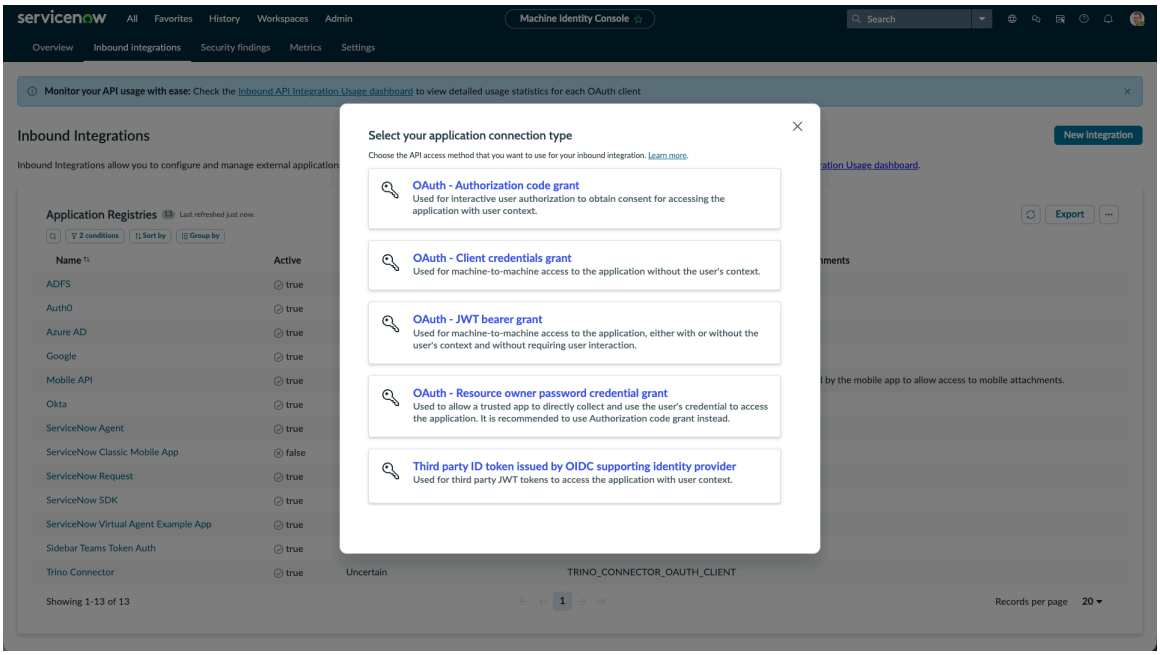

To learn more about the different grant types and how to configure them, refer to the following topics:

- Authorization code grant
- Client credentials grant
- Third party token grant
- JSON Web token bearer grant
- Resource owner password credential grant

Authorization code grant

The OAuth authorization code grant is a secure and widely used flow for web, mobile, or desktop apps that access user data with user consent. It supports both private clients (using a client secret), and public clients (using PKCE).

In the Authorization code grant flow, ServiceNow functions as both the authorization server (handling user authentication and token issuance) and the resource server (hosting the APIs). If SSO is enabled, ServiceNow redirects the user to the configured Identity Provider (IdP) for authentication. After the IdP successfully authenticates the user, control returns to ServiceNow, which then issues the authorization code. This process ensures that even with external authentication, ServiceNow remains the authority for issuing tokens and managing API access.

Note: If you want to use your own identity provider (such as Azure AD or Okta) as the authorization server, consider using the flow.

Related topics

- Authorization code grant workflow
- Configure an OAuth authorization code grant

Authorization code grant workflow

The OAuth authorization code grant is a secure and widely used flow for web, mobile, or desktop apps that access user data with user consent. It supports both private clients (using a client secret), and public clients (using PKCE).

Before you begin

Role required: `oauth_admin, mi_admin, admin`

About this task

This topic collection provides information on how a client application can use the Authorization code grant flow to obtain a token from ServiceNow and make API calls with that token. Private clients use client secret, while public clients use PKCE code challenge.

Authorization workflow

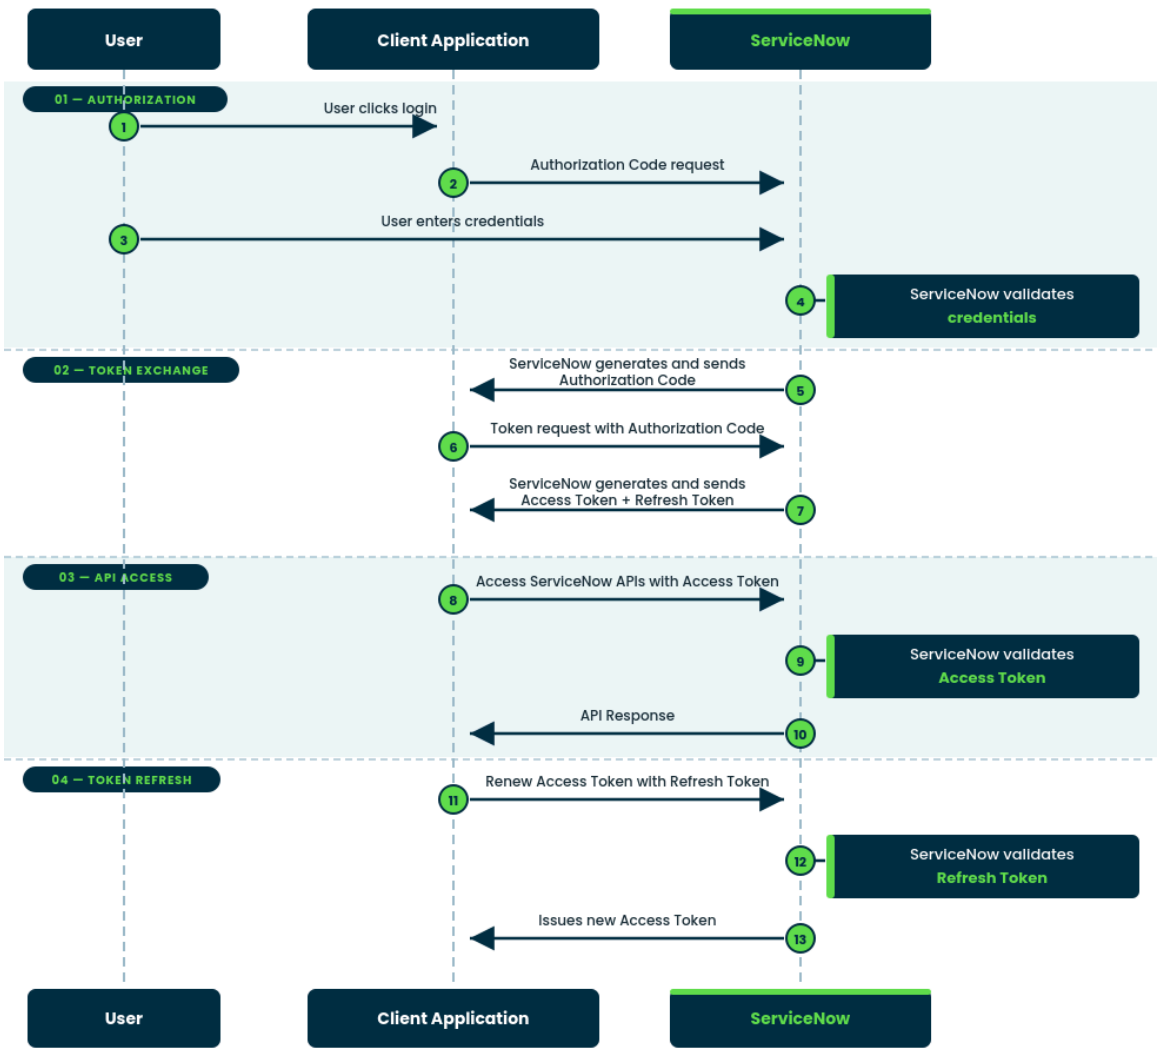

Procedure

1. Log in from the client application.

   The user begins the login process from the client application interface.

2. Initiate the authorization request.

   The client redirects the user to the ServiceNow authorization endpoint. The method of initiating the authorization request depends on the type of client: Public or Private.

   #### Public Clients

   Public clients (Example: Mobile or Single-page applications) cannot securely store a client secret. Therefore, they must use **Proof Key for Code Exchange (PKCE)** to enhance security.

   - In the authorization request, include the **PKCE code challenge** and specify the **code challenge method**.
   - During the token request, the client must send the **code verifier** to validate the authorization code.

   Perform a GET request to the authorization endpoint with the following parameters:

   ```
   Method: GET
   Endpoint: https://<servicenow_base_url>/oauth_auth.do
   ```

   Authorization Request Parameters (Public Client - PKCE)

Parameter	Required	Description
`response_type`	Yes	Set the value to `code` to initiate the authorization code flow.
`client_id`	Yes	The unique identifier for your client application. Format: YOUR_CLIENT_ID
`redirect_uri`	Yes	The URI to which ServiceNow sends the authorization code. Example: https://yourapp.com/callback
`code_challenge`	Yes	A base64url-encoded SHA-256 hash of the code verifier. This is used as part of the PKCE flow.
`code_challenge_method`	Yes	Specifies the transformation method used for the code challenge. Set to `S256`.

Parameter	Required	Description
`scope`	Optional	A space-delimited list of requested scopes. Example: `incident_read incident_write`.
`state`	Yes	A client-generated value used to avoid CSRF attacks. The value is returned unchanged in the redirect URI, enabling the client to validate it.

Note: Starting with the Madrid release, the system property `glide.oauth.state.parameter.required` mandates the use of the `state` parameter in the OAuth requests. The `state` property is set to `true` by default in the new instances, and `optional` in upgraded instances. In case of missing `state` parameter, the authorization request fails and the following error is displayed: `Missing State parameter in request`.

Private Clients

Private clients (Example: Server-side applications) can securely store a client secret and do not require PKCE.

- The authorization request is initiated by redirecting the user to the authorization endpoint. **No client secret or PKCE code challenge is required** in this step.
- During the token request, the client includes the **client secret** along with the **authorization code** to obtain the access token.

Perform a GET request to the authorization endpoint with the following parameters:

```
Method: GET
Endpoint: https://<servicenow_base_url>/oauth_auth.do
```

Authorization Request Parameters (Private Client-Client Secret)

Parameter	Required	Description
`response_type`	Yes	Set the value to `code` to initiate the authorization code flow.
`client_id`	Yes	The unique identifier for your client application. Format: YOUR_CLIENT_ID
`redirect_uri`	Yes	The URI to which ServiceNow sends the authorization code.

Parameter	Required	Description
		Example: https://yourapp.com/callback
`scope`	Optional	A space-delimited list of requested scopes. Example: `incident_read incident_write`.
`state`	Yes	A client-generated value used to avoid Cross-Site Request Forgery (CSRF) attacks. The value is returned unchanged in the redirect URI, enabling the client to validate it.

3. Log in and grant access consent to the client application.
   Log in to ServiceNow (or IdP, if SSO is enabled), and grant access consent to the client application.
4. ServiceNow (or IdP, if SSO is enabled) validates the credentials and ServiceNow returns an authorization code to the client.

   After the successful authentication, the browser is redirected to the `redirect_uri`, and the authorization code is included in the query string:

   ```
   https://yourapp.com/callback?code=AUTH_CODE&state=xyz123
   ```

5. Initiate the authorization request.

   The client redirects the user to the ServiceNow authorization endpoint for an access token. The method of initiating the authorization request depends on the type of client: Public or Private.

   Public Clients

   Public clients (Example: Mobile or Single-page applications) cannot securely store a client secret. Therefore, they must use **Proof Key for Code Exchange (PKCE)** to enhance security.

   - In the authorization request, include the **PKCE code challenge** and specify the **code challenge method**.
   - During the token request, the client must send the **code verifier** to validate the authorization code.

   The client sends a POST request to token endpoint with the following parameters:

   ```
   Method: POST
   Endpoint: https://<servicenow_base_url>/oauth_token.do
   Headers: Content-Type:
    application/x-www-form-urlencoded
   ```

Token Request Parameters (Public Client-PKCE)

Parameter	Required	Description
`grant_type`	Yes	Set the value to `authorization_code` to exchange the code for a token.
`code`	Yes	The authorization code received from the authorization endpoint.
`redirect_uri`	Yes	The URI used in the initial authorization request. Example: https://yourapp.com/callback
`client_id`	Yes	The unique identifier for your client application.
`code_verifier`	Yes	The original string used to generate the PKCE `code_challenge`.
`state`	Yes	A client-generated value used to help prevent CSRF attacks. The value is returned unchanged in the redirect URI, enabling the client to validate it.

Private Clients

Private clients (Example: Server-side applications) can securely store a client secret and do not require PKCE.

- The authorization request is initiated by redirecting the user to the authorization endpoint. **No client secret or PKCE code challenge is required** in this step.
- During the token request, the client includes the **client secret** along with the **authorization code** to obtain the access token.

Perform a POST request to the authorization endpoint with the following parameters:

```
Method: POST
Endpoint: https://<servicenow_base_url>/oauth_token.do
Headers: Content-Type:
 application/x-www-form-urlencoded
```

Token Request Parameters (Private Client-Client Secret)

Parameter	Required	Description
`grant_type`	Yes	Set the value to `authorization_code`

Parameter	Required	Description
		to exchange the code for a token.
`code`	Yes	The authorization code received from the authorization endpoint.
`redirect_uri`	Yes	The URI used in the initial authorization request. Example: https://yourapp.com/callback
`client_id`	Yes	The unique identifier for your client application.
`client_secret`	Yes	The client's secret used to authenticate with the token endpoint.
`state`	Yes	A client-generated value used to help prevent CSRF attacks. The value is returned unchanged in the redirect URI, enabling the client to validate it.

6. Access the ServiceNow APIs with the access token.

   Example:

   Make a GET request to the APIs using the access token. Include the access token in the `Authorization` header.

```
Method: GET
End Point: https://<servicenow_base_url>/api/now/incident
Authorization: Bearer YOUR_ACCESS_TOKEN
```

7. Renew the access token if it has expired.
   Make a POST request to refresh the access token (private clients only) with the following parameters:

```
Method: POST
Endpoint: https://<servicenow_base_url>/oauth_token.do
Headers: Content-Type: application/x-www-form-urlencoded
```

Refresh Token Request Parameters (Private Client)

Parameter	Required	Description
`grant_type`	Yes	Set the value to `refresh_token` to request a new access token.
`refresh_token`	Yes	The refresh token previously issued by the token endpoint.

Parameter	Required	Description
`client_id`	Yes	The unique identifier for your client application.
`client_secret`	Yes	The client secret used to authenticate with the token endpoint.

Configure an OAuth authorization code grant

Configure the OAuth authorization code grant to enable secure and interactive user authentication to enable applications to access resources on behalf of users. The OAuth authorization code grant verifies that the API access is granted based on the user identity and permissions.

Before you begin

Role required: `oauth_admin, mi_admin, admin`

Procedure

1. Navigate to **Machine Identity Console** > **Inbound integrations** > **New integration** > > **OAuth authorization code grant**.
   The New Record page appears.
2. Update the text fields in the Details form with the appropriate information.

   Details form

Field	Description
Name of OAuth entity	Name of the OAuth entity.
Provider name	Enter the name of the service provider you want to integrate with. Example: Microsoft, Google, Zoom, SAP, etc
Redirect URL	URL to which the authorization code should be sent after authentication.
Client ID	Unique ID assigned to identify the application.
Client Secret	The secret key that only the application and the authorization server can identify. The application uses this key to authenticate and obtain access tokens.

   Select the **This is a public client check box** if the application can't securely store credentials, and doesn't require a secret key to prove its identity during authorization. The client secret information is processed for public clients.
3. Update the text fields in the **Auth scope** form with the appropriate information. The authentication scope defines the level of access an application has to a resource. Select the authentication scope for the specific REST APIs you want to access.

Auth scope form

Field	Description
Auth scope	The level of access an application has to a resource. The authentication scope restricts the actions that an access token can perform on APIs or data.
Limit authorization	The names of the APIs for which you want to restrict authorization.
Allow access only to APIs in selected scope	Enable the option for the integration to only access APIs that are explicitly listed in the selected scopes.

4. Update the text fields in the Advanced options (optional) form with the appropriate information.

Advanced options form

Field	Description
Enforce token restriction	The Enforce token restriction option limits the client to accessing only the APIs specified in the REST API Access Policies. If you unselect it, the client can access other REST APIs based on the user ACL permissions.
Token Format	Format of token to generate. Options: ○ JWT ○ Opaque **Note:** ○ The jwks url is available in the location: `api/now/oauth/jwks`. ○ The rotated (inactive keys) from jwks response after is removed after 105 days default.
Access token lifespan	Duration (in seconds) for which the OAuth access token remains valid before it expires. **Note:** The default value is 1800 seconds.
Refresh token lifespan	Duration (in seconds) for which the OAuth refresh token remains valid before it expires. **Note:** The default value is 8,640,000 seconds.
Login URL	HTTP redirection endpoint to authenticate with the authorization server.

Field	Description
Logo URL	Web address of an image that represents the application during the authentication and authorization process. It's displayed on the authorization server's consent screen to help you recognize the requesting application.

Enforcing token restriction applies limitations on how an OAuth access token can be used, enhancing security by verifying that tokens are valid only under specific conditions. Enable the **Enforce token restriction** check box to limit OAuth access tokens to specific APIs defined in the API access policy. If the Enforce token restriction is turned off, the token can be used across other REST APIs.

5. Select **Create new auth scope** to add a new auth scope.
6. Select **Save**.
   A new OAuth authorization code grant is created.
7. Go to **All** > **Inbound integrations** > **Application Registries** to view the newly created OAuth authorization code grant.

Client credentials grant

Use the OAuth client credentials grant type for back-end services or automated integrations that access ServiceNow® APIs without user interaction. The client application authenticates directly using its client ID and secret, and receives an access token that represents the application itself, and not the user.

Related topics

Client credentials grant workflow

Configure an OAuth Client credential grant

Client credentials grant workflow

Authenticate a client application using a client credentials workflow. The client credentials grant workflow is used by back-end services or system integrations to access ServiceNow® APIs without user involvement.

Before you begin

Role required: `oauth_admin, mi_admin, admin`

About this task

This workflow describes how a client application (back-end service or system integration) authenticates directly with ServiceNow using its client credentials without user interaction. The application requests an access token using its client ID and client secret, which ServiceNow validates before issuing the token. The client then uses this token to access ServiceNow APIs. ServiceNow validates each request before returning the appropriate response.

Client credentials grant workflow

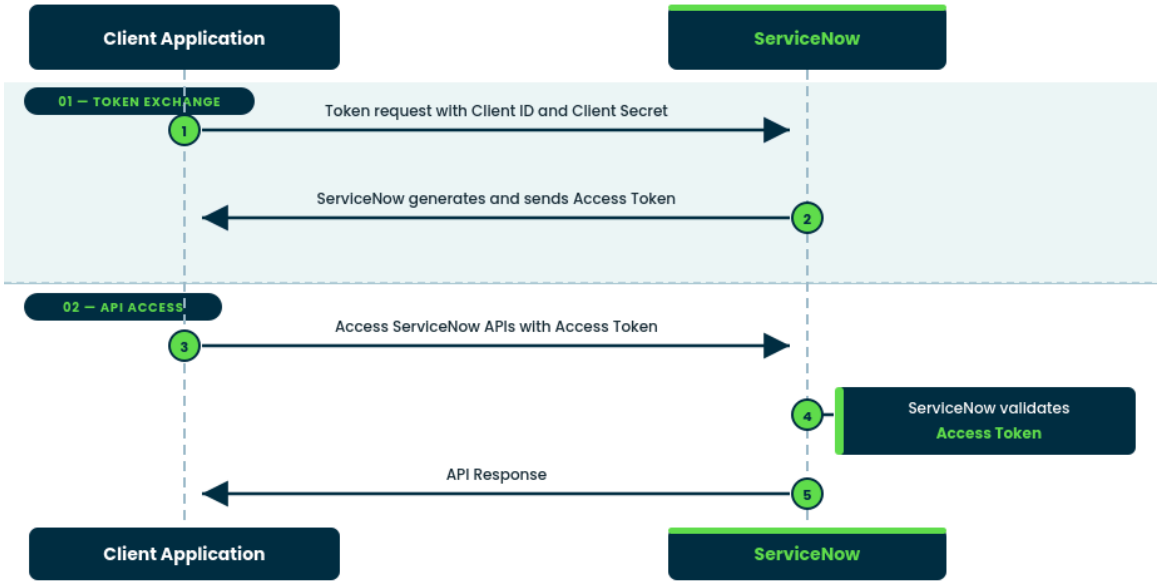

Procedure

1. **The client application makes a token request** to the ServiceNow end point with the following parameters:

```
Method: POST
Endpoint: https://<servicenow_base_url>/oauth_token.do
```

Token Request Parameters

Parameter	Required	Description
`grant_type`	Yes	OAuth 2.0 grant type. Example: `client_credentials`
`client_id`	Yes	Unique identifier for the client application. Example:`YOUR_CLIENT_ID`
`client_secret`	Yes	Secret associated with the client ID. Example:`YOUR_CLIENT_SECRET`
`scope`	Optional	Requested permissions for the access token. Example: `incident_read incident_write`

2. ServiceNow validates the credentials and returns the access token.
3. Make an API request with the access token.

   Include the access token in the `Authorization` header of each API request.

```
Method: POST
Endpoint: https://<servicenow_base_url>/api/now/incident
Authorization: Bearer YOUR_ACCESS_TOKEN
```

4. ServiceNow validates the token and returns the appropriate API response.

> **Note:** Use the client credentials grant workflow only with trusted, server-side applications. Maintain the `client_secret` securely. Ensure that you don't use the `client_secret` in client-side environments such as browsers or mobile apps.

Configure an OAuth Client credential grant

Configure the OAuth Client Credentials Grant for secure machine-to-machine authentication without user interaction. It authenticates applications using client credentials and grants-controlled API access with scoped permissions.

Before you begin

Role required: `oauth_admin, mi_admin, admin`

Procedure

1. Navigate to **Machine Identity Console** > > **Inbound integrations** > > **New integration** > **OAuth Client credential grant**.
   The OAuth Client credential configuration page appears.
2. Update the text fields in the Details form with the appropriate information.

Details form

Field	Description
Name	The name of the OAuth entity.
Provider name	Enter the name of the service provider you want to integrate with. Example: Microsoft, Google, Zoom, SAP, etc
Client ID	The unique ID assigned to identify the application.
Client Secret	The secret key that only the application and the authorization server can identify. The application uses this key to authenticate and obtain access tokens.

   Select the **Active** check box.

3. Update the text fields in the Advanced options (optional) form with the appropriate information.
4. Update the text fields in the **Auth scope (optional)** form with the appropriate information.
   The authentication scope defines the level of access an application has to a resource. Select the authentication scope for the specific REST APIs you want to access.

Auth scope form

Field	Description
Auth scope	The level of access an application has to a resource. The authentication scope restricts

Field	Description
	the actions that an access token can perform on APIs or data.
Limit authorization	The names of the APIs for which you want to restrict authorization.
Allow access only to APIs in selected scope	Enable the option for the integration to only access APIs that are explicitly listed in the selected scopes.

a. Select **Create new auth scope** to add a new auth scope.

5. Update the text fields in the Advanced options (optional) form with the appropriate information.

Advanced options form

Field	Description
Enforce token restriction	The Enforce token restriction option limits the client to accessing only the APIs specified in the REST API Access Policies. If you unselect it, the client can access other REST APIs based on the user ACL permissions.
Token Format	Format of token to generate. Options: ○ JWT ○ Opaque **Note:** ○ The jwks url is available in the location: `api/now/oauth/jwks`. ○ The rotated (inactive keys) from jwks response after is removed after 105 days default.
Access token lifespan	Duration (in seconds) for which the OAuth access token remains valid before it expires. **Note:** The default value is 1800 seconds.

6. Select **Save**.
   A new OAuth Client credential grant is created.
7. Go to **All** > **Inbound integrations** > **Application Registries** to view the newly created client credential grant.

Third party token grant

The third party token grant enables ServiceNow® to accept identity tokens from trusted external identity providers, such as Azure AD or Okta. Third party token grant provides secure, token-based access. This method supports secure access and single sign-on (SSO) in federated authentication scenarios.

The client application requests an ID or access token from a trusted external identity provider, such as Azure AD or Okta, and includes it in the `Authorization` header of API requests to ServiceNow®. ServiceNow® validates the token and, if trusted, grants access based on the asserted identity.

You can use accounts from a third-party identity provider (IdP) to access the ServiceNow® API for:

- Third party token workflow for user accounts
- Third party token workflow for service accounts

Third party token workflow for user accounts

This workflow is based on the token federation concept. It allows client applications to obtain tokens directly from an IdP and use them to access ServiceNow APIs.

Before you begin

Role required: `oauth_admin, mi_admin, admin`

About this task

The third-party client application requests tokens directly from your identity provider (IdP). The authentication method between the client and the IdP is flexible and can be configured to meet your specific requirements. After successful authentication, the IdP issues an ID token or access token, and optionally a refresh token. These tokens are sent directly to the client application, which then uses them to access ServiceNow APIs.

Note: ServiceNow validates the token using the public key configured during setup and grants access to the requested APIs. Ensure that the token is in JSON Web Token (JWT) format.

User Account Workflow

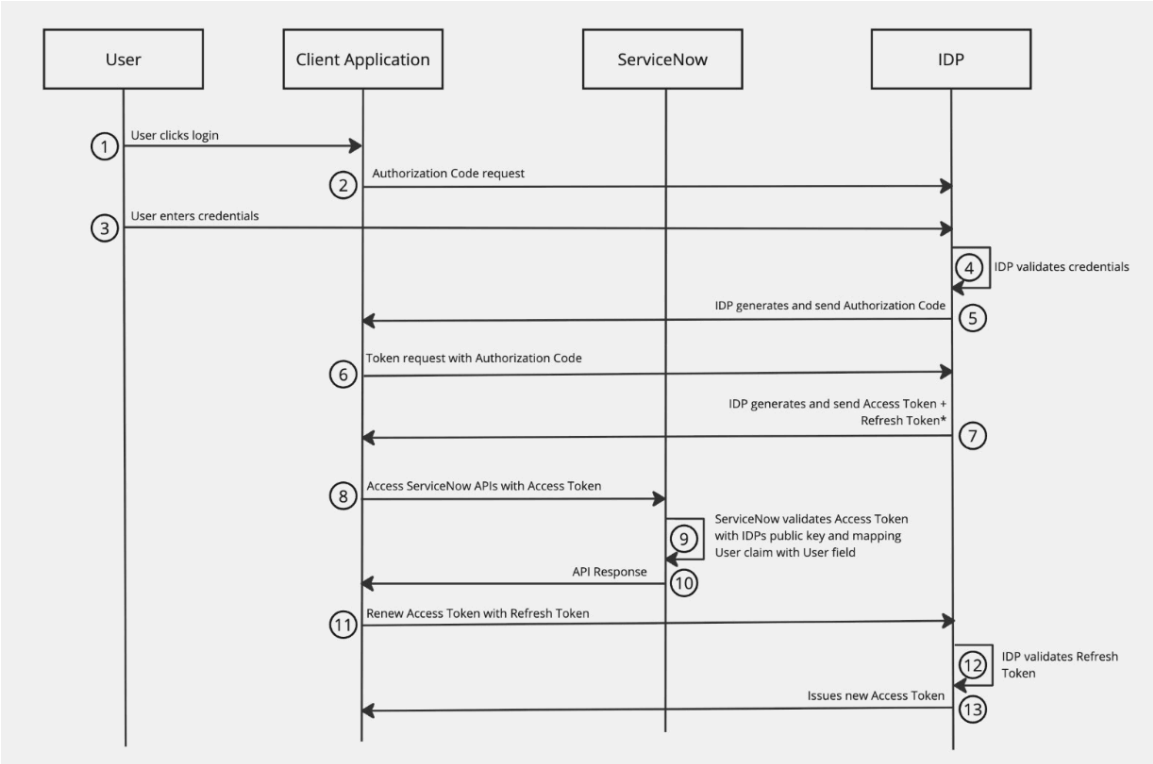

Note: This diagram is for illustrative purpose. It shows the Authorization code grant flow between your client application and the identity provider. The workflow is flexible. You can use a different flow based on your requirements.

Procedure

1. Configure your third party client application.

   Set up your third party client application to request tokens directly from your identity provider (IdP). Select an authentication method that best fits your security and integration requirements.

2. Create an OAuth client in ServiceNow.

   Provide the required details to enable validation of incoming tokens from your identity provider (IdP). For more information on how to configure, see Configure a third party ID token

Third party token workflow for service accounts

Create a service account in ServiceNow® to represent the identity of a third-party application accessing APIs through a trusted identity provider (IdP). This account maps the token claims to a user record and manages access with roles and permissions.

Before you begin

Role required: `oauth_admin, mi_admin, admin`

About this task

When a third-party application authenticates using a token from an external identity provider (IdP), ServiceNow needs a corresponding user record to map the identity and apply access controls.

Create a corresponding `sys_user` account in ServiceNow for your service account. The value of the claim configured during the initial setup in the token issued by your Idp is mapped to the user field specified. This account represents the service identity in ServiceNow. You can restrict this account to API access only, and assign the necessary permissions by adding the appropriate roles and groups.

Service Account Workflow

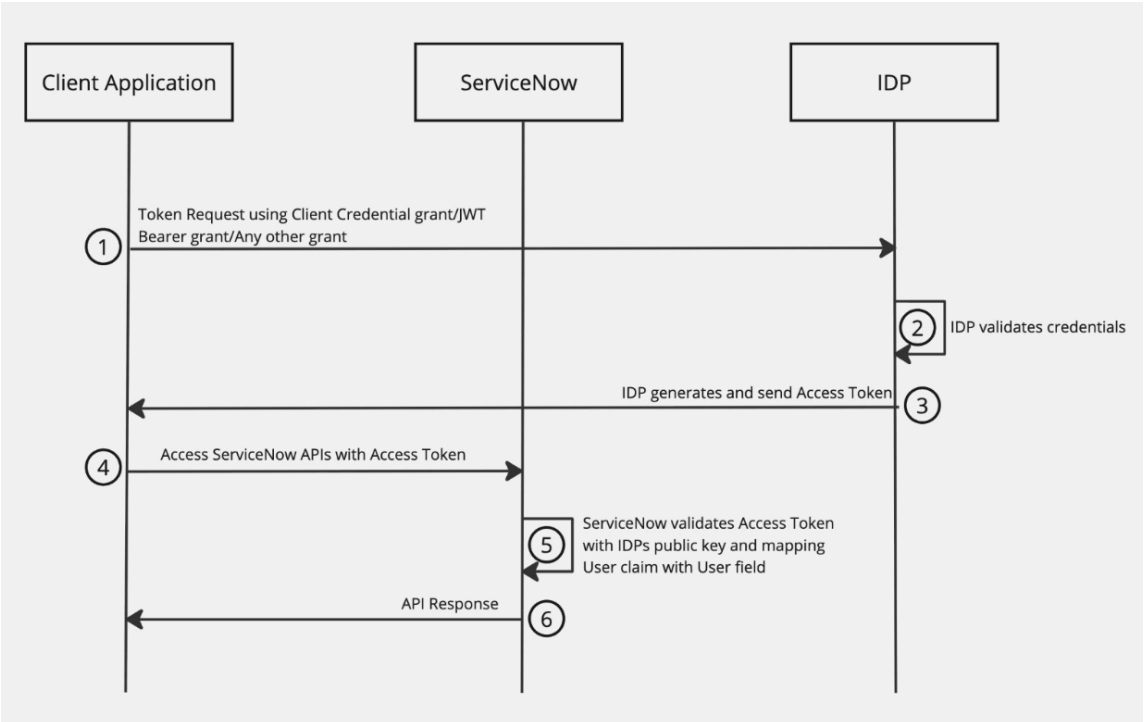

Procedure

1. Follow the Third party token workflow for user accounts to create a user account.
2. Create a `sys-user` account in ServiceNow to represent your service account identity. Ensure that the token claim value matches with that of the value in the mapped user field (such as user_name or email) in the user record. Example: user_name, email.
   a. Select the Web service access only option to restrict the account to API access.
   b. Assign the required roles and groups to grant the appropriate permission.

   The ServiceNow platform maps the configured claim to the specified user field in the `sys_user` record. It enforces access based on that user's assigned roles and groups.
3. Make a GET request with the authorization header to the following endpoint:

```
Method: GET
Endpoint: https:// <servicenow_base_url> /api/now/incident
Authorization: Bearer YOUR_THIRD-PARTY_TOKEN
```

Configure a third party ID token

Configure a third-party ID token to enable secure authentication by verifying user identities through an external IdP. The third-party ID token improves security by reducing stored credentials, confirms seamless authentication, and supports interoperability with industry standards like OpenID Connect (OIDC).

Before you begin

Role required: `oauth_admin, mi_admin, admin`

Procedure

1. Navigate to **Machine Identity Console** > > **Inbound integrations** > **New integration** > **Third party ID token**.
2. Update the text fields in the **Details** form with the appropriate information.

Details form

Field	Description
Name	The name provided by the resource owner (user) during authentication.
Provider name	Enter the name of the service provider you want to integrate with. Example: Microsoft, Google, Zoom, SAP, etc
Client ID	The unique ID assigned to identify the application.
Client secret	The secret key that only the application and the authorization server can identify. The application uses this key to authenticate and obtain access tokens.

Enforcing token restriction applies limitations on how an OAuth access token can be used, enhancing security by verifying tokens are valid only under specific conditions. Enable the **Enforce token restriction** check box to limit OAuth access tokens to specific APIs defined in the API access policy. If **Enforce token restriction** is turned off, the token can be used across other REST APIs.

3. Update the text fields in the **Auth scope (optional)** form with the appropriate information. The authentication scope defines the level of access an application has to a resource. Select the authentication scope for the specific REST APIs you want to access.

Auth scope form

Field	Description
Auth scope	The level of access an application has to a resource. The authentication scope restricts the actions that an access token can perform on APIs or data.
Limit authorization	The names of the APIs for which you want to restrict authorization.
Allow access only to APIs in selected scope	Enable the option for the integration to only access APIs that are explicitly listed in the selected scopes.

   a. Select **Create new auth scope** to add a new auth scope.

4. Update the text fields in the **Advanced options (optional)** form with the appropriate information.

Advanced options form

Field	Description
Access token lifespan	The duration (in seconds) for which the OAuth access token remains valid before it expires. **Note:** The default value is 1800 seconds.
Refresh token lifespan	The duration (in seconds) for which the OAuth refresh token remains valid before it expires. **Note:** The default value is 8,640,000 seconds.

5. Select **Save**.
   A new third-party ID token is created.
6. Go to **All** > **Inbound integrations** > **Application Registries** to view the newly created third party ID token.

JSON Web token bearer grant

Use this flow when a client application needs secure, unattended access to ServiceNow resources, either as itself or on behalf of a user.

The client application generates a signed JWT that includes identity-related claims, such as the user or system it represents. It sends it to the ServiceNow instance to request an access token.

JWT Structure

The JWT must be signed using the client's private key. It must include the following standard claims:

- iss – Issuer (client ID)
- sub – Subject (user or system identity)
- aud – Audience (ServiceNow token endpoint)
- exp – Expiration time
- iat – Issued at

Note: ServiceNow uses the public key (uploaded in the OAuth JWT profile) to validate the signature and maps the sub claim to a user record.

Related topics

JSON Web token grant workflow

Configure an OAuth JSON web token bearer grant

JSON Web token grant workflow

Use this flow when a client application needs secure, unattended access to ServiceNow resources, either as itself or on behalf of a user.

Before you begin

Role required: `oauth_admin, mi_admin, admin`

About this task

The client application generates a signed JWT with identity-related claims such as the user or system it represents. The client application sends the JWT to the ServiceNow instance to request an access token.

- **When acting on behalf of a user:**

  The token represents a previously authenticated user. It enables secure, seamless access without prompting the user for credentials or consent. ServiceNow trusts the request by validating the user's identity from the signed token, eliminating the need for real-time user interaction.

- **When acting as itself:**

  The token identifies and authenticates the client application. Instead of using a shared secret, the application signs the token with a private key. This offers a more secure alternative to the client credentials grant.

JWT Grant workflow

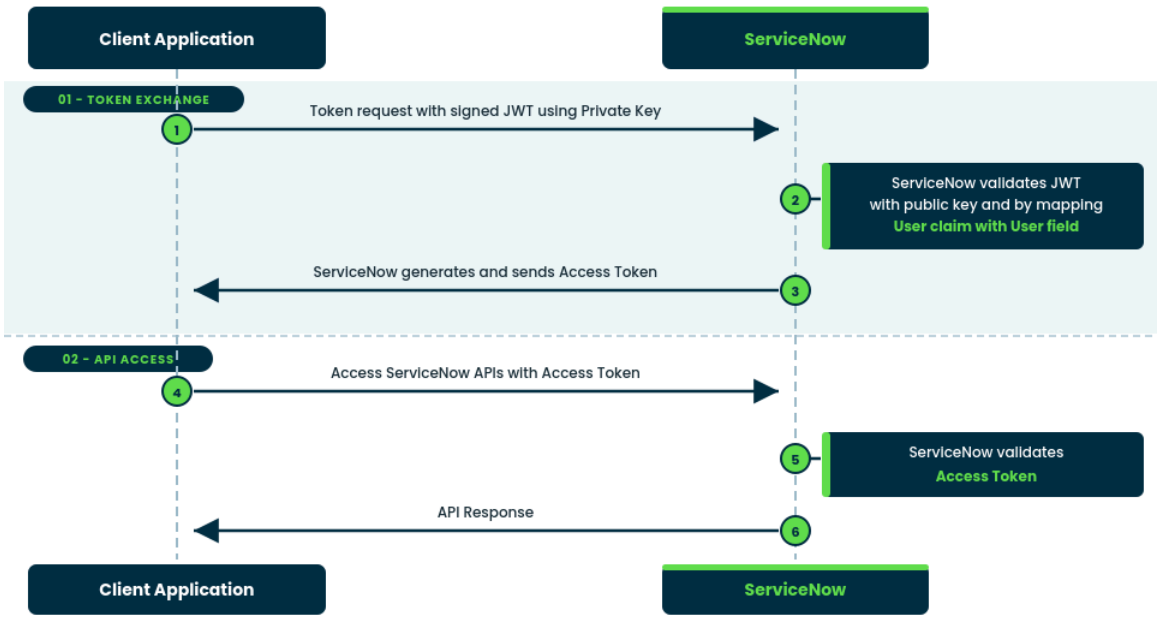

Procedure

1. The client application sends a token request to ServiceNow, with a JWT signed with its private key.
2. ServiceNow validates the JWT using the corresponding public key.
   It maps the `sub` (subject) claim in the token to a `sys_user` record.
3. ServiceNow validates the JWT, and issues the access token.
4. The client includes the access token in the API requests to ServiceNow.
5. ServiceNow validates the access token, and returns the appropriate API response.

Configure an OAuth JSON web token bearer grant

Configuring an OAuth JSON Web Token (JWT) bearer grant secures token-based authentication without user interaction. It enhances security with signed JWTs and reduces authentication overhead by eliminating repeated login attempts.

Before you begin

Role required: `oauth_admin, mi_admin, admin`

The supported algorithms for JSON Web Token (JWT): RS256, RS384, RS512, ES256, ES384, ES512, HS256, HS384, and HS512.

Procedure

1. Navigate to **Machine Identity Console** > > **Inbound integrations** > > **New integration** > **JWT bearer grant**.
2. Update the text fields in the Details form with the appropriate information.

Details form

Field	Description
Name	A unique name that identifies the application that you require JWT OAuth access for.

Field	Description
Provider name	Enter the name of the service provider you want to integrate with. Example: Microsoft, Google, Zoom, SAP, etc
Client ID	The auto-generated unique ID of the application. The system uses the value of this field to retrieve the public or shared key and validate the JWT. The value of this field should match the value of the issuer and audience claims in the JWT.
Client secret	The shared secret string that both the instance and the client application or website use to authorize communications with one another. Leave this field empty to have the instance auto-generate a client secret. To display existing client secrets, select the lock icon.
User field	Field in the User (sys_user) table that the system uses to match the value of the subject claim in the JWT. Example: If you add a token that has a subject claim value of user.name@example.com, then you would set the **User Field** to **Email**. This field tells the system to search the email field for the user.name@example.com value and use the matching user record in the inbound request.
Enable JTI Verification	Select to require a new token every token exchange. Default: Selected.
JWKS URL	The JSON Web key set URL. Its is a collection of public keys in JSON format. Identity providers publish a JWKS at a well-known URL so that client applications and services can retrieve the keys and validate the signatures of JSON Web Tokens (JWTs).
JTI claim	Unique identifier for each token. ServiceNow uses this claim to detect and prevent token replay by verifying that a token is not reused.
Comments	Add relevant comments.
JWT verifier map	Specify the identity provider, the verification method (such as a JWKS URL or certificate), and the mapping of JWT claims to ServiceNow user fields. Click on the Plus icon to add or edit the maps. Provide the following details in the **JWT verifier map** page:

Field	Description
	○ Name- The unique name of the JWT verifier map configuration. ○ Application - The application where the verifier map is used. ○ Kid (Key ID)- The identifier of the key used to validate the JWT signature. ○ Sys certificate - The ServiceNow certificate record used for token verification ○ Shared key - A symmetric key used to validate tokens.
Active	Select the check box to make the OAuth application active.

3. Update the text fields in the Advanced options (optional) form with the appropriate information.

Advanced options (Optional) form

Field	Description
Enforce token restriction	Select to only enable tokens to be used with APIs set to enable the authentication profile. You can set grant access using an API access policy. Default: Unselected.
JWKS cache lifespan	The duration (in minutes) for which ServiceNow caches the JSON Web Key Set (JWKS) from the identity provider.
Access token lifespan	The duration (in seconds) for which the access token remains valid before it expires.
Clock skew	Small differences in the system clocks of servers or devices involved in generating and validating a token can lead to issues when validating time-sensitive claims. Adjust the time above. Default value is: 0 seconds.

4. Update the text fields in the Auth scope (optional) form with the appropriate information.

Note: When you select an **Auth scope**, all the associated APIs are automatically populated in the **Limit authorization** text box.

Auth scope form

Field	Description
Auth scope	Access level of an application. The authentication scope restricts the actions that an access token can perform on APIs or data.

Field	Description
Limit authorization	Names of the APIs for which you want to restrict authorization.
Allow access only to APIs in selected scope	Enable the option for the integration to only access APIs that are explicitly listed in the selected scopes.

a. Select **Add another row** to add auth scopes.

b. Select **Create new auth scope** to add a new auth scope.
Enter the name of the Auth scope in the **Scope field** text box to select the newly created Auth scope. You can manually add and edit the APIs that needs to be associated with the new Auth scope.

> **Note:** Adding or editing APIs from the **Auth scope** menu affects all OAuth entities that are associated with the same authorization scope.

5. Update the text fields in the Advanced options (optional) form with the appropriate information.

Advanced options (Optional) form

Field	Description
Enforce token restriction	Select to only enable tokens to be used with APIs set to enable the authentication profile. You can set grant access using an API access policy. Default: Unselected.
JWKS cache lifespan	The duration (in minutes) for which ServiceNow caches the JSON Web Key Set (JWKS) from the identity provider.
Access token lifespan	The duration (in seconds) for which the access token remains valid before it expires.
Clock skew	Small differences in the system clocks of servers or devices involved in generating and validating a token can lead to issues when validating time-sensitive claims. Adjust the time above. Default value is: 0 seconds.
Token Format	Format of token to generate. Options: ○ JWT ○ Opaque **Note:** ○ The jwks url is available in the location: `api/now/oauth/jwks`. ○ The rotated (inactive keys) from jwks response after is removed after 105 days default.

6. Select **Save**.
   A new OAuth JSON Web Token bearer grant is created.
7. Go to **All** > **Inbound integrations** > **Application Registries** to view the newly created JWT bearer grant.

Resource owner password credential grant

Configuring an OAuth Resource Owner Password Credential (ROPC) grant enables applications to authenticate users by directly using their credentials to obtain an access token.

Security Considerations

The ROPC flow exposes user credentials directly to the client application, making it inherently less secure than modern alternatives. It should only be used in scenarios where the client is fully trusted, tightly controlled, and securely managed.

Avoid using this grant in modern applications unless absolutely necessary. For secure user-based access, it is strongly recommended to use the Authorization Code Flow with PKCE, which keeps credentials out of the client and leverages secure redirection and token handling practices.

Related topics

- Resource owner password credential grant workflow
- Configure an OAuth resource owner password credential grant

Resource owner password credential grant workflow

This flow is used in legacy or highly controlled environments where secure alternatives aren't feasible. The client app directly collects and sends user credentials to ServiceNow to obtain an access token, making it suitable only for trusted internal use.

Before you begin

Role required: `oauth_admin, mi_admin, admin`

About this task

This procedure outlines how a trusted client application obtains an access token by directly handling user credentials and uses it to access ServiceNow APIs.

The user logs in through the app, which sends both its own credentials and the user's to ServiceNow. ServiceNow validates the credentials and issues an access token that the app uses to call APIs.

Resource owner password credential grant workflow

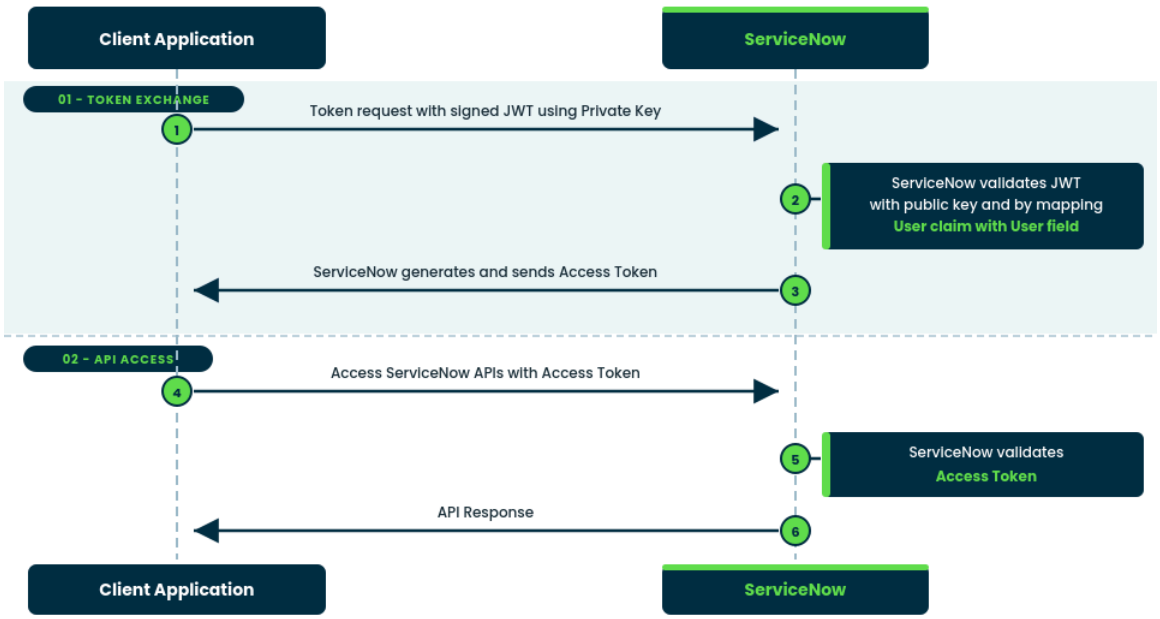

Procedure

1. The user logs in to the client application.
2. The client application sends a token request to with the following parameters:
   - Client ID and client secret.
   - Username and password of the user.

Example

```
Method: POST
Endpoint: https://<servicenow_base_url>/oauth_token.do
Headers: Content-Type: application/x-www-form-urlencoded
```

Token Request Parameters

Parameter	Required	Description
`grant_type`	Yes	Specifies the OAuth grant type.
`client_id`	Yes	The unique identifier for your client application. Format: YOUR_CLIENT_ID
`client_secret`	Yes	The client application's secret key. Format: YOUR_CLIENT_SECRET
`username`	Yes	The user's ServiceNow username.
`password`	Yes	The user's ServiceNow password.

Parameter	Required	Description
`scope`	Optional	Defines the level of access requested. Example: ○ incident_read ○ incident_write

3. ServiceNow validates both the client and user credentials, and returns the access token.
4. The client uses the access token to call ServiceNow APIs.

Example

```
Method: GET
Endpoint: https://<servicenow_base_url/api/now/incident
Authorization: Bearer YOUR_ACCESS_TOKEN
```

5. ServiceNow validates the access token, and returns the API response.

Configure an OAuth resource owner password credential grant

Configuring an OAuth resource owner password credential (ROPC) grant enables applications to authenticate users by directly using their credentials to obtain an access token. This method is ideal for trusted applications and legacy systems that require authentication without browser-based flows, enabling secure token validation and controlled API access.

Before you begin

Role required: `oauth_admin, mi_admin, admin`

Procedure

1. Navigate to **Machine Identity Console** > > **Inbound integrations** > > **New integration** > **OAuth Resource owner password credential grant**.
2. Update the text fields in the **Details** form with the appropriate information.

Details form

Field	Description
Name	The name provided by the resource owner (user) during authentication.
Provider name	Enter the name of the service provider that you want to integrate with. Example: Microsoft, Google, Zoom, SAP, and so on
Client ID	The unique ID assigned to identify the application.
Client secret	The secret key that only the application and the authorization server can identify. The application uses this key to authenticate and obtain access tokens.
Active	Select the check box to make the OAuth application active.

3. Update the text fields in the **Advanced options (optional)** form with the appropriate information.

   Enforcing token restriction applies limitations on how an OAuth access token can be used, enhancing security by verifying that tokens are valid only under specific conditions. Enable the Enforce token restriction check box to limit OAuth access tokens to specific APIs defined in the API access policy. If the Enforce token restriction is turned off, the token can be used across other REST API.

   Details form

Field	Description
Access token lifespan	The duration (in seconds) for which the access token remains valid before it expires.
Refresh token lifespan	The duration (in seconds) that a refresh token remains valid after it's issued is specified in the Lifespan field.

4. Update the text fields in the Auth scope (optional) form with the appropriate information. The authentication scope defines the level of access an application has to a resource. Select the authentication scope for the specific REST APIs you want to access.

   Note: When you select an **Auth scope**, all the associated APIs are automatically populated in the **Limit authorization** text box.

   Auth scope form

Field	Description
Auth scope	Access level of an application. The authentication scope restricts the actions that an access token can perform on APIs or data.
Limit authorization	Names of the APIs for which you want to restrict authorization.
Allow access only to APIs in selected scope	Enable the option for the integration to only access APIs that are explicitly listed in the selected scopes.

   Note: Adding or editing APIs from the **Auth scope** menu affects all OAuth entities that are associated with the same authorization scope.

   a. Select **Create new auth scope** to add a new auth scope.

5. Update the text fields in the **Advanced options (optional)** form with the appropriate information.

   Enforcing token restriction applies limitations on how an OAuth access token can be used, enhancing security by verifying tokens are valid only under specific conditions. Enable the Enforce token restriction check box to limit OAuth access tokens to specific APIs defined in the API access policy. If Enforce token restriction is turned off, the token can be used across other REST API.

Advanced options form

Field	Description
Enforce token restriction	The Enforce token restriction option limits the client to accessing only the APIs specified in the REST API Access Policies. If you unselect it, the client can access other REST APIs based on the user ACL permissions.
Token Format	Format of token to generate. Options: ○ JWT ○ Opaque **Note:** ○ The jwks url is available in the location: `api/now/oauth/jwks`. ○ The rotated (inactive keys) from jwks response after is removed after 105 days default.
Access token lifespan	Duration (in seconds) for which the OAuth access token remains valid before it expires. **Note:** The default value is 1800 seconds.
Refresh token lifespan	Duration (in seconds) for which the OAuth refresh token remains valid before it expires. **Note:** The default value is 8,640,000 seconds.

6. Select **Add another row** to create another Auth scope with the associated APIs.
7. Select **Save**.
   A new OAuth resource owner password credential grant is created.
8. Go to **All** > **Inbound integrations** > **Application Registries** to view the newly created OAuth Resource owner password credential grant.

Old inbound integrations experience

Old experience - Inbound integrations.

Note:

You can perform the OAuth inbound configuration, depending on the following type of grant type:

- OAuth authorization code grant flow

  Note: For authorization code flow, user needs to complete the Authentication by local login, SSO or MFA and then provide consent.

- Password grant
- JWT bearer grant flow
- ID token flow
- OAuth implicit grants
- Client Credentials

Configure OAuth integration that includes the following enhancements from Zurich release:

- Increase client secret length up-to 4096 characters to meet security requirements of third-party systems.
- Provide a JSON Web Key Set (JWKS) URL to automatically manage and update the public key for JSON Web Tokens (JWT) signature validation.
- Request OAuth tokens using the JWT grant type signed with Elliptic Curve Digital Signature Algorithm (ES) signing algorithms, including ES256, ES384, and ES512, for inbound JSON Web Tokens (JWT).
- Customize the JWT ID (JTI) claim name in both inbound OpenID Connect (OIDC) and JWT Bearer flows.

OAuth authorization code grant flow

Authorization code grant flow allows a user to access a resource by authenticating directly with an OAuth server that trusts the resource, in contrast with authenticating with username/password credentials.

This implementation of OAuth authorization code flow allows access to a resource via REST. The authorization code framework gets the access token through the authorized URL that the user configures rather than requiring the user to enter a username/password. The username/password are never exposed to the client that is requesting access to the resource.

A ServiceNow instance as the authorization server

The OAuth server is typically a third-party authorization server. You can also specify a ServiceNow instance as the authorization server that issues the tokens for authorization code flow.

The user who owns the restricted resource must authorize access. The user can also revoke the issued access token at any time to terminate access.

Authorization code grant flow process

The Authorization code grant flow process consists of these three steps:

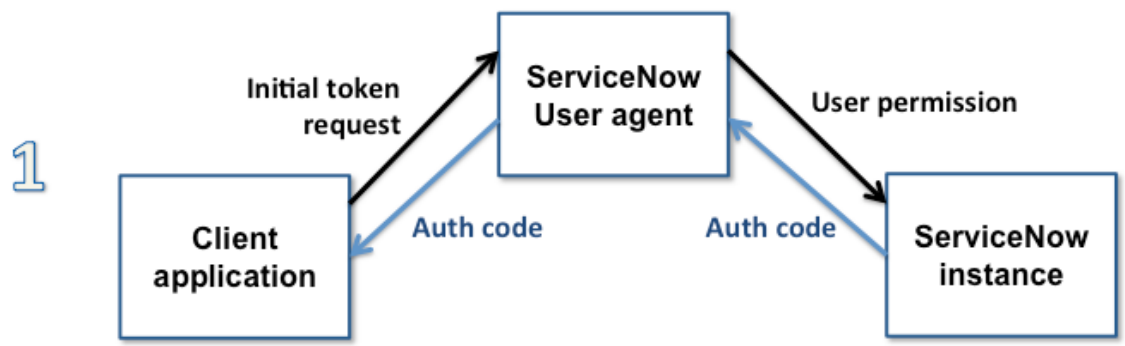

In step one, the client application or website initiates a REST API call in the form of a GET request to the instance via the user agent. Typically, the REST call is initiated when the end user clicks a button or a link on the client application or website to request an access token. In the client application, the end user also has to specify the authorization URL, token URL, client ID, and client secret. For an explanation of these items, see the field descriptions in this topic: Use a third-party OAuth provider. If the client asks for a grant type, the end user must select **Authorization Code**.

Example GET request from the client application to the instance:

```
https://myinstance.service-now.com/oauth_auth.do?response_type=code&redirect_uri={the_redirect_url}&client_id={the_client_identifier}
```

Note: The **response_type** must be **code** to use the standard OAuth code grant flow.

The end user must manually allow access to the restricted resource on the instance. In the ServiceNow implementation, the end user must be logged into the instance. The instance prompts the end user with a UI page that has **Allow** and **Deny** buttons.

The item that the client application is actually requesting the token from is the OAuth provider application registry record that you created, also known as the authorization endpoint (see Use a third-party OAuth provider). The auth code is sent from the authorization endpoint to the client. It does not go to the client directly but to the **Redirect URL** that you specify on the authorization endpoint form. This URL is also known as a callback URL. You can obtain this URL from the client application or website.

Example response from the instance to the client application, providing an authorization code:

```
https/http://{callbackURL}?code={the actual auth code}
```

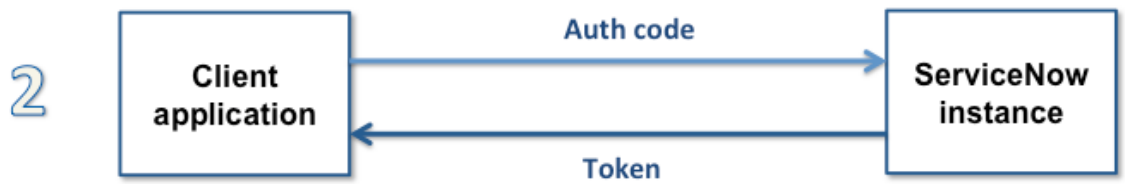

Now that the client application has the authorization code, the client uses the code to request the access token. The authorization code proves that the user has consented in step 1.

Example POST request from the client application to the ServiceNow instance that provides the auth code and requests the access token:

```
https://myinstance.service-now.com/oauth_token.do?grant_type=authorization_code&code={the auth
 code}&redirect_uri={the_same_redirect_url}&client_id={the_same_client_identifier}&client_secret={client_secret_value}
```

The endpoint on the instance returns an access token and a refresh token. The refresh token can be used to request additional access tokens.

You can manage the tokens, including revoking the token, in the instance. See Manage OAuth tokens.

The client application uses the access token to authenticate to the REST API. After authenticating the client application, the REST API returns the requested data in a JSON payload.

Example GET request for the JSON payload of data for the Incident [incident] table:

```
https://myinstance.service-now.com/api/now/table/incident?access_token={the_token}
```

Note: The system also supports OAuth implicit grants, also known as implicit grant code flow.

Integration support

Authorization code flow supports the following integrations on the instance:

- Multi-SSO
- SAML 2.0 Update 1
- Multifactor authentication

The mobile interface is also supported.

Authorize access to an OAuth endpoint using auth code flow

End users who own a protected resource on the ServiceNow instance must authorize access to the resource before the instance can provide the access token.

Before you begin

Role required: none. You must already be logged in to the instance that holds the protected resource. Alternatively, you can log in using the authentication method (such as multi-factor authentication or SAML) that your ServiceNow administrator already set up.

Procedure

1. Click the link or button on the client application where you are requesting access to the protected resource on the instance.
   This kicks off the token request. If you are making a REST call from one instance to another, this link is **Get OAuth Token** on the REST Message form.
2. If you are not logged in, log in now.
   If you are not the same user as the user specified in the upper-right corner, click **Not You?** and log in.
3. Click account permissions to open the list of access tokens that you have already issued.
   This is the same as the **Self-Service** > **My connected** apps token list.
4. Click **Allow** to allow access and have the instance issue the authorization code (if using auth code flow) or the access token (if using implicit grant type).

If you click **Deny**, the authorization is not allowed, but you are not logged out of the instance.

A message that confirms access should appear. If you are requesting access from the REST Message form on an instance, the following message appears at the top of the form: OAuth Refresh token is available and will expire at {date}.

Authorization code flow state parameter requirement

The glide.oauth.state.paramater.required system property enables the State parameter to be required in an OAuth request for authorization code flow.

State parameter

Role required: none.

Beginning in the Madrid release, the system property *`glide.oauth.state.parameter.required`* adds a State parameter for an OAuth request. For zbooted instances, the property is true. For upgraded instances, the property is not present, so the State parameter is not enabled. The State parameter is a string value, and should not contain special characters. The State parameter cannot be empty or " ".

Validating the state parameter

Create an endpoint for clients to access the instance. Initiate an authorization code flow for an oauth_auth.do. For example:

```
http://myinstance.service-now.com/oauth_auth.do?grant_type=autho
rization_code&client_id=e9dba45b380d1300e676ccc91cef468f&respons
e_type=code
```

If you do not specify the state parameter in the request, you get an error and the authorization code is not returned.`Missing State parameter in request.`

Adding the State parameter to the request:

```
http://myinstance.service-now.com/oauth_auth.do?grant_type=autho
rization_code&client_id=e9dba45b380d1300e676ccc91cef468f&respons
e_type=code&state=123
```

Adding the State parameter redirects you to the login screen and the regular authorization code flow returns the authorization code.

Note: The response URL contains the state parameter passed in the request. In the example, the added parameter is `state=123`.

If the authorization code flow starts from oauth_initiator.do:

```
http://myinstance.service-now.com/oauth_initiator.do?oauth_reque
stor_context=sys_rest_message&amp;oauth_requestor=eab8341fec0d13
00964f214a2c2fcf67&amp;oauth_provider_profile=dfa8f01fec0d130096
4f214a2c2fcf51&amp;response_type=code
```

The State parameter is automatically added when redirected by oauth_auth.do.

```
http://myinstance.service-now.com/oauth_auth.do?response_type=co
de&amp;state=-790938844&amp;redirect_uri=http://10.11.95.5:160
01/oauth_redirect.do&amp;client_id=e9dba45b380d1300e676ccc91cef4
68f
```

Authorization code flow example: ServiceNow instance as authorization server

You can use an instance as an authorization server to issue tokens to a client using authorization code flow.

Before you begin

Role required: none.

This example uses two instances: one as the authorization server and the other as the client. One instance uses a REST call to request tokens from another instance.

You must Activate OAuth on both instances.

Procedure

1. On the authorization server instance (running the Istanbul or later release), navigate to **System OAuth** > **Application Registry** and then click **New**.
2. Click **Create an OAuth API endpoint for external clients**.
3. Fill out the form fields for the OAuth application record as described in Create an endpoint for clients to access the instance.
   Completing these steps sets up an authorization server. Follow the next steps to set up the client server.
4. On the client instance, navigate to **System OAuth** > **Application Registry** and then click **New**.
5. Click **Connect to a third party OAuth Provider**.
6. Fill out the form fields for the OAuth application record as described in .
   Note the following field values:
   - **Name**: A unique name that identifies the application that you require OAuth access for.
   - **Client ID**: Client ID of the application registry record that you created for the authorization server.
   - **Client Secret**: [Read-Only] The auto-generated unique ID of the application. The instance uses the client ID when requesting an access token.
   - **Default Grant type**: Select **Authorization code**.
   - **Authorization URL**: URL of the instance that is the authorization server. Remember to append `oauth_auth.do` at the end of the URL.
   - **Logo URL**: The URL that contains an image to use as the application logo. The logo appears on the approval page when the user receives a request to grant a client application access to a restricted resource on the instance.
   - **Token URL**: URL of the instance that is the authorization server. Remember to append `oauth_token.do` at the end of the URL.
   - **Redirect URL**: URL of the instance: the client server instance. Remember to append `oauth_redirect.do` at the end of the URL.
7. Create a profile for the record with the **Authorization code** grant type.
   The client server is setup. You can now create an outbound REST message and get an OAuth token.

Create an endpoint for clients to access the instance

Create an OAuth application endpoint for external client applications to access the ServiceNow instance.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System OAuth** > **Application Registry** and then click **New**.
2. On the interceptor page, click **Create an OAuth API endpoint for external clients** and then fill in the form.

Field	Description
Name	A unique name that identifies the application that you require OAuth access for.
Client ID	[Read-Only] The auto-generated unique ID of the application. The instance uses the client ID when requesting an access token.
Client Secret	[Required] The shared secret string that both the instance and the client application or website use to authorize communications with one another. The instance uses the client secret when requesting an access token. Leave this field blank to have the instance auto-generate a client secret. To display existing client secrets, click the lock icon.
Redirect URL	The callback URL that the authorization server redirects to. Enter the full URLs of the clients requesting access to the resource, appended by `/oauth_redirect.do`. For example, `http://token_consumer:port/oauth_redirect.do`. Enter as many URLs as needed for all possible token consumers. The instance matches the URL of the incoming request to one of the redirect URLs. If no match is made, the instance uses the first redirect URL.
Logo URL	The URL that contains an image to use as the application logo. The logo appears on the approval page when the user receives a request to grant a client application access to a restricted resource on the instance.
Active	Select the check box to make the application registry active.
Refresh Token Lifespan	The number of seconds that a refresh token is valid. The instance uses the lifespan value when requesting a refresh token. By default, refresh tokens expire in 100 days (8640000 seconds).
Enforce Token Restrictions	Select to only allow tokens to be used with APIs set to allow the authentication profile. You can set grant access using an API access policy. For more information, see Create REST API access policy. Default: Unselected.
Mobile Client	Represents the entity for mobile app or web. This information is used to analyze the login information with mobile or web.
Access Token Lifespan	The number of seconds that an access token is valid. The instance uses the lifespan value when requesting an access token. By default, access tokens expire in 30 minutes (1800 seconds).
Token Format	Format of the token to generate. The format determines the structure of a token and the information it includes.
Subject Claim	Field in the User (sys_user) table that's used to populate the value of the subject (sub) claim of a JWT token. The sub claim is a piece of information that identifies the subject, or user, of the JWT token. This field only applies if the **Token Format** is JWT.
Comments	Additional information to associate with the application.
Client Type	Choose the client type, based on the type of your client. Options:

Field	Description
	○ **Iframe Embedded** ○ **Integration as a User** ○ **Integration as a Service** To know more, see Configure client type for OAuth and SSO records.

3. Click **Submit**.

Result

The system creates a record in the Application Registries [oauth_entity] table with of type OAuth Client. When the instance actually issues tokens and authorization codes, they are stored in the table. See Manage OAuth tokens for more information.

OAuth API response parameters

The OAuth 2.0 API produces a JSON response containing the following parameters as name:value pairs.

Access token response parameters

Response parameter	Description
scope	Amount of access granted by the access token. The scope is always **useraccount**, meaning that the access token has the same rights as the user account that authorized the token. For example, if Abel Tuter authorizes an application by providing login credentials, then the resulting access token grants the token bearer the same access privileges as Abel Tuter.
token_type	Type of token issued by the request as defined in the OAuth RFC. The token type is always **Bearer**, meaning that anyone in possession of the access token can access a protected resource without providing a cryptographic key. See RFC6750 for more information about how OAuth 2.0 uses bearer tokens.
expires_in	Lifespan of the access token in seconds.
refresh_token	String value of the refresh token.
access_token	String value of the access token. Access requests made within the access token expiration time always return the current access token.
format	[Optional] Output format of the response. This value is always JSON.

Note: If any OAuth provider sends the response body as "content-type" instead of "Content-Type", the OAuth HTTP client may not parse the response correctly. To correct this issue, create a system property using these parameters.

Field	Value
Name	glide.oauth.inhouse.httpclient.enabled
Type	true|false
Value	false

For details on creating system properties, see Add a system property

The following example illustrates the JSON string returned by an access token request. (Spaces have been added to improve readability).

```
{"scope":"useraccount","token_type":"Bearer","expires_in":1800,
"refresh_token":"w599voG89897rGVDmdp12WA681r9E5948c1CJTPi8g4HGc4
NWaz62k6k1KOFMxHW4OH8yOO3Hoe",
"access_token":"FOjh9korTyzd9kaZqZOSzjKZuS3ut0i4P46Lc52m2JYHiLIc
qzFAumpyxshU9mMQ13gJHtxD2fy"}
```

OAuth API request parameters

Learn about the OAuth API request parameters that access token requests use.

Note: The content-type of the OAuth API should be *application/x-www-form-urlencoded*. A content-type of *application/json* results in an unspecified error.

Access token request parameters

Request parameter	Description
grant_type	[Required] The type of credentials authorizing the request for an access token. This parameter must have one of the following values: • **password**: A set of user credentials to authorize the access token request. Specify the user credentials in the username and password parameters. • **refresh_token**: An existing refresh token authorizes the access token request. Specify the refresh token in the refresh_token parameter.
client_id	[Required] Auto-generated unique ID of the client application requesting the access token.
client_secret	[Required] Shared secret string that the instance and the OAuth application use to authorize communications with one another.
username	User account name that authorizes the access token request. This parameter is required for access token requests with a **grant_type** of **password**.
password	Password for the user account that authorizes the access token request. This parameter is required for access token requests with a **grant_type** of **password**.
refresh_token	Existing refresh token that authorizes the access token request. This parameter is required for access token requests with a **grant_type** of **refresh_token**.

Requests Using User Credentials

The instance requires clients to provide user login credentials when first authorizing the client or when authorizing the creation of a new refresh token. This type of request always returns two tokens:

- An access token
- A refresh token

The instance verifies that the user is active, not currently locked out, and has an interactive session. If any of these conditions are false, the instance does not produce an access token.

Access requests made within the expiration time of the access token always return the current access token.

Note: This type of authorization grant relies on TLS encryption to protect the user credentials during transmission.

The following example illustrates requesting an access token with a set of user credentials (Spaces have been added to improve readability).

```
$ curl
 -d"grant_type=password&client_id=be3aeb583ace210011c15b24a43e2
5d8
&client_secret=client_password
&username=admin&password=admin"
https://instancename.service-now.com/oauth_token.do
```

Requests Using a Refresh Token

The instance can use an existing refresh token to create a new access token. This type of request returns only an access token. The instance confirms that the refresh token has not expired before generating a new access token. Access requests made within the refresh token expiration time always return the current refresh token. Transmitting refresh tokens is generally more secure than transmitting user credentials. The following example illustrates requesting an access token with an existing refresh token (Spaces have been added to improve readability).

```
$ curl
 -d"grant_type=refresh_token&client_id=be3aeb583ace210011c15b24a
43e25d8
&client_secret=client_password
&refresh_token=w599voG89897rGVDmdp12WA681r9E5948c1CJTPi8g4HGc4NW
az62k6k1K0FMxHW40H8yOO3Hoe"
https://instancename.service-now.com/oauth_token.do
```

Create an OAuth JWT API endpoint for external clients (machine to machine integration)

OAuth JWT bearer token enables the client web applications to authenticate with your instance seamlessly using the inbound JWT grant type instead of requiring the end user to manually log in or share the password.

Before you begin

The supported algorithms for JSON Web Token (JWT): RS256, RS384, RS512, ES256, ES384, ES512.

Generate a JWT with the following claims at the client side:

- `aud`: Must match the value of the Client ID.
- `sub`: Must be a user identifier, such as the user's mail that you want to associate the token with.
- `iss`: Recommended matching the value of the Client ID. If the `aud` and `iss` isn't matching, then add the `iss` value in the claim validation.
- `exp`: Any desired expiration.

Example decoded JSON Web Token

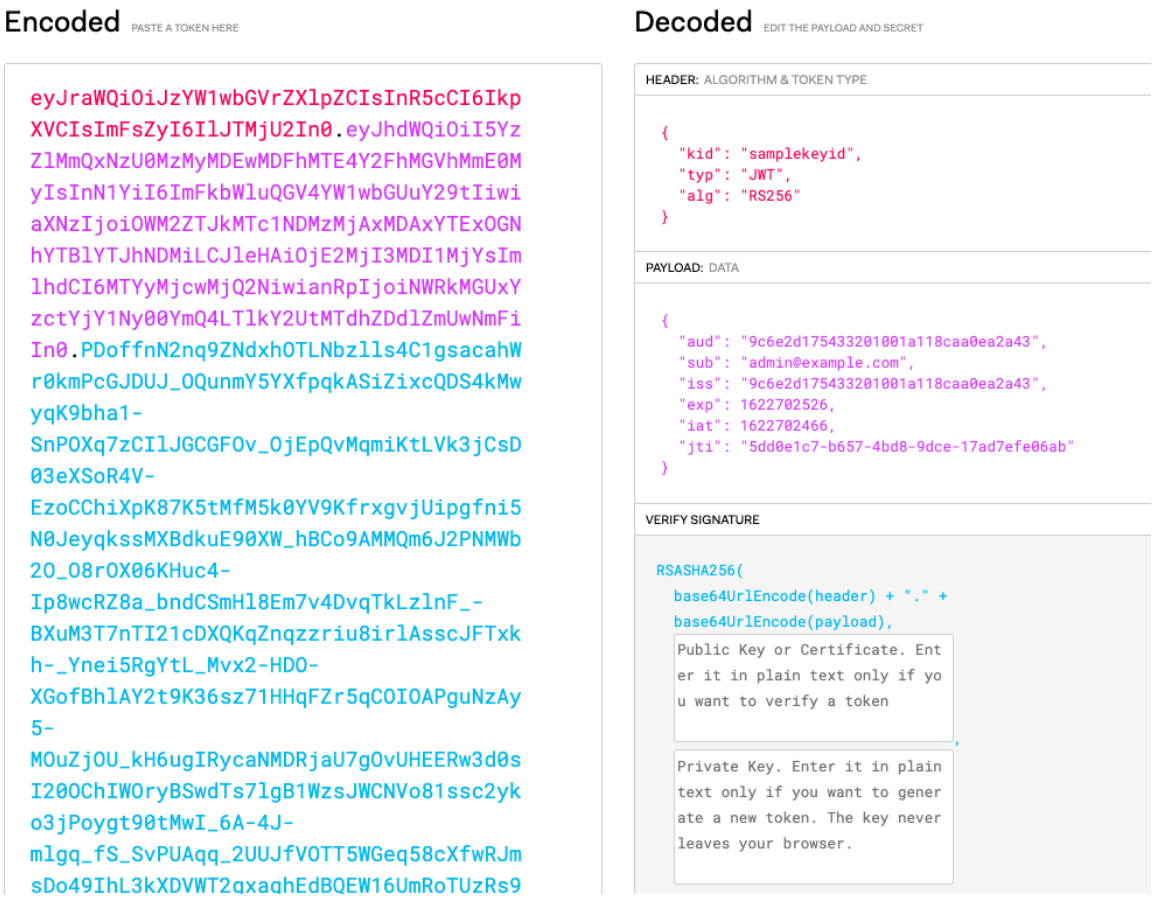

Configuration in ServiceNow

About this task

Since using the JWT grant type doesn't include the password in the request, it enables a greater security between web services. For example, you can develop an external application and use tokens to authenticate inbound requests to your ServiceNow instance.

Role required: admin

For more information about JSON Web Tokens, see https://jwt.io/ .

Procedure

1. Add the public key of the client app to the **sys_certificate** table.
2. Set up the configuration in your ServiceNow instance to verify the incoming JWT.
   a. Navigate to **System OAuth** > **Application Registry**.
   b. Select **Create an OAuth JWT API endpoint for external clients**.
   c. Complete the form with information about your token.

OAuth JWT table

Field	Description
Name	A unique name that identifies the application that you require JWT OAuth access for.
Client ID	The auto-generated unique ID of the application. The system uses the value of this field to retrieve the public or shared key and validate the JWT. The value of this field should match the value of the issuer and audience claims in the JWT.
Client Secret	The shared secret string that both the instance and the client application or website use to authorize communications with one another. Leave this field empty to have the instance auto-generate a client secret. To display existing client secrets, select the lock icon. **Note:** If **Public Client** is selected, you can omit the **Client Secret**.
User Field	Field in the User (sys_user) table that the system uses to match the value of the subject claim in the JWT. For example, if you add a token that has a subject claim value of user.name@example.com, then you would set the **User Field** to **Email**. This field tells the system to search the email field for the user.name@example.com value and use the matching user record in the inbound request.
Enable JTI Verification	Select to require a new token every token exchange. Default: Selected.
Application	Read-only application scope. This field is auto-populated.
Accessible from	Cross-scope access policy. For more information, see Application access settings.
Access Token Lifespan	Amount of time that the token is valid. Unit: Seconds
Token Format	Format of the token to generate. The format determines the structure of a token and the information it includes.
Subject Claim	Field in the User (sys_user) table that's used to populate the value of the subject (sub) claim of a JWT token. The sub claim is a piece of information that identifies the subject, or user, of the JWT token. This field only applies if the **Token Format** is JWT.
Clock Skew	Allowed time difference between the server and client clocks when validating the `exp` and `nbf` claims in the JWT. Unit: Seconds Default: 300
Enforce Token Restrictions	Select to only enable tokens to be used with APIs set to enable the authentication profile. You can set grant access using an API access policy. For more information, see Create REST API access policy.

Field	Description
	Default: Unselected.
Comments	Additional information to associate with the application.
Public Client	Add this field to the form if the JWT client is public. When selected, you don't need to include a **Client Secret**. Default: Unselected.
Client Type	Choose the client type, based on the type of your client. Options: ▪ **iframe Embedded** ▪ **Integration as a User** ▪ **Integration as a Service** To know more, see Configure client type for OAuth and SSO records.

d. Save the form.

e. Add records to the JWT Verifier Maps related list to verify the JWT signature.

JWT Verifier Maps table

Field	Description
Name	Name of the JWT-mapping record.
Kid	Key ID from the JWT.
Shared Key	The shared key for the specified key ID.
Application	Read-only application scope.
Sys certificate	Certificate record in the X.509 Certificates (sys_certificate) table. The certificate that was uploaded in step 1.

f. Add any custom claims associated with your JWT to the OAuth JWT Claim Validations related list.

You don't need to add records for the following required claims:

- iss
- aud
- sub
- exp

Note:

- If the `aud` and `iss` isn't matching, then add the `iss` value in the claim validation.
- For certificates, you can add multiple verifier maps associated with multiple keys.

OAuth JWT Claim Validations table

Field	Description
My external client	Auto-populated with the OAuth JWT record.
Claim Value Type	Data type of the claim value.
Claim Name	Name of the claim you want to add.
Claim Value	Value of the claim.
Application	Read-only application scope.

3. Send a cURL request containing the JWT token to obtain an access token from your instance.

Example

The following is a sample cURL command requesting an access token:

```
$ curl
 -d"grant_type=urn:ietf:params:oauth:grant-type:jwt-bearer
&client_id=be3aeb583ace210011c15b24a43e25d8
&client_secret=client_password
&assertion=
 eyJraWQiOiJzYW1wbGVrZXlpZCIsInR5cCI6IkpXVCIsImFsZyI6IlJTMjU2In
0.eyJhdWQiOiI5YzZlMmQxNzU0MzMyMDEwMDFhMTE4Y2FhMGVhMmE0MyIsInN1Y
iI6ImFkbWluQGV4YW1wbGUuY29tIiwiaXNzIjoiOWM2ZTJkMTc1NDMzMjAxMDAx
YTExOGNhYTBlYTJhNDMiLCJleHAiOjE2MjI3MDI1MjYsImlhdCI6MTYyMjcwMjQ
2NiwianRpIjoiNWRkMGUxYzctYjY1NyOOYmQ4LTlkY2UtMTdhZDdlZmUwNmFiIn
0.PDoffnN2nq9ZNdxhOTLNbzlls4C1gsacahWr0kmPcGJDUJ_OQunmY5YXfpqkA
SiZixcQDS4kMwyqK9bha1-SnPOXq7zCIlJGCGFOv_OjEpQvMqmiKtLVk3jCsDO3
eXSoR4V-EzoCChiXpK87K5tMfM5k0YV9KfrxgvjUipgfni5NOJeyqkssMXBdkuE
90XW_hBCo9AMMQm6J2PNMWb2O_O8rOXO6KHuc4-Ip8wcRZ8a_bndCSmHl8Em7v4
DvqTkLzlnF_-BXuM3T7nTI21cDXQKqZnqzzriu8irlAsscJFTxkh-_Ynei5RgYt
L_Mvx2-HDO-XGofBhlAY2t9K36sz71HHqFZr5qCOIOAPguNzAy5-MOuZjOU_kH6
ugIRycaNMDRjaU7gOvUHEERw3d0sI2OOChIWOryBSwdTs7lgB1WzsJWCNVo81ss
c2yko3jPoygt90tMwI_6A-4J-mlgq_fS_SvPUAqq_2UUJfVOTT5WGeq58cXfwRJ
msDo49IhL3kXDVWT2gxaqhEdBQEW16UmRoTUzRs9A9sOm18y3skmOVtnEOm-MlJ
MFQZ754UMzbiHOZsMmk1ivCGIjex5JO_lDjKElWF5RHGz3YShCoa4JKDZsqYMvI
k1SvzyQXjuFqPdS2vzg2m1eKGUwr3m6uNs_HflcDystwVdMZ7nLlBG4"
https://instancename.service-now.com/oauth_token.do
```

If the JWT client is a public client, such as the Mobile SDK, then you can omit the client_id and client_secret parameters from the request. The following is a sample cURL command requesting an access token that omits the client_id and client_secret:

```
$ curl
 -d"grant_type=urn:ietf:params:oauth:grant-type:jwt-bearer
&assertion=
 eyJraWQiOiJzYW1wbGVrZXlpZCIsInR5cCI6IkpXVCIsImFsZyI6IlJTMjU2In
0.eyJhdWQiOiI5YzZlMmQxNzU0MzMyMDEwMDFhMTE4Y2FhMGVhMmE0MyIsInN1Y
iI6ImFkbWluQGV4YW1wbGUuY29tIiwiaXNzIjoiOWM2ZTJkMTc1NDMzMjAxMDAx
YTExOGNhYTBlYTJhNDMiLCJleHAiOjE2MjI3MDI1MjYsImlhdCI6MTYyMjcwMjQ
2NiwianRpIjoiNWRkMGUxYzctYjY1NyOOYmQ4LTlkY2UtMTdhZDdlZmUwNmFiIn
0.PDoffnN2nq9ZNdxhOTLNbzlls4C1gsacahWr0kmPcGJDUJ_OQunmY5YXfpqkA
SiZixcQDS4kMwyqK9bha1-SnPOXq7zCIlJGCGFOv_OjEpQvMqmiKtLVk3jCsDO3
eXSoR4V-EzoCChiXpK87K5tMfM5k0YV9KfrxgvjUipgfni5NOJeyqkssMXBdkuE
90XW_hBCo9AMMQm6J2PNMWb2O_O8rOXO6KHuc4-Ip8wcRZ8a_bndCSmHl8Em7v4
```

```
DvqTkLzlnF_-BXuM3T7nTI21cDXQKqZnqzzriu8irlAsscJFTxkh-_Ynei5RgYt
L_Mvx2-HDO-XGofBhlAY2t9K36sz71HHqFZr5qCOIOAPguNzAy5-MOuZjOU_kH6
ugIRycaNMDRjaU7gOvUHEERw3dOsI2OOChIWOryBSwdTs7lgB1WzsJWCNVo81ss
c2yko3jPoygt90tMwI_6A-4J-mlgq_fS_SvPUAqq_2UUJfVOTT5WGeq58cXfwRJ
msDo49IhL3kXDVWT2gxaqhEdBQEW16UmRoTUzRs9A9sOm18y3skmOVtnEOm-MlJ
MFQZ754UMzbiHOZsMmk1ivCGIjex5JO_lDjKElWF5RHGz3YShCoa4JKDZsqYMvI
k1SvzyQXjuFqPdS2vzg2m1eKGUwr3m6uNs_HflcDystwVdMZ7nLlBG4"
https://instancename.service-now.com/oauth_token.do
```

The instance returns the access token in its response:

```
{
    "access_token":
 "KynMY2HOuwWkRc8g8YLXjnQxWbH5_wbnSiLsnaOoKw61GZkkV0ytZP74uF7hJ
yjfsWfaaFijqQzq2kcABNJxNA",
    "scope": "useraccount",
    "token_type": "Bearer",
    "expires_in": 1799
}
```

Note: The inbound JWT grant type doesn't include refresh tokens.

4. Make a REST API call to access a resource using the access token.

Example

The following is a cURL command to access the incident table using the token.

```
$ curl -H "Authorization: Bearer
 KynMY2HOuwWkRc8g8YLXjnQxWbH5_wbnSiLsnaOoKw61GZkkV0ytZP74uF7hJy
jfsWfaaFijqQzq2kcABNJxN"
https://instancename.service-now.com/api/now/v1/table/incident
```

Result

The system retrieves the access token in the REST call and enables access to the requested resource.

Configure an OAuth OIDC provider for accepting third-party token

Configure an OAuth OpenID Connect (OIDC) provider to accept identity tokens generated by a third-party OIDC provider using inbound API calls using Single Sign-On option (Multi-Provider SSO).

Before you begin

Role required: admin

About this task

The ServiceNow AI Platform supports OpenID Connect (OIDC) through the external Single Sign-On (SSO) implementation in addition to inbound API calls.

For an example of an OIDC provider configuration, see Set up Microsoft Entra ID spoke. For an SSO-specific example of an OIDC provider configuration, see Create an OpenID Connect (OIDC) configuration for Single Sign-On (SSO).

Procedure

1. Navigate to **All** > **System OAuth** > **Application Registry**.
2. Select **Submit**.

The record is saved in the Application Registries [oauth_entity] table.
When your instance issues tokens and authorization codes it creates a record in the Application Registries [oauth_entity] table with type **External OIDC Provider**. See for more information.

3. **Optional:** Go to the related list on the record OAuth Entity Profiles to validate a system-generated default profile for the new OAuth provider without any scope.
   You can change or add an OAuth provider profile including the name, grant type, and OAuth Scope.
4. **Optional:** Go to the related list on the record OAuth Entity Scopes to define all available OAuth scopes for this OAuth provider.
   The scopes defined can be selected when you create or update a profile. Each OAuth scope defined contains a name and a scope that you must get from the provider specification, such as a read-scope or a write-scope. Each scope must be defined separately.
5. **Optional:** Go to the related list on the record User Provisioning to enable automatic user provisioning.

Option	Description
Automatically provision users	Option to enable force authentication for users.
Provision data source	The data source to use to transform an OIDC token to a ServiceNow user. Use the **Lookup list** to select the predefined data source tem plate, then open the record to configure the Transformed table mapping. When configur ing the Transform mapping, the source fields are from the *JWT token*, the target fields are from the *sys_user* table.
User roles applied to provisioned users	The user roles applied to the newly provi sioned ServiceNow users.

Example: The following is an example of a cURL request to invoke a REST API call

Invoke a REST API call.

Perform the following steps:

- Register the app in the OpenID Connect Provider.
- Configure the OAuth OIDC Entity.
- Configure the OIDC Provider:

OIDC Provider

OIDC Provider	Name of the OIDC provider.
OIDC Metadata URL	Specify the OIDC Metadata URL (well-known configuration URL). This information is used to fetch the public keys to validate the token through the `jwks` endpoint.
User Claim	The the claim which is validated against user table.
User Field	User claim which identifies user record.

Enable JTI claim verification	When enabled, the ServiceNow JWT token validation will also validate the JTI sent by the provider. **Note:** If validation isn't checked, the `jti` can't be validated, regardless if it's present in the JWT token. The claim name in the token must be `jti`. This information is used to help prevent replay attacks.

- Get a JWT token.
- Invoke a REST API call.
  - The ID token in the Authorization header to access Table API or scripted web service.

```
curl -X GET --header "Accept:application/json"
 https://<instance_name>.service-now.com/api/now/table/incide
nt/897b04f2dbd4a300a135364e9d961952 -k
--header "Authorization: Bearer
 eyJraWQiOiJjNTZtZTlXU0xPVnY3UFMwcTg4Qzl1bOlzNjFQYTdmUG4yZFVF
OW9RNUg4IiwiYWxnIjoiUlMyNTYifQ.eyJzdWIiOiIwMHVnZDg1OD
VkczI1WXpUSjBoNyIsIm5hbWUiOiJpbXJhbiBhbGkiLCJsb2NhbGUiOiJlbi1
VUyIsImVtYWlsIjoiaW1yb241NDNAZ21haWwuY29tIiwidmVyIjoxLCJpc3Mi
OiJodHRwczovL2RldiO5MzQ
xMjEub2t0YXByZXZpZXcuY29tIiwiYXVkIjoiMG9hZ2Q4bzk3a2lCT3dwd0Iw
aDciLCJpYXQiOjE1Mzc5MzMzMjYsImV4cCI6MTUzNzkzNjkyNiwianRpIjoiS
UQueThVdXpWNUg2bm16SzRs
OTI1RFVrQnJoR1o1MmJzVVpGVHRVTEphQjg3ayIsImFtciI6WyJwd2QiXSwia
WRwIjoiMDBvZ2Q4NTgycEFqZDZTemcwaDciLCJub25jZSI6InNub3ciLCJwcm
VmZXJyZWRfdXNlcm5hbWUiO
iJpbXJvbjU0MOBnbWFpbC5jb2OiLCJnaXZlbl9uYW1lIjoiaW1yYW4iLCJmYW
1pbHlfbmFtZSI6ImFsaSIsInpvbmVpbmZvIjoiQW1lcmljYS9Mb3NfQW5nZWx
lcyIsInVwZGFOZWRfYXQiOj
E1Mzc5MzAxOTcsImVtYWlsX3ZlcmlmaWVkIjpOcnVlLCJhdXRoX3RpbWUiOjE
1Mzc5Mjk2NjF9.OG87SYxWFgHGlhBYby2H79diRm9rlYZTeEkIINRUatwg-p4
739htB8xEY-5_t6yU_6k5w1
Opdgtt5M5QFZRPXVbQZNoGtY-BxnOBjaimcFgoWfhY_OldnGTkzN2RYyIHvr
f9-yhxg347zvczmLrgMMa_VwG4rxrtE6rUXaIpIeIK5b-Deq8ADz8UTUTKpF_
5RWk4X-oh5xK6BLniFHk4ShO
Zq2v_mjproXwKk5euJKrVrar2lQ4adZCOSTRuTf3ThMO5WDhOsel-82LngXtL
zRJJ51IqxAsXnsOkJHLLqLtH1hXNRKfwT1ScQoE_OfWm4tOKryI2j4wSMEanF
tLXIw"
```

  - If the user is authenticated, a valid application/json response is returned. Otherwise, a "User Not Authenticated" error message is returned.

```
User Not Authenticated
{"error":{"message":"User Not
 Authenticated","detail":"Required to provide Auth
 information"},"status":"failure"}
```

Configure client type for OAuth and SSO records

Configure the **Client Type** field for OAuth and SSO record related configurations.

When establishing sessions for various login types such as Web UI (interactive login), Iframe Embedded, Embedded or Integration, you can configure the client type for the OIDC (OAuth Entity), SAML, and Digest records that are used for various login.

The client type choices are as follows:

- **Iframe Embedded**: Can be used for interactive ServiceNow instance that resides in an Iframe in the third party websites. For example, if there is a sensitive table (sys_user table - phone number of the user), as admin you can configure the ACL with the sec attribute (Iframe Embedded) set to false, to avoid the user accessing the data (table information) on the Iframe Embedded session in the third-party.
- **Integration as a User**: Can be used for Virtual Agent chat bot that are installed in desktop apps such as Slack, Teams etc.
- **Integration as a Service**: Can be used for machine to machine integration (communication between services).

Note:

- For OIDC (OAuth Entity): All the choices are available.
- For SAML and Digest: Only Iframe embedded. You must edit the form and add the **Client Type** field to the record and select the **Iframe Embedded** client type.
- If the **Client Type** selected is none, then there is no classification of the session.

It is recommended to use client type field to every records created for OIDC (OAuth Entity), SAML, and Digest. This enables you to have a better control over each login methods that has the same configurations but distinguished with the client type.

After configuring the field, whenever a user logs in from the corresponding configuration (OAuth or SSO), once the authentication is successful, the session is considered based on the configured client type and accordingly the session timeout is leveraged.

For the current session, the corresponding security attribute are included or can be leveraged to prevent users from accessing table specific information within the selected client type. For more information, see OOB (Out-of-Box) Security Attributes.

Session time out for the Client Types

Following are the system properties related to session time out for the various client types:

- `glide.session_timeout.iframe_embedded`
- `glide.session_timeout.integration_as_a_user`
- `glide.session_timeout.integration_as_a_service`

OAuth implicit grants

ServiceNow instances support the implicit grant of an access token.

The implicit grant type, also known as *`implicit grant code flow`*, allows the access token to be given directly to the client application via the user agent, which is typically the web browser or mobile device. No refresh tokens are granted. The end user must still grant access to the protected resource on the instance, just as with standard.

OAuth implicit grant flow process

Just as with the standard authorization code flow process, the client application makes a request to use the restricted resource on the instance and the end user approves it. The request is in the form of a URL sent to the instance. The URL must include the following parameters:

- `client_id=<the necessary client ID>`. This is mandatory to identify which protected resource the client application wants access to.
- `response_type=token`. This is mandatory to request the access token directly (as opposed to asking for an authorization code). The value must be `token` for implicit grants. In the standard authorization code flow example, the response type is `code`.
- `redirect_uri=<a URL>`: The location where the token is sent.

The authorization server sends the access token, rather than an authorization code, to client application via the user agent.

Here is an example GET request to receive the JSON payload of data for the Incident [incident] table:

```
https://myinstance.servicenow.com/oauth_auth.do?response_type=token&redirect_uri={the_redirect_url}&client_id={the_client_identifier}
```

If the user grants access, the token is included in the redirect (callback) URL:

```
https/http://{callbackURL}?access_token={the_token}
```

Client Credentials

Use the OAuth client credentials grant type for Inbound Integrations from a third party OAuth client to the ServiceNow® platform.

The administrators can use the client credential (CC) grant type to enable integration from a third party OAuth client to the ServiceNow platform.

Inbound client credentials grant type is a capability that can be controlled through a system property. By default the system property is false.

To use the client credentials grant type, you must perform the following steps:

- Create the `glide.oauth.inbound.client.credential.grant_type.enabled` system property.
- Add the **OAuth Application User** field to the OAuth Entity form.

Create the Client Credentials system property

Create the `glide.oauth.inbound.client.credential.grant_type.enabled` system property to use Client Credentials grant type for OAuth inbound integrations.

Before you begin

Role required: admin

Plugin required: OAuth 2.0.

Procedure

1. In the navigation filter, enter `sys_properties.list`.
   The entire list of properties in the System Properties [sys_properties] table appears.
2. Select **New**.

3. On the form, fill the following fields.

Field	Description
Name	Name of the property you're creating. In this case, `glide.oauth.inbound.client.credential.gran`
Description	Type a brief, descriptive phrase describing the function of the property.
Type	Select the appropriate data type from the list. In this case, `true| false`.
Value	Set the desired value for the property. In this case, `true` to enable the client credentials grant type for OAuth inbound integrations.

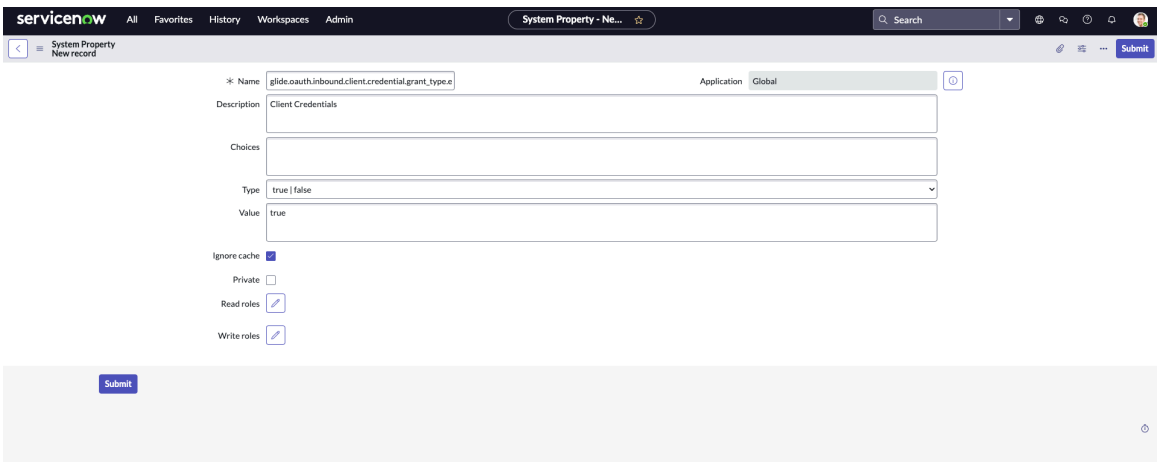

Note: Other fields in the form such as Choices, Ignore cache, Private, Read roles, and Write roles can be configured according to your requirements.

4. Select **Submit**.

Note: If the **Ignore cache** check box is selected, the system flushes the server cache when the parameter is changed.

Next, you must create an OAuth client (OAuth API endpoint for external client) and add OAuth Application User field to the OAuth client record.

Add the OAuth Application User

Add the OAuth Application User field on the OAuth Entity form to use the Client Credentials grant type for OAuth inbound integrations.

Before you begin
Role required: admin

Plugin required: OAuth 2.0.

You must create an OAuth client. For more information, see Create an endpoint for clients to access the instance.

Procedure

1. Open the OAuth client record that was created.
2. Select more options icon on the page header.
3. Select **Configure** > **Form Design**.
4. On the Form Design page, add **OAuth Application User** from the list of fields.
5. Save or Update the form.
6. Select the user for the **OAuth Application User**.
   For example, System Administrator.

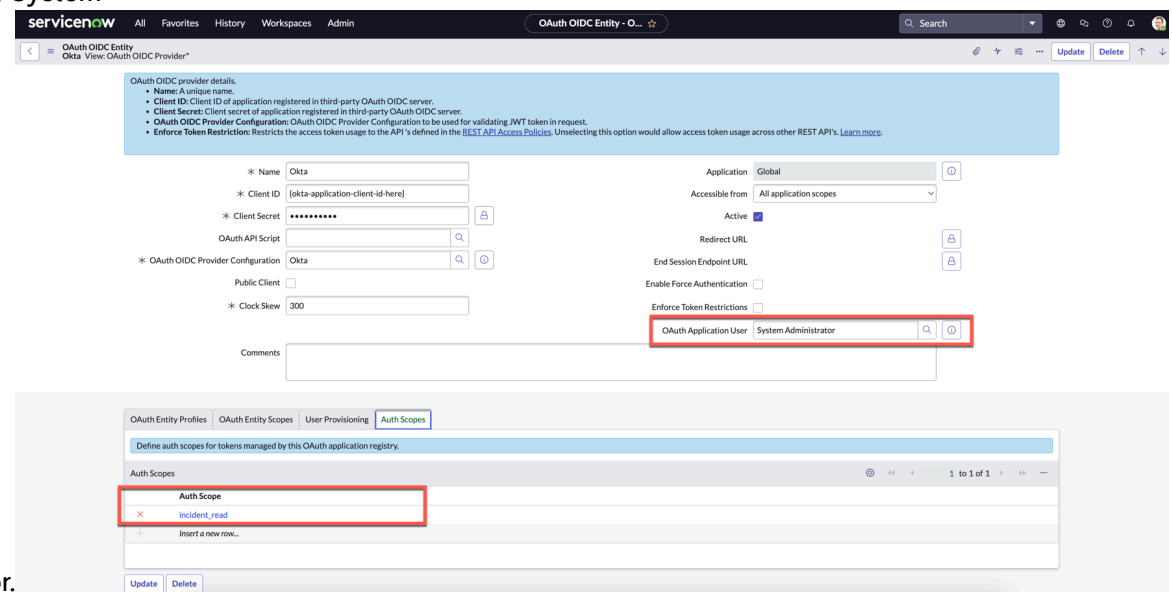

> **Note:** You must use the REST API Auth Scope with client credentials grant type to control the access provided to the 3rd party client.

7. Save or Update the form.

   Any authorization request with the **Grant Type** as **Client Credentials**, **Client ID**, and **Secret** is passed for the associated OAuth Application User in ServiceNow®.

> **Note:** If the OAuth Application User isn't selected on the OAuth client record or the Client Credentials property is set to `false`, then the authorization request isn't passed.

Manage OAuth tokens

Open OAuth tokens to provide access to restricted resources.

Before you begin

Role required: any user or admin

About this task

OAuth tokens issued by the instance and third party OAuth provider are stored in oauth_credential table.

Some of the important columns in this table:

- Token: Value of the token issued by ServiceNow instance.
- Type: Determines if the token is Access Token or Refresh Token.

- Expires: Data/Time when the Access or Refresh Token expire.
- Token Received: Value of the token issued by a 3rd party OAuth Provider. This value is in encrypted format.

Token Expiration and Validity is as follows:

- Access Token: By default, an instance issues access tokens with a 30-minute lifespan in the scenario where the instance is the OAuth provider.
- Refresh Token: By default, an instance issues refresh tokens with a 100-day lifespan in the scenario where the instance is the OAuth provider.

Procedure

1. Navigate to one of the following menu options:
   - **Self-Service** > **My Connected Apps** to see the tokens that the instance created when you granted access to a resource on the instance.
   - **System OAuth** > **Manage Tokens** to see all tokens. Only administrators can access this module.
2. Click the **Name** to open the token.
3. Click **Revoke Access** to prevent access to the restricted resource.

   You can also view other information about the token, including the scope it allows access to and the expiration date.

   You can select the **Clean Expired OAuth Credentials** record from the Schedule page (`sys.trigger.list`) and configure the following:

   - **com.snc.platform.security.oauth.is.active**: By default, the value is set true.
   - **com.snc.platform.security.oauth.hours.expired.credential.is.kept**: Set the value based on your requirement to determine the number of hours you want the keep the expired `oauth` credential in the system.
   - **com.snc.platform.security.oauth.day.old.credential.is.kept**: Set the value based on your requirement to determine the number of days you want the keep the expired `oauth` credential in the system.

Revoke an OAuth token

You might want to revoke an OAuth access or refresh token for security reasons.

Before you begin
Role required: admin

About this task
Revoking the token pertains to the situation where your instance acts as the OAuth resource server. You can revoke the token through a URL.

Procedure

Access your instance using `oauth_revoke_token.do` and append the access or refresh token.

For example: `https://[Your_ServiceNow_Instance]:[port]/oauth_revoke_token.do?token=[access or refresh token]` without the brackets `[ ]`.

Result
This endpoint access does not require authentication. The token in this request is marked as expired.

OAuth outbound

OAuth outbound enables you to pull data from a third-party provider to your instance.

You must have the security_admin role to manage the OAuth integration.

You can configure outbound OAuth 2.0 for the following grant types:

- **Connect to third-party provider**: Use the client ID and secret to send it to the OAuth provider. For more information, see Connect to a third-party OAuth provider.
- **JWT Bearer**: An authorization server validates a JWT token which enables identity and security information to be shared across security domains. For more information, see Set up OAuth provider with JWT Bearer grant type.
- **SAML2 Bearer**: Generates the SAML2 assertion and then exchanges the assertion for the access tokens wit the provider.

  Note: For outbound request to SuccessFactors use the **SAML2 Bearer** as the Default Grant Type. To know more about how to configure **SAML2 Bearer**, refer the example in Set up the v4.x.x.

- **Authorization code**: The code that is granted to the client to obtain an access token, which is then used to obtain access to the resource. If you select this option, then you need an authorization URL (the URL of the authorization server).
- **Resource owner password credentials**: The user name and password of the user that is trying to obtain access to the resource.
- **Client Credentials**: The client ID and client secret, which are both used to get the access token. This method does not provide refresh tokens.
- **MID Server**: The MID Servers facilitate communication and data movement between a single ServiceNow® instance and external applications, data sources, and services. Use the Authorization code, resource owner password credential, SAML bearer, and JWT bearer OAuth grant types of OAuth for outbound integration requests through the MID Server. Personal Auth is also supported through the MID server.

You can configure **OAuth provider scenario** (Outbound): Your instance pulls data from a third-party provider.

Note: You must user authenticate for the first time to fetch the token post which, you don't need to authenticate using a user account before the token expiry.

Connect to a third-party OAuth provider

Configure how the client ID and secret are sent to your OAuth provider.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System OAuth** > **Application Registry** and then click **New**.
2. On the interceptor page, click **Connect to a third-party OAuth provider** and then fill in the form.

Field	Description
Name	Unique name for the third-party OAuth connection.
Client ID	The client ID of application registered in the third-party OAuth server.
Client Secret	The client secret of the application registered in the third-party OAuth server.
OAuth API Script	The script used to customize request and response to the external OAuth provider.
Logo URL	The OAuth application logo URL.
Default Grant type	The default grant type used to establish the token. Choices include: ○ **Authorization code**: The code that is granted to the client to obtain an access token, which is then used to obtain access to the resource. If you select this option, then you need an authorization URL (the URL of the authorization server). ○ **Resource owner password credentials**: The user name and password of the user that is trying to obtain access to the resource. ○ **Client Credentials**: The client ID and client secret, which are both used to get the access token. This method does not provide refresh tokens. ○ **JWT Bearer**: An authorization server validates a JWT token which enables identity and security information to be shared across security domains. ○ **SAML2 Bearer**: Generates the SAML2 assertion and then exchanges the assertion for the access tokens wit the provider. **Note:** For outbound request to SuccessFactors use the **SAML2 Bearer** as the Default Grant Type.
Refresh Token Lifespan	Time, in seconds, that the refresh token is valid. The default time is 8,640,0000 seconds.
Public Client	Enables public clients to require PKCE for an authorization. **Note:** You can use only Authorization Code as the *`Default Grant type`* when PKCE is enabled.
Code challenge method	The code challenge method used in OAuth PCKE workflow. Choices include: ○ S256 [Default] ○ Plain ○ None
Comments	Add any comments regarding the OAuth app.
Application	Application and scope that contain this record.
Accessible from	Make this app accessible from all application scopes or from this scope only.
Active	Select the check box to make the app active.
Authorization URL	The OAuth authorization code endpoint.
Token URL	The OAuth server token endpoint.

Field	Description
Token Revocation URL	The OAuth server token revocation endpoint.
Redirect URL	The OAuth callback endpoint. If blank, the instance auto-generates an entry.
Use mutual authentication	Check the box to use mutual authentication for token request and revocation. This feature requires a mutual authentication profile to be specified.
Send Credentials	The OAuth client populates the client credentials in the request: ○ In Request Body (Form URL-Encoded) ○ Basic Authorization header ○ As Private Key JWT
Client Type	Choose the client type, based on the type of your client. Options: ○ **Iframe Embedded** ○ **Integration as a User** ○ **Integration as a Service** To know more, see Configure client type for OAuth and SSO records.

The system creates a record in the Application Registries [oauth_entity] table with type OAuth Provider.

3. **Optional:** Select OAuth Entity Profiles to validate a system-generated default profile for the new OAuth provider without any scope.
   You can change or add an OAuth provider profile including the name, grant type, and OAuth Scope.
4. **Optional:** Select OAuth Entity Scopes to define all available OAuth scopes for this OAuth provider.
   You can select the scopes when you create or update a profile. Each OAuth scope contains a name and a scope that you must get from the provider's specification, such as a read scope or a write scope. Each scope must be defined separately.
5. **Optional:** Select OAuth Entity Resources to define all OAuth resources this OAuth provider.
6. Select **Submit**.

JWT Bearer

JSON Web Tokens (JWTs) enable the capability to configure server-to-server API interactions between ServiceNow and external API providers without requiring any user intervention.

The JSON Web Token(JWT) bearer grant is a JSON string contains claim values which are evaluated and validated by the JWT Grant Handlers at the Authorization Server end, before issuing an access token.

Using the JWT Bearer grant type, you can configure OAuth 2.0 JWT bearer grant flow for outbound rest message.

Set up OAuth provider with JWT Bearer grant type

JSON Web Tokens (JWTs) enable the capability to configure server-to-server API interactions between ServiceNow and external API providers without requiring any user intervention. This support enables Integration Hub or other automated tasks using JWTs to configure API and Service integrations with different providers.

Before you begin
Role required: admin

About this task
The following tasks show how ServiceNow can be set up to use JWTs for OAuth 2.0 client authentication and authorization grants. ServiceNow is the OAuth client, and you can configure an OAuth provider, such as Box or Docusign.

Procedure
1. Upload Java Key Store certificate
   Attach a JKS certificate to your instance to use to enable the JWT client authentication.
2. Configure a JWT signing key
   Create a JWT signing key to assign to your Java KeyStore (JKS) certificate.
3. Create a JWT provider with a JWT signing key
   Add a JWT provider to your ServiceNow instance.
4. Connect to a third-party OAuth provider
   Create a third-party OAuth provider with a JWT Bearer as the default grant type in the ServiceNow Application Registry.
5. Specify an OAuth profile
   Open the OAuth entity profile of the OAuth provider and assign a JWT provider.

Upload Java Key Store certificate

You can attach a Java KeyStore (JKS) certificate to your instance to use to enable the JWT client authentication.

Before you begin
Role required: admin

Procedure
1. Navigate to **All** > **Multi-Provider SSO** > **x509 Certificate**.
2. Fill in the form as needed.

Option	Description
Name	A unique name for your certificate.
Notify on expiration	Designate whom to notify when the certificate expires.
Warn in days to expire	Send an email notification to your certificate manager before your certificate expires.
Active	Enables the certificate to use for token re quests.
Type	The type of certificate you are uploading.
Expires in days	The amount of days until the certificate ex pires.
Key store password	The password associated with the certificate.

Option	Description
Short description	

3. Click **Submit**.

Configure a JWT signing key

Create a JSON Web Token (JWT) signing key to assign to your Java KeyStore (JKS) certificate,

Before you begin

Role required: admin

Note: If you want to add **X.509 Certificate SHA-1 Thumbprint int (x5t)** to the header as part of the JWT Key, you must configure the form and add the **X.509 Certificate SHA-1 Thumbprint int (x5t)** field.

Procedure

1. Navigate to **All** > **System OAuth** > **JWT Keys**.
2. Fill in the form as needed.

Option	Description
Name	A unique name for your JWT Key signing con figuration.
Signing Keystore	The keystore designated when signing the JWT.
Key ID	The Key ID (kid) helps identify which key is used when multiple keys are used to sign to kens. **Note:** If you configure this field, the Key ID claim is included in the JWT. If you do not configure this field, your JWT will not have a Key ID claim.
Signing Algorithm	The algorithm to use to sign with the JWT key. RSA 256 is the only algorithm available.
Signing Key Password	The password associated with the signing key.
Active	Designate that the JWT key alias is actively referenced from a JWT provider.

3. Click **Submit**.

Create a JWT provider with a JWT signing key

Add a JSON Web Token (JWT) provider to your ServiceNow instance.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System OAuth** > **JWT Provider**.
2. Fill in the form and click **Submit**.

Option	Description
Name	A unique name for your JWT provider configu ration.
Expiry Interval (sec)	The lifespan of the tokens, in seconds, gener ated by the JWT provider.
Signing Configuration	The ServiceNow JWT signing key configura tion to apply.

Generate a JSON Web Token (JWT)

Create a JSON Web Token (JWT) for representing claims securely between two parties on the ServiceNow AI Platform.

The GlideJWT API is a scoped, scriptable API which generates a JWT. There are three arguments necessary before generating the JWT:

- Sys_id of JWT Provider
- JSON serialized header
- JSON serialized payload

There are two JWT API scripts, JWTTokenInternal and JWTTokenRestricted, that you can use when configuring a JWT Provider. The JWTTokenRestricted script enables administrators to configure who can generate a JWT. The JWTTokenInternal script is read-only and enables only logged in users to generate a JWT.

To generate a JWT:

- Create a JWT Key with a shared key (HMAC) or a signing keystore (RSA)
- Associate a JWT provider with the signing configuration referring a JWT key

You can use the API to create your token.

You can use standard and custom claims when configuring a JWT provider. You can pass dynamic header and payload claims as part of the generateJWT API signature.

Sample script to test API:

```
var jwtAPI = new sn_auth.GlideJWTAPI();
var headerJSON = {  "kid": "a1234"  };
var header = JSON.stringify(headerJSON);

var payloadJSON = { "jti": "testjti", "iss": "testiss", "sub":
 "testsub" };
var payload = JSON.stringify(payloadJSON);

var jwtProviderSysId = "7a40dde2d5303300964fb7c8f3c14ab5";
var jwt = jwtAPI.generateJWT(jwtProviderSysId, header, payload);

gs.info("JWT:" + jwt);
```

OAuth client APIs

The OAuth client API provides methods to request and revoke OAuth tokens.

The OAuth client provides these classes:

- GlideOAuthClient: Methods for requesting and revoking the refresh and access tokens.
- GlideOAuthClientRequest: Methods for handling client requests.
- GlideOAuthClientResponse: Methods for handling client responses.
- GlideOAuthToken: Methods for retrieving the access token and information about the access token.

You can also customize the OAuthUtil script include to intercept the request parameters and also parse the responses from external OAuth providers.

When using OAuth classes in a scoped script, use the `sn_auth` namespace identifier.

OAuth parameters for default profile support

The default profile feature requires a set of parameters that you can use with the *`setParameter()`* API to specify the OAuth requestor, a context for the request, and the provider profile.

In the OAuth provider scenario, you must set three parameters that tell the OAuth provider which OAuth profile to use by default. When these three parameters are set, the access token is saved in the instance database. Use the parameters with *`GlideOAuthClientRequest`*.

OAuth parameters for default profile support

Parameter	Description
`oauth_requestor`	The sys_id of the object, which can be a user record or an email account.
`oauth_requestor_context`	Descriptor that provides context for the OAuth requestor. As a good practice, use the name of the table where the `oauth_requestor` object is saved.
`oauth_provider_profile`	The sys_id of the OAuth profile record that is the default (see Specify an OAuth profile).

You do not need to use parameters to set the grant type and scope because the values are configured in the OAuth profile record. If you do not use the parameters, you can use the GlideOAuthClientRequest API methods *`setScope`* and *`setGrantType`*. For additional information, see setScope and setGrantType.

Private Key JWT Support for OAuth 2.0 Client Authentication

Support JWT Support for OAuth 2.0 Client Authentication.

Private Key JWT Client Authentication is an authentication method that can be used by clients to authenticate to the authorization server when using the token endpoint.

In this authentication mechanism, only the clients that have registered a public key and signed a JWT using that key can authenticate.

The JWT must contain REQUIRED claim values and may contain OPTIONAL claim values. To know more about the claim values needed for the JWT for `private_key_jwt` authentication, refer the Client Authentication section in the OpenID Connect core documentation.

Note: The authentication token must be sent as the value of the client_assertion parameter. The value of the#**client_assertion_type** parameter must be `#urn:ietf:params:oauth:client-assertion-type:jwt-bearer`.

Plugins required for OAuth 2.0 Client Authentication using JWT token:

- **OAuth 2.0 (com.snc.platform.security.oauth)**: This plugin is active on new and upgraded instances. If the plugin is not active on your instance, you can activate it.
- **Integration - Multiple Provider Single Sign-On Installer (com.snc.integration.sso.multi.installer)**: For OIDC based single sign-on use case.

You can use the OAuth 2.0 Client Authentication using Private Key JWT for the following:

- OIDC based single sign-on
- Outbound OAuth integrations

Configure Private Key JWT for OIDC based SSO

Configure Private Key JWT for OIDC based SSO integrations.

Before you begin

Role required: admin

You must perform the following tasks before choosing Private Key JWT for OIDC based SSO.

- Upload Java Key Store certificate: Attach a JKS certificate to your instance to use to enable the JWT client authentication.
- Configure a JWT signing key: Create a JWT signing key to assign to your Java KeyStore (JKS) certificate.

  Note: If you want to add **X.509 Certificate SHA-1 Thumbprint int (x5t)** to the header as part of the JWT Key, you must configure the form and add the **X.509 Certificate SHA-1 Thumbprint int (x5t)** field.

- Create a JWT provider with a JWT signing key: Add a JWT provider to your ServiceNow instance.

To include a JWT Key for OIDC based Identity Provider, you must:

- Install the **Integration - Multiple Provider Single Sign-On Installer** (`com.snc.integration.sso.multi.installer`) plugin.
- Enable the properties for **Multiple Provider SSO Properties**. For more information, see Multi-Provider SSO properties, tables, and scripts.
- Create an OIDC Identity Provider. For more information, see Create an OpenID Connect (OIDC) configuration for Single Sign-On (SSO).

Procedure

1. Navigate to **All** > **System OAuth** > **Application Registry**.
2. Select the OIDC Identity Provider that you had created.
3. On top of the form, select **Configure** > **Form Design**.

   Note: You must add **Send Credentials** and **JWT Provider** fields to the form to use the Private Key JWT for OIDC based Identity Provider authentication requests.

4. Choose the **As Private Key JWT** for Send Credentials.
5. Select the **JWT Provider**.

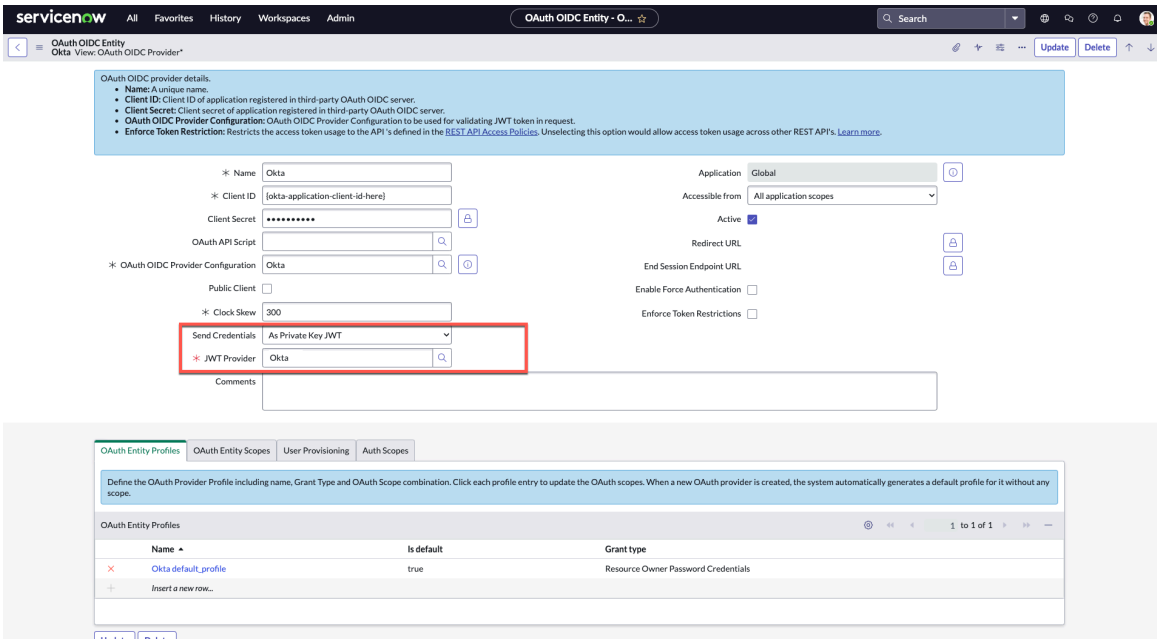

When the user authenticates, the authentication page has options to login through Okta.

Configure Private Key JWT for Outbound OAuth

Configure Private Key JWT for outbound OAuth integrations.

Before you begin

Role required: admin

Before configuring Private Key JWT for outbound OAuth integrations, you must perform the following tasks:

- Upload Java Key Store certificate: Attach a JKS certificate to your instance to use to enable the JWT client authentication.
- Configure a JWT signing key: Create a JWT signing key to assign to your Java KeyStore (JKS) certificate.

> **Note:** If you want to add **X.509 Certificate SHA-1 Thumbprint int (x5t)** to the header as part of the JWT Key, you must configure the form and add the **X.509 Certificate SHA-1 Thumbprint int (x5t)** field.

- Create a JWT provider with a JWT signing key: Add a JWT provider to your ServiceNow instance.

Procedure

1. Navigate to **All** > **System OAuth** > **Application Registry** and then click **New**.
2. On the interceptor page, click **Connect to a third-party OAuth provider** and then fill in the form.

> **Note:** You must add **Send Credentials** and **JWT Provider** fields to the form to use the Private Key JWT for Outbound OAuth authentication requests.

Field	Description
Name	Unique name for the third-party OAuth connection.
Client ID	The client ID of the application registered in the third-party OAuth server.
Client Secret	The client secret of the application registered in the third-party OAuth server.
OAuth API Script	The script used to customize requests and responses to the external OAuth provider.
Logo URL	The OAuth application logo URL.
Default Grant type	Choose: **Client Credentials**: The client ID and client secret, which are both used to get the access token. This method does not provide refresh tokens.
Refresh Token Lifespan	Time, in seconds, that the refresh token is valid. The default time is 8,640,000 seconds.
Public Client	Enables public clients to require PKCE for an authorization. **Note:** You can use only Authorization Code as the *`Default Grant type`* when PKCE is enabled.
Comments	Add any comments regarding the OAuth app.
Application	Application and scope that contain this record.
Accessible from	Make this app accessible from all application scopes or from this scope only.
Active	Select the check box to make the app active.
Authorization URL	The OAuth authorization code endpoint.
Token URL	The OAuth server token endpoint.
Token Revocation URL	The OAuth server token revocation endpoint.
Redirect URL	The OAuth callback endpoint. If blank, the instance auto-generates an entry.
Use mutual authentication	Check the box to use mutual authentication for token request and revocation. This feature requires a mutual authentication profile to be specified.
Send Credentials	Choose: **As Private Key JWT**
JWT Provider	JWT Provider details. You can use the lookup to select the JWT provider.

The system creates a record in the Application Registries [oauth_entity] table with type OAuth Provider that can be used for Private JWT Key authentication.

Create an outbound REST message

Create outbound rest message to authorize instance as authorization server.

Before you begin
Role required: admin

Procedure

1. Navigate to **System Web Services** > **Outbound** > **REST Message** and then click **New**.
2. Fill out the form fields for the OAuth application record as described in .
   Note the following field values:
   - **Endpoint**: URL of the instance that is the authorization server.
   - **Authentication type**: **OAuth 2.0**.
   - **OAuth profile**: OAuth profile that you created for the client server.
3. On the REST message record, click **Get OAuth Token**.
4. Authenticate with the instance that provides the token—the method depends on the single-sign on integration.
   You might use:
   - Your username and password that you use to authenticate to the instance.
   - The username and password for the IdP if is enabled. Click **Use External Login** to access the IdP login screen.

     Note: In order to automatically get redirected to the IdP login page, you must set the glide.authenticate.external system property.

   - Your code, if MFA is enabled.
5. Click **Allow** or **Deny** to complete the authorization and issue the tokens.
   The process that follows is outlined in OAuth authorization code grant flow.

Personal authentication

Personal authentication enables you to securely connect and manage your OAuth-based integrations like Microsoft OneDrive or Google Drive.

OAuth 2.0 Credentials

The OAuth 2.0 Credentials module in ServiceNow (`oauth_2_0_credentials`) enables you to configure and manage access tokens used to connect with external OAuth 2.0-compliant systems. Use the Personal integration type when each user must authenticate with their own identity to an external system.

Use the Credential form to configure OAuth 2.0 credentials with `integration_type = Personal`. For more information, see OAuth 2.0 credentials.

Personal authentication is supported only for the following OAuth 2.0 grant types:

- Authorization Code
- Resource Owner Password Credentials (ROPC)

Grant types such as Client Credentials, and JWT Bearer Grant aren't supported with `integration_type = Personal`.

Note: Personal authentication is also supported via a MID Server for the Authorization Code and ROPC grant types. For more information, see OAuth token fetching via MID.

Personal authentication dashboard

Use your personal credentials to connect to third-party integrations. View, authenticate, revoke, and renew your personal authentications through a simplified, consolidated interface. For more information, see Using the Personal Authentication dashboard.

Note: The personal authentication dashboard can only be accessed by users who are assigned to the role: `sn_personal_auth.personal_auth_user`

Configure Personal Authentication

You can configure personal OAuth authentication with the REST step in Flow Designer.

Before you begin

Role required: admin

Ensure that you install the IntegrationHub Starter Pack Installer (com.glide.hub.integrations) or a higher version.

About this task

This task guides you through configuring personal OAuth authentication for REST steps in ServiceNow® Flow Designer. It enables REST calls to run using the session user's credentials, ensuring secure and personalized API access.

For information on how to check if a personal OAuth token exists for a user, see Get Personal OAuth Token (using GlideOAuthClient).

For information to generate the initial token for a user, who doesn't have access to the credentials page, see Generate Personal Auth Initiator URL.

Procedure

1. Navigate to **Application Registries**, and create an OAuth application registry to connect to an external endpoint.
2. Navigate to **Connection & Credentials Aliases**, and create a connection alias.
   This alias is used in the REST steps. For more information, see Create a Connection & Credential alias.
3. Navigate to **HTTP(s) Connection**, and update the external end-point details for the connection record created in the previous step.
4. Create an **OAuth Credential**.
   a. Navigate to **OAuth 2.0 Credentials**.
   b. Create a new OAuth credential record and link it to the OAuth profile created in Step 1.
5. Add the **IntegrationType** field to the credential form.
6. Update the **IntegrationType** field of the credential created in Step 4 to **Personal**.
7. Generate a personal access token.
   a. As a logged-in user, open the credential record.
   b. Select **Get OAuth Token** to create a personal token.
   c. Select **Manage Tokens** to view and manage the tokens.

   Note: Add a UI action to your application so that the end users can generate tokens. Only administrators can open the credential form directly.
8. Navigate to **Action** to create an action for your use case.

9. Add a REST step to the action.
   Select the connection alias created in Step 2.

   Note: Test the action with the REST outbound call. The credential is marked to be used for personal integration. The REST step might display an error since the action runs with the System integration role.

10. Create a new subflow, and add the action created in Step 8.

    a. In the **Subflow properties** window, select **User who initiates session** in the **Run As** text field. Don't select **System User** instead from the **Run As** text field.

11. Test the subflow.
    The REST step uses the token created for the session user. The subflow can also be invoked using FlowAPIs.

    Sample script to invoke a subflow:

```
try {
    // Execute synchronously in the foreground. Allows
 access to subflow outputs.
    var result = sn_fd.FlowAPI.getRunner()
        .subflow('global.getpersonalincidentssubflow')
        .inForeground()
        .run();

    var outputs = result.getOutputs();
} catch (ex) {
    var message = ex.getMessage();
    gs.error(message);
}
```

12. Manage missing or expired tokens.

    If the session user has no access token, the REST request returns a **HTTP 401 Unauthorized status code** response. Ensure that the token is created before you initiate the flow.

    If the access token is expired but a valid refresh token exists, the system automatically renews the access token.

Get Personal OAuth Token (using GlideOAuthClient)

Check whether the user has a personal OAuth token. Use it to confirm valid access before running REST steps or integrations that require personal OAuth credentials.

Before you begin

Role required: admin

About this task

Use the `GlideOAuthClient` API to check if a personal OAuth token exists for the currently logged-in user. You can also use `ScopedPersonalAuthAPI` to get the personal OAuth token. For more information, see PersonalAuthAPI - Scoped.

Procedure

1. Use the following sample script to check for a personal access token associated with the current session user:

```
function dumpToken(token) {
  if (token) {
    gs.info("Access token: " + token.getAccessToken());
    gs.info("Expires in: " + token.getExpiresIn());
    gs.info("Refresh token: " + token.getRefreshToken());
  }
}

var oAuthClient = new sn_auth.GlideOAuthClient();
oAuthClient.setPersonal(true); // Returns the token for the
 logged-in user

var token = oAuthClient.getToken('<credential_sys_id>',
 '<oauth_profile_sys_id>');
dumpToken(token);
```

2. Replace `<credential_sys_id>` and `<oauth_profile_sys_id>` with the appropriate record values.
   The `setPersonal(true)` method confirms that the token returned belongs to the currently logged-in user.

> **Note:** For more information on the methods for requesting and revoking OAuth refresh and access tokens, see: Glide OAuth Client API Documentation.

Generate Personal Auth Initiator URL

Generate the initial token for a user who doesn't have access to the credentials page to configure personal authentication.

Before you begin
Role required: connection_admin

About this task
Users without the **connection_admin** role can't access the Credentials page to generate OAuth tokens. These users must generate a personal token using the `oauth_initiator` URL with additional parameter indicating that the token is personal and requested for session user.

You can also use the scoped `PersonalAuthAPI` with the `sn_personal_auth` plugin to generate the initiator URL. For more information, see PersonalAuthAPI - getInitiatorURL(String aliasId)

> **Note:** If the personal authentication plugin (`com.snc.sn_ihub_personal_auth`) is activated, use the scoped API to generate the initiator URL. This. API is available only if the plugin is installed.

Procedure

1. Use the following format to construct the token generation URL for Password Grant Type:

```
https://<instance-name>.service-now.com/oauth_password_input.
do?
sysparm_oauth_requestor_context=oauth_2_0_credential&sysparm_oa
uth_requestor=
```

```
<credential sys_id>&sysparm_oauth_provider_profile=<OAUTH profile sys_
id>&sysparm_oauth_personal=true
```

2. Use the following format to construct the token generation URL for Authorization Code Grant Type:

```
https:// ://<instance-name>.service-now.com /
oauth_initiator.do?
oauth_requestor_context=oauth_2_0_credentials&oauth_requestor=
<credential sys_id>&oauth_provider_profile=<OAUTH profile sys_id>&resp
onse_type=code&personal=true
```

Activate Personal Authentication Dashboard

You can activate the Personal Authentication plugin (com.snc.sn_ihub_personal_auth) for Integration Hub if you have the admin role. The application includes demo data and installs related ServiceNow® Store applications and plugins if they are not already installed.

Before you begin

Integration Hub requires a separate subscription from the rest of the ServiceNow AI Platform.

To purchase a subscription, contact your ServiceNow account manager. When you purchase a subscription, certain plugins are activated automatically. If a paid plugin isn't activated automatically, you can manually activate it from the All Applications list in your instance.

Note:

Before purchasing a subscription, you can evaluate this feature on a non-production instance without charge by requesting it from the Now Support Service Catalog.

Role required: admin

About this task

The following items are installed with Integration Hub:

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Personal Authentication plugin (com.snc.sn_ihub_personal_auth) using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Self-register to ServiceNow instance

Use external user self-registration to on-board a large volume of external users to your instance. This feature enhances identity verification to improve customer experiences and supports commonly used registration flows.

Explore  Learn the features and business value of Self-register.	Activate  Understand how to Activate External Self-register.
Configure  Configure the Self-register.	Enable  Know about the properties in Self-register.

Explore Self-register

Use external user self-registration to on-board a large volume of external users to your instance. This feature enhances identity verification to improve customer experiences and supports commonly used registration flows.

External user self-registration enables a large group of users to register to a ServiceNow app without the help of an administrator. For example, a university with a large group of students needing on-campus parking could register themselves to the campus parking app. Each student would be given a limited set of privileges specific to parking, and would complete an automated registration process. This system could generate a parking number, unique to each student, and insure they had a parking space.

A self-registration flow consists of a custom ServiceNow app, in this case, a campus parking app. After configuring an app with necessary tables, an administrator would configure the registration, including a pre-registration flow, field mapping, post registration mapping, captcha, role mapping and post registration flow.

Activate External User Self-Registration

You can activate the External User Self-Registration plugin (com.snc.external_user_self_registration) if you have the admin role.

Before you begin
Role required: admin

About this task

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the plugin using the filter criteria and search bar.

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

3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

External roles in self-registration

To prevent inadvertently providing access to external users, you can assign the snc_external role to all external users.

External users that self-register must be assigned the snc_external role, which has the least privileges. The snc_external role indicates that the user is external to your organization and should not have any access to resources unless explicitly allowed through ACLs for the snc_external role or additional roles that inherit the snc_external role.

By default, users with the snc_external role cannot access:

- Tables without the role that inherits the snc_external role or the public role.
- Non-record type resources, such as processors and UI pages without the snc_external role or a role that inherits the snc_external role.
- Platform Analytics dashboards.

Beginning with the Paris release, you must enable an exclude-list property to enforce the explicit assignment of snc_external roles. For information about enabling the property, see Prevent future internal role assignments for external users.

Configure a user registration configuration for external users

Create a user registration configuration record to bootstrap the onboarding process of external users to custom ServiceNow applications. This form guides the external users through the self-registration process.

Before you begin

- Role required: admin
- Activate External User Self-Registration

Procedure

1. Navigate to **All** > **External User Self-Registration** > **User Registration Configurations** and click **New**.
2. On the form, fill in the fields.

User registration configuration form

Field	Description
Name	Name of the registration form. For example, a student parking registration might be Parking Management Students.
Roles assigned to provisioned users	Roles assigned to the provisioned users. The specified roles have to extend or contains the `snc_external` role. The specified role can also be `snc_external` role. For external users, each role must have an `snc_external` to notate an external user. If you have pre-configured roles, the roles should be accessible when you unlock the roles and search for users.
Enable terms and conditions	Option to add a terms and conditions URL to the registration page.
Terms and conditions URL	The publicly accessible URL which contains the terms and conditions of the registration form. This field appears only when **Enable terms and conditions** is selected.
Description	The description of the registration form. This field appears only when you save or submit the configuration. You can add more information about the registration form in this field. **Note:** You can only add a description after you save or update a user registration configuration.
Application	Application containing this record. The application is automatically set to `Global`.
Active	Option which makes the user registration configuration active. This option is selected by default.
Enable CAPTCHA	Option to add a CAPTCHA to the registration form. The default CAPTCHA provider is Google reCAPTCHA. **Note:** To enable CAPTCHA for user registration, follow the steps in Configure Google reCAPTCHA for external user self-registration.

3. Click **Submit**.
   A user registration configuration with the default settings is created.

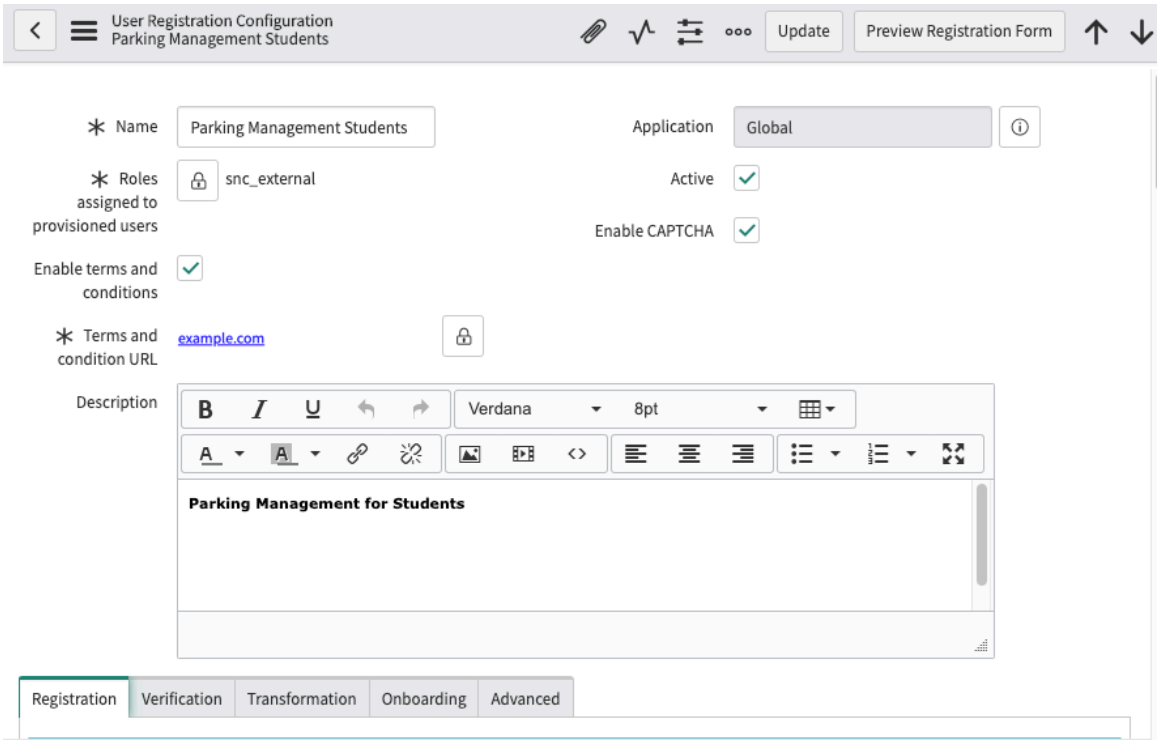

4. **Optional:** Configure the **Registration** tab to show fields and the order that they appear:

Registration Form Fields

Column	Description
Display in Registration Form	Set any field you want to display as `True`.
Order	Set an order number to display fields on your form.
Mandatory	Set any field you want to make it mandatory as `True`.
Validation only field	Set any field you want to use only for validation. For example, registration code.

You can choose to display the default form fields or you can add custom form fields to the Registration Form. For more information, see Default registration form fields.

Registration tab

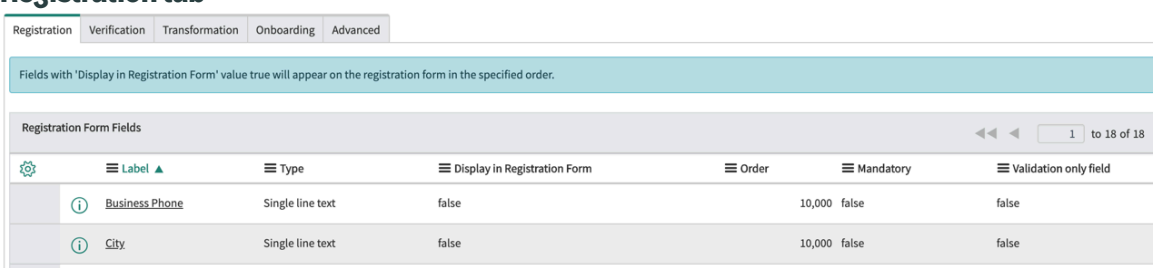

You can also add custom registration form fields. For more information, see Add a custom registration form field.

5. **Optional:** Configure the **Verification** tab to verify the identity of the registered users. When the user verification flow triggers, an activation link is sent to the user's registered email address.

Verification tab fields

Field	Description
Requires user verification	Option to trigger a user verification subflow which runs after user registration. The subflow is for user identity verification.
User verification flow	The subflow used to verify the identity of the user. The subflow triggers only when you enable user verification. (Optional) The **External User Verification** subflow is available by default. You can create a copy of the default subflow in Workflow Studio and modify it according to your requirements. For more information, seeFlow Designer . **Note:** To preview the **External User Verification** subflow in a new tab use the following shortcuts: ○ Macintosh: Command + Click ○ Windows: Control + Click
Activation link expiry time (in hours)	The number of hours after which an activation link expires. The default value is 24.

Verification tab

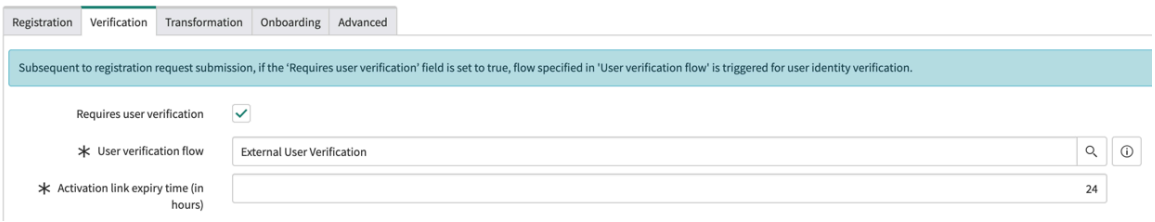

6. **Optional:** Configure the **Transformation** tab to map self-registered users and activated users. There are two transformation maps (u_reg_xmap_[number]) which automatically map the registered users from the User Acti Req [number] table to the Self Registered User [number] table. You can create a copy of these default transformation maps and modify the map according to your requirements. For more information, see Transform maps .

Transformation tab

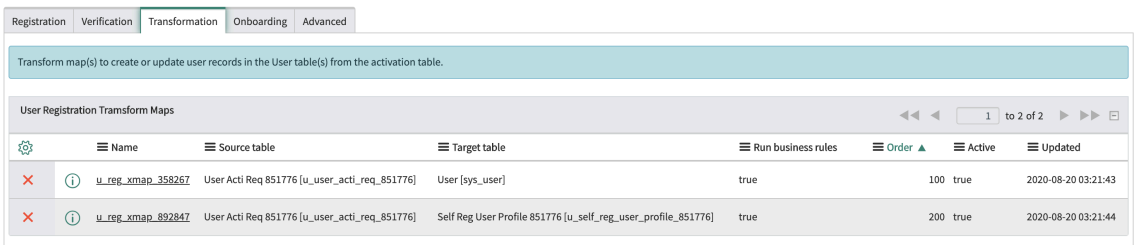

7. Configure the **Onboarding** tab to trigger subflows for onboarding activated users.

The default **External User Onboarding** subflow sends an email to the user that contains a link to reset their password. You can create a copy of the default subflow and modify it according to your requirements.

Note: When the **External User Onboarding** subflow triggers, the subflow sends an email to the user that contains a link to reset the password.

Onboarding tab

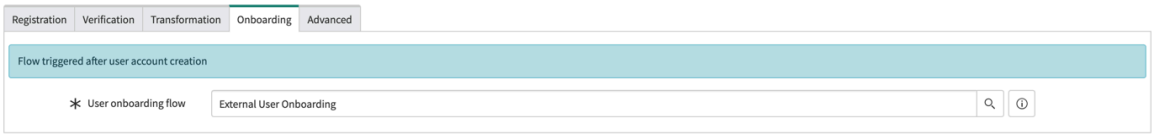

8. **Optional:** Configure the **Advanced** tab to map user tables and redirection pages of the registration form.

Advanced tab

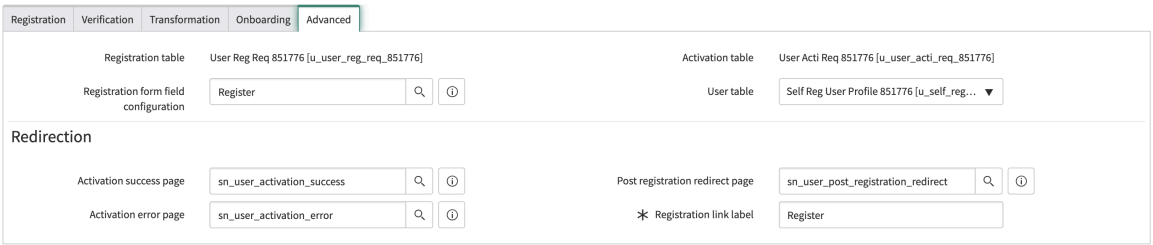

Advanced tab

Field	Description
Registration table	Name of the table where registration form information is saved.
Registration form field configuration	The record associated to the registration form in the Record Producer.
Activation table	Label and name of the table used for user activation. The activation table contains records of the users who have completed verification.
User table	Label and name of the user profile table.
Activation success page	The page to which a user redirects after the activation is successful.
Activation error page	The page to which a user redirects when the activation fails.
Post registration redirect page	The page to which user is redirected after registration.
Registration link label	The button name used for registration from the service portal. The default value is Register.

While making changes or after completing all changes in the user registration configuration, you can use the **Preview Registration Form** button to preview changes in the registration form.

Registration Form Preview

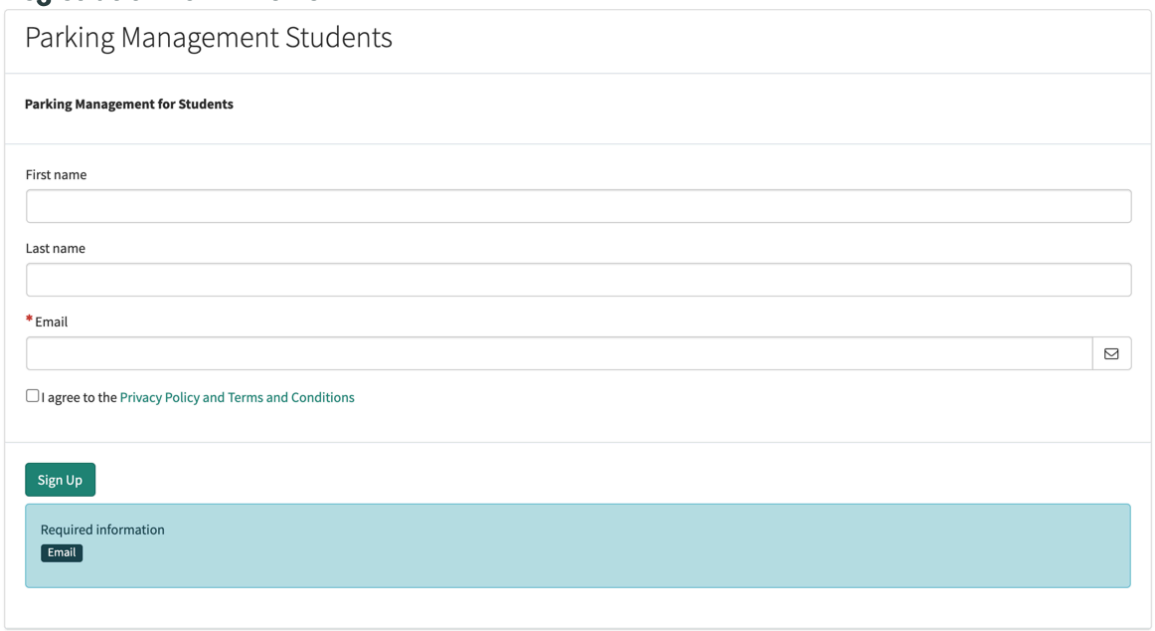

Configure Google reCAPTCHA for external user self-registration

To use the Google reCAPTCHA service, you must request an API key pair from Google and then configure the related system properties.

Before you begin

- Request an API key pair (a site key and a secret) from Google at https://www.google.com/recaptcha.
- Role required: admin

About this task

Procedure

1. Navigate to **All** > **System Properties** > **All Properties**.
2. Search for the following properties and set the values:

Property	Value
glide.user.registration.google.recaptcha.secret	The secret authorizes communication between the application and the reCAPTCHA server. Type: password2
glide.user.registration.google.recaptcha.site_key	The site key used to invoke the reCAPTCHA service on your registration page. Type: string

Property	Value
glide.user.registration.captcha.widget	The Sys_ID of the Captcha Widget. Type: String

Default registration form fields

You can use the default registration form fields or create custom registration form fields.

Registration form fields

The following form fields are added in the Registration form by default.

Field Label	Type
Business Phone	Single line text
City	Single line text
Country	Select Box
EDU Status Used to differentiate staff, students, and visitors. **Note:** This field is useful when the registration form is for an educational institution.	Select Box
Email	Email
First name	Single line text
Gender	Single line text
Home phone	Single line text
Language	Select Box
Last name	Single line text
Middle name	Single line text
Mobile phone	Single line text
Name	Single line text
Prefix Title of the user. For example, Mr. , Dr. , and so on.	Select Box
State / Province	Single line text
Street	Wide Single Line Text
Title - Job title	Select Box
Zip / Postal code	Single line text

Note: You cannot delete the default registration form fields.

Add a custom registration form field

You can add custom fields in the user self-registration form.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **External User Self-Registration** > **User Registration Configurations**.
2. Open the record for the required user registration configuration.
3. Navigate to the end of the Registration Form Fields section and click **Insert a new row...**.
4. Enter the field name under the **Label** column and click check mark.
   A new row is added with the default values. You can configure the custom registration form field based on your requirements.

Registration Form Field Columns

Column	Description
Label	The field name which appears on the registration form.
Type	The type of the user interface element. Following are the supported types: ○ Single line text ○ Email ○ Date ○ Date/time ○ Yes/No ○ Wide Single Line Text ○ Multiple Choice ○ Select Box
Display in Registration Form	Option to display the field in the registration form.
Order	The sequence in which the form fields appear on the registration form. The field with the lowest order value appears first and the field with the highest order value is appears last. The default value is 10,000.
Mandatory	Option to make a field mandatory on the registration form.
Validation only field	Option to use a field only for validation purposes. For example, registration code. When set to true, this field is not saved in the User Table (sys_user).

5. Save or Update the changes.

Enable external user self-registration for Service Portal

Enable external users to register to a ServiceNow app through Service Portal.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **Service Portal** > **Portals**.
2. Open a portal record.
3. On the form, fill in the **External user registration configuration** field.

   Select a user registration configuration.
4. Click **Update**.

Result
The login widget includes a link to the registration form that you previously configured.

Verify user self-registration requests

After a user registers from the Service Portal , a user record is added to the Registration Requests module. You can view the list of registered users who have successfully registered in the Service Portal .

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **External User Self-Registration** > **Registration Requests**.
   The list of user records are displayed.
2. (Optional) Review the individual user records and modify according to your requirements. For example, you can change the status of a user record form **Pending** to **Processed**.

Token based authentication (User logins)

Enhance the security mechanism for users to access a network using token based authentication.

Time limited authentication

Digest Token Authentication

The digest token authentication passes user credentials and a digest token within an unencrypted HTTP header.

Token based authentication is a protocol which allows users to verify their identity, and in return receive a unique access token.

It helps to enhance the security mechanism for users to access a network.

Time limited authentication

Support time limited authentication for your ServiceNow instance.

Explore

Learn the features and business value of Time limited authentication.

Activate

Understand how to Activate Time limited authentication.

Tutorial: Time limited authentication

Know about the properties in Zero Trust Access.

Explore Time limited authentication

Support time limited authentication for your ServiceNow instance.

Note: Time limited authentication is very specific to ServiceNow instance, the customized links for users can only be created within ServiceNow.

Admins can configure link based authentication on the ServiceNow instance. The configured link can be shared with the user through Email or SMS and user can use those links to log in to the instance.

Login to instance with this authentication scheme is controlled through Adaptive Authentication policies configured on the ServiceNow instance.

Time based authentication enables you to perform the following:

- Admin can configure link based auth to use it within an expiry time.
- Application team can fetch nonce to be used with link for the specific configuration item using exposed scriptable API. Link generation will be taken care by Application team.
- Application team can generate the link and send to the user through an existing channel.
- User with the link can log in to instance one time and within expiry time specified as part of configuration.
- Admins can configure the low privileged roles for the authentication scheme.
- Admins can enforce MFA for the authentication scheme as a second factor for authentication with the TLA link.

Activate time limited authentication

Time limited authenication activates through the Integration - Multiple Provider Single Sign-On Installer plugin.

Before you begin

Role required: admin

Note: Time limited authentication is very specific to ServiceNow instance, the customized links for users can only be created within ServiceNow.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Time Limited Authentication (*com.snc.authenticate.time_limited_authentication*) plugin using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Time limited authentication with SMS - Twilio Tutorial

Set up time limited authentication with MFA factors such as SMS using Twilio.

Before you begin

Role required: admin

Plugins required:

- *com.snc.authenticate.time_limited_authentication* (Time Limited Authentication)
- *com.snc.authentication.sms_mfa* (Multi-factor authentication with SMS)

Note: Time limited authentication (TLA) is very specific to ServiceNow instance, the customized links for users can only be created within ServiceNow.

The tutorial instructions provided enable the admin to provide a link-based login with SMS as second factor (MFA) for users with a specific role.

After a successful configuration, the system generates a link, and then shares the link with the user through notification (Email/SMS) channel. Clicking the link, the user is prompted to specify the OTP sent to Email or SMS factor based on user role (configuration).

Note:

- TLA should always be followed by MFA and MFA should be enabled by admin using Adaptive Authentication for TLA login. To know more about how to configure MFA with Adaptive Authentication, see Multi-factor Authentication context.
- TLA should be used for the users who have limited privileges.

Procedure

1. Creating a Twilio configuration.
   a. Create a Twilio test account. For more information, see Twilio.
   b. Navigate to **All** > **Notify** > **Administration** > **Twilio Direct Configuration**.
   c. Provide the **Account SID** and **Auth token** (created from Twilio) and save the record.

   Note: You can create your own provider configuration and use that for TLA. In this example, it is Twilio. To know more about how to create a MFA provider configuration, see Configure MFA Provider.

2. Configuring and enabling Time limited authentication(TLA) record.
   a. Navigate to **All** > **Time Limited Authentication Config records** and click **New**.
   b. On the form, fill the fields.

Time Limited Authentication Properties

Field	Description
Name	Name of the record.
One time use	Enable to use the TLA link once.
Expiry	Specify the seconds for the link expiry. The default is 45 minutes.
Failed Redirect	Enter the URL to redirect users after a failed authentication.
Single Sign-On Script	Details of the SSO script that you want to use.
Active	Option to make the configuration active.
Max login attempts	Specify the number of attempts allowed with the generated TLA link for login. Un-check the **One time use** checkbox to provide the max number of attempts.

Field	Description
External logout redirect	Enter the URL to redirect users after logout.

c. Click **Submit**.

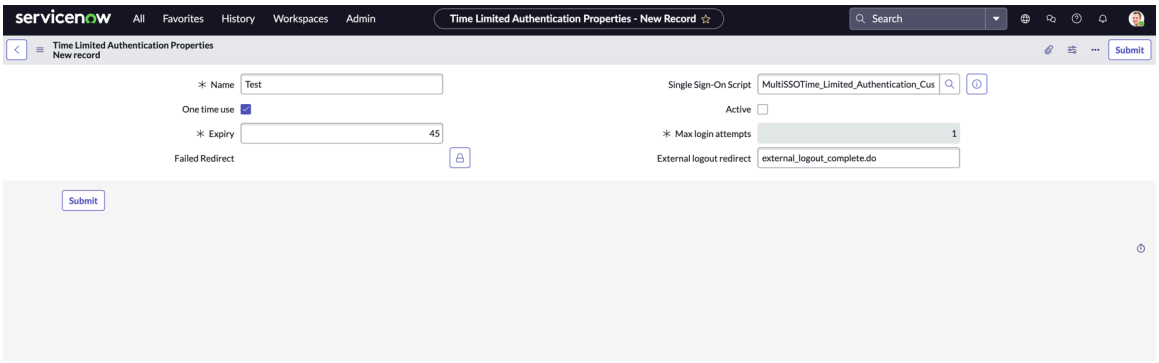

d. Navigate to **All** > **Multi-Provider SSO** > **Administration** > **Properties** and enable the **Enable multiple provider SSO** property and **Save**.

3. Allowing TLA to only a specific user persona using the post-authentication context policy.

a. Navigate to **Roles** and create a role. For example: remote_worker.

b. Create a user with a valid email id and mobile number. To know how to create a user see, Create a user.

c. Assign the role to the user. To know how to assign the role to the user, see Assign a role to a user.

d. To create a role filter criteria, navigate to **All** > **Adaptive Authentication** > **Role Filter Criteria**, create a new filter **remoteworkerrole** and condition **Role is remote_worker**.

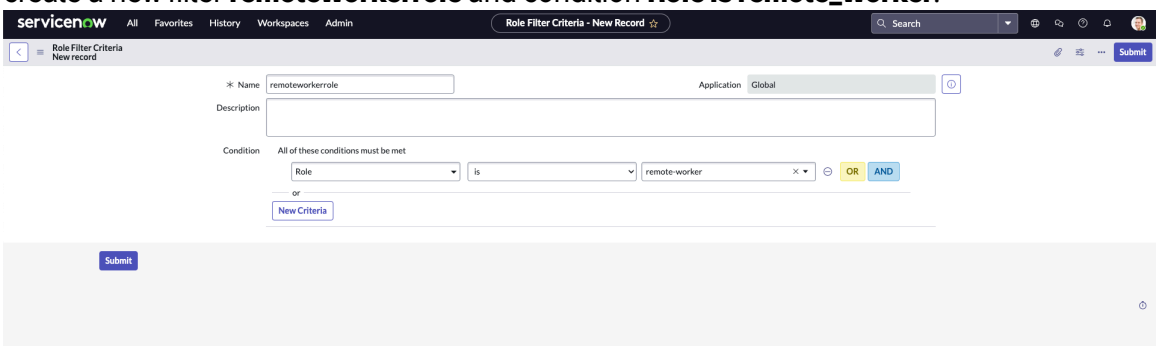

e. To add policy condition based on the deny policy context based on the IdP and role filter criteria, navigate to **All** > **Adaptive Authentication** > **Post Authentication Context**.

f. Click the information icon and **Open the Record**.

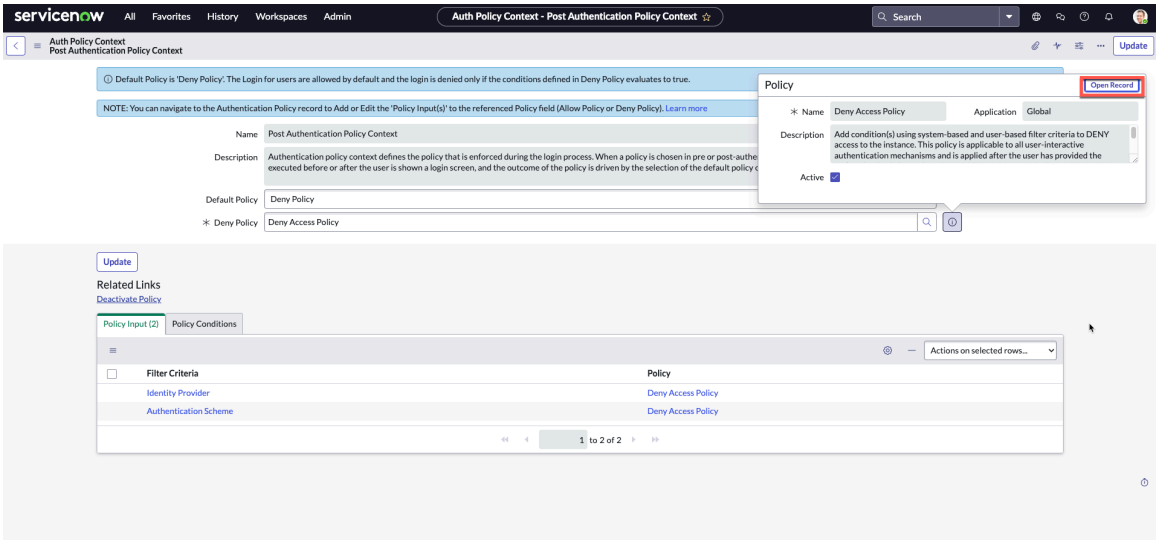

g. In the Policy Input, click **Edit** and add the role (remoteworkerrole) and **Save**.

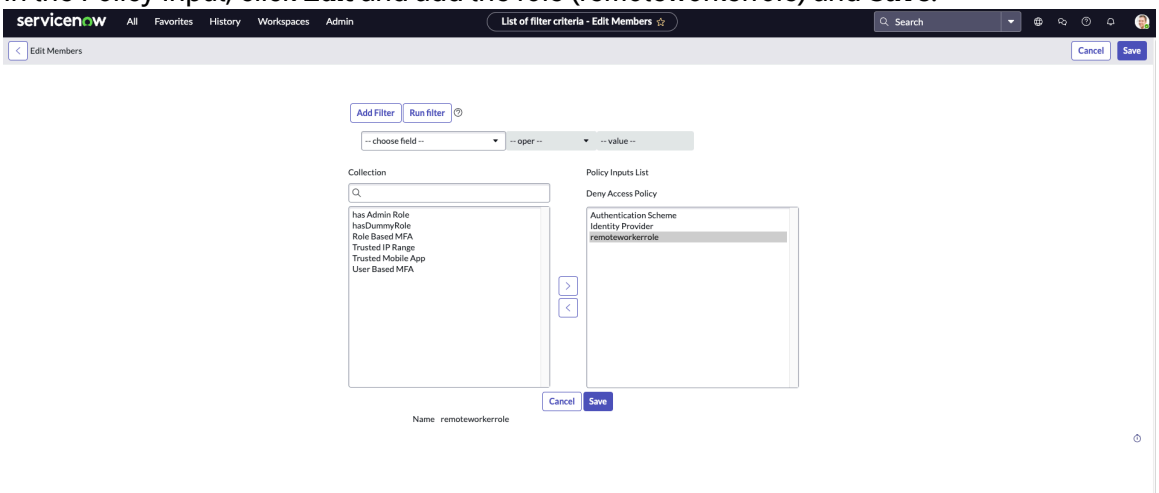

h. In the Policy Condition, add the condition for the policy input and **Submit** the record.

4. Configuring the step-up authentication policy - MFA context.

a. Navigate to **All** > **Multi-factor Criteria**.

b. Select the **Role based multi-factor authentication** and add the role under Multi-factor Roles section and **Update**. In this example: **remote_worker**.

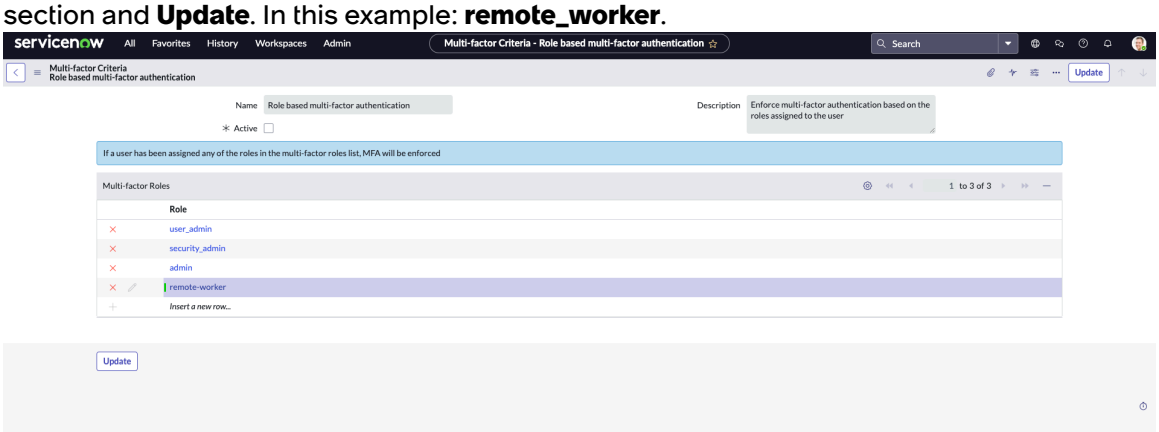

c. Navigate to **All** > **Adaptive Authentication** > **MFA Context**.

d. Make sure the following:

- Default Policy field is **Step-up MFA Policy**
- Step-Up MFA Policy is **Step-up MFA Policy**

e. Click the Information icon and **Open Record**.

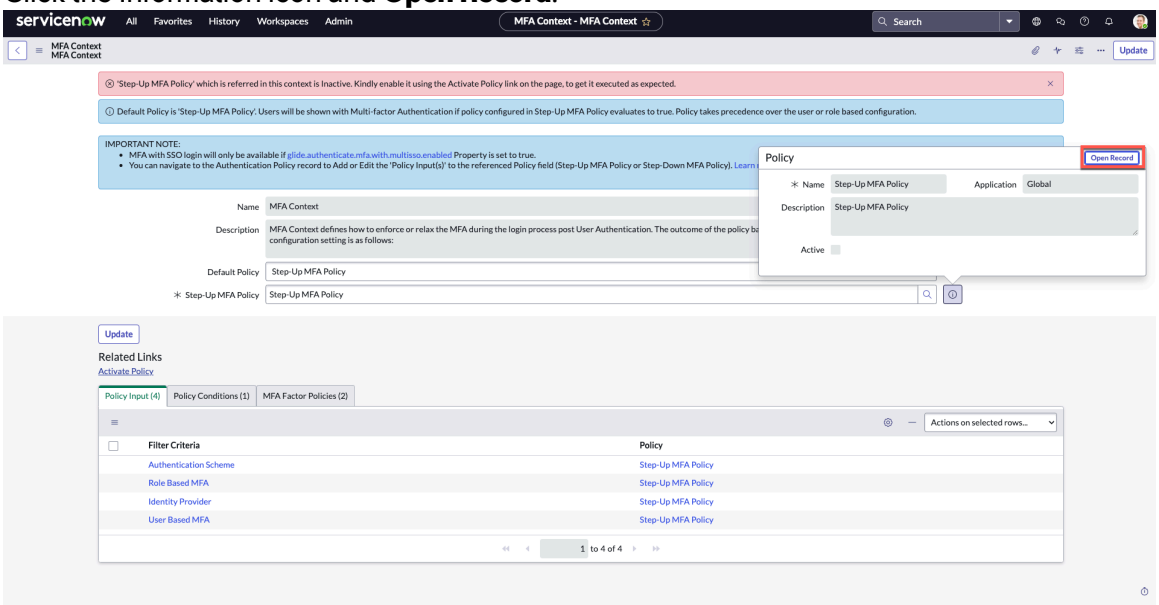

f. On the Step-Up MFA Policy form, in the Policy Inputs, click **Edit**.

g. Add the **Role based multi-factor authentication** to the list and **Save**. In this example, **remoteworkerrole**.

h. In the Policy Condition, click **Enforce MFA if Role based or User based MFA settings are true**.

i. In the Enforce MFA if Role based or User based MFA settings are true page, make sure **Role Based MFA** is **true**.

5. Enforcing MFA to use SMS as a MFA factor policy.

a. Navigate to **All** > **Adaptive Authentication** > **MFA Context**.

b. On MFA Context page, click **MFA Factor Policies** and click on policy **Display SMS OTP as an MFA Factor Policy**.

c. Click **Edit** and add **remoteworkerrole** in the **Policy Inputs**.

d. Click **Policy Conditions** and create a Policy condition.

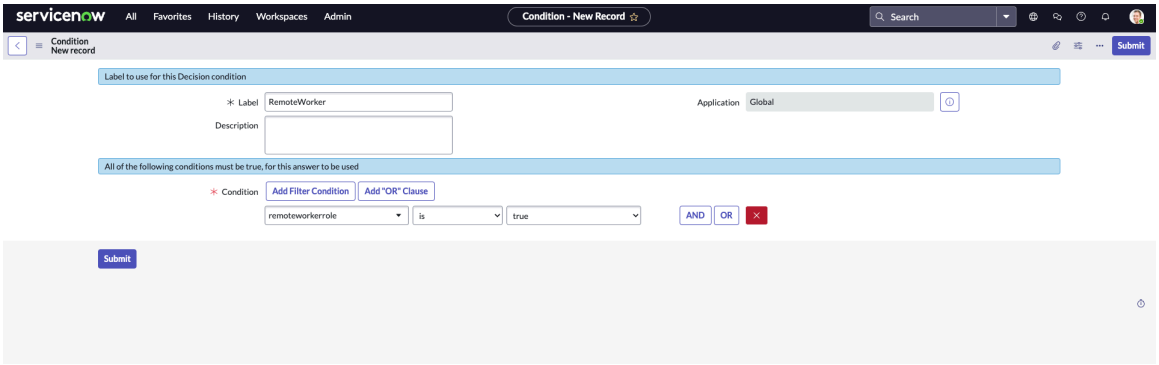

e. Click **Submit**.

The TLA link generated and shared to the users assigned with **remoteworkerrole** as role will be promoted to use the SMS code as a second factor to login the instance.

6. Enabling the other required properties.

a. Navigate to **All** > **Multi-factor Authentication** > **Properties**.

b. Enable the following check-boxes.

- **Enable Multi-factor authentication**
- **Enable Multi-factor Authentication with SSO**

c. Save the record.

d. Navigate to **All** > **Adaptive Authentication** > **Authentication Policies** > **Properties**.

e. Enable the **Enable Authentication Policy** checkbox.

f. Save the record.

7. Generating a TLA link – Example.

a. Navigate to **All** > **System Definition** > **Scripts – Background**.

b. Use the following API by providing user sysid and config id.

```
var tla=new global.TimeLimitedAuthentication();
gs.info(tla.generateNonce("user_sysid",
"config1_sys_id","IAR2"));
```

Note: The source (IAR2) is not a mandatory parameter.

c. Query Parameter is returned as shown:

```
nonce=VCeinfboDtOM&glide_sso_id=b3277f1b44351110f8779b5a2d9909f3&user=3b02
```

d. Create a URL in the below format:

```
https://<instance-url> /login_with_sso.do?uri=<encoded url>&
nonce=2olIQSxdgkjs&glide_sso_id=0c15bf09c3711110c5ec4e483c40dd7a&user=6282
```

8. Click the URL, the following MFA screen is displayed for login.

Digest token authentication

The digest token authentication passes user credentials and a digest token within an unencrypted HTTP header.

Explore	Configure
	
Learn the features and business value of Digest token authentication.	Understand how to configure Digest token authentication.

Explore Digest token authentication

The instance reads the HTTP header value and compares its computed hash value of the digest token. If the computed hash value matches the digest token value, then the instance searches for a matching value in the User table. If there is a matching value in the User table, the instance considers the user pre-authenticated and logs the user in.

Digest token authentication is more secure than simple unencrypted HTTP headers because any accidental or intentional change to the unencrypted HTTP header produces a different hash value. If the hash value fails to match, the instance denies the user access to the requested instance. This prevents users from attempting to login with another user's credentials.

To know more about digest link expiry, see this KB article.

Note: Use Time Limited Authentication (TLA) to configure time based expiry links. To know more, see Time limited authentication.

Integration requirements

A Digest Token Authentication integration requires:

- A web server
- SiteMinder or another single sign-on application to pre-authenticate the user on the local network
- A web page or portal that passes user credentials to the target instance in one of these formats
  - HTTP Header
  - URL parameter
  - Cookie
- A web page or portal that creates and passes a digest token to the target instance using one of these encoding techniques
  - SHA1
  - MD5
  - SHA 256 (recommended)

Configure the digest properties for multi-provider single sign-on (SSO)

After enabling a digest installation exist script, configure properties for multi-provider SSO.

Before you begin

Role required: admin

About this task

If you are not using multi-provider single sign-on, configure standard single sign-on properties.

Procedure

1. Navigate to **All** > **Multi-Provider SSO** > **Identity Providers**.
2. Fill in the fields of Digest Properties form.

Option	Description
Name	Enter the name of the digest token.
User	Enter the sys_user field that contains the matching data for the incoming header.
HTTP Digest header name	Enter the HTTP header you generated. For ex ample, DE_USER.
HTTP header name	Enter the HTTP header you generated for your created digested token. For example, SM_USER.
Secret Passphrase	Enter the secret key to use for encoding digest keys. For example, `32 or more characters`.
Failed SSO Redirect field	Enter the URL to redirect users after a failed authentication.
External logout redirect	Enter the URL to redirect users after a logout.
Single Sign-on Script	Select **MultiSSO_DigestedToken**.
Client Type	Choose the client type, based on the type of your client. Options:**Iframe Embedded**. **Note:** If client type field is required for your configuration, you can edit the form and add the field. To know more, see Configure client type for OAuth and SSO records.

3. Click **Update**.
4. Set your Digested Token default to true.
   When you set the default to true, this overwrites the system default digest token record associated to SSO. Once the first multi-provider SSO related IdP record activates, only records associated to multi-provider SSO will be used.

   Digest token records which exist in digest properties table can be individually called by appending the Sys_ID of the IdP. For example, a

digest token record in the following authentication URL: `https://<instance_name>.service-now.com/login_with_sso.do?glide_sso_id=<sys_id_of_Digest_token_record>&SM_USER=<user_name>&DE_USER=<d`

Sample digest token implementations

Here are several samples of creating a digest token.

Sample digest authentication implementations

Digest built with	Secret key value	Hash method	Example
Java	32 and more characters	SHA256	Sample Java Digest Algorithm for Encryption
C	Value of sharedKey parameter	Value of strEncryptionMethod parameter (SHA256 or MD5)	Sample C

Sample Java digest algorithm for encryption

This Java algorithm illustrates creating a digest token from an HTTP header.

This sample assumes:

- The web server supports Java
- The hash computation method is SHA1
- The secret key value is abc123
- The unencrypted HTTP header name is user_name

Change the Java code to use another hash computation mechanism (such as MD5), change the secret key value, or HTTP header name.

```
public class DigestTest
{
    private static final String MAC_ALG = "HmacSHA256";
       fKey  = {32 or more characters};
    public byte[] getDigest(String acct) {
        try {
            byte[] bkey = fKey.getBytes();
            byte[] data = acct.getBytes();
            Mac mac = null;
            try {
                mac = Mac.getInstance(MAC_ALG);
                mac.init(new SecretKeySpec(bkey, MAC_ALG));
            } catch (Exception e) {
                e.printStackTrace();
            }
            byte[] sig = mac.doFinal(data);
            String signature = new String(sig);
            System.out.println("value:" + acct);
            System.out.println("digested value:  " + signature);
            return sig;
        } catch (IllegalStateException e) {
            e.printStackTrace();
        }
        return null;
```

```
    }
    public static void main(String[] args) {
        BASE64Encoder encoder = new BASE64Encoder();
        DigestTest test = new DigestTest();
        String userName = "user_name";
        System.out.println("base 64 digest username: " +
 encoder.encode(test.getDigest(userName)));
    }
}
```

Sample C

This C class illustrates creating a digest token from three input parameters.

- strEncryptionMethod – lists the hash computation method (SHA1, SHA256 or MD5)
- message – lists the value to be converted into a digest token
- sharedKey – lists the secret key

Note: Use stronger secure hash algorithm like SHA256 for digest token generation.

This sample assumes:

- The web server supports C
- Other code calls this class and passes the expected parameters

Sample Code

```
 private stringdigestData(string strEncryptionMethod , string
message , string sharedKey ) {
          UnicodeEncoding myUnicodeEncoding  = new
UnicodeEncoding ( ) ;

          byte [ ] messageBytes  = System. Text. Encoding.
ASCII. GetBytes (message ) ;
          byte [ ] sharedKeyBytes  = System. Text. Encoding.
ASCII. GetBytes (sharedKey ) ;
          byte [ ] hashedMessage ;

          string b64SHA1Message ;

           if (this. DEBUG ) {
              TextBoxMessage. Text = message ;
              TextBoxSecret. Text = sharedKey ; }

           switch ( (strEncryptionMethod ) )

           { case "SHA1" :

                 HMACSHA1 hmacsha1  = new HMACSHA1 ( ) ;
                 hmacsha1. Key = sharedKeyBytes ;
                 hashedMessage  = hmacsha1. ComputeHash
(messageBytes ) ;
                 b64SHA1Message  = Convert. ToBase64String
(hashedMessage ) ; if (this. DEBUG ) TextBoxDigest. Text =
Convert. ToString (hashedMessage ) ; break ;
```

```
                case "MD5" :

                   HMACMD5 hmacmd5  = new HMACMD5
(sharedKeyBytes ) ;
                   hashedMessage  = hmacmd5. ComputeHash
(messageBytes ) ;
                   b64SHA1Message  = Convert. ToBase64String
(hashedMessage ) ; if (this. DEBUG ) TextBoxDigest. Text =
Convert. ToString (hashedMessage ) ; break ;

                default :

                   b64SHA1Message  = "Unknown Encryption
Method" ; break ;

             }


            TextBoxBase64. Text = b64SHA1Message ; return
b64SHA1Message ;

         }
```

Web service security

Enforce security using basic authentication, mutual authentication, or WS-Security.

Explore web security	Configure mutual authentication
Learn the features and business values of web service security.	Use mutual authentication to establish trust by exchanging secure sockets layer (SSL) certificates.

Explore Web service security

Enforce security using basic authentication, mutual authentication, or WS-Security.

Basic Authentication

To enforce basic authentication on each request for a WSDL document or posting of SOAP messages, you may set the property *glide.basicauth.required* to `true`. If you do so, each WSDL or SOAP request would have to contain the "Authorization" header as specified in the Basic Authentication protocol. Because the request is non-interactive, the **Authorization** header is always required during a request.

Supplying basic authentication information whether or not it is required has the added advantage that the data created or updated as a result of the Web Service invocation is done on behalf of the user supplied in the basic authentication credentials. As an example, when creating

an Incident record, the journal fields have the user id of the basic authenticated user, instead of the default **Guest** user.

To make the authorization header ignore the capitalization rules, use the *glide.security.script.include.name.case.insensitive.list* property. You can modify this property in the System Properties [sys_properties] table and add the script includes that are necessary to process the authentication. By default, this property has these values:

- BasicAuth
- CustomAuth

Add other script includes as needed.

To supply basic authentication when using Perl and the SOAP::Lite libraries, you can implement the following function:

```
sub SOAP :: Transport :: HTTP :: Client :: get_basic_credentials
 { return 'user_name' => 'password' ; }
```

- When using C# .NET VS 2005 or older, you can take advantage of the Credentials object, for example:

```
System.Net . ICredentials cred  = new System.Net .
 NetworkCredential ( "user_name",  "password" ) ;

service . ServiceNow proxy  = new service . ServiceNow ( ) ;
service . get getService  = newservice . get ( ) ;
service . getResponse getServiceResponse  = new service .
 getResponse ( ) ;

 try {
  proxy . Credentials = cred ;
  getService . sys_id = "bf522c350a0a140701972dbf876f1610" ;
  getServiceResponse  = proxy . get (getService ) ; catch
 (Exception ex ) { }
```

- When using C# .NET VS 2008, you can take advantage of the ClientCredentials object, for example:

```
Demo_Incident. ServiceNowSoapClient client  = new
 Test08WebService . Demo_Incident . ServiceNowSoapClient ( ) ;
client . ClientCredentials . UserName . UserName = "admin" ;
client . ClientCredentials . UserName . Password = "admin" ;
```

  Then in your app.config file look for the following and change `None` to `Basic`:

```
<transport clientCredentialType= "None" proxyCredentialType=
 "None" realm= "" />
```

- When using VB .NET taking advantage of the Credentials object would look like the following:

```
Sub Main()
         Dim cred  As New
 System.Net.NetworkCredential( "user_name",  "password")

         Dim proxy  As New VB_Democm.incident.ServiceNow
         Dim getIncident  As New VB_Democm.incident. get Dim
 getResponse  As New VB_Democm.incident.getResponse
```

```
        proxy.Credentials = cred

        getIncident.sys_id =  "[your sysID here]"

        getResponse = proxy. get(getIncident)

    End Sub
```

The resulting response when Basic Authentication is turned on and no credentials are supplied looks like this:

```
<html> <head > <title >Apache Tomcat/5.0.28
 - Error report </ title > <style > <!--H1
 {font-family:Tahoma,Arial,sans-serif;color:white;background-co
lor:#525D76;font-size:22px;}    H2
 {font-family:Tahoma,Arial,sans-serif;color:white;background-co
lor:#525D76;font-size:16px;}    H3
 {font-family:Tahoma,Arial,sans-serif;color:white;background-co
lor:#525D76;font-size:14px;}    BODY
 {font-family:Tahoma,Arial,sans-serif;color:black;background-co
lor:white;}    B
 {font-family:Tahoma,Arial,sans-serif;color:white;background-co
lor:#525D76;}    P
 {font-family:Tahoma,Arial,sans-serif;background:white;color:bl
ack;font-size:12px;}   A {color : black;}   A.name
 {color : black;}   HR {color : #525D76;}--> </
 style > </ head > <body > <h1 >HTTP Status 401 -\ </ h1 >
 <HR size = "1" noshade = "noshade" > <p >< b >type </ b >
 Status report </ p > <p >< b >message </ b > <u >< / u >< /
 p > <p >< b >description </ b > <u >This request requires
 HTTP authentication (). </ u >< / p > <HR size = "1" noshade
 = "noshade" > <h3 >Apache Tomcat/5.0.28 </ h3 > </ body > </
 html >
```

Configure mutual authentication

Mutual authentication establishes trust by exchanging secure sockets layer (SSL) certificates.

During the SSL handshake, the server presents its certificate to the client. Subsequently, depending on the server's configuration, it may request a certificate from the client. Both server and client perform certificate validation procedures to ensure the authenticity and integrity of the presented certificates.

Following a successful validation, acknowledgments are exchanged before the initiating the HTTPS connection.

Note: For information about using a custom HTTPS protocol profile to enable mutual authentication, see Create a protocol profile.

Administrators do the preliminary work of setting up the client Key Store and generating certificates before certification requests are fulfilled.

Warning: This feature only enables mutual authentication on outbound https connections.

Creating the Key Store

The instance currently supports uploading a Java Key Store file, which contain a private key and a public certificate pair, the public certificate includes complete chain including the root certificate.

To set up your client Key Store:

- You require a certificate signed by a trusted certificate authority (CA).
- Your API endpoint provider may assist you in generating the Key Store.

If you need to generate the Key Store, the process involves several steps using command-line interface commands to generate a new Key Store file, create a certificate signing request (CSR), and import signed certificates. Before importing your domain's primary certificate, any root or intermediate certificates should be imported first. Here's a step-by-step instruction:

1. Generate a Java keystore and key pair.

```
keytool -genkey -alias mydomain -keyalg RSA -keystore
 my.keystore
```

2. Generate a CSR for an existing Java keystore.

```
keytool -certreq -alias mydomain -keystore my.keystore -file
 mydomain.csr
```

3. Send the CSR to your CA signing authority to sign and return the certificate files, which includes intermediate and root certificates along with signed certificate.
4. Import a root or intermediate certificate authority CA certificate to an existing Java keystore.

```
keytool -import -trustcacerts -alias root -file Thawte.crt
 -keystore my.keystore
```

Note: You can bundle all the certificates in one file and import. This is preferable option. If you do this way you can skip 5

```
keytool -import -alias mydomain -file merged.crt -keystore
 my.keystore
```

5. Import a signed primary certificate to an existing Java keystore.

```
keytool -import -trustcacerts -alias mydomain -file
 mydomain.crt -keystore my.keystore
```

Setting up the Key Store

Now that the key store has been created, it can be uploaded to the Certificates table. On the **System Definition** > **Certificates** page, click **New** and set the following fields:

- Enter a certificate **Name**.
- Store the key store as **Active**.
- Set **Type = Java Key Store**.
- Provide a **Key store password**. This is the password that was used to create the keystore.

Click **Submit** to create the Java Key Store entry.

Key Store

X.509 Certificate | = Required field — Update | Delete

Attachments: my.keystore [view]

Name:	Key store	Type:	Java Key Store
Active:	☑	Key store password:	••••••••••••••••••••••
Short description:	key store used for mutual authentication		

Specifying a Trusted Server Certificate

During an outbound SSL connection, which is an HTTPS Web Service call, it is possible to specify a certificate provided by the service provider that ensures the validity of the service provider during the SSL connection. For example, a browser attempting to connect to a secure service which identifies itself by a certificate.

By uploading the trusted server certificate, ServiceNow ensures that the service it is connecting to is valid and secure.

Create a new Certificate entry with the type of "Trust Store Cert" and attach a DER formatted certificate, or copy and paste its PEM format into the **PEM Certificate** field.

Protocol Profile

Certificate Exchange

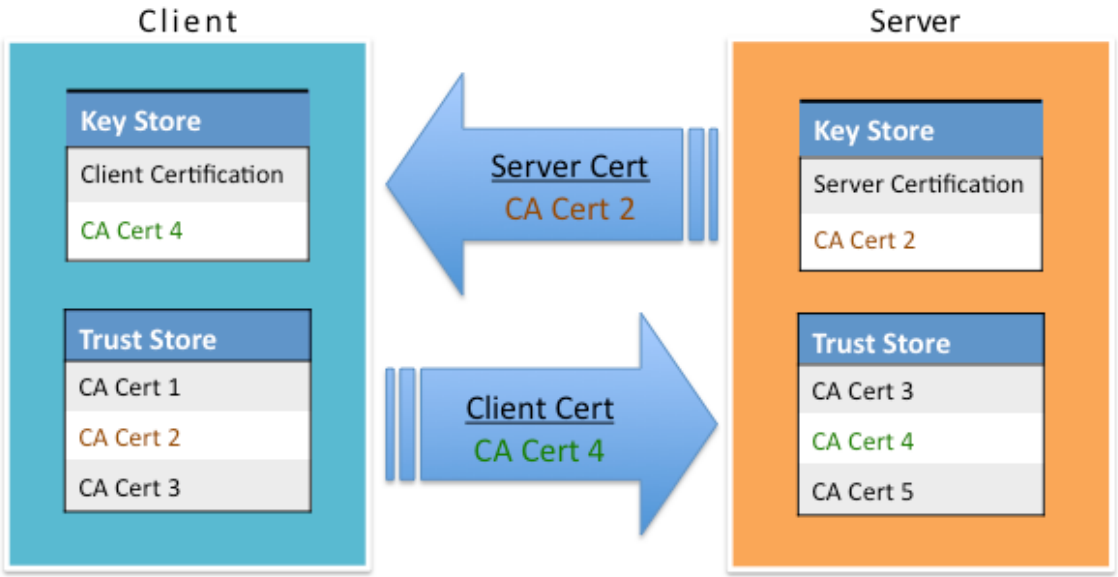

- When a client requests the server certificate for authentication, a certificate signing request (CSR) is generated.
- To respond to a CSR, the server generates two unique cryptographic keys: A public key, which is used to encrypt messages to the server and a private key, which is used to decrypt messages. Both keys are kept in the Key Store.
- Keys are used to decrypt the client secure messages so they can be read by the server. Any outgoing connection that is going to be HTTPS verifies the certification by checking the Key

Store, offering its public certification, and uses the trust store certificates to verify mutual trust back.

- To complete the secure link between the client and the server, the server matches the certificate to the corresponding private key. Because only the server has access to the private key, the server can decrypt the data from the client.

Note: For information about using a custom HTTPS protocol profile to enable mutual authentication, see Create a protocol profile.

The server responds by sending a certificate. Is this a certificate that the client accepts? If yes, a message is sent to the server accepting the certificate and a secure channel is initiated. If the certificate is not accepted, it may mean that the root authority is needed for certification.

Access Control List Rules

Rules for access control lists (ACLs) restrict access to data by requiring users to pass a set of requirements before they can interact with it.

Explore ACLs  Learn about Access Control Lists(ACLs).	Configure ACLs  Configure ACLs.
Contextual Security Manager  Learn about the about Contextual Security Manager.	Advanced ACL Configuration  Learn about Advanced ACL configurations and tools.

Explore Access Control Lists

Explore Access Control Lists (ACLs).

All access control list rules specify:

- The decision type, rule type and operation which defines the ACL
- The object being secured
- The conditions required to access the object

Components of ACL

The decision type defines whether users are allowed to access the object if conditions are met or denies access the object unless conditions are met.

Decision type	Description
Deny-Unless	Restrict access to resource by explicitly denying access unless conditions are passed. See #unique_1501_Connect_42_section_qnd_snl_zbc for more information.
Allow-If	Allow access to resource if conditions are passed.

The object is the target to which access needs to be controlled. Each object consists of a type and name that uniquely identifies a particular table, field, or record. With the Applies-to field users have granular control over which specific records this ACL will apply to.

For example, all these entries specify an object:

Type	Name	Object secured
record	[incident].[--None--]	The Incident table.
record	[incident].[active]	The Active field in the Incident table.
record	[incident] Applies-To: Priority = P1	Only Priority 1 incidents in the incident table.
REST_Endpoint	user_role_inheritance	The record for the user_role_inheritance Scripted REST API.

Each operation describes a valid action the system can take on the specified object. Some objects, such as records, support multiple operations, while other objects, such as a REST_Endpoint, only support one operation.

For example, all these entries specify an operation:

Type	Name	Operation	Operation secured
record	[incident].[-- None --]	create	Creating records in the Incident table.
record	[incident].[active]	write	Updating the Active field in the Incident table.
REST_Endpoint	user_role_inheritance	execute	Running the user_role_inheritance scripted REST API.

The conditions specify when someone can access the named object and operation. Security administrators can specify condition requirements by adding:

- One or more user roles to the **Requires role** list.
- One or more security attributes need to be evaluated to be true.

- One or more data conditions.
- A script that evaluates to true or false or sets the `answer` variable to true or false.

To gain access to an object and operation, a user must pass all conditions listed in an access control. For example, this access control restricts access to view operations on the incident table.

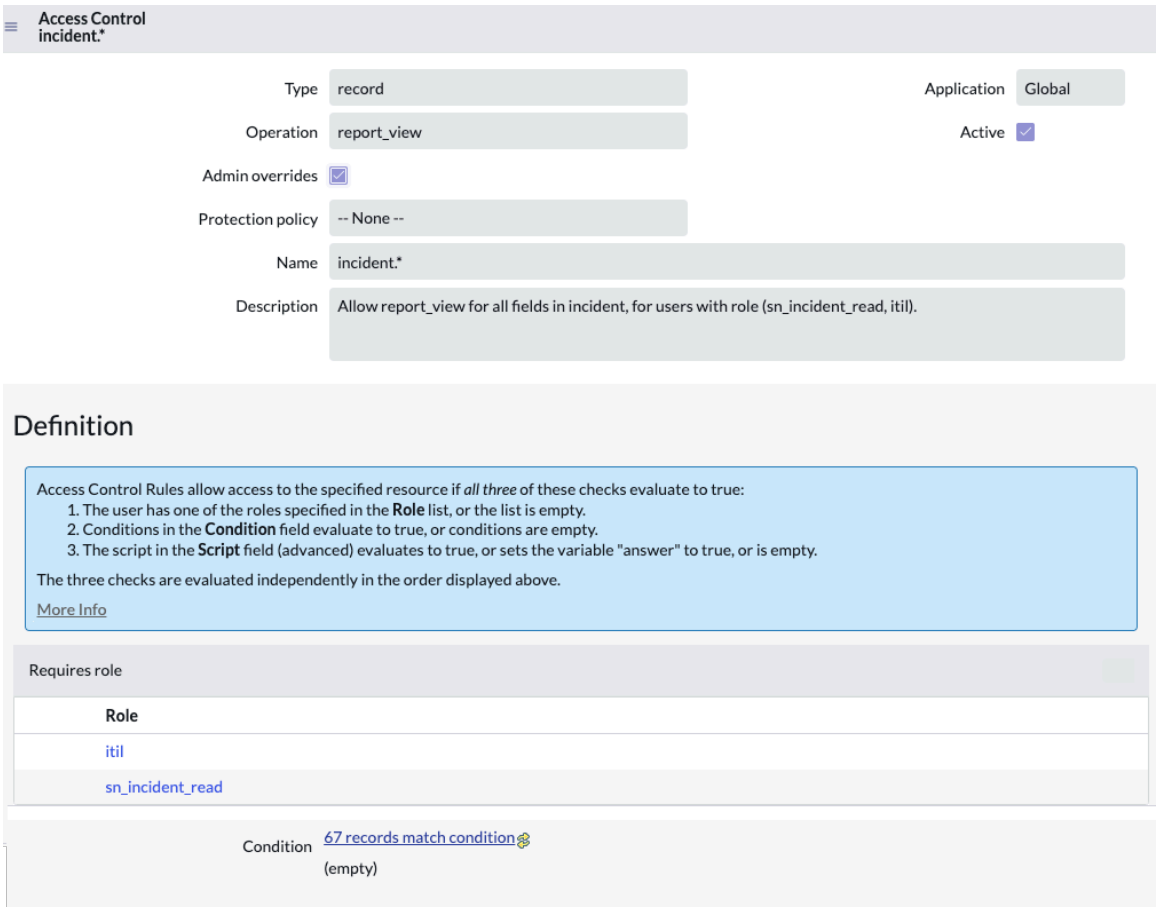

To update a record in the incident table, a user must have the listed roles and the record must meet the condition.

Condition type	Requirement	Description
Requires role	**Requires role**:itil	Only allow users with the itil role to update incidents.
Security Attributes	UserIsAuthenticated	Only authenticated users to update incidents.
Data Condition	[Incident state] [is not] [Closed]	Only allow updates to active incident records.

Applies-to behavior

The Applies-to field determines whether an ACL applies to records, whereas data condition evaluates an ACL that's already applied. Applies-to specifies whether the ACL affects a specific record; if it's empty, the ACL applies to all records. Applies-to can be used for granular ACL enforcement whereas "Data Condition" is an evaluation criteria.

Note: Applies-to is case sensitive.

Deny by default behavior

By default the ACL engine completely denies access if an ACL is empty or invalid. Empty ACLs are defined as ACLs without one or more of these components:

- Defined role
- Security attribute
- Data condition
- Script

Invalid ACLs are defined as:

- ACLs with roles that do not exist (e.g. have no row in the database)
- ACLs with Security Attributes that do not exist (e.g. have no row in the database)
- ACLs with a script that contains "answer=true" or "true"

If the system detects the user creating an ACL it will prompt the user to select a role or an existing security attribute.

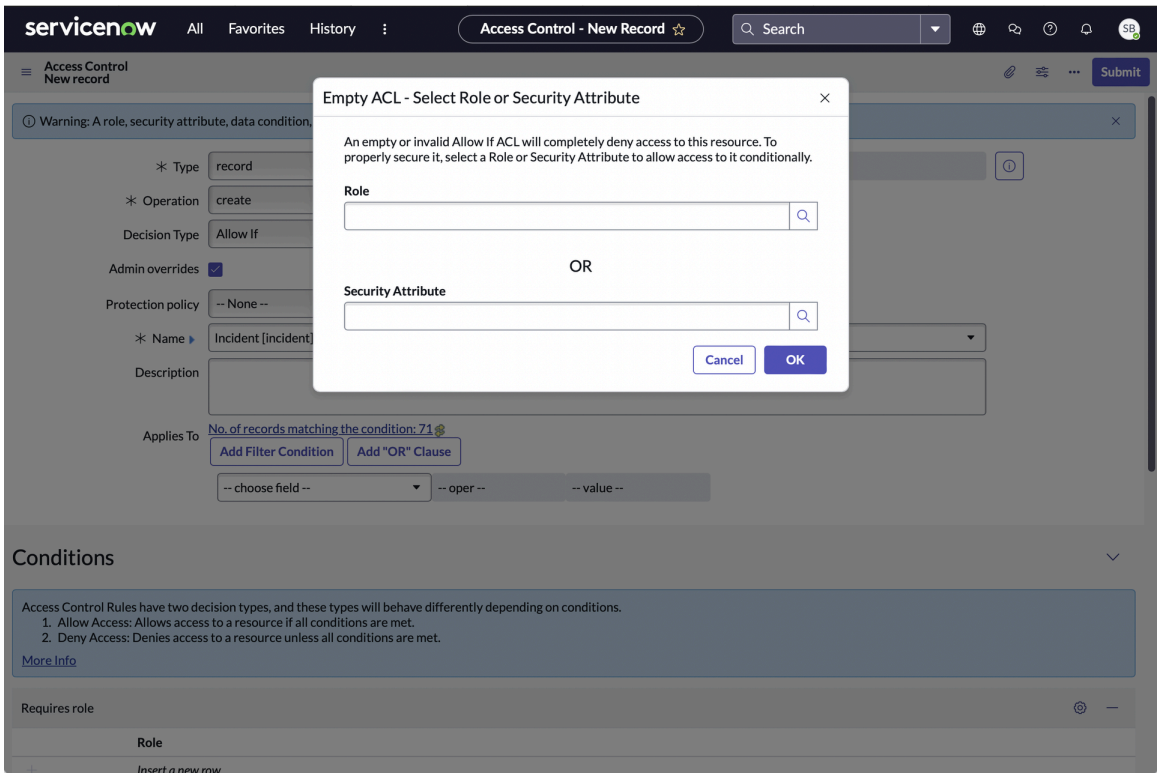

ACL evaluation process

An ACL rule only grants a user access to an object if the user meets all conditions required by the matching ACL rule.

- The condition must evaluate to *true*.
- The script must evaluate to *true* or return an answer variable with the value of *true*.

- The user must have one of the roles in the required roles list. If the list is empty, this condition evaluates to *true*.
- [Record ACL rules only] The matching table-level and field-level ACL rules must both evaluate to *true*.

ACL evaluate conditions

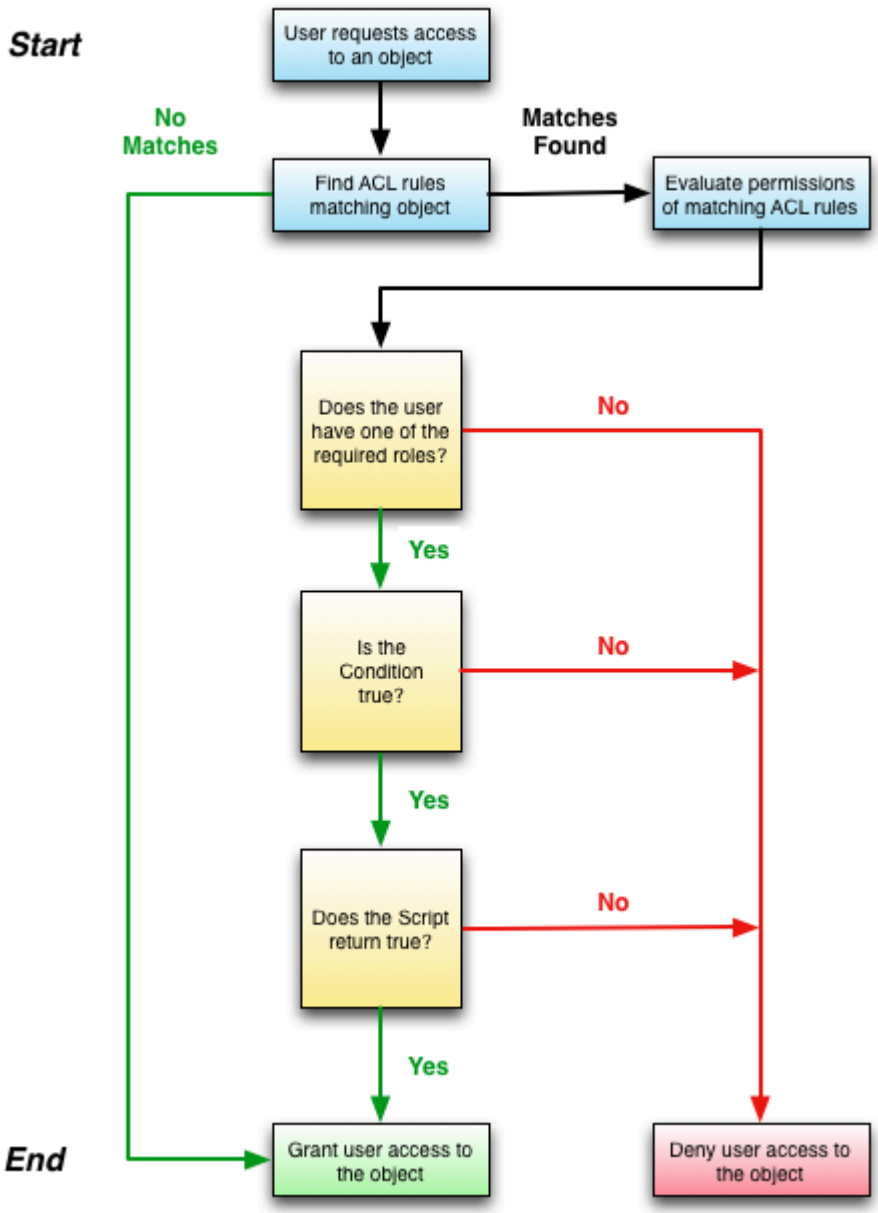

Whenever a session requests data, the system searches for access control rules that match the requested object and operation. If there's a matching access control rule, then the system evaluates if the user has the conditions required to access the object and operation. If an access control rule specifies more than one condition, then the user must meet all conditions to gain access to the object and operation. Failing any one condition check prevents the user from accessing the matching object and operation.

If a user does not meet the conditions of the first matching rule, the system evaluates the conditions of the next matching access control rule as specified by the access control processing order. If the user fails to meet the conditions of any matching access control rule, the system denies access to the requested object and operation.

Note: If there are no matching access control rules for the requested object and operation, then the system grants the user access to it. In practice, it is rare for the system to find no matching rules because the system has a set of default access control rules that protect all record operations.

The effects of being denied access to an object depend on the ACL rule that the user failed. For example, failing a read operation ACL rule prevents the user from seeing the object. Depending on the object secured, the ACL rule hides a field on a form, hides rows from a list, or prevents a user from accessing a UI page. The following table contains a complete list of results of failing an ACL rule for a given operation and object type.

Pre and post query ACL checks

Your instance checks ACL rules both before and after a user makes a query. Because different information is available before and after a query, results can be different.

Pre-query ACL check

Before your instance runs a database query, it checks the ACL rules for each field in the queried table to determine which fields a user may access. This check only looks at the user's roles, and checks to see if these roles allow access to fields. Because this check runs before the query, the ACL doesn't have access to the records on the table, so it can't take that data into account. Scripts and conditions that rely on knowing the contents of a record aren't evaluated.

If the user doesn't have read access at this point, the value for the field isn't shown to the user.

Post-query ACL check

After the query, your instance checks each record returned by the query. During this check, there's context for the ACL, so the role, condition, and script portions of the ACL are evaluated. If the user doesn't have read access at this point, the value for the field isn't shown to the user, however the user sees the field label if their roles allow access to the field.

Operation	Results of failing an ACL rule on object
execute	User can't execute scripts on a record or UI page.
create	User can't see the **New** UI action from forms. The user also cannot insert records into a table using API protocols such as web services. A **create** ACL with a condition requiring that a field contain a specific value may evaluate as false. Fields on new records are considered empty until the record is saved.
read	User can't see the object in forms or lists. The user also can't retrieve records using API protocols such as web services.
write	User sees a read-only field in forms and lists, and the user can't update records using API protocols such as web services.
delete	User cannot see the **Delete** UI action from forms. The user also can't remove records from a table using API protocols such as web services.
edit_task_relations	User cannot define relationships between task tables.
edit_ci_relations	User cannot define relationships between Configuration Item [cmdb_ci] tables.

Operation	Results of failing an ACL rule on object
save_as_template	Used to control the fields that should be saved when a template is created.
add_to_list	User can't view or personalize specific columns in the list mechanic.
list_edit	User can't update records (rows) from a list.
report_on	User can't create a report on the ACL table. For more information, see Restrict report creation with an ACL rule.
report_view	User can't view the content of a report on the ACL table or on the ACL field. For more information, see Reporting.
personalize_choices	User can't right-click a list field and select **Configure Choices**.

ACL matching requirements for objects

Object Type	Matching ACL Rules Required to Access Object	Existing wild-card ACL Rules
Client-callable script includes Processors	Users must meet the conditions of two ACL rules: **1.** All wild-card ACL rules for the object (if any ACL rule exists for the operation). **2.** The first ACL rule that matches the object's name (if any ACL rule exists for the operation).	By default, there are no wild-card (*) rules for these object types. If you create a wild-card ACL rule for one of these objects, then the ACL rule applies to all objects of this type.
UI pages Record	Users must meet the conditions of two ACL rules: **1.** The first ACL rule that matches the record's field (if any ACL rule exists for the operation). **2.** The first ACL rule that matches the record's table (if any ACL rule exists for the operation).	By default, there are wild-card table rules (*) for the create, read, write, and delete operations and wild-card field rules (*.*) for the personalize_choices, create, and save_as_template operations. When you create a table, create ACL rules for the table unless you want to use the provided wild-card ACL rules.

Note: The Security manager default behavior (`glide.sm.default_mode`) property determines whether users can access objects that only match against wild-card table ACL rules. When this property is set to *Deny access*, only administrators can access objects that match the wild-card table ACL rules.

Note: The wild-card field ACL rule (*.*) for the create operation reuses the same conditions as the write operation. This means that the create conditions are the same as the write conditions unless you define an explicit create operation ACL rule.

Multiple ACL rules at the same point in the processing order

If two or more rules match at the same point in the processing order, the user must pass any one of the ACL rules conditions to access the object. For example, if you create two field ACL rules for *incident.number*, then a user who passes one rule has access to the number field regardless of whether the user failed any other field ACL rule at the same point in the processing order.

Required role

Normal admin users can view and debug access control rules. However, to create or update existing access control rules, administrators must elevate privileges to the security_admin role. See Elevate to a privileged role for instructions.

ACL rules in scoped applications

You can create ACL rules for objects in the same scope as the ACL rule. You can also create ACL rules for tables with at least one field that is in the same scope as the ACL rule.

For tables that are in a different scope than the ACL rule record, the types of rules are limited.

- You can create an ACL rule for any table, UI page, or other object that is in the same scope as the ACL rule.
- You can create an ACL for a field that is in the same scope as the ACL rule.
  - If the table is in the same scope, you can use a script to evaluate conditions.
  - If the table is in a different scope, you can't use a script to evaluate conditions.
- You can't create or modify ACL rules for objects that are in a different scope than the application you've selected in the application picker, including adding a role to an ACL in a different scope.
- You can create wild-card table rules (*) only in the global scope.
- You can create wild-card field rules (*) only for tables in the same scope as the ACL rule.

ACL rule types

Create ACL rules on different components of the system.

Record ACL rules

Record ACL rules consist of table and field names.

- The table name is the table that you want to secure. If other tables extend from this table, then the table is considered a parent table. ACL rules for parent tables apply to any table that extends the parent table.
- The field name is the field that you want to secure. Some fields are part of multiple tables because of table extension. ACL rules for fields in a parent table apply to any table that extends the parent table.

ACL rules can secure the following record operations:

Operation	Description
execute	Enables users to execute client callable script includes and REST endpoint execution.

Operation	Description
query_match	Enables users to submit match queries("is", "is not", "is empty", etc).
conditional_table_query_range	Enables users to give partial ACL-access based on read ACLs. Created for the tables that have the read ACLs without Data condition and script.
query_range	Enables users to submit range queries("starts with", "ends with", "contains", etc) and sorting is unrestricted.
create	Enables users to insert new records (rows) into a table.
read	Enables users to display records from a table.
write	Enables users to update records in a table.
delete	Enables users to remove records from a table or drop a table.
edit_task_relations	Enables users to extend the Task [task] table.
edit_ci_relations	Enables users to extend the Configuration Item [cmdb_ci] table.
save_as_template	Enables users to save a record as a template.
add_to_list	Prevents users from viewing or personalizing specific columns in the list mechanic. **Note:** Conditions and scripts are not supported.
list_edit	Enables users to update records (rows) from a list.
report_on	Enables users to report on tables.
report_view	Enables users to report on field ACLs.
personalize_choices	Enables users to configure the table or field.
data_fabric	Allows a data fabric table to reference a local table.

Record ACL rules are processed in the following order:

- Match the object against table ACL rules.
- Match the object against field ACL rules.

This processing order ensures that users gain access to more specific objects before gaining access to more general objects. A user must pass both table and field ACL rules to access a record object.

- If a user fails a table ACL rule, the user is denied access to all fields in the table, even if the user passes a field ACL rule.
- If a user passes a table ACL rule, but fails a field ACL rule, the user cannot access the field described by the field ACL rule.

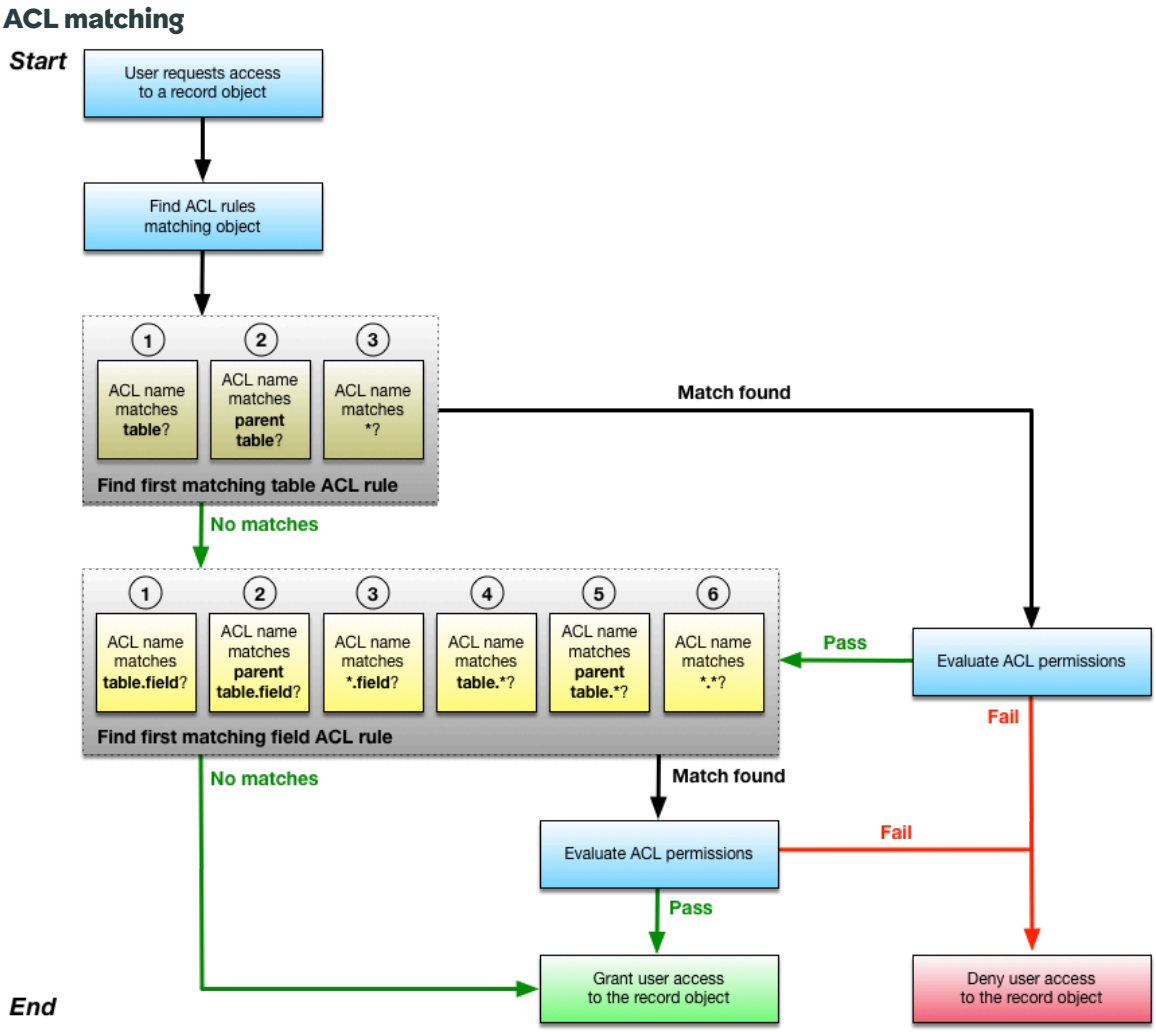

Processor ACL rules

Processor ACL rules specify the processor you want to secure. For a list of available processors, navigate to **System Definition** > **Processors**.

By default, an ACL rule for the EmailClientProcessor is included to restrict the email client to users with the itil role.

Processor ACL rules honor the STAR (*) rule if they cannot find a more specific ACL for those resources.

Table ACL rules

The user must first pass the table ACL rule. Since the base system includes STAR (*) table ACL rules that match every table, the user must always pass at least one table ACL rule. The base system provides additional table ACL rules to control access to specific tables.

Table ACL rules are processed in the following order:

1. Match the table name. For example, incident.
2. Match the parent table name. For example, task.
3. Match any table name (*). For example, *.

If a user fails all table ACL rules, the user cannot access any fields in the table. If a user passes a table ACL rule, the system then evaluates the field ACL rules.

Field ACL rules

After a user passes a table ACL rule, field ACL rules are processed in the following order:

1. Match the table and field name. For example, incident.number.
2. Match the parent table and field name. For example, task.number.
3. Match any table (*) and field name. For example, *.number.
4. Match the table and any field (*). For example, incident.*.
5. Match the parent table and any field (*). For example, task.*.
6. Match any table (*) and any field (*). For example, *.*.

A user must pass the table ACL rule to be granted access to the table's fields. For example, the user must first pass the table ACL rule for the incident table to access the **Number** field in the incident table.

The first successful field ACL evaluation stops ACL rule processing at the field level. When a user passes a field ACL rule, the system stops searching for other matching field ACL rules. For example, if a user passes the field ACL rule for incident.number, the system stops searching for other ACL rules that secure the **Number** field in the incident table.

Access to query information of inferred data is restricted for protected fields, therefore preventing return of predictive information.

UI page ACL rules

UI page ACL rules specify the UI page to be secured. For a list of available UI pages, navigate to **System UI** > **UI Pages**. When defining an ACL rule for a UI page, use the fully scoped page name. For example, **x_myapp_mypage**.

Note: You can use the STAR (*) character in the **Name** field on **ui_page** type ACLs to match any UI pages.

UI page ACL rules honor the STAR (*) rule if they cannot find a more specific ACL for those resources. For example, if you have a UI page named `mysecretpage` but do not define an ACL for this UI page, the STAR (*) rule for the UI page processor is used for access check.

ACL rules can secure the following UI page operation:

Operation	Description
read	Allows users to display the UI page.

Client-callable script include ACL rules

Script include ACL rules specify the client-callable script include to be secured. For a list of available script includes, navigate to **System Definition** > **Script Includes**. You can personalize the list to show the *Client callable* column.

The base system does not include any ACL rules for client-callable script includes.

Client-callable script include ACL rules honor the STAR (*) rule if they cannot find a more specific ACL for those resources.

Datatype ACL

The datatype ACL enables you to write ACL rules that apply to all fields of a specific type.

Datatype ACLs provide a targetted approach to access controls by restricting table fields based on data type. This allows for broader security constraints than wildcard(*) ACLs. The syntax for datatype ACLs follows the format *.[(`field to be restricted`)].

Conventional field ACLs are limited to specific table-field identifiers, all fields in a table, or all tables with a specific field name, in contrast the datatype ACLs enable you to apply security uniformly across fields that share certain metadata. This helps prevent creating additional ACLs to apply security uniformly across all fields.

When implementing datatype ACLs, its essential to validate all affected fields before and after adding them to avoid unexpected security issues. See Create a datatype ACL to create your own datatype ACL.

To review existing datatype ACLs navigate to **All** > **System Security** > **Access Controls** and use the **Name** field to search for ACLs that start with `*.[`.

Note: Scripting Governance uses datatype ACLs by default for scripting restrictions, see _ for more details.

Create a datatype ACL

Learn how to create a datatype ACL.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Security** > **Access Controls**.
2. Select the **New** button.
3. Select the **Change mode** icon near the **Name** field
4. Enter your datatype in the form of `*.[User Datatype]` in the **Name** field.
5. Select a **Role** or **Security Attribute** to associate with the ACL.
6. **Optional:** Fill in any additional details on the **Create an ACL** form.
   See Configure an ACL rule for additional information.
7. Select **Submit**.

8. Verify the security rules for the ACL in the popup **Verify Security Rules** window.

Example of Verify Security Rules

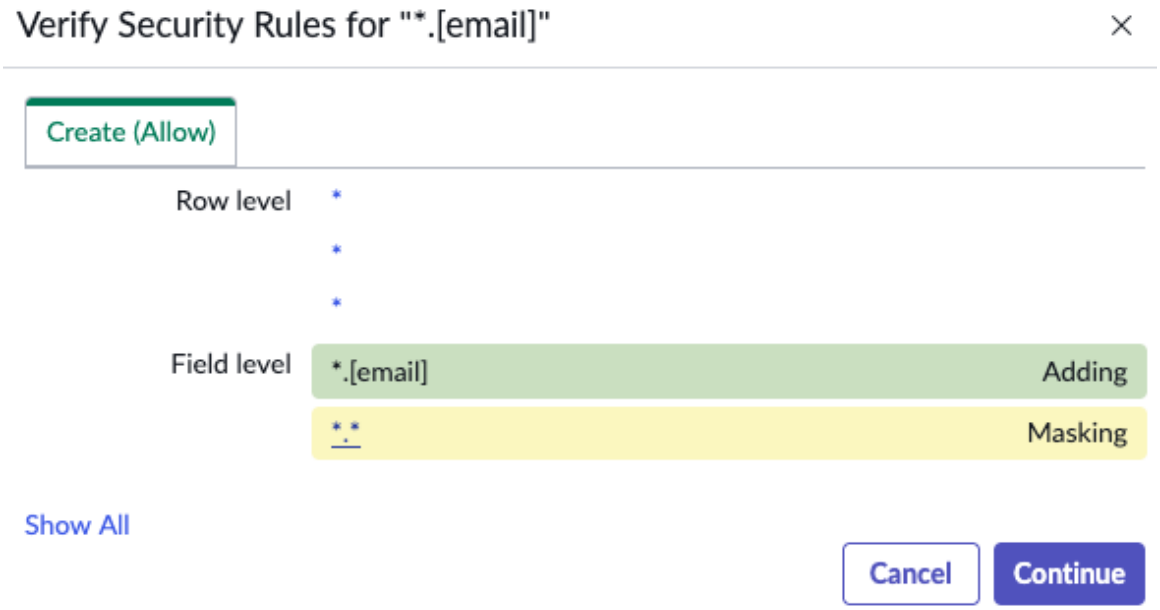

ACL control of function fields

When evaluating access to a function field, in addition to checking access to the function field itself, the system also checks access to the function's contributing fields. Contributing fields are those used as the arguments in a given function definition.

For more information about function fields, see Function field.

In Rome and earlier, the system simply checks access to the function field itself (as with any other field). If the ACLs on that field allow access, the user receives the resulting value, regardless of whether the user has access to the contributing fields.

In Zurich and later, the system also requires access to all contributing fields in order to allow access to the function field. If one or more of the contributing field ACLs refuse access, the function field also refuses access.

The only operations affected by the new requirement are read and report_view. Report_view has its own additional requirements.

Operation	Description
read operation	A user has read access to a function field only if both of the following are true: • The user has read access to the function field. • The user has read access to all of the contributing fields used in the function.
report_view operation	A user has report_view access to a function field only if all of the following are true: • The user has report_view access to the function field. • The user has report_view access to each of the contributing fields.

Operation	Description
	• There is a role-only read ACL without conditions and without a script, and the user has that role. • The user has role-only read access to the contributing fields, such that only ACLs without condition or script can allow.

Examples

Given:

- Table: salary
- Columns: base, bonus, total (all are Integers in this example)
- Function field: The total column is marked as a function field, with function definition `glidefunction:add(base, bonus)`.
- Contributing fields: base and bonus, since they're used in the function definition
- Roles: salary_admin, bonus_admin

Example 1: All fields allow access

ACLs	Result
total, base, bonus: read and report_view for role salary_admin, with no conditions or scripts	A user with the salary_admin role is granted read and report_view access to total because they have the required role.

Example 2: Contributing field refuses read access

ACLs	Result
• total, base: read and report_view for role salary_admin, with no conditions or scripts • bonus: report_view for role salary_admin , with no conditions or scripts • bonus: read for role bonus_admin, with no conditions or scripts	A user with the salary_admin role is refused read and report_view access to total, because bonus refuses read access to their role.

Example 3: Contributing field ACL has script

ACLs	Result
• total, base: read and report_view for role salary_admin, with no conditions or scripts • bonus: report_view for role bonus_admin, with no conditions or scripts	A user with the salary_admin role is granted read access to total, because they have the required role for all fields. But the same user with the salary_admin is refused report_view access, because the read ACL with the script refuses access by default for this case, even though they have the required role.

Example 3: Contributing field ACL has script (continued)

ACLs	Result
• bonus: read for role salary_admin, with a script (note that it doesn't matter what's in the script, only that it's there)	

Security jump-start - ACL rules plugin

The Security jump-start access control level (ACL Rules) plugin is installed automatically on all new instances. Use this plugin to quickly secure multiple system tables and expedite the production launch process for your organization.

This plugin isn't intended for existing instances, as it may alter security access to tables already in use in a production environment. If an admin strongly desires to install this plugin on an existing instance, you should test it thoroughly in a test instance first. Doing so helps to verify compatibility with the organization's current implementation.

If an admin is interested in the new ACL rules provided by this plugin, they can manually create one or more in an existing instance, using the list of ACLs as a guideline.

The following ACLs are included in this plugin. Select the icon in a header row to sort that column in ascending or descending order. The Operation key is as follows:

- R=read
- W=write
- D=delete
- C=create

Name	Operation	Description
`cmdb_ci`	WCD	asset or itil role required to write/create/delete Configuration Item records
`cmn_department`	WD	user_admin role required to write/delete Department records
`cmn_location`	WC	user_admin role required to write/create Location records
`core_company`	WD	user_admin role required to write/delete Company records
`kb_knowledge`	create	knowledge role required to create Knowledge records
`ldap_ou_config`	RWCD	user_admin role required to read/write/create/delete LDAP OU Definition records
`ldap_server_config`	RWCD	user_admin role required to read/write/create/delete LDAP Server records
`process_guide`	WCD	admin role required to write/create/delete Process Guide records

Name	Operation	Description
process_step	WCD	admin role required to write/create/delete Process Step records
sc_category	create	catalog_admin role required to create Service Catalog Category records
sc_category	delete	catalog_admin role required to delete Service Catalog Category records
sc_category	write	catalog_admin role required to write to Service Catalog Category records
sc_cat_item	write	catalog_admin role required to write to Catalog Item records
sc_cat_item	delete	catalog_admin role required to delete Catalog Item records
sc_cat_item	create	catalog_admin role required to create Catalog Item records
sysevent_email_action	read	all users can read Email Notification records (for subscription purposes)
sysevent_register	RWCD	admin role required to read/write/create/delete Event Registry records
sysevent_script_action	RWCD	admin role required to read/write/create/delete Script Action records
syslog	RWCD	admin required to read/write/create/delete Log Entry records
sysrule	RWCD	admin required to read/write/create/delete Rule records (Email Notifications, Inbound Email Actions, Approval Rules, and so on)
sysrule	read	all users can read Email Notification records for (subscription-based notifications)
sys_app_application	WCD	admin required to write/create/delete Application records
sys_app_category	WCD	admin role required to write/create/delete Application Category records
sys_app_module	WCD	admin required to write/create/delete Module records
sys_audit	RWCD	admin required to read/write/create/delete Audit records
sys_dictionary	RWC	personalize_dictionary role required to read/write/create Dictionary records
sys_dictionary.*	read	personalize_dictionary role can read Dictionary fields
sys_documentation	delete	personalize_dictionary role required to delete Field Label records
sys_documentation	create	personalize_dictionary role required to create Field Label records

Name	Operation	Description
sys_documentation	write	personalize_dictionary role required to write to Field Label records
sys_gauge	RWCD	admin role required to read/write/create/delete Gauge records
sys_gauge_count	RWCD	admin role required to read/write/create/delete Gauge Count records
sys_group_has_role	read	itil role required to see Group Role records
sys_home	WCD	itil_admin role required to write/create/delete Welcome Page Section records
sys_installation_exit	WCD	admin role required to write/create/delete Installation Exit records
sys_job	WCD	admin role required to write/create/delete Sys Job records
sys_nav_link	WCD	admin role required to write/create/delete Navigation Link records
sys_perspective	WCD	admin role required to write/create/delete Menu List records
sys_portal	RWCD	admin role required to read/write/create/delete Portal records
sys_portal_page	RWCD	admin role required to read/write/create/delete Homepage records
sys_portal_preferences	RWCD	admin role required to read/write/create/delete Portal Preferences records
sys_processor	WC	admin role required to write/create Processor records
sys_properties	WC	admin role required to write/create System Property records
sys_properties_category	WCD	admin role required to write/create/delete Property Category records
sys_report	delete	roles that can delete Report records (doesn't restrict deleting through Report UI)
sys_report	write	roles that can write to Report records (doesn't restrict editing through Report UI)
sys_report	read	users can read their own Report records, those of their groups, and GLOBAL ones (doesn't affect viewing through Report UI)
sys_report	read	roles that can read Report records (doesn't restrict viewing through Report UI)
sys_reportroles	read	admin role required to read Report Roles records
sys_script	WCD	admin role required to write/create/delete Business Rule records

Name	Operation	Description
sys_script_ajax	WCD	admin role required to write/create/delete AJAX Script records
sys_script_client	WCD	admin role required to write/create/delete Client Script records
sys_script_include	WCD	admin role required to write/create/delete Script Include records
sys_security_acl	write	admin role required to write to Access Control records
sys_security_acl_role	create	admin role required to create Access Roles records
sys_security_acl_role	delete	admin role required to delete Access Roles records
sys_security_acl_role	write	admin role required to write to Access Roles records
sys_security_operation	delete	admin role required to delete Security Operation records
sys_security_operation	create	admin role required to create Security Operation records
sys_security_operation	write	admin role required to write to Security Operation records
sys_security_operation	query_range	
sys_security_operation		
sys_security_type	write	admin role required to write to Security Type records
sys_security_type	create	admin role required to create Security Type records
sys_security_type	delete	admin role required to delete Security Type records
sys_status	create	admin role required to create System Status records
sys_status	delete	admin role required to delete System Status records
sys_status	write	admin role required to write to System Status records
sys_template	write	template_editor role required to write to Template records
sys_template	create	emplate_editor role required to create Template records
sys_template	delete	template_editor role required to delete Template records
sys_template	read	template_editor role required to read Template Roles records

Name	Operation	Description
sys_ui_action	create	admin role required to create UI action records
sys_ui_action	delete	admin role required to delete UI action records
sys_ui_action	write	admin role required to write to UI action records
sys_ui_action_view	write	admin role required to write to UI View Action records
sys_ui_action_view	create	admin role required to create UI View Action records
sys_ui_action_view	delete	admin role required to delete UI View Action records
sys_ui_policy	create	admin role required to create UI Policy records
sys_ui_policy	delete	admin role required to delete UI Policy records
sys_ui_policy	write	admin role required to write to UI Policy records
sys_ui_policy_action	create	admin role required to create UI Policy Action records
sys_ui_policy_action	delete	admin role required to delete UI Policy Action records
sys_ui_policy_action	write	admin role required to write to UI Policy Action records
sys_ui_script	write	admin role required to write to UI Script records
sys_ui_script	delete	admin role required to delete UI Script records
sys_ui_script	create	admin role required to create UI Script records
sys_user	write	Users with no role can't update any user record but their own
sys_user_grmember	delete	user_admin role required to delete Group Member records
sys_user_grmember	write	user_admin role required to write to Group Member records
sys_user_group	create	Only itil and above can create group records
sys_user_group	write	Only itil and above can write to group records
sys_user_has_role	read	itil role required to see User Role records
sys_user_role	create	admin role required to create Role records
sys_user_role	delete	admin role required to delete Role records
sys_user_role	write	admin role required to write to Role records

Name	Operation	Description
`sys_user_role_contains`	read	itil role required to see Contained Role records
`sys_user_role_contains`	write	admin role required to write to Contained Role records
`sys_user_token`	RWCD	admin role required to read/write/create/delete User Token records

Note: To learn more about this plugin, see Enable security jump start plugin (ACL Rules) [Updated in Security Center 1.3] in Instance Security Hardening Settings.

Configure an ACL rule

Configure custom access control list (ACL) rules to secure access to new objects or to change the default security behavior.

Before you begin

Role required: security_admin

About this task

To create ACL rules, you must elevate privileges to the security_admin role.

For tables that are in a different scope from the ACL rule record, the types of rules are limited. For Scope Master tables to derive scope and execute scoped ACLs, set the *glide.enforce_security_scope.<scope_name>* property to **true**. This ensures ACLs in the global scope don't match when there are scope-specific ACLs created on the relevant table. Examples are when securing data within shared application tables in the Global scope, such as sys_attachment or sys_question_answer tables.

Procedure

1. Elevated privilege roles to the security_admin role.
2. Navigate to **System Security** > **Access Control (ACL)**.
3. Select **New**.

   Tip: When creating an ACL, it's helpful to review the Deny-Unless ACL.

4. Complete the form.

Access control fields

Field	Description
Type	Select what kind of object this ACL rule secures. The type of object determines how the object is named and what operations are available. This field becomes read only after the ACL rule is created. If you want to change the type, you must delete the ACL and create one with the correct type.
Operation	Select the operation that this ACL rule secures. Each object type has its own list of operations. An ACL rule can only secure one operation. To secure multiple operations, create a separate ACL rule for each. If you're creating a rule for a report_view operation, see also Report_view access control.

Field	Description
Decision Type	Select the decision type of the ACL. **Allow If** allows access upon successful evaluation. **Deny Unless** denies access unless there's successful evaluation.See Deny-Unless ACL for more information.
Admin overrides	Select this check box to have users with the admin role automatically pass the permissions check for this ACL rule. Admin users pass regardless of what script or role restrictions apply. However, the **nobody** role, which only ServiceNow personnel can assign, takes precedence over the admin override option. If an ACL is assigned the **nobody** role, admin users can't access the resource even when **Admin overrides** is selected. See Base system roles. Clear this check box if administrators must meet the permissions defined in this ACL rule to gain access to the secured object. Since administrators always pass role checks (see the description of the **Requires role** field), use the condition builder or **Script** field to create a permissions check that administrators must pass.
Protection Policy	Select this to set the protection policy on the ACL
Name	Enter the name of the object being secured, either the record name or the table and field names. The more specific the name, the more specific the ACL rule. You can use the wildcard character asterisk (*) in place of a record, table, or field name to select all objects that match a record type, all tables, or all fields. You can't combine a wildcard character and a text search. For example, inc* isn't a valid ACL rule name, but incident.* and *.number are valid ACL rule names. **Note:** Select the blue triangle to manually enter the record name or the table and field names of the object being secured. Use this option to secure an object that doesn't appear in the dropdown.
Description	Enter a description of the object or permissions that this ACL rule secures.
Applies To	Select **Add Filter Condition** to create a filter for the table selected in the **Name** field. The ACL applies only to the records matching this filter condition. For example, you could create an ACL for the incident table that applies only to incidents with the Critical priority. **Note:** You must select a table in the **Name** field to use this option. ACLs beginning with * cannot use the **Applies To** filter.
Active	Select this check box to enforce this ACL rule.
Advanced	Select this check box to display the **Advanced Condition** fields. See step 6.

5. **Optional:** To narrow the scope of the ACL fill in the **Conditions** fields as necessary.

Requires role	Use this list to specify the roles a user must have to access the object. If you list multiple roles, a user with any one of the listed roles can access the object. The *Requires role* list appears as a related list. **Note:** Users with the admin role always pass this permissions check because the admin role automatically grants users all other roles.

Security Attribute Condition	Use this section to define what the user can access based on user and environment criteria. For more information, see Security Attributes fundamentals. **Note:** The Condition field is case-sensitive.
Data Condition	Use this condition builder to select the fields and values that must be true for users to access the object. **Note:** The Condition field is case-sensitive.

6. **Optional:** If the **Advanced** box is checked, fill in the **Advanced Conditions** fields as necessary.

Controlled by References	Enforces the ACL on related records. See Related record access for more details.
Script	Enter a custom script describing the permissions required to access the object. The script can use the values of the *current* and *previous* Global variables in business rules as well as system properties. The script must generate a true or false response in one of two ways: ○ return an *answer* variable set to a value of true or false ○ evaluate to true or false In either case, users only gain access to the object when the script evaluates to true and the user meets any conditions the ACL rule has. Both the conditions and the script must evaluate to true for a user to access the object. If there's script in the **Script** field. This script executes even if the field isn't displayed on the form. **Note:** If the evaluated item is in a related list, **current** points to the item the related list is on, not to the current item the ACL is for. However, If the item you're evaluating the ACL for isn't in a related list, **current** points to the actual item.

7. Select and hold (or right-click) the form header and select **Save**.

Deny-Unless ACL

Learn details about Deny-Unless ACLs.

Deny-Unless ACLs are evaluated with a "deny-unless" approach. The ACL defines the users that will NOT be denied. Said another way, the user will be denied access **unless** the role, condition, and script requirements are met.

Important:

Deny-Unless ACLs will take priority against Allow-If ACLs in ACL Evaluation, as it will be evaluated first.

A Deny-Unless ACL produces two outcomes

Evaluation outcome	Result
Pass	The defined roles, data conditions, security attributes, and script requirements are met. The ACL proceeds to further evaluation **Important:** Even if a Deny-Unless ACL matches, access is only granted when an Allow-If ACL explicitly permits it. If no Allow-If ACL is matched and the Deny-Unless ACL passes, the system grants access by default.
Fail	The Deny-Unless ACL is marked as failing and access will be denied.

The following is an explained example of a Deny-Unless ACL:

- ACL has roles `sn_hr_core.manager` and `itil`
- Condition has active = `true`
- script has answer = `gs.isLoggedIn();`

The user is denied access unless all three requirements for this ACL are satisfied. In order for this Deny-Unless ACL to pass, a user needs either the `sn_hr_core.manager` or `itil` roles, be accessing a record that has active field = `true`, and be logged in. The Deny-Unless ACL will fail if any of the three requirements isn't met.

Query ACLs

Query ACLs allow you to define more granular access control by explicitly defining who can query the data.

What is a query ACL

A query ACLs have their operation set to either `query_range` or `query_match`. Query ACLs allow for more specific control of user queries, restricting or enabling access based on their setup. Query ACLs are powerful tools against blind query attacks, where an attacker blindly queries the data to extract information from results, even when they can't see the values.

When to use a query ACL

Wherever a column contains sensitive values, and allows partial/conditional access to data a query ACL should be considered and implemented as necessary based on the sensitivity of the data. Wherever there is a partial/conditional access to rows and their columns in tables, especially where that access is not enforced by data filters, query ACLs should be implemented as necessary based on the sensitivity of the data.

Note: Consider query ACLs when some users have access to some rows or columns and not others .

Example: Payroll query control

I can see one row in payroll table with my salary, but there is no reason for me to be able to issue range queries to query users with a salary contained within 2 boundaries. A `query_range` ACL on salary would prevent me from issuing that query.

Example: HR query control

I can see all hr_profiles, but can only see SSN for myself. I have no business querying SSN, and query ACLs should prevent me from running queries against SSN of other hr profiles to try to extract SSN mappings.

Query ACL behavior

Query ACLs use `query_match` and `query_range` operations for secure and granular table querying behavior. Their behaviors are described below:

`query_match`

`query_match` is composed of: EQUALS, NOT_EQUALS, IN, NOT_IN, SAMEAS, NSAMEAS, ANYTHING, ISEMPTYSTRING, ISEMPTY, ISNOTEMPTY, ISNULL, ISNOTNULL. `query_match` is made of the "safe operators", in a sense that they are built to fetch specific record(s), and can't be exploited to return others.

Evaluation outcome	Result
Pass	User can submit match queries
Fail	User will not be able to submit match queries: • EQUALS • NOT_EQUALS • IN • NOT_IN • SAMEAS • NSAMEAS • ANYTHING • ISEMPTYSTRING • ISEMPTY • ISNOTEMPTY • ISNULL • ISNOTNULL

`query_range`

`query_range` is composed of all the others (STARTS_WITH, CONTAINS, >=, <= etc) which are more dangerous as they allow users to query for more records by adjusting the boundary values.

Evaluation outcome	Result
Pass	User can submit range queries and sorting is unrestricted
Fail	The user will not be able to submit range queries with (STARTS_WITH,

Evaluation outcome	Result
	CONTAINS, >=, <=, etc. Sorting by column is restricted

Important:

Query ACLs (both query_match and query_range) default to a star.star ACL that delegates to read access. This means, where ACLs are enforced on queries, if no query ACL was created then read access to the column is evaluated ; if query ACLs are defined then they override the default behavior.

Secure records in an embedded list

To apply security to the records in embedded lists, limit editing and deleting records in embedded lists to specific roles.

Before you begin

Role required: security_admin

Procedure

1. Navigate to **All** > **System Security** > **Access Control (ACL)**.
2. Open the **Write** or **Delete** record for the appropriate table.
3. In the Requires Role section of the form, add the roles that have write or delete permission for that table.
4. Save the changes.
   When records from the associated table appear in an embedded list, the edit and delete options are available only to users with the specified roles.

Related record access

Related record access enable consistent control over what records users are able to access between related tables.

Related record access details

Related record access enables you to configure access to related record, determining whether the related parent table's ACL should be enforced on data that is allowed access from the ACL. Consider the following use-case, a user has access to a record via ACL, then related record access enables the users' access to all other records related to it, via either reference or bidirectional relationship.

Contextual Security Manager

Contextual Security Manager protects your data by controlling read, write, create, and delete authorization.

Key advantages

The Contextual Security Manager is aware of the system table hierarchy, enabling you to create specific security rules for a field based on where in the hierarchy it is displayed. Benefits of the Contextual Security Manager include:

- **Contextual security**: Secure a record based on the contents of the record.
- **Hierarchical security**: Apply security rules to any level in the object hierarchy.

Securing fields and tables

With the legacy Simple Security Manager, you could secure fields and tables by adding roles to the appropriate dictionary entry. With the Contextual Security Manager, these dictionary roles are no longer tested. Instead, the system looks for ACL rules on fields and tables.

Warning: After you install the Contextual Security Manager, you must secure fields and tables via ACL rules. Even if you Configuring the form layout the dictionary form and add roles to a dictionary entry, no change in rights occurs.

Contextual security and roles

You can grant roles to users or groups. However, after installing the Contextual Security Manager, the **roles** field on the user record is no longer checked and no longer appears on your user and group forms. Instead, you must add roles to the Roles related list instead of to the user or group record.

Applications and modules contain lists of the roles required to view them. For example, to view the System Definition application, the admin role is required. Security rights for applications and modules are still defined using role arrays.

Both catalog items and catalog variables contain lists of the roles required to view them. Security rights for catalog items and catalog variables are still defined via these role arrays.

Under the Contextual Security Manager, a group still automatically inherits any role granted to the group when the inherits flag for the role is set to *true*.

Activating the Contextual Security Manager

The Contextual Security Manager is active in the base system. If there are many duplicate entries in the User Roles table, you may need to upgrade to Contextual Security: Role Management V2 to eliminate duplicate roles. Plugins include:

Contextual Security: Role Management [com.glide.role_management]

Provides contextual security functionality. This plugin is automatically installed.

Contextual Security: Role Management V2 [com.glide.role_management.inh_count]

Prevents duplicate entries caused by inherited roles in the User Roles [sys_user_has_role] table. This plugin is automatically installed on new instances and can be activated for upgrades. The Contextual Security: Role Management Enhancements plugin is a previous version of this plugin. The Role Management Enhancements plugin does not include the RoleManagementVerify() script. This script returns a list of changes that an upgrade will perform, enabling you to monitor changes made by the plugin.

Note: After activating Role Management V2, you must set the glide.role_management.v2.audit_roles system property to allow the Audit Roles table to create audit records related to user roles. To learn more about setting this property and about the Audit Roles table, see:

- Enable role auditing with Contextual Security: Role Management V2.
- Hardening settings in Instance Security Hardening Settings.
- Audit user roles

Security Attribute Conditions

Starting with the Vancouver release, ACL rules support a Security Attribute Condition type in addition to the standard role, condition, and script fields. Security Attribute Conditions evaluate properties of the current user or session, such as group membership, authentication state, or network location, to determine whether access is granted.

Security Attribute Conditions can be defined as local or existing:

Local

The condition is saved only to the ACL where it is created. Conditions default to local.

Existing

The condition references a Security Attribute condition already defined on another ACL, enabling reuse across multiple rules.

For a complete list of built-in security attribute types and configuration details, see Security Attributes fundamentals and Security Attribute Scope.

Prevent duplicate entries with Contextual Security: Role Management V2

Roles inherited from other roles are added as individual entries in the User Roles table [sys_user_has_role], potentially causing one role to have duplicate entries. Contextual Security: Role Management V2 eliminates these duplicate entries and prevents future duplicates.

Eliminate duplicate entries through inheritance count

Contextual Security: Role Management V2 uses the Inheritance Count (inh_count) column to track the number of times a role is inherited from another role or group. In the User Roles [sys_user_has_role] table, a user can inherit a specific role only one time, eliminating duplicate entries. The Inheritance Count (inh_count) column is read-only and calculates the number of times the user inherits a role.

Activation changes

Contextual Security: Role Management V2 is automatically installed on new instances and can be activated for upgrades. When activated, Contextual Security: Role Management V2 replaces both Contextual Security and Contextual Security: Role Management Enhancements.

When Contextual Security: Role Management V2 is activated, the following columns are deprecated, but remain in the User Roles table for backward compatibility:

- granted_by (used only by Role Delegation)
- included_in_role
- included_in_role_instance

Warning: If these columns are in use in any custom scripts on your instance, do not upgrade to Role Management V2.

Visualize role inheritance through the Role Inheritance Map

The Role Inheritance Map displays a visual representation of inherited roles. You can use this map to understand the roles represented in the Inheritance Count (inh_count) column. To view the Role Inheritance Map, configure the User Roles [sys_user_has_role] table to display the Role Inheritance Map column.

Role Inheritance Map

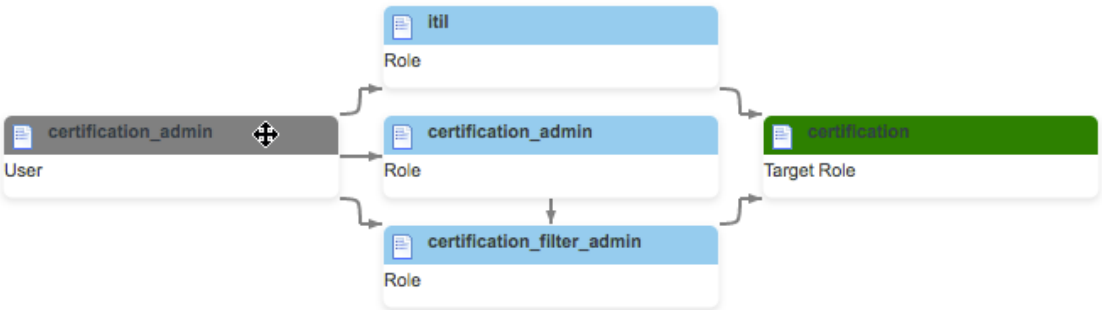

Note: Concurrent update to group or role assignment may result in incorrect inheritance count. You must enable the `glide.security.inh_count_patcher.enabled` property to get the exact inheritance count.

Upgrade to Contextual Security: Role Management V2

Contextual Security: Role Management V2 is automatically installed on new instances. You can upgrade from Contextual Security: Role Management to Contextual Security: Role Management V2 to eliminate duplicate roles in the User Roles table and prevent future duplicates.

Before you begin

Role required: admin, security_admin
You must have the admin role and Elevate to a privileged role to obtain the effective security_admin role.

Note: Before upgrading to Contextual Security: Role: Management V2, you must enable the Audit Roles table to create audit records related to user roles. To learn more about setting the required system property, see Enable role auditing with Contextual Security: Role Management V2.

About this task

This procedure describes how to upgrade your Contextual Security role and how to activate the related plugins described in the following table.

Plugins for Contextual Security: Role Management V2

Plugin	Description
Contextual Security: Role Management V2 [com.glide.role_management.inh_count]	Prevents duplicate entries in the User Roles [sys_user_has_role] table. The security_admin role or a user with elevated privileges is required to activate the plugin, or contact Customer Service and Support.
Contextual Security: Role Management V2 REST API [com.glide.role_management.inh_count.rest_api]	Enables API functionality for role management.

Before upgrading from Contextual Security: Role Management to Contextual Security: Role Management V2, test the results of an upgrade by running the script. The script returns a list of changes that an upgrade will perform. If the changes are acceptable, install the Contextual Security: Role Management V2 plugin. If the changes are not acceptable, do not install the Contextual Security: Role Management V2 plugin. Alternatively, you can perform the upgrade and then manually make any necessary changes.

Procedure

1. Test the impact of an upgrade prior to upgrading by running the following script.

   a. Navigate to **System Definition** > **Scripts - Background**.

   b. Run the following script in global scope.

   ```
   new RoleManagementVerify().verifyInheritedRoles();
   ```

   For large sys_user_has_role tables, the execution may take up to several hours to complete. Do not edit or add user roles during this time.
   Example result based on test data:

   ```
   *** Script: 2016-12-01 19:58:54 Starting checking of
    inherited roles for all users...
   *** Script: User: itam, inherited roles to be ADDED:
    financial_mgmt_user
   *** Script: User: bernard.laboy, inherited roles to be
    DELETED: api_analytics_read,pa_viewer,rest_api_explorer,a123
   *** Script: User: bernard.laboy, inherited roles to be ADDED:
    dependency_views
   *** Script: Number of inherited-role records in sys_user_has
    role, current: 260, after re-calculation: 258
   *** Script: Number of users with discrepancies for inherited
    roles: 2
   *** Script: 2016-12-01 19:58:55 Finished checking of
    inherited roles for all users!
   ```

   c. Evaluate the script results to determine whether the proposed changes are acceptable.

2. Activate the Contextual Security: Role Management V2 plugin.

   Important: The security_admin role or a user with elevated privileges is required to activate the plugin, or contact Customer Service and Support.

   a. Navigate to **System Definition** > **Plugins**.

   b. Find and click the plugin name.

   c. On the System Plugin form, review the plugin details and then click the **Activate/Upgrade** related link.

   d. Click **Activate**.

Result

After activating Role Management V2, the changes outlined in the script result are enacted. The Inheritance Count (inh_count) column in the User Roles table is read-only and automatically reflects the number of times the user inherits a role.

Enable role auditing with Contextual Security: Role Management V2

Set a system property to enable the Audit Roles table to create audit records related to user roles.

Before you begin

Role required: admin

About this task

When enabled, the Audit Roles [sys_audit_role] table maintains changes to user records. For more information about role audits, see Audit user roles. If the Contextual Security: Role Management V2 [com.glide.role_management.inh_count] plugin is installed, you must set a system property to **true** to enable role auditing.

Procedure

1. Navigate to the System Properties [sys_properties] table.
2. Add the *glide.role_management.v2.audit_roles* system property and set it to **true**.

   If the Contextual Security: Role Management V2 [com.glide.role_management.inh_count] plugin is installed, setting this property to **true** enables the Audit Roles [sys_audit_role] table to create records when user roles change. The table records role changes that occur after the property is set. Existing role assignments are not backfilled.

Double-check form submission

When the system determines that a particular field (such as task.number) should not be written to by the current user, the system renders that field in a read-only mode, which is why the number field is not writable on most incidents.

If you set the system to double-check the values of any incoming fields for writability, then the system applies the same set of security rules to the inbound leg of a transaction. When you submit an incident, for example, the system double-checks to determine if the number field can be written to before posting any changes.

If you tell the system not to double-check inbound transactions, then the system allows you to write to a nominally read-only field if that is the transaction the client sends back. In many deployments this is actually a desirable behavior if, for example, you are using client scripts to set nominally read-only fields in response to user selections in other, writable fields.

Property	Location	Default
Double check security on inbound transactions during form submission (rights are always checked on form generation)	**System Properties** > **Security**	Disabled (no double-checking)

Default deny property

The default deny property (`glide.sm.default_mode`) controls the security manager default behavior when the only matching ACL rules are the wildcard table ACL rules.

A set of wildcard table ACL rules for the most common record-based operations are available: read, write, create, and delete. A significant number of ACLs to provide role-based access to

system tables are also available. For example, there are ACLs that grant sys_script access to the business_rule_admin role because that role is documented as being able to manage business rules.

Use the `glide.sm.default_mode` property to deny or allow these operations on all tables:

- **Deny Access**: The wildcard table ACL rules restrict the read, write, create, and delete operations on all tables unless the user has the admin role or meets the requirements of another table ACL rule. Other operations, such as report_on and personalize_choices, are unaffected by this setting.
- **Allow Access**: The wildcard table ACL rules allow the read, write, create, and delete operations on all tables unless there are specific table ACL rules in place to restrict such operations.

You cannot reset `glide.sm.default_mode` to **Allow Access** once it has been set to **Deny Access**.

Note: By default, the wildcard table ACL rules are the only ACL rules that check for the value of the `glide.sm.default_mode` property. If you want to control other operations with this setting, create your own ACL rules to check for this property value.

To learn more about this property, see Deny by default with empty ACLs [Updated in Security Center 1.3] in Instance Security Hardening Settings.

Advanced ACL configuration

In addition to creating new ACLs or modifying existing ones, you can configure other aspects of ACL functionality.

Deny ACL rules

Deny-Unless

Query ACL rules

Task	Description
Apply ACL script conditions to reference fields	Enable a property to allow script conditions to apply to reference fields if you want to control access to the data that a reference field displays on a form or in a list. There might be an impact to the performance of your instance if you enable this.
Apply ACLs to AJAXGlideRecord (client-side Glide record)	Apply ACLs to *`GlideAjax`* API calls so that the system queries only the data that the currently connected user has rights to access.
Evaluate the admin override at the access level	Force ACL evaluation for admin overrides at the access level. By default, users with the admin role automatically pass the permissions check for this ACL rule when the Admin Overrides option is selected on the ACL rules form.
Use ACL debugging and troubleshooting tools	Use tools like the ACL watcher, field level debugging, and access ACL rule output messages to help you troubleshoot and debug ACLs.

Provide external users access to a table

To enable users with only the snc_external role to access the list view of a table, you must create a series of ACLs.

Before you begin

Role required: security_admin

Procedure

1. Elevate to a privileged role.
2. Create an ACL rule with the following settings:
   - **Type**: ui_page
   - **Operation**: read
   - **Name**: {table_name}_list
   - **Required role**: snc_external
3. On the default **read** ACL for the table, add **snc_external** in the Required role list.
   Create the ACL if it does not already exist.
4. Use these settings to create another ACL:
   - **Type**: ui_page
   - **Operation**: read
   - **Name**: {table_name}
   - **Required role**: snc_external
5. Use these settings to create another ACL to give the user write access to a field in the table:
   - **Type**: record
   - **Operation**: create
   - **Name**: {table_name} {column_name}
   - **Required role**: snc_external

   Repeat this step for every field that you want to give the user write access to. Use an asterisk * instead of the column name to provide access to all fields at once.

Apply ACL script conditions to reference fields

Use the *glide.sys_reference_row_check* system property to enable scripted conditions for reference fields.

The default behavior is intended to improve instance performance. If you want to enable script conditions for reference fields, add the following system property.

Note: For more information on creating system properties, see Add a system property.

System property

Property	Description
glide.sys_reference_row_check	Controls whether the script conditions of Access Control Rules apply to a table's reference fields.

System property (continued)

Property	Description
	• Type: true | false • Default value: **false** • Location: Add a system property to the System Properties [sys_properties] table

Note: If the *glide.sys_reference_row_check* system property is not present, or has been set to false, script conditions for Access Control Rules are not applied. This means an ACL containing scripted conditions will pass it's check as long as the other ACL criteria are met (such as role requirements).

Apply ACLs to AJAXGlideRecord (client-side Glide record)

Use a system property to perform access control list (ACL) rule validation when server-side records (for example, tables) are accessed using GlideAjax APIs within a client script.

If you choose to apply access control lists (ACL) to *GlideAjax* API calls, then you can only query data to which the currently connected user has rights to access. For example, if the user is logged in as an ESS user who has no rights to read the `cmn_location` table, then any *GlideAjax* API call by the user will fail.

If the ServiceNow AI Platform is running without GlideAjax ACL call checking, an API can return information that the currently logged in user could not otherwise access.

Use GlideRecordSecure when querying data to ensure the highest level of security. GlideRecord relies on ACL enforcement through configurations whereas GlideRecordSecure applies stricter security controls. GlideRecordSecure offers a more secure, out-of-the-box solution for handling sensitive data.

Note: Set this property in **System Properties** > **Security**.

Property	Default
Apply standard security ACLs to *AJAXGlideRecord* calls	ACL checking enforced

Warning: The `sys_class_name, sys_id` and `sys_domain` are ignored for ACL checks.

To learn more about this property, see Require AJAXGlideRecord ACL checking [Updated in Security Center 1.3] in Instance Security Hardening Settings.

Evaluate the admin override at the access level

If you want to force ACL evaluation for admin overrides at the access level, you can add a system property.

Before you begin

Role required: security_admin

About this task

ACLs are evaluated cumulatively. If there are number of ACLs on any given field and the **Admin Overrides** option is **false** (not selected) on one of them, then the effective admin overrides for all the ACLs are considered to be **false**. This causes admins to be unable to pass even the ACL where the override should be in effect.

Procedure

Add the following property to the system properties table:

Property	Description
glide.security.admin.override.accessterm	Evaluates the admin override condition at the access term level. • Type: true | false • Default value: **true** for new instances, **false** for upgrades • Location: Add to the System Properties *[sys_properties]* table **Note:** If the property is not defined on the Instance, the value evaluates as false.

ACL debugging tools

Field level debugging and access ACL rule output messages are available to help you troubleshoot and debug ACLs. The ACL configuration watcher lets you know what related ACLs exist when you modify one.

Field level debugging

When debugging is enabled, a small bug icon ( ) appears beside each field with an ACL rule. Clicking the icon lists the ACL rules that apply for the field and the evaluation results. Enable debugging by navigating to **System Security** > **Debugging** > **Debug Security Rules**.

Field level security on an incident

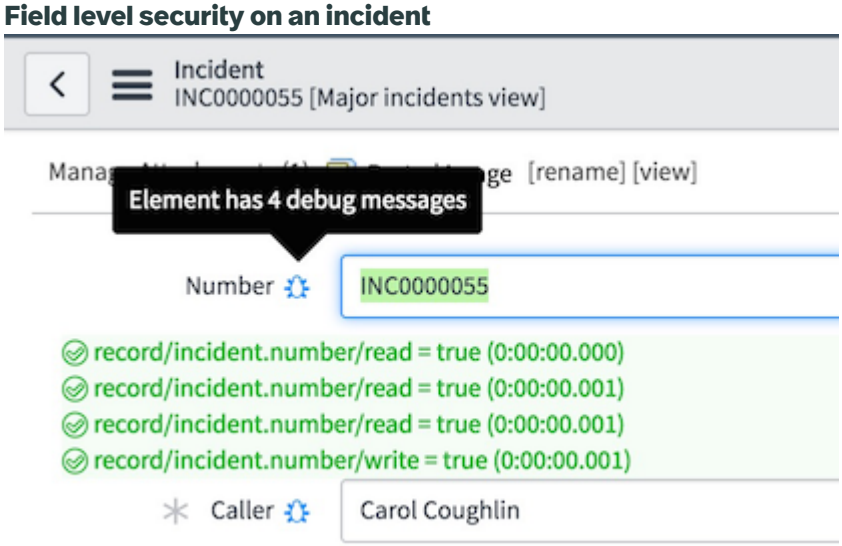

After enabling ACL debugging, you can impersonate another user to see what ACL rules the user passes and fails. When you impersonate a user, you can only see what that user is allowed to see. For example, you cannot view a record that an ACL prevents the user from seeing. To make debugging easier, read-only access to certain ACL-related tables is enabled by default, even when impersonating a user that does not have read access to the tables. To change this functionality, set the following property to **false**.

To enable ACL rule debugging, navigate to **System Security** > **Debug Security Rules**.

System property	Description	Default setting
glide.security.access_allow_read_impersonator	Allows read access to the following tables while impersonating a user: sys_security_acl, sys_security_operation, sys_security_type, and sys_user_role. As a result, the impersonating user can read data that the impersonated user cannot read.	true **Note:** When the property is set to false, the impersonated user might be prevented from reading ACL-related data. In this case, a second session logged in as admin or security_admin might be required to debug ACLs.

ACL rule output messages

ACL debugging displays ACL rule output messages at the bottom of each list and form. The output message displays the following:

Message element	Description
TIME	The total time used to process this ACL rule.
PATH	Information that uniquely identifies each ACL rule in the format: <ACL rule type>/<ACL rule name>/<Operation>.
CONTEXT	The object being evaluated by the ACL rule.
RC	The return code of the ACL rule. A true value passes the ACL rule. A false value fails the ACL rule.
RULE	A brief summary of processors and scripts, followed by ACL results for each table-level and field-level ACL evaluation. Most ACL evaluations show an overall pass or fail result followed by a breakdown of the results for each type of ACL criteria: • iAccessHandler: An internal system check using hidden source code on the platform. This is a system security check that you cannot modify. IAccessHandler can grant or deny access to a resource without evaluating ACLs. If IAccessHandler is ignored, then the ACLs are evaluated. You cannot modify the IAccessHandler checks in any way. For example, an IAccessHandler implementation is used for access checks on application resources and this cannot be changed. This is available starting with the Istanbul release. • Roles: Verification that the user has the correct role.

Message element	Description
	• Condition: Verification that the user passed the condition specified on the ACL rule (if any). • Script: Verification that the user passed the script specified on the ACL rule (if any).

The icons that appear show how the ACL was evaluated:

Icon	Description
A green checkmark ()	Indicates the table or field passed the criteria.
A red x icon ()	Indicates the table or field did not pass.
An empty gray circle icon ()	Indicates the ACL evaluation did not need to be performed.
A blue checkmark, x, or empty circle	Indicates that the ACL was taken from a cached result of a previous ACL check. The icons mean the same as the above.

You can perform these actions on the ACL debug output:

- Select or clear these check boxes at the top of the debug output:
  - **Security rules**: Show or hide the results of the ACL checks.
  - **Others**: Show or hide other warnings or messages.
- Click the name of the ACL next to any of the output messages to open that ACL record.

- Hover the cursor over any of the icons for the four ACL checks to see more information.

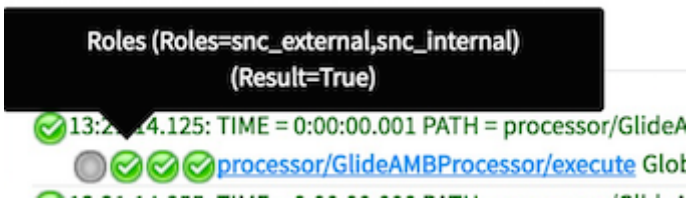

ACL troubleshooting reference

ACL troubleshooting includes identifying ACL rule errors and use the debugging tools to fix the ACL related problems.

Access analyzer

Access analyzer helps the administrators to view permissions for the selected user, role, or group. It is a diagnostic security tool that provides comprehensive visibility into resource permissions and access controls at the Access Control List (ACL) level, enabling you to understand who has access to their resources, identify overly permissive configurations, and maintain least-privilege access principles. To learn more about how to use the tool, see Access analyzer.

Enable debugging

Enable debugging to help troubleshoot an issue.

Troubleshoot

Error or symptom	Solution
You cannot access records from a custom table.	Create a table ACL rule for the custom table granting users access to the table. Without an explicit table ACL rule, users must pass the permissions in the table wildcard (*) ACL rule, which by default restricts access to administrators only. Enable debugging and determine what ACL rules are evaluated for the custom table.
You create a custom ACL rule that does not work properly.	The most likely problems are that another rule takes precedence over your custom rule in the processing order or that the user does not meet all the permission requirements for the object type. Enable debugging and verify that the ACL rule is being evaluated.
Your field ACL rule does not work properly.	There is likely a table ACL rule that the user has not met. Enable debugging and determine what ACL rules are evaluated for the field. Verify that there is not a conflicting table ACL rule or duplicate field ACL rule.
Your table ACL rule does not work properly.	There is either an ACL rule higher in the processing order or a duplicate table ACL rule interfering with the table ACL rule. Enable debugging and determine what ACL rules are evaluated for the table.
You can see a field in a list but not in form.	It is possible that the ACL rule conditions or script are being triggered in the list but not in the form. Enable debugging and determine when the ACL rules evaluate to true. Update the conditions or script to have the same behavior on the list and form.
You receive an error message when trying to execute a processor or client-callable script include.	There is an ACL rule for the processor or client-callable script include that the user has not met. If the user should have access to the object, enable debugging and determine what ACL rules are evaluated for the processor or script include. Update the ACL rule or the user roles as needed to access the object.

ACL configuration watcher

The ACL configuration watcher lets you know what related ACLs exist on a table when you insert, update, or delete an ACL on the same table.

The ACL configuration watcher is an interceptor window that displays every time you make important changes on the Access Control [sys_security_acl] table. It displays a security rules summary window where you can view ACLs related to the one you are modifying. You cannot modify any ACLs from the security rules window. To make any modifications, close the watcher window and go to those ACLs.

The ACL configuration watcher does not appear in the following situations:

- If you save or update an ACL record without actually making any changes.
- If you make minor updates (not an insert or delete), such as updating scripts, conditions, and the admin-overrides option.
- If the ACL record is not active.

ACL Security Rules window

The configuration watcher shows the ACL execution plan. The execution plan is displayed in the security rules pop-up window. You can view this kind of information:

ACL configuration window elements

Item	Description
red highlight	An ACL that is deleted or deactivated.
blue highlight	An ACL that is modified.
green highlight	An ACL that is added or becomes active.
Masked	An ACL that was effective until you made a change.
Unmasked	An ACL that was just made effective when you made a change.

Configuration watcher example

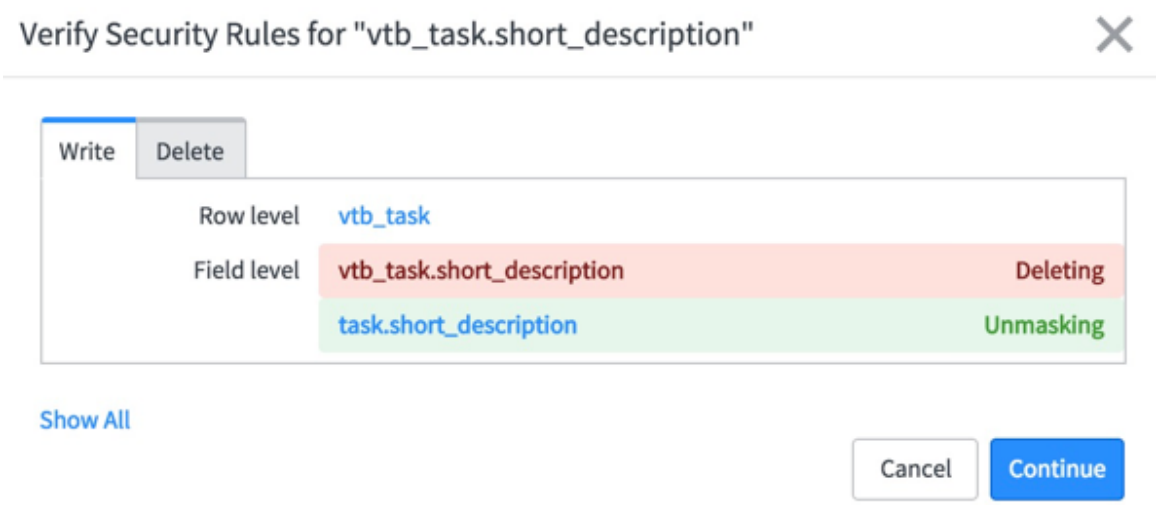

Show ACL execution plan

Administrators can view how ACLs relate to each other by viewing an execution plan for any ACL in the instance.

Before you begin
Role required: security_admin

Procedure

1. Elevate to a privileged role.
2. Navigate to **System Security** > **Access Control (ACL)**.
3. Open an ACL with the type *record*.

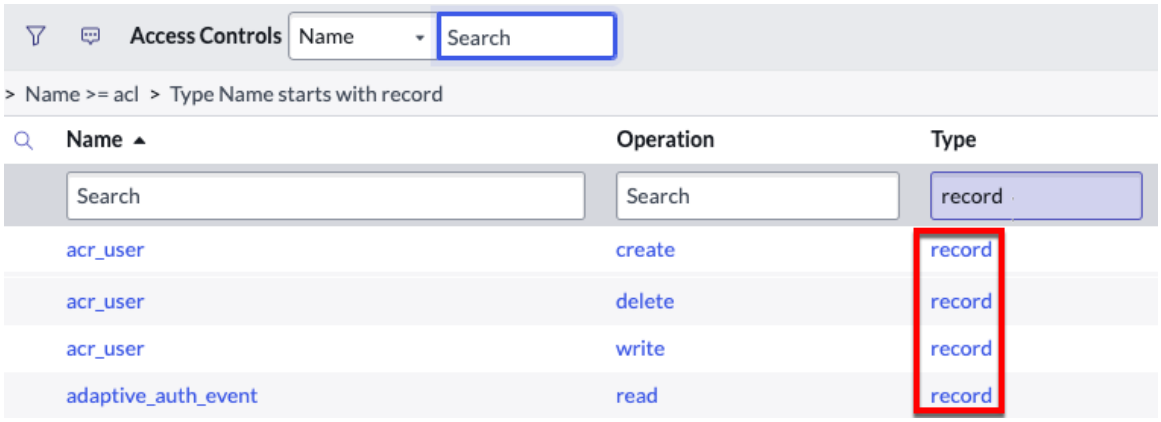

4. Click **Show ACL Execution Plan** in the Related Links section.

   The security rules window appears for the ACL. The example displays the execution plan for "vtb_task".

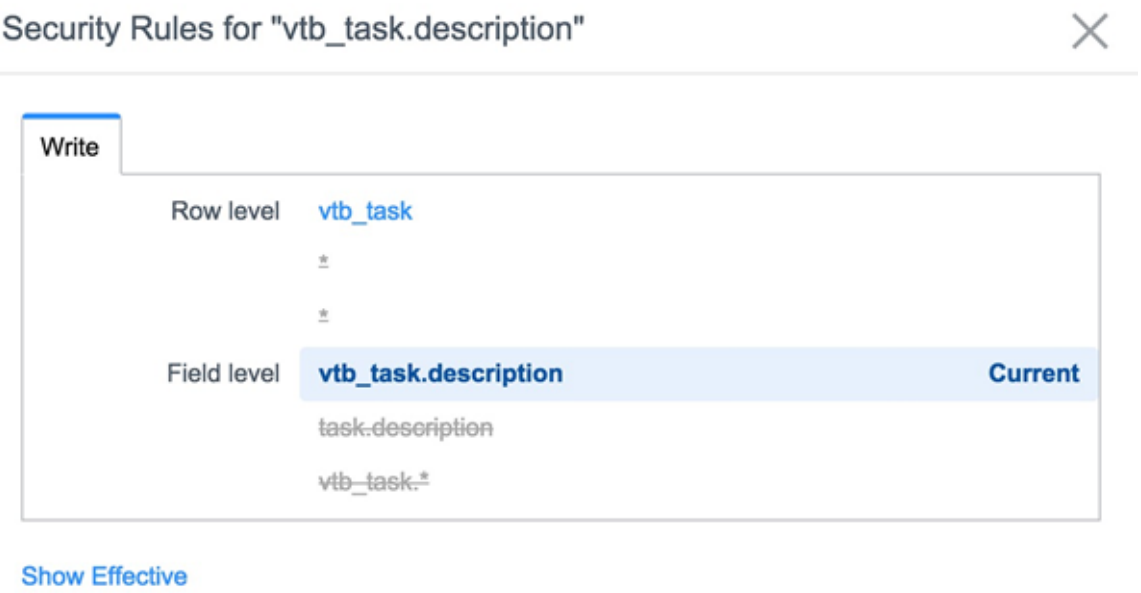

ACL execution plan window

UI item	Description
Title	The name of the ACL.
Tab name	If the ACL is create, read, write, or delete.
Row level	Row-level ACLs that run on this table.
Field level	Field-level ACLs that run only on this field (or column in the table).

5. Click **Show all** to show all related ACLs, including those ACLs that are overridden and generic ACLs that apply to all records.
   Overridden ACLs have a line through the name, and generic ACLs have the wildcard character asterisk (*) for the name.
6. Show only the immediate ACLs related to the one you are viewing and hide the ACLs on tables from which the ACL table is extended and the generic wildcard (*) ACLs by clicking **Show Effective**.

Use the ACL configuration watcher

Use the ACL configuration watcher after you elevate to security_admin role.

Before you begin

Role required: security_admin
Elevate to a privileged role

Procedure

1. Open an ACL that is a record-type ACL.
2. Perform an action on the ACL, such as modifying it, or selecting an option from the context menu like **Insert**.
3. If you modified any values on the Access Control form, right-click the header and select **Save** or click **Update** or **Delete**.

   The Security Rules window appears. The system did not yet perform the database action on the ACL, so the changes are not yet saved.

   These are examples of security rules on the Visual Task Board application's Private Task [vtb_task] table. See ACL configuration watcher for a description of the items on this window.

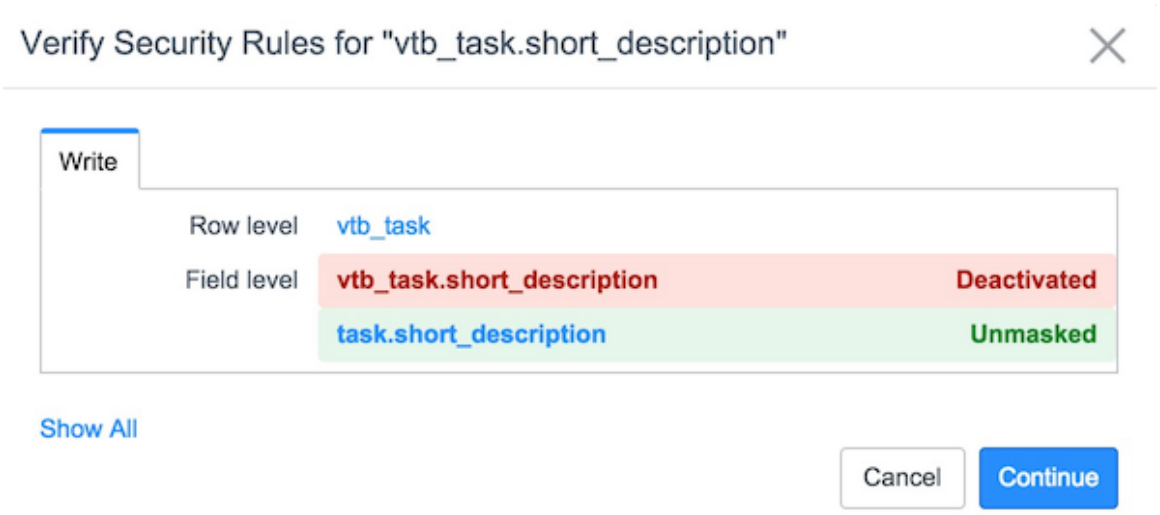

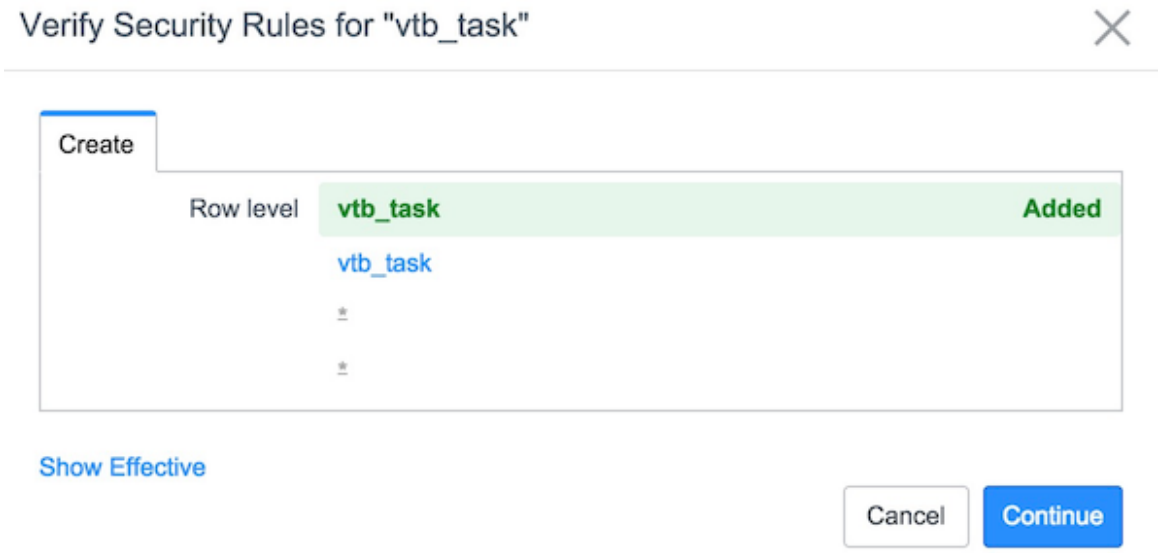

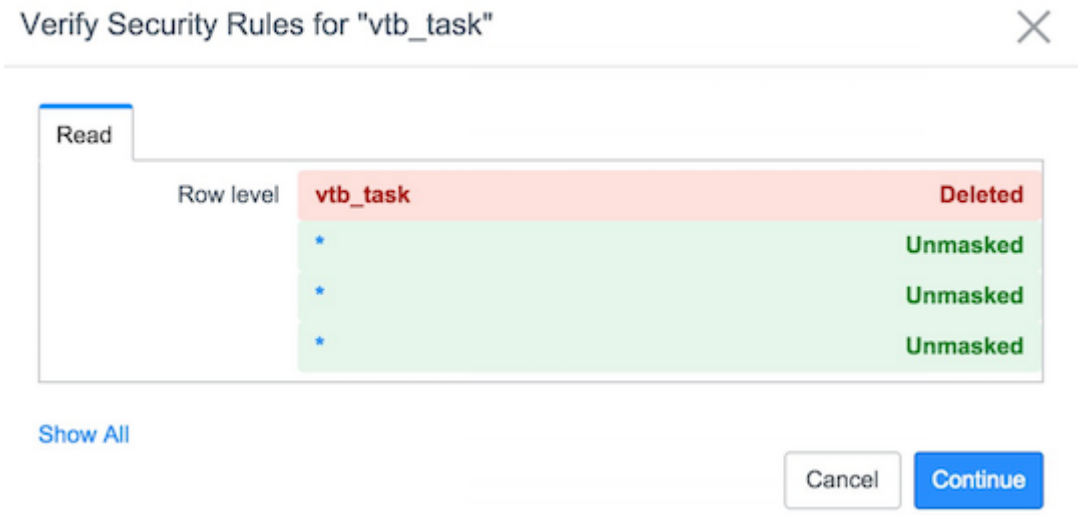

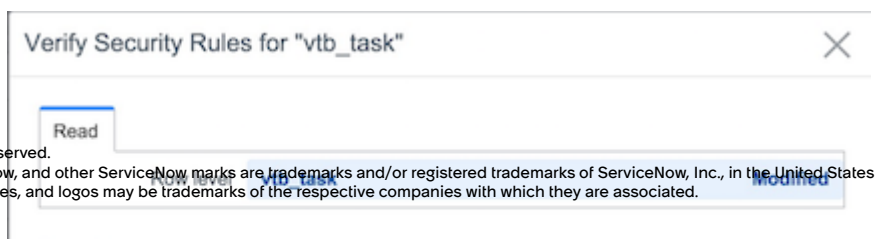

4. Just as with the execution plan, you can click **Show all** to show all related ACLs, including those that are overridden and generic ACLs that apply to all records, or click **Show Effective** to show only the immediate ACLs related to the one you are viewing.
5. Hover your mouse over any of the ACLs to see a description.

Access analyzer

ServiceNow® Access analyzer is an access diagnostic tool that helps to view the permissions of an identity for a resource.

Explore  Learn the features and business value of Access analyzer.	Use  Understand how to use Access analyzer.
Permission Evaluation  Know how the permissions are evaluated.	Frequently asked questions  Get details about commonly asked questions about Access analyzer.

Explore Access analyzer

Analyze identities on the ServiceNow instance.

ServiceNow Access analyzer is an application that helps the administrators to view permissions for the selected user, role, or group.

Note:

- Access analyzer is a ServiceNow Store product. Visit the ServiceNow Store website to view all the available apps and for information about submitting requests to the store.
- Access analyzer impersonates the identity record to retrieve details about the permissions and doesn't read or store any personal or sensitive data of the identity.
- Access analyzer evaluation results are the same irrespective of any access policies defined for the users such as Zero trust access (ZTA). The policies are only evaluated during the actual user login and aren't evaluated during the access analyzer flow.
- Access analyzer has limitations in accurately evaluating access of the resources related to managed scope resources and delegated developer.

Evaluate Access

Evaluate Access is a capability of the ServiceNow Access Analyzer, which helps the administrators to view permissions for the selected user, role, or group.

It enables you to analyze and view the permissions of users, groups, roles for a table, client callable script includes, UI pages, and REST endpoints.

Using Access Analyzer, organizations can improve their security posture, identity governance, risk management, achieve their compliance goals, and understand who (identity) has access to what (resources).

Compare Access

Compare Access is a capability of the ServiceNow Access analyzer V2, which enables administrators to compare user access and determine the right level of access for the users on your ServiceNow instance.

Compare Access can be perform between the users for the user records and access control.

Compare Access enables you to perform the following analysis:

- Level 1: Compare user records to understand the attributes, roles, and groups.
- Level 2: Compare access controls to run the root cause analysis by finding access issues.

Benefits

The following are some of the benefits of using the Access analyzer:

- Analyze access to resources (tables).
- Compare the access of 2 users.
- Compare the roles and groups of 2 users.
- Generate a report showing whether an identity has access to a resource (table).
- Understand who has access for critical security hygiene.
- Help to prevent from over-provisioning permissions.
- Achieve the least privilege principals when implementing access controls.
- Limit access to certain data, which includes applications, tables, rows or columns, and other resources.
- Provide reporting capabilities for the analyzer results.
- Compare access between user records and access controls.
- Determine the right level of access for users on your ServiceNow instance.

Use Access analyzer

Analyze identities and their access on the ServiceNow® instance.

Before you begin
Role required: admin, access_analyzer_admin

The following procedure describes the steps for accessing Access analyzer and using various features within Access analyzer.

Note: Access Analyzer is a ServiceNow® Store product.

Procedure

1. Navigate to **All** > **Access Analyzer** > **Analyze Permissions**.
   The Analyze access and permissions homepage is displayed.

   Access Analyzer has the following capabilities:

   - Evaluate Access
   - Compare user records
   - Compare access controls
2. Select the respective tab to use Access analyzer based on your requirement.

Use Evaluate access

Analyze identities on the ServiceNow® instance.

Before you begin

Role required: admin, access_analyzer_admin

The following procedure describes the steps for accessing Evaluate Access in the Access Analyzer and using its various features.

Note: Access analyzer is a ServiceNow Store product.

Procedure

1. Navigate to **All** > **Access Analyzer** > **Analyze Permissions**.
   The Analyze access and permissions homepage is displayed.
2. Select the **Evaluate access** tab.
3. Select your criteria as follows:

Select your criteria for analyzing access and permissions

Field	Description
Analyze by *	Analyze access for a user, a role, or a group
Select user *	Specify a user name to select from the list.
Rule type *	Analyze access for a table, a UI page, a REST Endpoint, or a client callable script include.
Select table *	Specify a table name to select from the list.
Select record	Specify a record name to select from the list.
Select field	Specify a field name to select from the list.

4. Specify the description in the **Description** field.
5. Select **Analyze permissions**.

   The access results for the user are displayed. Similarly you can analyze the permissions of a Group or Role for the following rule types:

   - Table (record)
   - Client callable scripts include
   - REST endpoints

The access results are displayed.

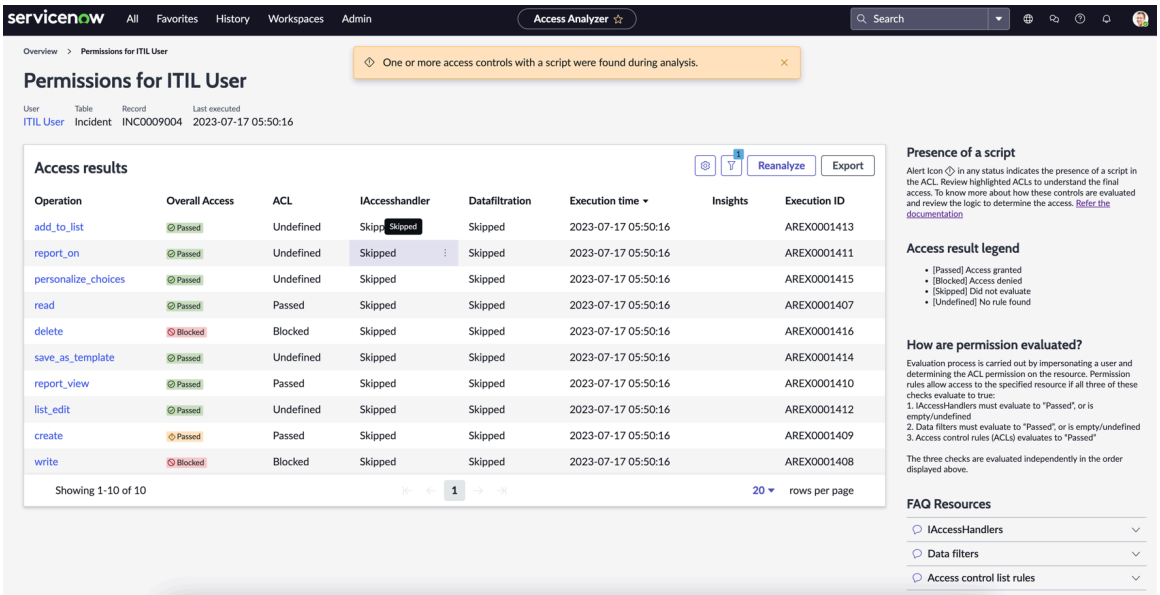

The Access results table includes the following fields:

Access results

Fields	Description
Operation	The type of operation that the user, group, or role can perform for the selected table, record, or field.
Overall Access	Result of the overall access. The results are as follows: ○ [Passed] Access granted ○ [Blocked] Access denied ○ [Skipped] Didn't evaluate ○ [Undefined] No rule found
ACL	Whether an ACL is defined for the selected operation.
Access Handler	An internal system check using hidden source code on the platform. IAccessHandler can grant or deny access to a resource without evaluating ACLs. If IAccessHandler is ignored, then the ACLs are evaluated.
Data filtration	A data filter is a form of access control designed to work along with the existing Access Control rules (ACLs) on your instance.
Execution time	The time at which the access results were executed.
Insights	More information about the selected operation.
Execution ID	A unique ID for each access result execution.

6. Select the Operation for more information about the ACL.
   For example, if you select **read**, the access control related to read is displayed.

ACL Details

Field	Description
Name	Name of the ACL.
Decision Type	Decision type configured for the ACL. **Allow access** or **Deny access**.
Applies to condition	Whether the ACL is applied to a condition.
ACL Applies to	Details about the resource the ACL is applied.
Status	Status of the ACL or Access result.
Required ACL Roles	Details of the role that is required to access the resource.
Role	Status of the role. Passed, Skipped, or Blocked.

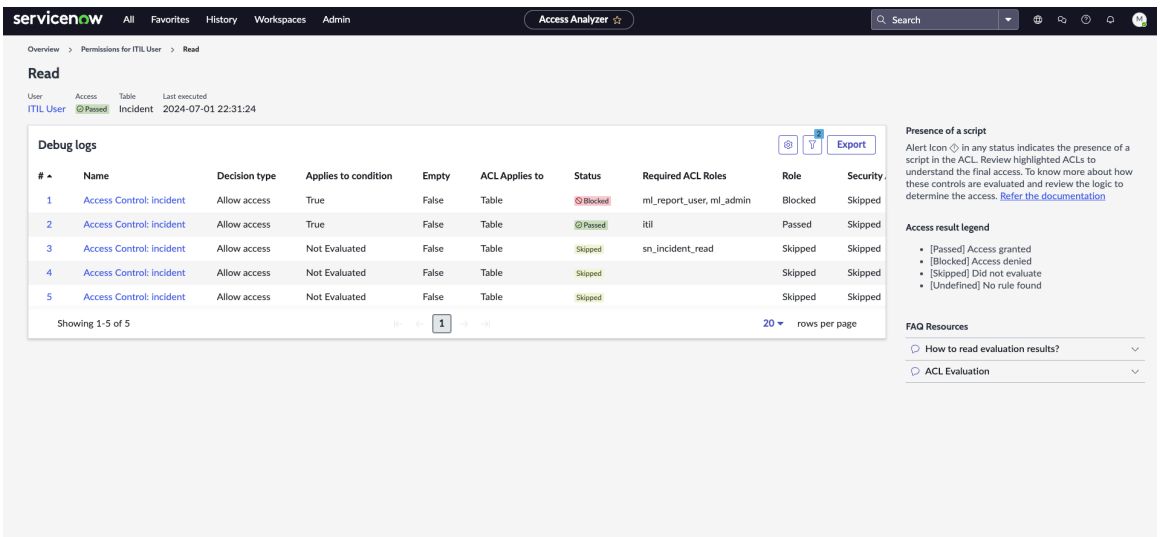

View permissions for a user

Use Access Analyzer to view permissions for a selected user.

Before you begin

Role required: admin, access_analyzer_admin

The following procedure describes the sample steps to view permissions for a selected user (**ITIL User**) to view the permissions of the **Incident** table using the evaluate access in Access Analyzer.

Note: Access Analyzer is a ServiceNow® Store product.

Procedure

1. Navigate to **All** > **Access Analyzer** > **Analyze Permissions**.
   The Analyze access and permissions homepage is displayed.
2. Select your criteria as follows:

Select your criteria for analyzing access and permissions

Field	Description
Analyze by *	Select **User**.
Select user *	Specify a user name to select from the list. In this example, **ITIL User**.
Rule type *	Analyze access for a **Table**, **UI page**, **REST Endpoint**, **client callable script include**, **AI agent**, or **Agentic workflow**. In the example, `Table`.
Select table *	Specify a table name to select from the list. In this example, **Incident**.
Select record	Specify a record name to select from the list. In this example, **INC0000001**.
Select field	Specify a field name to select from the list. This field can be used to analyze permission even at a field level. For example, Active, Created By, and so on.

3. Specify the description in the **Description** field.
4. Select **Analyze permissions**.

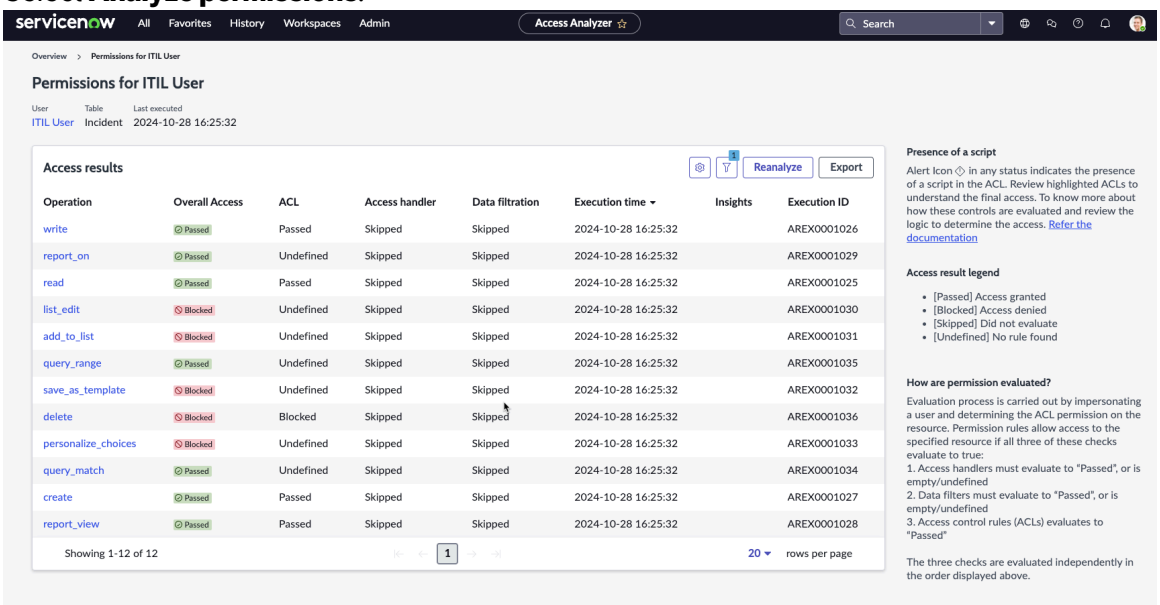

The **Access results** for the **ITIL User** is displayed.

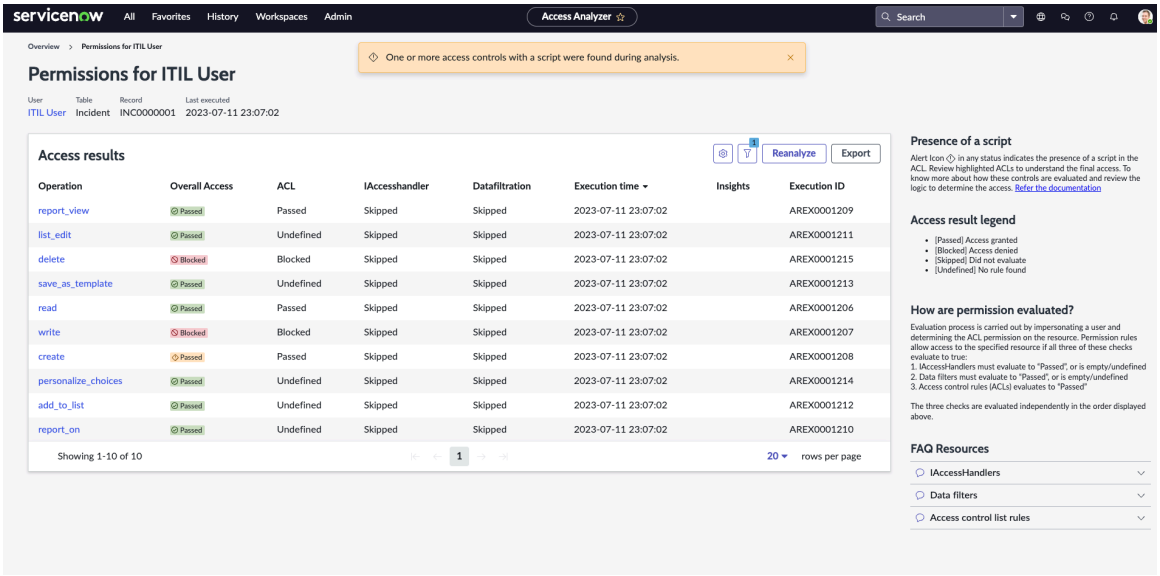

The results can be read, by referring to the Legends, access control list (ACL), IAccesshandler, and Data filters.

Let's take the example of **read** operation. For the **ITIL User** overall access is Passed, which means the user is able to read the record with the correct permissions (ACL).

Similarly for **create** operation, the overall access is passed with an alert icon, which means that there could be a presence of script for the ACL evaluation.

Note: In the example, **write** and **delete** operations are blocked for the selected user and the user can't edit or delete the selected record (INC0000001).

5. Select read operation to know more about the Debug logs.

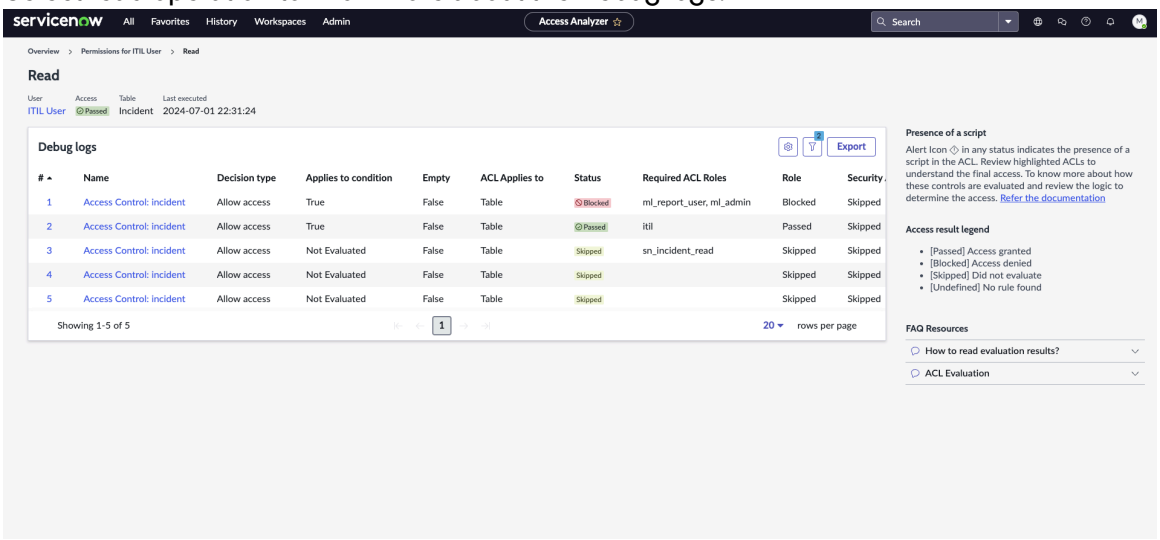

The Debug logs page displayed the business rule and associated ACLs that are required to perform the **read** operation for the record.

The Debug logs shows that there's a business rule and 4 ACLs associated for the read operation.

There's a status **Passed** for one of the ACL, which means to read the selected record, the user has the required ACL and can read the record. Since one of the ACL is **Passed**, the other ACL evaluations are **Skipped**.

6. Select the Access Control that is Passed to know the details of the ACL.

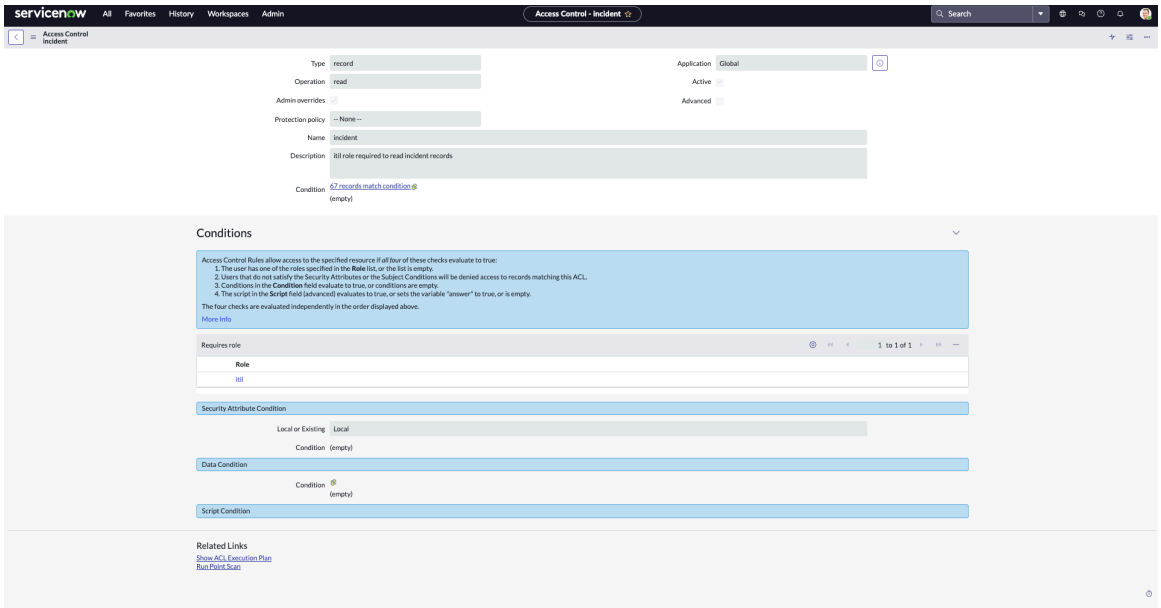

The details of the Access Control for the selected ACL are displayed.

For a selected operation with the **Passed** and a presence of a script. The Access Control page displays the associated script for the record.

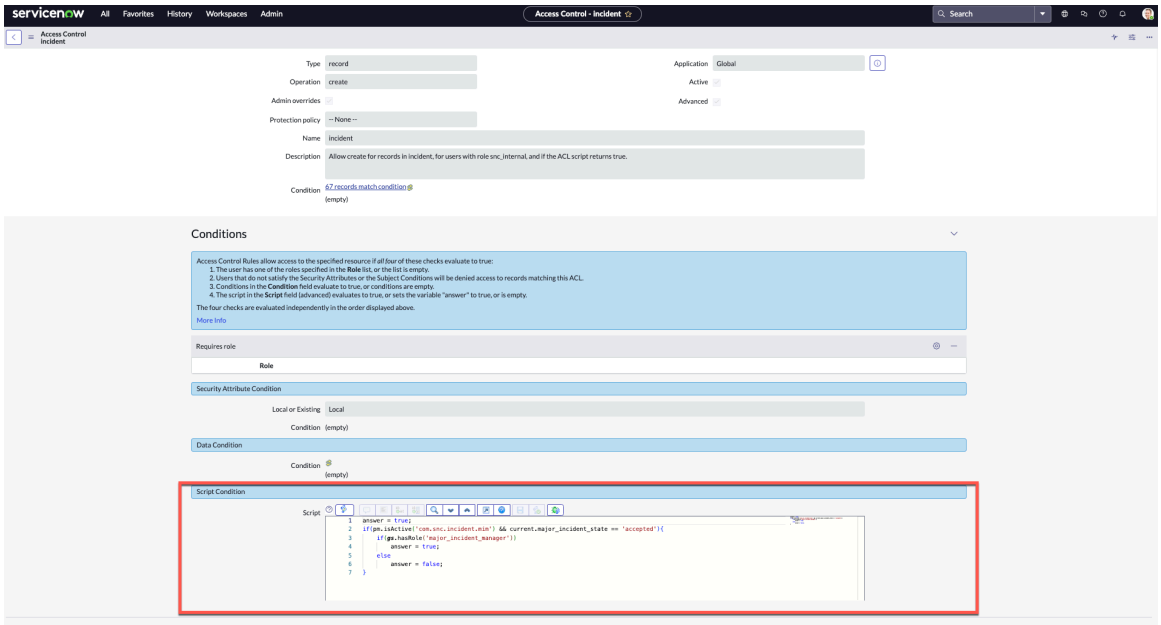

View permissions for a role

Use Access Analyzer to view permissions for a selected role.

Before you begin

Role required: admin, access_analyzer_admin

The following procedure describes the sample steps to view permissions for a selected role (**user_admin**) to view the permission for a REST API Endpoint using the evaluate access in Access Analyzer.

Note: Access Analyzer is a ServiceNow® Store product.

Procedure

1. Navigate to **All** > **Access Analyzer** > **Analyze Permissions**.
   The Analyze access and permissions homepage is displayed.
2. Select your criteria as follows:

Select your criteria for analyzing access and permissions

Field	Description
Analyze by *	Select **Role**.
Select user *	Specify a role to select from the list. For example, **user_admin**.
Rule type *	Analyze access for a table, a UI page, a REST Endpoint, or a client callable script include. For example, **REST Endpoint**.
REST endpoint*	Specify a REST endpoint. For example, `/api/global/user_role_inheritance`. **Note:** The complete REST endpoint must be used when using the selected field.
REST endpoint method *	Specify a REST endpoint method. For example, GET.

3. Specify the description in the **Description** field.
4. Select **Analyze permissions**.

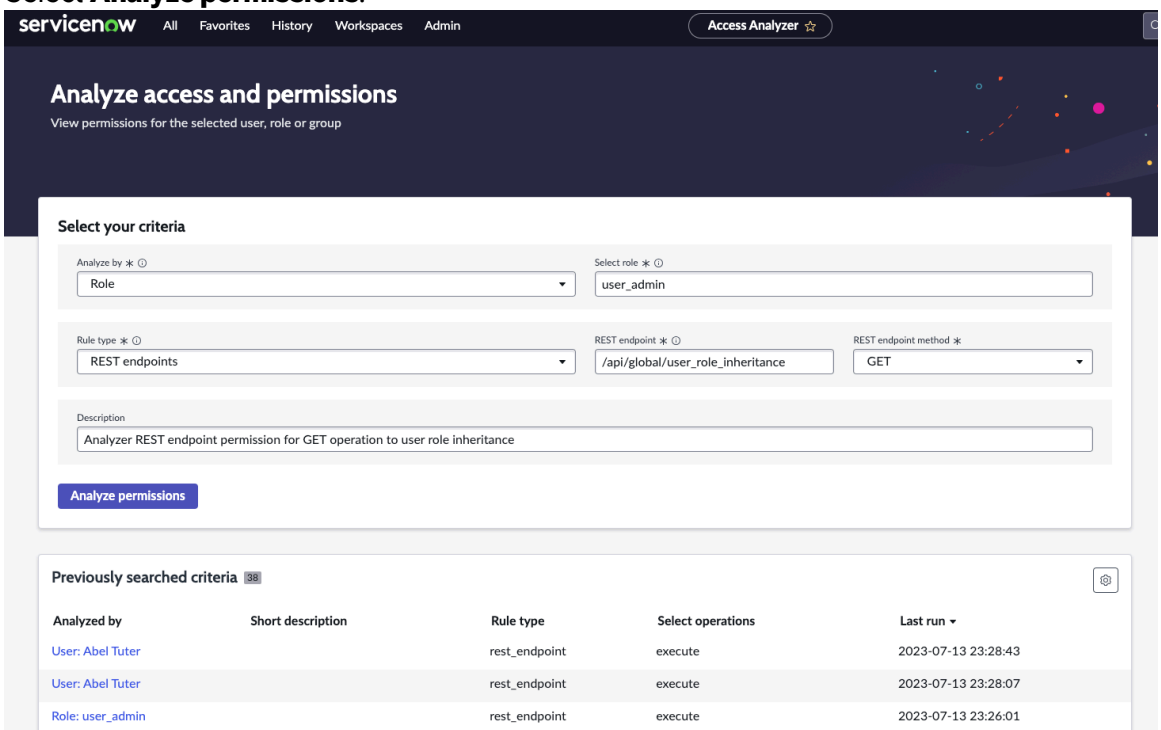

The **Access results** for the **user_admin** role is displayed.

The results can be read, by referring to the Legends, access control list (ACL), IAccesshandler, and Data filters.

The overall access for the role is passed, which means that the role (**user_admin**) is able to access the **REST endpoint** for the selected **GET** method.

View permissions for a group

Use Access Analyzer to view permissions for a selected group.

Before you begin

Role required: admin, access_analyzer_admin

The following procedure describes the sample steps to view permissions for a selected group (**Incident Management**) to view the permissions of an **Incident UI** page using the evaluate access in Access Analyzer.

Note: Access Analyzer is a ServiceNow® Store product.

Procedure

1. Navigate to **All** > **Access Analyzer** > **Analyze Permissions**.
   The Analyze access and permissions homepage is displayed.
2. Select your criteria as follows:

Select your criteria for analyzing access and permissions

Field	Description
Analyze by *	Select **Group**.
Select user *	Specify a user name to select from the list. For example, **Incident Management**.
Rule type *	Analyze access for a table, a UI page, a REST Endpoint, or a client callable script include. For example, **UI page**.
UI Page *	Specify the UI page. For example, **incident.do**.

3. Specify the description in the **Description** field.
4. Select **Analyze permissions**.

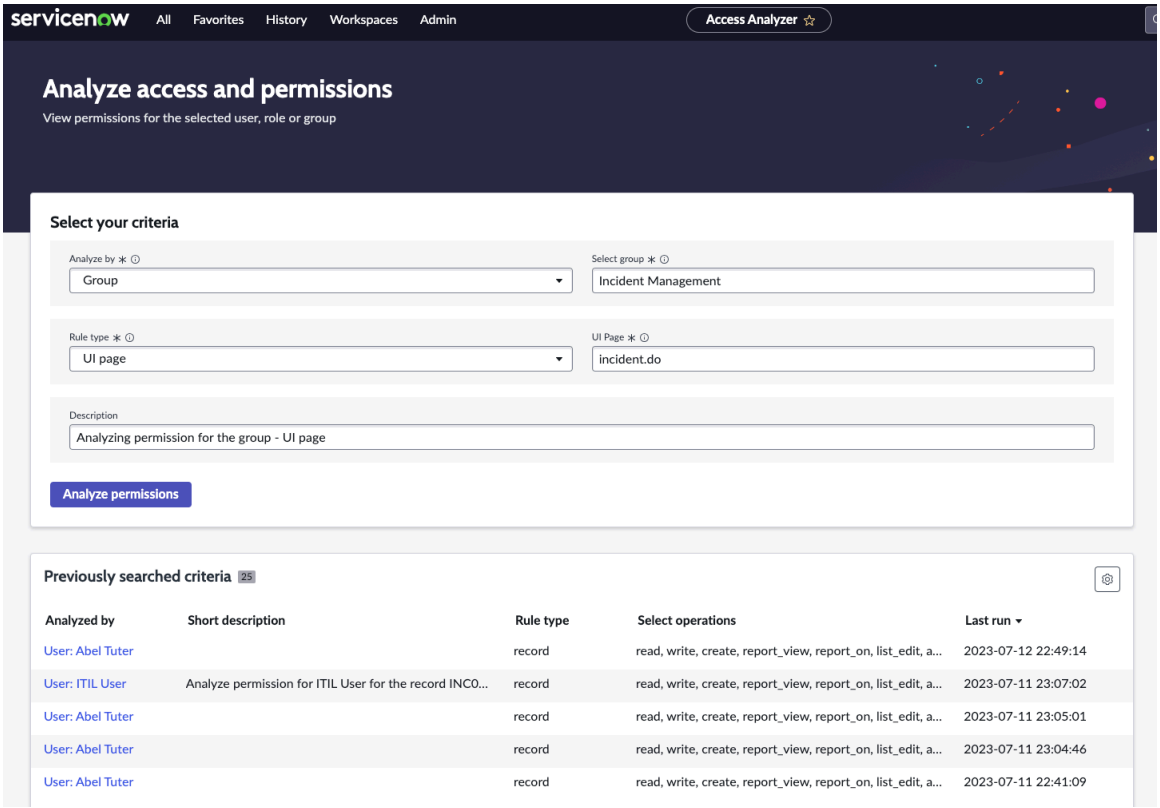

The **Access results** for the **Incident Management** group is displayed.

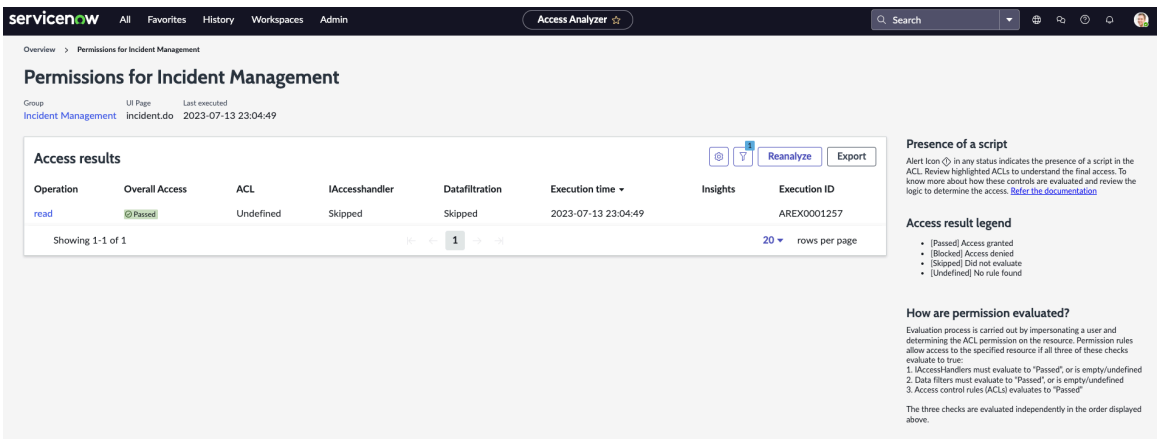

The results can be read, by referring to the Legends, access control list (ACL), IAccesshandler, and Data filters.

The overall access for the group is passed, which means that the users within the group (**Incident Management**) are able to access the Incident record.

Export Access Analyzer queries

Export the queries analyzed using Access Analyzer.

Before you begin

Role required: admin, access_analyzer_admin

The following procedure describes the steps for accessing Access Analyzer and using various features within Access Analyzer.

Note: Access Analyzer is a ServiceNow® Store product.

Procedure

1. Navigate to **All** > **Access Analyzer** > **Analyze Permissions**.
   The Analyze access and permissions homepage is displayed.
2. Select your criteria as follows:

Select your criteria for analyzing access and permissions

Field	Description
Analyze by *	Analyze access for a user, a role, or a group
Select user *	Specify a user name to select from the list.
Rule type *	Analyze access for a table, a UI page, a REST Endpoint, or a client callable script include.
Select table *	Specify a table name to select from the list.
Select record	Specify a record name to select from the list.
Select field	Specify a field name to select from the list.

3. Specify the description in the **Description** field.
4. Select **Analyze permissions**.

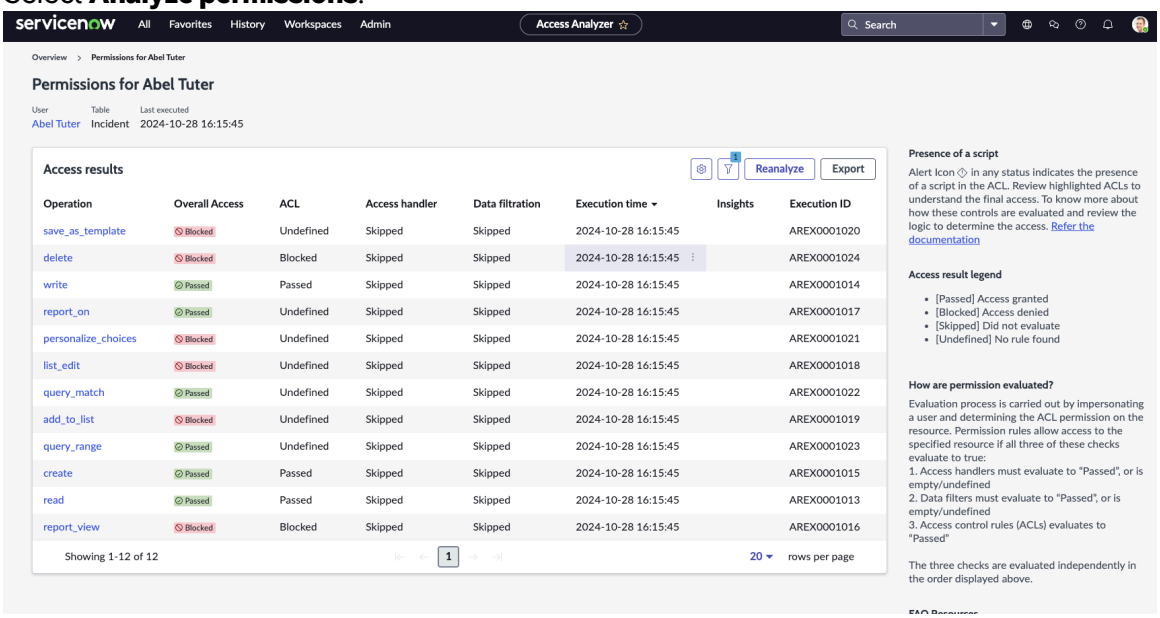

The access results for the user are displayed. Similarly you can analyze the permissions of a Group, Role for the following rule types:

- Table (record)
- Client callable script include
- REST endpoints

The **Access results** for the selected rule type is displayed.

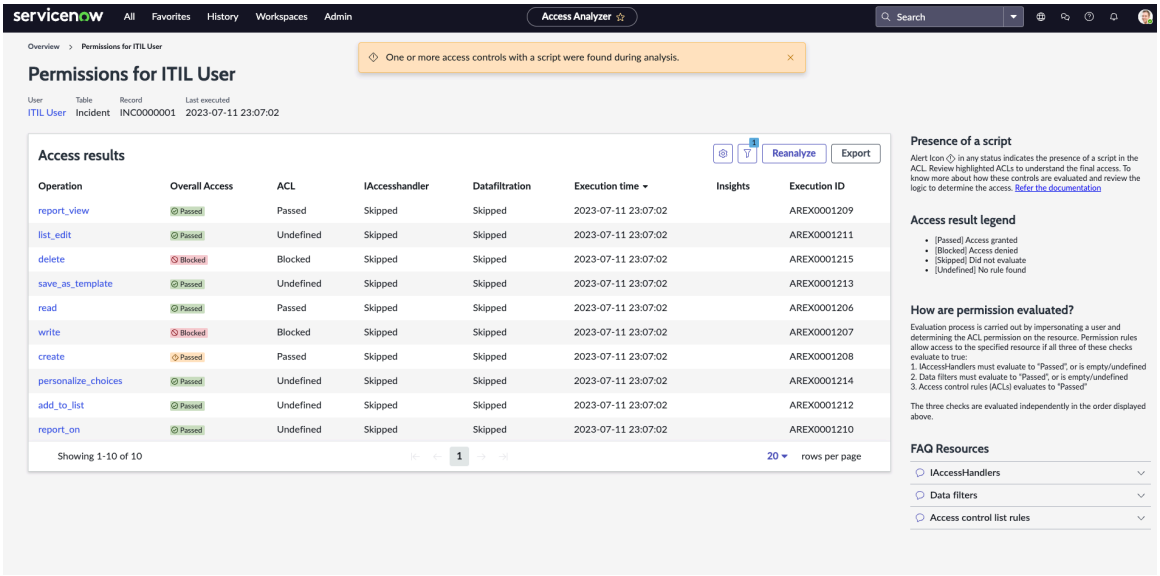

5. Click **Export**.

a. Choose the File Type.

Available file types are Excel, CSV, JSON, PDF.

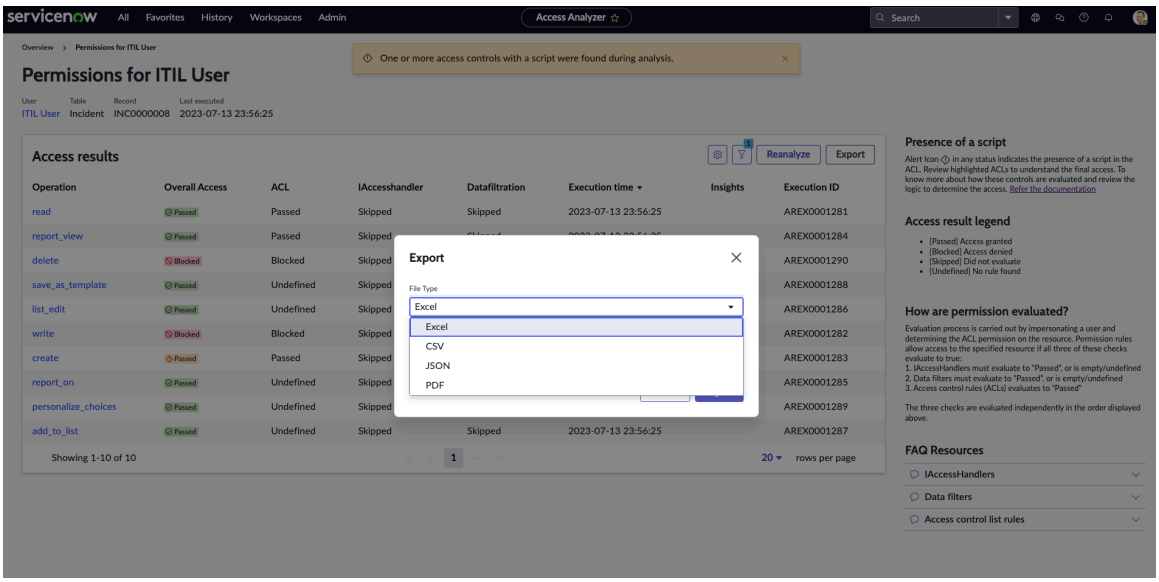

b. Choose the Delivery Type.

Available delivery types as Download and Email.

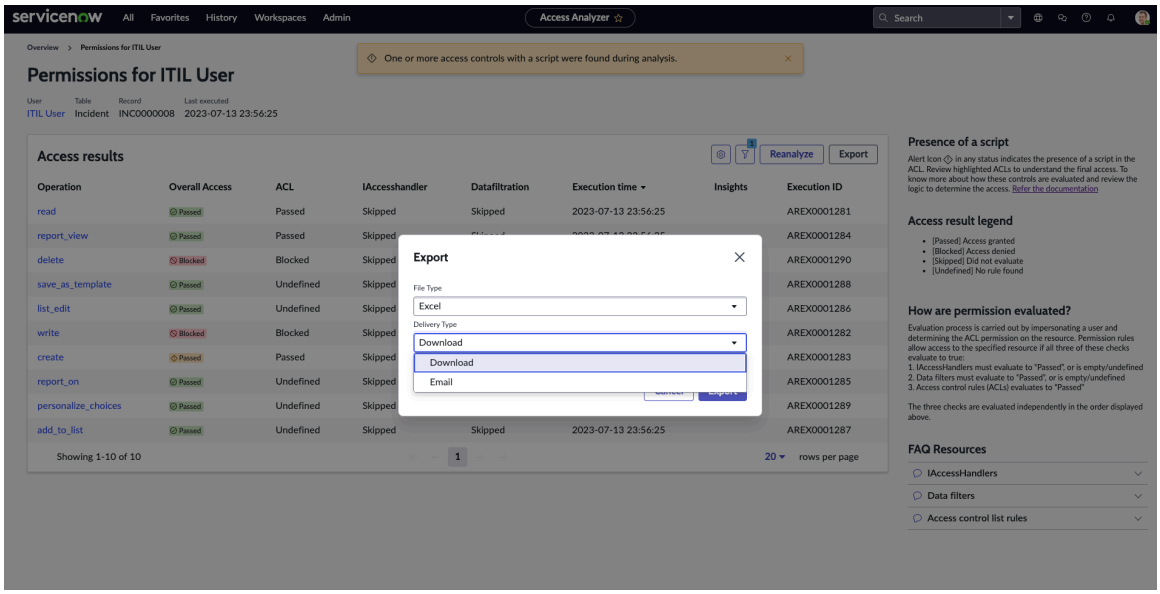

Compare user records

Compare user records to understand the access between two users.

Before you begin

Role required: admin, access_analyzer_admin

The following procedure describes the steps for comparing user records using the Access Analyzer.

Note: Access Analyzer is a ServiceNow® Store product.

Procedure

1. Navigate to **All** > **Access Analyzer** > **Analyze Permissions**.
   The Analyze access and permissions homepage is displayed.
2. Select the **Compare user records** tab.
3. Select **user 1** and **user 2** for comparison.

   For example, `ITIL User` as **user 1** and `Abel Tuter` as **user 2**.

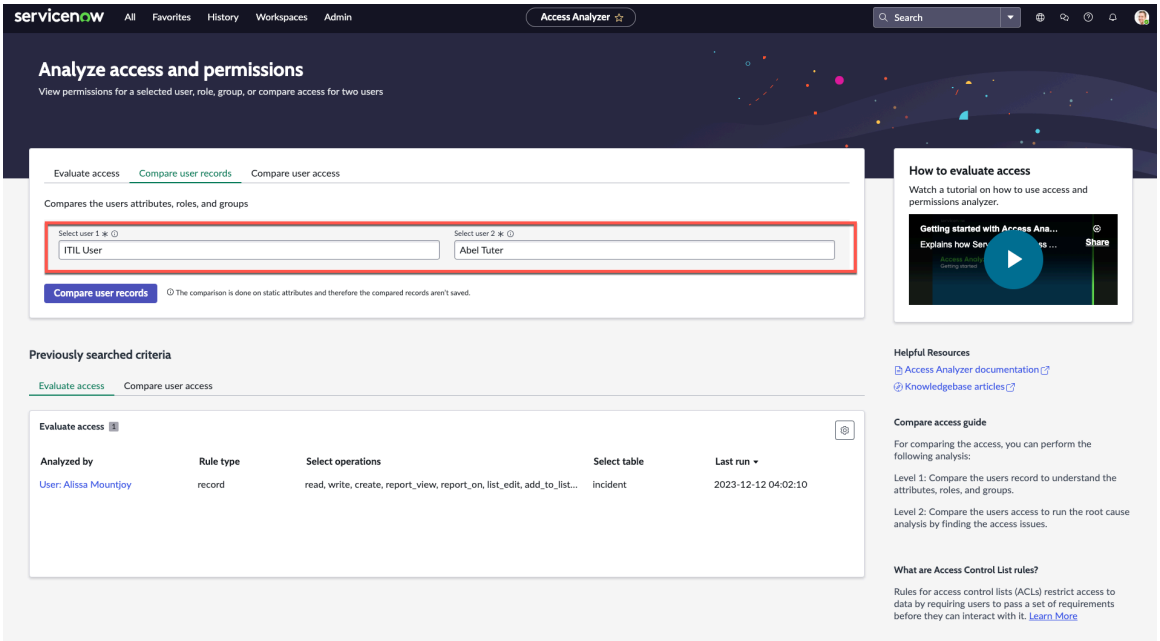

4. Select **Compare user records**.

   The results are displayed with the following tabs:

- **Details**: Display the user's

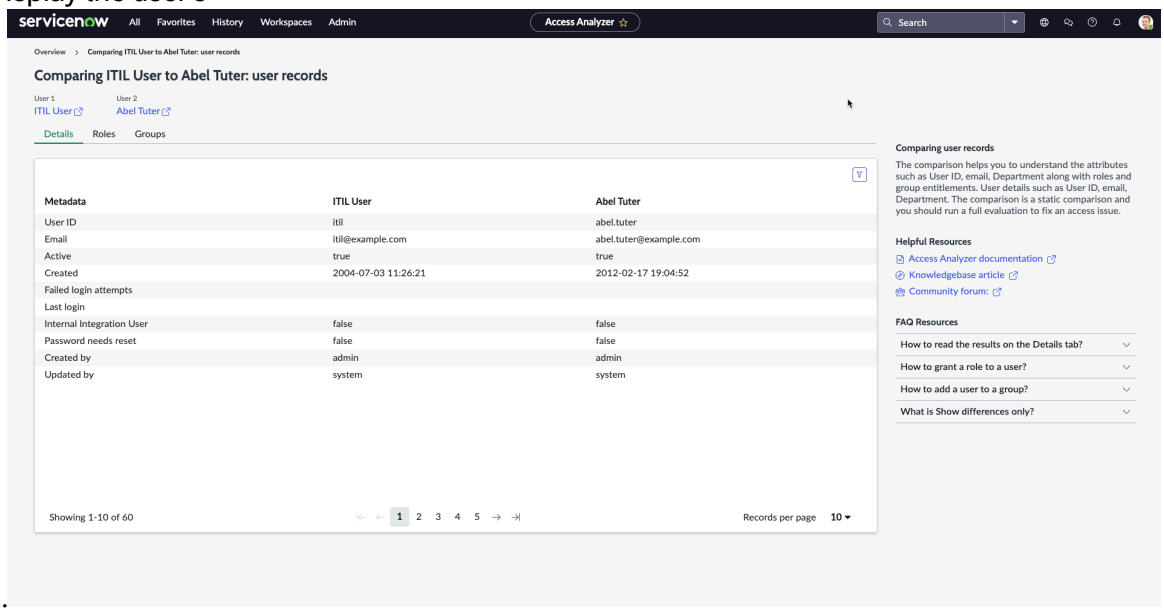

metadata.

- **Roles**: Display the roles that are assigned to the user. You can select the role to know more about the role and its entitlements.

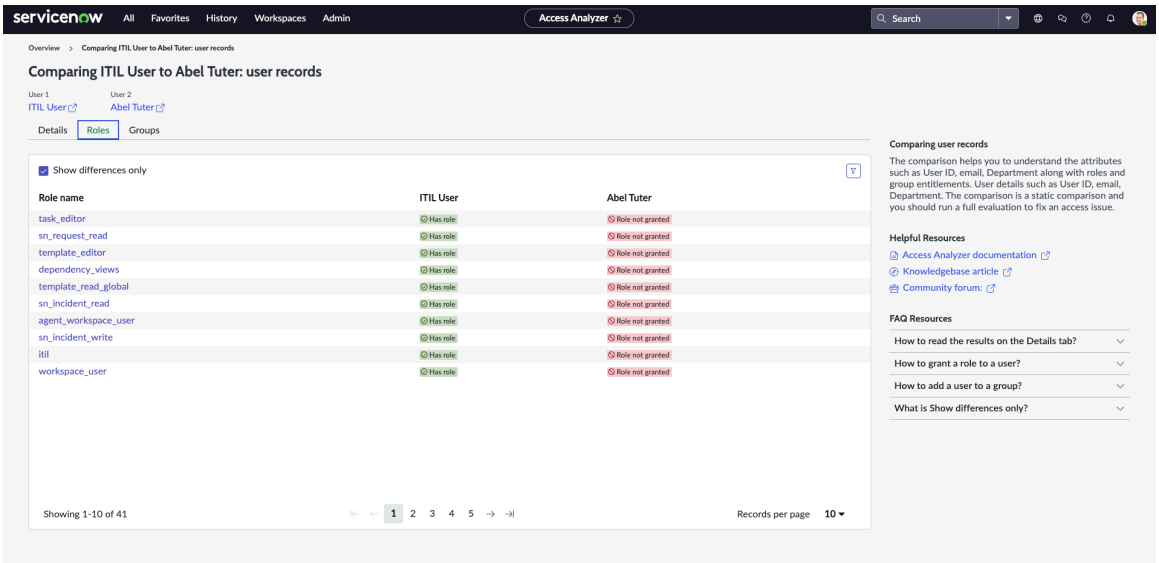

- **Groups**: Display the groups that are assigned to the user. You can select the group to know more about the group and its entitlements.

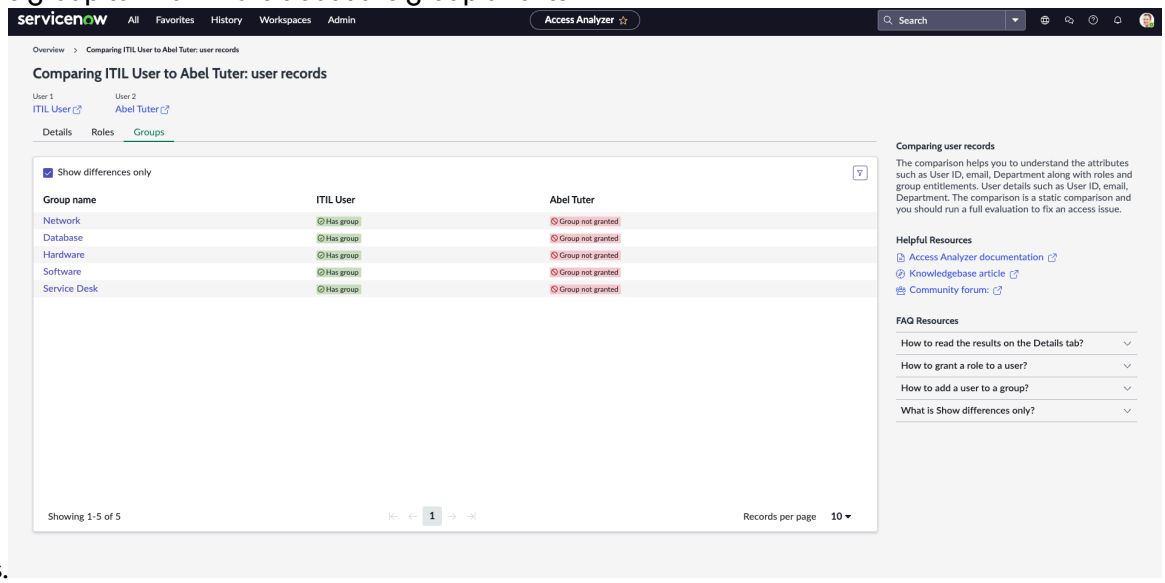

Similarly you can compare different users in the ServiceNow instance to understand the access that is assigned to the users.

Compare user access

Compare the user's access control using the Access Analyzer.

Before you begin

Role required: admin, access_analyzer_admin

The following procedure describes the steps for comparing the access control between the users using the Access Analyzer.

Note: Access Analyzer is a ServiceNow® Store product.

Procedure

1. Navigate to **All** > **Access Analyzer** > **Analyze Permissions**.
   The Analyze access and permissions homepage is displayed.
2. Select the **Compare user access** tab.
3. Fill in the following fields:

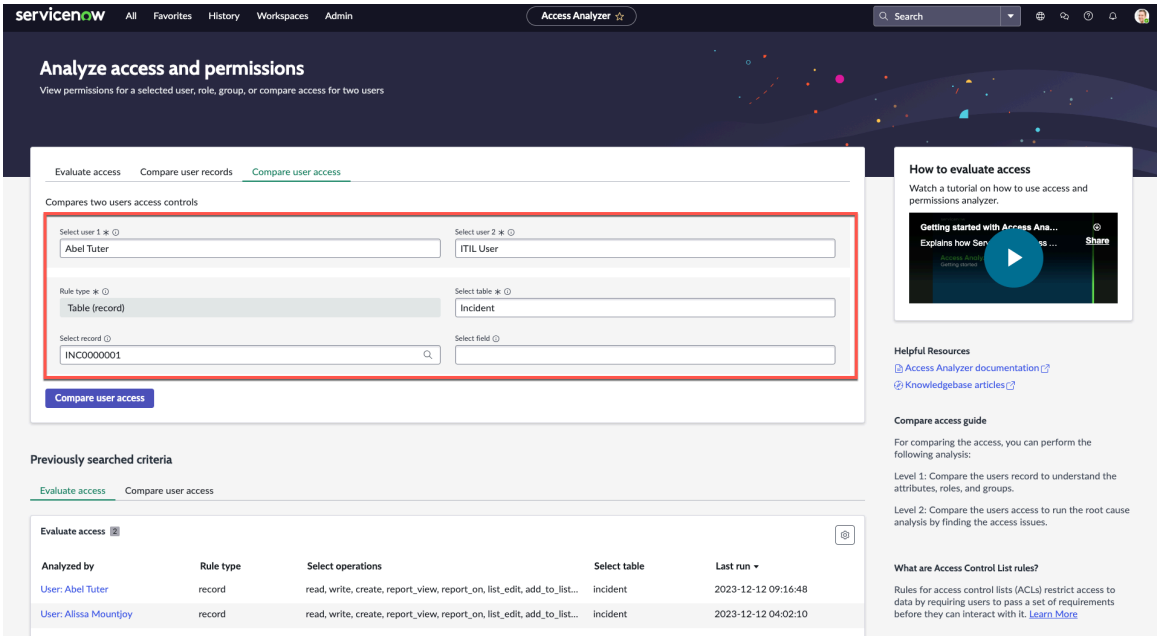

Compare user access

Field	Description
Select user 1*	Specify a user name to select from the list for the comparison.
Select user 2*	Specify a user name to select from the list to compare with the user 1.
Rule Type*	Analyze access permissions for a table. **Note:** Only access permissions for a table can be used in the **compare user access**.
Select table*	Specify a table name to select from the list.
Select record	Specify a record name to select from the list.
Select field	Specify a field name to select from the list.

4. Select **Compare user access**.

   The **compare user access** results for the selected users are displayed.

   The compare user access results show the operation and the access evaluation status for the users. For example, Abel Tuter and ITIL User.

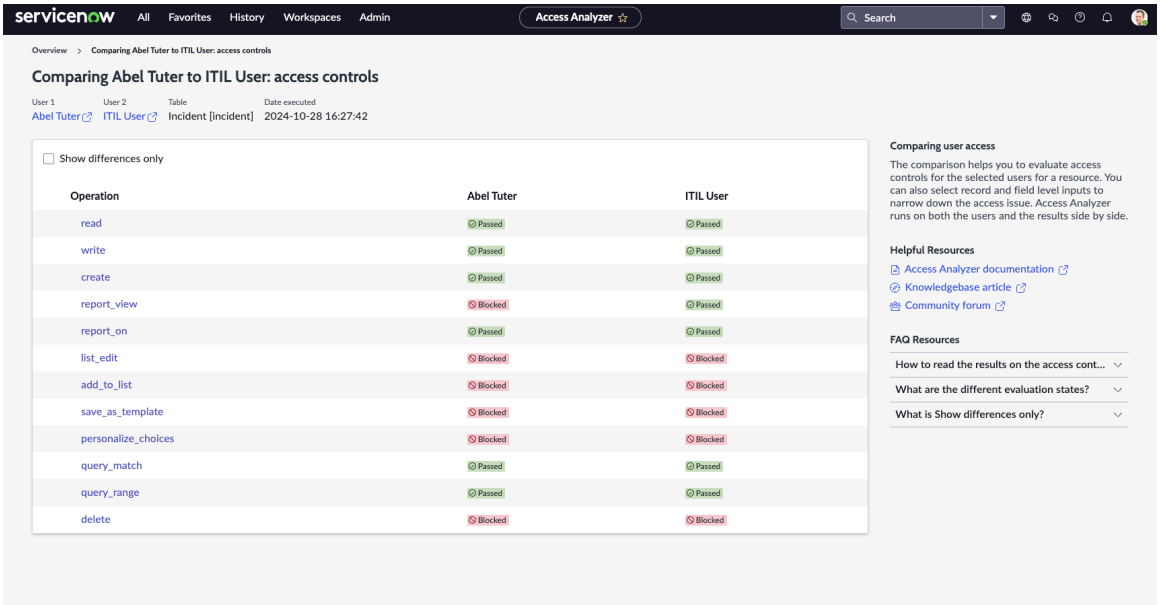

5. Select the Operation to know more about the permission evaluation and the roles the users are assigned to.
   For example, **read** operation.
6. Select any of the **Access Control** to know more about the access.

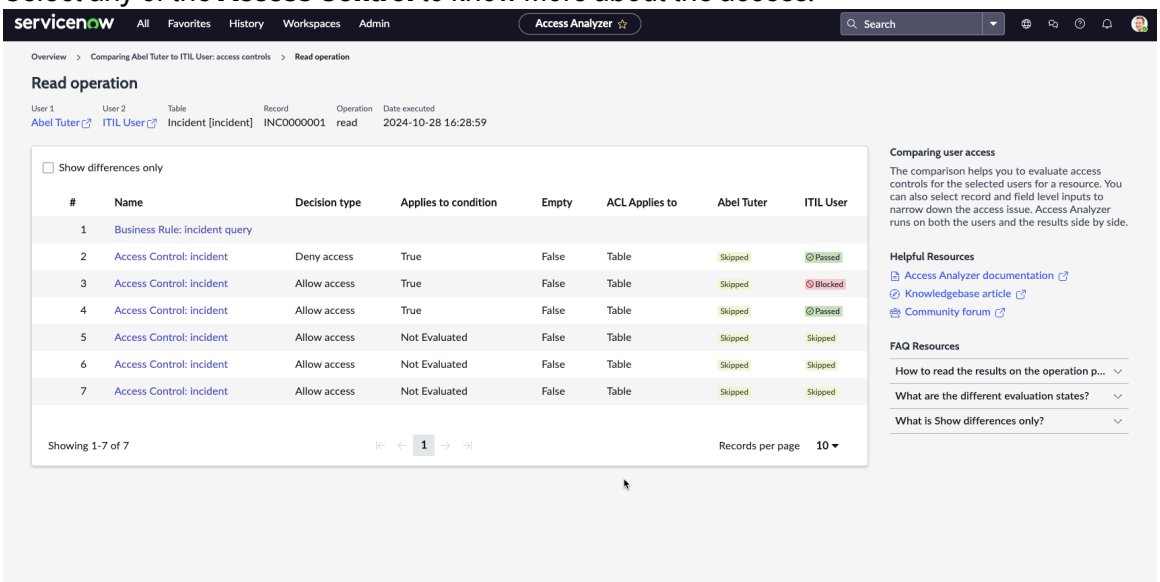

The Access Control details such as Roles, Security Attribute, Condition, and Script evaluation status are displayed.

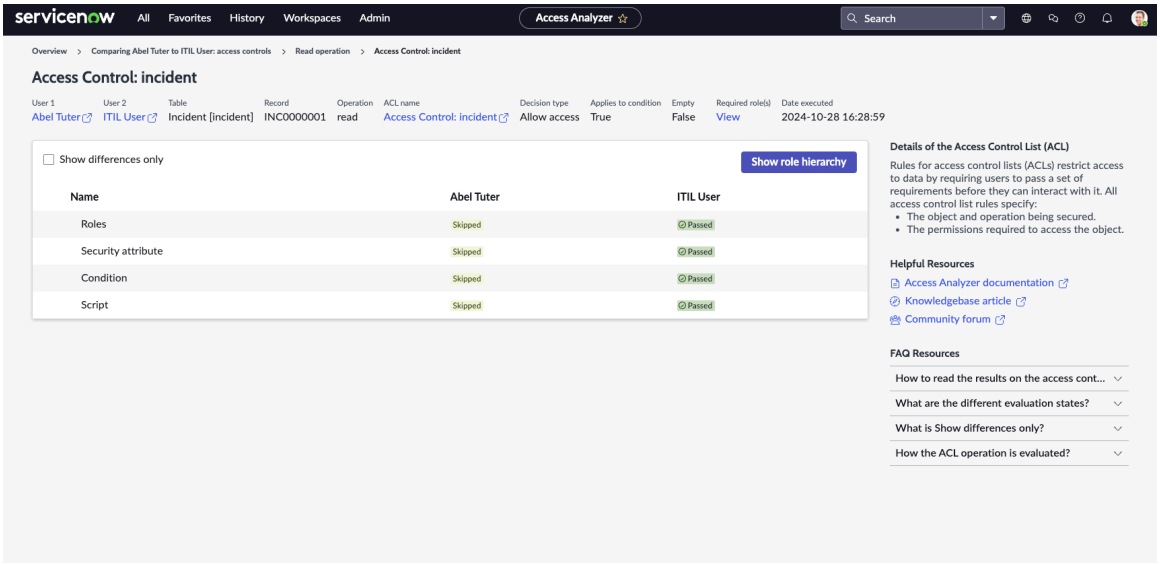

7. Select **Show role Hierarchy** to know more about the roles and groups the user is assigned and compare both the users.

   Based on the role hierarchy, you can assign the necessary role and group assignments to the user to have access to the resources (table).

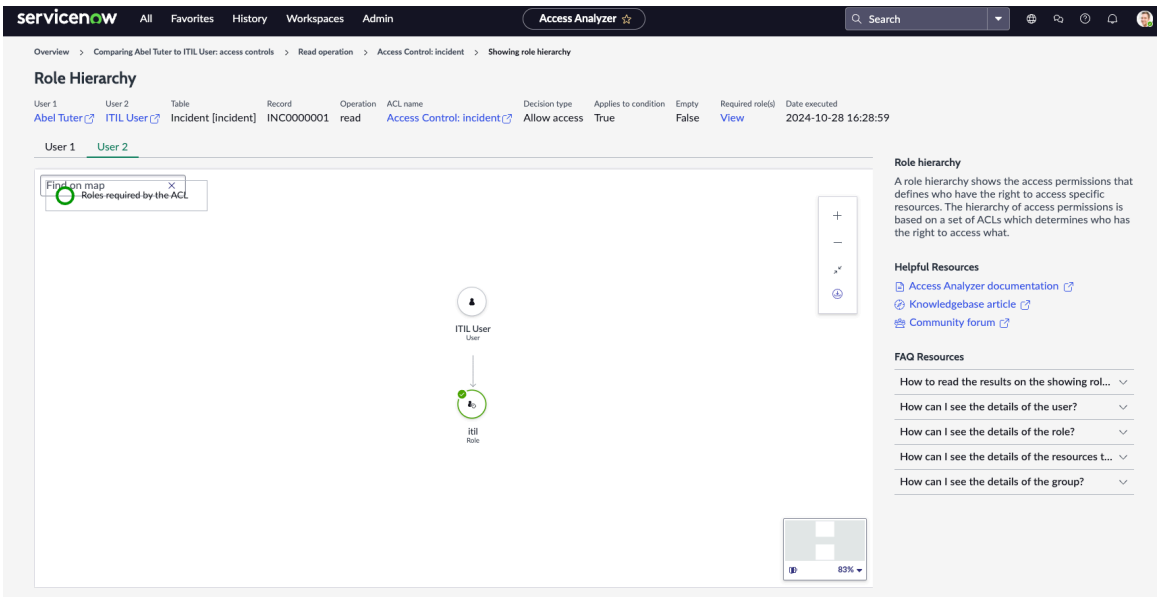

   In the example, **Abel Tuter** doesn't have `itil` assigned. You can determine the necessary role and group assignments to the user by looking into the role hierarchy.

   You can select the node to learn more about the role, resources the role can access, or group.

View the previously searched criteria in Access Analyzer

View the previously searched criteria in Access Analyzer.

Before you begin

Role required: admin, access_analyzer_admin

Note: Access Analyzer is a ServiceNow® Store product.

Procedure

1. Navigate to **All** > **Access Analyzer** > **Access Analyzer Queries**.

   Previously searched criteria has the following sections:

   - **Evaluate access**: Displays results based on queries made through the **Evaluate access** features.
   - **Compare user access**: Displays results based on queries made through the **Compare user access** features.

   Note: The previously searched criteria are not stored when using the compare user records feature.

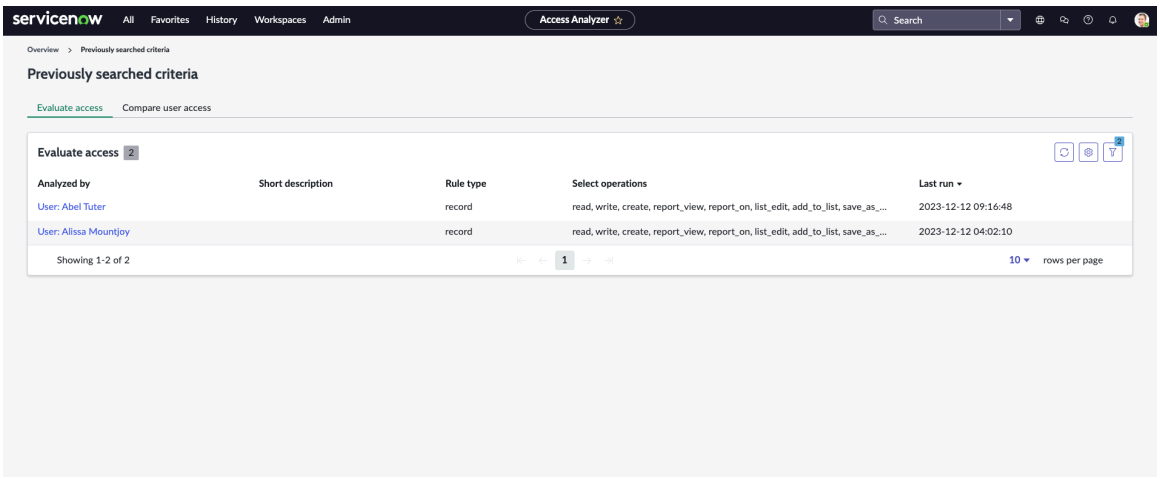

2. Select the **Analyzed by** links to view previously searched criteria from the Evaluate access section.

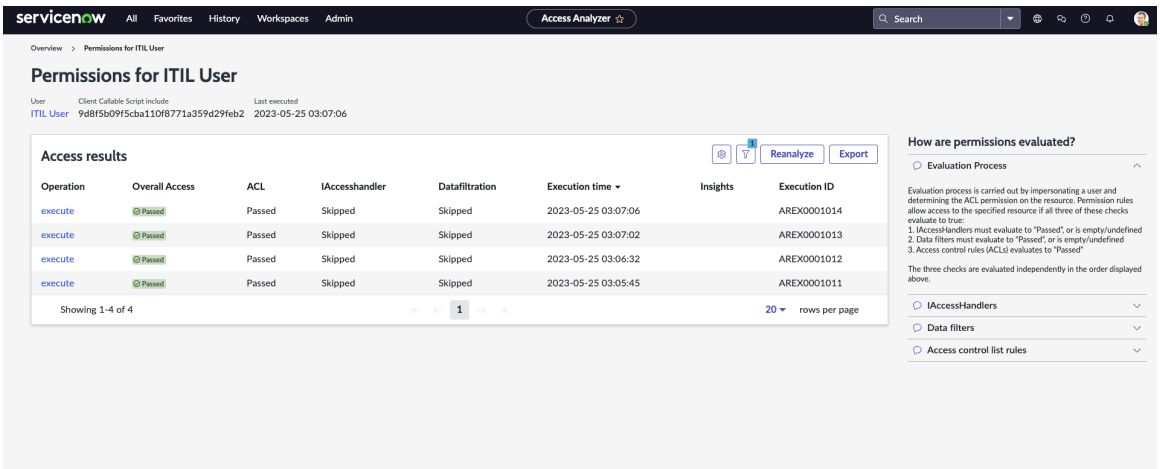

You can select **Reanalyze** the access for the user. Export the details using the **Export** option.

3. Select the **Compared users** links to view the previously searched criteria from the Compare user access section.

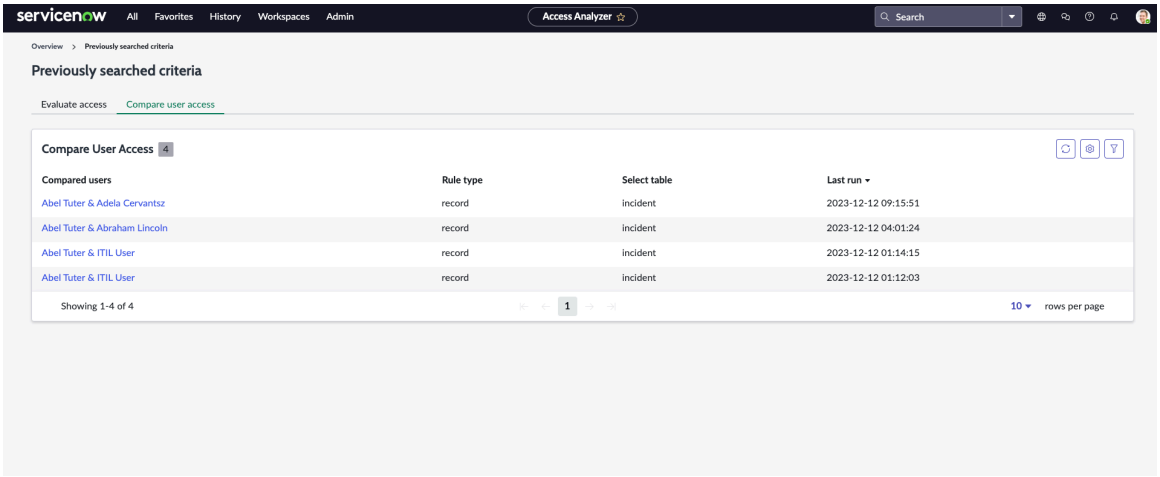

Permission evaluation

Permission evaluation criteria when using the Access analyzer.

Evaluation hierarchy

Permission for the selected user, group, or role is evaluated in the following hierarchy:

- Business rule: A business rule is a server-side script that runs when a record is displayed, inserted, updated, or deleted, or when a table is queried.
- Access Handler: An internal system check using hidden source code on the platform.
- Data Filtration: Data filter is a form of access control designed to work along with the existing Access Control rules (ACLs) on your instance. Data filter support only read operation.
- Access control list (ACL): Rules for access control lists (ACLs) restrict access to data by requiring users to pass a set of requirements before they can interact with it. Within an ACL, the following hierarchy is evaluated:
  - Role
  - Security Attribute
  - Condition
  - Script

You can analyze access and permissions for the selected user, role, or group using the Access Analyzer. The permissions are evaluated based on the following rule types:

- **Table Level Evaluation**: Role and security attribute ACLs are used for Table level evaluation.
- **Record or Field level Evaluation**: Role, security attribute, condition, and script level ACLs are used for Record or Field level evaluation.
- **UI page**: Support Only ready operations. Only read level ACLs are evaluated.
- **REST Endpoint**: Support only execute operation. Only execute level ACL are evaluated.

Details about the important fields in the Access Results are as follows:

- Presence of a script
- Access result legend

- Evaluation process
- IAccessHandlers
- Data filters
- Access control list rules

Presence of a script

Alert Icon in any status indicates the presence of a script in the ACL. Review highlighted ACLs to understand the final access. To know more about how these controls are evaluated and review the logic to determine the access, see Access Analyzer Debug logs.

Legend in Access Analyzer

When Analyzing the access and permissions, legends are displayed as part of the evaluation process. Following are the legends:

- [Passed] Access granted
- [Blocked] Access denied
- [Skipped] Did not evaluate
- [Undefined] No rule found

Evaluation process

Evaluation process is carried out by impersonating a user and determining the access control list (ACL) permission on the resource. Permission rules enable access to the specified resource if the following checks are evaluated to true:

- IAccessHandlers must evaluate to “Passed”, or is empty or undefined
- Data filters must evaluate to “Passed”, or is empty or undefined
- Access control rules (ACLs) evaluate to “Passed”

IAccessHandlers

An internal system check using hidden source code on the platform. IAccessHandler can grant or deny access to a resource without evaluating ACLs. If IAccessHandler is ignored, then the ACLs are evaluated.

You can’t change the IAccessHandler checks. For example, an IAccessHandler implementation is used for access checks on application resources such as read access.

Data filter

Data filter is a form of access control designed to work along with the existing Access Control rules (ACLs) on your instance.

Access control list rules

Rules for access control lists (ACLs) restrict access to data by requiring users to pass a set of requirements before they can interact with it.

Frequently Asked Questions

Frequently asked questions while using the Access analyzer.

Evaluate Access

The following are some of the frequently asked questions while using the Evaluate Access feature in the Access Analyzer:

Frequently asked questions

Questions	Explanation
How to read the evaluation results displayed by the Access Analyzer?	Each row represents an individual access control list (ACL). The sequence (#) in the results the order in which ACLs are evaluated. The status shows whether overall access is granted (passed) or denied (blocked).
How are ACLs Evaluated?	At a table level, ACLs are evaluated only for roles and security attributes, conditions and scripts aren't evaluated. Roles are evaluated first. If Roles are blocked, conditions and scripts are skipped. For more information, see Configure an ACL rule.
What are the legends in Access Analyzer?	When Analyzing the access and permissions, legends are displayed as part of the evaluation process. The following are the legends: • [Passed] Access granted • [Blocked] Access denied • [Skipped] Didn't evaluate • [Undefined] No rule found
What is the Alert icon in the Access results mean?	Alert Icon in any status indicates the presence of a script in the ACL. Review highlighted ACLs to understand the final access. To know more about how these controls are evaluated and review the logic to determine the access, see Access Analyzer Debug logs.
What is **IAccesshandler**?	An internal system check using hidden source code on the platform. It's a system security check that you can't modify. **IAccessHandler** can grant or deny access to a resource without evaluating ACLs. If this **IAccessHandler** is ignored, then the ACLs are evaluated. You can't modify the **IAccessHandler** checks in any way. For example, an **IAccessHandler** implementation is used for access checks on application resources such as read-only access.

Frequently asked questions (continued)

Questions	Explanation
What are data filters?	Data filters are a form of access control designed to work along with the existing Access Control rules (ACLs) on your instance.
What is an ACL rule?	Rules for access control lists (ACLs) restrict access to data by requiring users to pass a set of requirements before they can interact with it.

Time limited role assignments found for the user due to which the results may be impacted. You can review the time-limited roles assigned for the user here.

Compare user records

The following are some of the frequently asked questions while using the Compare user record feature in the Access Analyzer:

Frequently asked questions

Questions	Explanation
How to read the results on the Details tab?	The Details tab displays the metadata associated to the user 1 and user 2
How to grant a role to a user?	From the Users tab, you can check the role that must be granted for the user and assign that role.
How to add a user to a group?	From the Groups tab, you can check the group the user must be added and add the user to the group.
What is Show difference only?	When you enable the Show differences only check box, only the roles or group that are different between the user 1 and user 2 is displayed.

Compare user access

The following are some of the frequently asked questions while using the Compare user access feature in the Access Analyzer:

Frequently asked questions

Questions	Explanation
How to read the results on the access control comparison page?	The access control comparison page displays the evaluation states for different ACL operations.
What are the different evaluation states?	When comparing access controls between the users, following are the different evaluation states:

Frequently asked questions (continued)

Questions	Explanation
	• Passed • Blocked
What is show differences only?	When you enable the Show differences only check-box, only the operation evaluation state that are different between the user 1 and user 2 is displayed.
How the ACL operation is evaluated?	Access control list (ACL) is the rule for access control lists (ACLs) that restrict access to data by requiring users to pass a set of requirements before they can interact with it. Within an ACL, the following hierarchy is evaluated: • Role • Security Attribute • Condition • Script
How to read the results on the showing role hierarchy page?	The showing role hierarchy page displays the role that is assigned for user 1 and user 2. You can understand the role that is required for the user for a particular ACL operation.
How can I see the details of the user?	You can select the **User (node)** > **More actions** > **View user** details to know more about the user.
How can I see the details of the role?	You can select the **Role (node)** > **More actions** > **View role** details to know more about the role.
How can I see the details of the resources the role can access?	You can select the **Role (node)** > **More actions** > **View all resources that the role has access** to know the resources the role can access.
How can I see the details of the group?	You can select the **Group (node)** > **More actions** > **View group** details to know more about the group.

Access Analyzer Debug logs

Debug logs display the details of the select access result operation.

Fields in Debug logs

The Debug logs in the Access Analyzer displays information about the selected operation to understand the permissions, business rules, and ACLs associated with the operation.

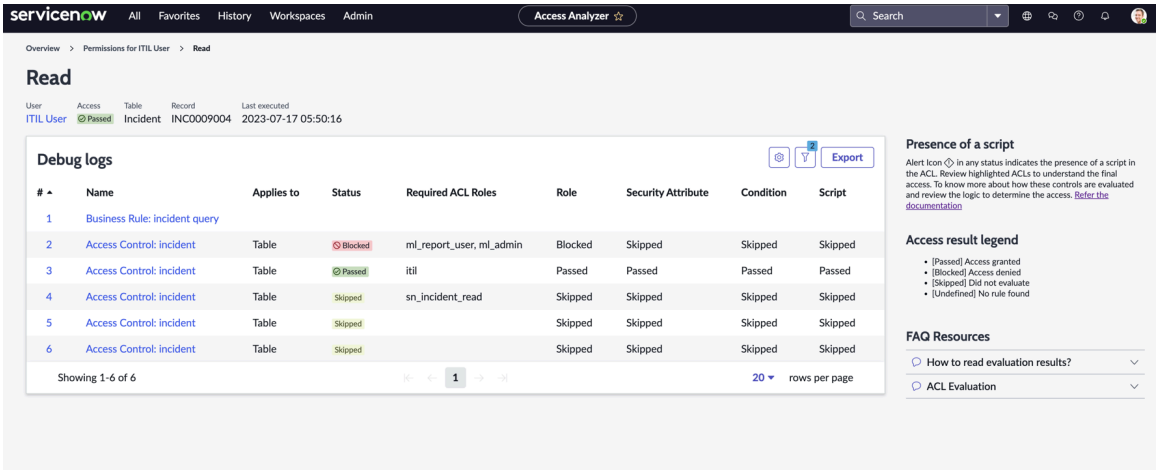

Following are the fields and their description in the Debug logs:

Debug logs

Fields	Description
Name	The details about the business rule or ACL. You can select the business rule of ACL for more information.
Applies to	The details about the application of ACL at a field, record, or table level.
Status	Status of the ACL for the associated role and permission.
Requires ACL	The role that is required for accessing the field, record, or table.
Role	The details about the role being Blocked, Passed, Skipped for the Access Control.
Security Attribute	The details about the security attribute being Blocked, Passed, Skipped for the Access Control.
Condition	The detail about the condition being Blocked, Passed, Skipped for the Access Control.
Script	The details about the script being Blocked, Passed, Skipped for the Access Control.
Customized	The details about the customized ACL if any for the Access Control.
Application	Status of the Application. Global or Store.

Evaluation hierarchy

Permission for the selected user, group, or role is evaluated in the following hierarchy:

- Business rule: A business rule is a server-side script that runs when a record is displayed, inserted, updated, or deleted, or when a table is queried.
- Access Handler: An internal system check using hidden source code on the platform.

- Data Filtration: A data filter is a form of access control designed to work along with the existing Access Control rules (ACLs) on your instance. Data filters support only read operation.
- Access control list (ACL): Rules for access control lists (ACLs) restrict access to data by requiring users to pass a set of requirements before they can interact with it. Within an ACL, the following hierarchy is evaluated:
  - Role
  - Security Attribute
  - Condition
  - Script

Access control list evaluation

ACLs for the operations are evaluated in the sequence as follows:

- Role
- Security Attribute
- Condition
- Script

Presence of a script

Alert Icon in any status indicates the presence of a script in the ACL. Review highlighted ACLs to understand the final access.

Note: During an Access analyzer query, business rules are executed first and then the access control list.

Sequence of execution

The sequence of access result execution in different scenarios is as follows:

- **Presence of an inherited or wildcard ACL**: During the sequence of execution the inherited ACLs are evaluated first and then wildcard ACL.
- **One ACL is passed the others are skipped**: During execution and evaluation of permission if one ACL is passed the other ACL execution and evaluation is skipped. Because the overall permission for the selected operation requires one ACL to access a field, record, or table for an identity.
- **Field level ACL and table level ACLs execution**: During execution field level ACLs are executed first followed by table level ACL to provide more granular results when analyzing the access for an identity.
- **Evaluation in the presence of scripted ACL**: When there's a presence of a script, the overall access for the operation is passed with an Alert icon to indicate the script in the ACL.

Access Simulator

The Access Simulator helps you simulate how access to the specified table would change once a role or group is assigned or removed from the user.

Explore  Learn the features and business value of Access Simulator.	Use  Understand how to use Access Simulator.
Configure  Know how to configure Access Simulator.	Frequently asked questions  Get details about commonly asked questions for Access Simulator.

Exploring Access Simulator

Access Simulator helps you simulate how access to the specified table would change once a role or group is assigned or removed from the user.

Administrators can use the Access Simulator to simulate the access requirements of their users on a specified table and understand how the access controls changes for a user. It helps in quickly understanding the present access of the users and understand the impacts of providing or removing roles to the users.

https://player.vimeo.com/video/988627830?badge=0&autopause=0&player_id=0&app_id=58479

Access Simulator can be used for the simulating access on a user for the following scenarios:

- Adding Roles to users
- Removing Roles from users
- Adding users to Groups
- Removing users from Groups

To know more about the Access Simulator settings and enable take actions, see Configuring the Access Simulator (Take actions).

Configuring the Access Simulator (Take actions)

Configure the take actions options to complete the role and group assignments that you wish to add after the simulation.

Before you begin
Role required: admin, access_analyzer_admin

Procedure

1. Navigate to **All** > **Access Analyzer** > **Settings**.

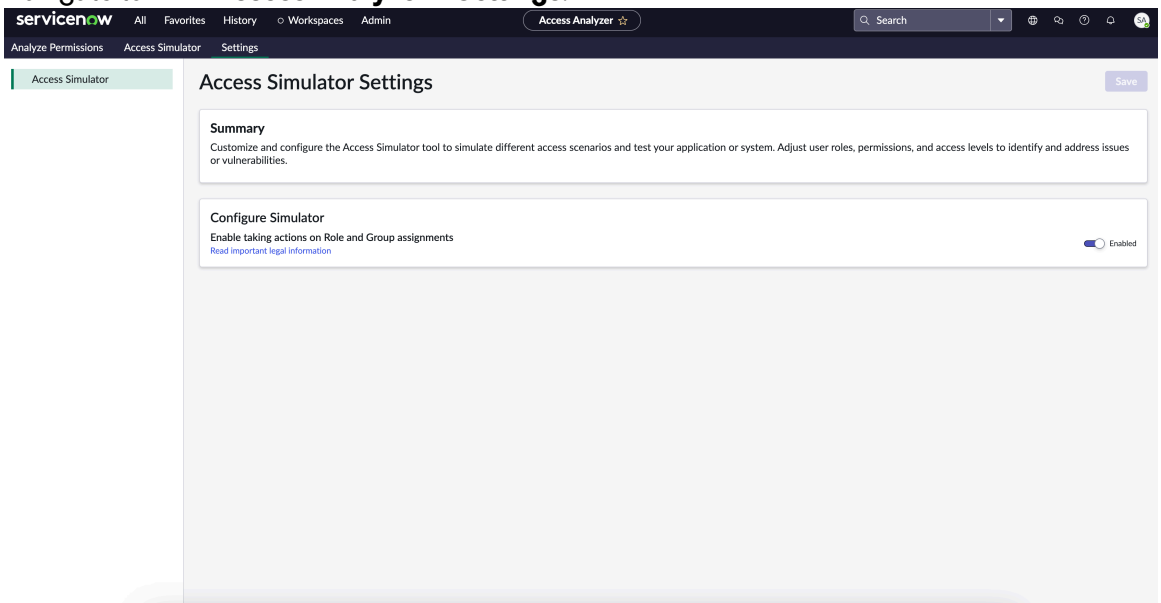

2. Select the toggle switch to enable.
   Read through the legal information before enabling the Access Simulator (Take actions).
3. Select **Accept**.
4. Select **Save** to save the configuration.

Using the Access Simulator

Use the Access Simulator to simulate the user's access to the specified table would change once a role or group is assigned or removed.

Before you begin
Role required: admin, access_analyzer_admin

Enable the take actions. For more information, see Configuring the Access Simulator (Take actions).

Procedure

1. Navigate to **All** > **Access Analyzer** > **Access Simulator**.
2. Select the **Simulate** option under each section based on the following scenarios:
   - Adding Roles to users
   - Removing Roles from users
   - Adding users to Groups
   - Removing users from Groups

Adding Roles to users

Use the **Simulate Add Role** for simulating the user's access changes for a resource (tables).

Before you begin

Role required: admin, access_analyzer_admin

Enable the take actions. For more information, see Configuring the Access Simulator (Take actions).

Procedure

1. Navigate to **All** > **Access Analyzer** > **Access Simulator**.
2. Select **Simulate** from the Add a Role to the user section.
3. Specify the following fields for your simulation criteria:

Add a Role to the user

Field	Description
Select user *	Specify a user name to select from the list. In this example, **Abel Tuter**.
Select role *	Specify a role to select from the list. In this example, **itil_admin**.
Rule type *	Rule type is auto-populated and it can't be changed.
Select table *	Specify a table name to select from the list. In this example, **Incident**.
Select record	Specify a record name to select from the list (Optional).
Select field	Specify a field name to select from the list. This field can be used to analyze permission even at a field level. For example, Active, Created By, and so on.

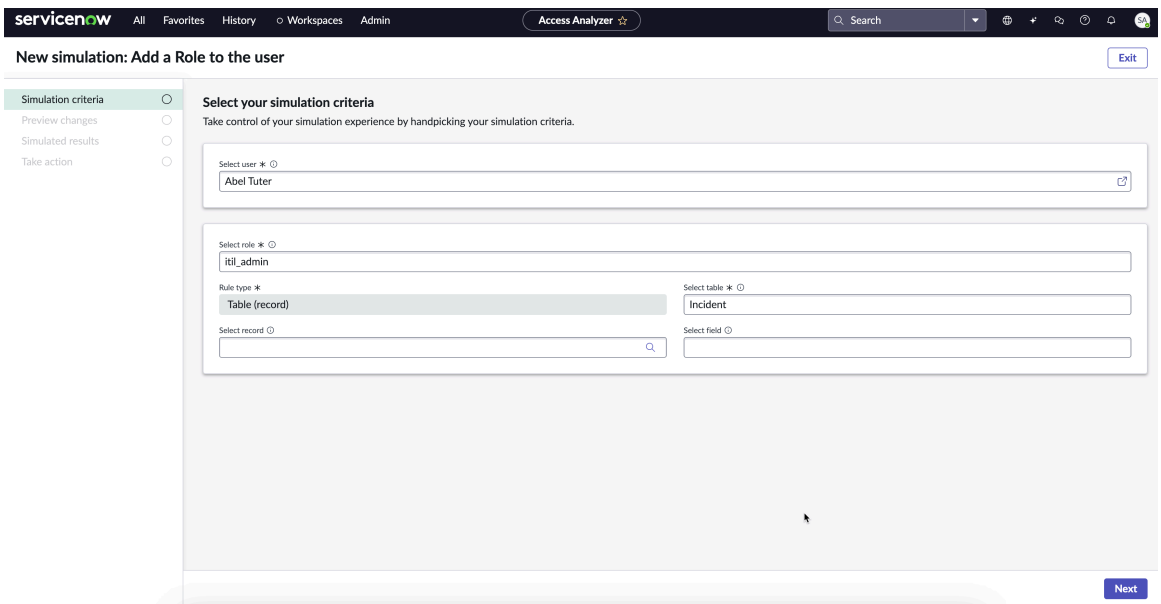

4. Select **Next**.
5. Preview the changes.
   The roles that are newly assigned are simulated in the Preview changes. You can validate the roles and child roles that are being added to the user before moving to the next step.

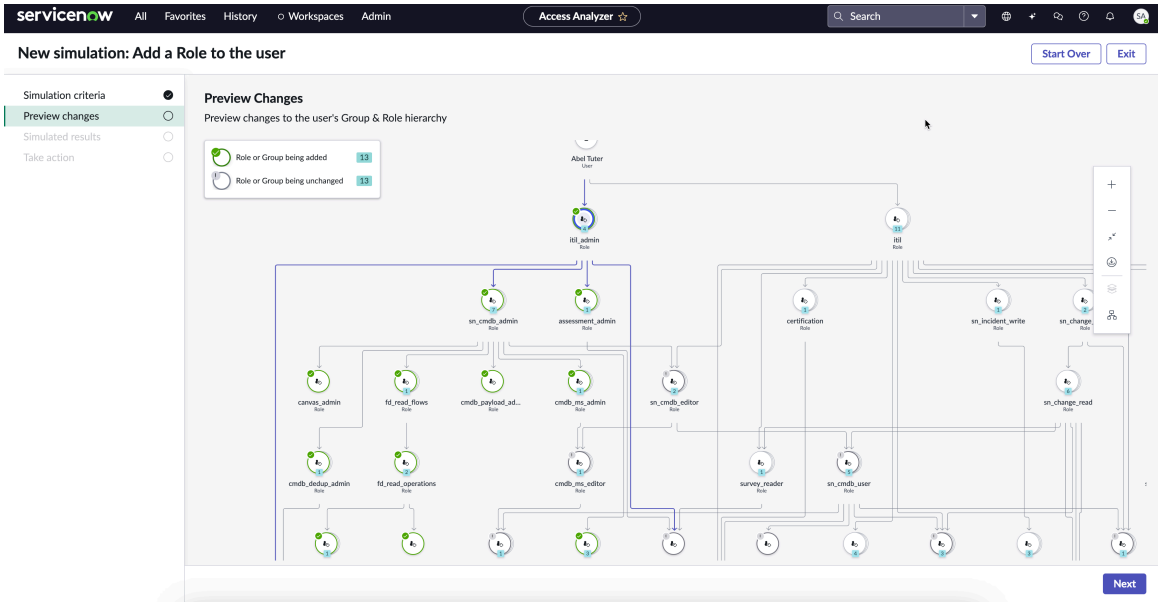

You can see the new roles that are being added as part of **itil_admin** role is displayed along with the existing roles of the user.

6. Select **Next**.
7. Validate the **Present status** and **Simulated status** to verify the access that is being **Passed** or **Blocked** to the simulated user.

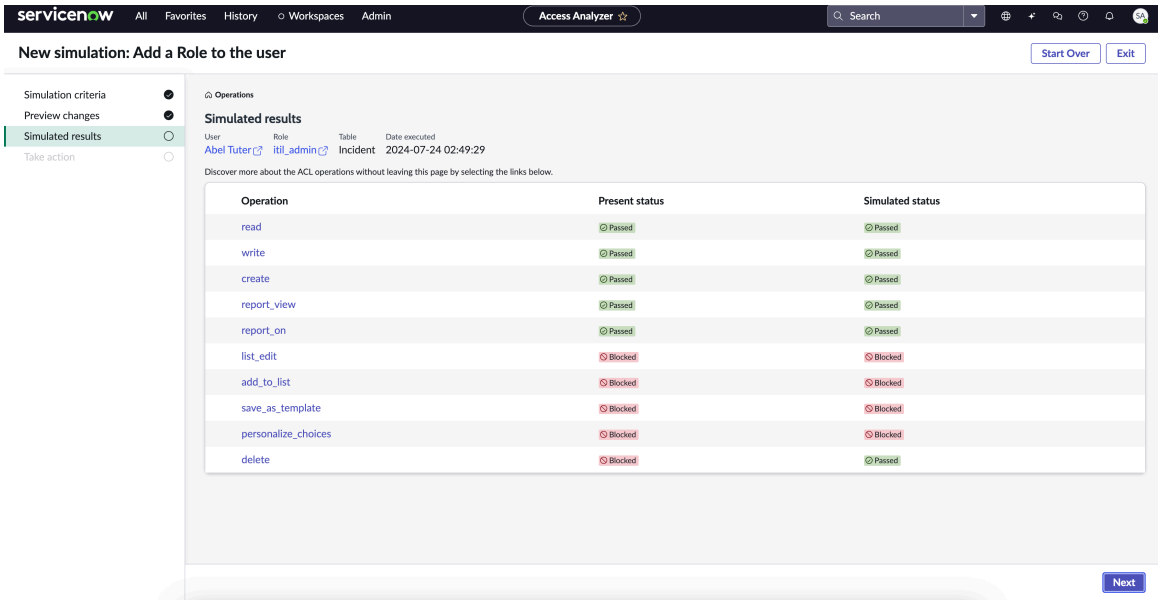

Adding the **itil_admin** enables the user to perform delete operation on the table (Incident).

Note:

- If you want to know more about the ACL (operations), select the operation links for each record.
- If you want to start the simulation again for a different role, select **Start over**.
- If you want to exit the simulation, select **Exit**.

8. Select **Next**.
9. Select **Add and complete**.

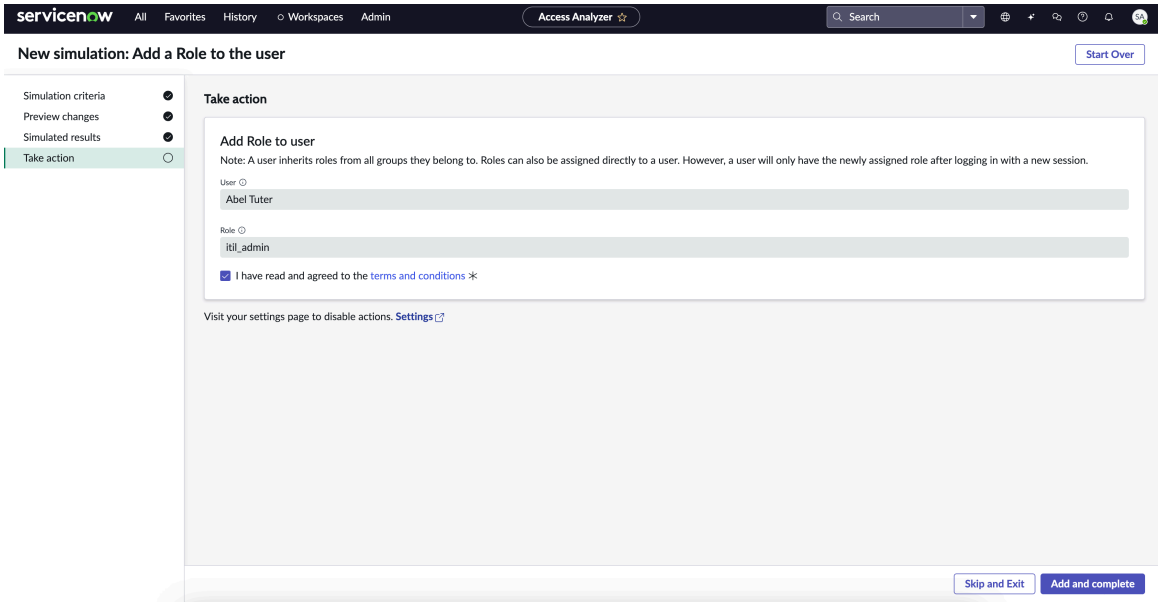

Note:

- If the Access Simulator isn't enabled, you can't complete the simulation. To enable, select **Enable actions** and accept the legal information.
- If you want to hide the simulation, select **Hide actions**. To unhide and enable actions, go to the settings. For more information, see Configuring the Access Simulator (Take actions).
- If you want to exit the simulation, select **Skip and Exit**.

The roles are successfully added to the user. You can verify the access using the Access Analyzer. For more information about how to verify the access of a user, see Access analyzer.

Removing Roles from users

Use the **Simulate Remove Role** for simulating the user's access changes for a resource (tables).

Before you begin

Role required: admin, access_analyzer_admin

Enable the take actions. For more information, see Configuring the Access Simulator (Take actions).

Procedure

1. Navigate to **All** > **Access Analyzer** > **Access Simulator**.
2. Select **Simulate** from the Remove a Role from the user section.
3. Specify the following fields for your simulation criteria:

Remove a Role from the user

Field	Description
Select user *	Specify a user name to select from the list. In this example, **Abel Tuter**.
Select role *	Specify a role to select from the list. In this example, **itil**.

Field	Description
Rule type *	Rule type is auto-populated and it can't be changed.
Select table *	Specify a table name to select from the list. In this example, **Incident**.
Select record	Specify a record name to select from the list.
Select field	Specify a field name to select from the list. This field can be used to analyze permission even at a field level. For example, Active, Created By, and so on.

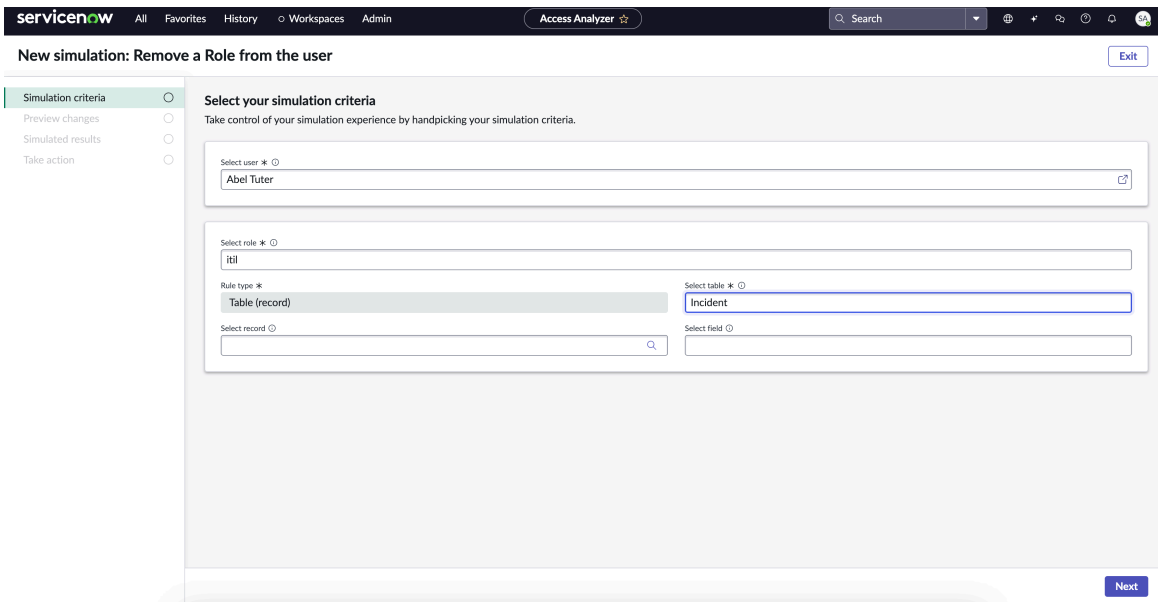

4. Select **Next**.
5. Preview the changes.
   The roles that are removed are simulated in the Preview changes. You can validate the roles and child roles that are being removed to the user before moving to the next step.

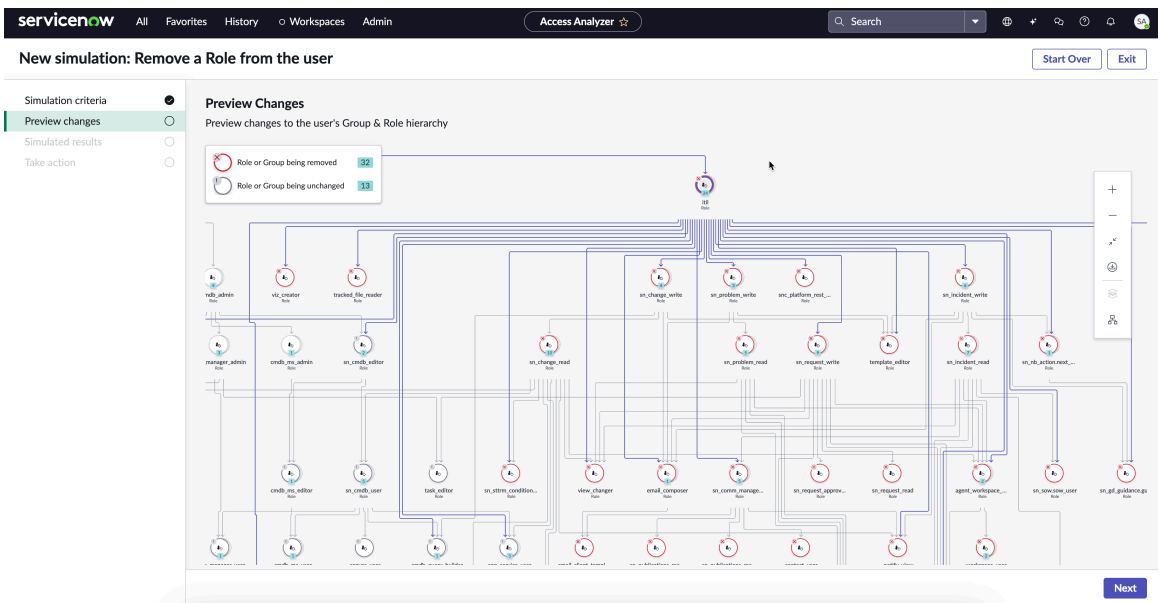

6. Select **Next**.
7. Validate the **Present status** and **Simulated status** to verify the access that is being **Passed** or **Blocked** to the simulated user.

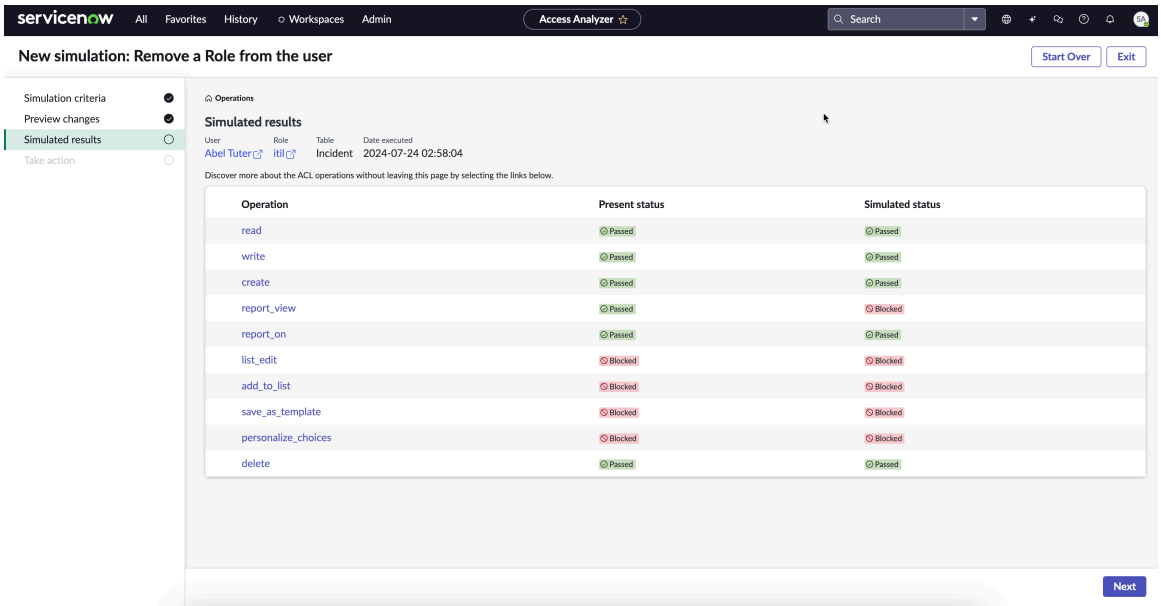

Note:

- If you want to know more about the ACL (operations), select the operation links for each record.
- If you want to start the simulation again for a different role, select **Start over**.
- If you want to exit the simulation, select **Exit**.

8. Select **Next**.
9. Select **Remove and complete**.

Note:

- If the Access Simulator isn't enabled, you can't complete the simulation. To enable, select **Enable actions** and accept the legal information.
- If you want to hide the simulation, select **Hide actions**. To unhide and enable actions, go to the settings. For more information, see Configuring the Access Simulator (Take actions).
- If you want to exit the simulation, select **Skip and Exit**.

The roles are successfully removed from the user. You can verify the access using the Access Analyzer. For more information about how to verify the access of a user, see Access analyzer.

Adding users to Groups

Use the **Simulate Add to Group** for simulating the user's access changes for a resource (tables) when the user is added to a group.

Before you begin

Role required: admin, access_analyzer_admin

Enable the take actions. For more information, see Configuring the Access Simulator (Take actions).

Procedure

1. Navigate to **All** > **Access Analyzer** > **Access Simulator**.
2. Select **Simulate** from the Add the user to a Group section.
3. Specify the following fields for your simulation criteria:

Add the user to a group

Field	Description
Select user *	Specify a user name to select from the list. In this example, **Abel Tuter**.
Select group *	Specify a group to select from the list. In this example, **Incident Management**.
Rule type *	Rule type is auto-populated and it can't be changed.
Select table *	Specify a table name to select from the list. In this example, **Incident**.
Select record	Specify a record name to select from the list.
Select field	Specify a field name to select from the list. This field can be used to analyze permission even at a field level. For example, Active, Created By, and so on.

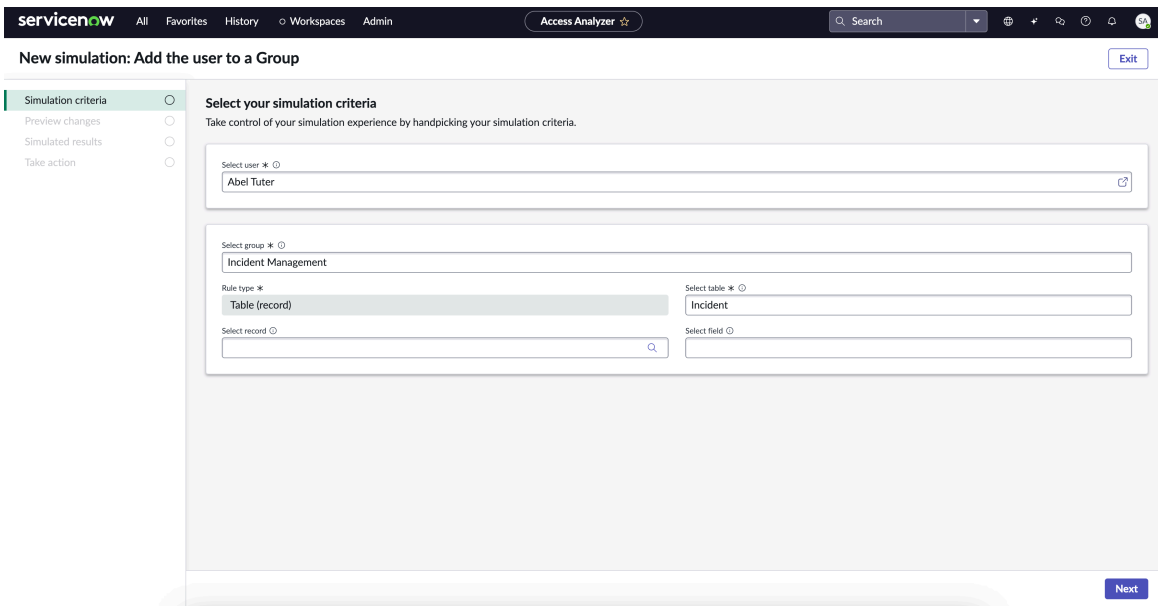

4. Select **Next**.
5. Preview the changes.

   The group that the user is assigned is simulated in the Preview changes. You can validate the changes before moving to the next step.

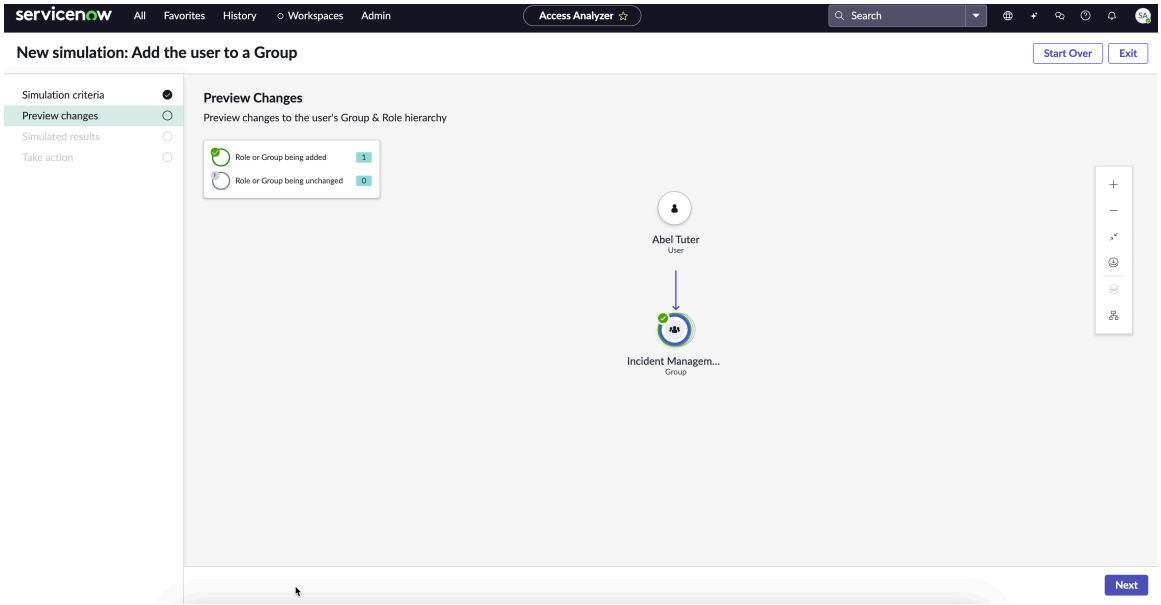

6. Select **Next**.
7. Validate the **Present status** and **Simulated status** to verify the access that is being **Passed** or **Blocked** to the simulated user.

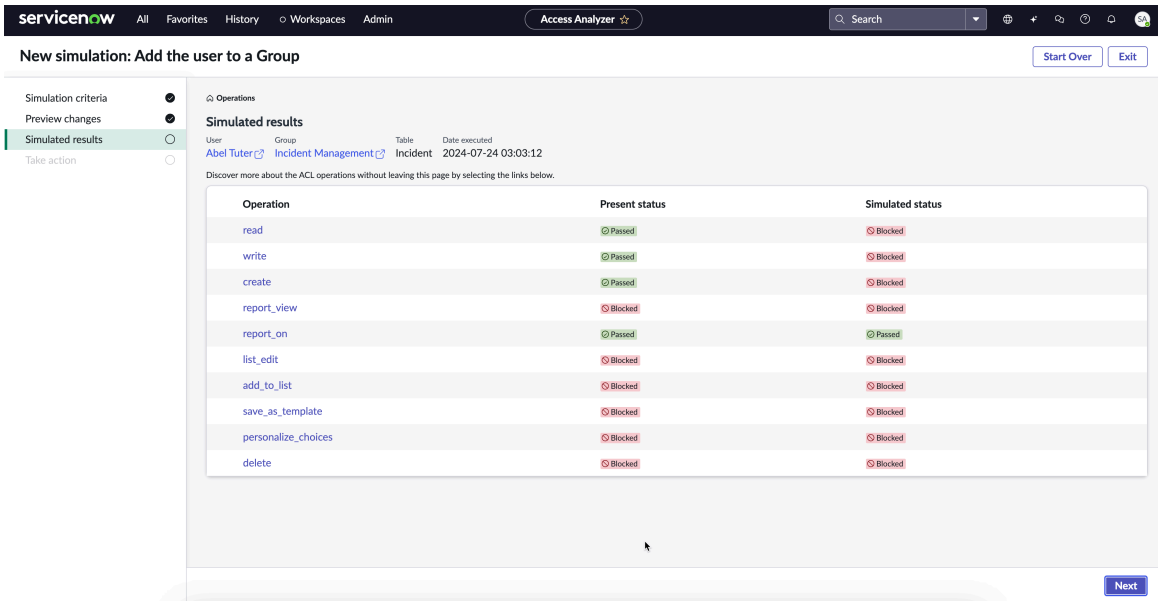

Note:

- If you want to know more about the ACL (operations), select the operation links for each record.
- If you want to start the simulation again for a different role, select **Start over**.
- If you want to exit the simulation, select **Exit**.

8. Select **Next**.
9. Select **Add and complete**.

Note:

- If the Access Simulator isn't enabled, you can't complete the simulation. To enable, select **Enable actions** and accept the legal information.
- If you want to hide the simulation, select **Hide actions**. To unhide and enable actions, go to the settings. For more information, see Configuring the Access Simulator (Take actions).
- If you want to exit the simulation, select **Skip and Exit**.

The user is successfully added to the group.

Removing users from Groups

Use the **Simulate Remove from Group** for simulating the user's access changes for a resource (tables) when the user is removed from a group.

Before you begin

Role required: admin, access_analyzer_admin

Enable the take actions. For more information, see Configuring the Access Simulator (Take actions).

Procedure

1. Navigate to **All** > **Access Analyzer** > **Access Simulator**.
2. Select **Simulate** from the Remove the user from a Group section.
3. Specify the following fields for your simulation criteria:

Remove the user from a group

Field	Description
Select user *	Specify a user name to select from the list. In this example, **ITIL User**.
Select group *	Specify a group to select from the list. In this example, **Incident Management**.
Rule type *	Rule type is auto-populated and it can't be changed.
Select table *	Specify a table name to select from the list. In this example, **Incident**.
Select record	Specify a record name to select from the list.
Select field	Specify a field name to select from the list. This field can be used to analyze permission even at a field level. For example, Active, Created By, and so on.

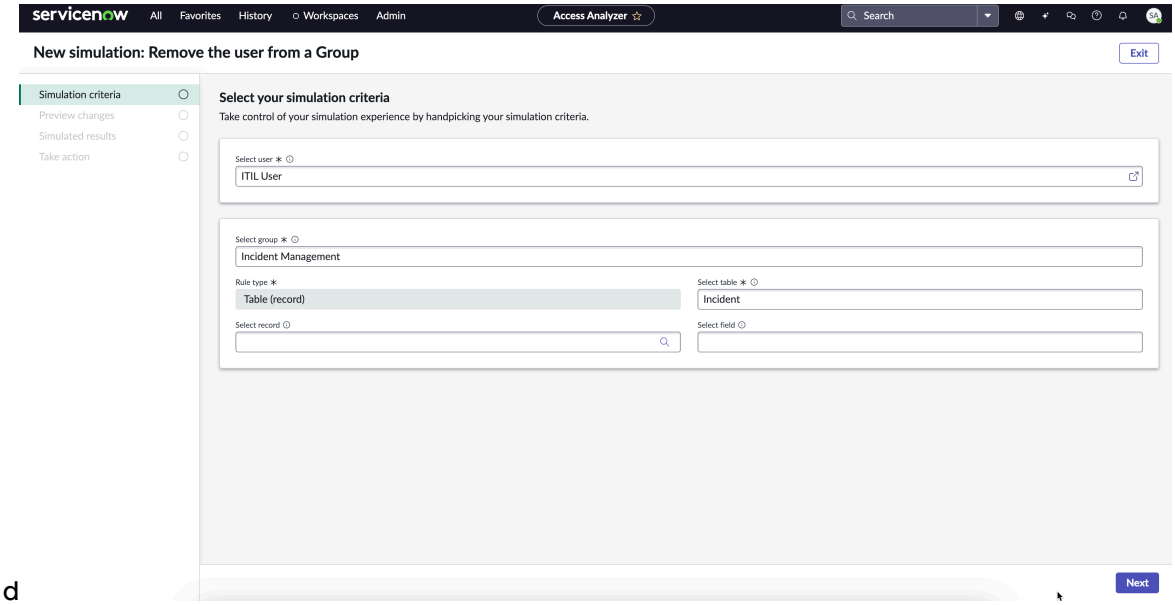

d

4. Select **Next**.
5. Preview the changes.

   The group from which the user is removed is simulated in the Preview changes. You can validate the changes before moving to the next step.

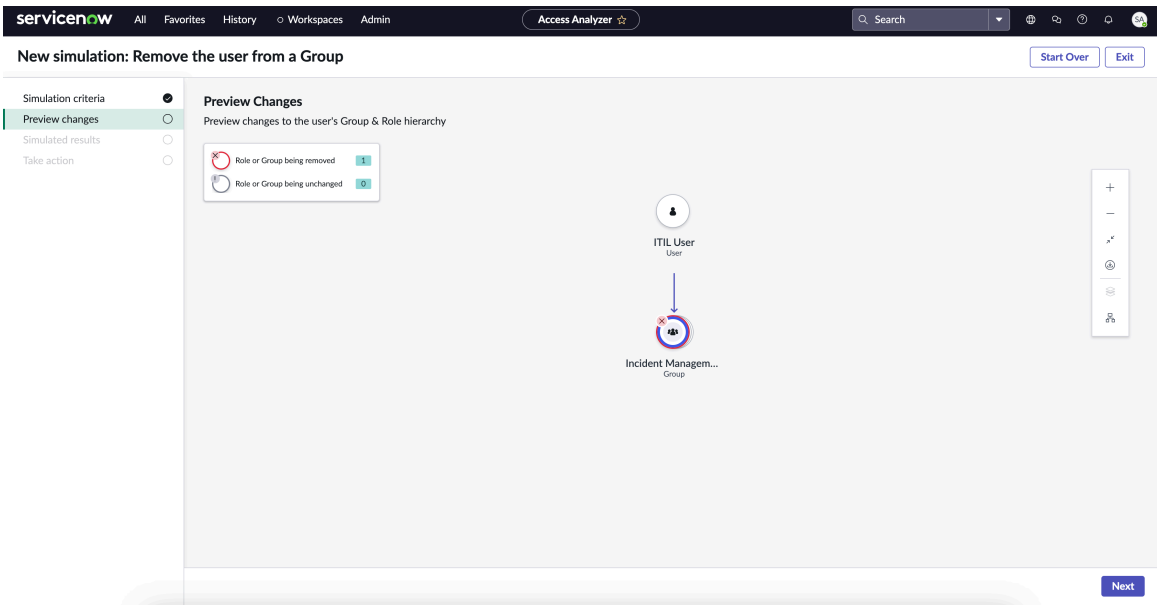

6. Select **Next**.
7. Validate the **Present status** and **Simulated status** to verify the access that is being **Passed** or **Blocked** to the simulated user.

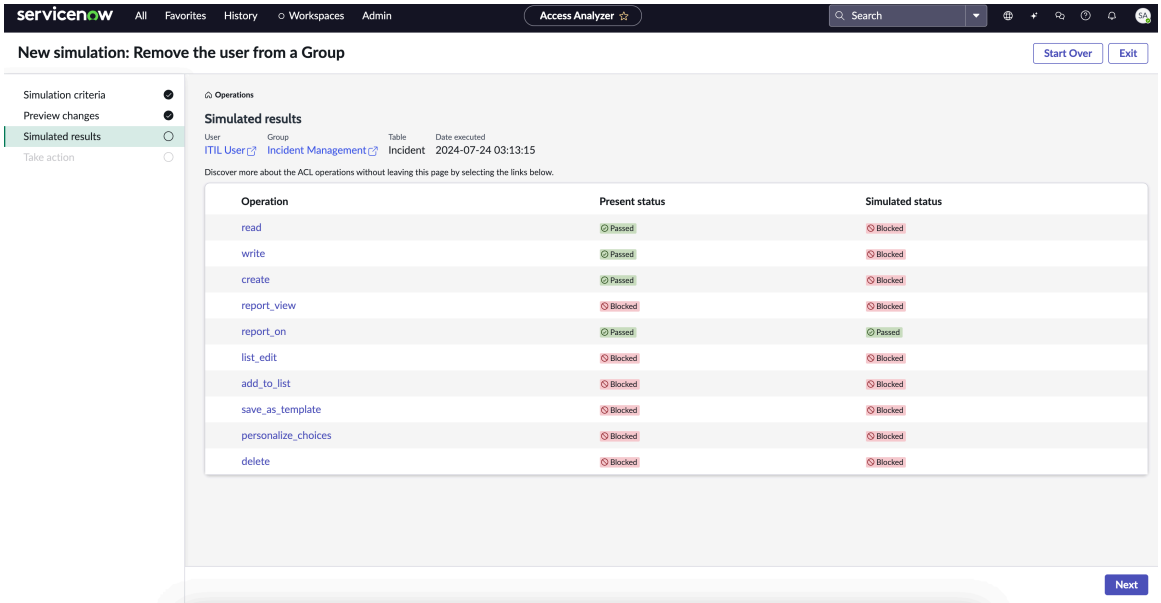

Note:

- If you want to know more about the ACL (operations), select the operation links for each record.
- If you want to start the simulation again for a different role, select **Start over**.
- If you want to exit the simulation, select **Exit**.

8. Select **Next**.
9. Select **Remove and complete**.

Note:

- If the Access Simulator isn't enabled, you can't complete the simulation. To enable, select **Enable actions** and accept the legal information.
- If you want to hide the simulation, select **Hide actions**. To unhide and enable actions, go to the settings. For more information, see Configuring the Access Simulator (Take actions).
- If you want to exit the simulation, select **Skip and Exit**.

The user is successfully removed from the user.

Frequently Asked Questions

Frequently asking question about using Access Simulator.

Frequently Asked Questions

Questions	Explanation
How do I get started with the Simulator?	You can select the Simulate option for the different scenario that you would like to simulate for the users and follow through the steps.
How does Simulator work?	The Access Simulator simulates access to the specified table to the selected user and

Frequently Asked Questions (continued)

Questions	Explanation
	previews the access changes to you to take decisions around access provisioning.
How to enable or disable Actions in Simulator?	You can enable or disable Actions in Access Simulator by navigating to Settings or during the Take Actions steps in the simulation process.

Access Insights

The Access Insights helps you gain confidence in your Role or Group assignments.

Explore

Learn the features and business value of Access Insights.

Configure

Understand how to configure Access Insights.

Use

Know how to use Access Insights.

Explore Access Insights

The Access Insights in the Access Analyzer helps you gain confidence in your Role or Group assignments.

The Access Insights displays the statistics of your user's peer and organizational level access to a resource. It enables you to see if the role or group can be assigned to the user that you're currently reviewing for additional entitlements. It helps you to compare the role and group assignments withing the same peer group.

The Access Insights is displayed only while comparing user's access using the compare user access feature in the Access Analyzer. To learn more about using Access Insights, see Use Access Insights.

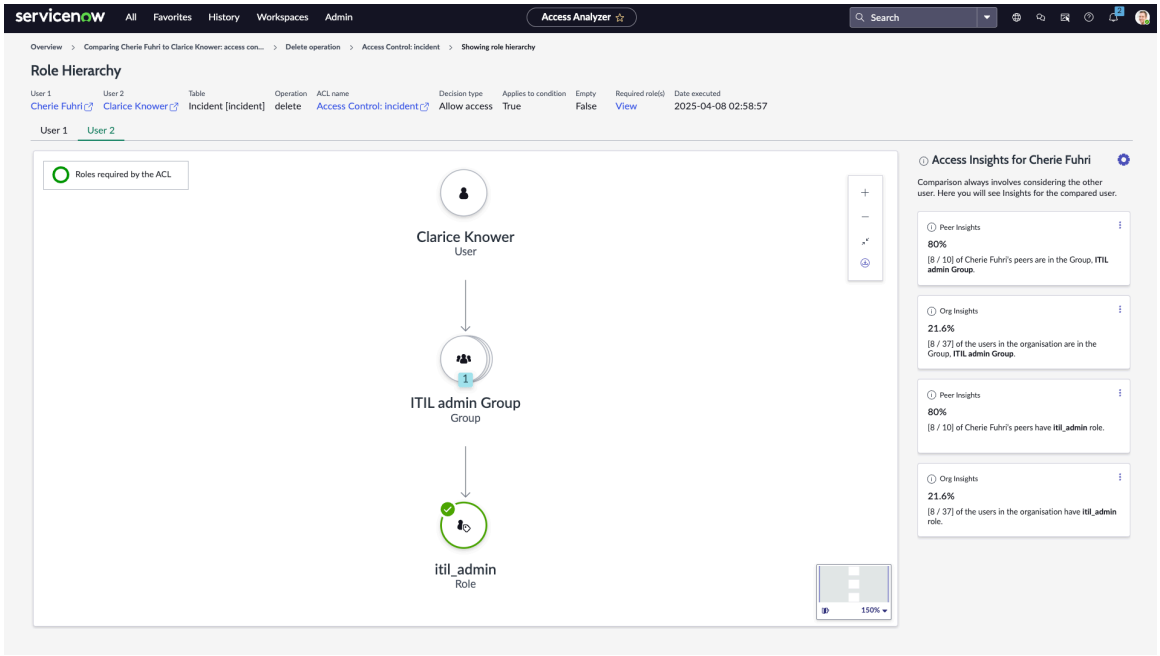

Based on the statistics, you can decide to provide the correct role or group assignments to the users you're comparing within their peers.

Note:

- Access Insights is available with Access Analyzer V4.
- You should compare users at a peer-level (same organization, same department, same manager), while viewing the statistics within the Access Insights feature.

To enable Access Insights in the Access Analyzer, you must navigate to the Settings. To learn more, see Configure Access Insights.

Configure Access Insights

Enable the Access Insights feature in the Access Analyzer.

Before you begin

Role required: admin, access_analyzer_admin

The following procedure describes the steps for enabling Access Insights.

Procedure

1. Navigate to **All** > **Access Analyzer** > **Analyze Access**.
2. Select **Settings** tab.
3. Select **Access Insights**.
4. Enable **Access Insights**.

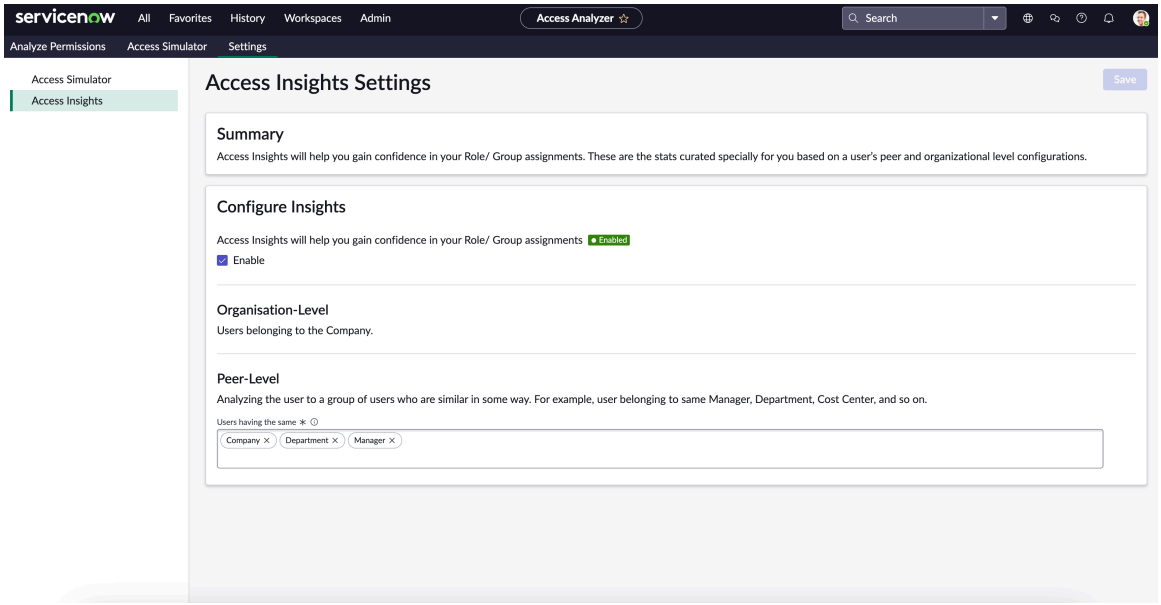

Note: By default, Peer-Level fields are selected with **Company**, **Department**, **Manager**. You can add more fields such as **Cost Center**, **Country Code**, **Location**, and **Title**.

5. Select **Save**.

Result

The Access Insights feature is enabled and displayed when you're viewing the **Role Hierarchy** while comparing user access to a resource. To know more about using Access Insights, see Use Access Insights.

Use Access Insights

Use Access Insights to understand the users peer level access to a selected resource.

Before you begin

Role required: admin, access_analyzer_admin

Note: To view the details of though Access Insights feature, you must use the Compare user access feature.

The following procedure describes the steps for viewing the peer level role or group assignments statistics within the Access Insights feature.

Procedure

1. Navigate to **All** > **Access Analyzer** > **Analyze Permissions**.
   The Analyze access and permissions home page is displayed.
2. Select the **Compare user access** tab.
3. Fill in the following fields:

Compare user access

Field	Description
Select user 1*	Specify a user name to select from the list for the comparison.

Field	Description
Select user 2*	Specify a user name to select from the list to compare with the user 1.
Rule Type*	Analyze access permissions for a table. **Note:** Only access permissions for a table can be used in the **compare user access**.
Select table*	Specify a table name to select from the list.
Select record	Specify a record name to select from the list.
Select field	Specify a field name to select from the list.

4. Select **Compare user access**.

   The **compare user access** results for the selected users are displayed.

   The compare user access results show the operation and the access evaluation status for the users. For example, Charlier Fuhri and Charlie Knower.

5. Select an operation.
   For example, **read** operation.
6. Select any of the **Access Control** to learn more about the access.
7. Select **Show role Hierarchy** to view the access comparison at a peer level access.

   The Access Insights is displayed as follows:

   - **Peer Insights**: Displays how many users have the same role assigned within their peers (the same manager or department).
   - **Org Insights**: Displays how many users have the same role assigned within the organization.

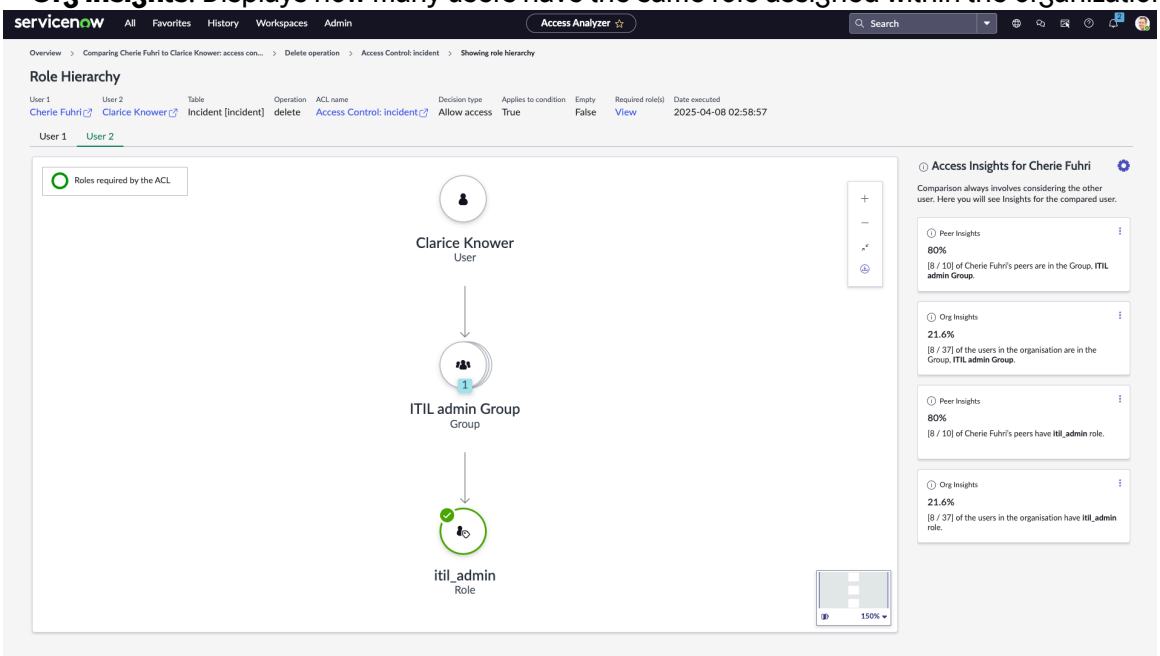

Based on this information, you can choose to provide or revoke access to the user that you'd compared.

Security Attributes

Security Attributes offer a flexible alternative to access control lists.

Get started

Explore Security Attributes

Learn fundamentals of Security Attributes

Create Security Attributes

Create new Security Attributes

OOB(Out-of-Box) Security Attributes

Explain OOB Security Attributes

Security Attributes Logging

Review Security Attribute logging

Security Attributes fundamentals

A Security Attribute is a highly configurable piece of information about a Subject or its Environment. When used in access controls, they enable fine grain security configuration in a non-complex way.

Security Attributes offer an alternative method for access control via role definitions to currently practiced ACL(Access Control Lists) configuration. Security Attributes offer several advantages to ACL-based configurations:

Better security

Improved permissions evaluation obfuscation ensures that your organization remains secure.

Human readable

Security Attributes are designed to simplify the creation and ease of use of security permissions.

Flexible security

Build persona definitions from both Out-of-Box configuration in combination with customer defined personas.

Logging and auditing

Security Attributes offer detailed auditing and logging to give more insight to security measures and theory.

Create Security Attributes

Create new Security Attributes with a step-by-step guide.

Before you begin

Role required: security_admin

Procedure

1. Navigate to **System Security** > **Security Attributes**.
2. In the Security Attributes list, select **New**.
3. On the Security Attributes fields form, fill in the fields.

Security Attribute fields

Field	Description
Label	Label of Security Attribute.
Name	Name of Security Attribute.
Type	The type of Security Attribute. ○ compound **Note:** For further information about Compound Security Attributes see Compound Security Attributes ○ integer ○ list ○ string ○ boolean(true|false)
Is dynamic	If the Security Attribute value needs to be re-evaluated per each
Description	User generated description of Security Attribute
Application	Static field, application Scope.
Lookup Table	Reference an external table for evaluation.
Lookup Table Column	Reference a table column for evaluation.
Script	Derive a value from a script.

OOB (Out-of-Box) Security Attributes

Commonly used, generalized Security Attributes roles ready for use.

Overview

The OOB(Out-of-Box) Security Attributes are an easy way to begin using and learning the capabilities of Security Attributes with a series of preconfigured Security Attribute roles. The OOB Security Attribute roles are commonly used access control roles.

To create your own Security Attribute, or expand an OOB Security Attributes capabilities see Compound Security Attributes

OOB Security Attributes

Attribute	Description
Group	User is member of a specified group
GroupExplicit	User is an explicit member of a specific group
HasAdminRole	User has the admin persona
Impersonating	User is impersonated
InteractiveSession	Current session interactive
LoggedIn	User is logged-in and authenticated
NetworkCriteria	Additional Network Criteria
Role	User has specific role
RoleExplicit	User has specific role explicitly defined

Security attributes for client session

Following security attributes are added for client session (plugin: `com.glide.client_session_security_attributes`):

- **IsIframeEmbeddedSession**: The attribute is used in embedded iframe in third party portals. For example, Engagement messenger, Virtual Agent Web Client, Embedded Session, and so on.
- **IsIntegrationAsAServiceSession**: The attribute is used in messaging apps like Virtual Agent in teams, Virtual Agent in slack, Virtual Agent in whatapp.
- **IsIntegrationAsAUserSession**: The attribute is used in integration account user or web service user.
- **Is Servicenow Web Session**: The attribute is used in web interactive session.
- **Is Mobile App Session**: The attribute is used if the **mobile_client** property is true in the **oauth_entity** record.

Note: You must use the **IsIframeEmbeddedSession**, **IsIntegrationAsAServiceSession**, and **IsIntegrationAsAUserSession** only while configuring client type for OAuth and SSO records. To learn more, see Configure client type for OAuth and SSO records.

Non-explicit and explicit behavior explained

Security Attributes address nuanced permission needs with an explicit vs. non-explicit (inherited) evaluation of roles permissions.

Example: Separate an IT persona access from HR role file access

We will define a Security Attribute that enables Admin to see HR related files but does not allow the Admin file access. Presume the Admin role is not defined as part of the HR group, but is _____. The **Role** sees _____

Compound Security Attributes

Compound Security Attributes enable you to create consistent and reusable Security Attribute profiles suit your business needs

Overview

Compound Security attributes are defined from one or more pre-existing Security Attributes to create a single reference combination of Security Attributes for permissions evaluation.

Compound Security Attribute Behavior

Create compound Security Attributes

Create a compound Security Attribute for easy reuse.

Before you begin
Role required: security_admin

Procedure

1. Navigate to **System Security** > **Security Attributes**.
2. In the Security Attributes list, select **New**.
3. Select `Compound` in the **Type** field.
4. On the Security Attributes fields form, fill in the fields.

Security Attribute fields

Field	Description
Label	Label of Security Attribute.
Name	Name of Security Attribute.
Description	User generated description of Security Attribute
Condition	The specific Security Attribute conditions used to the define the Security Attributes evaluation.
New Criteria	Add additional set of **Or** conditions for evaluation.

Security Attribute Scope

Security Attributes supports scoping capabilities.

Security Attribute Scoping

Security Attributes scoped behavior is consistent with Platform's scoping behaviors. Security Attributes created within a scope are only available in access controls in the same scope.

Local and Existing Security Attributes

Existing and Local Security Attributes enable customers to reuse Security Attribute condition sets.

Local & Existing Attributes

The ACL Security Attribute Condition Builder enables customers to specify if a Security Attribute is either existing or local.

Note: Security Attributes conditions are defaulted to local.

A local defined Security Attribute is saved only to the single ACL where it is created in.

The existing option enables users to reference an existing Security Attributes conditions to the condition builder.

Field Query Roles and Restrictions

The **Field Query Roles** and **Field Query Restrictions**attribute enable you to better control what information is available to users in tables.

Overview

The **Field Query Roles** and **Field Query Restrictions** attributes enable more fine tuned control over what actions users can take on tables.

Field Query Roles

Prevents querying sensitive fields based on user role.

Field Query Restrictions

Restrict access to the record if the access to the applicable field is denied.

Configure a Field Query Role

Learn how to enable and configure a Field Query Role attribute.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Tables**.
2. Select a **Table**.
3. Select a **Column Label** to view its **Dictionary** entry.
4. Select the **Attributes** tab.
5. **Tip:** If you do not have the **Field Query Roles** field, select **New** to create one

   Select the **Field Query Roles** attribute.
6. Enter the roles required to view this table column in the **Value**
7. Select **Update**

Result

The table column can only be queried by the roles defined in the **Field Query Roles** attribute. If a user without the defined roles queries the table column they will be denied access to the operation.

Configure Field Query Restrictions

Learn how to configure a Field Query Restriction

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Tables**.
2. Select a **Table**.
3. Select a **Column Label** to view its **Dictionary** entry.
4. Select the **Attributes** tab.
5. **Tip:** If you do not have the **Field Query Restrict Record Access** field, select **New** to create one

   Select the **Field Query Restrict Record Access** attribute.
6. Enter `true` in the **Value** field to restrict record access to only users who can who can view the field, otherwise enter `false`.
7. Select **Update**

Result

The table column will restrict the results of any queries. Restricted users will be informed how many results a query returned, but no further information.

Scripting Governance Tool

Use the Scripting Governance Tool to provide a single, centralised control for managing scripting access across your ServiceNow AI Platform.

Explore	Use
	
Learn the features and business value of Scripting Governance Tool.	Understand how to use the Scripting Governance Tool.

Explore Scripting Governance Tool

The Scripting Governance Tool provides a single, centralised control for managing scripting access across your ServiceNow AI Platform.

The Zurich release introduces a feature that gives administrators centralised control over who can edit script fields on the ServiceNow AI Platform. The feature adds a new permission layer built on the **Conditional Script Writer** group and its child role, `snc_required_script_writer_permission`. Users must be members of this group to edit any script field, regardless of their existing ACL-based access. This permission layer is enforced through data-type ACLs and maintained by scheduled jobs and system properties.

To manage this feature, ServiceNow AI Platform provides the **Scripting Governance Tool** — a dashboard where administrators can see which users have scripting access, scan the instance for users who have edited script fields, and directly add or revoke scripting access from users.

Why it matters

Scripting can read and write data across tables, bypass business logic, and alter platform behaviour at a fundamental level making scripting access one of the most important permissions to govern on any instance.

Before Zurich, scripting access was controlled solely by ACLs on individual script fields. To determine which users could edit scripts, administrators had to check each ACL individually. There was no single place to view or manage scripting access across the instance.

Scripting governance feature addresses it by adding a mandatory second layer of access control on top of existing ACL-based permissions. Satisfying a field-level ACL alone is no longer sufficient to edit a script field. Now, security administrators can manage scripting access centrally by adding or removing users from the **Conditional Script Writer** group.

The two-layer access model

Scripting governance feature enforces a two-layer access model. Both layers must pass independently before a user can edit any script field. Passing one layer alone is not sufficient.

Layer 1 — Existing field-level ACL

The user must pass the existing ACL on the scripting field (out-of-the-box or custom). This check is unchanged from pre-Zurich behaviour.

Layer 2 — Scripting Governance Tool role check

The user must also hold the `snc_required_script_writer_permission` role, which is granted through membership in the **Conditional Script Writer** group.

Important:

A user who could edit script fields before the Zurich upgrade — because their roles satisfied the field-level ACL — will be blocked after the upgrade unless they also satisfy Layer 2. The `snc_required_script_writer_permission` role does not grant new scripting access on its own. It only unlocks access for users who already pass Layer 1.

The following table summarises access outcomes under the two-layer model:

Passes field-level ACL (Layer 1)?	Holds `snc_required_script_wr` (Layer 2)?	Can edit script field?
Yes	Yes	# Yes
Yes	No	# No
No	Yes	# No
No	No	# No

Data type ACLs

Scripting governance feature introduces 9 data type ACLs to enforce Layer 2. These are **deny unless** ACLs that require the `snc_required_script_writer_permission` role for the following data types:

Data type

- script
- script_client
- script_plain
- script_server
- email_script
- html_script
- html_template
- xml
- condition_string

Note: The admin role does not have scripting access by default. Admin users are subject to the same two-layer check and cannot edit script fields unless they are members of the **Conditional Script Writer** group or explicitly hold the `snc_required_script_writer_permission` role.To know more, see Datatype ACL.

Scheduled jobs and properties

As part of the scripting governance feature, the **Conditional Script Writer** group is introduced, which has the `snc_required_script_writer_permission` role as a child role. When the Zurich upgrade completes, a one-time scheduled job runs to auto-populate all eligible users to this group ensuring no one loses scripting access after the upgrade. Once that job completes, a recurring weekly job is created to keep the group membership up to date. Both jobs and their eligibility criteria are described as follows:

Add users to Conditional Script Writer group

A job that runs once immediately after the Zurich upgrade completes. It adds all users who meet the eligibility criteria to the **Conditional Script Writer** group, ensuring no user loses scripting access after the upgrade. After the job completes, the platform disables it.

The job uses the property `glide.security.scripting_role.provisioning_job_running`, which is set to true while the job is running and false after it completes. When the job finishes, it creates the property `glide.security.scripting_role.auto_provisioning` and sets it to **true**.

Update users in Conditional Script Writer group

A weekly job that adds new users who meet the eligibility criteria to the group and removes users who no longer meet the eligibility criteria. It is controlled by the `glide.security.scripting_role.auto_provisioning` property:

Update behaviour based on the property

Value	Behaviour
true (default)	The weekly job runs on schedule and adds qualifying users to the group automatically.

Update behaviour based on the property (continued)

Value	Behaviour
false	The job does not run. Users are only added to the group manually by an administrator.

To disable automatic provisioning, set `glide.security.scripting_role.auto_provisioning` to **false** in System Properties.

Note: The `glide.security.scripting_role.auto_provisioning` property is created at runtime and has a different **sys_id** on every instance. Always reference it by a property name.

Eligibility criteria

Both scheduled jobs use the same criteria to determine which users are added to the **Conditional Script Writer** group:

Eligibility criteria

Explicit Role plugin	User type	Requirement	Added to group?
Enabled	External	—	No
Enabled	Internal	Must have snc_internal and at least one additional role	Yes
Disabled	Any user	Must have at least one role	Yes

Note: The **snc_external** role and the `snc_required_script_writer_permission` role are conflicting roles. ServiceNow AI Platform does not allow both role to the same user. External users cannot have scripting access.

Scripting Governance Tool

The Scripting Governance Tool is a dashboard that helps administrators manage scripting access on the ServiceNow platform. It provides visibility into users who are members of the **Conditional Script Writer** group — giving a clear picture of how many users can script on your instance.

You can also view which groups contain the scripting role and which roles contain it as a child role. There are two ways to manage scripting access through the tool:

- **Manual configuration**: Manually add or remove users from the **Conditional Script Writer** group to control who has scripting access.
- **Scan for users who have scripted**: Scan your instance to find users who have scripted within a specific time frame. The scan queries the audit logs and identifies any user who has performed write or update to a table having script field.

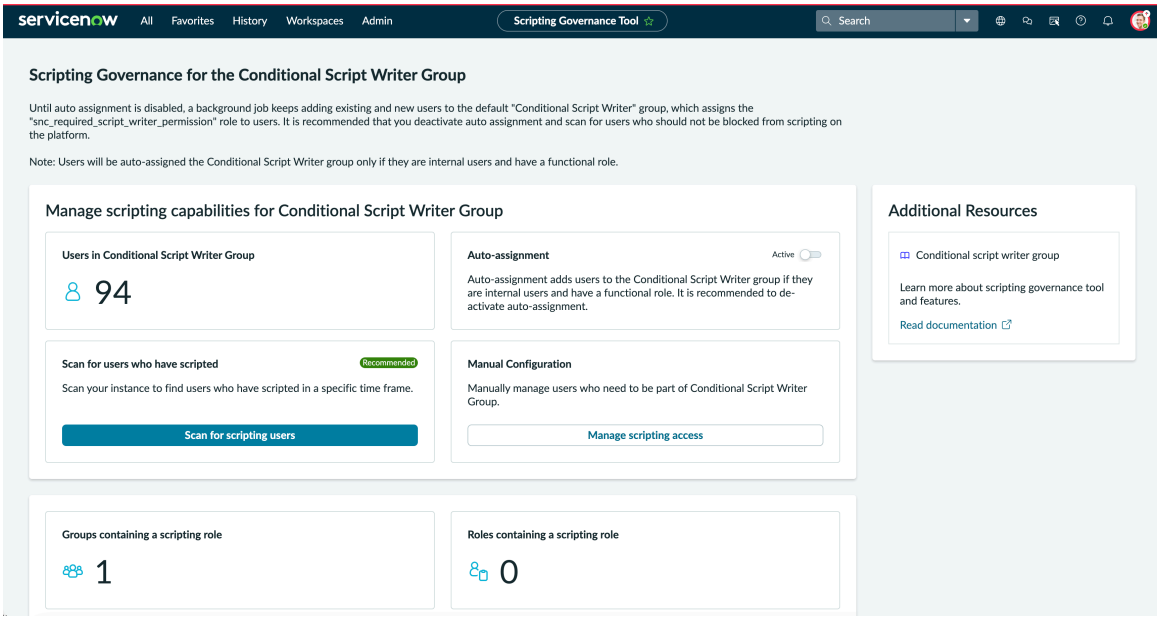

Note: It is recommended to manage scripting access exclusively through the **Conditional Script Writer** group. Adding the `snc_required_script_writer_permission` role as a child role to other roles or groups reduces your ability to centrally control who can script on your instance.

Use Scripting Governance Tool

Use the Scripting Governance Tool to provide a single, centralised control for managing scripting access across your ServiceNow AI Platform.

Before you begin

Role required: security_admin

Important: You must have elevate your role `security_admin`. See Elevate to a privileged role.

The Scripting Governance Tool helps you review user scripting governance your ServiceNow AI Platform. The **Conditional Script Writer** group grants scripting permissions to its members via the `snc_required_script_writer_permission` role. Users are added to the group by either an automated assignment or manual configuration. You can manage both of these settings from the Scripting Governance Tool.

Procedure

1. Navigate to **All** > **System Security** > **Scripting Governance Tool**.
2. Use the following features to learn more:

Scripting Governance Tool

Features	Description
Users in Conditional Script Writer Group	Displays the number of users in the **Conditional Script Writer** group. These users are granted the `snc_required_script_writer_permission` until removed from the group.

Features	Description
Auto-assignment	Auto-assignment assigns new users to the **Conditional Script Writer** group if the users are `internal users` and has one functional role. You can select the slider to de-activate. **Note:** It is recommended to de-activate auto-assignment.
Scan for users who have scripted (Recommended)	Scans for users in the ServiceNow AI Platform who have modified records containing script fields such as business rules, script includes, or client scripts. When running the scan, you can define a time period for which the scan checks.
Manual Configuration	You can manually select users that stay in the **Conditional Script Writer** group. Select the **Manage scripting access** button and enter the users into the text field for manually managing the user removal process from the **Conditional Script Writer** group.
Groups containing a scripting role	Displays the number of groups that contain a scripting role. By default the **Conditional Script Writer** group has a scripting role. **Note:** It is recommended to manage scripting access exclusively through the **Conditional Script Writer** group. Adding the `snc_required_script_writer_permission` role as a child role to other roles or groups reduces your ability to centrally control who can script on your instance.
Roles containing a scripting role	Displays the number of roles that contain a `snc_required_script_writer_permission` role. **Note:** It is recommended to manage scripting access exclusively through the **Conditional Script Writer** group. Adding the `snc_required_script_writer_permission` role as a child role to other roles or groups reduces your ability to centrally control who can script on your instance.
View scans	After you run the scan, the details of the scan are displayed on the View scans.

Features	Description
View removals	When you schedule for removal of user from the **Conditional Script Writer** group, the details are displayed on the View removals.

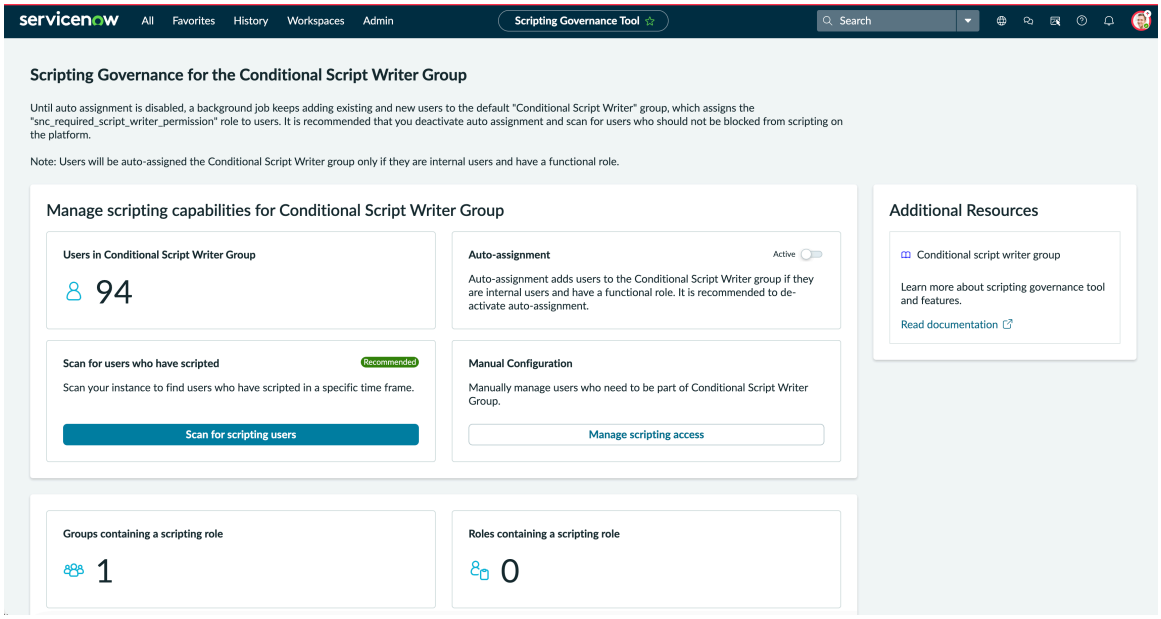

Result

You can view the following topics to understand how you can:

- Scan for users who have scripted
- Remove users from the Conditional Script Writer group

Scan for users who have scripted

Scan your instance to find users who have scripted within a specific time frame. The scan queries the audit logs and identifies any user who has performed write or update to a table having script field.

Before you begin

Role required: security_admin

Important: You must have elevate your role `security_admin`. See Elevate to a privileged role.

Procedure

1. Navigate to **All** > **System Security** > **Scripting Governance Tool**.
2. Select **Scan for scripting users**.
3. Select a time period from the drop-down list.
   - 7 days
   - 15 days
   - 30 days
   - 90 days

- 120 days
- 180 days

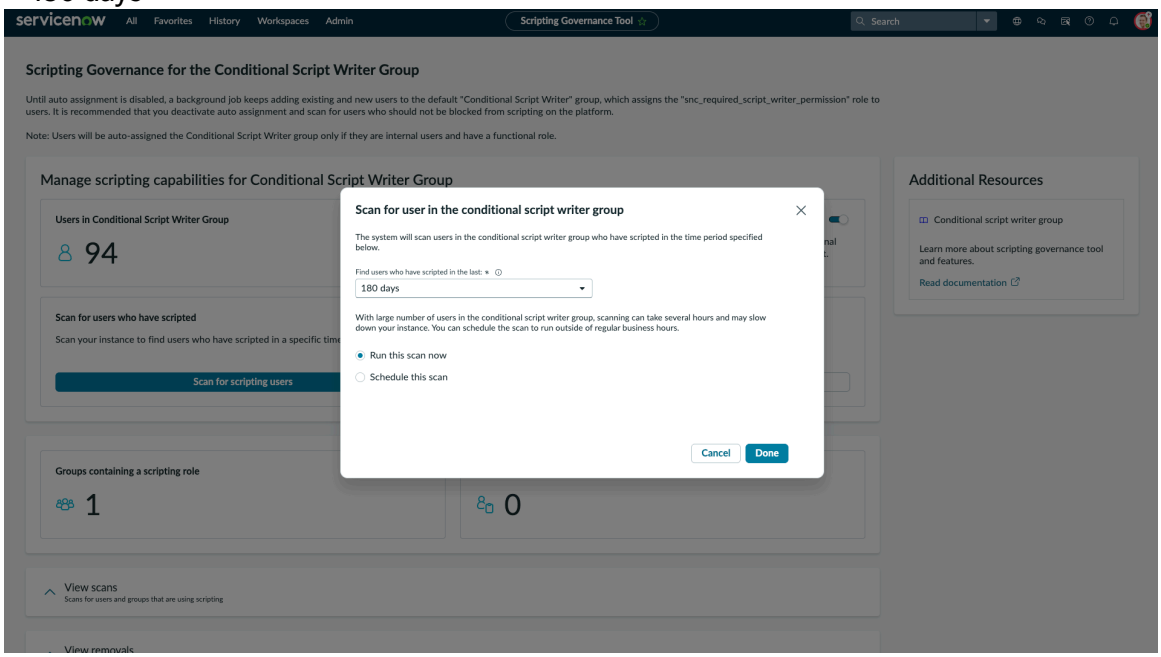

The default value is **180 days**.

4. Select how you want to run the scan.
   - Run this scan now — starts the scan immediately.
   - Schedule this scan — runs the scan at a time you specify.
5. Select **Done** to scan.
6. Select the scan record from **View scan** section.
   The following details are available on the **View scan** section.

Scripting scan table fields

Label	Description
Scan number	A unique identifier for the scan, for example `SGTSCAN0000001`.
Status	Status of the scan
Completed	Completion state of the scan
Scheduled	Scheduled time of execution for the scan
Created by	User who initiated the scan
Action	Actions the scan took

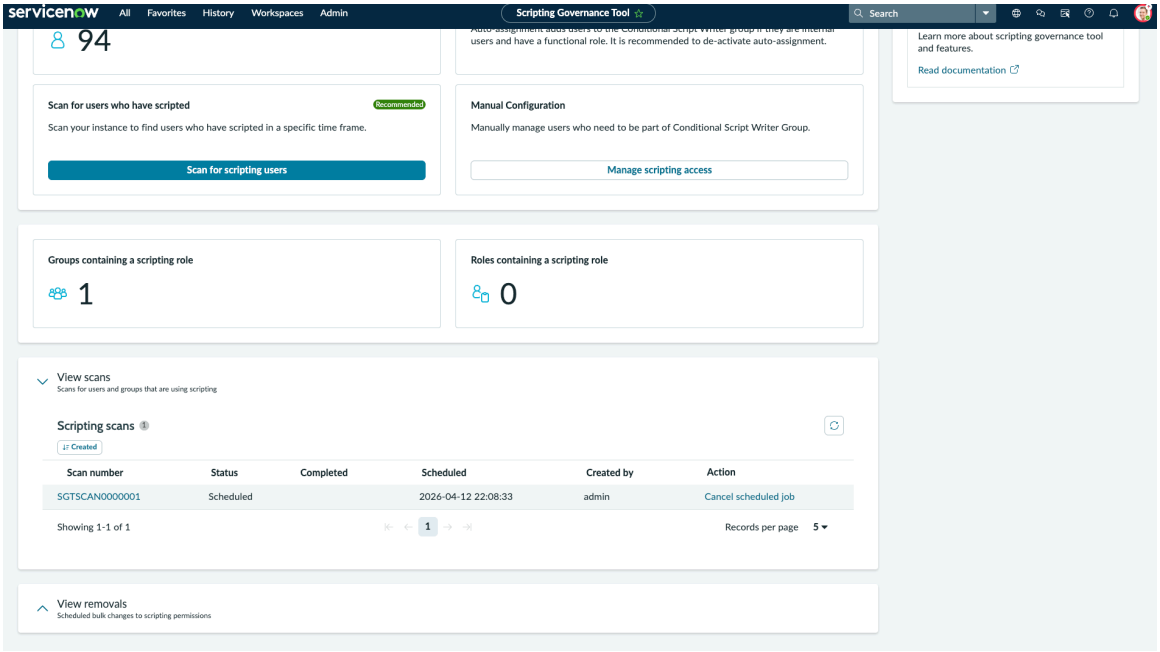

7. Select the scan record to open the scan results page.

   The scan results page opens and displays the following information:

   - Scan record number — a unique identifier for the scan, for example `SGTSCAN0000001`.
   - Date range — the start and end timestamps of the period that was scanned.
   - Users that have scripted — the total count of users in the **Conditional Script Writer** group and has <> role who modified at least one script field within the scanned period.
   - Users that have not scripted — the total count of users in the group who made no scripting changes within the scanned period. These users are candidates for access review and removal.
   - Users that have scripted table — a detailed list of active scripters, showing each user's Name, Last scripted date, Title, and Department.

   Note: If the **Users that have scripted count** is zero and the table shows `No data to display`, no users in the ServiceNow AI Platform have scripted within the selected time period.

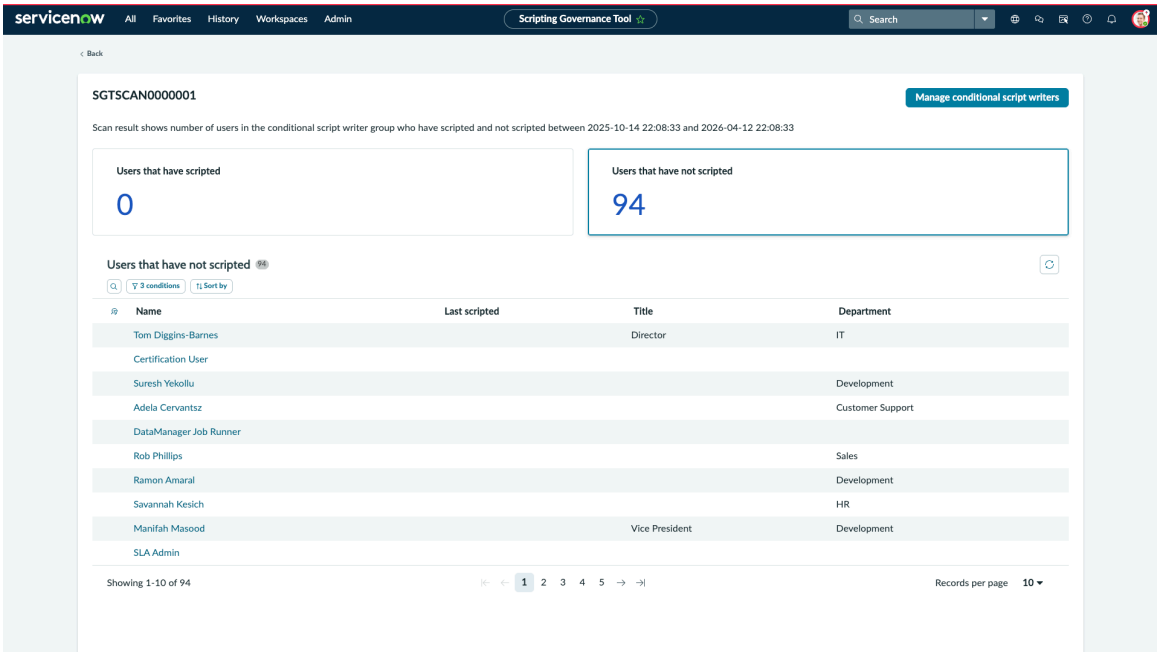

8. Review the results and identify users who require action.
9. To act on the results, select **Manage conditional script writers** on the scan results page. The **Manage users in the conditional script writer group** opens, where you can review group membership and remove users who no longer require scripting access using the **Schedule removal** option. To learn more about how to schedule removal, see Remove users from the Conditional Script Writer group.

Result
The scan produces a clear picture of active and inactive scripters on the ServiceNow AI Platform.

What to do next
After reviewing the scan results, remove any users from the **Conditional Script Writer** group who no longer require scripting access. Removing a user from the group does not affect any other permissions they hold on the ServiceNow AI Platform.

Remove users from the Conditional Script Writer group

Use the Manage scripting access to manually add or remove users from the **Conditional Script Writer** group to control who has scripting access.

Before you begin
Role required: security_admin

Important:

- You must have elevate your role `security_admin`. See Elevate to a privileged role.
- Until auto-assignment is turned off, a background job continues adding existing and new users to the **Conditional Script Writer** group, which assigns the `snc_required_script_writer_permission` role. It is recommended that you deactivate auto-assignment and manage the removal process to confirm only users with a legitimate scripting need remain in the group.

Procedure

1. Navigate to **All** > **System Security** > **Scripting Governance Tool**.
2. Select **Manage scripting access**.
3. Add the users who should retain scripting access in the **Users to keep in the conditional script writers group** field.

   Only users added to this field are retained in the group after the removal job runs. All other group members are removed. Search and select users by name.

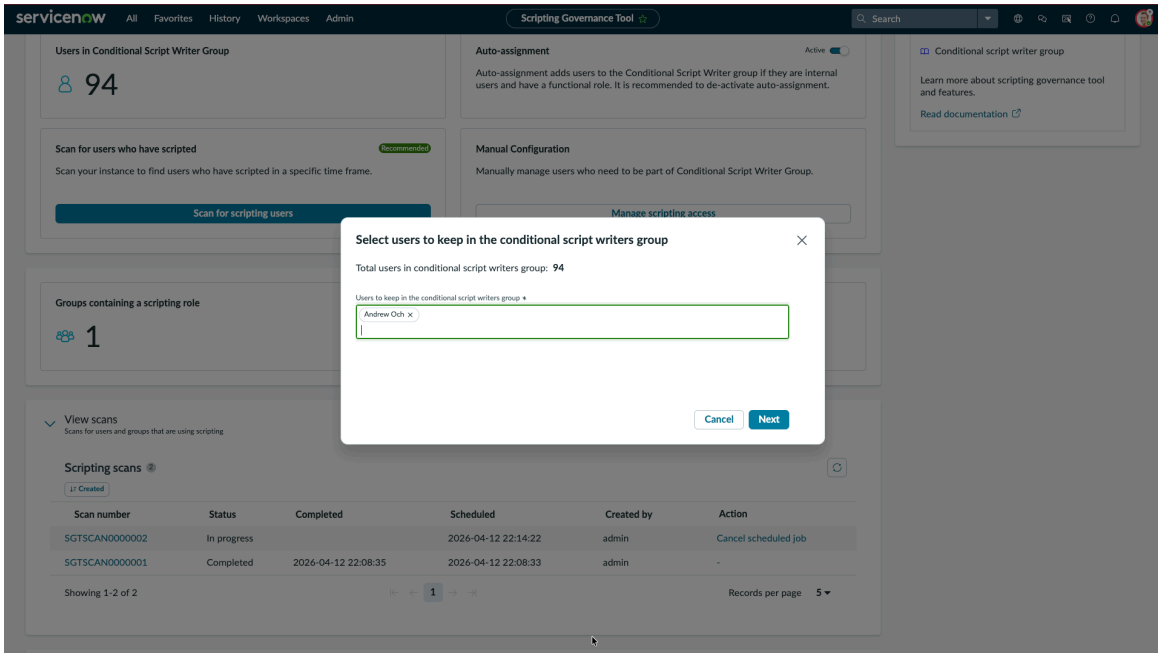

4. Select **Next**.
   The Manage users in the conditional script writer group page opens with the **Keep list** card displays the number of users retained, and the **Keep in the conditional script writer group** table lists the users you selected.
5. Review the keep list to confirm it is correct, then select **Schedule removal**.

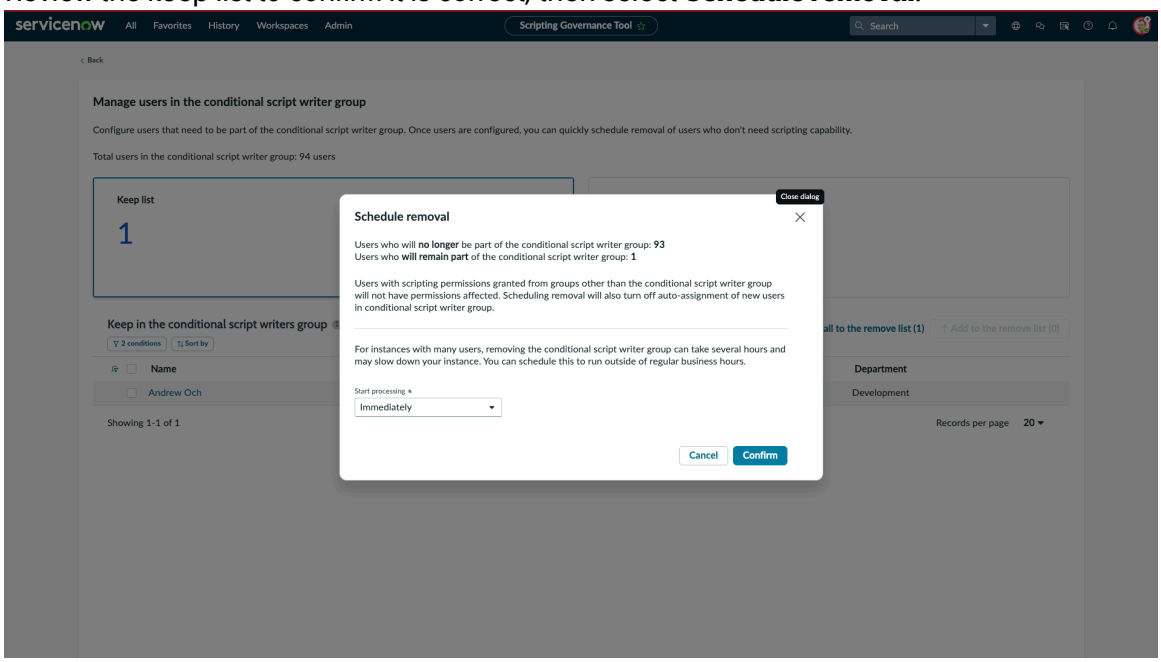

   The **Schedule removal** dialog opens and displays a summary of the pending changes:

6. Select when the removal job should run in the **Start processing** field:

   The following options are available:

   - Immediately — runs the removal job as soon as you confirm.
   - In 1 hour
   - In 2 hours
   - In 3 hours
   - In 5 hours
   - In 8 hours
   - In 13 hours

7. Select **Confirm** to schedule the removal.
   A banner confirms that the `Conditional script writer group member removal job is running or scheduled`. The removal job removes all users not on the keep list from the **Conditional Script Writer** group and disables auto-assignment.

Result

Users not included in the keep list are removed from the **Conditional Script Writer** group at the scheduled time and lose the `snc_required_script_writer_permission` role. Removal does not affect any other permissions those users hold on the instance. Auto-assignment is also disabled, preventing the background job from re-adding users to the group automatically.

What to do next

After the removal job completes, verify the outcome using the **View removals** section at the bottom of the Scripting Governance Tool dashboard.

Manage Scripting Governance Tool

Enable or disable the Scripting Governance Tool on your instance by running the appropriate script. Only users with the `security_admin` role can run these scripts and modify the associated properties.

Scripting Governance Tool states

Scripting Governance Tool operates in one of two states. The active state determines whether scripting governance policies are enforced and whether users are provisioned to the Conditional Script Writer group.

Note:

- Scripting Governance Tool is enabled by default. You can choose to disable.
- You must elevate your role to `security_admin` to enable or disable Scripting Governance Tool.

Scripting Governance Tool states

States	Behavior of Scripting Governance Tool
Enabled	• Scripting Governance Tool and all associated ACLs are active on the instance. • Users are evaluated against scripting access rules and assigned to the appropriate script writer groups. • The `security_admin` can run scans to identify users with scripting access and manage group membership. • Scripting governance policies are enforced across all applicable records and transactions. • Audit logs and visibility into scripting access are available to the security admin.
Disabled	• Scripting Governance Tool and all associated ACLs are deactivated on the instance. • No scripting governance policies are enforced. Users are not evaluated or assigned to script writer groups. • Existing group memberships from a prior enabled state are preserved but have no enforcement effect. • The Scripting Governance Tool interface remains accessible to the security admin but scanning and access management actions are inactive. • Scripting Governance Tool can be re-enabled by the `security_admin` at any time without data loss.

Disable scripting governance

To disable Scripting Governance, navigate **All** > **Scheduled Script Executions** (`sysauto_script_list.do`) and run the **Disable Scripting Governance** script to deactivate Scripting Governance Tool on your instance.

Running this script performs the following actions:

- Disables the `glide.security.scripting_role.provisioning_job_running` property.
- Disables the `glide.security.scripting_role.auto_provisioning` property.
- Disables the `glide.security.scripting_governance.enabled` property.

- Disables the **Add Users** to **Conditional Script Writer Group** and **Update Users** in **Conditional Script Writer Group** scheduled jobs.
- Removes all users from the **Conditional Script Writer Group** through a scheduled job.

Enable scripting governance

To enable Scripting Governance, navigate **All** > **Scheduled Script Executions** (`sysauto_script_list.do`) and execute the **Enable Scripting Governance** script to activate Scripting Governance Tool on your instance.

Running this script performs the following actions:

- Enables the `glide.security.scripting_role.provisioning_job_running` property.
- Enables the `glide.security.scripting_governance.enabled` property.
- Enables the **Add Users** to **Conditional Script Writer Group** and **Update Users** in **Conditional Script Writer Group** scheduled jobs.
- Schedules the **Add Users** to **Conditional Script Writer Group** job to run.

Machine identity access controls

Define and enforce granular access control policies on specific resources for integration users.

Machine identity access controls enable administrators to create granular access control policies for integration users. This additional layer of security allows admins to specify the exact resources (REST APIs, SOAP APIs, and tables) an integration user can access, ensuring tighter access governance.

Note: Applicable to users flagged with Web Service Access Only.

See Create a machine identity access control to learn how to create a machine identity access control

Create a machine identity access control

Enable administrators to define and enforce granular control for integration users by introducing User Access Profiles. This feature provides an additional layer of security and control, allowing admins to specify the exact resources (REST APIs and SOAP APIs) that an integration user can access, ensuring tighter governance and minimizing security risks.

Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **System Security** > **Machine Identity Access Controls**.
2. Select the **New** button.
3. Fill in the fields of the form.

Machine Identity Access Control fields

Field	Description
Name	Name of the access control record.

Field	Description
Application	Application containing the record.
Description	Description of the record.
Active	Determines if the policy is active
REST API Policy	Select the target REST API policy. **Note:** Select the and the icon to add a policy.
SOAP API Policy	Select the target SOAP API policy. **Note:** Select the and the icon to add a policy.
Tables	Select the tables this policy applies to
Applies to Child Table	Check this to apply the policy to child tables of the **Tables** field

4. Select the `Insert a row below` prompt and add users to apply the control to.
   You can add multiple users to the access control.

   Note: You can only select users with Web Service Access.

5. Select **Submit**.

Result
The following is an example of a machine identity access control form that has been filled out:

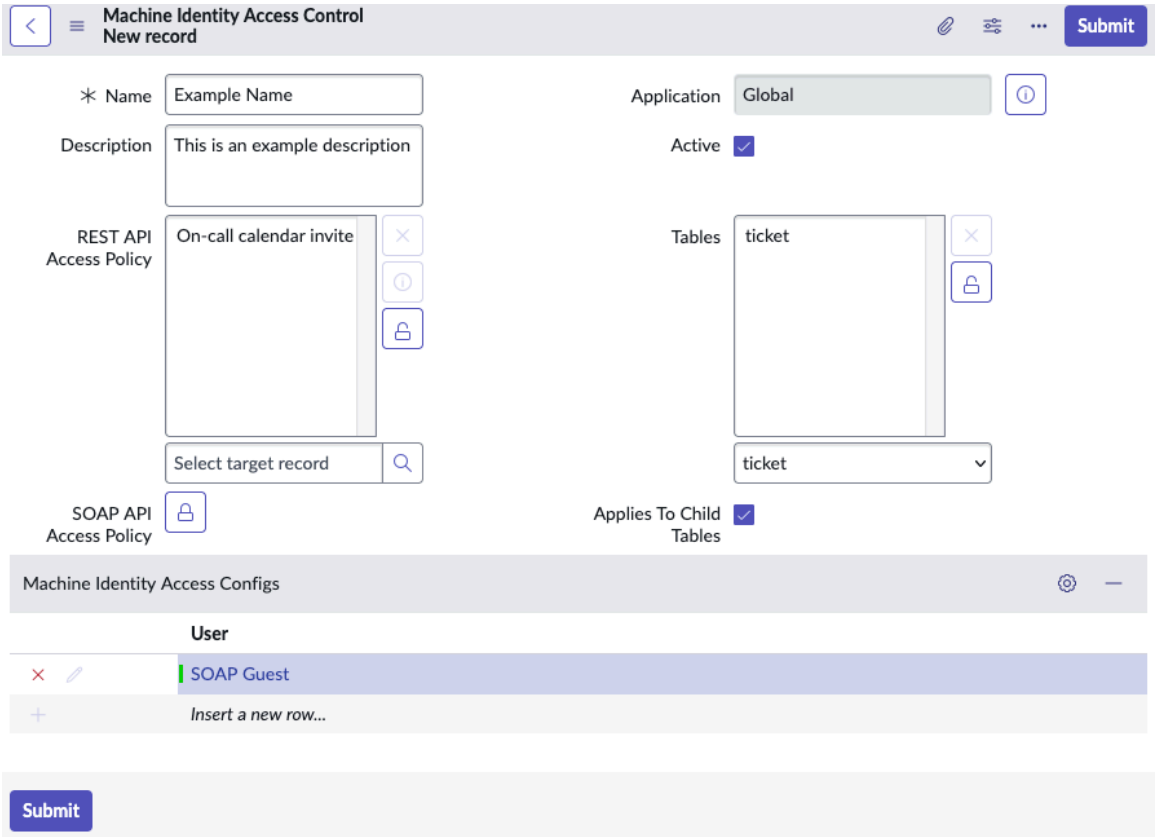

A user with an machine identity access control cannot access any other APIs (REST or SOAP) and will only be able to access the resources explicitly stated in the access control, even if they have the required roles.

Data filtration

Use data filtration to control access to tables and records based on subject attributes when performing read queries.

Important:

Starting with the Yokohama release, Data filtration is being prepared for future deprecation. It will be hidden and no longer activated on new instances. Consider using Security data filters to restrict access to records based on role, or security-attribute related assertions. For details see Data filtration.

Explore Data filtration	Create Data filtration Rules
Learn about Data filtration.	Create your own Data filtration rules to suit your needs.

Debug Data

Learn how to debug your Data filtration results.

Explore Data filtration

Use Data filtration to control access to tables and records based on subject attributes when performing read queries.

Data filtration is a separate form of access control designed to work along with the existing Access Control rules (ACLs) on your instance. Data filtration denies access to tables and records that do not match subject attributes defined by an administrator. Data filtration is designed to make auditing, reporting, and troubleshooting easier.

This is an optional feature that administrators can activate on their instance.

Data filtration features

Data Filters

Use data filters to grant access based on information within a record. Data filters use data in a tables field to determine whether a record is available to your users.

Subject attribute based condition builder

Use subject attributes to evaluate user role, group, subject criteria, or IP network address.

Data filtration uses a deny based model

Data filtration uses a deny based model to control access to records. With Data filtration, your instance denies access to records unless a record meets the criteria defined by Data filtration.

Data filtration enforcement

Data filtration rules run after the database query for read operations and are evaluated before ACLs. A record denied by any Data filtration rule will not proceed and be evaluated by ACL rules. Data filtration rule enforcement is consistent with that of read ACLs.

Data filtration and reporting

Data filtration and ACL's are both applied only when creating list view reports. Reporting does not apply access control when collecting aggregated data. In this case, neither Data filtration nor ACLs are checked.

For aggregated reports, Data filtration works in conjunction with existing *`Report_view access control list`* behaviors. See Report_view access control for further details on configuring these report controls.

Session debugging

Data filtration supports session debugging. Use session debugging to see which Data filtration records apply for a given query. Admins can use this information to troubleshoot user access to records.

Components of Data filtration

Data filtration works using the following record types:

Data filtration records

Create a Data filtration [sys_df_data_filtration] record to grant table access on your instance. The Data filtration record contains the **Data filter** and **Subject attribute** conditions described above to limit the scope of the rule and the affected users.

Subject criteria records

Subject criteria [sys_df_subject_criteria] records represent specific user attributes you can use to determine whether to grant access with a Data filtration rule. These attributes can be a user's groups, roles, or IP address. To create a subject criteria, you must create the subject criteria record, as well as criteria input and criteria conditions records. For details on this process, see Create subject criteria.

After creating a subject criteria records, you can apply them to a rule. This is done on the **Subject Condition** tab of your Data filtration rule.

Criteria input records examples

Example criteria input for all roles containing admin

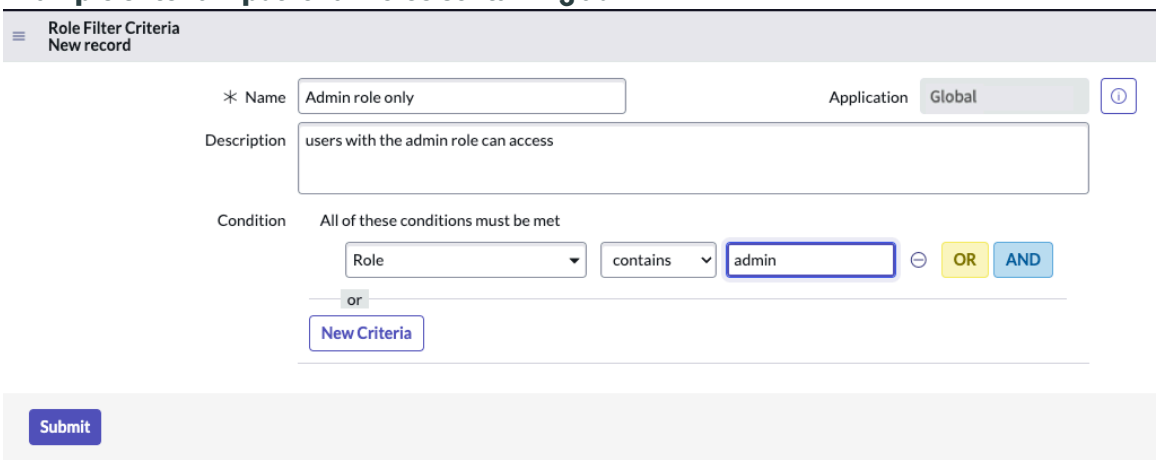

Criteria inputs [sys_df_subject_filter_criteria_m2m] are records that contain criteria to compare with the user. This can be a list of user groups or roles, an IP address range, or an IP address subnet. These records are used along with subject criteria condition records to evaluate against a user's groups, roles, or IP address to determine access to a table or it's records.

Subject criteria condition records

Example criteria condition using the Admins Only criteria input

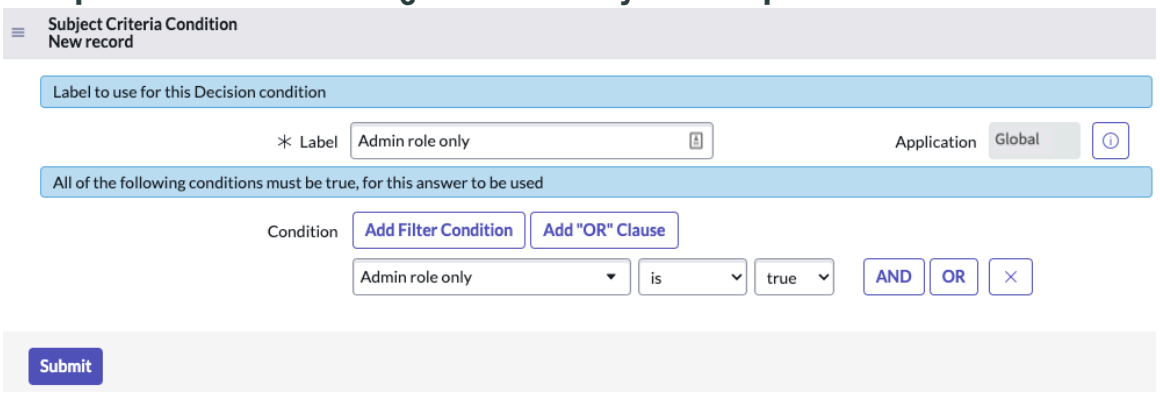

Use subject criteria condition [sys_df_subject_criteria_condition] records to define how to compare user attributes with the roles, groups, or IP addresses defined in you criteria inputs. You can use multiple criteria inputs in a single subject criteria condition to further narrow down access to your records.

Activate data filtration

Learn how to activate data filtration on your instance.

Before you begin

Important:

Starting with the Yokohama release, Data filtration is being prepared for future deprecation. It will be hidden and no longer activated on new instances. Consider using Security data filters to restrict access to records based on role, or security-attribute related assertions. For details see Data filtration.

Role required: security_admin

Procedure

1. Navigate to **All** > **System Definition** > **Plugins**.
2. Use the search field to find the Data Filtration (com.glide.data_filtration) plugin.
3. Click **Install**, and then in the Activate Plugin dialog box, click **Activate**.

Create data filtration rules

Learn how to create data filtration rules to grant your users' access to records are tables.

Before you begin

Role required: security_admin

Note: To create or modify data filtration rules you, must elevate to the privileged role. For details on this process, see Elevate to a privileged role.

Procedure

1. Navigate to **All** > **Data Filtration** > **Data Filtration Records**.
2. Click **New** in the **Data Filtration** list.
   A new data filtration form displays.
3. In the form, fill in the fields as needed.

Data filtration form

Field	Description
Table	Table to which this data filtration rule applies. **Note:** Non-maint users cannot create data filtration on some tables, to work-around this remove the tableChoicesScript=DataFiltrationTableList attribute, but make sure that no filters will be created on any of the sys_df_xxx tables or any tables in the sys_df_table_exclusion.
Active	Sets the data filtration rule as active. **Note:** Keep data filtration rules inactive until you are ready to test to avoid unintentionally locking users out of records.
Description	Description of the data filtration rule.
Cascading	Select to set the data filtration rule to apply to extended tables. For example, you select the Task[task] table, and enable cascading. In this case, the data filtration rule also applies to all tables extended from task, such as Incident[incident] and Change Request[change_request]. For detail on table extension, see Table extension and classes **Note:** This field is enabled by default.

4. **Optional:** To narrow the scope of the rule fill in the **Conditions** fields as needed.

Field	Description
Subject Condition	All conditions must be met for access.
Security Attribute Condition	All conditions must be met for access. **Local** The attribute is defined only in the scope of the data filtration rule. **Existing** The attribute is defined by reference to an already existing **Security Attribute**
Data condition	Defines the conditions for data to be subject to the rule. **Note:** An empty Data Condition will apply to all records in the selected table.

5. Select **Save** from the form menu.
   After you have saved your data filtration rule, this rule automatically applies to all records on the selected table, unless specified otherwise by the data condition.

Add a data filter for your data filtration rule

You can optionally use a data filter to narrow the scope of your data filtration rule to apply only to specific records on a table.

Before you begin

Role required: admin

Procedure

1. On your data filtration record, open the **Data Filter** tab.
2. Use the condition builder to filter the table records their field values.
   The data filter uses the same condition builder used in other parts of the platform. For details on using this interface, see Condition builder.

   Important: The **Data Filter** tab appears empty until you select a table in the **Table** field.

3. Use the **Preview** button to see a count of how many records match your data filter.
4. Select the number of records to open a list of the matching records.

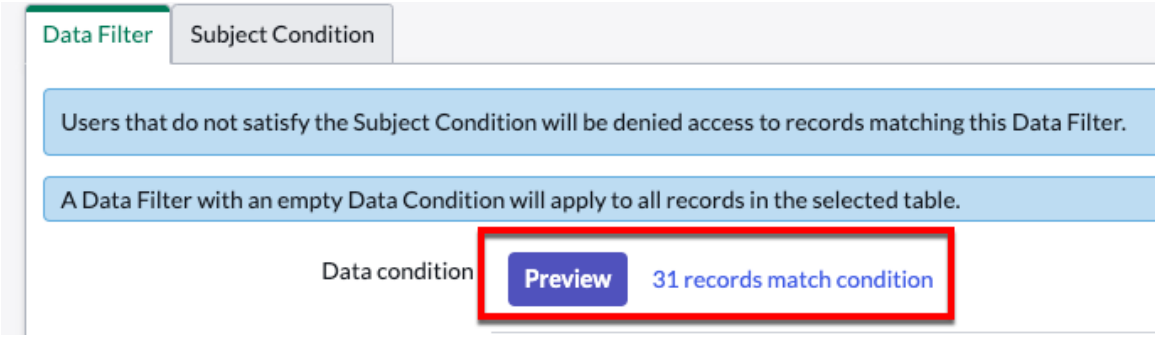

5. Select **Save**.

Example:

This example shows a data filtration rule for the Incident[incident] table. The data filter is set to select all active records that are not in the **Security** category. With this rule active, users can see these records. See the section below to further using criteria outside the contents of the record.

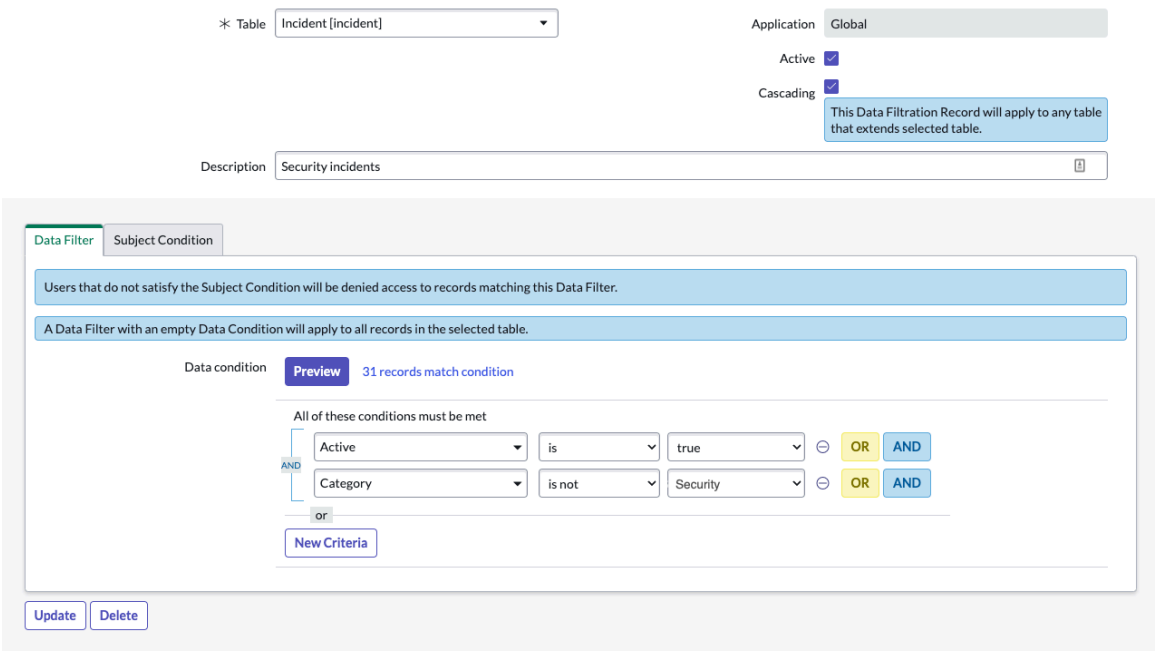

Important:

The **not** operation in your conditions may return unexpected results, depending on the type of database your instance uses. For example, take the following condition:

For this condition, the expected result would be that the result set would be all records where the company is not `ServiceNow` and all records that do not have a value in the **company** field. Instances using databases other than MySQL and Maria do not return values records with an empty **company** field. When using **not** queries for these instances, include conditions to ensure empty values are returned.

Add subject attributes to your data filtration rule

Optionally, use subject attributes to narrow the scope of your data filtration rule based on attributes such as IP network address, user groups and roles, or subject criteria.

Before you begin

Role required: admin

Procedure

1. On your data filtration record, open the **Subject Criteria** tab.
2. Use the condition builder to filter the table records based on one or more of the following criteria.

Subject criteria options

Option	Description
Network Criteria	Allows access to records based on network IP range or IP subnet.
Subject Criteria	Allows access to records based on subject criteria. Select a subject criteria record to apply it's conditions to your data filtration rule. For details on creating subject criteria records, see Create subject criteria.

Option	Description
Subject Group	Allow access if the user is a member of a specific group. Select a group from the **Groups**[sys_user_group] table.
Subject Role	Allow access if the user is a member of a specific group. Select a group from the **Roles**[sys_user_role] table.

Important: Subject criteria conditions only support the **is** operator.

3. After adding your subject criteria click **Save**.

Create subject criteria

Create subject criteria records for use in data filtration rules.

Before you begin

Role required: security_admin

Important: To create or modify data filtration rules you, must elevate to a privileged role. For details on this process, see Elevate to a privileged role.

Procedure

1. Navigate to **All** > **Data Filtration** > **Subject Criteria**.
2. In the **Subject Criteria** list, click the **New** button.
   A new subject criteria form displays.
3. In the form, fill in the fields as needed.

Subject criteria fields

Field	Description
Name	Name for the subject criteria.
Application	Scoped application for the subject criteria. This field is read-only, and automatically populates with the current scoped application.
Description	Description for the subject criteria.

4. Right-click the form header and click **Save** from the context menu.
   After saving, the **Criteria Inputs** and **Criteria Conditions** fields appear.

Create a subject criteria input

Create a subject criteria input to define criteria that your data filtration rules filter against.

Before you begin

Role required: security_admin

Procedure

1. On your subject criteria record, open the **Criteria Inputs** tab.
2. In the **Criteria Inputs** list, click **New**.

3. Select the policy Input for the criteria you want to create.

Policy inputs

Policy input	Description
IP filter criteria	Create a policy input based on IP address
Role filter criteria	Create a policy input based on use role
Group filter criteria	Create a policy input based on user group

A **IP Filter Criteria**, **Role Filter Criteria**, or **Group Filter Criteria** form displays, depending on your selection in this step.

4. In the filter criteria form, fill in the fields as needed.

Filter criteria fields

Field	Description
Name	Name for the filter criteria
Application	Scoped application for the subject criteria. This field is read-only, and automatically populates with the current scoped application.
Description	Description of the filter criteria

5. Below the form fields are tabs used to define the IP addresses, groups, or roles for the filter criteria inputs.

Filter criteria fields for specific input types

Policy input type	Description	Creation
IP filter criteria	User IP filter criteria to create a range or subnet of IP addresses. Your subject criteria can then compare with the user's IP address against this range or subnet.	Use the **IP Range** or **Subnet (CIDR)** to define IP addresses for your input. **IP Range** In the **IP Ranges** list, double-click **Insert a new row**, and enter a starting IP address in the **Start IP**. Then, press the tab key, and enter an ending IP address in the **End IP** field. Finally, press Enter to save the list entry. **Subnet (CIDR)** In the **Subnets** list, double-click **Insert a new row**, and enter a network IP in the **Network IP** field. Then, press the tab key, and

Policy input type	Description	Creation
		enter a netmask in the **Netmask** field. Finally, press Enter to save the list entry. **Note:** Scheduled jobs being triggered by a scheduler are not intended to have data filtered using a network criteria since they do not have the context of the requesting client IP. A more appropriate type of filtration may be the Role/Group subject condition.
Role filter criteria	Use role filter criteria to create a selection of roles. Your subject criteria can then compare with the user's assigned roles against this selection.	Use the condition builder in the **Condition field** to select roles for your input.
Group filter criteria	Use group filter criteria to create a selection of user groups. Your subject criteria can then compare with the user's assigned groups against this selection.	In the **Groups for criteria** table, double-click **Insert a new row**, and select a user group. Press Enter or click the green check mark icon to save the group. Click the **Insert a new row** text below the first entry to create additional entries.

6. After defining your input, click **Submit**.

Note: In addition to creating criteria inputs for your subject criteria, you can also add existing ones. Click **Edit** in the **Criteria Inputs** tab, and select from any existing input.

Create a subject criteria condition

Create a condition to compare a user information against existing subject criteria input.

Before you begin

Role required: security_admin

Criteria conditions compare user attributes against existing criteria inputs to determine whether to allow access to records. In order to create a criteria condition, you need to have created subject criteria. For details on that process see, Create a subject criteria input.

Procedure

1. On your subject criteria record, open the **Criteria Conditions** tab.
2. In the **Criteria Conditions** list, click **New**.
3. On the **Subject Criteria Condition** form, fill in the fields as needed.

Subject criteria condition form

Field	Description
Label	A descriptive label for your condition
Application	Scoped application for the subject criteria. This field is read-only, and automatically populates with the current scoped application.

4. Create a condition for your subject criteria condition using the condition builder by selecting from the following condition options.
   These condition options include any subject criteria inputs you have created.
5. **Optional:** Create more conditions by clicking the **Add filter condition** or **Add "OR" clause** buttons.

   Note: Unless your conditions are separated by an **or** clause, all conditions must evaluate to true in order for the subject criteria condition to evaluate to true.

6. Click **Submit** to save the subject criteria condition.

What to do next

Use subject criteria in your data filtration rules to limit access to tables and records. For details on how to use subject criteria in data filtration rules, see Create subject criteria.

Data filtration debugging

Use the session log to see how data filtration affects your records and debug user access issues.

Before you begin

Role required: admin

Output information appears in the session logs when users access records. You can use this logging information along with impersonation to learn why users see or do not see records. You can then use that information to adjust your data filtration rules and ensure that users only see what you intend.

Procedure

1. Navigate to **All** > **System Security** > **Debugging** > **Debug All Security**.
   The **Script Debugger** opens in a new browser tab or window.

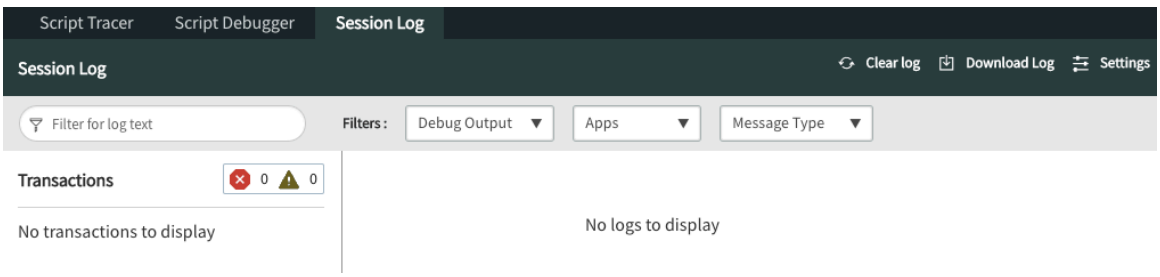

2. In the **Script Debugger** window, select the **Session Log** tab.
3. In another browser tab or window, impersonate a user to troubleshoot that user's access.

Note: Impersonation allows an admin to see an instance with another users settings and access. For details on impersonation, see Impersonate a user.

4. While impersonating a user, access a list or record where you see unexpected behavior. This behavior may be a record the user sees, but should not see, or a list of records that are not appearing as expected for a user.
   After accessing records with the impersonated user, you should begin to see output in the session debugger.
5. Look for data filtration information in the session debugger.

This example shows two log messages where a data filter denied access to records. The log entries appear as red text, and include why the data filter denied access, as well as the sys_id of the data filter. You can click on this sys_id to open the data filtration record.

This example shows a log message where a data filter allowed access to a record. These log entries appear as green text. As with the first message, you can click this sys_id to open the data filtration record.

6. Use this information to make any adjustments to your data filtration rules.
   Repeat these steps to refine your rules and give users the access they need.

Security data filters

Security data filters restrict access to records based on role, or security-attribute related assertions.

Exploring security data filters

Security data filters enable access restriction to records based on a users' role, or other security attribute related assertions. Security data filters ensure only authorized users can view records regardless of how data is accessed.

Security data filters are applied before a query is executed so restricted data never leaves the database. In contrast conditional ACLs filter data after a query is executed possibly leaking data.

Note: Pair security data filters with Deny-Unless ACL to ensure consistent security

Features of security data filters

The key features of security data filters are:

- Security data filters are applied in-query.
- Security data filter conditions AND to the query on the target table and with each other.

- Security data filters are not checked by `canRead`. See When to use security data filters for more details
- Data filter scoping rules are based on the scope of the table, data filters do not follow ScopeMaster or sys_scope scope rules

Security data filter application and enforcement

Generally security data filters are applied after absolute ACLs (also called table-level ACLs), and after row ACLs. Security data filters are applied by default, and impact system behavior if not used carefully. See Default security filters for a list of the default security data filters.

Security data filters are applied only to GlideRecordSecure, GlideRecordSandbox and GlideAggregateSandbox queries by default. There are two new GlideRecord APIs `enableSecurityFeature` and `disableSecurityFeature` that can be used in both Java and server-side scripts to enable or disable data filters for a specific query.

Important: You should explicitly enable data filters for user-facing queries that are not using GlideRecordSecure.

When to use security data filters

Security data filters are best suited to:

- Prevent sensitive data from leaving the database
- Suppress the "rows hidden by security" message
- Prevent sensitive data from leaking through reports

When not to use security data filters

Security data filters should not be used:

- As visibility control
- As replacement for ACLs
- With a large number of filter conditions
- On unindexed columns

Security data filter behavior

Multiple security data filters combine together for evaluation, like an `AND` combines operands. As an example,given three security data filters:

- Filter 1: `active=true
- Filter 2: `priority=1`
- Filter 3: `state=open`

And an initial query:

```
SELECT * FROM task WHERE state != closed AND active = true AND
        priority = 1
```

The final query is:

```
SELECT * FROM task WHERE state != closed
        AND active = true AND priority = 1 AND state = open
```

See the diagram below for a visual representation of this example:

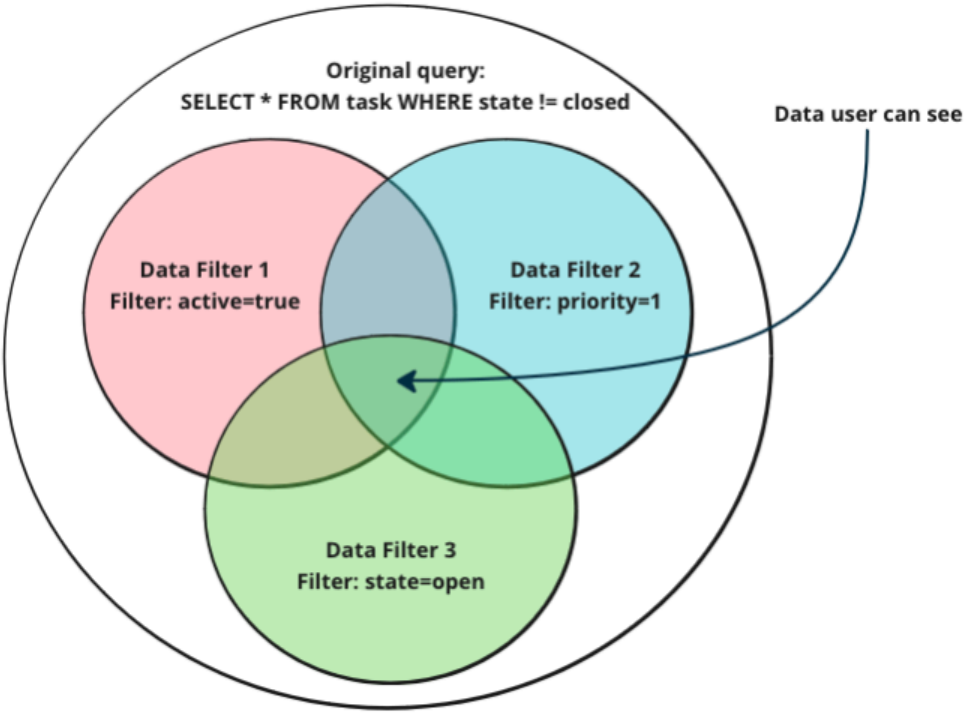

One important difference in how security data filters and ACLs are applied is, data filters on a child table do not apply to the parent table when data is queried from the parent table. Add a data filter on both child and parent tables to restrict access to records in the parent table. The diagram below highlights the

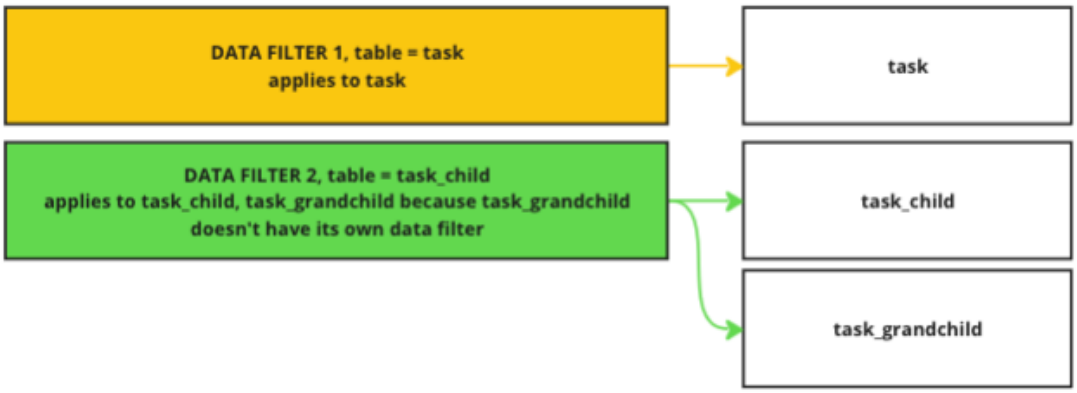

hierarchy:

Note: A common solution for this is to add a data filter on the parent that completely hides child records in the parent table.

Data filters are applied with scoping rules similar to ACLs, but with some key differences because data filters apply before-query.

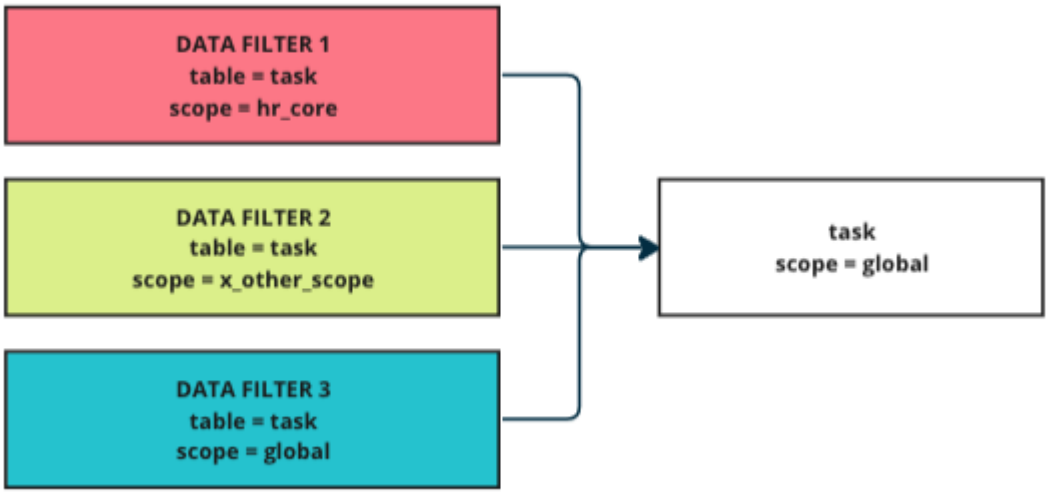

Create a security data filter

Learn how to create security data filter rules to grant your users' access to records are tables.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Security** > **Security Data Filters**.
2. Click **New** in the **Security Data Filters** list.
3. In the form, fill in the fields as needed.

Data filtration form

Field	Description
Table name	The table that the security data filter applies to.
Description	Description of the security data filter rule.
Active	Sets the security data filter rule as active. **Note:** Keep security data filter rules inactive until you are ready to test to avoid unintentionally locking users out of records.
Show in UI	Determines whether a notification will be displayed in the UI if the security data filter applies to a query
Application	The application scope of the security data filter.
Mode	The mode of the data filter.
Filter	The filter condition that determines which records the data filter applies to
Security Attributes	The security attributes that control if the data filter will apply or not

4. Select **Save** from the form menu.

Result
After you have saved your security data filter rule, this rule automatically applies to all records on the selected table, unless specified otherwise by the data condition.

Default security filters

Generally data filters are applied after absolute ACLs (also sometimes called table-level ACLs), and after row-level ACLs. They are applied by default, and can be impactful to system behavior if not used carefully.

Default applied security data filters

Security data filters are applied by default to the following places:

Category	Application
List and Forms	• UI16 • Workspaces • Service Catalog • Service Portal • Mobile
Reports and Dashboards	• Reports • Data visualizations • Core UI Dashboards • Platform Analytics experience dashboards
Data Export	• Export XML • CSV • Excel • JSON • PDF
Flow Engine	Record Lookup steps
Search	• AI Search • Text Search
GlideRecord	• GlideRecordSecure • GlideRecordSandbox • GlideAggregateSandbox

Category	Application
REST and GraphQL APIs	• REST table API • REST stats API • GraphQL table and stats API equivalent

Security Roles

Security Roles provide added security, every user must have at least one role so that the instance can distinguish between internal and external users.

Explore Explicit Roles

Learn the key features and business value of Explicit Roles.

Explore Elevated Privilege

Learn how the Elevated privilege role enables session based privileges.

Explicit Roles

You can give both internal users and external users access to your instance. However, you might not want both types of users to have the same level of access. To provide added security, every user must have at least one role so that the instance can distinguish between internal and external users.

As of the Paris release, no user can have both of the explicit roles (snc_internal and snc_external). Groups and role containment cannot include both roles, either, since that would cause any group member or user who is assigned to such a group or such a role to automatically have both roles. The ServiceNow AI Platform aborts any operation that would create such a scenario.

Note: You cannot add any other roles as contained roles to the "*snc_external*" or "*snc_internal*" roles.

External users must obtain, at minimum, the snc_external role. The snc_external role indicates that the user is external to your organization. They should not have any access to resources unless explicitly allowed through ACLs for the snc_external role or additional roles that inherit the snc_external role. By default, users with the snc_external role cannot access:

- Tables without the role that inherits the snc_external role or the public role.
- Non-record type resources, such as processors and UI pages without granting access to the snc_external role or a role that inherits the snc_external role.
- Platform Analytics dashboards.

Do not mark the snc_internal role as elevated. Otherwise, internal users cannot access the instance.

Explicit Roles plugin

When the Explicit Roles plugin is activated:

- All users must have the snc_internal role to access internal resources or the snc_external role to access external resources. Users without either explicit role can access only public resources.
- All existing users are automatically assigned the snc_internal role. This role does not change existing access levels or system behavior. Instead, it provides a category to differentiate internal users from external users. All internal users maintain the same level of access as before the plugin was activated.

  Tip: To prevent changing existing functionality for users, activating the Explicit Roles plugin assigns the snc_internal user role to all existing users in the instance. This includes any external users added before the Explicit Roles plugin was activated. After the Explicit Roles plugin is activated, do the following for all external users added before the Explicit Roles plugin was activated:

  - Remove the snc_internal role.
  - Add the snc_external role.

  Doing the preceding ensures that external users added before activating the Explicit Roles plugin do not have access to internal resources that should be available only to internal users.

- Newly created users are automatically assigned the snc_internal role when they first attempt to log in to the instance, unless they have been explicitly assigned the snc_external role. You can add the snc_external role to a new user before they first log in to the instance to provide external user rights.

  Important:

  Activate this plugin during a maintenance window or when few users are logged in. Users currently logged in when the plugin is activated will not be dynamically assigned the snc_internal role. Rather, users must log out and log back in to be assigned the snc_internal role. Once the plugin is activated, you can add or remove the snc_internal and snc_external roles at any time to change user rights.

  After the plugin is activated, any time a user logs in, the user is given the snc_internal role if the account does not already have that role, or the snc_external role. This includes users logged in via impersonation.

- All existing ACLs that do not have a role requirement are automatically assigned the snc_internal role. Because both existing ACLs and users are assigned the snc_internal role, existing access levels do not change.
- Newly created ACLs that do not have a role requirement are automatically assigned the snc_internal role. This role assignment does not apply to a newly created ACL with a role assigned.
- For all existing Processor [sys_processor] records or newly created Processor [sys_processor] records with **Type=script**, the snc_internal role is automatically added to the **Roles** field if the field is empty.
- To restrict access to UI pages to internal users, the plugin automatically assigns the snc_internal role to the * ACL with a **Type** of **ui_page**.

- To restrict access to processors to internal users, the plugin automatically assigns the snc_internal role to the * ACL with a **Type** of **processor**.
- External users must obtain, at minimum, the snc_external role to access the instance. This role must be manually granted to external users. Access to records is granted through ACLs.

Do not move system update sets among instances with and without the Explicit Roles plugin enabled. For more information, see System update sets.

Note: This plugin also requires the Contextual Security Manager plugin.

glide.security.explicit_roles.do_not_fix behavior

As of the Xanadu release the `glide.security.explicit_roles.do_not_fix` has been adjusted with changes to snc_internal. The snc_internal role is now the same both in memory and in sys_user_has_role. The new behavior for `glide.security.explicit_roles.do_not_fix` is:

glide.security.explicit_roles.do_not_fix new behavior

Value	Result
False	Add snc_internal role both in memory and sys_user_has_role
True	Do not add snc_internal role in memory or sys_user_has_role

To exclude snc_internal role for certain users, use the `glide.security.explicit_roles.ignore.snc_internal.exclude_role_list` property.

Tip: To revert to previous `glide.security.explicit_roles.do_not_fix` behavior use the `glide.security.explicit_roles.do_not_fix_in_memory` property.

Providing table access to external users

You can provide external users access to a table by adding a role to the table that inherits the snc_external role. For more information, see Provide external users access to a table.

The hasRoles() method

The `hasRoles()` method is still available, but is deprecated in the Geneva release. Use the `hasRole(role name)` method instead.

If you do use the `hasRoles()` method, note these changes:

- This method automatically excludes the default snc_internal role when it checks for roles. This means that if a user has only the snc_internal role, the `hasRoles()` method still returns **false**.
- If the user has the snc_external role, the method returns **false** because the instance considers external users to be without a role.

Mutual exclusion: snc_external versus snc_internal

The ServiceNow AI Platform prevents users from having both the snc_external role and the snc_internal role. The ServiceNow AI Platform applies this mutual exclusion everywhere in the system and writes error messages to the logs for each conflict.

Note: ACLs can have both roles if the ACL resources should be accessible to all users.

Example: Adding both explicit roles to a user (direct collision):

1. Assign user Abel Tuter the snc_internal role.
2. Assign user Abel Tuter the snc_external role.

Result: Adding the snc_external role fails because Abel Tuter has the snc_internal role.

Example: Adding both explicit roles to a group (direct collision):

1. Consider a group called Test Group that currently has no explicit roles assigned to the group.
2. Add Abel Tuter to the Test Group.
3. Add the snc_external role to Test Group.

Result: Adding the snc_external role fails because Abel Tuter already has the snc_internal role and can't have both roles.

Example: Adding an explicit role to a group where a group member has the conflicting explicit role (indirect collision):

1. Assign user Abel Tuter the snc_internal role.
2. Consider a group called Test Group that currently has no explicit roles assigned to the group.
3. Add Abel Tuter to the Test Group.
4. Add the snc_external role to the Test Group.

Result: Adding the snc_external role to the group fails because Abel Tuter would inherit the snc_external role through group membership. Both explicit roles would be assigned to the same user, which isn't allowed.

For other examples, see the following table:

Role	Attempted action	Result
Direct collision		
The user has the snc_internal role.	Add the snc_external role.	The action is aborted.
The user has the snc_external role.	Add the snc_internal role.	The action is aborted.
The user has no explicit role.	Add the snc_internal or snc_external role.	The role is added.
The user has both explicit roles (existing collision).	Add the user to a group with no roles.	The action is aborted.
A role not associated with any users	Add the snc_external role.	The action is aborted.

Role	Attempted action	Result
contains the snc_internal role.		
A role not associated with any users contains the snc_external role.	Add the snc_internal role.	The action is aborted.
A role contains both explicit roles (existing collision).	Add the role to a user, role, or group.	The action is aborted.
A group with no members has the snc_internal role.	Add the snc_external role.	The action is aborted.
A group with no members has the snc_external role.	Add the snc_internal role.	The action is aborted.
A group with no members has no roles.	Add the snc_internal or snc_external role.	The role is added.
Indirect collision		
Role containment with collision	1. Grant a role called Test Role to a user with the snc_internal role. 2. Add the snc_external role Test Role.	The action is aborted.
Role containment without collision	1. Grant a role called Test Role to a user with no roles. 2. Add the snc_external role to Test Role.	The role is added to both the user and Test Role.
Group containment with collision	1. Add a user who has the snc_internal role to a group called Test Group 2 (child of Test Group 1). 2. Add the snc_external role to Test Group 2. 3. Add the snc_external role to a parent group called Test Group 1 (parent of Test Group 2).	The action is aborted.

Role	Attempted action	Result
Group containment without collision	1. Add a user with no roles to a group called Test Group 2 (child of Test Group 1). 2. Add the snc_external or snc_internal role to Test Group 1 (parent of Test Group 2).	The role is added to the parent group, the child group, and the user.
Group containment plus role containment with collision	Add contains_external to Test Group 1, the parent of Test Group 2.	Test Group 1 and Test Group 2 both get contains_external, but don't explicitly get the snc_external role.
	Add the snc_internal role to Test Group 2, the child of Test Group 1.	The action is aborted.
Group parent change plus group containment	1. Remove Test Group 1 as the parent of Test Group 2. 2. Add the snc_internal role to Test Group 1. 3. Add the snc_external role to Test Group 2. 4. In Test Group 2, set Test Group 1 as the parent group and save.	The action is aborted. Repeat for already nested groups, with the same expectation.

The cause of an aborted action appears in the error message and must be addressed before another attempt succeeds.

For direct cases, such as adding an explicit role to an individual user, verify which explicit role the user should have. If the user has the wrong explicit role, it must first be removed, and then the correct explicit role must be added.

For indirect cases,such as adding an explicit role to a group (so that a group member would be assigned both explicit roles), evaluate whether that user should be in the group. Also determine whether the group should be given the explicit role, including any inheritance through group hierarchy and role containment.

Note that the ServiceNow AI Platform reports only the first potential collision encountered. If repeated attempts continue to fail after remediation, with a new root cause each time, re-evaluate the relevant user/group/role interdependence more broadly. You may want to rethink how groups and role containments are structured.

Request Explicit Roles

Activate Explicit Roles by requesting the Explicit Roles plugin (com.glide.explicit_roles) through the Now Support Service Catalog.

Before you begin

Role required: admin

Important: Activate this plugin during a maintenance window or when few users are logged in. Users currently logged in when the plugin is activated will not be dynamically assigned the snc_internal role. Rather, users must log out and log back in to be assigned the snc_internal role. Once the plugin is activated, you can add or remove the snc_internal and snc_external roles at any time to change user rights.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Select **Request plugin** to open the **Activate Plugin** form on Now Support.
3. On the **Activate Plugin** form, provide the following information.

Activate Plugin form

Field	Description
What is your target instance	Select the instance that you want to activate the plugin on.
Which plugin would you like to activate	Select the name of the plugin to activate. **Note:** If the system doesn't list the plugin you want or if you're activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	Select the date and time to activate the plugin.

Example

For example, see the following form to activate the Event Management plugin on an instance named SNC Instance.

4. Select **Submit**.
   After the maintenance window, the system installs the plugin on your instance. To confirm the installation, go to the Installed tab in the Application Manager.

Elevated privilege roles

Elevated privilege roles require you to manually accept the responsibility of using the role before you can access the features of the role.

By default, you do not have elevated privilege roles upon login. You must manually elevate to the privilege of the role. An elevated privilege role lasts only for the duration of your user session. Session timeout or logout removes the role.

You can designate any role as an elevated privilege role, and then assign that role to one or more users. Do this when you want to restrict users from having access to the rights that the role provides immediately after login. You can designate the privilege role on the Role form. See Create a role for instructions.

To use an elevated role, you must meet these conditions:

- The elevated role must be assigned to you.
- You must manually elevate to a specific elevated role to get its privileges, even if you are already elevated to a second elevated role that contains the first elevated role.

  For example, if elevated role A contains elevated role B, even if you elevate to role A, you must still elevate to role B to get its privileges.

The admin role

To grant the admin role to a user, the granting user must also have the admin role. For example, a user with only the user_admin role cannot grant the admin role to other users.

- Non-admin users cannot add a user to a group that contains the admin role.
- To grant the security_admin role to a user, the granting user must also have the admin role and must elevate to the security_admin role before granting the security_admin role to other users. A user with only the admin role cannot grant the security_admin role to other users.
- A user without the security_admin role cannot add a user to a group that contains the security_admin role.

Warning: The use of elevated privilege on an admin role is not supported. Instead, require admins to manually elevate, see Force administrators to manually elevate

The security_admin role

In the base system, the security_admin role is the only role that has elevated privileges. This role is automatically assigned to the user who is the default System Administrator (admin) user. It provides access to ACLs and High Security Settings.

Roles assigned to the System Administrator (admin) user

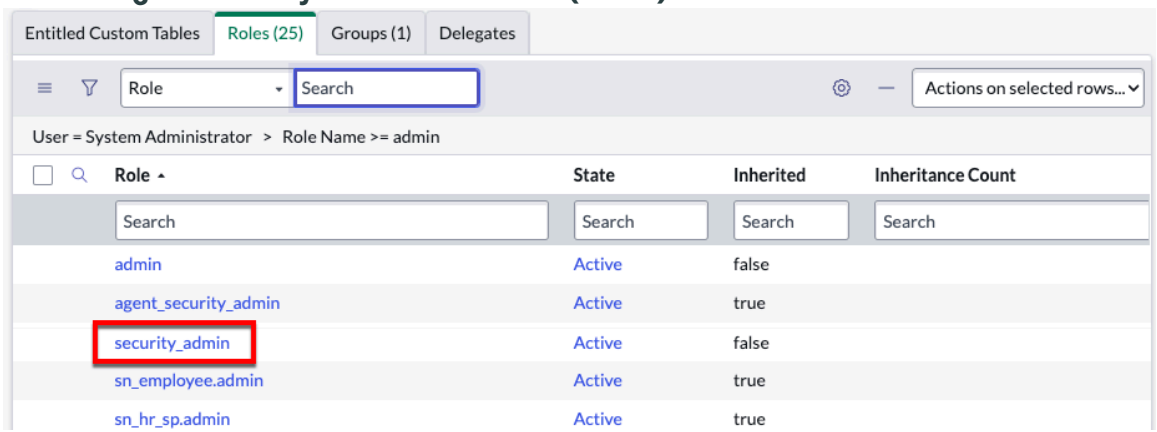

Note: To see this role, you must actually elevate to the security_admin role first. If you are logged in as the System Administrator (admin) user only, you cannot see the security_admin record in the list of roles.

Security_admin role

The security_admin role is an elevated privilege role provided with High Security Settings that lets users create and change access controls and change High Security Settings.

In the base system, only the default System Administrator (admin) user has the security_admin role. Since it requires elevating privileges, the admin user does not have this role at login. After

elevating privileges, the admin user has the security_admin role for the duration of the user session. See Elevate to a privileged role for more information.

To maintain high security, the security_admin role requires elevating privileges. Limit the users and groups to which you assign this role.

Elevate to a privileged role

The base system admin can elevate to a privileged role to have access to the features of High Security Settings.

Before you begin

Role required: admin

Note: If you grant additional users the admin role, they cannot elevate to a privileged role. Only the base system admin can elevate.

Procedure

1. In the banner frame, click your image, or click your initials if you do not have an image uploaded.
2. Select **Elevate Roles**.
   A dialog box appears that contains available roles for elevation.
3. Select the elevated roles to be assigned and click **OK**.
   This role grants the user elevated privileges to all resources controlled by the role for the remainder of the session. When the user logs out, the elevated privileges are terminated with the session but may be reestablished at the next login.
4. End role elevation by returning to the dialog box in step 2, and deselecting the role.

Force administrators to manually elevate

A property is available to force all users with the administrator role to manually select the role that they want to elevate to.

Before you begin

Role required: security_admin

Procedure

1. Log in as admin.
2. Elevate your role to security_admin.
3. Navigate to **sys_properties.list**.
4. Search for and select the *glide.security.strict_elevate_privilege* property.
5. Set the Value field to **true** and click **Submit**.

Result

When the user logs in, they are presented with a dialog window to select the role to which they can elevate.

Connections and Credentials

Credentials and connection information are required to gain access to a computer or network device for Discovery, Service Mapping, and Cloud Management or to perform work using Orchestration. When adding content to Share or AppStore, you can configure connections and credentials relevant to your environment without modifying built content.

Explore	Configure
 Learn about Credentials.	 Configure Credentials.
Reference	**Troubleshoot**
 Get details about Credentials.	 Learn how to troubleshoot issues in connections and credentials.

Explore credentials, connections, and aliases

All application integrations in the ServiceNow AI Platform use connections, credentials, and aliases to enable applications to access resources.

Before you can execute an application integration in the ServiceNow AI Platform, you must create and configure connection information, corresponding credentials, and add an alias. To understand how ServiceNow defines these terms:

Connection

A connection is an integration with a system, such as an IP address or endpoint with protocols. It contains specific details, such as database particulars, when integrating with a database.

Credential

A credential is the authentication data required to make the connection, such as an ID and password.

Alias

An alias is a naming convention, or tag, that ties to a set of connections or credentials on your instance. An alias contains the necessary connection and credential information to make an application integration. Rather than enter that information every time you integrate, you can use an alias. For example, you can designate an alias to house your QA, development, and production credentials for the same application integration. The alias resolves the application integration for each environment.

The ServiceNow AI Platform distinguishes different types of aliases:

Credential Alias

This alias associates to credential data only, and resolves during runtime.

Connection and Credential Alias

This alias associates to connection information and the credential data required to complete the integration, and resolves during runtime.

Within connection and credential aliases, you can also create additional aliases called child aliases. Child aliases allow you to create multiple connections within the same application integration. When you create a child alias, the alias you created it under becomes a parent alias. While child aliases inherit properties from their parent alias, child aliases carry their own connection and credential information.

Benefits to using Connections, Credentials, and Aliases

- Central location to store and manage credentials to an external service
- Define once and reuse for multiple platform features
- Minimize configuration of other platform features
- Allow non-administrators to use predefined connections and credentials
- Increased security

Features using Connections, Credentials, and Aliases

The following features use connections, credentials, and aliases:

- Flow Designer
- IntegrationHub
- Cloud Management
- Discovery
- Orchestration
- Service Mapping

You can set up aliases on the ServiceNow AI Platform in one of two ways:

- Using the Connections and Credentials module. See Create a Connection & Credential alias.
- In the Connections dashboard of Integration Hub. See Add a connection.

Note: Integration Hub requires a separate subscription. For more information, see Request Integration Hub.

Credential synchronization on MID Servers

Each MID Serverin your network synchronized with the instance keeps a copy of every credential that you create. The Management, Instrumentation, and Discovery (MID) Server is a Java application that enables communication and the movement of data between a ServiceNow instance and external applications, data sources, and services. This synchronization speeds up the reading of credentials when applications like Discovery or Service Mapping need to access multiple devices on the network. The MID Servers synchronize when they find a *credentials_reload* job in the ECC Queue. The reload job instructs the MID Server to make a SOAP call to the instance to get the entire list of credentials in the Credentials [discovery_credentials] table, including all the field values. To learn more, see MID Server.

The SOAP response that your instance sends to each MID Server also includes custom fields that you added to any credential form that you customized. If you added reference fields, the data in the referenced table is also sent as part of the SOAP response. This can lead to performance issues when credential synchronization occurs with multiple MID Servers. To control this, manually add these properties to the System Properties [sys_properties] table:

Note: To change the values in these properties, add them to the System Properties [sys_properties] table. If you do not add them, the system uses the default value.

Property	Description
`com.snc.credentials_user_fields`	Includes all customized fields in credential sync. Set this property to false if you do not want to include the fields that you added to credential forms. • **Type**: true | false • **Default value**: true
`com.snc.credentials_recursion_depth`	Defines the number of tables to traverse when the credential-sync mechanism collects fields from reference tables. Lower this number if you are experiencing performance issues and you have customized credential forms that include reference fields to tables that also have reference fields. • **Type**: integer • **Default value**: 3

Scope protections for Credentials and Connections

You can classify certain types of Connection & Credential records as belonging to a scope, and extend scope protections to them. These scope policies protect records you create in a table, and prevent interactions with records that are private to another scope.

An **Application** field is available in the Connection [sys_connection] and Discovery Credentials [discovery_credentials] tables for associating these types of records to specific scopes. It is not visible on UI forms in Zurich, but you can easily add it to them. To learn more about these record types, and adding the field to their UI forms, see:

- Get started with connections
- Get started with credentials
- Credential aliases for Discovery
- Configuring the form layout

Restricting the use of a Connection & Credentials record to a specific scope is important for managing applications that require enforced security. These applications include HR Service Delivery or Security Operations Connection & Credentials records created in scoped administered applications are not visible to admin users. Associating a Connection & Credentials record to a specific application scope affords the following protections:

- Applies Access Control List rules (ACLs) to restricted scopes. To learn more about scoped ACLs, see Access control list rules.

  Note: Some applications using scope administration and enforced security may require additional setup. To learn more, see Manage HR roles

- Protects records when queried using a script. If you do a query from the Global scope, and the Connection & Credential record is in a protected scope, it doesn't appear in the query, unless given access to it.

You can customize and grant access to query-restricted records by using Restricted Caller Access. To learn more, see Restricted caller access privilege settings. Scoping restrictions also apply to all children tables of the Connection [sys_connection] and Discovery Credentials [discovery_credentials] tables. Empty fields and other scopes are not restricted.

Note: Scope protections are only enabled for specific secured scopes to avoid confusion when setting up new records. If someone makes a connection in their scoped application scope, it doesn't have automatic scope restriction.

Domain separation and Credentials and Connections

Domain separation is supported in Credentials and Connections. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can control several aspects of this separation, including which users can see and access data.

Support level: Standard

- Includes all aspects of **Basic** level support.
- Application properties are domain-aware as needed.
- Business logic: The service provider (SP) creates or modifies processes per customer. The use cases reflect proper use of the application by multiple SP customers in a single instance.
- The instance owner must configure the minimum viable product (MVP) business logic and data parameters per tenant as expected for the specific application.

Sample use case: An Admin must be able to make comments required when a record closes for one tenant, but not for another.

For more information on support levels, see Application support for domain separation.

Overview

Credentials are tied to various ServiceNow features which access systems outside the instance. Credentials follow the domain separation tied to the feature employing the credentials.

Connections are protocol-specific information referencing a target host outside the instance. A connection can specify the domain to run an activity.

How domain separation works in Credentials and Connections

Credentials access resources outside of the instance, and are used by the Discovery, Orchestration, Service Mapping, and Cloud Provisioning and Governance applications. These credentials are not tied to a specific domain, rather, they can be bound to an application and then follow the domain separation that the application uses. Credentials can also be

assigned to a MID Server, and then follow the domain separation specified by the MID Server configuration.

Connections access a target host using a JMS, JDBC, or HTTP(s) connection. You can specify global or a specific domain to which the connection belongs.

Related topics

Domain separation for service providers

Connection & Credential configuration templates

Enable users with the admin and flow_designer roles to set up spoke integrations with third-party systems using a single, customizable form.

For example, you can set up an OAuth integration, which registers an OAuth provider, generates a token, and creates connection and credential records. An action designer or developer can use a configuration template to set up the spoke in one place and the system creates the associated records.

Benefits

Configuration templates enable:

- Admins or flow designers to set up a complex integrations using a single form.
- Developers to set static values in an integration, simplifying the setup process for admins and flow designers.

Supported credential types

You can create configuration templates for integrations with these credential types:

- Basic auth
- API key
- OAuth JWT Bearer grant type
- OAuth Authorization Code grant type
- Custom authentication

Configuration template components

Default Data Template

Sets static information that applies to every integration. For example, you can set the API and token URL if the value applies to every integration.

Dynamic Data Template

Defines the information that the user must complete to set up the integration. For example, you can add user name and password key-value pairs to gather user-defined values.

Post Processing Script

Creates additional records required by the integration. For example, if your spoke includes custom tables, you can create records in those tables based on user input in the configuration template. This script executes after the connection and credential records are created.

Pre Edit Script

Pre-populates the custom fields in the **Additional Information** section when you edit an existing connection. Pre-populating the custom fields enables you to view the current value associated to the custom field.

Test Action

Enables you to test a connection directly from an integration action in the flow view of Workflow Studio. The test action uses an action definition to test the alias that the template is currently attached to. For details, see Create a test action to test a connection alias from a configuration template.

Demo data

The Connection & Credential Templates [sys_alias_templates] table includes example templates to demonstrate how to set up templates for common authentication types. Use these examples as a guide when creating your own.

Configure a template for OAuth JWT Bearer grant type

This example configuration template sets up Credential and Connection records using the JWT Bearer grant type to authenticate requests to Docusign.

Default data template

Each top-level item in the default data template creates an associated record. The template includes these sections:

- `Credential`: Creates a record in the Credentials table.
- `Connection`: Creates a record in the Connections [sys_connection] table and any associated connection records.
- `Additional`: Optionally creates records in a custom table. The post processing script tells the system what to do with these records.

The following example creates the records required for OAuth JWT Bearer grant type authentication.

```
{
  "credential": {
    "oauth_entity": {
      "oauth_entity_profile": [
        {
          "grant_type":
 "urn:ietf:params:oauth:grant-type:jwt-bearer",
          "name": "Docusign Profile",
          "default": true,
          "oauth_entity_profile_scope": [
            "users:read.email"
          ]
        }
      ],
      "code_challenge_method": "S256",
      "type": "consumer",
      "oauth_entity_scope": [
        {
          "oauth_entity_scope": "users:read.email",
          "name": "email"
        }
```

```
      ],
      "client_id": "<provider-client-id>",
      "use_mutual_auth": false,
      "revoke_token_url":
"https://<provider-domain-name>.com/oauth2/revoke",
      "default_grant_type":
"urn:ietf:params:oauth:grant-type:jwt-bearer",
      "public_client": false,
      "oauth_api_script": "3e3a3a11c333210016194ffe5bba8f70",
      "name": "Docusign Spoke OAuth",
      "client_secret": "<provider-client-secret>",
      "auth_url":
"https://<provider-domain-name>.com/oauth2/auth",
      "token_url":
"https://<provider-domain-name>.com/oauth2/token",
      "redirect_url":
"https://<instance-name>.service-now.com/oauth_redirect.do"
    },
    "jwt_provider": {
      "jwt_keystore_aliases": {
        "kid": "<provider-key-id>",
        "name": "Docusign Spoke JWT Key",
        "signing_keystore": "<signing-keystore-sys-id>",
        "signing_algorithm": "rsa_256",
        "signing_key_password": "password"
      },
      "jwt_claim_validation" : [ {
        "name" : "iss",
        "is_standard" : true,
        "data_type" : "string",
        "value":"<docusign-iss-claim>"
      }, {
        "name" : "sub",
        "is_standard" : true,
        "data_type" : "string",
        "value":"<docusign-sub-claim>"
      }, {
        "name" : "aud",
        "is_standard" : true,
        "data_type" : "string",
        "value":"<docusign-aud-claim>"
      }, {
        "name" : "scope",
        "is_standard" : false,
        "data_type" : "string",
        "value" : "signature impersonation"
      } ],
      "name": "Docusign Spoke JWT Provider",
      "jwt_api_script": "9ef6af86ff10330001d3cd6bd53bf144"
    },
    "name": "Docusign Spoke Credential",
    "table": "oauth_2_0_credentials"
  },
  "connection": {
    "use_mid": false,
    "connection_url": "https://<provider-domain-name>.com",
    "name": "Docusign Spoke Connection",
```

```
    "table": "http_connection"
  },
"additional":{
    "docusign_account_name": "<docusign-account-name>",
    "docusign_account_email": "<docusign-account-email>"
}
}
```

Dynamic data schema

The dynamic data schema defines what the user sees when they create a Connection & Credential alias and collects their input. Use dot-walking syntax to map user input to fields created in the default data template. For example, `connection_fields` maps user input to the `connection_url` field in the `connection` object created by the default data template.

```
{
  "connection_fields": [
    {
      "name": "connection.connection_url",
      "label": "Connection URL",
      "type": "text",
      "defaultValue": "https://demo.docusign.net",
      "hint": "Connection URL for Docusign"
    }],
"additional_fields":[
    {
      "name": "additional.docusign_account_id",
      "label": "Docusign Account Number",
      "type": "text",
      "hint": "Docusign Account Number"
    },
    {
      "name": "additional.docusign_account_name",
      "label": "Docusign Account Name",
      "type": "text",
      "hint": "Name to identify the Docusign account"
    },
    {
      "name": "additional.docusign_account_email",
      "label": "Docusign Account Email",
      "type": "text",
      "hint": "Docusign Account Email"
    }
  ],
  "credential_fields": [
    {
      "name": "credential.oauth_entity.client_id",
      "label": "OAuth Client ID",
      "type": "text",
      "hint": "Client ID for Docusign"
    },
    {
      "name": "credential.oauth_entity.redirect_url",
      "label": "OAuth Redirect URL",
      "type": "text",
      "defaultValue":
 "https://<instance-name>.service-now.com/oauth_redirect.do",
```

```
      "hint": "Callback URL for Docusign"
    },
 {
      "name":
 "credential.jwt_provider.jwt_claim_validation[0].value",
      "label": "Issuer (iss) Claim value",
      "type": "text",
      "hint": "The integrator key (also known as client ID) of
 the application"
    },
 {
      "name":
 "credential.jwt_provider.jwt_claim_validation[1].value",
      "label": "Subject (sub) Claim value",
      "type": "text",
      "hint": "The user ID of the user to be impersonated"
    },
 {
      "name":
 "credential.jwt_provider.jwt_claim_validation[2].value",
      "label": "Audience (aud) Claim value",
      "type": "text",
      "defaultValue": "account-d.docusign.com",
      "hint": "The URI of the authentication service instance to
 be used e.g. account.docusign.com"
    },
    {
      "name":
 "credential.jwt_provider.jwt_keystore_aliases.kid",
      "label": "Key ID (kid)",
      "type": "text",
      "hint": "Indicates which key was used to secure the JWS"
    },
    {
      "name":
 "credential.jwt_provider.jwt_keystore_aliases.signing_keystor
e",
      "label": "Key Store",
      "type": "file"
    }
  ]
}
```

Post processing script

The following post processing script maps user input to fields in the sn_docusign_spoke_accounts table.

```
(function execute(aliasId, connectionSysId, jsonDefaultData,
 jsonDynamicData) {
 var jsonDynamicDataP = JSON.parse(jsonDynamicData);
 var accountGR = new GlideRecord("sn_docusign_spoke_accounts");
    accountGR.setValue("account_name",
 jsonDynamicDataP["additional.docusign_account_name"]);
    accountGR.setValue("alias", aliasId);
    accountGR.setValue("email",
 jsonDynamicDataP["additional.docusign_account_email"]);
```

```
    accountGR.setValue("id",
 jsonDynamicDataP["additional.docusign_account_id"]);
    accountGR.insert();
})(aliasId, connectionSysId, jsonDefaultData, jsonDynamicData);
```

Resulting Docusign Connection and Credential configuration form

When the user navigates to the associated Docusign Connection & Credential Alias and selects **Create New Connection & Credential**, the following dialog appears.

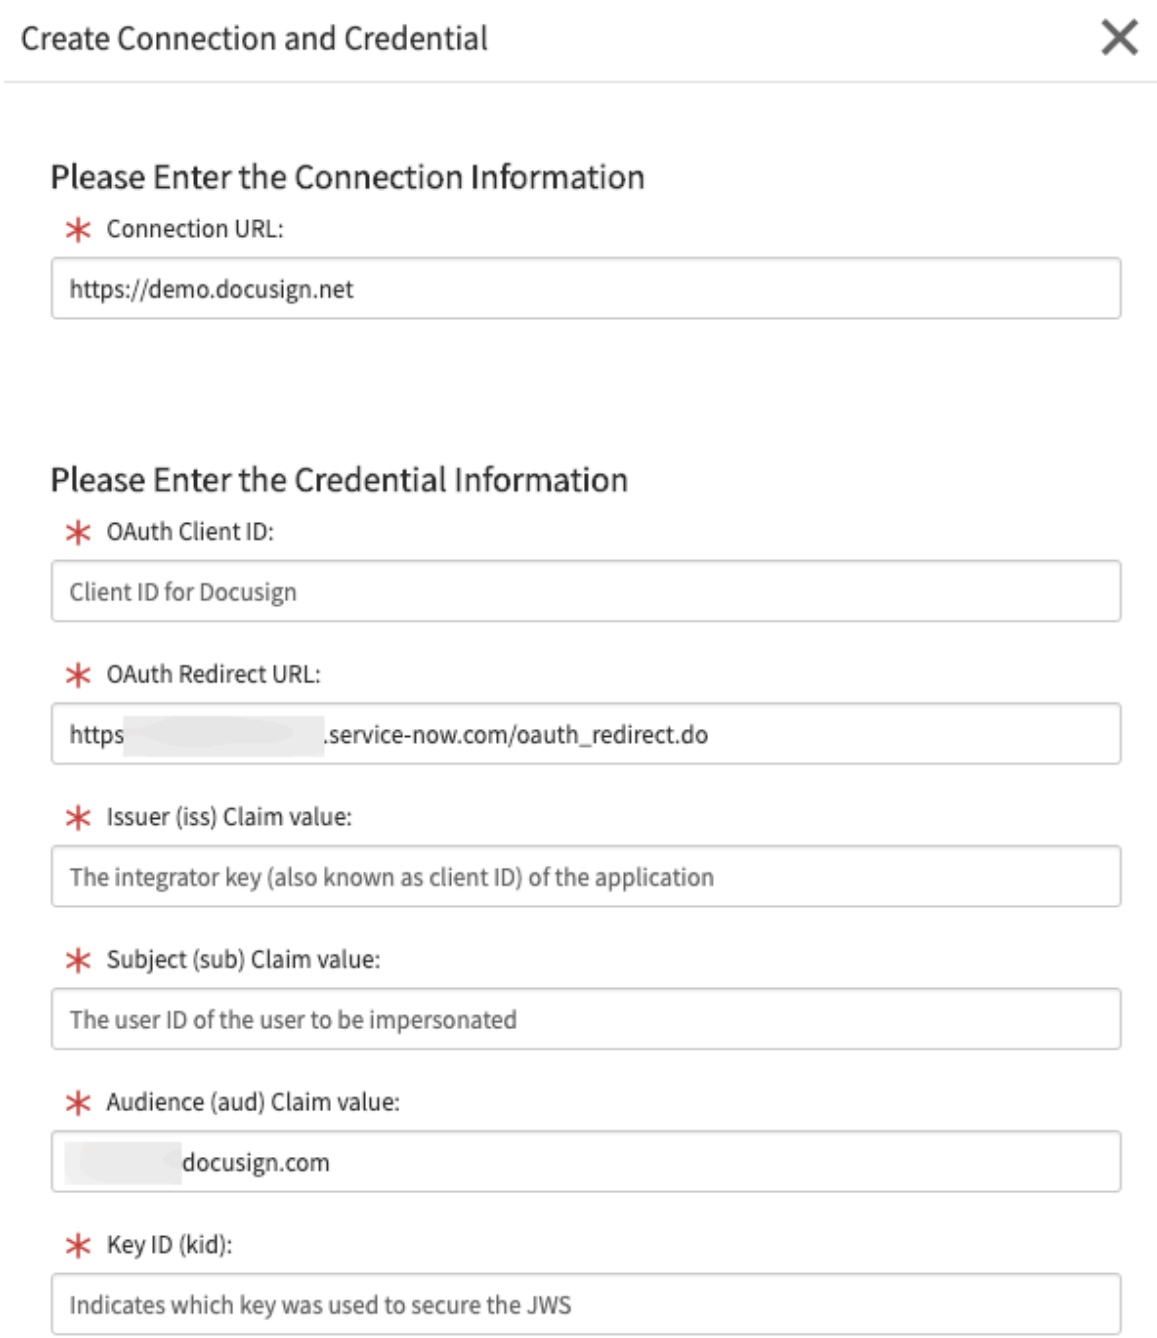

Create a configuration template

Create a template that defines the inputs required to set up a spoke. Set static key-value pairs to create records and set values that apply to every integration. Set dynamic key-value pairs

to gather user input and set field values that may vary. Using this template, admins and flow designers can set up the spoke from a single form.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **IntegrationHub** > **Connections & Credentials** > **Configuration Templates**.
2. Select **New**.
3. Select the type of configuration template that you want to create.

Configuration type	Description
HTTP Connection with OAuth Authorization Code grant type	Creates a template to register the third-party system as an OAuth provider using an authorization code.
HTTP Connection with OAuth JWT Bearer grant type	Creates a template to register the third-party system as an OAuth provider using a JSON Web Token.
HTTP OAuth with Client Credentials grant type	Creates a template to integrate with a third-party application via the OAuth 2.0 authentication that uses the client credentials grant type.
HTTP Connection with OAuth Client Credentials grant type (External Storage)	Creates a template that enables the sending of OAuth token request that comprises client credentials stored by an external storage to an OAuth server via the MID Server.
HTTP Connection with Basic Auth Credential	Creates a template to integrate with the third-party system using basic authentication.
HTTP Connection with API Key Credential	Creates a template to integrate with the third-party system using an API key.
Other Configuration	Creates a blank template, enabling you to set up a template for custom authentication.

4. In the **Name** field, add a name to identify the template.
5. In the **Default Data Template** field, make any required changes.

   This field sets static information that applies to every integration. For example, you can set the API and token URL if the value applies to every integration.

   These objects in the default data template are required:

   - `credential`: Creates a credential record with the required fields.
   - `connection`: Creates a connection record with the required fields. Access connection attributes using the `extended_attributes` child object. For example:

```
"connection": {
            "extended_attributes": {
                "api_version": "v1"
```

```
            },
        "connection_url":
"https://<provider-domain-name>.com",
        "name": "Spoke Connection",
        "table": "http_connection"
    }
```

You can use the `additional` object to set up data for a custom table, and use the post processing script to insert the data into the table.

Note: In the default data template for the OAuth Authorization Code grant type, the values for the `oauth_entity_profile_scope` and the `oauth_entity_scope` keys must match. In the following example, both keys have the `Read user's email` value.

```
"oauth_entity_profile": [
        {
          "grant_type": "authorization_code",
          "name": "<provider-name> Profile",
          "default": true,
          "oauth_entity_profile_scope": [
            "Read user's email"
          ]
        }
      ],
      "code_challenge_method": "S256",
      "type": "consumer",
      "oauth_entity_scope": [
        {
          "oauth_entity_scope": "Read user's email",
          "name": "email"
        }
      ],
```

6. In the **Dynamic Data Schema** field, make any required changes.

   This field Defines the information that the user must complete to set up the integration. For example, you can add user name and password key-value pairs to gather user-defined values.

   Fields in the dynamic data schema include these properties:

   - `name`: The field that the user's input maps to. For example, to map user input to the Connection URL field in the Connection record, enter `connection.connection_url`.
   - `label`: The field label that the user sees when completing the template.
   - `type`: The field type. Make sure that this data type matches the data type of the field you are mapping the value to.
   - `defaultValue`: Optional. The field's default value. If no default is provided, the hint displays.
   - `hint`: Optional. Hint text to display when there is no default value.

Note: If setting up a template for OAuth JWT Bearer grant type authentication, you may want user input for a single key-value pair in the `jwt_claim_validation` array. You can refer to a single key-value pair in the dynamic data schema by referring to its index in the array. For example, your default data template might include this snippet.

```
"jwt_claim_validation" : [ {
        "name" : "iss",
        "is_standard" : true,
        "data_type" : "string",
        "value":"<docusign-iss-claim>"
      }, {
        "name" : "sub",
        "is_standard" : true,
        "data_type" : "string",
        "value":"<docusign-sub-claim>"
      }, {
        "name" : "aud",
        "is_standard" : true,
        "data_type" : "string",
        "value":"<docusign-aud-claim>"
      }, {
        "name" : "scope",
        "is_standard" : false,
        "data_type" : "string",
        "value" : "signature impersonation"
      } ],
```

Refer to the `iss` key-value pair using the item's zero-based index: `credential.jwt_provider.jwt_claim_validation[0].value`.

7. **Optional:** In the **Post Processing Script** field, add a script that c Creates additional records required by the integration. For example, if your spoke includes custom tables, you can create records in those tables based on user input in the configuration template. This script executes after the connection and credential records are created.
   The post-processing script has access to these global objects.

Global object	Description
aliasId	Sys_id of the alias record from the Connection & Credential Alias [sys_alias] table.
connectionSysId	Sys_id of the connection record created by the template.
jsonDefaultData	JSON content from the **Default Data Template** field in String format.
jsonDynamicData	JSON content from the **Dynamic Data Template** field in String format.

8. In the **Pre-Edit Script** field, add a script to pre-populate the additional fields when you edit a connection.

   This script returns an array of objects. Each object has a name-value pair for populating the additional fields. For example, if the connection requires fields that are in a custom table, you can map these fields to the custom table.

   The **Pre-Edit Script** has access to the following global objects:

Global object	Description
aliasId	Sys_id of the alias record from the Connection & Credential Alias [sys_alias] table.
connectionSysId	Sys_id of the connection record created by the template.
jsonDefaultData	JSON content from the **Default Data Template** field in String format.
jsonDynamicData	JSON content from the **Dynamic Data Template** field in String format.

Each object in the script has the following properties:

- `name`: Name of the custom field to specify the value in the connection.
- `value`: Value that you want to map to populate the custom field. You can map the field either by using a function, variable, or by hard-coding it.

The following data types are supported for the fields:

Supported data types for the fields

Type	Description
Text	String value.
Boolean	Selection box. Selection indicates true value and unselection indicates false value.
Number	Numerical value.
Date	Date value in yyyy-mm-dd format. You can also use the GlideDate object.
Choice	List of valid choices defined in the **Dynamic Data Schema** field.
Reference	Valid GlideRecord.
Radio group	Groups that contains a different set of fields. These groups are available as a choice in a drop-down list when you edit a connection. The fields in each group appear when you select the required group from the drop-down list. For example, consider the following structure of radio group defined in the **Dynamic Data Schema** field: `{ "name": "radio_groups", "label": "Radio Groups", "type": "radio", "groups": [{ "name": "radio_group1", "label": "Radio Group 1", "fields": [{ "name": "radio_field1", "label": "Radio Field 1", "type": "text", "defaultValue": "efgh",`

Type	Description
	<pre>"mandatory": true }] }, { "name": "radio_group2", "label": "Radio Group 2", "fields": [{ "name": "radio_field2", "label": "Radio Field 2", "type": "text", "defaultValue": "abcd", "mandatory": true }], "default_group": true }</pre> For this example, you can use the following code snippet to understand how radio groups are used in the script by using dot-walking: <pre>{ name: "radio_field.first_radio_group.radio_field1", value: "radio field 1" }, { name: "radio_field.second_radio_group.radio_field2", value: "radio field 2" }, { name: “radio_groups”, value: gr.getValue(‘radio_groups’) }</pre> For more information on how to use dot-walking, see Dot-walking.

Tip: If the pre-populated values are not appearing in the fields while editing a connection, navigate to **System Diagnostics** > **Session Debug** > **Debug Log** to diagnose the issue.

9. **Optional:** In the **Test Action** field, enter the name of the test action.
   Creating a test action enables you to test a connection directly from an integration action in the flow view of Workflow Studio. The test action uses an action definition to test the alias that the template is currently attached to. For details, see Create a test action to test a connection alias from a configuration template.
10. Add the template to a Connection & Credential alias.
    a. Navigate to **IntegrationHub** > **Connections & Credentials** > **Connection & Credential Aliases**.
    b. Open the alias record for the spoke.
    c. In the **Configuration Template** field, click the Lookup icon.

d. Select the template that you created from the list.

e. Click **Update**.

Result

When the user navigates to the associated Connection & Credential alias and selects **Create New Connection & Credential**, a dialog appears to collect their input. If you created a template for the OAuth Authorization Code grant type, you can also retrieve an OAuth token from this dialog.

What to do next

Test the template by navigating to the associated Connection & Credential alias and selecting **Create New Connection & Credential**. Verify that the dialog collects the expected data and creates the required records in the system.

Get started with connections

Use the connections table to set up a Basic, JMS, JDBC, or HTTP(s) connection to a target host.

Connection Table

The Connection table (sys_connection) is the base table for all connection tables. You can set up connections for the following protocols:

- Basic connection for PowerShell and SSH
- JDBC
- JMS
- HTTP(s)

The connection table references the connection alias table, which couples the connection alias to connection information. Every connection records the following information:

Base connection properties

Field	Description
Name	Name of the connection. This field must be unique on the table.
Credential	Specify the credential to use with this connection. This is optional.
Connection alias	The connection alias resolves your connection and credentials at run time. Only one connection is active per Connection alias at any one time.
Active	Check to make the current connection active.
Domain	Domain to which the connection belongs.

Credential is unique across active connections, if not empty.

Upgrading connection information

- The JDBC connection [jdbc_connection] and JMS connection [orch_jms_ds] tables are existing Orchestration connection tables that now extend from the Connection [sys_connection] table. The tables originally extended from sys_metadata. The sys_metadata related data is removed.

- The tables move from the Orchestration run time plugin [com.snc.runbook_automation.runtime] to the Credentials & Connections plugin.
- The upgrade process obtains JDBC and JMS connection information and creates corresponding connection aliases and assigns the alias to its corresponding connection.
- JDBC field name changes:
  - JDBC server is renamed to host
  - Database port is renamed to port
  - Data of the JDBC server and database migrates to host and port during the upgrade

Create a basic connection for PowerShell and SSH

Configure connection information to use with a custom activity or action that uses the PowerShell or Secure Shell (SSH) protocol.

Before you begin

Role required: admin or connection_admin

Procedure

1. Navigate to **All** > **Credentials & Connections** > **Connections**.
2. Click **New**.
3. Select **Basic Connection for PowerShell & SSH**.
4. Complete the form.

Field	Description
Name	Unique name of the connection record.
Credential	Select the credential record used to authorize the connection.
Connection alias	Select the alias record to associate with this connection. Using an alias enables you to update the connection record without having to reconfigure any actions or activities that use the alias.
Host	Fully qualified domain name of the target host where the system runs the activity or action. For example, host.domain.com.
Active	Select to make this connection active.
Domain	Determines the domain the activity runs in. Flow Designer does not support domain separation and ignores this field.
Override default port	Target port used by the connection. If you leave this field blank, the system uses the default port value.
Use MID Server	Select to connect to the target host through a MID Server. If selected, define fields in the Advanced MID Server Configuration section. **Note:** PowerShell requires a MID Server.
MID Selection	Option to select a specific MID Server or MID Cluster. Choose any one of the following options.

Field	Description
	○ **Auto-Select MID Server**: Your ServiceNow instance selects the MID Server without manual input. ○ **Specific MID Server**: Your ServiceNow instance uses MID Server that you specify. ○ **Specific MID Cluster**: Your ServiceNow instance uses the MID Cluster that you specify. A MID Cluster is a group of MID Servers that enables your ServiceNow instance to handle multiple integrations, and improve integration speed. For more information, see Configure a MID Server cluster. This field is available when **Use MID Server** is checked. **Note:** Ensure that the Integration Hub connection record is referred, not an Orchestration connection record.
Capabilities	The capabilities the MID Server must support to be eligible for selection. The system runs the action or activity from a MID Server that supports the selected capabilities. Only displays if **Use MID server** is selected. Required capabilities determine which MID Server is selected at runtime. To learn more about how a MID Server is selected during runtime, see MID Server selection. This field is visible only when **Auto-Select MID Server** is selected from the MID Selection list.
MID Application	The application the MID Server must support to be eligible for selection. The system runs the action from a MID Server that supports the selected application. Only displays if **Use MID server** is selected. To learn more about how a MID Server is selected during runtime, see MID Server selection.
MID Server	Specific MID Server on which the step runs. This field is visible only when **Specific MID Server** is selected from the MID Selection list.
MID Cluster	The specific MID Cluster you want to use. This field is available when **Use MID Server** is checked, and **Specific MID Cluster** is selected from the MID Selection list.

5. Click **Submit**.

Create an HTTP(s) connection

The HTTP(s) connection provides the information custom HTTP(s) actions or activities use to connect.

Before you begin

Role required: connection_admin

Procedure

1. Navigate to **All** > **Credentials & Connections** > **Connections**, click **New**, and select **HTTP(s) Connection**.
2. Add the following connection information and click **Submit**:

Field	Description
Name	Unique name of this HTTP(s) connection.
Credential	Select the credential record used to authorize the connection.
Connection Alias	Select the alias record to associate with this connection. Using an alias enables you to update the connection record without having to reconfigure any actions or activities that use the alias.
URL builder	Either manually enter the connection URL or use system to build the URL based on the inputs. Default is unchecked. If checked, the connection URL is calculated from the following fields: ○ Mutual authentication — Check box if mutual authentication is used. ○ Protocol — If mutual authentication is not used, enter protocol. The default is HTTPs. ○ Protocol profile — If mutual authentication is used, enter protocol profile from sys_protocol_profile. ○ Host ○ Port ○ Base path — Path of the connection string. **Note:** If mutual authentication is checked, connection URL is built: Protocol + :// + host:port +URL. If mutual authentication is unchecked, connection URL is built: Protocol profile + :// + host:port +URL
Connection URL	If URL builder is unchecked, enter the connection URL into this field. **Note:** If mutual authentication is checked, connection URL is built: Protocol + :// + host:port +URL. If mutual authentication is unchecked, connection URL is built: Protocol profile + :// + host:port +URL
Active	Check the box to make this connection active.
Domain	Determine the domain the action or activity runs in.
Use MID Server	Check to use a MID Server for this action or activity. If selected, define fields in the Advanced MID Server Configuration section.
MID Selection	Option to select a specific MID Server or MID Cluster. Choose any one of the following options. ○ **Auto-Select MID Server**: Your ServiceNow instance selects the MID Server without manual input. ○ **Specific MID Server**: Your ServiceNow instance uses MID Server that you specify.

Field	Description
	○ **Specific MID Cluster**: Your ServiceNow instance uses the MID Cluster that you specify. A MID Cluster is a group of MID Servers that enables your ServiceNow instance to handle multiple integrations, and improve integration speed. For more information, see Configure a MID Server cluster. This field is available when **Use MID Server** is checked. **Note:** Ensure that the Integration Hub connection record is referred, not an Orchestration connection record.
Capabilities	The capabilities the MID Server must support to be eligible for selection. The system runs the action or activity from a MID Server that supports the selected capabilities. Only displays if **Use MID server** is selected. Required capabilities determine which MID Server is selected at runtime. To learn more about how a MID Server is selected during runtime, see MID Server selection. This field is visible only when **Auto-Select MID Server** is selected from the MID Selection list.
MID Application	The application the MID Server must support to be eligible for selection. The system runs the action from a MID Server that supports the selected application. Only displays if **Use MID server** is selected. To learn more about how a MID Server is selected during runtime, see MID Server selection.
MID Server	Specific MID Server on which the step runs. This field is visible only when **Specific MID Server** is selected from the MID Selection list.
MID Cluster	The specific MID Cluster you want to use. This field is available when **Use MID Server** is checked, and **Specific MID Cluster** is selected from the MID Selection list.
Connection Timeout	Number of milliseconds the system waits for a successful host connection. If a successful connection does not occur during this time, the connection request times out. Leave this field empty to use the system default connection timeout value. **Note:** Avoid setting the **Connection Timeout** value to zero, as this may cause a stale connection.
additional_http_headers (Attributes tab)	Include any additional headers required, such as sap-client=100. If adding multiple headers, separate them with a semi-colon.
odata_ping_url (Attributes tab)	The default OData ping URL is /sap/bc/ping. To edit the ping URL for OData heartbeat, select the lock icon, edit the URL, and select the lock icon again.

3. Click **Submit**.
   You are ready to create a custom HTTP(s) action or activity.

Create a JDBC connection

The JDBC Connection provides the information custom JDBC actions or activities use to connect to various target databases.

Before you begin

You must have an appropriate JAR file, whether it is supplied with the instance or a custom JAR file.

Note: The ServiceNow instance supplies `mysql-connector-java-5.1.21.jar`, `sql-server-jdbc-4.0.jar`, and `ojdbc6.jar` files as part of the current release, which supports MySQL, SQLServer, and Oracle databases. Other databases, such as Sybase or DB2 Universal, must use a custom JAR file that must be uploaded to the instance before setting the JDBC connection.

Role required: connection_admin

About this task

JDBC credentials are retrieved separately by the activity designer template and support external credential storage, such as CyberArk.

Procedure

1. Navigate to **All** > **Credentials & Connections** > **Connections**, click **New** and select **JDBC Connection**.
2. Complete the form using the fields in the table.
   The database selection in the **Format** field determines which fields are available.

JDBC connection fields

Field	Database Format	Description
Name	All	Unique name of this JDBC connection. For example, you might enter **JDBC MySQLProd**.
Credential	All	Add credentials for JDBC provider.
Connection alias	All	Select the alias record to associate with this connection. Using an alias enables you to update the connection record without having to reconfigure any actions or activities that use the alias.
Query timeout	All	Maximum elapsed time the JDBC query is allowed to run without a response.
Connection timeout	All	Number of seconds the system waits before closing a JDBC connection or connection request. For example, if the **Connection Timeout** value is 10 seconds, the system waits 10 seconds for a successful connection. If a successful connection does not occur during this time, the connection request times out. If a successful connection does occur, the connection remains open until there is a 10-second period of inactivity. Once the connection has been inactive for 10 seconds, the connection is closed.

Field	Database Format	Description
		Note: Avoid setting the **Connection Timeout** value to zero, as this may cause a stale connection.
Active	All	Check the box to make this an active connection.
Domain	All	Domain for this table. By default, the JDBC Connection [jdbc_connection] table runs in the **global** domain.
Format	All	Database type for this connection. The default choices are: ○ **MySQL** ○ **Oracle** ○ **SQLServer** ○ **None** You can add **Sybase** or **DB2 Universal** to the choice list by uploading the appropriate JDBC driver JAR file to the instance. Orchestration automatically recognizes these drivers when they are loaded into the system and adds them to this list.
Host	Oracle, MySQL, SQLServer	Host name or IP address of the database server.
Oracle sid	Oracle	The Oracle database site identifier. The default value is **orcl**.
Oracle port	Oracle	Port that the Oracle database is using. The default value is **1521**.
Database name	MySQL, SQLServer	Name of the database.
Port	My SQL, SQLServer	Port that the selected database is using.
Instance name	SQLServer	Instance name for the selected SQLServer
Connection URL	All	URL that the MID Server uses to connect to the specified database. The URL is created automatically when you save the form, and is read-only for the default databases. **Note:** If the format selected is not one of the default databases, you must create the connection URL manually so that the MID Server knows how to create the connection.
JDBC driver	None, DB2 Universal, Sybase	The JDBC driver to use for this connection when it is not a default database. **Note:** If you add a Sybase or DB2 Universal database, you must enter the driver name in this field and upload the driver JAR file to the instance.

Field	Database Format	Description
Use MID server	All	Check to use a MID server for this action or activity. If selected, define fields in the Advanced MID Server Configuration section.
MID Selection	All	Option to select a specific MID Server or MID Cluster. Choose any one of the following options. ○ **Auto-Select MID Server**: Your ServiceNow instance selects the MID Server without manual input. ○ **Specific MID Server**: Your ServiceNow instance uses MID Server that you specify. ○ **Specific MID Cluster**: Your ServiceNow instance uses the MID Cluster that you specify. A MID Cluster is a group of MID Servers that enables your ServiceNow instance to handle multiple integrations, and improve integration speed. For more information, see Configure a MID Server cluster. This field is available when **Use MID Server** is checked. **Note:** Ensure that the Integration Hub connection record is referred, not an Orchestration connection record.
Capabilities	All	The capabilities the MID Server must support to be eligible for selection. The system runs the action or activity from a MID Server that supports the selected capabilities. Only displays if **Use MID server** is selected. Required capabilities determine which MID Server is selected at runtime. To learn more about how a MID Server is selected during runtime, see MID Server selection. This field is visible only when **Auto-Select MID Server** is selected from the MID Selection list.
MID Application	All	The application the MID Server must support to be eligible for selection. The system runs the action from a MID Server that supports the selected application. Only displays if **Use MID server** is selected. To learn more about how a MID Server is selected during runtime, see MID Server selection.
MID Server	All	Specific MID Server on which the step runs. This field is visible only when **Specific MID Server** is selected from the MID Selection list.
MID Cluster	All	The specific MID Cluster you want to use. This field is available when **Use MID Server** is checked, and

Field	Database Format	Description
		Specific MID Cluster is selected from the MID Selection list.

3. Click **Submit**.

Related topics

JDBC credentials

Create a JMS connection

Configure your system to use Java Messaging Service (JMS) with a custom JMS activity or action.

Before you begin

Role required: connection_admin

About this task

The MID Server must have the correct JMS connection factories for your organization. Configure those values in the *mid.property.jms.command.allowed_factory_names* property, found in **MID Server** > **Properties**. The default values for this property can be changed to any value or comma-separated list of values that the third-party JMS provider advertises.

Procedure

1. Navigate to **Credentials & Connections** > **Connections**.
2. Click **New**, select **JMS Connection**, fill in the form, and then click **Submit**.

Option	Description
Name	Unique name of this connection factory.
Credential	Add credentials for JMS provider.
Connection Alias	Select the alias record to associate with this connection. Using an alias enables you to up date the connection record without having to reconfigure any actions or activities that use the alias.
Initial Context Factory	Name of the JNDI class that is used to create the InitialContext. **Note:** For example, to connect to Ac tiveMQ V5.10 (JMS Provider), the value is `org.apache.activemq.jndi.ActiveMQInitialContextFac tory`.
Provider URL	Location of the running JMS provider installa tion. **Note:** For example, to connect to Ac tiveMQ V5.1: `tcp://ipAddressOr HostName:61616`.

Option	Description
Active	Check the box to make this an active connec tion.
Domain	Determine the domain the action or activity runs in.
Use MID server	Check to use a MID Server for this action or activity. If selected, define fields in the Ad vanced MID Server Configuration section.
MID Selection	Option to select a specific MID Server or MID Cluster. Choose any one of the following op tions. ○ **Auto-Select MID Server**: Your ServiceNow instance selects the MID Server without manual input. ○ **Specific MID Server**: Your ServiceNow in stance uses MID Server that you specify. ○ **Specific MID Cluster**: Your ServiceNow in stance uses the MID Cluster that you speci fy. A MID Cluster is a group of MID Servers that enables your ServiceNow instance to handle multiple integrations, and improve integration speed. For more information, see This field is available when **Use MID Server** is checked.
MID Cluster	The specific MID Cluster you want to use. This field is available when **Use MID Server** is checked, and **Specific MID Cluster** is select ed from the MID Selection list.
Capabilities	The capabilities the MID Server must support to be eligible for selection. The system runs the action or activity from a MID Server that supports the selected capabilities. Only dis plays if **Use MID server** is selected. To learn more about how a MID Server is se lected during runtime, see MID Server selec
MID Application	The application the MID Server must support to be eligible for selection. The system runs the action from a MID Server that supports the selected application. Only displays if **Use MID server** is selected. To learn more about how a MID Server is se lected during runtime, see MID Server selec

3. Navigate to **Connections & Credentials** > **Credentials**.
4. Click **New**, select **JMS Credentials**, and then provide the user name and password the MID should use to communicate with the JMS provider.
   For more information, see JMS credentials.
5. Click **Submit**.
   You are ready to create a custom JMS action or activity.

Create connection attributes for IntegrationHub

Define connection-specific variables that you can use in Integration Hub integration steps.

Before you begin

Role required:

- The admin role is required to create connection attributes.
- The connection_admin or admin role is required to assign attribute values.
- The action_designer or admin role is required to use attributes in a custom action.

Connection attributes are only used by integration steps, which require a subscription to Integration Hub. For more information about activating Integration Hub, see Request an Integration Hub plugin.

About this task

When using an integration step, you must establish a connection with an external system. Use a Connection & Credential alias instead of defining the connection inline. An alias enables you to update the connection details once without having to reconfigure every action. Any action step that uses an alias inherits the attributes associated with it. Workflow Studio displays attributes as data pills that you can drag into your action step. For example, you can create a page size attribute that becomes a REST step query parameter.

For more information about building custom Workflow Studio actions, see Workflow Studio.

Procedure

1. Navigate to **All** > **Credentials & Connections** > **Connection & Credential Aliases**.
2. Create or select an alias record.
3. From the Connection Attributes related list, click **New**.
4. Define the attribute label and field type.
   For a list of field types, see field types.
5. Click the Advanced view related link to set advanced dictionary preferences for the attribute. For example, to create an attribute with a dynamically calculated value. See Dictionary entry form.
6. Click **Submit**.
7. Define the attribute values in the connection record.
   a. Navigate to **Credentials & Connections** > **Connections**.
   b. Create or select a connection record with the same connection type as the alias.
   c. From **Connection alias**, select the alias with connection attributes.
   d. Save the record.

The Attributes tab populates with the connection attributes defined in the alias record.

e. Set values for the attributes.
If the alias has **Support Multiple Active Connections** enabled, you can associate more than one connection record with an alias and define attribute values in each connection record. If there are multiple connection records with attribute values for the same alias, the connection used when the flow executes determines the attribute values. For example, suppose that you have one action that uses an alias with two active connections endpoints: production and test. The attribute resolves to the value defined by the connection used at runtime.

8. Add the alias to an integration step in Workflow Studio.
   a. Navigate to Workflow Studio and create or select an action.
   b. Add an integration step to the action.
   c. Under Connection Details, add the alias you created attributes for.

   The connection attributes associated with the alias display as data pills in the Data pane.

   Note: The system does not track changes to connection attribute labels and data types after you associate the alias to a step. To refresh the connection attribute label or data type, delete the alias from the step and add it again.

Get started with credentials

The MID Server uses the credentials you create in the Credentials [discovery_credentials] table to access resources for Discovery, Orchestration, Service Mapping, and Cloud Management.

How MID Servers use credentials

By default, Windows MID Servers use the login credentials of the MID Server service on the host machine to discover Windows devices in the network. You should Configure Windows MID Server service credentials so that they have at least local administrator privileges. For Linux and UNIX machines and network devices, the MID Server uses the SSH and SNMP credentials configured in the instance in **Discovery** > **Credentials**.

MID Servers that Orchestration uses must have access to the necessary credentials to execute commands on computers in the network, as specified by the Workflow activities. Orchestration can use the same SSH and SNMP credentials as Discovery, but has two additional credentials designed for specific Workflow activities: Windows (for PowerShell activities) and VMware.

Encryption and decryption

The platform stores credentials in an encrypted field on the Credentials [discovery_credentials] table. Once they are entered, they cannot be viewed.

When the MID Server requests credentials, the ServiceNow AI Platform decrypts the credentials using the following process:

1. The credentials are decrypted on the instance with the password2 fixed key.
2. The credentials are re-encrypted on the instance with the MID Server's public key.
3. The credentials are encrypted on the load balancer with SSL.
4. The credentials are decrypted on the MID Server with SSL.
5. The credentials are decrypted on the MID Server with the MID Server's private key.

Note: The platform does not have separate encryption keys for multi-tenant instances.

Credential order

Credentials can be assigned an order value in the Credentials Form, which forces the application to try all the credentials at their disposal in a certain sequence. If you do not specify an order value, the application tries the credentials in the Credentials [discovery_credential] table randomly, until it finds one that works. For example, when:

- Orchestration attempts to run a command on an SSH server, such as a Linux or a UNIX machine.
- Discovery attempts to query an SNMP device, such as a printer, router, or UPS.

After identifying the credentials for a device, Discovery and Orchestration create an affinity between the credentials and the device using the Credential Affinity `[dscy_credentials_affinity]` table. All subsequent discoveries or Orchestration activities attempt to match the credentials in this table with a device for which an affinity exists. If credentials for a device change, Discovery and Orchestration try all available credentials again until they create a new affinity.

Note: If Orchestration and Discovery are installed, and credential alias is enabled, multiple affinities can exist. In this case, the platform looks up credentials for each affinity and inserts the credential for the affinity with the lowest order into the probe.

Ordering credentials is useful in the following situations:

- The credentials table contains many credentials, with some used more frequently than others. For example, the table contains 150 SSH credentials, and five of those credentials are used to log in to 90% of the devices. It is good practice to configure those five credentials with low-order numbers, which place them at the top of the execution list. Discovery and Orchestration work faster when they try these common credentials first. After the first successful connection, the ServiceNow AI Platform knows which credentials to use the next time for each device.
- The ServiceNow AI Platform has aggressive login security. For example, configure database credentials with a low-order value if Solaris database servers in the network only provide three failed login attempts before locking out the MID Server.

Credential aliases

Credential aliases are available for Discovery and Orchestration.

Aliases for Discovery enable an administrator to:

- Employ a credential filtering behavior with configurable levels of compliance.
- Assign multiple credential aliases to a Discovery schedule.
- Prevent the creation of credential affinities that use inappropriate or sensitive credentials. To learn more, see credential affinities.

Aliases for Orchestration enable workflow creators to:

- Assign individual credentials to any activity in an Orchestration workflow
- Assign individual credentials to any action in Flow Designer
- Assign different credentials to each occurrence of the same activity type in an Orchestration workflow.
- Assign different credentials to each occurrence of the same action in designer flow.

External credential stores

If you do not want credentials stored in your instance, you can use external credential repositories. External credential stores save the credentials in an external site that your instance can access. CyberArk is the only supported external credential store. However, other external stores can be configured using the ServiceNow API.

Create a Connection & Credential alias

Define an alias to label a credential or connection record.

Before you begin

Role required:

- The admin role is required to create an alias.
- The credential_admin and connection_admin roles have read access to the alias record.

About this task

The Connection & Credential alias defines an alias that labels a credential or connection record. The alias contains these fields.

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
2. Select **New**.
3. Complete the fields on the form.

Connection & credential aliases

Field	Description
Name	Name of the alias. An alias can only contain alpha, number, and underscore characters. During an upgrade, the tag in the credential record migrates to a Connection & Credential Alias. If the credential tag contains special characters other than alphabets, numbers, and underscores, it preserves the tag name after the upgrade. You can still use this migrated alias, but you can't update the alias until you change the name to meet the naming restrictions.
ID	Unique identifier for the Connections & Credentials alias, based on the format `scope_name.alias_name`. ○ If the scope is Global, the ID is the alias name. For example, if you create a Workday alias in the global scope, it sets the ID to `workday`. ○ If you create a workday alias in the HR app scope, it sets the ID to `x_hr_app.workday`.
Type	Select either **Credential** or **Connection and Credential**. The default is Connection and Credential.
Application	Application scope against which the Connection & Credentials alias is assigned. The current session scope you last selected in the application picker appears.

Field	Description
	○ For example, **Global** appears if it is the current scope for this session. ○ You can change the scope in the application picker before creating an alias. To learn more application scopes and how to select them, see: ▪ Application scope ▪ Select an application from the application picker
Connection type	Name of the connection type, either Basic, HTTP, JDBC, JMS, or Kafka. The default is HTTP.
Support Multiple Active Connections	Designator that indicates whether the alias supports multiple active connections. Add connections using the Connections table and associated them to the alias using the Connections related list.
Default Retry Policy	Retry policy for the alias. For more information, see Retry policy.
Configuration Template	Configuration template to use to create a connection and credential record.

4. Right-click the header and select **Save**.
   The Connections and Connection Attributes related list appears.

Related List	Description
Connections	Related connection records associated with this alias. After creating the alias, you can define connection records and associate them with the alias. If **Support Multiple Active Connections** is selected, you can associate more than one connection with an alias.
Connection Attributes	Attributes for the connection. Define data specific to a connection and use it in an Integration Hub integration step. For more information, see Create connection attributes for IntegrationHub.
Child Aliases	Child aliases associated with the parent alias. After creating a connection and credential alias, you can create child alias to configure multiple connections for the same application integration.

5. **Optional:** If you want to create a credential and connection associated with your credential alias, under *Related Links*, select **Create New Connection & Credential**.
   The resulting connection and credential records are based on a pre-defined configuration template. See connection and credential configuration templates.
6. **Optional:** If you want to create a child alias for your connection and credential alias, under the **Child Aliases** related list, select **New**.

   a. Enter a name for the child alias and select **Submit**.
      The child alias inherits properties from the parent alias. You can then configure a child alias to contain its own set of connection and credential information.

What to do next

Create one or more connection records to associate with the alias or child aliases. For more information about creating connections, see Get started with connections. Add connection attributes to the alias to make connection meta data available to flows in Workflow Studio.

Set up OAuth integration via MID Server

Create a connection record that enables the sending of an OAuth token request to a third-party server via a MID Server.

Before you begin

Confirm that you have subscribed to the ServiceNow IntegrationHub Standard Pack Installer. For more information, see https://www.servicenow.com/content/dam/servicenow-assets/public/en-us/doc-type/legal/snc-addendum-integrationhub.pdf.

Role required: Admin

Procedure

Configure the connection with a configuration template.

a. Navigate to **All** > **IntegrationHub** > **Configuration Templates**.

b. Select **New**.

c. Select **HTTP Connection with OAuth Client Credentials grant type**.

d. Update the form, as required.
For example, you can provide the scope in the `oauth_entity_profile_scope` and `oauth_entity_scope` fields in the Default Data Template section. To learn to create a configuration template, see Create a configuration template.

e. Select **Submit**.
You have created the configuration template.

f. Navigate to **All** > **IntegrationHub** > **Connection & Credential Aliases**.

g. Update the form.

Connection and credential alias

Field	Description
Name	Name of the alias. An alias can only contain alpha, number, and underscore characters. During an upgrade, the tag in the credential record migrates to a Connection & Credential Alias. If the credential tag contains special characters other than alphabets, numbers, and underscores, it preserves the tag name after the upgrade. You can still use this migrated alias, but you cannot update the alias until you change the name to meet the naming restrictions.
Application	Unique identifier for the Connections & Credentials alias, based on the format `scope_name.alias_name`.

Field	Description
	○ If the scope is Global, the ID is the alias name. For example, if you create a Workday alias in the global scope, it sets the ID to `workday`. ○ If you create a workday alias in the HR app scope, it sets the ID to `x_hr_app.workday`.
Parent Alias	Option to select the alias under which you're creating this connection and credential alias. The connection and credential alias that you're creating is a child alias. A child alias is listed under the **Child Aliases** tab on the parent connection and credential alias page. 
Type	Option to indicate the type of alias you're creating. Choose from the following options. ○ Credential: Alias that contains a credential record. ○ Connection and Credential: Alias that contains both connection and credential record. This option is selected by default. Confirm that Connection and Credential is selected.
Support Multiple Active Connections	Designator that indicates whether the alias supports multiple active connections. Add connections using the Connections table and associated them to the alias using the Connections related list.
Default Retry Policy	Retry policy for the alias. For more information, see Retry policy.
Configuration Template	Option to select the configuration template based on which you're creating the connection and credential alias. Select the template of the type **HTTP Connection with OAuth Client Credentials grant type** that you had created.

h. Select **Submit**.
You have created the connection and credential alias record.

i. Navigate to **All** > **IntegrationHub** > **Connections Dashboard**.

j. In the Search all connections field, enter the name of the connection and credential alias record that you created.

k. On the connection and credential alias record, select **View Details**.

l. Select **Configure**.

m. Fill the form.

Connection record

Field	Description
Connection Name	Name of the connection record. You can't update the name.
Connection URL	Option to provide the URL that connects with the third-party server.
Use MID	Option to specify that you want to send OAuth token requests via a MID Server. **Important:** Confirm that the option is selected.
MID Selection	Option to specify whether you want to use a specific MID Server or enable auto-selection of MID Server, or use a MID Cluster. Choose one of the following options. ○ Auto-Select MID Server ○ Specific MID Server ○ Specific MID Cluster
Capabilities	Option to select one or more MID Server capabilities. Capabilities define the specific functions of a MID Server within an IP address range, allowing an application to select the most appropriate MID Server. Select the MID capabilities icon ( ) to select one or more capabilities. **Note:** This option appears if you select **Auto-Select MID Server** in the MID Selection field.
MID Application	Option to specify a MID application or accept the default application choice. **Note:** This option appears if you select **Auto-Select MID Server** in the MID Selection field. By default, the **ALL** option is selected.
MID Server	Option to select a MID Server.

Field	Description
	Note: This option appears if you select **Specific MID Server** in the MID Selection field.
MID Cluster	Option to select a MID Cluster. **Note:** This option appears if you select **Specific MID Cluster** in the MID Selection field.
OAuth Client ID	Option to specify the client ID.
OAuth Client Secret	Option to specify the client secret.
Connect to Auth Server via MID Server	Option to specify that the connection between the ServiceNow instance and the Auth server takes place via the MID Server. **Important:** Confirm that the option is selected.
OAuth Token URL	Option to specify the OAuth token URL that is used to request OAuth tokens.

n. Select **Configure and Get OAuth Token**.
The connection and credential record is created.

Credential aliases for Discovery

Credential aliases for Discovery allow an administrator to use specific credentials on Discovery schedules. You can configure behaviors for your aliases that determine how strictly the system enforces their use.

Without credential aliases, Discovery schedules can access all credentials that are defined in the instance. This behavior might not be desirable in some circumstances, particularly for credentials with elevated privileges. Credential aliases provide more control over which credentials a Discovery schedule is allowed to use and prevents the unnecessary exposure of credentials with elevated privileges.

How credential aliases work

A business rule called *Insert Discovery Affinity & Cred Aliases* (previously named *Insert Discovery Affinity*) runs when a record (a task for performing Discovery) is inserted into the ECC Queue.

The business rule attaches the credential aliases defined in the Discovery schedule to the probe, so when the probe reaches the MID Server on its way to performing discovery, the MID Server knows exactly which credentials it can use to attempt to access the device the probe was sent to scan.

The MID Server filters credentials by affinity and then by tags, if any exist. Credentials must match all credential tags. The MID Server iterates until it finds a credential that works.

If the business rule determines that an affinity exists for the device, the rule identifies the proper *credential_id* to use. This is the *sys_id* of the record in the Credentials `[discovery_credentials]` table.

When the platform encounters an affinity with a credential alias value, defined as *credential_alias* in the business rule, the business rule determines whether or not the credential referenced by the affinity has the specified alias. If it does, the business rule selects the *credential_id* of the alias and passes that value to the MID Server.

If a credential alias is defined for a schedule and the schedule is configured to use that alias, the schedule will ignore any previously existing credential-to-target affinity—but only if the credential itself is not associated with any other credential alias. If the credential does not have any credential alias, any other affinities that exist for the target system are checked.

Create a Discovery credential alias

Create the alias and then add that alias to a credential in the credential record. You can add a credential to multiple aliases and add multiple credentials to a single alias.

Before you begin

Role required: admin, credential_admin (read access only), connection_admin (read access only)

About this task

A Discovery schedule only uses credentials that are contained in the aliases defined for that schedule.

Note: If a credential alias is defined for a schedule, it will ignore any previously existing credential affinity between the credential and the target that is discovered in a schedule that is setup to use that credential alias.

Procedure

1. Create an alias.
   a. Navigate to **Connections & Credentials** > **Connection & Credential Aliases**.
   b. Click **New**.
   c. Enter a unique name for the alias and select **Credential** for the alias **Type**.
   d. Click **Submit**.

The **Credentials** related list appears. You can add new credentials for this alias in this list but not existing credentials.

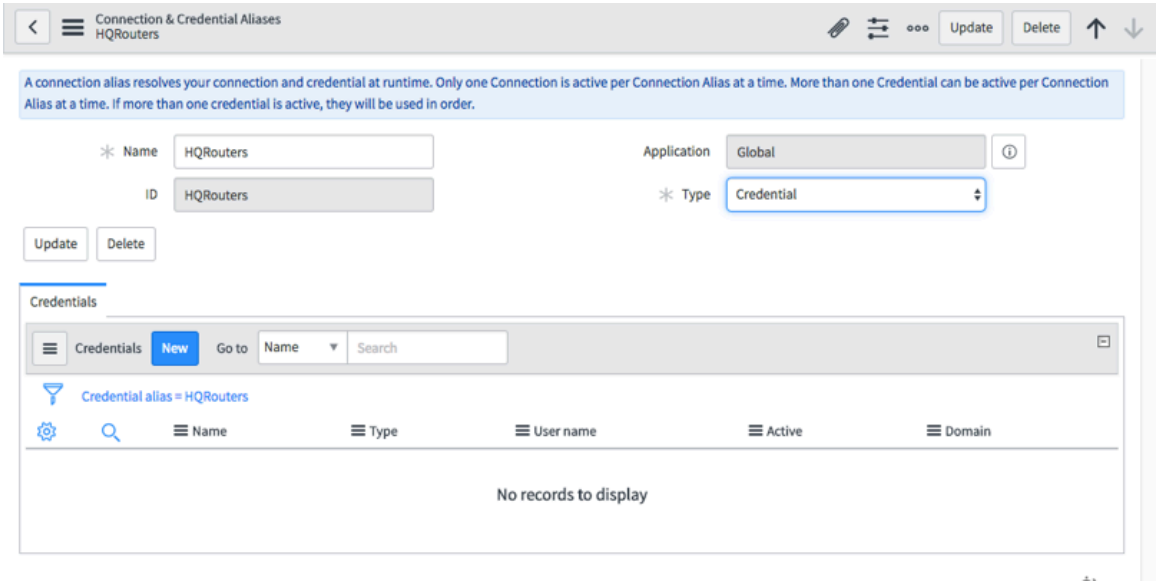

2. Configure a credential for the new alias.

   a. Navigate to **Connections & Credentials** > **Credentials**.

   b. Select an existing credential from the list or click **New** to create a new credential.

   c. In the credential record, unlock the **Credential alias** field and select the alias you created.

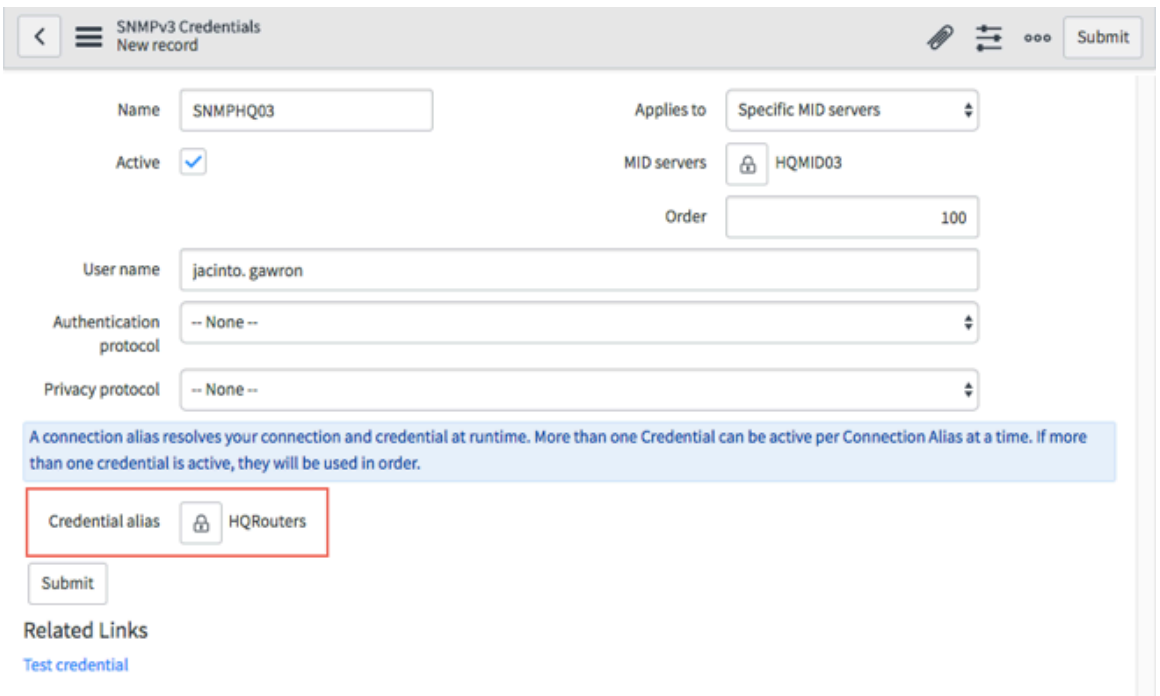

   d. Save or submit the record.

3. Return to **Connections & Credentials** > **Connection & Credential Aliases** and open your new alias.

The credential you attached to the alias now appears in the related list.

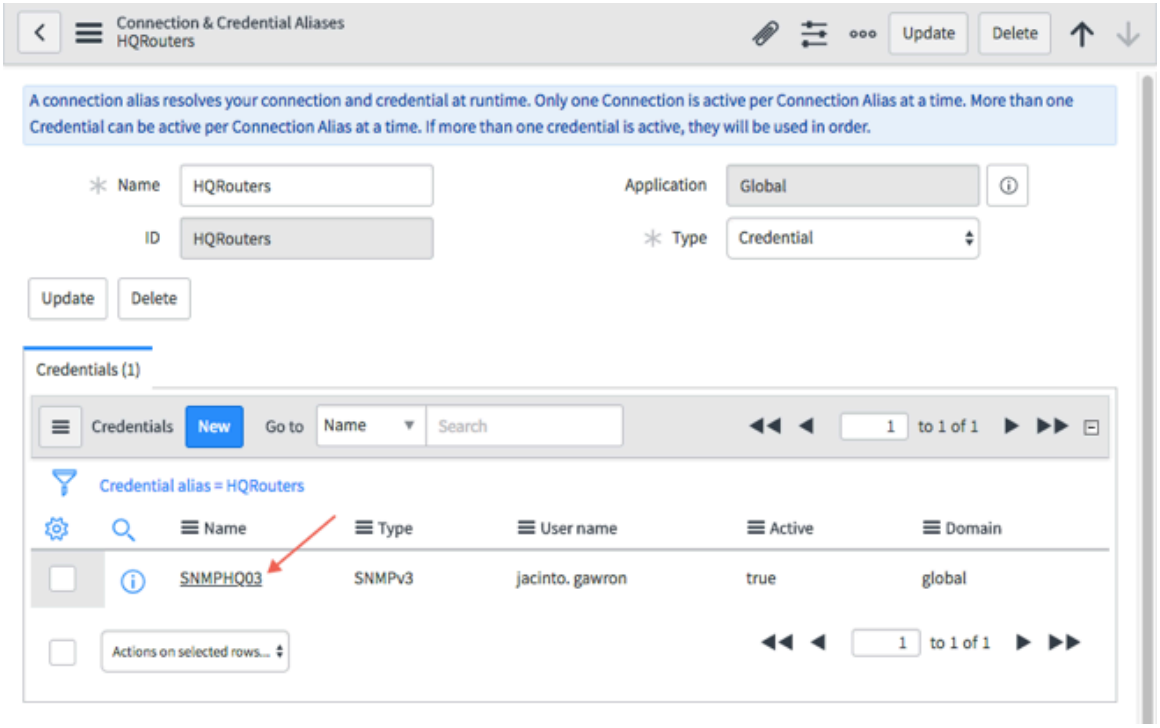

4. To create an additional credential for this alias, click **New** in the related list and select a credential type.
   The alias name is pre-populated in the **Credential alias** field of the credential record.
5. Complete the fields in the form and submit the record.

Credential aliases for Orchestration activities

Credential alias gives an administrator more control over the credentials used in Orchestration activities.

This is useful when an activity requires specific credentials to perform a task. You can use a credential tag to assign individual credentials to any activity in a Orchestration workflow or assign different credentials to each occurrence of the same activity type in an Orchestration workflow.

Credential alias interacts with credential affinity to determine which credentials should be used for an Orchestration activity.

How credential alias works

A business rule called *Insert Discovery Affinity* (renamed from Insert Credential Affinity in the Geneva release) runs when a record is inserted into the ECC Queue. This rule determines whether a credential affinity exists for the device and identifies the proper *credential_id* (the *sys_id* of the record in the Credentials `[discovery_credentials]` table) to use. When the platform encounters an affinity with a credential alias value defined (*credential_alias* in the business rule), the business rule determines if the credential referenced by the affinity has the specified alias. If it does, the business rule selects the *credential_id* of the credential alias and passes that value to the MID Server. If the credential does not have the specified credential alias, any other affinities that exist for the target system will be checked. If no affinity references an appropriately tagged credential, the MID Server iterates through the Credentials *[discovery_credentials]* table and selects the credential with the appropriate tag. The MID Server then creates a new affinity for this credential.

Create and test your credentials

Create and test the credentials that Discovery, Service Mapping, Cloud Management, and Orchestration require to access hardware and software in your network.

Before you begin

Role required: admin
Review your security policy and options with your organization's security team.

About this task

This task contains general procedures for creating credentials. Refer to the documentation for your credential type for details on specific fields and requirements.

Supported credential types

Applicative credentials	Basic authentication credentials	Chef server credentials
CIM credentials	Cloud credentials	Container image repository credentials
Infoblox credentials	JDBC credentials	JMS credentials
OAuth 2.0 credentials	SAP credentials	SNMP credentials
SSH credentials	VMware credentials	Windows credentials

Note: To improve security, limit the scope of credentials to a specific MID Server or schedule to avoid unnecessary credentials.

Procedure

1. Navigate to one of these modules:
   - **Discovery** > **Credentials**
   - **Service Mapping** > **Credentials**
   - **Orchestration** > **Credentials**
2. Click **New**.
3. On the Credentials page, click a link for the credential type and complete the form.

   Refer to the documentation for the credential type you selected for details.

   You can submit a credential record first and then test it later, or test the credential immediately before saving it.

   Credential testing is supported for these credential types:
   - SSH private keys
   - Windows
   - SNMP v3
   - VMware
   - JDBC
   - JMS
4. Under **Related Links**, click **Test credential**.

Note: Credentials are encrypted at all times during the test.

5. Complete the fields in the Test Credential dialog box.

Test credentials dialog box

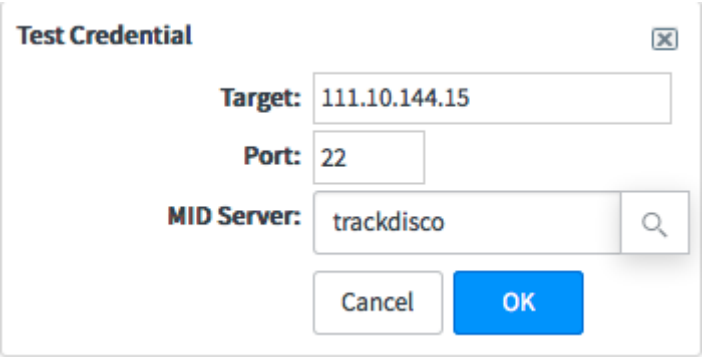

Credential test fields

Field	Description	Credential type
Target	Target host on which these credentials are run. This value must be an IP address for all credential types except VMware, which can be the host URL. You can not target any MID servers. **Note:** For JMS, this is the provider URL. The information in this URL tells JNDI how to find and access the JMS Provider. An example value for connecting to ActiveMQ V5.1, is `tcp://ipAddressOrHostName:61616`.	All
Port	Port on the target to use for this test. The system pre-populates this field with the default port for the selected credential type.	All
MID Server	MID Server to use for this test. You must use a Windows MID Server to test Windows credentials. Only **Up** and **Validated** MID Servers are available.	All
DB Type	Type of database on which to test these credentials.	JDBC

Field	Description	Credential type
DB Name	Name of the database on which to test these credentials.	JDBC
Initial Context Factory	Name of the JNDI class that is used to create the InitialContext. Using this **Initial Context Factory**, various JMS objects, such as JMS Connection, are created. For example, to connect to ActiveMQ V5.10, (JMS Provider), the value in this field would be `org.apache.activemq.jndi.ActiveMQInitialContextFactor`	JMS

6. Click **OK** to begin the test.
   An indicator appears, showing that the system is attempting to contact the target using the credentials you have provided. When the instance connects to the target it displays a success message. If the instance encounters a problem with the test inputs you have provided, it displays the appropriate error message. The following are some common error messages.
   - Incorrect target or port number:

     TCP connection failure

     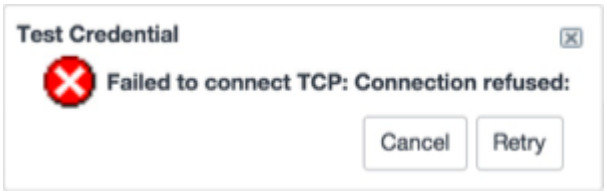

   - Incorrect user name or password:

     Authentication failure

     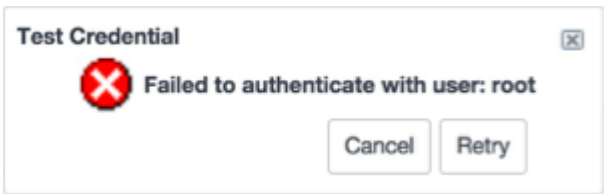

   - Incorrect MID Server for Windows credentials:

     MID Server error

     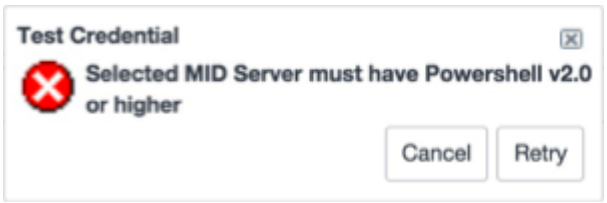

7. Click **Retry** to open the test Credential dialog box and correct the input error.

8. When your credentials test is successful, click **Submit** to save the record.

> **Important:** Testing credentials does not ensure that the credentials have the necessary privileges required for the intended Discovery or Orchestration workflow tasks.

Ansible Tower credentials

Ansible Tower credentials are required to access your Ansible configuration management account. Use these credentials to manage Ansible resources through the Cloud Management application.

To integrate Cloud Provisioning and Governance with the Ansible configuration management account, you must configure the user name and password for the administrator account in Ansible.

Form fields for Ansible Tower credentials

Field	Description
Name	Provide a descriptive name.
User name/ Password	Enter the authentication credentials for the Ansible Tower user with administrator rights.

> **Note:** You do not need to configure the other fields.

API key credentials

An API key is a unique code that is passed in to an API to identify the calling application or user.

API key credentials

API key credentials form fields

Field	Input value
Name	Enter a unique and descriptive name for this credential.
Active	Enable or disable these credentials for use.
API Key	Enter the API key.
Credential alias	Allow flow and workflow creators to assign individual credentials to any activity in a flow or workflow or assign different credentials to each occurrence of the same activity type in a flow or workflow.
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.

Applicative credentials

Some applications require credentials in addition to the credentials the that host machine requires. Credentials required to access these applications are referred to as applicative credentials.

A typical credential contains a user name and a password for logging in to a device or application. While most applications require only one credential for accessing them, sometimes hosts and applications have separate credentials for extra security. For example, ABAP SAP Central Services (ASCS) requires applicative credentials in addition to the SSH or Windows host credentials for the server that hosts ASCS.

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

As with host credentials, you assign applicative credentials to MID Servers.

You create applicative credentials per CI type, for example, the CI type for ASCS is SAP ASCS Application [cmdb_ci_appl_sap_ascs]. The preconfigured pattern for discovering CIs belonging to this CI type contains commands that require a MID Server to use the applicative credential for this CI type. If there's more than one credential configured for this CI type, the MID Server tries using these credentials in the order you define until it finds the credential that fits.

Check the Discovery requirements information in the ServiceNow documentation to determine if you need to configure applicative credentials for specific application CIs. There's no need to configure applicative credentials, if Discovery prerequisites don't mention it.

Applicative credentials form fields

Field	Description
Name	Name of the credential. Use a descriptive name like Oracle DB or London Oracle DB (for an Oracle database). Don't use spaces or special characters for the credential name.
Active	Select the check box to use the credential.
User name	Enter the actual user name of the applicative credential.
Password	Enter the actual password of the applicative credential. Don't use spaces or special characters for the credential name.
CI type	Select a CI type to which the CI belongs.
Credential Alias	Create an alias to assign specific credentials for specific discovery schedules. When assigning an alias, you must identify the table name for the CI type whose applicative credentials the application uses. Applications may use applicative credentials of a CI type different from their own. For a specific application, see the list for the appropriate table: • ABAP SAP Central Services (ASCS): cmdb_ci_appl_sap_ascs • IBM Security Access Manager appliance: cmdb_ci_app_server_webseal • SAP Central Instance: cmdb_ci_appl_sap_ascs • SAP Central Services (SCS): cmdb_ci_appl_sap_ascs • SAP Evaluated Receipt Settlement (ERS): cmdb_ci_appl_sap_ascs • SAP Java Cluster: cmdb_ci_appl_sap_ascs • SAP NetWeaver Dialog Instance: cmdb_ci_appl_sap_ascs

Applicative credentials form fields (continued)

Field	Description
	• Microsoft Exchange Mailbox (for Microsoft Exchange): cmdb_ci_exchange_mailbox • Microsoft SQL Database: cmdb_ci_db_mssql_instance • MySQL Server: cmdb_ci_db_mysql_instance • Oracle Advanced Queue Queue: cmdb_ci_db_ora_instance • Oracle Database: cmdb_ci_db_ora_instance • Oracle E-Business Suite: cmdb_ci_db_ora_instance • Oracle WebLogic Module: cmdb_ci_app_server_weblogic • Tibco Enterprise Message Service (EMS): cmdb_ci_appl_tibco_message
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.

Basic authentication credentials

The basic authentication credential type manages access to store basic authentication credentials.

These fields are available in the Credentials form for basic authentication.

Basic Auth credentials form

Field	Input value
Name	Enter a unique and descriptive name for this credential. For example, you might call it **Basic Authentication**.
User Name	Enter the user name.
Password	Enter the password.
Credential ID	Enter the unique key configured for these credentials in the CyberArk external credential storage system. The credential ID may be used as a safe override when multiple safes are in use. By default, the syntax in the **Credential ID** field is this: *<safe name>:<Credential ID>*. If the safe name is omitted, there must be a safe name defined in the `config.xml` file. To change the separator character from the default colon to another character, override the value with the optional *ext.cred.safe_name* parameter. The **Credential ID** field has a limit of 40 characters.

Basic Auth credentials form (continued)

Field	Input value
	This field is only visible when the **External storage** check box is selected.
External credential store	Select this check box to use an external credential storage system. When you select this option the **User name** and **Password** fields are replaced with the **Credential ID** field. Currently, the only supported external storage system is CyberArk.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.

Chef server credentials

Chef server credentials access chef integrations with the instance.

These fields are available on the Credentials form for Chef server type credentials. This information comes from the settings you configured when you performed Chef server installation.

Field	Description
Name	Enter a unique and descriptive name for this credential.
Active	Enable or disable these credentials for use.
Admin Name	Provide the administrator name that you created during Chef server installation.
Admin Key	Enter the RSA private key that the Chef server generated when you created the administrator.
Validator Name	Enter the validator.
Validator Key	Enter the RSA private key that the Chef server generated when you created an organization.
Cert Name	Enter the certification name.
Cert Key	Enter the certification key.

CIM credentials

The CIM credential type manages access to a CIM server (also referred to as a CIMOM - Common Information Model Object Manager) for information about VMware ESX servers. This credential type is available for Discovery.

These fields are available in the Credentials form for CIM.

Field	Description
Name	Enter a unique and descriptive name for this credential.

Field	Description
Active	Enable or disable these credentials for use.
User name	Enter the user name to create in the Credentials table. Avoid leading or trailing spaces in user names. A warning appears if the platform detects leading or trailing spaces in the user name. For CIM discovery, the user must have the admin role.
Password	Enter the password.
Credential ID	Enter the unique key configured for external credentials in the JAR file uploaded to the MID Server for an external credential system. The **Credential ID** field has a limit of 40 characters. This field is only visible when the **External credential store** check box is selected.
Credential alias	Allow workflow creators to assign individual credentials to any activity in an Orchestration workflow or assign different credentials to each occurrence of the same activity type in an Orchestration workflow. To use the credential for discovering CIs not belonging to this CI type using Service Mapping and Discovery patterns, enter the table name for the CI type to which the CI belongs, for example cmdb_ci_apache_web_server.
External credential store	Select this check box to use an external credential storage system. When you select this option the **User name** and **Password** fields are replaced with the **Credential ID** field. External credential storage is only available when the External Credential Storage plugin in activated. **Note:** Currently, the only supported external storage system is CyberArk.
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
MID servers	Select one or more MID Servers from the list of available MID Servers. The credentials configured in this record are available to the MID Servers in this list. This field is available only when you select **Specific MID servers** from the **Applies to** field. **Note:** Selecting Specific **Specific MID servers** doesn't affect mid server selection. It's used only to decide which mid servers should have visibility to the credential. **Specific MID servers** isn't supported in Orchestration activities.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.
Windows MID Server	When active, the defined credential represents the MID Server service account.

Field	Description
Service Account	

Configure NetApp storage devices for CIM credentials

NetApp storage devices require additional configuration for Discovery to explore them.

Before you begin
Role required: admin

Procedure

1. Install the SMI-S Provider on the storage device host.

   See the Download the NetApp SMI-S Provider software package for instructions and requirements.

   Note:

   ServiceNow doesn't maintain the documentation on this site. This document can change without notice.

2. Create a user account and password for the SMI-S agent.
3. Create a credential record for the SMI-S agent credentials.
   Set the credential type to **CIM**.

Cloud credentials

Cloud credential types manage access to cloud-based applications, including Amazon Web Services and the Microsoft Azure cloud.

AWS Identity and Access Management (IAM) roles

If you have a MID Server installed on Amazon EC2 in an AWS cloud, and if that MID Server is configured to discover resources within the cloud, you can use security credentials provided by AWS Identity and Access Management (IAM) roles rather than credentials managed on your instance. These AWS credentials grant permissions in the cloud through an instance profile, based on roles. These credentials are temporary and rotate automatically on a configurable interval. When an IAM role is defined on the MID server. For details, see Configure the MID Server for AWS IAM roles.

Discovery ignores any credentials stored on the instance in favor of the credentials granted by the role in the instance profile. For more information on AWS instance profiles, see IAM Roles for Amazon EC2.

AWS credentials

AWS Credentials form fields

Field	Input value
Name	Unique and descriptive name for the AWS credentials.

AWS Credentials form fields (continued)

Field	Input value
Active	Option to use the credential.
Access Key ID	The **Access key ID** that you generated on the AWS Management Console, such as: APIAIOSFODNN7EXAMPLE.
Secret access key	The **Secret access key** that you generated on the AWS Management Console, such as: wPalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY.

Azure Service Principal credential form fields

Field	Value
Name	Enter the name of the service principal to register with the instance.
Tenant ID and	Paste the Azure **Directory ID** value from the Azure portal into the Cloud Management **Tenant ID** field.
Client ID	Paste the Azure **Application ID** value of the application that you registered in Azure into the Cloud Management **Client ID** field.
Authentication Method	Select **Client secret**. **Note:** **Client assertion** is not supported.
Secret key	Paste the secret key that was generated while creating the Azure Service Principal. This field appears when **Authentication method** is **Client secret**.

Azure Enterprise Agreement credentials

Azure Enterprise Agreement credentials are necessary for the billing functionality that the Cloud Management application provides.

Azure Enterprise Agreement credentials form fields

Field	Description
Name	Enter a descriptive name.
Enrollment number	Enter the enrolment number from Azure.
Access Key	Paste the access key that Azure provides.

Cloud Management credentials

This credential type is available for Orchestration.

Cloud Management credentials form fields

Field	Input value
Name	Enter a unique and descriptive name for this credential. For example, you might call it **Cloud Atlanta**.
Active	Enable or disable these credentials for use.
Type	Specify **AWS**.
User name	Enter the CIM user name to create in the Credentials table. Avoid leading or trailing spaces in user names. A warning appears if the platform detects leading or trailing spaces in the user name.
Password	Enter the CIM password.
SSH Passphrase	Enter a memorable phrase for key generation. For example, you might enter **Friday is a good day**.
SSH private key	Enter the SSH private key.
Authentication protocol	Select the **MD5**or **SHA** authentication protocol that was used to generate the **Authentication Key**.
Authentication Key	Enter a SSH-generated authentication key.
Privacy protocol	Enter one of the following privacy protocols that describes encryption for the **Privacy Key**: • **3DES** for Triple Data Encryption Standard (DES) • **AES128** for Advanced Encryption Standard (AES) with 128 bit encryption • **AES192** for AES with 192 bit encryption • **AES256** for AES with 256 bit encryption • **DES** for legacy DES encryption
Enter an additional privacy key.	
Credential alias	Allow workflow creators to assign individual credentials to any activity in an Orchestration workflow or assign different credentials to each occurrence of the same activity type in an Orchestration workflow.
External credential store	Select this check box to use an external credential storage system. When you select this option the **User name** and **Password** fields are replaced with the **Credential ID** field. Currently, the only supported external storage system is CyberArk.
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
Classification	Enter the Application Classification for CI discovery.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the

Cloud Management credentials form fields (continued)

Field	Input value
	credentials have the same order number (or none), the instance tries the credentials in a random order.

Cloud Management (CMP) node credentials

Cloud Management (CMP) node credentials associate credentials for a virtual server that Cloud Management provisions. The Cloud Management application automatically creates these credentials.

Note: You might need to deactivate these credentials if you no longer want them used, change the order precedence, or select a MID Server that is allowed to access them. Otherwise, you do not need to manually create or modify this type of credential.

CMP node credentials form fields

Field	Description
Name	The automatically generated name based on the datacenter where the virtual machine is located.
Active	If the credentials are active.
Applies to	Choose whether this credential is available to a specific MID Server or a all MID Servers.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.
User Name and Password	The virtual server user name and password.
SSH passphrase and SSH private key	The private key and the passphrase that protects the key if the virtual server requires it.
Authentication Protocol and Authentication Key	The private key and the passphrase that protects the key if the virtual server requires it.
Privacy Protocol and Privacy Key	The encryption protocol used with the virtual server and enter the privacy key.
Credential alias	Allow workflow creators to assign individual credentials to any activity in an Orchestration workflow or assign different credentials to each occurrence of the same activity type in an Orchestration workflow.

Cloud Management (CMP) SSH key pair credentials

Cloud Management (CMP) SSH key pairs store the keys that the Cloud Management application automatically generates when users provision stack resources.

Note: You might need to deactivate these credentials if you no longer want them used. Otherwise, you do not need to manually create or modify this type of credential.

CMP SSH key pair credentials form fields

Field	Description
Name	The automatically generated name.
Active	If the credentials are active.
SSH Public Key	The public key.
SSH Private Key	A secure private key that can be used instead of a password for SSH logins.

Container image repository credentials

The container image repository credentials manage access to private repositories for container image scanning. This credential type is available for Discovery.

These fields are available in the Credentials form for credential image repository type credentials.

Container Image Repository Credentials form

Field	Description
Name	User-assigned credential name.
User name	User name with read permissions for the repository.
Password	Password with read permissions for the repository.
Repository	Fully qualified domain name of the repository. For example: `docker.io/snow_images`.

Infoblox credentials

Infoblox credentials are required to set up IP pools (IPAM) in the Cloud Management application.

These fields are available on the Credentials form for Infoblox type credentials.

Field	Description
Name	Enter a unique and descriptive name for this credential.
Active	Enable or disable these credentials for use.
Applies to	Choose whether this credential is available to a specific MID Server or a all MID Servers.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have

Field	Description
	the same order number (or none), the instance tries the credentials in a random order. Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.
wAPI Version	Enter the version of wAPI you are using.
User Name and Password	Enter the InfoBlox user name and password.

JDBC credentials

The JDBC credential type manages access to a Java Database Connectivity (JDBC) connection. This credential type is available for Discovery and Orchestration.

These fields are available in the Credentials form for JDBC type credentials.

Field	Description
Name	Enter a unique and descriptive name for this credential.
Active	Enable or disable these credentials for use.
User name	Enter the user name to create in the Credentials table. Avoid leading or trailing spaces in user names. A warning appears if the platform detects leading or trailing spaces in the user name. For CIM discovery, the user must have the admin role.
Password	Enter the password.
Credential ID	Enter the unique key configured for external credentials in the JAR file uploaded to the MID Server for an external credential system. The **Credential ID** field has a limit of 40 characters. This field is only visible when the **External credential store** check box is selected.
Credential alias	Allow workflow creators to assign individual credentials to any activity in an Orchestration workflow or assign different credentials to each occurrence of the same activity type in an Orchestration workflow. To use the credential for discovering CIs not belonging to this CI type using Service Mapping and Discovery patterns, enter the table name for the CI type to which the CI belongs, for example cmdb_ci_apache_web_server.
External credential store	Select this check box to use an external credential storage system. When you select this option the **User name** and **Password** fields are replaced with the **Credential ID** field. External credential storage is only available when the External Credential Storage plugin in activated.

Field	Description
	Note: Currently, the only supported external storage system is CyberArk.
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
MID servers	Select one or more MID Servers from the list of available MID Servers. The credentials configured in this record are available to the MID Servers in this list. This field is available only when you select **Specific MID servers** from the **Applies to** field. **Note:** Selecting Specific **Specific MID servers** doesn't affect mid server selection. It's used only to decide which mid servers should have visibility to the credential. **Specific MID servers** isn't supported in Orchestration activities.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.
Windows MID Server Service Account	When active, the defined credential represents the MID Server service account.

JMS credentials

The JMS credentials type manages access to a Java Message Service (JMS). This credential type is available for Discovery and Orchestration.

These fields are available in the Credentials form for JMS.

Field	Description
Name	Enter a unique and descriptive name for this credential.
Active	Enable or disable these credentials for use.
User name	Enter the user name to create in the Credentials table. Avoid leading or trailing spaces in user names. A warning appears if the platform detects leading or trailing spaces in the user name. For CIM discovery, the user must have the admin role.
Password	Enter the password.
Credential ID	Enter the unique key configured for external credentials in the JAR file uploaded to the MID Server for an external credential system. The **Credential ID** field has a limit of 40 characters.

Field	Description
	This field is only visible when the **External credential store** check box is selected.
Credential alias	Allow workflow creators to assign individual credentials to any activity in an Orchestration workflow or assign different credentials to each occurrence of the same activity type in an Orchestration workflow. To use the credential for discovering CIs not belonging to this CI type using Service Mapping and Discovery patterns, enter the table name for the CI type to which the CI belongs, for example cmdb_ci_apache_web_server.
External credential store	Select this check box to use an external credential storage system. When you select this option the **User name** and **Password** fields are replaced with the **Credential ID** field. External credential storage is only available when the External Credential Storage plugin in activated. **Note:** Currently, the only supported external storage system is CyberArk.
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
MID servers	Select one or more MID Servers from the list of available MID Servers. The credentials configured in this record are available to the MID Servers in this list. This field is available only when you select **Specific MID servers** from the **Applies to** field. **Note:** Selecting Specific **Specific MID servers** doesn't affect mid server selection. It's used only to decide which mid servers should have visibility to the credential. **Specific MID servers** isn't supported in Orchestration activities.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.
Windows MID Server Service Account	When active, the defined credential represents the MID Server service account.

OAuth 2.0 credentials

OAuth 2.0 credentials enable ServiceNow to obtain access to user accounts on an HTTP service.

These fields are available in the Credentials form for OAuth 2.0.

OAuth 2.0 credentials form

Field	Input value
Name	Enter a unique and descriptive name for this credential. For example, you might call it **OAuth2 credential**.
Active	Specify whether this credential is active.
OAuth Entity Profile	An OAuth profile is a combination of a grant type and at least one scope.
Connect to Auth Server via MID Server	Connects your ServiceNow instance to an on-premise OAuth server that resides behind a firewall through a MID Server. It can also connect your ServiceNow instance to a cloud-based OAuth server through a MID server. When this option is enabled, the request for an OAuth token is sent through the MID Server. **Important:** • The option appears when the value in the **Grant type** field in the OAuth Entity Profile is set to either**Client Credentials**, **Authorization Code**, or **Resource Owner Password Credentials**. To learn how to set an OAuth entity profile for a third-party OAuth provider, see Connect to a third-party OAuth provider. • If you select the **Connect to Auth Server via MID Server** checkbox, you must identify the required MID Server or MID Servers from the **Applies to** list.
Applies to	Specify if the credential record is applicable for all MID Servers, or a specific MID Server. If specific, add the MID servers as necessary.

OAuth 2.0 credentials form (continued)

Field	Input value
	Important: Ensure that you are aware of these considerations if you have selected the **Connect to Auth Server via MID Server** check box. • Ensure that all the MID Servers selected in Applies to can communicate with the Auth server. This is required to execute the token request against the **Token URL** mentioned in the OAuth provider record (that is linked to the OAuth entity profile specified in the OAuth 2.0 credential record). 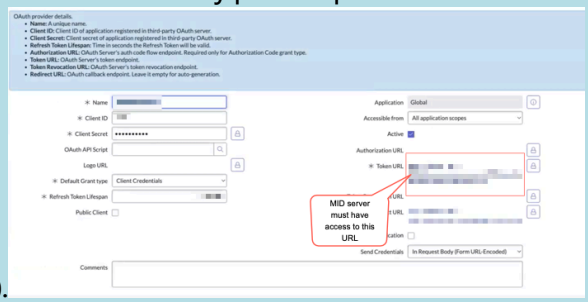 • Ensure that there is at least one MID Server (in the MID Servers selected in **Applies to**) with these configurations: ○ The value of the **Status** field is **Up**. ○ The value of the **Validated** field is **Yes**. ○ The capability of the MID Server is set to **REST** or **ALL**. To learn how to configure the MID server, see https://www.servicenow.com/docs/access?context=configure-capabilities&version=zurich&pubname=zurich-integrate-applications&ft:locale=en-US 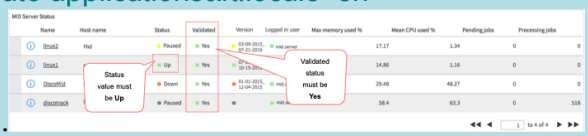 To learn more about these statuses, see MID Server dashboard.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.
Credential alias	Specify the credential alias that you want to tie to the OAuth 2.0 credential.
Integration Type	Indicates the integration type for the credential. Invoke an API of a third-party with an OAuth request that generates an OAuth token that is system or user specific. Following are the integration types:

OAuth 2.0 credentials form (continued)

Field	Input value
	• **System**: Pull the token information based on the requester profile. The **System** integration type supports the following authentication mechanisms: **1.** Security Assertion Markup Language (SAML) **2.** JSON Web Token (JWT) • **Personal**: Pull the token information that is user-specific. The MID Server user must have the `oauth_admin` role. The **Personal** and **System** integration types support the following grant types: **1.** Authorization Code **2.** Resource Owner Password Credentials If this **Personal** is selected on the OAuth Requestor Profile page, an additional flag called as **Personal** is displayed. **Note:** • Any information that is related to a user can only be accessed with user-specific OAuth tokens with the Integration Type as **Personal**. • To use the session user-related token, you have to select the **Run As** filed in the Flow properties as **User who initiates session**.

SAP credentials

The SAP credential type manages access to SAP JCo systems. This credential type is available for Discovery and Orchestration.

These fields are available in the Credentials form for SAP type credentials.

Field	Description
Name	Enter a unique and descriptive name for this credential.
Active	Enable or disable these credentials for use.
User name	Enter the user name to create in the Credentials table. Avoid leading or trailing spaces in user names. A warning appears if the platform detects leading or trailing spaces in the user name. For CIM discovery, the user must have the admin role.
Password	Enter the password.
Credential ID	Enter the unique key configured for external credentials in the JAR file uploaded to the MID Server for an external credential system. The **Credential ID** field has a limit of 40 characters. This field is only visible when the **External credential store** check box is selected.
Credential alias	Allow workflow creators to assign individual credentials to any activity in an Orchestration workflow or assign different credentials to each occurrence of the same activity type in an Orchestration workflow.

Field	Description
	To use the credential for discovering CIs not belonging to this CI type using Service Mapping and Discovery patterns, enter the table name for the CI type to which the CI belongs, for example cmdb_ci_apache_web_server.
External credential store	Select this check box to use an external credential storage system. When you select this option the **User name** and **Password** fields are replaced with the **Credential ID** field. External credential storage is only available when the External Credential Storage plugin in activated. **Note:** Currently, the only supported external storage system is CyberArk.
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
MID servers	Select one or more MID Servers from the list of available MID Servers. The credentials configured in this record are available to the MID Servers in this list. This field is available only when you select **Specific MID servers** from the **Applies to** field. **Note:** Selecting Specific **Specific MID servers** doesn't affect mid server selection. It's used only to decide which mid servers should have visibility to the credential. **Specific MID servers** isn't supported in Orchestration activities.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.
Windows MID Server Service Account	When active, the defined credential represents the MID Server service account.

SNMP credentials

Discovery explores many kinds of devices (switches, routers, printers, and so on) using the SNMP protocol. Credentials for SNMP don't include a user name, just a password, called the *`community string`*.

The default read-only community string for many SNMP devices is public, and Discovery will try that automatically. Enter the appropriate SNMP credentials if they differ from the *`public`* community string.

Discovering SNMP uses all community strings that are configured. This behavior does not apply to discovering SNMPv3.

The default Orchestration activity SNMP Query returns the object identifier (OID) of a device and requires SNMP credentials.

SNMP community credentials

The SNMP Community credential type manages access to discover many kinds of devices (switches, routers, printers, etc.) using the SNMP protocol. This credential type is available for Discovery, Service Mapping, and Orchestration.

Credentials for SNMP do not include a user name, just a password (the community string). The default read-only community string for many SNMP devices is public, and the system will try that automatically. Enter the appropriate SNMP credentials if they differ from the public community string.

These fields are available in the Credentials form for SNMP community.

Field	Description
Name	Enter a unique and descriptive name for this credential.
Active	Enable or disable these credentials for use.
User name	Enter the user name to create in the Credentials table. Avoid leading or trailing spaces in user names. A warning appears if the platform detects leading or trailing spaces in the user name. For CIM discovery, the user must have the admin role.
Password	Enter the password.
Credential ID	Enter the unique key configured for external credentials in the JAR file uploaded to the MID Server for an external credential system. The **Credential ID** field has a limit of 40 characters. This field is only visible when the **External credential store** check box is selected.
Credential alias	Allow workflow creators to assign individual credentials to any activity in an Orchestration workflow or assign different credentials to each occurrence of the same activity type in an Orchestration workflow. To use the credential for discovering CIs not belonging to this CI type using Service Mapping and Discovery patterns, enter the table name for the CI type to which the CI belongs, for example cmdb_ci_apache_web_server.
External credential store	Select this check box to use an external credential storage system. When you select this option the **User name** and **Password** fields are replaced with the **Credential ID** field. External credential storage is only available when the External Credential Storage plugin in activated. **Note:** Currently, the only supported external storage system is CyberArk.
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
MID servers	Select one or more MID Servers from the list of available MID Servers. The credentials configured in this record are available to the MID Servers in this list. This field is available only when you select **Specific MID servers** from the **Applies to** field.

Field	Description
	Note: Selecting Specific **Specific MID servers** doesn't affect mid server selection. It's used only to decide which mid servers should have visibility to the credential. **Specific MID servers** isn't supported in Orchestration activities.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.
Windows MID Server Service Account	When active, the defined credential represents the MID Server service account.

SNMPv3 credentials

SNMPv3 credentials accept a privacy protocol and an additional privacy key and are available for Discovery and Orchestration. For external storage in CyberArk, you can select a privacy account key.

These fields are available in the Credentials form for SNMPv3.

SNMPv3 credential fields

Field	Input value
Name	Unique and descriptive name for this credential. For example, you might call it **SNMP Community Atlanta**.
Active	Enable or disable these credentials for use.
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
MID servers	Select one or more MID Servers from the list of available MID Servers. The credentials configured in this record are available to the MID Servers in this list. This field is available only when you select **Specific MID servers** from the **Applies to** field.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.

SNMPv3 credential fields (continued)

Field	Input value
User name	Enter the SNMP user name. Avoid leading or trailing spaces in user names. A warning appears if the platform detects leading or trailing spaces in the user name.
Authentication protocol	Select the authentication type to use for this credential. The choices are: • MD5 • SHA-1 • SHA-224 • SHA-256 • SHA-384 • SHA-512
Authentication Key	Enter the authentication key to use for this credential.
Privacy protocol	Select the encryption protocol for this credential. The choices are: • 3DES • AES128 • AES192 • AES256 • DES
Privacy key	Enter the key associated with the selected privacy protocol.
Credential ID	Enter the unique key configured for external credentials in the JAR file uploaded to the MID Server for an external credential provider. The **Credential ID** field has a limit of 40 characters. This field is only visible when the **External credential store** check box is selected.
Privacy Credential ID	Enter the privacy account key configured for SNMPv3 credentials in CyberArk. If you are using a privacy protocol in CyberArk, this field must have the same value as the **Name** field for the CyberArk SNMPv3 Privacy Account. This field is only visible for SNMPv3 when you select **CyberArk** from the **Credential Store Type** field. If you are not using a privacy key for CyberArk, leave this field empty.
Credential alias	Allow workflow creators to assign individual credentials to any activity in an Orchestration workflow or assign different credentials to each occurrence of the same activity type in an Orchestration workflow.
External credential store	Select this check box to use an external credential storage system. When external storage is enabled, the **Credential ID** field appears. External credential storage is only available when the External Credential Storage plugin is activated and the **External Storage** view is selected.

SNMPv3 credential fields (continued)

Field	Input value
Credential Store Type	Select the external storage provider. Select **CyberArk** only if you are using a CyberArk SNMPv3 privacy key. The **Privacy Credential ID** field appears to allow entry of the key.
Use Context	Select this check box to add a context value for this credential. This field is visible in the **Discovery** view. Contexts are not currently supported for external credential storage. **Note:** A context is a collection of management information accessed by an SNMPv3 credential that references a specific OID. Contexts are sometimes referenced to collect information about the device that cannot be accessed by the normal credential. A context can be provided by the manufacturer or configured separately. If you have multiple SNMPv3 credentials with the same user name and keys, but some of your devices need a context and some do not, then you will need to create separate records for each device.
Context Name	Enter the context name value for this credential. This should only be used if you have devices that require this value for full access. This field is visible when the **Use Context** check box is selected.

SSH credentials

Discovery, Orchestration, and Integration Hub explore UNIX and Linux devices by using SSH credentials to execute commands over Secure Shell (SSH). SSH commands must run with root privileges, either with root credentials or through the use of sudo. SSH private key credentials provide additional security.

Granting root privileges

Before granting root privileges, review your security policy and options with your organization's security team.

Use either of these approaches to allow users to run SSH commands with root privileges:

- Give other credentials for Discovery, Orchestration, or Integration Hub, but grant the user in those credentials the right to execute certain commands with root privileges, using sudo. This is a secure way to grant limited privileges. Discovery, Orchestration, or Integration Hub use sudo on any probe that has the `must_sudo`parameter set to **true** (it defaults to **false**). However, each system must be configured to allow sudo to work. This is done by editing the **/etc/sudoers** file using the **visudo** command.
- Give **root** credentials. These are obviously the most powerful credentials, but may not be desirable from a security perspective. If Discovery, Orchestration, or Integration Hub have the root credentials to any UNIX or Linux system, no further configuration is required.

Privileged commands

The platform provides default privileged commands for the MID Server to use and the ability to add additional commands to the system. For details about using sudo and other privileged commands, see MID Server privileged commands.

SSH private key credential type

Note: SSH private key credentials should be used in most cases. They provide better security than SSH password credentials, including against MitM (man-in-the-middle) attacks in which communications between two parties are intercepted.

Field	Input value
Name	Unique and descriptive name for this credential. For example, you might call it **SSH Atlanta**.
Active	Enable or disable these credentials for use.
User name	Enter a UNIX or Linux user name. Avoid leading or trailing spaces in user names. A warning appears if the platform detects leading or trailing spaces in the user name.
Password	Enter the UNIX or Linux password. For **SSH Private Key** type credentials, enter the sudo password if one is required for the user name.
SSH passphrase	Type a secure SSH passphrase. This field is available only for **SSH Private Key** credentials.
SSH private key	Enter a secure, RSA, DSA, ECDSA or ED25519 private key. The private key must be entered in the proper format to ensure it is correctly encrypted. The private key must start with the string `-----BEGIN`. Here is an example of a correctly formatted RSA private key: `-----BEGIN RSA PRIVATE KEY-----` `MIIEogIBAAKCAQEAsEK65scPssPSobpDFMpR+Btv3MS4Q7NP8ERaSt` `RZsh3IWz+x...` `...7hrxV2dbSug6OFahyupGWBGtPnXm5PaE2X5WPLuUj94ue48i1Fs` `-----END RSA PRIVATE KEY-----` An example of a DSA key: `-----BEGIN DSA PRIVATE KEY-----` `MIIEogIBAAKCAQEAsEK65scPssPSobpDFMpR+Btv3MS4Q7NP8ERaSt` `RZsh3IWz+x...` `...7hrxV2dbSug6OFahyupGWBGtPnXm5PaE2X5WPLuUj94ue48i1Fs` `-----END DSA PRIVATE KEY-----` An example of a ECDSA key: `-----BEGIN EC PRIVATE KEY-----` `MIIEogIBAAKCAQEAsEK65scPssPSobpDFMpR+Btv3MS4Q7NP8ERaSt` `RZsh3IWz+x...` `...7hrxV2dbSug6OFahyupGWBGtPnXm5PaE2X5WPLuUj94ue48i1Fs` `-----END EC PRIVATE KEY-----` And an example of an ED25519 private key: `-----BEGIN OPENSSH PRIVATE KEY-----` `b3BlbnNzaC1rZXktdjEAAAAABG5vbmUAAAAEbm9uZQAAAAAAAAABAA` `AAMwAAAAtzc2gtZW` `QyNTUxOQAAACAlY1qhcdwx8VQzZ5XaIC5ltQpjRr3lIlq/aE66mufm` `iwAAAKDQUtxZOFLc`

Field	Input value
	<pre>WQAAAAtzc2gtZWQyNTUxOQAAACAlYlqhcdwx8VQzZ5XaIC5ltQpjRr 3lIlq/aE66mufmiw AAAECuvsTkFUPdpTh0kw23i8TYx19qsFOZ3TRgowkkHBh6wSViWqFx 3DHxVDNnldogLmW1 CmNGveUiWr9oTrqa5+aLAAAAGmFiaGluYXYuc3V0YXJATVJFTUE3OT AzMkI3AQID -----END OPENSSH PRIVATE KEY-----</pre> **Note:** For an ED25519 private key, only the OpenSSH key format is supported, which is generated using OpenSSH SSH-keygen utility. The ServiceNow AI Platform supports private keys in the PEM format generated by the OpenSSH ssh-keygen utility. To convert PPK keys that were generated by PuTTY: • Open your private key in PuTTYGen. • Export it in OpenSSH format from the menu **Conversions** > **Export OpenSSH key**. • Save the new OpenSSH key.
SSH Certificate	Enter an RSA or ED25519 based OpenSSH certificate. When the certificate is entered, a private key is used for certificate based authentication. This authentication is supported from OpenSSH 7.8 onwards.
Credential alias	• Allow flow designers to use aliases to manage connection and credential information. Using an alias eliminates the need to configure multiple credentials and connection information profiles when using multiple environments. If the connection or credential information changes, you do not need to update any actions that use the connection. For more information, see Connections and Credentials. • Allow workflow creators to assign individual credentials to any activity in an Orchestration workflow or assign different credentials to each occurrence of the same activity type in an Orchestration workflow.
External credential store	Select this check box to use an external credential storage system. When you select this option the **User name** and **Password** fields are replaced with the **Credential ID** field. Currently, the only supported external storage system is CyberArk.
MID servers	Select one or more MID Servers from the list of available MID Servers. The credentials configured in this record are available to the MID Servers in this list. This field is available only when you select **Specific MID servers** from the **Applies to** field.
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
Order	The order (sequence) in which the platform tries this credential as it attempts to log onto devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the

Field	Input value
	credentials have the same order number (or none), Discovery or Orchestration tries the credentials in a random order.

SSH credential type

These fields are available in the SSH credentials form.

Field	Description
Name	Enter a unique and descriptive name for this credential.
Active	Enable or disable these credentials for use.
User name	Enter the user name to create in the Credentials table. Avoid leading or trailing spaces in user names. A warning appears if the platform detects leading or trailing spaces in the user name. For CIM discovery, the user must have the admin role.
Password	Enter the password.
Credential ID	Enter the unique key configured for external credentials in the JAR file uploaded to the MID Server for an external credential system. The **Credential ID** field has a limit of 40 characters. This field is only visible when the **External credential store** check box is selected.
Credential alias	• Allow flow designers to use aliases to manage connection and credential information. Using an alias eliminates the need to configure multiple credentials and connection information profiles when using multiple environments. If the connection or credential information changes, you do not need to update any actions that use the connection. For more information, see Connections and Credentials. • Allow workflow creators to assign individual credentials to any activity in an Orchestration workflow or assign different credentials to each occurrence of the same activity type in an Orchestration workflow. To use the credential for discovering CIs not belonging to this CI type using Service Mapping and Discovery patterns, enter the table name for the CI type to which the CI belongs, for example cmdb_ci_apache_web_server.
External credential store	Select this check box to use an external credential storage system. When you select this option the **User name** and **Password** fields are replaced with the **Credential ID** field. External credential storage is only available when the External Credential Storage plugin in activated. **Note:** Currently, the only supported external storage system is CyberArk.
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
MID servers	Select one or more MID Servers from the list of available MID Servers. The credentials configured in this record are available to the MID Servers in this list. This field is available only when you select **Specific MID servers** from the **Applies to** field.

Field	Description
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.

Commands that require root privileges for Discovery, Orchestration, and Integration Hub

These examples assume that the user name is **Disco**. Substitute the actual user name and ensure that the paths for the commands match the paths on the system.

Note: Sudo commands do not work with private key credentials, because there is no password to supply to the sudo command. A solution is to add the *NOPASSWD* option to the sudo configuration. For example, you might enter: `disco ALL=(root) NOPASSWD: /usr/sbin/dmidecode,/usr/sbin/lsof,/sbin/ifconfig`.

UNIX and Linux commands requiring root privileges

Command	Purpose
HP-UX	
adb	Gathers CPU speed and memory. • **/etc/sudoers line example**: Disco ALL=(root) /usr/bin/adb • **Used by**: Discovery
All Linux and UNIX versions	
chage	Changes the number of days between password changes and the date of the last password change. • **/etc/sudoers line example**: Disco ALL=(root) /usr/bin/chage • **Used by**: Orchestration and Integration Hub
chpasswd	Changes user passwords. • **/etc/sudoers line example**: Disco ALL=(root) /etc/chpasswd • **Used by**: Orchestration and Integration Hub
All Linux	
dmidecode	Gathers several pieces of information about the hardware, including the serial number embedded within the motherboard. • **/etc/sudoers line example**: Disco ALL=(root) /sbin/dmidecode • **Used by**: Discovery
fdisk	Gathers the disks and size information on the system.

UNIX and Linux commands requiring root privileges (continued)

Command	Purpose
	• **/etc/sudoers line example**: Disco ALL=(root) /usr/bin/fdisk -l • **Used by**: Discovery
multipath	Gathers device mappings for MPIO. • **/etc/sudoers line example**: Disco ALL=(root) /usr/bin/multipath -ll • **Used by**: Discovery
ls	Gathers the contents of a directory. • **/etc/sudoers line example**: Disco ALL=(root) /usr/bin/ls, /bin/ls • **Used by**: Discovery
Linux and Solaris	
dmsetup	Examines a low level volume. • **/etc/sudoers line example**: ○ Disco ALL=(root) /usr/bin/dmsetup table * ○ Disco ALL=(root) /usr/bin/dmsetup ls • **Used by**: Discovery
All UNIX versions	
lsof	Determines the relationship between processes and the connections being made to the system. • **/etc/sudoers line example**: Disco ALL=(root) /sbin/lsof • **Used by**: Discovery
oratab	Grants read access to the oratab file for locating the Oracle Home and pfile. • **/etc/sudoers line example**: N/A • **Used by**: Discovery
Solaris	
iscsiadm	Gets iSCSI IQNs • **/etc/sudoers line example**: ${sudo:iscsiadm list target -S} • **Used by**: Discovery
fcinfo	Gets WWPNs for ports. • **/etc/sudoers line example**: ${sudo:fcinfo remote-port -sl -p $port} • **Used by**: Discovery

UNIX and Linux commands requiring root privileges (continued)

Command	Purpose
prtvtoc	Reports information about disk partitions. • **/etc/sudoers line example**: Disco ALL=(root) /usr/bin/prtvtoc • **Used by**: Discovery
pfiles	Used for gathering TCP connections information. • **/etc/sudoers** line example: Disco ALL=(root) /usr/bin/pfiles • **Used by**: Discovery
pgrep	Used for listing process IDs of a particular region to run pfiles on. • **/etc/sudoers line** example: Disco ALL=(root) /usr/bin/pgrep • **Used by**: Discovery
/usr/bin/ps	Lists running process. As an alternative to running with root access, add a proc_owner role. • **/etc/sudoers line example**: Disco ALL=(root) /usr/bin/ps • **Used by**: Discovery
/usr/ucb/ps	Lists running process. As an alternative to running with root access, add a proc_owner role. The use of the `/usr/ucb/ps` command is deprecated as of Solaris 11. Because Discovery, Orchestration, and Integration Hub require the use of this command for all Solaris versions, you must install the ucb utility manually on Solaris 11 systems. For instructions, see KB0564262. • **/etc/sudoers line example**: Disco ALL=(root) /usr/ucb/ps • **Used by**: Discovery

For a list of privileged commands that you need for Discovery and Service Mapping, see Service Mapping commands requiring a privileged user. This list includes commands that require elevated rights to discover and map Unix-based hosts in your organization.

Access Requirements for Non-Root Credentials

If you do not provide Discovery with root access credentials, you must provide credentials with the following access requirements.

Application	File or Directory	Access Required
Apache	httpd.conf	Read
Hbase	hbase-site.xml	Read
JBoss	jboss-service.xml	Read
	JBoss home directory	Read
	web.xml	Read
MySQL	my.cnf	Read
NGINX	nginx.conf	Read
Oracle	oratab	Read
	Associated (s) pfiles	Read
Oracle Listener	lsnrctl	Execute
	listener.ora	Read
Tomcat	catalina.jar	Read
	server.xml	Read
	web.xml	Read
Unix	/etc/*release	Read
	/etc/bashrc	Read
	/etc/profile	Read
	/proc/cpuinfo	Read
	/proc/vmware/sched/ncpus	Read
	/var/log/dmesg	Read
	APD directory	Read
WebSphere	cell.xml	Read
	server.xml	Read
	serverindex.xml	Read

VMware credentials

The VMware credentials type manages access to vCenter credentials.

Applications that access VMware cloud resources need access to VMware credentials. For example, the VMware credential type enables Discovery to explore VMware's vCenter running on a Windows machine to discover ESX machines, virtual machines, and resource pools. The VMware Discovery and automation API (vCenter API) now provides the globally unique serial number for computer CIs. CIM credentials aren't needed to enable access to each VMware host.

Note: Windows credentials aren't necessary for vCenter Discovery, when valid VMware credentials are used.

Important: Do not use **VMware** Type credentials for Orchestration activities that perform work on the individual virtual machines cloned by vCenter (for example, restarting a Linux VM). For these activities, the credential **Type** depends on the operating system of the virtual machine (either **SSH** or **Windows**).

VMWare credentials form

Field	Description
Name	Enter a unique and descriptive name for the VMware credentials.
Active	Enable or disable these credentials for use.
User name	Enter the user name that you use for your VMware account. Avoid leading or trailing spaces in user names. A warning appears if the platform detects leading or trailing spaces in the user name. The VMware credentials must have the read-only role in vCenter.
Password	Enter the password for the VMware account.
Applies to	Select one or more MID Servers from the list of available MID Servers. The credentials configured in this record are available to the MID Servers in this list. This field is available only when you select **Specific MID servers** from the **Applies to** field.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.

Windows credentials

Windows credentials provide access to Windows computers. This credential type is available for Discovery and Orchestration.

Credential requirements

Discovery and Orchestration have the following requirements for Windows credentials:

- Install a MID Server on a Windows host as a service.
- Add Windows credentials to one of these locations:
  - An entry in the Credentials `[windows_credentials]` table
  - A MID Server service account to run as a specific Windows user or domain account.

Granting proper permissions

To provide sufficient permissions, Windows credentials must be one of the following:

- A domain user with local administrator access on the target Windows hosts.

> **Important:** If User Account Control (UAC) is enabled on the Windows operating system where discovery runs, and the user account is part of the local Administrators group, administrator tasks may fail. To avoid interruptions, we recommend disabling UAC. For more information, see the Why does the User Access Control (UAC) need to be disabled for Windows Discovery? article in the Now Support Knowledge Base

- A local account that has administrator privileges and UAC disabled on the same target host.

- A user who meets the requirements of Windows probes and permissions (Discovery only).
- A user who meets the requirements of the Orchestration activity to be run (Orchestration only).

Note: No logon privileges are needed. Account does NOT need to be interactive.

Security around granting privileged access can be enhanced by using JEA profiles to run Discovery. For more information, see Microsoft Just Enough Administration (JEA) for Discovery.

Workgroup computers

To run Powershell commands to discover a Workgroup computer, configure the MID Server credentials for either of these users:

- Built-in administrator account on the Workgroup computer.
- Domain user on the Workgroup computer.

Multi-domain configuration

To enable Windows credentials to function across multiple domains, make sure to sure to use the correct name formats and MID Server configuration.

Discovery and Orchestration support Windows domain credentials in both **User Principal Name** and **Down-Level Logon Name** user name formats. For example, **Domain\UserName** or **UserName@example.domain.com**. You can provide Windows workgroup credentials in the following format: WORKGROUP\UserName.

Note: You can also provide a local account by using the **.** user name.

These additional actions are required to enable credentials to function across multiple Windows domains.

Condition	Additional actions required
MID Server host on the same domain as the Windows target.	None
MID Server host on a different domain than the Windows target.	Ensure that PowerShell 3.0 (or higher up to 5.1) is installed on the MID Server host.
MID Server host on a different domain than the Microsoft SQL Server target.	See MSSQL server discovery .

Windows credentials type

These fields are available in the Credentials form for Windows:

Field	Description
Name	Enter a unique and descriptive name for this credential.
Active	Enable or disable these credentials for use.

Field	Description
User name	Enter the user name to create in the Credentials table. Avoid leading or trailing spaces in user names. A warning appears if the platform detects leading or trailing spaces in the user name. For CIM discovery, the user must have the admin role.
Password	Enter the password.
Credential ID	Enter the unique key configured for external credentials in the JAR file uploaded to the MID Server for an external credential system. The **Credential ID** field has a limit of 40 characters. This field is only visible when the **External credential store** check box is selected.
Credential alias	Allow workflow creators to assign individual credentials to any activity in an Orchestration workflow or assign different credentials to each occurrence of the same activity type in an Orchestration workflow. To use the credential for discovering CIs not belonging to this CI type using Service Mapping and Discovery patterns, enter the table name for the CI type to which the CI belongs, for example cmdb_ci_apache_web_server.
External credential store	Select this check box to use an external credential storage system. When you select this option the **User name** and **Password** fields are replaced with the **Credential ID** field. External credential storage is only available when the External Credential Storage plugin in activated. **Note:** Currently, the only supported external storage system is CyberArk.
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
MID servers	Select one or more MID Servers from the list of available MID Servers. The credentials configured in this record are available to the MID Servers in this list. This field is available only when you select **Specific MID servers** from the **Applies to** field. **Note:** Selecting Specific **Specific MID servers** doesn't affect mid server selection. It's used only to decide which mid servers should have visibility to the credential. **Specific MID servers** isn't supported in Orchestration activities.
Order	Order (sequence) in which Discovery tries this credential as it attempts to log on to devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), the instance tries the credentials in a random order.
Windows MID Server Service Account	When active, the defined credential represents the MID Server service account.

Configure Windows credentials for the MID Server

Configure the MID Server to use either the credentials of its own Windows service or credentials from the Credentials [discovery_credentials] table.

Before you begin

Role required: admin

Procedure

1. Configure the MID Server to use credentials from the MID Server service account.
   a. Set the Configure Windows MID Server service credentials to a user who meets the permission requirements.
   b. Verify the user name meets the name format requirements.
   c. Fill in the fields on the form, as appropriate.
   d. Verify the credentials meet domain requirements.
2. Configure the MID Server use credentials from the Credentials [discovery_credentials] table.
   a. Add individual Windows credentials to the Credentials `[windows_credentials]` table.
      - Verify each credential meets the permission requirements.
      - Verify each username meets the name format requirements.
      - Verify each credential meets the Windows domain requirements.
   b. **Optional:** Configure the MID Server to use Powershell by setting the `mid.use_powershell` parameter to **true**.
      See Configuring MID Servers.
   c. Select the Windows MID Server Service Account check box to create a credential that represents the windows MID Server service account to run as a specific Windows user or domain account.

Credential affinity for Discovery and Orchestration

Credential affinity is an association between a set of credentials and a device on your network.

When Discovery or Orchestration first attempts to access a device, they try all available credentials until they find the correct ones. After identifying the credentials for a device, Discovery and Orchestration create an affinity between the credentials and the device using the Credential Affinity `[dscy_credentials_affinity]` table. All subsequent discoveries or Orchestration activities attempt to match the credentials in this table with a device for which an affinity exists. If credentials for a device change, Discovery and Orchestration try all available credentials again until they create a new affinity.

Credential Affinity diagram

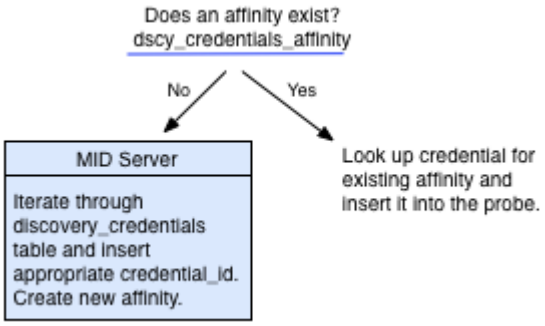

Note: If Orchestration and Discovery are installed, and credential alias is enabled, multiple affinities can exist. In this case, the platform looks up credentials for each affinity and inserts the credential for the affinity with the lowest order into the probe.

Credentials troubleshooting

Review the `<credentials_debug>` section of the ECC queue payload to troubleshoot issues with credentials.

Certain probes support credential debugging. Credential debugging inserts a `<credentials_debug>` section in the payload that the MID Server returns to the instance ECC queue. You can view the `<credentials_debug>` section to see detailed information about the credential lookup.

The `<credentials_debug>` section appears in the payload if:

- Credentials fail for WMIRunner probe, PowerShell probe, JMS, or SSHCommand probe.
- You set the *credentials_debug* parameter to **true** for the WMIRunner, PowerShell, or SSHCommand probes. If you set the parameter to true, the `<credentials_debug>` section appears even if the credential lookup is successful.

The `<credentials_debug>` section shows:

- Information about the credential search, such as the credential types, tags, and affinities.
- The IP address targeted.
- Information about each credential (in order) that the MID Server used, including the type, classification, tag, name, Sys ID, and external credential ID if present.

Sample payload showing invalid credentials

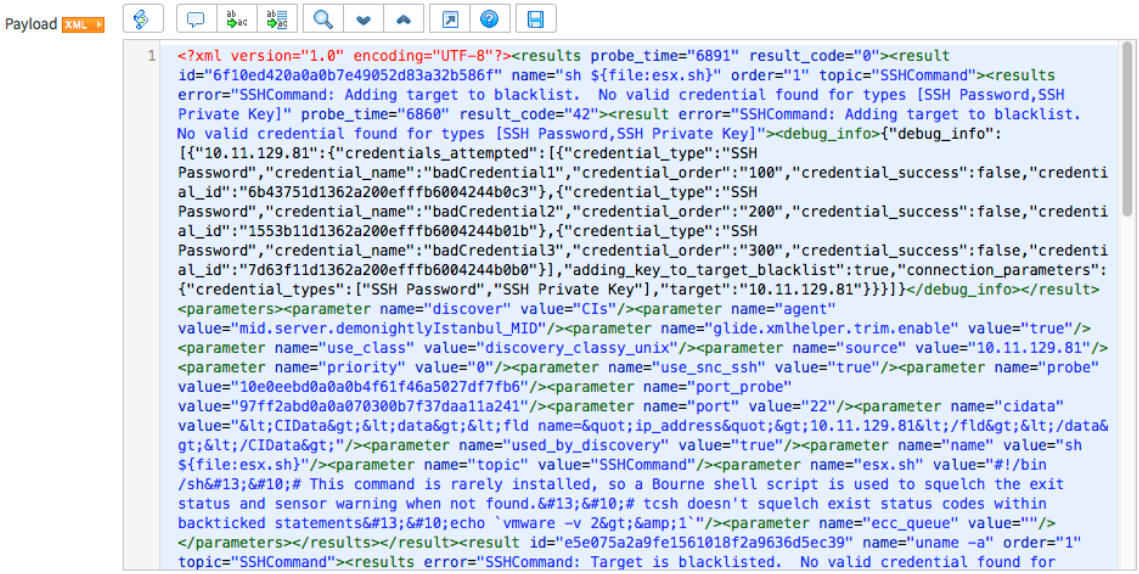

Details appear for the PowerShell parameter:

- If the local MID Server credentials were used after all the Windows credentials failed.
- If the credentials were skipped because you are trying to discover the same machine that the MID Server is on.
- If the mid.powershell.use_credentials parameter is set to `true`.

Details appear for the SSHCommand:

- If the credential search was skipped because the target IP is excluded.
- If the target IP was added to the exclusion list.

Note: The MID Server saves IP addresses for failed credential searches in an exclusion list in cache memory. This exclusion list specifies which devices the MID Server should stop trying to access. IP addresses are added to the exclusion list after every credential has failed. The IP addresses are cleared from the exclusion list cache either after five minutes, if the MID Server is restarted, or if the credential records on the instance are updated.

External credential storage

An instance can store credentials used by Discovery, Orchestration, and Service Mapping in an external credential repository rather than directly in a ServiceNow credentials record.

The instance maintains a unique identifier for each credential, the credential type (such as SSH, SNMP, or Windows), and any credential affinities. The MID Server obtains the credential identifier from the instance, and then uses a customer-provided JAR file to resolve the identifier from the repository into a usable credential. Currently, the ServiceNow® platform supports the use of the CyberArk vault or BeyondTrust for external credential storage.

External credential storage architecture

External credential storage architecture

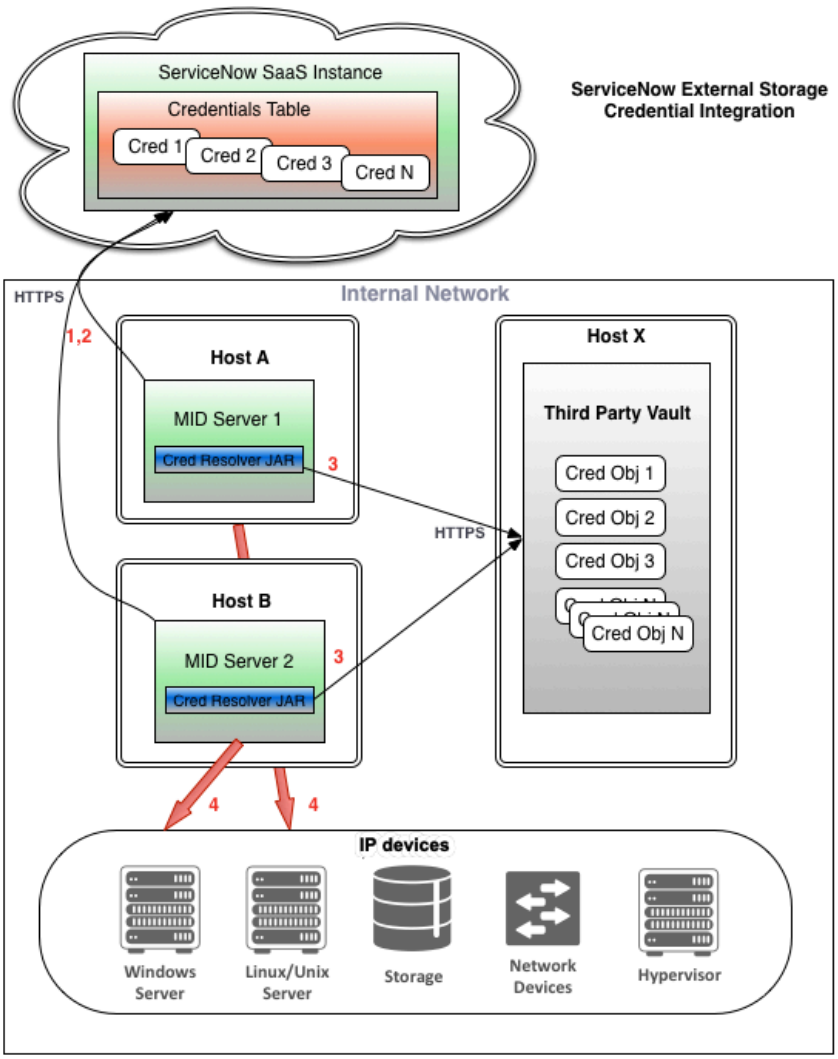

Credential process flow

The MID Server retrieves credentials from an external store using this process:

1. MID Server downloads credential objects from the ServiceNow Credentials [discovery_credentials] table that contain the corresponding credential ID from the target vault.
2. As each probe or pattern runs from Discovery or Orchestration jobs, the MID Server requests the credential by passing information such as credential ID, target IP address, and credential type to the Credential Resolver Java Jar file. The details about the correct credential object to retrieve from the vault are determined by the Credential Resolver.

   Many Credential Resolvers such CyberArk call an application supplied by the third-party vault vendor running on same machine as the MID Server. That application can often be configured to cache credentials and knows to update the cache when a credential changes in the vault, which is very important to avoid unnecessary network calls to the vault each time MID Server requests a credential. The Credential Resolver (using optional vendor application if present) makes a call to the vault to get the actual user name, password, etc.

For Credential Resolvers supplied out-of-box (only CyberArk today), the MID Server only caches a credential for up to several seconds using encryption in MID Server process memory. This means the MID Server can make multiple requests to the Credential Resolver for the same credential even when discovering a single device. Contact third party vendors for information about caching implementations for other Credential Resolvers.

3. MID Server executes the probe with the appropriate credential.

Note: Credential affinity still applies. The mechanism remains the same, since the only real difference from the MID Server's perspective is that the real credential details (user name and password) come from the third party vault.

External credential storage logging

The MID Server posts log messages about external credential storage.

If the repository encounters an error while attempting to resolve a credentials request, the MID Server posts log messages with this prefix: `Problem with client's CredentialResolver:`

Components installed with External Credential Storage

Business rule

The External Credential Storage business rule performs the following tasks when an administrator makes any change to the Enable External Credential Storage property:

- Changes the view for the Credentials record list and form to the External Storage view. This view enables users to see the **Credential ID** column in the list.
- Instructs the MID Server to refresh its credentials cache in preparation for a change in the way credentials are obtained.

Property

A property called Enable External Credential Storage [com.snc.use_external_credentials] enables or disables the External Credential Storage plugin after it's activated. The property is located in **Discovery Definition** > **Properties** and **Orchestration** > **MID Server Properties**, and is enabled when you activate the plugin.

If you disable external credential storage with the system property, the system automatically sets all the external credentials to inactive in the instance. If you re-enable the feature with this property, the system doesn't reset the external credential records to active. You must reactivate each credential record manually.

Request external credential storage for Discovery and Orchestration

The External Credential Storage plugin is available by request.

Before you begin
Role required: admin

About this task
There are two ways to request a plugin:

- Access the Now Support Service Catalog directly by selecting **Automation Store** > **Service Catalog** > **Activate Plugin** on Now Support.
- Access the Now Support Service Catalog through the All Applications page on your instance by following these steps.

For additional details about requesting a plugin, see Requesting a Plugin from the Service Catalog [KB0751715] article in the Now Support Knowledge Base.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Select **Request plugin** to open the **Activate Plugin** form on Now Support.
3. On the **Activate Plugin** form, provide the following information.

Activate Plugin form

Field	Description
What is your target instance	Select the instance that you want to activate the plugin on.
Which plugin would you like to activate	Select the name of the plugin to activate. **Note:** If the system doesn't list the plugin you want or if you're activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	Select the date and time to activate the plugin.

Example

For example, see the following form to activate the Event Management plugin on an instance named SNC Instance.

4. Select **Submit**.
   After the maintenance window, the system installs the plugin on your instance. To confirm the installation, go to the Installed tab in the Application Manager.

External credential storage configuration

Configure your instance to obtain credentials from a remote repository.

These procedures assume that you already have an external repository configured with the credentials you want to protect. The credential identifier configured in the ServiceNow instance must be mapped to the actual credential in the repository through the JAR file.

Note: ServiceNow supports two external vaults at a time: one default CyberArk credential resolver, and one custom external credential resolver. Creating a custom CyberArk credential resolver still uses the second, custom external vault, and no additional custom external vaults can be used.

To configure External Credential Storage, complete the following tasks in order.

Create a JAR file to resolve credentials

Create a JAR file to resolve credential identifiers sent from the MID Server into actual credentials from the repository.

Before you begin
Role required: agent_admin or admin

Make sure to include all the credential elements that the instance expects, such as the private key.

To create a JAR file to resolve credentials:

Procedure

Use either the templates provided on the ServiceNow github or the sample Java file.

Caution: These samples are intended as a template only. Do **NOT** use this code in production without modifying it for your environment.

a. Download the open source JAR files with instructions from the ServiceNow github:

- HashiCorp External Credential Resolver
- CyberArk External Credential Resolver

b. Use the following sample Java file as a template and modify it to suit your environment:

Example

```
package com.snc.discovery;

import java.util.*;
import java.io.*;

/**
 * Basic implementation of a CredentialResolver that uses a
 properties file.
 */

public class CredentialResolver {

 private static String ENV_VAR = "CREDENTIAL_RESOLVER_FILE";
 private static String DEFAULT_PROP_FILE_PATH =
 "C:\\dummycredentials.properties";

 // These are the permissible names of arguments passed INTO
 the resolve()
 // method.
```

```
// the string identifier as configured on the ServiceNow
instance...
public static final String ARG_ID = "id";

// a dotted-form string IPv4 address (like "10.22.231.12") of
the target
// system...
public static final String ARG_IP = "ip";

// the string type (ssh, snmp, etc.) of credential as
configured on the
// instance...
public static final String ARG_TYPE = "type";

// the string MID server making the request, as configured on
the
// instance...
public static final String ARG_MID = "mid";

// These are the permissible names of values returned FROM the
resolve()
// method.

// the string user name for the credential, if needed...
public static final String VAL_USER = "user";

// the string password for the credential, if needed...
public static final String VAL_PSWD = "pswd";

// the string pass phrase for the credential if needed:
public static final String VAL_PASSPHRASE = "passphrase";

// the string private key for the credential, if needed...
public static final String VAL_PKEY = "pkey";

// the string authentication protocol for the credential, if
needed...
public static final String VAL_AUTHPROTO = "authprotocol";

// the string authentication key for the credential, if
needed...
public static final String VAL_AUTHKEY = "authkey";

// the string privacy protocol for the credential, if
needed...
public static final String VAL_PRIVPROTO = "privprotocol";

// the string privacy key for the credential, if needed...
public static final String VAL_PRIVKEY = "privkey";


private Properties fProps;

public CredentialResolver() {
}

private void loadProps() {
```

```
  if(fProps == null)
   fProps = new Properties();

  try {
   String propFilePath = System.getenv(ENV_VAR);
   if(propFilePath == null) {
    System.err.println("Environment var "+ENV_VAR+" not found.
 Using default file: "+DEFAULT_PROP_FILE_PATH);
    propFilePath = DEFAULT_PROP_FILE_PATH;
   }

   File propFile = new File(propFilePath);
   if(!propFile.exists() || !propFile.canRead()) {
    System.err.println("Can't open
 "+propFile.getAbsolutePath());
   }
   else {
    InputStream propsIn = new FileInputStream(propFile);
    fProps.load(propsIn);
   }

 //fProps.load(CredentialResolver.class.getClassLoader().getRes
ourceAsStream("dummycredentials.properties"));
  } catch (IOException e) {
   System.err.println("Problem loading credentials file:");
   e.printStackTrace();
  }
 }

 /**
  * Resolve a credential.
  */
 public Map resolve(Map args) {
  loadProps();
  String id = (String) args.get(ARG_ID);
  String type = (String) args.get(ARG_TYPE);
  String keyPrefix = id+"."+type+".";

  if(id.equalsIgnoreCase("misbehave"))
   throw new RuntimeException("I've been a baaaaaaaaad
 CredentialResolver!");

  // the resolved credential is returned in a HashMap...
  Map result = new HashMap();
  result.put(VAL_USER, fProps.get(keyPrefix + VAL_USER));
  result.put(VAL_PSWD, fProps.get(keyPrefix + VAL_PSWD));
  result.put(VAL_PKEY, fProps.get(keyPrefix + VAL_PKEY));
  result.put(VAL_PASSPHRASE, fProps.get(keyPrefix +
 VAL_PASSPHRASE));
  result.put(VAL_AUTHPROTO, fProps.get(keyPrefix +
 VAL_AUTHPROTO));
  result.put(VAL_AUTHKEY, fProps.get(keyPrefix + VAL_AUTHKEY));
  result.put(VAL_PRIVPROTO, fProps.get(keyPrefix +
 VAL_PRIVPROTO));
  result.put(VAL_PRIVKEY, fProps.get(keyPrefix + VAL_PRIVKEY));
```

```
 System.err.println("Error while resolving credential
id/type["+id+"/"+type+"]");

 return result;
}


/**
 * Return the API version supported by this class.
 */
public String getVersion() {
 return "1.0";
}

public static void main(String[] args) {
 CredentialResolver obj = new CredentialResolver();
 obj.loadProps();

 System.err.println("I spy the following credentials: ");
 for(Object key: obj.fProps.keySet()) {
  System.err.println(key+": "+obj.fProps.get(key));
 }

 }
}
```

Import a JAR file to resolve credentials

Import a JAR file created to resolve credential identifiers sent from the MID Server into actual credentials from the repository.

Before you begin

Role required: agent_admin or admin

After you create the JAR file, import it into the instance, where it becomes accessible to the MID Server.

Procedure

1. After creating the JAR and properties files, copy the properties file to the MID Server.
2. Navigate to **MID Server** > **JAR Files**.
3. Click **New**.
4. Complete the following fields:

Field	Description
Name	A unique and descriptive name for identifying the file in the instance.
Version	A version number for the file, if one is available.
Source	Location of the JAR file for reference purposes. Source information is not used by the system.
Description	Short description of the JAR file and its purpose in the instance.

5. Click the paper clip icon in the banner and attach the JAR file to the record.

Example

Attaching a JAR file

6. Click **Submit**.
7. Restart the MID Server service.
   The platform makes the JAR file available to any MID Server configured to communicate with the instance.

Configure the credential identifier

Configure the credential identifier in the instance.

Before you begin

Role required: admin
Verify the following items:

- The External Credential Storage plugin must be active.
- The Enable External Credential Storage Discovery property is enabled.

Procedure

1. Navigate to **All** > **Discovery** > **Credentials** or **Orchestration** > **Credentials**.
2. Click **New**.
3. Select a credential type.
4. Select the **External credential store** check box.
   The **User name** and **Password** fields disappear, and the **Credential ID** field and **Credential storage vault** menu appear.
5. From the **Credential storage vault** menu, select either None, the CyberArk vault, or a custom external credential storage vault.

   Note:

   If CyberArk vault is selected, the **Lookup key** menu appears with four lookup key choices: Credential ID, IP Address, FQDN, All of the above. Selecting All of the above can degrade performance because it requires accessing the vault multiple times.

   a. To use a custom external credential storage vault, navigate to Vault Configurations [vault_configuration.list] in the instance.

   b. Create a new record using a name associated an imported JAR file for a custom credential resolver.

   See the procedures Create a JAR file to resolve credentials and Import a JAR file to resolve credentials for information on creating a custom external credential storage vault.

6. Complete the Credentials form using the fields from the following table.

Field	Description
Name	Enter a unique and descriptive name for this credential.

Field	Description
Active	Enable or disable these credentials for use.
Credential ID	Enter the unique key configured for external credentials in the JAR file uploaded to the MID Server for an external credential system. This is the ID passed to the Java class in the parameter map: `public static final String ARG_ID = "id";` The MID Server uses this identifier to resolve the actual credentials on the repository. **Note:** This field is only visible when the **External credential store** check box is selected.
Tag	Allow workflow creators to assign individual credentials to any activity in an Orchestration workflow or assign different credentials to each occurrence of the same activity type in an Orchestration workflow.
External credential store	Select this check box to use an external credential storage system. When you select this option the **User name** and **Password** fields are replaced with the **Credential ID** field. External credential storage is only available when the External Credential Storage plugin in activated.
Credential storage vault	Choose the external credential storage vault from a list of available vaults. The menu is composed of records from the Vault Configurations [vault_configuration.list]. New records can be added and use names associated with custom credential resolver JAR files. See the procedures Create a JAR file to resolve credentials and Import a JAR file to resolve credentials for information on creating a custom external credential storage vault.
Applies to	Select whether to apply these credentials to **All MID servers** in your network, or to one or more **Specific MID servers**. Specify the MID Servers that should use these credentials in the **MID servers** field.
MID servers	Select one or more MID Servers from the list of available MID Servers. The credentials configured in this record are available to the MID Servers in this list. This field is available only when you select **Specific MID servers** from the **Applies to** field.
Order	Enter the order (sequence) in which the platform tries this credential as it attempts to log onto devices. The smaller the number, the higher in the list this credential appears. Establish credential order when using large numbers of credentials or when security locks out users after three failed login attempts. If all the credentials have the same order number (or none), Discovery or Orchestration tries the credentials in a random order.

7. Click **Submit**.

Configure the credential identifier for AWS

Configure your instance to obtain credentials from a remote repository.

Before you begin

Role required: cloud_admin

Verify that these plugins are active, and the MID Server has been installed:

- Discovery [com.snc.discovery]
- Cloud Provisioning and Governance [com.snc.cloud.mgmt]
- External Credential Storage [com.snc.discovery.external_credentials]

About this task

These procedures assume that you already have an external repository configured with the credentials you want to protect. The credential identifier configured in the ServiceNow instance must be mapped to the actual credential in the repository through the JAR file.

Procedure

1. Navigate to **All** > **Discovery** > **Credentials**.
2. Select a credential that your external credential storage provider supports.
3. Complete the form, using the fields from the table.

Field	Description
Name	A unique and descriptive name for this credential. For example, Amazon Web Services.
Active	Check box to enable or disable the credential.
Credential ID	Enter the Name under which this credential is stored in the external credential storage provider.
MID Servers	Select one or more MID Servers that can use these credentials.
External credential store	Select this check box to use an external credential storage system. When external storage is enabled, the **Credential ID** field appears. If this check box is not visible, click the menu icon in the header bar and select **View** > **External Storage** from the context menu.
Credential storage vault	Select **CyberArk**.

4. Click **Submit**.

CyberArk credential storage integration

The MID Server integration with the CyberArk vault enables ServiceNow® Orchestration, ServiceNow® Discovery, and ServiceNow® Service Mapping to run without storing any credentials on the instance.

Introduction to CyberArk

CyberArk Application Identity Management (AIM) product uses the Privileged Account Security solution to eliminate the need to store application passwords embedded in applications, scripts or configuration files, and allows these highly sensitive passwords to be centrally stored, logged, and managed within the CyberArk vault. This approach enables organizations to comply with internal and regulatory requirements of periodic password replacement and to monitor activities associated with all types of privileged identities, whether on-premise or in the cloud.

The instance maintains a unique identifier for each credential, the credential type (such as SSH, SNMP, or Windows), and any credential affinities. The MID Server obtains the credential identifier, credential type, and IP address from the instance, and then uses the CyberArk vault

to resolve these elements into a usable credential. The credential resolver can also look up the hostname, fqdn, and use reverse DNS lookup to get fqdn.

The CyberArk integration requires the ServiceNow® External Credential Storage plugin, which is available in **System Definitions** > **Plugins**. The MID Server and CyberArk AIM/API client must be installed on the same machine. CyberArk Application Access Manager (AAM) Credential Providers version 12.0.1 and later is supported.

Installed with CyberArk

- **Business rule**: The External Credential Storage business rule performs the following tasks when an administrator makes any change to the external credential storage property:
  - Changes the view for the Credentials record list and form to the External Storage view. This view enables users to see the Credential ID column in the list.
  - Instructs the MID Server to refresh its non-external credentials cache in preparation for a change in the way that credentials are obtained.
- **System property**: A property called Enable External Credential Storage [com.snc.use_external_credentials] enables or disables the External Credential Storage plugin after it is activated. This property is located in **Discovery Definition** > **Properties** and **Orchestration** > **MID Server Properties**, and is enabled when you activate the plugin.

  Note: If you disable external credential storage with the system property, the system automatically sets all the external credentials to inactive in the instance. If you re-enable the feature with this property, the system does not reset the external credential records to active. You must reactivate each credential record manually.

Supported credential types

The CyberArk integration supports these ServiceNow credential types:

- GCP
- Azure
- CIM
- JMS
- SNMP forum
- SNMPv3
- Basic Auth
- SSH Key Pair
- SSH Private Key (with key, pass phrase, and password)
- VMware
- Windows
- Applicative Credentials

Note: To use CyberArk integration with the GCP credential type, you must modify the external credential storage jar. For details see ServiceNow GCP Credential Resolver using CyberArk.

ServiceNow AI Platform features that use these network protocols also support the use of credentials stored on a CyberArk vault.

Credentials supported by network protocol

Network protocol	ServiceNow® Workflow Studio support	Orchestration support
SOAP	SOAP Step	Create a SOAP web service activity with basic authentication overrides
REST	REST Step	Create a REST web service activity with basic authentication overrides
JDBC	JDBC Step	JDBC activity
SSH	SSH Step	SSH activity
PowerShell	PowerShell Step	PowerShell activity
SFTP	SFTP Step	SFTP activity
JMS		JMS activity

Important: You cannot manage credentials stored on a CyberArk vault and a custom external credential storage system using the same MID Server. The MID Server and CyberArk AIM/API client must be installed on the same machine.

CyberArk architecture

CyberArk storage architecture

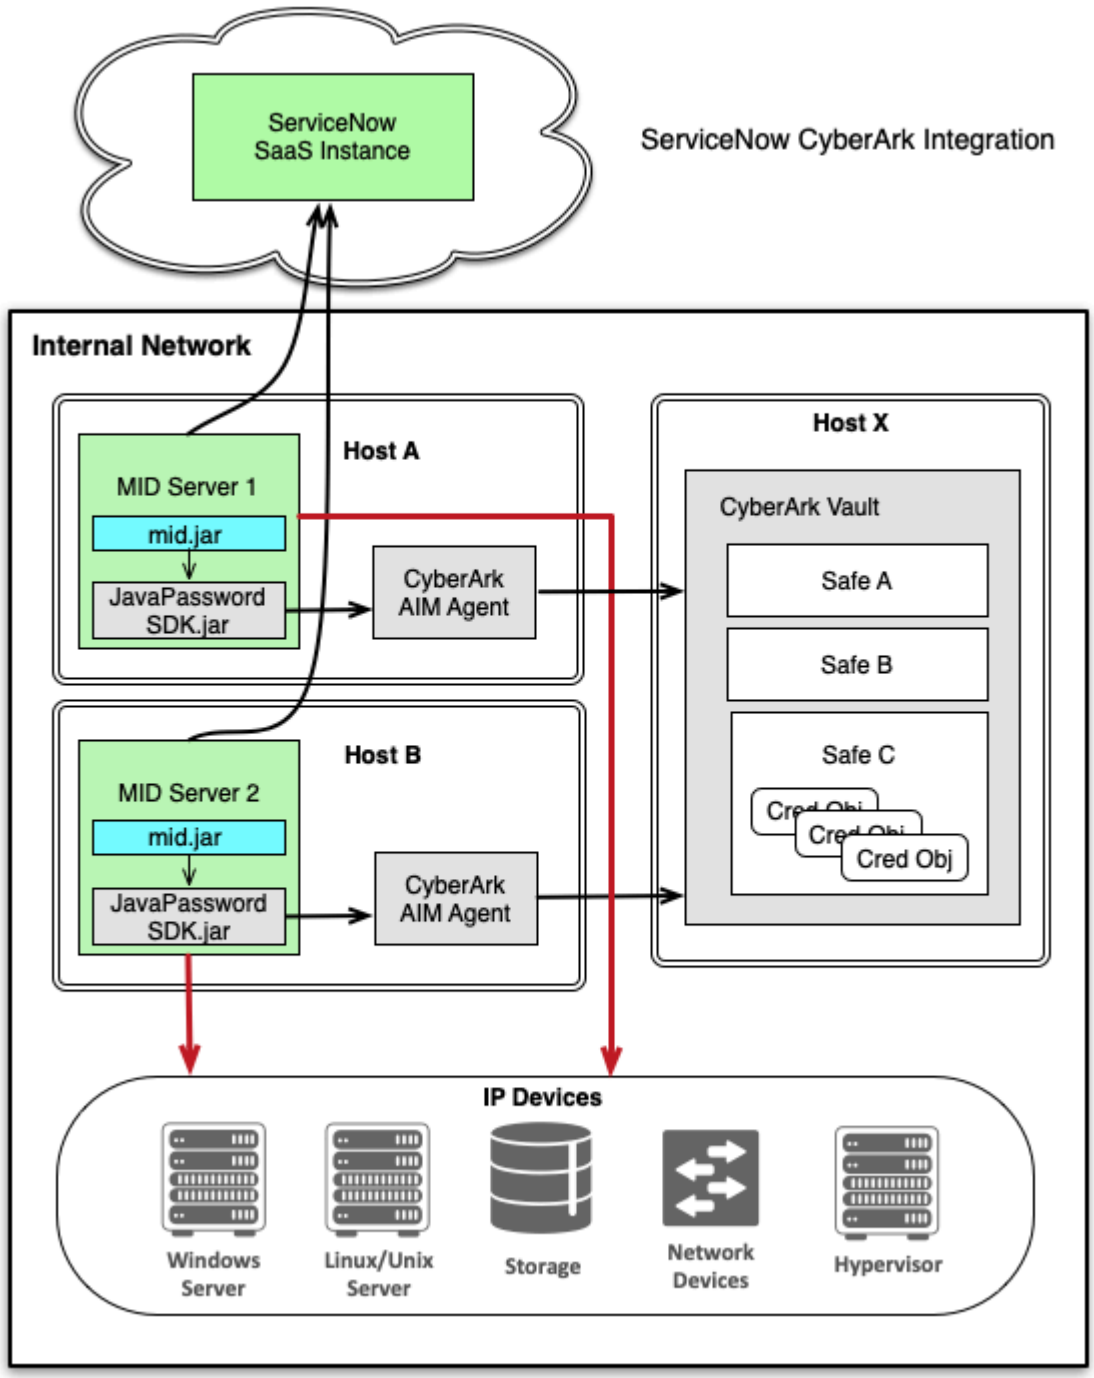

Note: CyberArk uses the base system `mid.jar` file for resolving credentials.

How the MID Server handles Windows accounts

Credential lookup initially attempts to match the specified credential ID to an existing value in the CyberArk vault **Name** field. If a match is found, that credential is returned. If no match is found, the credential lookup attempts to find a match using the IP address. If the IP address lookup

matches more than one credential, such as Windows and Tomcat on the same server, the lookup fails. To avoid this issue, set the *ext.cred.type_specifier* parameter in the MID Server config.xml file to **true** to force CyberArk to return credentials that match both the credential type and the IP address. For example, if an IP address is shared by both Windows and Tomcat, a credential type of Windows returns the Windows credential only.

Upgrade the CyberArk library

You can upgrade the CyberArk library if a secured configuration parameter is needed.

Check the following configuration parameter in the **config.xml**:
`<parameter name="mid.secure_config.provider" value="com.service_now.mid.services.config.CyberArkSecuredConfigProvider"/>`

Perform the following steps to perform the upgrade if a secured config parameter provider is configured.

1. Rename the CyberArk client version to `JavaPasswordSDK_MajorVersion_minorVersion_patchNum.jar`.
2. Create a new jar entry in the `ecc_agent` table where the rename jar can be attached. This new entry downloads to the MID Server. This step results in two jar (Passworsdk.jar and JavaPasswordSDK _12_X_X.jar).
3. Delete old ecc_agent entry from instance. This step deletes Passworsdk.jar from the MID Server, and the JavaPasswordSDK _12_X_X.jar remains in the system.

CyberArk integration configuration

These procedures include both CyberArk and ServiceNow configuration tasks, including references to the appropriate CyberArk documentation.

The credential identifier configured in the ServiceNow instance must be mapped to the credential name in the CyberArk vault. When looking up a credential, the MID Server finds the credential by matching the credential identifier to a name in vault, which must be unique. If the credential identifier is blank, then the MID Server finds the credential by IP address. To identify the credential by IP address, the system looks at the credential type to ensure that there is only one credential of that type at that address. An example of this might be when a Windows server and vCenter are both running on the same IP address. To support strict credential requirements like this in an SSH environment, a MID Server configuration parameter allows you to require that the credential type requested matches the type returned by CyberArk.

Note: The **Credential ID** field is the only field necessary to map your credentials to CyberArk, in all cases except for SNMPv3. The **Privacy Credential ID** field is optional and is required only when using SNMPv3 credentials and using a privacy protocol for the credentials. See Configure the CyberArk credential identifier for more details.

To configure your instance to obtain credentials from a CyberArk vault, complete these tasks in the order in which they appear below.

Configure the CyberArk vault and install the AIM API

Configure the CyberArk vault to allow MID Server access and install the CyberArk AIM API on the MID Server machine.

Before you begin
Role required: admin

Before starting this procedure, ensure that the External Credential Storage plugin is activated. CyberArk Application Access Manager (AAM) Credential Providers version 12.0.1 and later is supported.

Procedure

1. Configure the CyberArk vault with the application ID and authentication details that all MID Servers requesting credentials will use.
   For details, refer to the CyberArk Credential Provider and ASCP Implementation Guide.

   a. Ensure that CyberArk is configured to allow the MID Server to access the vault by creating an App-ID in CyberArk called *`ServiceNow_MID_Server`*.

   b. Make sure that every credential the MID Server needs is granted access to the *`ServiceNow_MID_Server`* App-ID.

   Note: You can override the default **ServiceNow_MID_Server** App-ID in the MID Server `config.xml` file using the ext.cred.app_id parameter. If you change the value in this parameter, make sure to configure a matching value in the vault.

2. Install the CyberArk Credential Provider, including the AIM API, on each machine that hosts a MID Server service that is used to access the credential store.

   Note: For information about installing AIM, refer to the CyberArk customer documentation.

3. Provision CyberArk accounts and set permissions for application access.
   For details, refer to the CyberArk Privileged Account Security Implementation Guide.

   a. In the CyberArk Password Safe, create the privileged accounts required by Discovery, Orchestration, or Service Mapping to access different devices and ensure that these accounts are members of the safes in which the necessary credentials are stored.

   b. Add the Credential Provider and application users as members of the Password Safes where the application passwords are stored.

Import the CyberArk JAR file

Import the CyberArk JavaPasswordSDK.jar file into the instance to make it accessible to the MID Server.

Before you begin

Role required: agent_admin or admin

Before starting this procedure, ensure that CyberArk is configured to allow the MID Server access to credentials. Ensure that the CyberArk AIM API is installed on each server hosting a MID Server that is used to access the vault.

Note: Two separate CyberArk integrations are not supported by the MID Server. The same version of the CyberArk AIM API must be installed on all MID Servers connected to the same instance.

About this task

Use this process even if the JavaPasswordSDK.jar file already exists on the MID Server.

Procedure

1. Navigate to **All** > **MID Server** > **JAR Files**.
2. Click **New**.
3. Complete the form using the fields in the table.

JAR File form fields

Field	Description
Name	Unique and descriptive name for identifying the file in the instance.
Version	Optional version number for the file, if one is available.
Source	Provider of the JAR file. Source information is not used by the system.
Description	Optional short description of the JAR file and its purpose in the instance.

4. Attach the JAR file to this record.
   The AIM `JavaPasswordSDK.jar` file comes with the AIM SDK installation files and is typically located on the MID Server in the AIM installation directory at `<install_dir>/CyberArk/ApplicationPasswordSdk`.

   Note: The attached AIM JavaPasswordSDK.jar must match the installed version of the CyberArk AIM API. Any mismatch in JavaPasswordSDK.jar files may lead to unpredictable behavior and potential functionality issues.

5. Click **Submit**.
6. Restart the MID Server service.
   The platform makes the JAR file available to any MID Server configured to communicate with the instance.

Configure the MID Server for CyberArk AIM

Configure the `config.xml` file to grant the MID Server access to the CyberArk vault with AIM API.

Before you begin

Role required: admin

Before starting this procedure, import the `JavaPasswordSDK.jar` file into the instance.

Procedure

Manually configure the MID Server Add a MID Server parameter file with these parameters.

This configuration cannot be done from the instance.

Required configuration parameters

Parameter	Value	Description
ext.cred.safe_folder	NameOfFolder	Folder to use for all credential lookups. For example, **root**.
ext.cred.use_cyberark	true	Boolean parameter indicating that this MID Server is integrated with CyberArk.

Optional configuration parameters

Parameter	Value	Description
ext.cred.safe_timeout	5 (sec)	Timeout of each credential lookup in the vault, specified in seconds.
ext.cred.safe_name	NameOfSafe	Default safe name used for all credential lookups. If parameters are in multiple safes, the credential ID may be specified in the format *<safeName>:<CredentialID>*. When configured like this, the **NameOfSafe** field is ignored. If all external credentials have their credential IDs specified in this format, then leave out the **NameOfSafe** field. **Note:** By default the separator character in this format is a colon. To assign any character you want as a separator, add this line to the `CredMap.properties` file: `safe.cred.split.string=<string>`.
ext.cred.app_id	ServiceNow_MID_Server	Specifies the App-ID used to grant permission to the MID Server to access the CyberArk vault. The default value, **ServiceNow_MID_Server**, must be defined in the CyberArk vault. You can use this parameter to override the default and specify your own App-ID. If you edit the App-ID in this parameter, make sure to configure CyberArk to match.
ext.cred.type_specifier	true	Forces an IP address lookup to return credentials that match both the CyberArk platform ID and the IP address. For example, if an IP address is shared by both Windows and Tomcat, a credential with a platform ID starting with **Win** returns the Windows credential only. When this parameter is set to true, CyberArk looks for platform IDs that begin with: • Win: Windows • Unix: SSH • VMWare: VMware
ext.cred.check_ssh_type	false	When set to true, requires that the type of SSH credential returned from CyberArk matches the type of credential requested. For example, if a normal SSH username/password credential is requested and only SSH keys are available, the credential lookup fails.

Configure the MID Server for CyberArk CCP

Configure the `config.xml` file to grant the MID Server access to the CyberArk CCP vault.

Before you begin

Role required: admin

About this task

Note:

In the Zurich family release, the instance needs an additional Update Set that can be downloaded from Enable CyberArk Central Credential Provider (CCP) Integration in Zurich Release [KB2682524]. Follow the provided installation steps. The Australia family release includes this script by default.

Procedure

Manually configure the MID Server Add a MID Server parameter file with these parameters.

This configuration cannot be done from the instance.

Required configuration parameters

Parameter	Value	Description
ext.cred.safe_folder	NameOfFolder	Folder to use for all credential lookups. For example, **root**.
ext.cred.use_cyberark	true	Boolean parameter indicating that this MID Server is integrated with CyberArk.
ext.cred.ccp_endpoint	CCPendpointURL	The CCP endpoint URL, which must use HTTPS. For example: `https:// /AIMWebService/api/Accounts`
ext.cred.cyberark.cert_path	/path/on/mid/agent/security/EXAMPLE-cyberark-client.pfx	Path to the certificate file.
ext.cred.cyberark.cert_password	example-password	Password for the certificate file.

Optional configuration parameters

Parameter	Value	Description
ext.cred.timeout	30	Timeout of each credential lookup in the vault, specified in seconds.
ext.cred.safe_name	NameOfSafe	Default safe name used for all credential lookups. If parameters are in multiple safes, the credential ID may be specified in the format *<safeName>:<CredentialID>*. When configured like this, the **NameOfSafe** field is ignored. If all external credentials have their credential IDs specified in this format, then leave out the **NameOfSafe** field.

Parameter	Value	Description
		Note: By default the separator character in this format is a colon. To assign any character you want as a separator, add this line to the `CredMap.properties` file: `safe.cred.split.string=<string>`.
ext.cred.app_id	ServiceNow_MID_Server	Specifies the App-ID used to grant permission to the MID Server to access the CyberArk vault. The default value, **ServiceNow_MID_Server**, must be defined in the CyberArk vault. You can use this parameter to override the default and specify your own App-ID. If you edit the App-ID in this parameter, make sure to configure CyberArk to match.
ext.cred.type_specifier	true	Forces an IP address lookup to return credentials that match both the CyberArk platform ID and the IP address. For example, if an IP address is shared by both Windows and Tomcat, a credential with a platform ID starting with **Win** returns the Windows credential only. When this parameter is set to true, CyberArk looks for platform IDs that begin with: • Win: Windows • Unix: SSH • VMWare: VMware
ext.cred.check_ssh_type	false	When set to true, requires that the type of SSH credential returned from CyberArk matches the type of credential requested. For example, if a normal SSH username/password credential is requested and only SSH keys are available, the credential lookup fails.
ext.cred.verify_ssl	true	The MID Server validates the CCP server certificate, verifying the server's identity. This setting is recommended for production environments. If set to false, the MID Server does not validate the CCP server certificate.
ext.cred.check_revocation	true	This parameter controls certificate revocation checking for the CCP server certificate chain. While true, enables CRL/OCSP checks.
ext.cred.snmpv2_community_property	Attribute Name	SNMPv2 is not natively supported in CyberArk. If your organization has created custom SNMPv2 credentials in which the community string does not appear in the password field of the credential, use

Parameter	Value	Description
		this property to map the attribute in the credential to the community string.

Configure CyberArk for SNMPv2 credentials

If your system uses SNMPv2, you can create a special file to map the attribute in a credential to the community string.

Before you begin

Role required: admin
Before starting this procedure, configure the MID Server to have access to the CyberArk vault.

About this task

Note: If the community string appears in the password field of the CyberArk credential, it is not necessary to perform this procedure.

SNMPv2 is not natively supported in CyberArk. If your organization has created custom SNMPv2 credentials in which the community string does not appear in the password field of the credential, use this procedure to map the attribute in the credential to the community string.

Procedure

1. For CyberArk CCP configurations:
   a. Open the `config.xml` file.
   b. Set the property `snmpv2_community_property` to the attribute name.
2. For CyberArk AIM API configurations:
   a. In a text editor, create a file called: `CredMap.properties`.
   b. Add the code: `SNMPv2.Community=attribute_name`
   c. Save the file to the `/agent` directory of your MID Server installation.

Result

On credential look-up, the MID Server attempts to find this attribute for the credential. If the attribute is not found, the MID Server then looks in the password field. If the password field is empty, the credential look-up fails.

Configure the CyberArk credential identifier

Create the unique key that CyberArk can use to identify specific credentials in the external repository.

Before you begin

Role required: admin
Before starting this procedure, ensure that the External Credential Storage plugin is activated, and the com.snc.use_external_credentials system property is set to true.

Procedure

1. Navigate to **All** > **Discovery** > **Credentials** or **Orchestration** > **Credentials**.
2. Click **New**.
3. From the list of credential types, select a type that supports CyberArk external storage.
4. Complete the form using the fields from your credential type.
5. Select the **External credential store** check box.

The **User name** and **Password** fields are replaced with the **Credential ID** field.

Note: If the check box is not visible, click the menu icon in the header bar and select **View** > **External Storage**.

6. In the **Credential ID** field, enter an expression using one of these formats:
   - If all your credentials are in the same safe, configure this safe name in the MID Server `config.xml` file using the ext.cred.safe_name parameter, and then specify the credential ID by name only, as **<credential ID>**.
   - To name credentials for a given platform that reside is a specific safe, define the credential ID as **<safe>:<credential ID>:<platform ID>**.
   - If your credentials are in multiple safes, specify the credential ID in this format: **<safe>:<credential ID>**.
   - If you want CyberArk to look up the credential by IP address, using an alternate safe, specify the credential ID in this format: **<safe>:**.
   - If you want CyberArk to look up the credential for an alternate platform ID in the same safe, use this format: **::<platform ID>**
   - If you want CyberArk to look up the credential in a configured safe by the IP address rather than the credential ID, leave this field blank. This is the best practice for handling installations in which each server has a unique credential. Without this type of lookup, you must create a credential ID record in your instance for every server in your environment.

   Note: The credential ID must match the value in the **Name** field for the CyberArk account. The **Credential ID** field has a limit of 160 characters.

7. If you are storing SNMPv3 credentials in CyberArk and are using the privacy protocol and privacy key, configure the ID as follows:

   a. In the **Credential Store Type** field, select **CyberArk**.
   The **Privacy Credential ID** field appears.

   b. Enter the **Name** of the CyberArk SNMPv3 privacy account in the **Privacy Credential ID** field.

   Note: Character limits for Credential ID and Privacy Credential ID vary by vault. ServiceNow supports up to 180 characters. CyberArk supports 160. Verify the limit in your vault.

8. Click **Submit**.

Configure AWS credentials on a CyberArk vault

Configure your CyberArk vault with the AWS credentials to be retrieved for use by your instance.

Before you begin
Role required: admin

About this task
Store the credentials as an Account on the CyberArk vault. When you configure access to the vault on your instance, the name you give to the Account must also be used as the credential ID.

Note: The procedure that follows references CyberArk Password Vault v14.2.1. If you are using a different version, set up configuration as per official CyberArk Password Vault documentation.

Procedure

1. In CyberArk, go to **Accounts** > **Accounts & Requests** > **Accounts View** > **Add Account**.
2. Select system type: **Cloud Service**.
3. Assign to platform: **Amazon Web Services - AWS - Access Keys**.
4. Store in Safe: Select a safe from the list.
5. Define properties: Enter the following information:

CyberArk credentials

Field	Value
AWS IAM Username	Enter the AWS Access Key, as provided by AWS.
AWS Access Key Secret (optional)	Enter the AWS Secret Access Key, as provided by AWS.
Customize account name	Toggle slider to enter a custom name for this key.
AWS Access Key ID	Enter the AWS Access Key again, as provided by AWS.
AWS Account ID Number	Enter the 12-digit AWS Account number.
AWS Account Alias Name (optional)	Enter an alias name for the account.

6. Choose **Add**.

What to do next
If you have not done so already, create a credential identifier on your instance to configure access to the CyberArk vault. For more details, see Configure access to external credential storage for AWS.

Configure AWS Credentials on a CyberArk Vault using Classic UI

Configure your CyberArk vault with the AWS credentials to be retrieved for use by your instance.

Before you begin
Role required: admin

About this task
CyberArk Password Vault v14.2.1 includes both Accounts View and Accounts View (Classic UI) configuration options to store the credentials as an Account. When you configure access to the vault on your instance, the name you give to the Account must also be used as the credential ID.

Note: The procedure that follows references CyberArk Password Vault v14.2.1. If you are using a different version, set up configuration as per official CyberArk Password Vault documentation.

Procedure

1. In CyberArk, go to **Accounts** > **Accounts & Requests** > **Accounts View (Classic UI)** > **Add Account**.
2. Enter the following information:

CyberArk credentials

Field	Value
Device Type	Select **Cloud Service**.
Platform Name	Select **Amazon Web Services - AWS - Access Keys**.
AWS IAM Username	Enter the AWS Access Key, as provided by AWS.
AWS Access Key ID	Enter the AWS Access Key again, as provided by AWS.
AWS Account ID Number	Enter the 12-digit AWS Account number.
AWS Access Key Secret	Enter the AWS Secret Access Key again, as provided by AWS.
Name	Enter a custom name for this key.

3. Choose **Save**.

What to do next
If you have not done so already, create a credential identifier on your instance to configure access to the CyberArk vault. For more details, see Configure access to external credential storage for AWS.

Configure Azure credentials on a CyberArk vault

Configure your CyberArk vault with the Azure credentials to be retrieved for use by your instance.

Before you begin
Role required: admin

About this task
To store an Azure credential, first create an Azure credential template in the CyberArk vault. This process only needs to be completed once for the vault.

Procedure
1. Log in to CyberArk in Administration mode.
2. Navigate to the **Administration** tab.
3. In **System Configuration**, edit **Platform Management**.
4. Navigate to **Cloud Provider Template** and duplicate it.
5. Edit the template for Azure credentials.
6. Add the following two properties:
   - *Name* as *Username* and *Display Name* as *Client ID*
   - *Name* as *Address* and *DisplayName* as *Tenant ID*
7. Apply the changes.
8. Navigate to the **Account** section and select **Add account**.

9. Select **Safe**.
10. Set the **Device Type** to **Cloud Service**.
11. Select the Azure template that was previously edited.
12. Fill in the information in for the **Client ID**, **Tenant ID**, and **Password** fields.
13. Select **Save**.

OAuth 2.0 authentication via MID Server using external credential storage

Store OAuth 2.0 credentials-client ID and client secret-in the CyberArk vault instead of the ServiceNow instance. The MID Server gets the credentials from the CyberArk vault, when required to get the OAuth token. The token is stored in the MID Server and refreshed automatically upon expiry.

The CyberArk Application Identity Management (AIM) product uses the Privileged Account Security solution to eliminate the need to store application passwords embedded in applications, scripts or configuration files, and allows these highly sensitive passwords to be centrally stored, logged, and managed within the CyberArk vault. You can configure the CyberArk vault to store OAuth 2.0 credentials rather than directly in a ServiceNow credentials record. To know more about CyberArk, see CyberArk credential storage integration.

Architecture of OAuth 2.0 authentication of MID Server request

The architecture has two parts: ServiceNow instance and the environment where the Application Identity Manager (AIM) client and the MID Server are configured. Examples of environment are the cloud or a customer environment.

The MID Server and the Application Identity Manager (AIM) client must be configured on the same environment and the Application Identity Manager (AIM) must be configured to interact with the CyberArk external vault. The CyberArk external vault could be hosted in the same environment as that of MID Server and the Application Identity Manager (AIM) or a different environment.

The ServiceNow instance maintains credential identifiers that map to specific OAuth 2.0 credentials stored in the CyberArk vault. Before sending an OAuth token request, the MID Server obtains the credential identifier from the instance, and then uses a customer-provided JAR file to send the identifier to the AIM client. The AIM client sends the request to CyberArk vault. The CyberArk vault sends the OAuth 2.0 credentials back to the MID server through the AIM client. After receiving the OAuth 2.0 credentials, the MID server sends the OAuth token request to the third-party authorization server. The token request comprises information such as client and client secret that CyberArk stores, and OAuth scope, and token URL that the instance stores. After the authorization server issues the OAuth token, the MID server stores it in its cache memory.

Note: This feature supports the Client Credentials Grant Type.

The image shows the MID server request authentication process.

Note: It's assumed that the third-party authorization server and the CyberArk vault are hosted in the customer network.

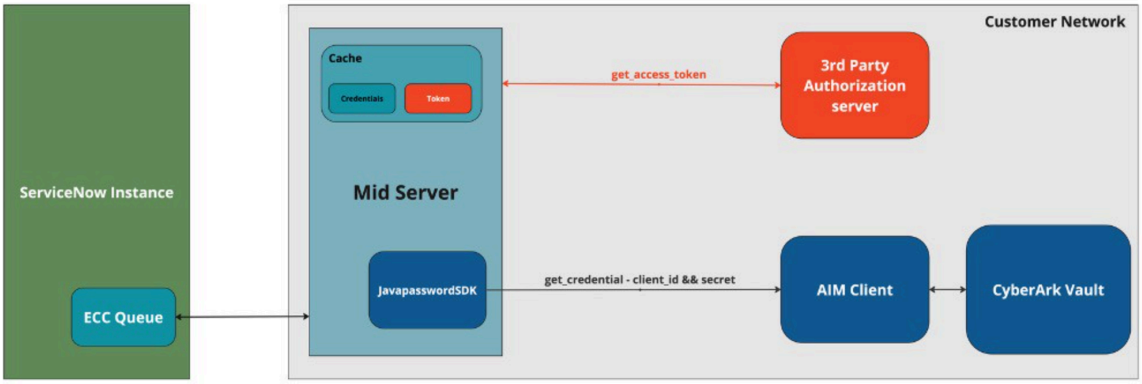

Configure a JAR file and credential identifiers

Configure a JAR file and credential identifiers so that the JAR file resolves the credential identifiers into actual credentials from the CyberArk external vault. The process enables the MID Server to get and include OAuth 2.0 credentials in the OAuth token request.

Before you begin

Role required: agent_admin or admin

The CyberArk external vault stores sensitive credentials and the ServiceNow instance stores credential identifiers that map to the specific credential names that the vault stores. The JAR file enables the mapping of a credential identifier to a specific credential name in the external vault. The mapping enables the MID Server to acquire and include the required credential in the OAuth token request. You must configure the credential identifiers and the JAR file in the ServiceNow instance and the MID Server respectively.

Procedure

1. Import the CyberArk JAR file.
2. Configure a connection to send OAuth request via the MID Server using external vault.

Configure CyberArk

Configure the CyberArk vault to store OAuth 2.0 credentials and respond to requests for OAuth 2.0 credentials from the MID Server.

Before you begin

Confirm that you have subscribed to the ServiceNow IntegrationHub Standard Pack Installer. For more information, see https://www.servicenow.com/content/dam/servicenow-assets/public/en-us/doc-type/legal/snc-addendum-integrationhub.pdf.

Role required: agent_admin or admin

Procedure

1. Configure the CyberArk vault and install the AIM API.
2. Import the CyberArk JAR file.
3. Configure OAuth 2.0 credentials on CyberArk.
4. Configure a connection to send OAuth request via the MID Server using external vault.

Configure OAuth 2.0 credentials on CyberArk

Configure your CyberArk vault with OAuth 2.0 credentials that the ServiceNow instance requests.

Before you begin
Role required: Admin

About this task
To store OAuth 2.0 credentials, first create an OAuth 2.0 credential template in the CyberArk vault. This process must only be completed once for the vault.

Procedure
1. Do the following steps, if you're configuring the OAuth 2.0 credentials for the first time.

   a. Log in to CyberArk in Administration mode.

   b. Navigate to the **Administration** tab.

   c. Select **Platform Management**.
   The Platform Management page displays the platforms under the Targets tab.

   d. Expand a platform type.

   e. Select the settings icon () that corresponds to a platform template and select **Duplicate**.

   f. In the Duplicate Platform window, enter a name for the template and select **Create**.

   Note: Note the system type under which you had created a duplicate template. For example, Cloud Service is a system type.

   g. Select the settings icon () that corresponds to the duplicate platform template that you created and select **Edit**.

   h. Add the property.
   Name as Username and Display Name as Client ID.

   i. Rename the other properties.

   Tip: Choose property names relevant to OAuth 2.0.

   j. Select **Apply**.
   The updates are applied.

   k. Select **OK**.

   l. On the left panel, navigate to Accounts.

   m. Select **Add Account**.

   n. Under the System Type heading, select the system type under which you had created the platform template.
   The System Type shows the platform templates.

   o. Under the Select Platform heading, select the platform template that you had created.

p. Select Safe.
Save the Safe value for later use.

q. Fill in the information in for the Client ID, Password (provide Client Secret as value), and other fields as required.

r. Select **Add**.
The account is added.

s. On the Account View page, open an account.

t. Select the **Details** tab.

u. Copy the value stored under the Account name section and save it for later use.
The value that you copied and stored is used as the credential identifier in the ServiceNow instance.

2. Do the following steps, if you're not configuring the OAuth 2.0 credentials for the first time.

a. Log in to CyberArk in Administration mode.

b. Navigate to the **Account** section and select **Add account**.

c. Select the system type from **System Type** under which the template for storing OAuth 2.0 exists.

d. Select the required platform template that you had created and select **Safe**.
Save the Safe value for later use.

e. Fill in the information in for the Client ID, Password (provide Client Secret as value), and other fields as required.

f. Select **Add**.
The account is added.

g. On the Account View page, open an account.

h. Select the **Details** tab.

i. Copy the value stored under the Account name section and save it for later use.
The value that you copied and stored is used as the credential identifier in the ServiceNow instance.

Configure a connection to send OAuth request via the MID Server using external vault

Configure a connection to send requests for OAuth 2.0 tokens to a third-party auth server via the MID Server. The MID Server adds OAuth 2.0 credentials (Client ID and Client Secret) from the CyberArk external vault, OAuth scope, and token URL from the instance to the request and sends them to the third-party auth server.

Before you begin
Verify the following items.

- Confirm that you have subscribed to the ServiceNow IntegrationHub Standard Pack Installer. For more information, see https://www.servicenow.com/content/dam/servicenow-assets/public/en-us/doc-type/legal/snc-addendum-integrationhub.pdf.
- You have created a JAR file that enables the MID Server to resolve credentials with the CyberArk external vault. See Create a JAR file to resolve credentials.
- You have imported the JAR file to the MID Server.
- The MID Server is set up for CyberArk. See Configure the MID Server for CyberArk AIM.
- OAuth 2.0 are configured on the CyberArk external storage. See Configure OAuth 2.0 credentials on CyberArk.

Role required: none

About this task

You can configure the connection and credential with a configuration template or manually. A configuration template defines the connection record form components and is reusable for other connection records. Manually, you can configure the connection for this purpose only. The topic shows both approaches.

Procedure

1. Configure the connection with a configuration template.

   a. Navigate to **All** > **IntegrationHub** > **Configuration Templates**.

   b. Select **New**.

   c. Select **HTTP Connection with OAuth Client Credentials grant type (External Storage)**.

   d. Update the form, as required.
   For example, you can provide the scope in the `oauth_entity_profile_scope` and `oauth_entity_scope` fields in the Default Data Template section. To learn to create a configuration template, see Create a configuration template.

   e. Select **Submit**.
   You have created the configuration template.

   f. Navigate to **All** > **IntegrationHub** > **Connection & Credential Aliases**.

   g. Update the form.

Connection and credential alias

Field	Description
Name	Name of the alias. An alias can only contain alpha, number, and underscore characters. During an upgrade, the tag in the credential record migrates to a Connection & Credential Alias. If the credential tag contains special characters other than alphabets, numbers, and underscores, it preserves the tag name after the upgrade.

Field	Description
	You can still use this migrated alias, but you cannot update the alias until you change the name to meet the naming restrictions.
Application	Unique identifier for the Connections & Credentials alias, based on the format `scope_name.alias_name`. ▪ If the scope is Global, the ID is the alias name. For example, if you create a Workday alias in the global scope, it sets the ID to `workday`. ▪ If you create a workday alias in the HR app scope, it sets the ID to `x_hr_app.workday`.
Parent Alias	Option to select the alias under which you're creating this connection and credential alias. The connection and credential alias that you're creating is a child alias. A child alias is listed under the **Child Aliases** tab on the parent connection and credential alias page. 
Type	Option to indicate the type of alias you're creating. Choose from the following options. ▪ Credential: Alias that contains a credential record. ▪ Connection and Credential: Alias that contains both connection and credential record. This option is selected by default. Confirm that Connection and Credential is selected.
Support Multiple Active Connections	Designator that indicates whether the alias supports multiple active connections. Add connections using the Connections table and associated them to the alias using the Connections related list.
Default Retry Policy	Retry policy for the alias. For more information, see Retry policy.
Configuration Template	Option to select the configuration template based on which you're creating the connection and credential alias. Select the template of the type **HTTP Connection with OAuth Client Credentials grant type (External Storage)** that you had created.

h. Select **Submit**.
You have created the connection and credential alias record.

i. Navigate to **All** > **IntegrationHub** > **Connections Dashboard**.

j. In the Search all connections field, enter the name of the connection and credential alias record that you created.

k. On the connection and credential alias record, select **View Details**.

l. Select **Configure**.

m. Fill the form.

Configure Connection form

Field	Description
Connection Name	Name of the connection. You can't update the name.
Connection URL	URL to the third-party server.
Use MID	Option to enable MID Server. **Important:** Confirm that the option is selected.
MID Selection	Option to specify MID Configuration for Connection. ▪ **Specific MID Server**: Option to indicate a manual selection of a specific MID Server. ▪ **Auto-Select MID-Server**: Option to indicate that you want the MID Server automatically selected. ▪ **Specific MID Cluster**: Option to indicate a manual selection of a MID cluster of MID Servers. **Important:** Confirm that the MID Server you select resides on the same machine as the CyberArk AIM Client that is configured to access the CyberArk vault. See Configure the CyberArk vault and install the AIM API.
MID Server	Option to manually specify a MID Server. This option appears if you select **Specific MID Server** in the **MID Selection** field.
MID Cluster	Option to manually specify a MID Cluster. This option appears if you select **Specific MID Cluster** in the **MID Selection** field.

Field	Description
External Credential Store	Option to use the CyberArk external credential storage. When you select the option, the MID Server gets the OAuth 2.0 credentials (Client ID and Client secret) from the external credential storage. **Important:** Confirm that the option is selected.
Credential ID	Identifier for the CyberArk account that holds the Client ID and Client Secret details. To view the procedure to get the credential identifier, see Configure OAuth 2.0 credentials on CyberArk. In the Credential ID field, enter an expression using one of these formats. ▪ If all your credentials are in the same safe, configure this safe name in the MID Server `config.xml` file using the ext.cred.safe_name parameter, and then specify the credential ID by name only, as **<credential ID>**. ▪ To name credentials for a given platform that reside is a specific safe, define the credential ID as **<safe>:<credential ID>:<platform ID>**. ▪ If your credentials are in multiple safes, specify the credential ID in this format: **<safe>:<credential ID>**.
OAuth Token URL	URL that specifies the endpoint to get the access token from the OAuth server.

n. Select **Configure Connection**.

2. Configure a connection manually.

a. Navigate to **All** > **System OAuth** > **Application Registry**.

b. Select **New**.

c. Select **Connect to a third party OAuth Provider using an external vault**.

d. Fill the form.

New application registry details

Field	Description
Name	Name to identify the application registry record. For example, enter: OAuth 2.0 token request via MID app.
Application	Option to specify the names of applications that can access this application registry. The field is read only.
Default Grant type	Default method to obtain access token by the client application from the OAuth server. The default and read-only grant type is Client Credentials.
Accessible from	Option to specify which applications can access this application registry.
Token URL	URL that specifies the endpoint to get the access token from the OAuth server.
Send Credentials	Option to specify the method of sending the OAuth 2.0 in the request body.
Comments	Enter related comments.

e. Under the **OAuth scope** column, create one or more OAuth scopes by doing the following steps.

i. Under the Name column, double-click the field and enter a name of the OAuth scope.

ii. Under the OAuth scope column, double-click the field and enter the scope.

f. Select **Submit**.
You have created an OAuth entity profile and the application registry.

g. Navigate to **All** > **IntegrationHub** > **Connections & Credentials** > **Credentials**.

h. Select **New**.

i. Select **OAuth 2.0 Credentials**.

j. Select **External Storage View**.

> **Important:** Select **External Storage View** only if the view is different from that of an OAuth 2.0 credentials external storage form.

k. Fill the form.

OAuth 2.0 Credentials

Field	Description
Name	Name of the credential record.
Applies to	Option to specify if the credential record is applicable for all MID Servers, or a specific MID Server. If specific, add the MID Servers.

Field	Description
	Choose from either of the following options. ▪ All MID Servers: The MID Server is automatically selected from a collection of MID Servers. ▪ Specific MID Servers: Option to specify one or more MID Servers.
MID Servers	Option to specify one or more MID Servers. **Note:** This field appears if you select **Specific MID Servers** in the Applies to field.
Active	Option to specify whether the credential record is available for use. By default, it's available.
OAuth Entity Profile	Option to specify an OAuth entity profile that the credential uses. Select the OAuth entity profile you had created above. See Configure an OAuth entity profile.
External credential store	Option to specify that the credentials are stored in an external storage and not in the ServiceNow instance. **Important:** Confirm that the option is selected.
Credential ID	Option to specify credential identifier for the CyberArk account which holds Client ID and Client Secret. To view the procedure to get the credential identifier, see Configure OAuth 2.0 credentials on CyberArk.
Credential storage vault	Option to specify the name of the external credential storage vault. Verify that you have selected CyberArk.

l. Select **Submit**.
You have created the credential record.

m. Create a connection and credential alias.
To know the steps, see Create a Connection & Credential alias.

n. Navigate to **All** > **IntegrationHub** > **Connections**.

o. Select **New**.

p. Select **HTTP(s) Connection**.

q. Fill the form.

Connection form

Field	Description
Name	Unique name of this HTTP(s) connection.
Active	Option to set the connection that you're creating active. The option is selected by default.
Credential	Select the credential record used to authorize the connection. Select the credential that you had created above.
Connection alias	Select the alias record to associate with this connection. Using an alias enables you to update the connection record without having to reconfigure any actions or activities that use the alias.
URL builder	Either manually enter the connection URL or use system to build the URL based on the inputs. Default is unchecked. If checked, the connection URL is calculated from the following fields: ▪ Mutual authentication — Check box if mutual authentication is used. ▪ Protocol — If mutual authentication is not used, enter protocol. The default is HTTPs. ▪ Protocol profile — If mutual authentication is used, enter protocol profile from sys_protocol_profile. ▪ Host ▪ Port ▪ Base path — Path of the connection string. **Note:** If mutual authentication is checked, connection URL is built: Protocol + :// + host:port +URL. If mutual authentication is unchecked, connection URL is built: Protocol profile + :// + host:port +URL
Connection URL	If URL builder is unchecked, enter the connection URL into this field. **Note:** If mutual authentication is checked, connection URL is built: Protocol + :// + host:port +URL. If mutual authentication is unchecked, connection URL is built: Protocol profile + :// + host:port +URL
Use MID server	Option to specify that you want to send OAuth token requests via a MID Server.

Field	Description
	Important: Confirm that the option is selected.
Connection timeout	Number of milliseconds the system waits for a successful host connection. If a successful connection does not occur during this time, the connection request times out. Leave this field empty to use the system default connection timeout value.
MID Selection	Option to specify one of the following options. ▪ Auto-Select MID Server: Selects from MID Servers based on MID Server criteria, regardless of whether they are members of a cluster. ▪ Specific MID Server: Manually select a MID Server. ▪ Specific MID Cluster: Automatic reassignment to another MID Server only selects from members of the specified cluster. **Important:** Confirm that the MID Server you select resides on the same machine as the CyberArk AIM Client that is configured to access the CyberArk vault. See Configure the CyberArk vault and install the AIM API.
Capabilities	Option to select one or more MID Server capabilities. Capabilities define the specific functions of a MID Server within an IP address range, allowing an application to select the most appropriate MID Server. Select the MID capabilities icon ( ) to select one or more capabilities. **Note:** ▪ Confirm that the MID Server you select resides on the same machine as the CyberArk AIM Client that is configured to access the CyberArk vault. See Configure the CyberArk vault and install the AIM API. ▪ This option appears if you select **Auto-Select MID Server** in the MID Selection field.

Field	Description
MID Server	Manually select a MID Server. This option appears if you select Specific MID Server from the MID Selection field.
MID Cluster	Manually select a MID cluster. This option appears if you select Specific MID Cluster from the MID Selection field.
MID Application	Option to specify a MID application or accept the default application choice. **Note:** This option appears if you select **Auto-Select MID Server** in the MID Selection field. By default, the **ALL** option is selected.

r. Select **Submit**.
You've created an HTTP(s) connection record.

s. Navigate to **All** > **Process Automation** > **IntegrationHub** > **Connections Dashboard**.

t. In the Search all connections field, enter the name of the connection record that you created.
The connection alias record appears.

OAuth 2.0 connection record is created

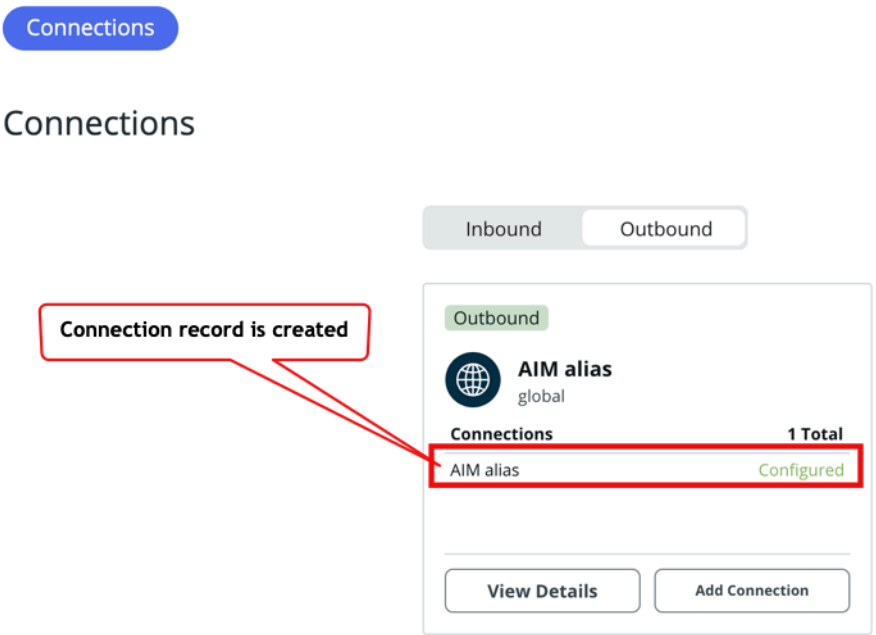

Authentication Algorithms

Verify the identity of the sender using Authentication Algorithms

Allow integration steps to authenticate with web services that require complex or non-standard connection or credential mechanisms. Associate authentication algorithms to credential and connection aliases to reduce or eliminate the need to manually configure integration steps.

You can use an authentication algorithm to generate custom authentication data for your integration steps. Integration steps can use this dynamic data to create any custom artifacts necessary to authenticate with the target web service. For example, a REST step could create an authentication header, query parameters, or a token.

Authentication algorithms support the following steps:

- Get Connection Info Step
- REST Step
- SOAP Step

For more information, see Integration steps

Authentication algorithm types

- **Amazon Signature Version 4**: This is a pre-built authentication algorithm to connect to Amazon Web Services.
- **Custom Authentication**: This is a template that developers can use to create their own authentication algorithms.

To know more, on how to configure the authentication algorithm, see Configure an authentication algorithm.

Scripts

Instance Authentication Scripts are on instance scripts part of *sys_script_include* table.

Instance Authentication Scripts

RequestAuthInternal	Read only script on instance, which supports generating AWS V4 signature or custom authentication that is sent with outbound request.
RequestAuthAWSV4Signer	Script extending RequestAuthInternal to implement signer to generate AWS V4 signature.
RequestAuthTwitterSigner	Script extending RequestAuthInternal to implement signer to generate Twitter signature using OAuth 1.0a.
RequestAuthSampleCustomSigner	Sample script extending RequestAuthInternal to understand how to write custom signer on instance.

MID Authentication Scripts are on MID scripts part of *ecc_agent_script_include* table.

MID Authentication Scripts

RequestAuthInternal	Read only script on MID, which supports generating AWS V4 signature or custom authentication that is sent with outbound request.
RequestAuthAWSV4MIDSigner	Script extending RequestAuthInternal to implement signer to generate AWS V4 signature.
RequestAuthTwitterSigner	Script extending RequestAuthInternal to implement signer to generate Twitter signature using OAuth 1.0a.
RequestAuthSampleMidCustomSigner	Sample script extending RequestAuthInternal to understand how to write custom signer on MID.

JavaScript API's

Following are the JavaScript API's for authentication algorithm.

- AuthCredential
- HttpRequestAuthedData
- HttpRequestData
- RequestAuthAPI

Configure an authentication algorithm

Configure an authentication algorithm so that you can sign outbound HTTP requests.

Before you begin

You should have a script include configured before you configure an authentication algorithm.
Role required: admin

Procedure

1. Navigate to **All** > **Credentials & Connections** > **Authentication Algorithms**, and click **New**.
2. On the form, fill in the fields.
   The database selection in the **Format** field determines which fields are available.

Authentication form

Field	Description
Name	Unique name of this algorithm.
Algorithm	Outbound request type.
Description	Description of what your algorithm does.
Application	Scope that your application runs in.
Instance Authentication Script	Script that you select from the Script Includes table.

Field	Description
MID Authentication Script	Script that you select from the MID Server Script Includes [Discovery view] table.

3. Click **Submit**.

Configure an Amazon Signature based Custom Algorithm

Generate the Amazon Signature based data needed to authenticate to a web service by running script.

Before you begin

- JavaScript knowledge
- REST knowledge
- Target web service API knowledge
- Connection, credential, and alias knowledge
- Role required: Developer

About this task

Use a connection and credential alias and Amazon Signature Version 4 based algorithm for authentication.

Procedure

1. Navigate to **All** > **Credentials & Connections** > **Authentication Algorithms**, and click **New**.
2. On the form, fill in the fields.
   The database selection in the **Format** field determines which fields are available.

Authentication form

Field	Description
Name	Unique name of this algorithm.
Algorithm	Outbound request type. Select **Amazon Signature Version 4**.
Description	Description of what your algorithm does.
Application	Scope that your application runs in.
Instance Authentication Script	Script that you select from the Script Includes table. In case of **Amazon Signature Version 4** algorithm, choose **RequestAuthAWSV4Signer**. The scripts available are as follows: ○ RequestAuthAWSV4Signer ○ RequestAuthInternal ○ RequestAuthSampleCustomSigner ○ RequestAuthTwitterSigner **Note:** To know more about the script click the information icon next to the field. The details of the script such as Name, API Name, Application, Accessible from, Script, and so on is displayed.

Field	Description
MID Authentication Script	Script that you select from the MID Server Script Includes [Discovery view] table. The scripts available are as follows: ○ RequestAuthAWSV4Signer ○ RequestAuthInternal ○ RequestAuthSampleCustomSigner ○ RequestAuthTwitterSigner

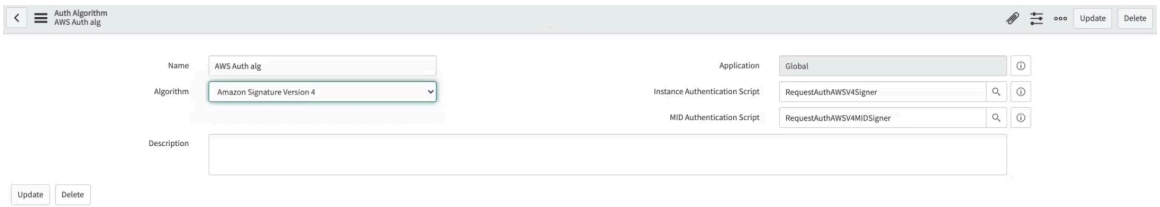

3. Click **Update**.
4. Navigate to **All** > **Connections & Credentials** > **Credentials**.
5. Click **New**.
6. Create AWS Credentials with Authentication Algorithm.
   In this case **AWS Auth alg**.
7. Specify the following:
   - Name
   - Active
   - Access Key ID
   - Secret Access Key
   - Credential alias
   - Authentication Algorithm

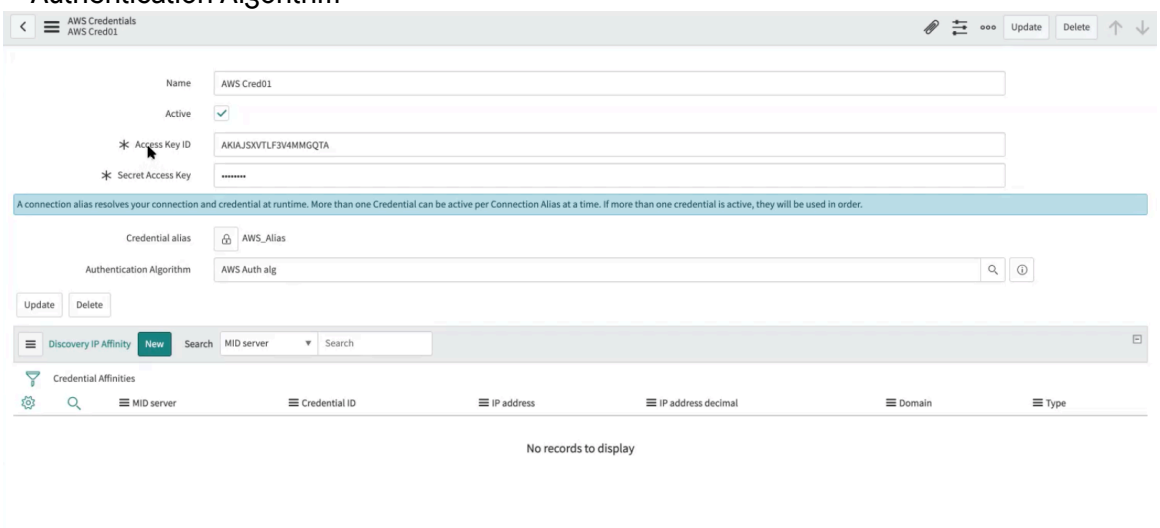

8. Click **Update**.

Result
Based on the selected scripts and authentication algorithm, the configured credentials (**Access Key ID** and **Secret Access Key**) or user's credentials (**Access Key ID**, **Secret Access Key**, and **Session Token**) generates a Amazon V4 signature that is sent as outbound request from ServiceNow to the provider (in this case AWS).

Example: REST step with AWS

Note: Amazon V4 signature based authentication can also be used from Script background.

Action: Get AWS Regions

Input REST step with AWS as follows:

- **Credentials Alias**: The alias that is created for AWS.
- **Base URL**: Base URL details from AWS.
- **HTTPS Method**: In this case it is GET method.
- **Query Parameters**: **Action** as **DescribeRegions**.

You can test the action, the associated regions are displayed. The response body is as follows:

Amazon V4 is defined with standard set of algorithm that supports authentication mechanism. This algorithm when used adds the signature as authorization header for authentication (HTTP request) using REST step.

Configure a custom authentication algorithm

Generate the custom data needed to authenticate to a web service by running script.

Before you begin

- JavaScript knowledge
- REST knowledge
- Target web service API knowledge
- Connection, credential, and alias knowledge
- Role required: Developer

About this task

Use a connection and credential alias and custom authentication based algorithm for authentication.

Procedure

1. Navigate to **All** > **Credentials & Connections** > **Authentication Algorithms**, and click **New**.
2. On the form, fill in the fields.
   The database selection in the **Format** field determines which fields are available.

Authentication form

Field	Description
Name	Unique name of this algorithm.
Algorithm	Outbound request type. Select **Custom Authentication**.
Description	Description of what your algorithm does.
Application	Scope that your application runs in.
Instance Authentication Script	Script that you select from the Script Includes table. The scripts available are as follows: ○ RequestAuthAWSV4Signer ○ RequestAuthInternal ○ RequestAuthSampleCustomSigner ○ RequestAuthTwitterSigner **Note:** ○ To know more about the script click the information icon next to the field. The details of the script such as Name, API Name, Application, Accessible from, Script, and so on is displayed. ○ In case of custom authentication with Twitter, you can choose **RequestAuthTwitterSigner**, since it uses an OAuth 1.0a method of authentication. This requires informations such as **API key and secret** and **Access token and secret** that can be used to create signatures to pass in an authorization header. For more information, see Authentication in Twitter.
MID Authentication Script	Script that you select from the MID Server Script Includes [Discovery view] table. The scripts available are as follows: ○ RequestAuthAWSV4Signer ○ RequestAuthInternal

Field	Description
	○ RequestAuthSampleCustomSigner ○ RequestAuthTwitterSigner

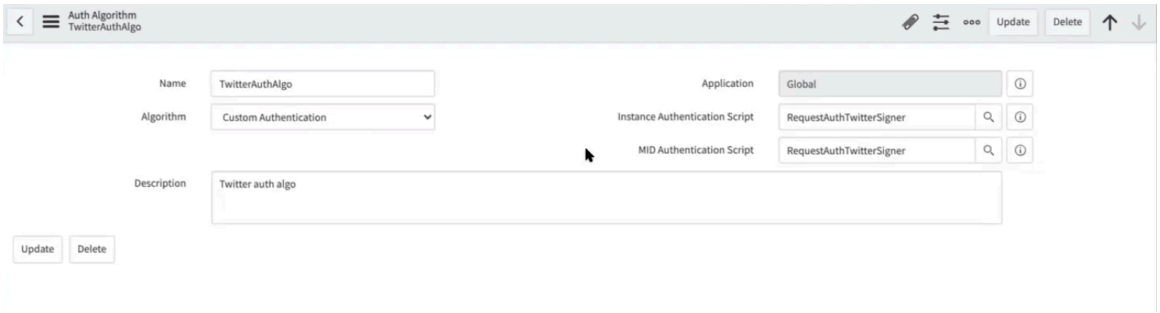

Based on the selected scripts and authentication algorithm, the configured credentials is sent as outbound request from ServiceNow to the provider.

3. Click **Update**.
4. Navigate to **All** > **Connections & Credentials** > **Credentials**.
5. Click **New**.
6. Create Twitter Credentials with Authentication Algorithm.
   In this case **TwitterAuthAlgo**.
7. Specify the fields:
   - Name
   - Active
   - Access token
   - Access token secret
   - Consumer key
   - Consumer secret
   - Credential alias
   - Authentication Algorithm

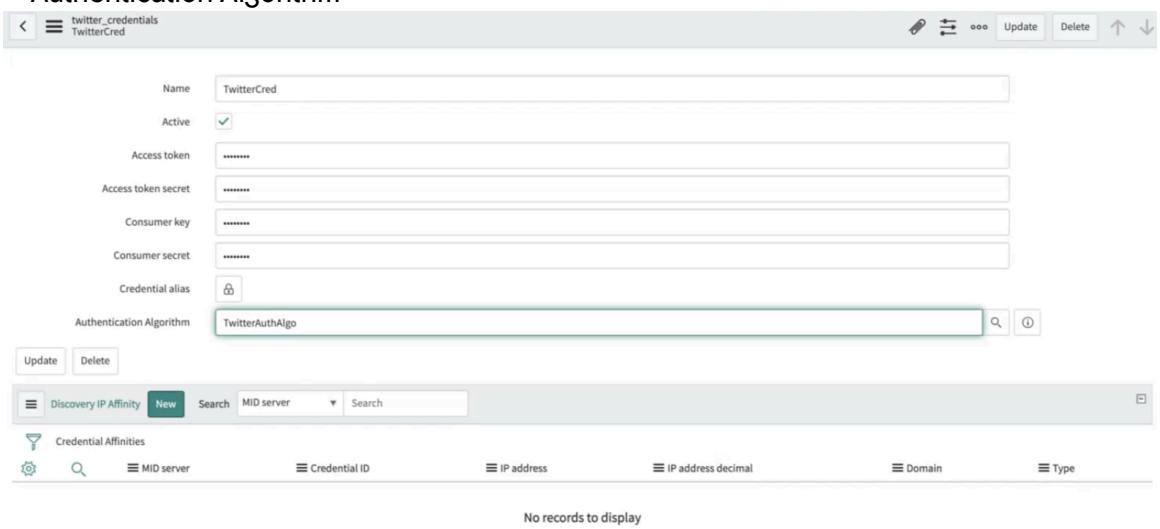

8. Click **Update**.

Example: REST step with Twitter

In case of Twitter, you must ensure the following spokes or credentials are available:

- Access token
- Access token secret
- Consumer key
- Consumer secret
- Authentication Algorithm

Action: TwitterAuthAlgo.

Input REST step with Twitter as follows:

- **Credentials Alias**: The alias that is created for Twitter.
- **Base URL**: Base URL details from Twitter.
- **HTTPS Method**: In this case it is POST method. Posting a tweet.
- **Query Parameters**: **Action** as **tweet**.

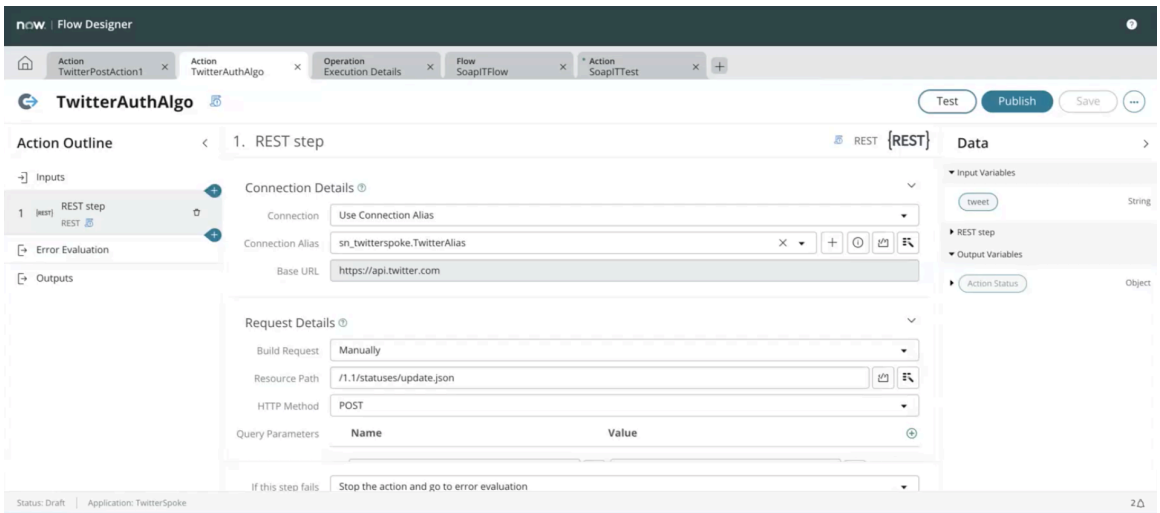

You can test the action. The tweet is posted on the Twitter page.

Check IP service affinity for Discovery and Orchestration

You can check the IP Services table for a list of IP addresses that are associated with a protocol.

Before you begin

Role required: admin

About this task

The IP Services table maps a port to a protocol. Several mappings are provided by default for commonly used port-protocol combinations, such as port 80 for HTTP, port 22 for SSH, and port 161 for SNMP.

A system property called *`glide.discovery.ip_service_affinity`* allows Discovery to remember the last port of the IP address that was discovered.

Important: You should not modify IP services unless your organization uses custom ports.

Procedure

1. Navigate to **All** > **Discovery Definition** > **IP Services**.
2. Filter the list to find the appropriate IP service.
3. Select the name of the service to go to that IP service page.
4. Select the **IP Service Affinities** tab for the list of IP addresses associated with that service.

ServiceNow® access control

The SNC Access Control plugin (com.snc.snc_access_control) enables you to control which Customer Service and Support employees can access your instance, and when.

Explore ServiceNow access control  Learn about the features of ServiceNow® access control.	Activate ServiceNow access control  Active ServiceNow® access control.
Configure ServiceNow access control  Understand how to configure ServiceNow® access control.	Audit Logging  Learn and review the ServiceNow® access control audit logging.

Explore ServiceNow® access control

The SNC Access Control plugin (com.snc.snc_access_control) enables you to control which Customer Service and Support employees can access your instance, and when.

When you first activate the plugin, Customer Service and Support employees cannot log into the instance. Any currently logged-in Customer Service and Support employees remain logged in. You create records in the SNC Access Control table that grant access to specific SNC employees or all employees.

The plugin prevents Customer Service and Support personnel from accessing the instances without your express permission. However, other authorized ServiceNow Operations personnel, in their capacity to support and manage the product, and verify usage are required to perform administrative actions on the underlying infrastructure. This infrastructure includes servers and

databases, among other infrastructure components that make up the SaaS solution. This access method is fully auditable and tracked.

This plugin enables you to restrict access to your instance without your express permission, so it may affect support service levels and the Availability SLA. Availability SLA is then measured from the time that Support staff personnel are granted access to your instance.

Login security

Security for authorized Customer Service and Support employee logins to instances employs encrypted tokens generated by a secure server. Only properly authenticated Customer Service and Support employees are granted access to an instance. Without the SNC Access Control plugin, the security server ensures that access rights are enforced on hi.service-now.com. When the plugin is enabled, the encrypted login tokens must match names in the plugin-provided access list, using the criteria defined in those records. This method of authentication enables you to determine precisely which Customer Service and Support employees may access their instances, and when these employees may do so.

The architecture chosen for this system has several features designed to enhance security for your instances:

Security server

The security server is a locked-down, Linux host that only ServiceNow security personnel can access. This server is the only system that has access to the critical private encryption key necessary to produce the login tokens. By using this compartmentalization (a standard security practice), the private key is protected, even in the unlikely event that an attacker compromises the HI instance.

Synthetic user

The facility on instances that enables authorized Customer Service and Support employees to log into their instance does not require an account to be provisioned on that instance. There is no user record provisioned, and no permanent or persisted credentials. Instead, a synthetic user is created for each Customer Service and Support employee logon. This user exists only in memory and provides no ongoing privileges. If the SNC Access Control plugin is enabled, you can deauthorize any Customer Service and Support employee at any time.

Tokens

The security tokens are specific to an instance and a particular Customer Service and Support employee. In addition, the mechanism that generates the tokens only works with actual Customer Service and Support employee logins to HI, not impersonated users. Once a security token is generated, only a specific Customer Service and Support employee can use it to log into an instance.

Time limit

Security tokens expire four hours after they are generated. This expiration limits the utility of hijacked tokens, which can only be used during this short window.

Logging

Logins by Customer Service and Support employees to instances are recorded as a login event.

- Every action taken by the logged-in Customer Service and Support employee is added to the transaction log in the database.
- It is also added to the instance log on the file system, which is inaccessible to most ServiceNow employees.
- Customer Service and Support employee logins and actions are readily identifiable, since the user names all end in @snc (like frodo.baggins@snc).

These actions provide you with easy-to-use, robust, and reliable security logging for non-employee access.

Security processing flow

When a Customer Service and Support employee wants to log into an instance, the security processing flow is as follows:

1. A Customer Service and Support technician requests a login for the instance through hi.service-now.com.
2. HI checks that the technician has the proper role authorizing access to instances.
3. If the user has the proper role, HI sends the request for access to the Security Server.
4. The Security Server verifies that the request came from the HI IP address, and evaluates the request (user, role, and IP address of the requester). If the request is valid, the Security Server approves it and constructs a token. This token contains the user name, roles, the instance ID, and the time (the start of the 4-hour token life span). Finally, the Security Server encrypts the token with the private encryption key.
5. The Security Server sends the encrypted token to HI.
6. Hi sends the token to the Support technician's browser.
7. The Support technician's browser initiates a login into the instance, using the special user name ending with `@snc`.
8. The instance uses the public key to decrypt the token. To verify the token, the instance matches it to the user name supplied in the previous step, the instance ID, and the authorized time window. If the SNC Access Control plugin is enabled, the instance verifies that the user is:
   - Listed
   - Active
   - Configured to access the instance in the current time window
9. If the user is authenticated, the instance creates a synthetic user in memory with the given roles. This user does not persist after the time limit expires, the user logs off, or the instance is restarted.

ServiceNow security access process flow

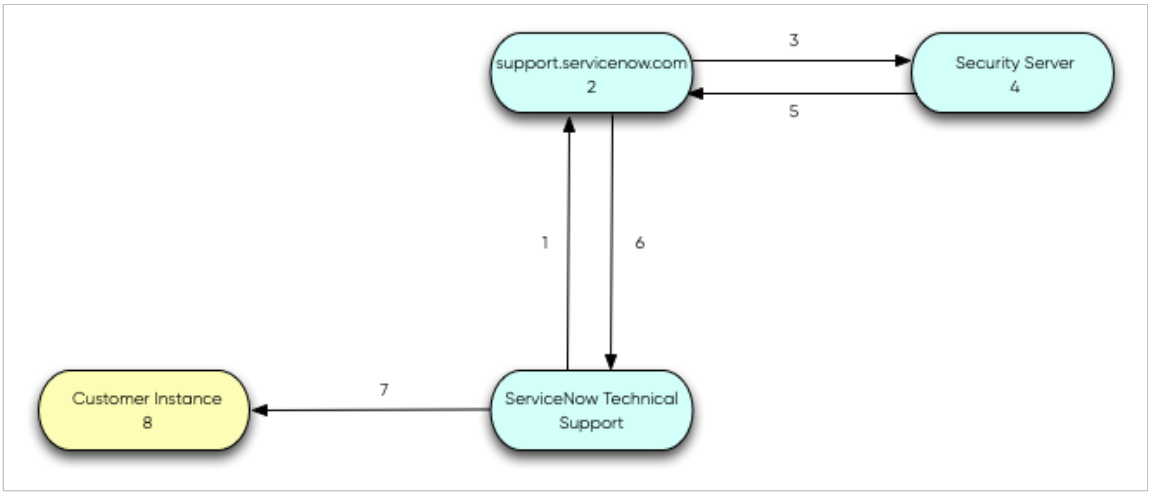

Audit logging

The following logging tracks logins and activity by Customer Service and Support employees:

- Event logs: The event logs show all Customer Service and Support logins to an instance.
- Transaction logs: The transaction logs show all activity on the instance, including any efforts to delete logs.

Note: To learn more about this plugin, see Enable SNC access control plugin [Updated in Security Center 1.3] in Instance Security Hardening Settings.

Activate ServiceNow® access control

You request activation of the SNC Access Control plugin (com.snc.snc_access_control).

Before you begin

Role required: admin

About this task

There are two ways to request a plugin:

- Access the Now Support Service Catalog directly by selecting **Automation Store** > **Service Catalog** > **Activate Plugin** on Now Support.
- Access the Now Support Service Catalog through the All Applications page on your instance by following these steps.

For additional details about requesting a plugin, see Requesting a Plugin from the Service Catalog [KB0751715] article in the Now Support Knowledge Base.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Select **Request plugin** to open the **Activate Plugin** form on Now Support.
3. On the **Activate Plugin** form, provide the following information.

Activate Plugin form

Field	Description
What is your target instance	Select the instance that you want to activate the plugin on.
Which plugin would you like to activate	Select the name of the plugin to activate. **Note:** If the system doesn't list the plugin you want or if you're activating the plugin on an OEM or on-premise instance, select the **Plugin I'm looking for is not listed** check box and then enter the name of the plugin.
Select Maintenance Date and Time	Select the date and time to activate the plugin.

Example

For example, see the following form to activate the Event Management plugin on an instance named SNC Instance.

4. Select **Submit**.
   After the maintenance window, the system installs the plugin on your instance. To confirm the installation, go to the Installed tab in the Application Manager.

Configure ServiceNow® access control

Configure an access control record to specify one or more Customer Service and Support employees who have permission to log in your instance.

Before you begin

Role required: admin

About this task

Note: The SNC Access Control (com.snc.snc_access_control) plugin prevents Customer Service and Support personnel from accessing the instances without your express permission. However, other authorized ServiceNow Operations personnel, in their capacity to support and manage the product, are required to perform administrative actions on the underlying infrastructure. This infrastructure includes servers and databases, among other infrastructure components that make up the SaaS solution. This access method is fully auditable and tracked.

This plugin enables you to restrict access to your instance without your express permission, so it may affect support service levels and the Availability SLA. Availability SLA is then measured from the time that Support staff personnel are granted access to your instance.

Procedure

1. Navigate to **All** > **System Security** > **SNC Access Control**.
2. Click **New**.
3. Fill in the form fields (see table).
4. Click **Submit**.

SNC Access Control

Form fields	Description
Name	Names each Customer Service and Support employee who has permission to log in this instance. ○ Express the names as firstname.lastname in lower case letters, separated by a period (for example, `john.smith`). Each name must have a corresponding user record in *`support.servicenow.com`*. ○ If more than one Customer Service and Support employee has permission to log in this instance, enter multiple names and separate them by commas.

Form fields	Description
	○ To enable all Customer Service and Support employees login rights to access the instance, enter an asterisk (*) in place of the name. ○ If you intend on restricting Customer Service and Support employee access to the instance, the values in the **Name** field must not have an asterisk (*) anywhere in the field.
Reason	Human-readable field that describes why you are granting access permission. This field is optional.
Start	Specifies the start date and time of the period during which the specified Customer Service and Support employees have login access. This field is mandatory.
End	Specifies the ending date and time of the period during which the specified Customer Service and Support employees have login access. This field is mandatory.
Active	Controls whether this permission record is active. The default is Active.

Audit logging

The following logging tracks logins and activity by ServiceNow employees.

Event Logs	The event logs show all ServiceNow logins to a customer instance.
Transaction logs	The transaction logs show all activity on the instance, including any efforts to delete logs.

Identity

Know more about the Identities in the instance.

Identity Center

Identity Center allows you to monitor, manage, and minimize identity-based risk and security gaps.

Global Identity

Global Identity is a ServiceNow® store app enables to update user attributes from one instance to instance (multiple instances).

System for Cross-domain Identity Management (SCIM)

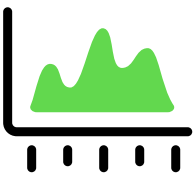

The System for Cross-domain Identity Management (SCIM) API provides endpoints to create, read, update, and delete operations on users and groups using the SCIM protocol.

Identity and Access Audit

Use the Identity and Access Audit to understand the changes made a user, group, role, and ACL.

Global Identity

Global Identity is a ServiceNow® product that enables to identify unique users across multiple instances.

Global Identity helps with multiple ServiceNow® instances to manage users, their attributes, and other foundational data across multiple ServiceNow® instances.

Explore Federated ID  Learn the key features and business value of Federated ID.	Accessing Federated ID  Access Federated ID.
Updating ID fields  Update and set-up the Federated ID.	

Explore Federated ID

Determine the users across multiple instances based on user name and email and provide a unique ID (Federated ID) to the user across instances.

Federated ID is used to identity users across the multiple ServiceNow® instance. Based on the federated id the user can be identified and the accurate number of users across multiple instances can be determined. For more information, see Federated ID.

Note: User name is required for generating Federated IDs.

Federated ID is a unique identifier for an identity using a hashing function across the ServiceNow® instances.

Federated ID for Users

By using the **User ID** and **Email** of the user across the instances, the Federated ID is created and displayed in the **sys_user** table.

After upgrading to the Zurich release, the Federated ID Generation (`com.glide.identity.globalid`) plugin is auto-installed on all instances.

Note:

- User name is required for generating Federated IDs.
- User name and email are used to generate Federated IDs by default. To update the fields for generating Federated IDs based on your requirement, see Update ID fields.
- **iamsync_admin** role is required to update the configuration.
- If there are users with duplicate user names and email, then the Federated ID is generated only for one user. If the user name is null or empty, then the Federated ID is null.

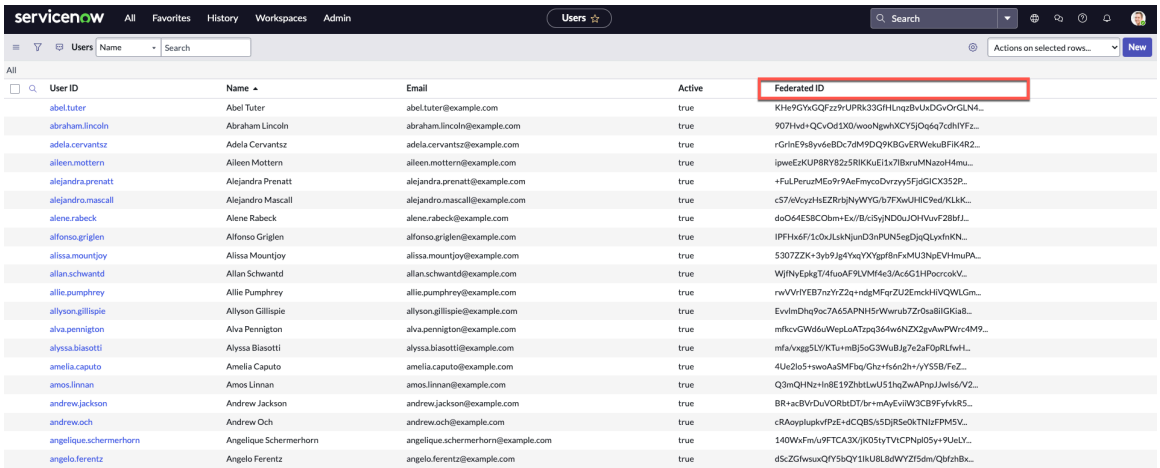

Schema changes after the plugin installed are as follows:

- New column federated_id in the `sys_user` table is created.
- New table - `iamsync_type` is auto populated with the default configuration for the `sys_user` table.

Federated ID is only supported for the `sys_user` table and all the tables that extend the sys_user table.

After upgrading to the Zurich release, the Federated ID Generation (`com.glide.identity.globalid`) plugin is auto-installed on all instances.

Federated ID for Roles

By using the `roles` of the user across the instances, the Federated ID is created and displayed in the **sys_user_role** table.

Note:

- Role is required for generating Federated IDs in the **sys_user_role** table.
- The Federated IDs in the **sys_user_role** table can be used to the over usages of the roles.
- Updating ID fields and Regenerate Federated IDs options are not available for Role type. There's a scheduled job that runs periodically to generate ID for roles in case its empty.

Access Federated ID Criteria

Access Federated ID Criteria to know about the ID fields used to generated Federated ID.

Before you begin
Role required: iamsync_admin

Procedure

1. Navigate to **All** > **Manage Federated ID** > **Federated ID Criteria**.
2. The Federated ID Criterias page, displays the record with the following details:
   - Type Name: **User**
   - Table Name: **User [sys_user]**
   - ID Fields: **user_name (User ID), email** (default fields that are used for Federated ID generation).
   - Status: **Completed** (Federated ID generation status). Available Status: **Ready**, **Running**, **Completed**, **Error**.

Note:
- User name is required for generating Federated IDs.
- User name and email are used to generate Federated IDs by default.

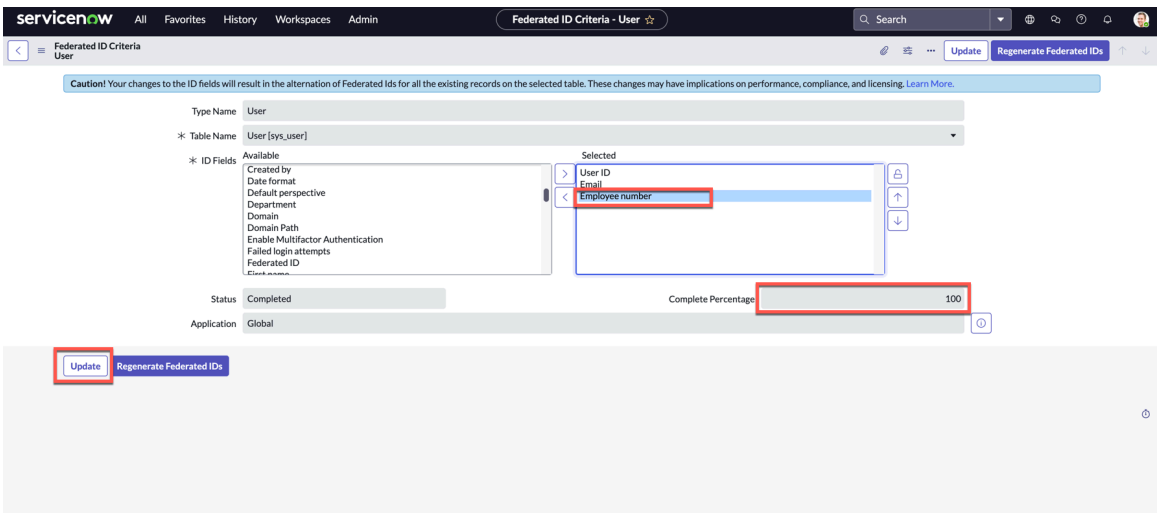

Note: Only the ID fields can be updated to generate a new Federated ID for the existing records. To know more, see Update ID fields.

Update ID fields

Update the ID fields to regenerate the Federated IDs based on the updated fields.

Before you begin
Role required: iamsync_admin

Note: Any changes to the ID fields result in the change of Federated Ids for all the existing records on the selected table. These changes may have implications on performance, compliance, and licensing.

Procedure

1. Navigate to **All** > **Manage Federated ID** > **Federated ID Criteria**.
2. Select the Type Name (**User**).

3. Select the new ID fields in the Federated ID Criteria User page that you want to add from the Available to Selected using the arrow buttons.

   For example, **Employee number**.

   Note:

   - User name is required for generating Federated IDs.
   - User name and email are used to generate Federated IDs by default.
   - If there are users with duplicate user names and email, then the Federated ID is generated only for one user. If the user name is null or empty, then the Federated ID is null.

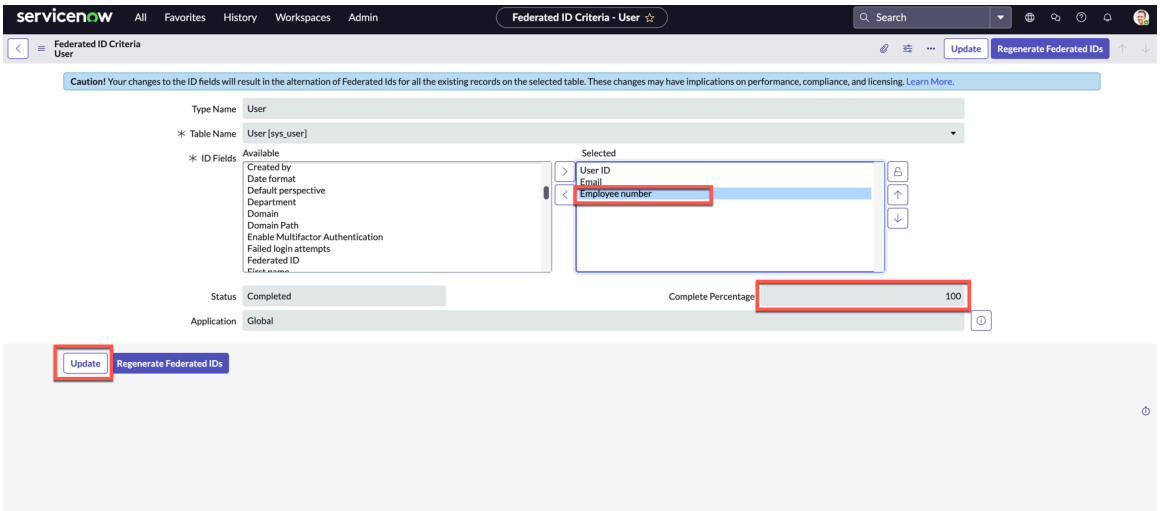

   Now, the **Employee number** selected becomes another attribute for generating the Federated ID.

4. Select **Update** to generate Federated IDs.

   Note: Select **Update** first and then check for the completion percentage (100) before initiating another update.

   The status percentage indicates the Federated ID generation for all the identities across instances.

   Note:

   - Until the previous update job is complete, don't change the field.
   - Fields that are updated should be a string type.
   - Fields that cannot be select as ID fields are as follows:
     - System level fields
     - Edge encryption fields
     - Password fields.

5. Navigate to the sys_user table to view the new Federated IDs that are generated due to updating the ID fields.

   Note: Select **Regenerate Federated IDs** if a user is created or updated via XML imports, low-level database updates, the instance not working correctly. Selecting **Regenerate Federated IDs**, regenerates the IDs for all the users, using the current id field criteria.

Identity and Access Audit

Use the Identity and Access Audit to understand changes made for users, groups, roles, and ACLs.

Explore  Learn about the features and business value of Identity and Access Audit.	Configure  Understand how to configure Identity and Access Audit.
Audit Results  View results from Identity and Access Audit.	

Explore Identity and Access Audit

Use the Identity and Access Audit to understand changes made for users, groups, roles, and ACLs.

Identity and Access Audit helps to understand the critical information about who has modified what, where and when in user accounts, groups and roles.

Helps to detect malicious users and track unusual activity in the ServiceNow® instance and adhere with compliance standards of being able to track access changes.

Identity and Access Audit (Identity Security Audit) is a plugin (`com.glide.security.audit`), which is auto-installed.

Auditing feature can be turned on or off by toggling the`glide.identity.security.audit.enabled` system property. By default, the property is set `true`.

Identity and Access Audit enables you to:

- View the changes made in the last 30 days to users, groups, role ACL attributes, role memberships, group memberships, and ACL roles.
- Track the changes in your ServiceNow instance.
- Help mitigate potential security and regulatory risks.
- Demonstrate compliance with auditors for different groups within the organization.
- Demonstrate that the organization isn't vulnerable to threats related to a lack of visibility in the user group and role changes.

User personas in Identity Access and Audit

Following are the different user personas in Identity and Access Audit:

- **Admin**: View the audit records and the configuration.
- **Security Admin**: View these audit trails. Modify the configuration to enable or disable auditing for a certain table or modify the fields that are being audited.

Audit Tables

The following tables can be audited using Identity and Access Audit:

- Group [sys_user_group]
- Role [sys_user_role]
- Access Control [sys_security_acl]
- User [sys_user]
- Group Role [sys_group_has_role]
- User Role [sys_user_has_role]
- Access Roles [sys_security_acl_role]
- Contained Role [sys_user_role_contains]
- Group Member [sys_user_grmember]

Modules in Identity and Access Audit

Identity and Access Audit has the following modules on the ServiceNow instance:

Module	Description
Audit Results	Displays the audits that occurred in the ServiceNow instance.
Configure Table & Fields	Configure the system tables and fields with the available fields from the Identity and Access Audit.

Module	Description
Configure Retention Period	Configure the retention period of the audited data. The maximum period that can be set is 30 days.
User Trails	Displays audits of users.
Group Trails	Displays audits of groups.
Role Trails	Displays audits of roles.
ACL Trails	Displays audits of ACLs.

Identity Audit Results

Displays audits that occurred in the ServiceNow® instance.

Audit Results displays changes made for users, groups, roles, and ACLs in the ServiceNow instance.

To access the Audit Results, navigate **All** > **System Security** > **Identity and Access Audit** > **Audit Results**. The Security Table Audits page is displayed with the following information.

Security Table Audits

Column Name	Description
Source Table	Details of the source table where the audit occurred.
Action	Describes the audit action.
Sys_id	Details about the sys_id for the audited record.
Created by	Details of the audit that was created by.
Transaction ID	A unique ID that represents every action on the particular audit performed.
Changed for user	Name of the user when the audit was performed.
Created	The time and date the audit was performed.

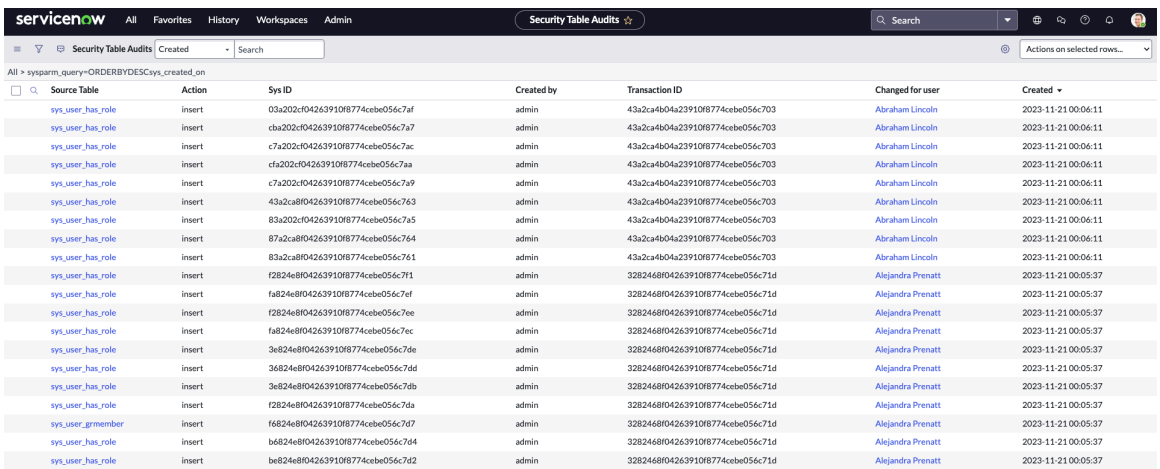

User Trails

Displays audits of users in the ServiceNow® instance.

User Trails display identity attribute changes, role membership changes, and group membership changes for a user.

To access User Trails, navigate to **All** > **System Security** > **Identity and Access Audit** > **User Trails**. The User Trails page is displayed with the following information.

User Trails

User Name	Name of the user.
User Sys id	Details about the User sys_id for the audited record.
Source Table	Details about the Source Table where the audit was performed.
Action	Describes the audit action.
Created By	The user who made the change.
Created	The time and date the audit was performed.

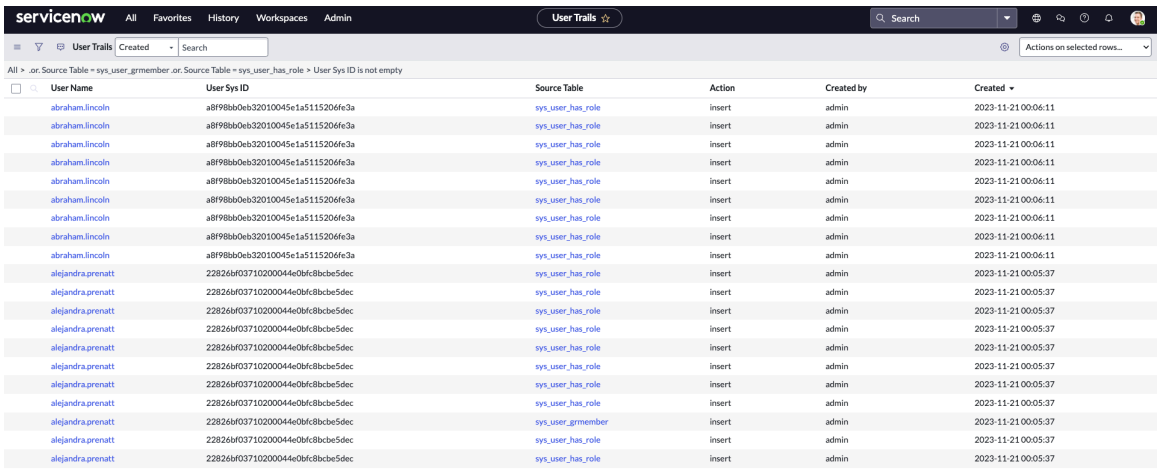

Group Trails

Displays audits of groups in the ServiceNow® instance.

Group Trails display attribute changes, membership changes, and roles changes for a group.

To access Group Trails, navigate to **All** > **System Security** > **Identity and Access Audit** > **Group Trails**. The Group Trails page is displayed with the following information.

Group Trails

Group Name	Name of the group.
Group Sys id	Details about the Group sys_id for the audited record.

Group Trails (continued)

Source Table	Details about the Source Table where the audit performed.
Action	Describes the audit action.
Created By	The user who made the change.
Created	The time and date the audit was performed.

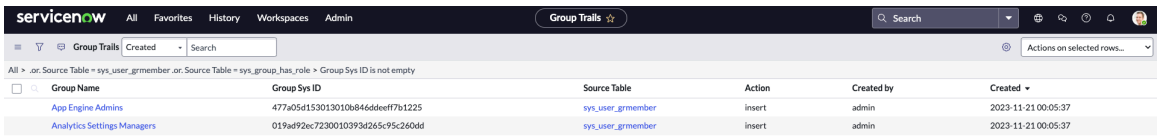

Role Trails

Displays audits of roles in the ServiceNow® instance.

Role Trails display attribute changes and parent-child relation changes for a role.

To access Role Trails, navigate to **All** > **System Security** > **Identity and Access Audit** > **Role Trails**. The Role Trails page is displayed with the following information.

Role Trails

Role Name	Name of the role.
Role Sys id	Details about the Role sys_id for the audited record.
Source Table	Details about the Source Table where the audit performed.
Action	Describes the audit action.
Created By	The user who made the change.
Created	The time and date the audit was performed.

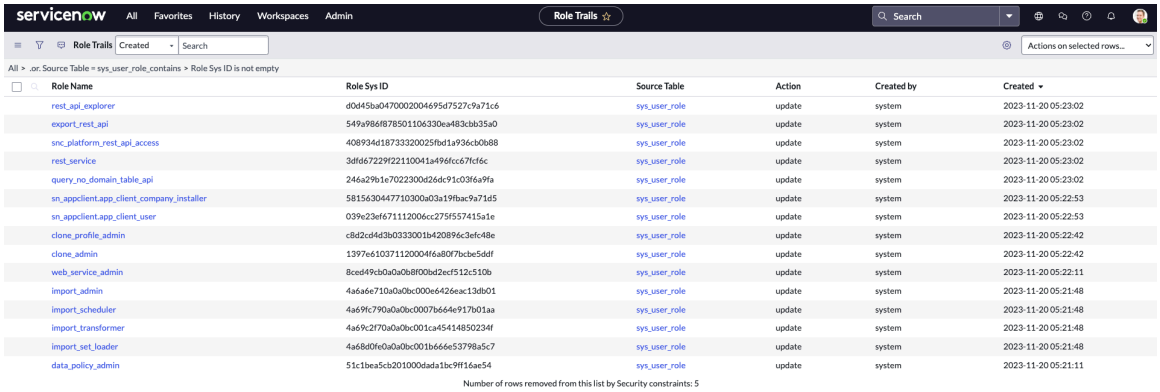

ACL Trails

Displays audits of ACLs in the ServiceNow® instance.

ACL Trails display attribute changes and required role relation changes for an ACL.

To access the ACL Trails, navigate to **All** > **System Security** > **Identity and Access Audit** > **ACL Trails**. The ACL Trails page is displayed with the following information.

ACL Trails

ACL Name	Name of the ACL.
ACL Sys id	Details about the ACL sys_id for the audited record.
Source Table	Details about the Source Table where the audit was performed.
Action	Describes the audit action.
Created By	The user who made the change.
Created	The time and date the audit was performed.

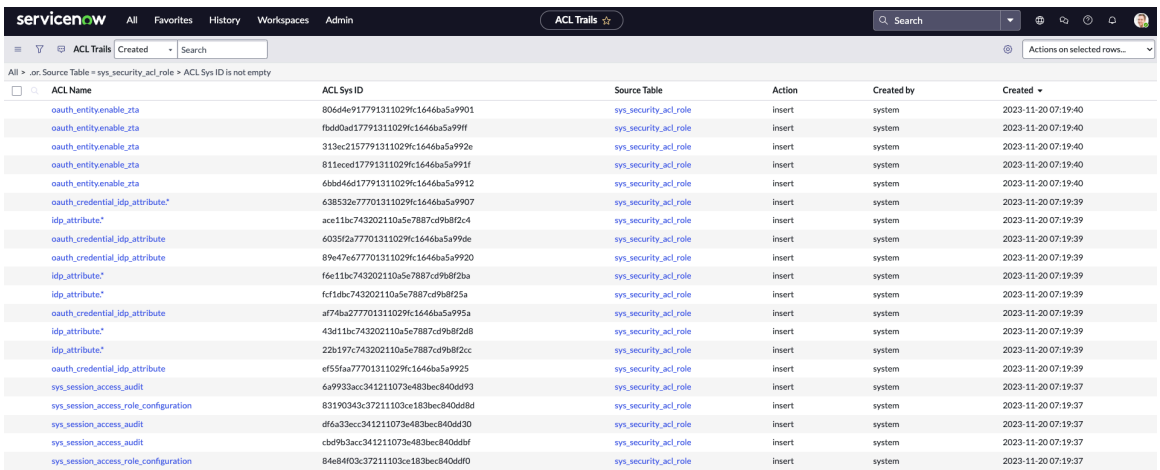

Security Auditable Fields

Displays table and field level details that will be audited in the ServiceNow® instance.

Security Auditable Fields displays the details of tables and fields that will be audited in the ServiceNow instance.

To access the Security Auditable Fields page, navigate to **All** > **System Security** > **Identity and Access Audit** > **Configure Tables & Fields**. The Security Auditable Fields page is displayed with the following information.

Security Auditable Fields

Column Name	Description
Table to Audit	Details of the table that is audited.
Audit Storage Destination	Details of the destination where the audit details are stored.
Field List	Auditing will be done for fields specified in the list.
Create	Audit the changes that are related to Create operation.
Update	Audit the changes that are related to Update operation.
Delete	Audit the changes that are related to Delete operation.
Active	Audit only if the configuration for the table is active.

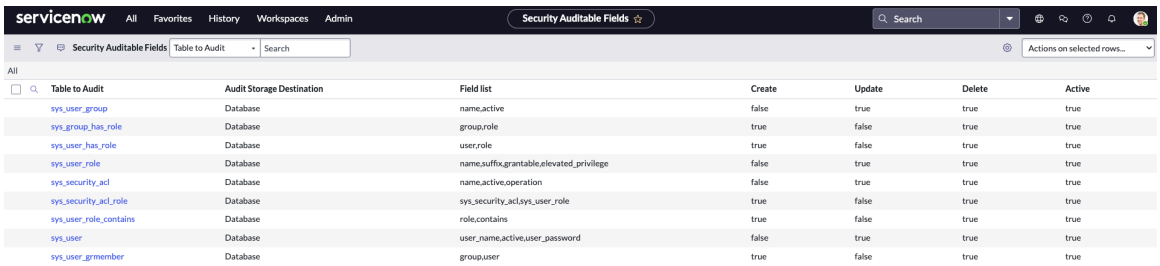

The following tables can be audited using the Identity and Access Audit:

- Group [sys_user_group]
- Role [sys_user_role]
- Access Control [sys_security_acl]
- User [sys_user]
- Group Role [sys_group_has_role]
- User Role [sys_user_has_role]
- Access Roles [sys_security_acl_role]

- Contained Role [sys_user_role_contains]
- Group Member [sys_user_grmember]

Configure Tables and Fields

Identity and Access Audit to understand the changes made for a user, group, role, and ACL.

Before you begin

Role required: security_admin

You must elevate your role to Security Admin to configure tables and fields for Identity and Access Audit.

The following tables can be configured for auditing:

- Group [sys_user_group]
- Role [sys_user_role]
- Access Control [sys_security_acl]
- User [sys_user]
- Group Role [sys_group_has_role]
- User Role [sys_user_has_role]
- Access Roles [sys_security_acl_role]
- Contained Role [sys_user_role_contains]
- Group Member [sys_user_grmember]

Note: To understand which fields can be configured for the tables, see Supported and unsupported fields in Identity Access and Audit.

Procedure

1. Navigate to **All** > **System Security** > **Identity and Access Audit** > **Configure Tables & Fields**.
2. Select the table that you want to audit a field from.

   For example, **sys_user**.

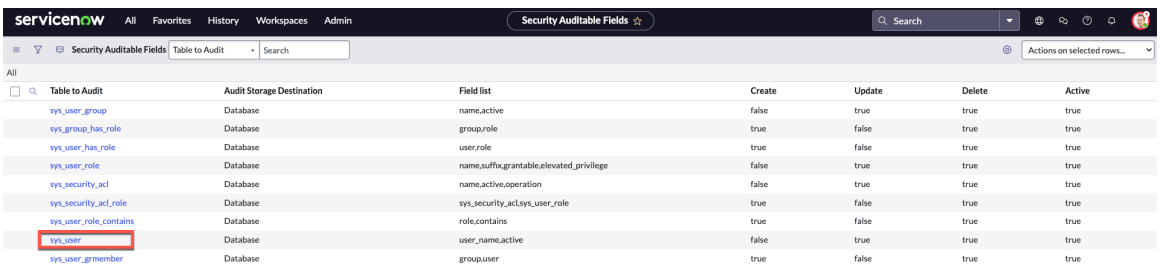

3. Add the field to be audited.

   For example, **Password**.

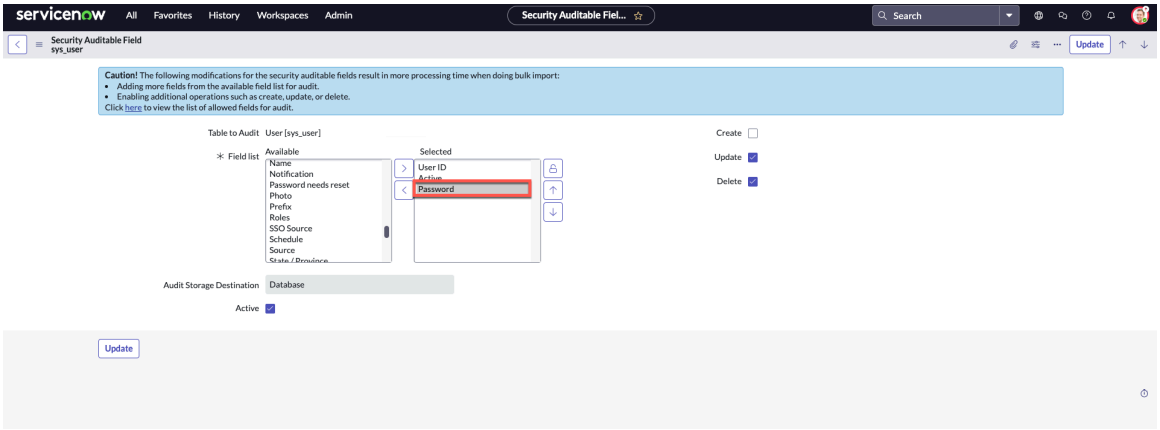

Note: The following modifications for the security auditable fields result in more processing time when doing bulk import:

- Adding more fields from the available field list for audit.
- Enabling additional operations such as create, update, or delete.

4. Update the record.

Any changes to the password field add a new record to the Security Table Audits. In this example, the audit shows a changed password field for the user **Abel Tuter**.

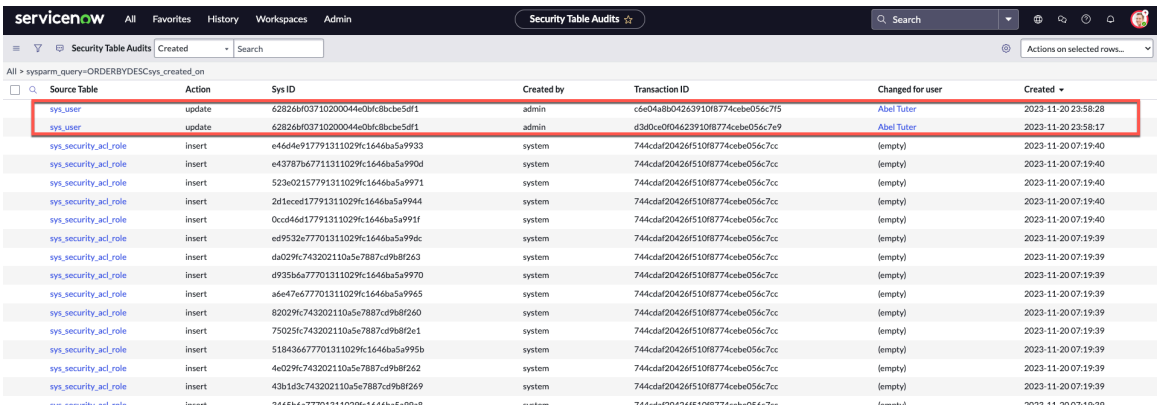

Selecting the created record displays the details of the changes.

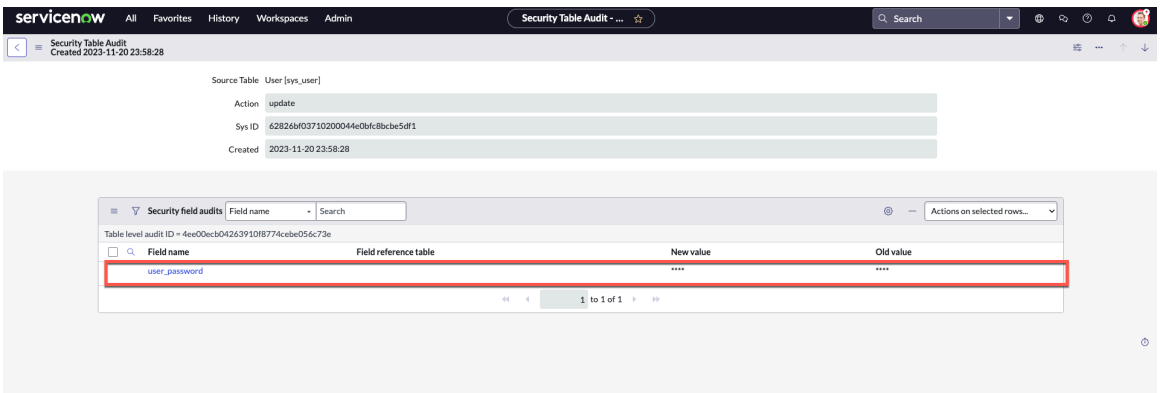

Configure Retention Period

Configure the retention period for Identity and Access Audit.

Before you begin
Role required: admin

Procedure
1. Navigate to **All** > **System Security** > **Identity and Access Audit** > **Configure Retention Period**.
2. On the form, you can change the Age in seconds.

   Note: The maximum number of days for Identity and Access Audit is 30 days (2,592,000 seconds).

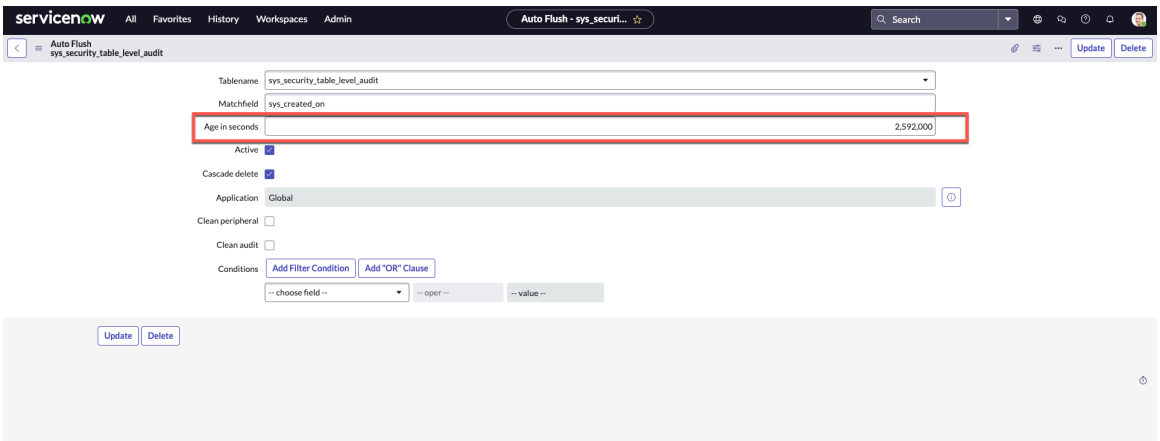

3. Save or Update the record.

Supported and unsupported fields in Identity Access and Audit

List of fields that are supported and not supported for audit.

Audit Fields Validation prevents some fields from being chosen in the field_list in the Security Field Audit Config (sys_sec_field_audit_config) table.

Field supported or not supported for Audit

Table	Fields supported or not supported
All tables	Fields that aren't supported: • Created On (sys_created_on) • Created By (sys_created_by) • Updated By (sys_updated_by) • Updated On (sys_updated_on)
User has role (sys_user_has_role)	Fields that are supported: • User • Role • Inherited • Count
User (sys_user)	Fields that aren't supported:

Field supported or not supported for Audit (continued)

Table	Fields supported or not supported
	• Last Login (last_login) • Last Login Time (last_login_time) • Last Login Device (last_login_device) • Enable Multi-factor Authentication (enable_multifactor_authn) • Default Perspective (default_perspective) • Calendar Integration (calendar_integration) • Federated ID (federated_id) • Password Needs Reset (password_needs_reset) • Failed Attempts (failed_attempts) • Last Password (last_password) • LDAP Server (ldap_server) • Locked Out (locked_out) • Notification (notification) • Roles (roles) • Domain (sys_domain) • Domain Path (sys_domain_path) • Time Format (time_format) • Hashed User ID (hashed_user_id) • Class Name (sys_class_name) • Mod Count (sys_mod_count)

Identity Center

Allows you to monitor, manage, and minimize identity-based risk and security gaps.

Explore	Activate
	
Learn the features and business value of Identity Center.	Understand how to activate Identity Center.

Identity Metrics for administrators

Know how the permissions are evaluated.

Identity Center for users

Get details about the commonly asked questions about the Access Analyzer.

Explore Identity Center

Identity Center is a collection of user attributes, access, devices, login history, security activity, and much more.

Identity Center offers capabilities to monitor, manage, and minimize identity-based risk and security gaps.

Identity Center is one stop to monitor, manage, and minimize identity-based risk and security gaps on ServiceNow.

To activate the Identity Center, install the Identity Center (`com.snc.identity_center`) plugin. Identity Center is available for the end user - to view the details about the active sessions, login history, and trusted devices with the Identity Center. For more information see, Identity Center for users.

Activate the Identity Center

For Identity Center, install the Identity Center (`com.snc.identity_center`) plugin.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Identity Center (`com.snc.identity_center`) plugin using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Identity Center for users

View the details about your active sessions, login history, and trusted devices with the Identity Center.

Identity Center is a collection of user attributes, devices, login history, security activity and much more. It provides a single pane view of all the data with additional security controls and notification capabilities.

With the Identity Center you can view the details about your active sessions, login history, and trusted devices.

To access the Identity Center for users, navigate to one of the following:

- On ServiceNow AI Platform, navigate to **All** > **Self-Service** > **My Profile** and select **View Identity Center** under the Related Links section.

  Note: You can also access your profile by selecting your user name in the instance header.

- On Now Support, select the profile, select **View Identity Center** at the bottom of the page.

The Identity Center page is displayed as follows:

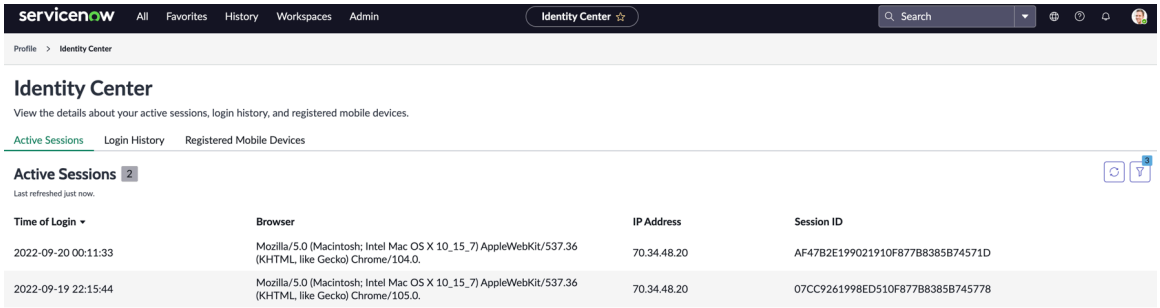

Identity Center has the following tabs:

- Active Sessions
- Login History
- Registered Mobile Devices

You can select these tabs to view the details such browser, IP Addressees, session relation information, login information, your registered mobile devices.

View Active Sessions in Identity Center

Displays the information about the user sessions.

Active sessions are the sessions that are open on the current#ServiceNow®#nstance with different browsers or devices.

The Active Sessions tab in the Identity Center helps you to identity your sessions based on the Browser, IP Address, and Sessions ID. Using this information you can take required actions such as extending or terminating the session.

Using this information can help determine if the sessions are real and not a security concern.

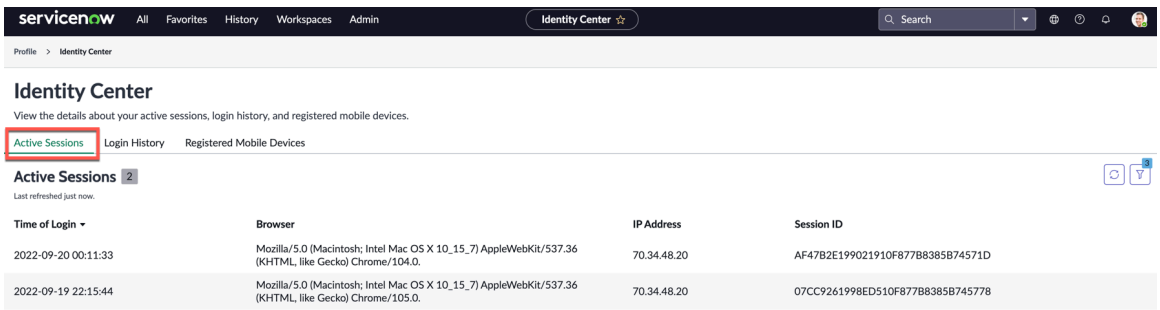

View Login History in Identity Center

Provides details about your login history.

The Login History tab in the Identity Center helps you with your login information and the status of your login.

You can use the filters to specify login actions to help security investigations. Filter can help determine if the activity was real or suspicious and report the information to your administrator.

Pagination in the login history is 20 by default and can be set to a maximum of 100.

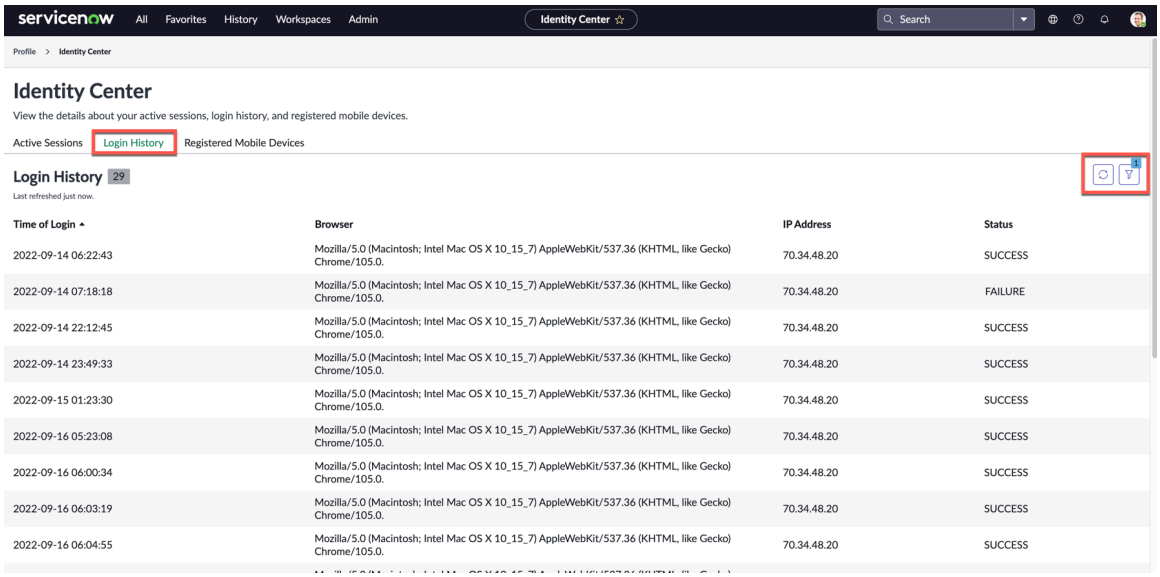

Following are some of the other details from the Login History:

- The Time of Login is the updated time when the mobile devices are used to access the instance.
- The records are stored in Identity Center for 30 days.

View Registered Mobile Devices in Identity Center

Provide details about your registered mobile on the ServiceNow instance.

The Registered Mobile Devices tab in the Identity Center displays the details about your registered devices that were used to access the ServiceNow instance.

If the Adaptive Authentication module is enabled and a registered mobile device is enrolled the Identity Center displays the details about your registered devices.

Further, it also displays details of the devices such as Operating System, Device Id, and Status of the device along with the Registration Time of the device.

To register your mobile device, you must make sure the Adaptive Authentication (*com.snc.adaptive_authentication*) plugin is installed and you've activated your Trusted Mobile App feature. For more information, see Activate Trusted Mobile App.

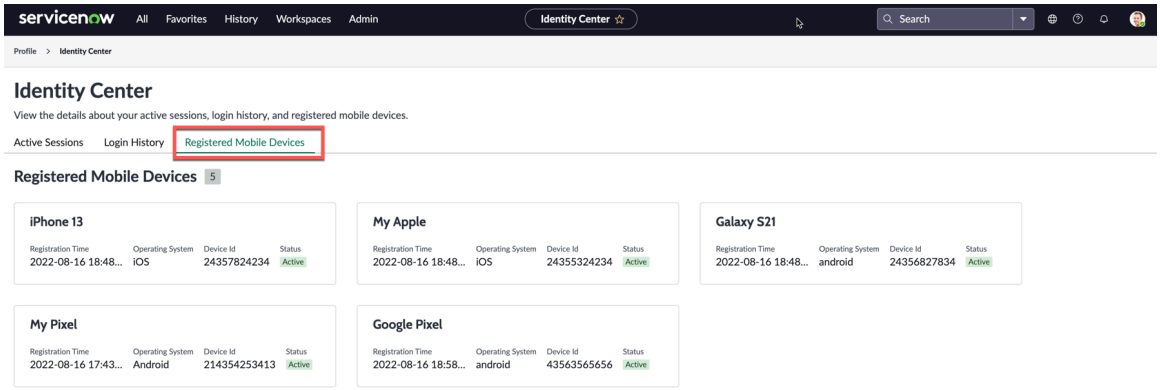

Identity Metrics for administrators

View trends of the users, privileged users, active sessions, and integrated account on your ServiceNow instance.

Identity Metrics for administrators has the trends for the following:

- Users
- Privileged users
- Integration or non-human accounts
- Active sessions
- Inactive sessions

To learn more, see the Security Center Metrics.

Machine Identity Console

Manage your service accounts which are used for integrations with ServiceNow.

Machine identities are digital identities that interact with systems and data. These identities are used to perform tasks autonomously. Machine Identity Console helps you to understand the security score of these identities and provide recommendations.

Explore	Activate	Configure
Learn the features and business value of Machine Identity Console.	Understand how to Activate Machine Identity Console.	Configure Inbound integrations using the Machine Identity Console.

Explore Machine Identity Console

Manage your service accounts, which are used for integrations with ServiceNow.

Machine Identity Console enables you to manage your non-human identities (NHIs) that are used to identify, authenticate, and authorize different software entities to access secured resources of ServiceNow. The entities that gain access include applications, workloads, APIs, bots, and automated systems. Unlike human identities, non-human identities (NHIs) aren't governed in the same manner and aren't directly associated with a human. Their identity and verification methods are distinct from the human users, and standard human security measures are not applicable to them.

https://player.vimeo.com/video/1101308622?h=b01d69f65a&badge=0&autopause=0&player_id=0&app_id=58479

Here's what is available on the Machine Identity Console Overview page:

- Total Machine identity (integration) accounts and the Accounts with high privilege roles.
- Unique API calls - Last 7 days
- Authentication method used - Last 7 days
- Machine Identity Security score and findings related to the identity.

Note: Accounts that have the **Internal Integration User** field set to `true` in their `sys_user` record will not populate data in the Machine Identity Console.

Related topics

Activate Machine Identity Console

Security findings

Metrics

Machine Identity Console Settings

Inbound integrations

Inbound integrations

Inbound Integrations in the machine identity console allows you to configure and manage external applications to access ServiceNow APIs.

The Inbound integrations within the ServiceNow's Machine Identity Console helps you manage and configure integrations and applications that connect with the ServiceNow platform. It acts as a central hub for storing application settings and API credentials, ensuring secure communication between ServiceNow and external systems. To know more, see New Inbound integrations experience.

Connections types:

- OAuth - Authorization code grant: Access a resource by authenticating directly with an OAuth server that trusts the resource. For more information, see Configure an OAuth authorization code grant.
- OAuth - Client credential grant: Used for Inbound Integrations from a third party OAuth client to the ServiceNow platform. For more information, see Configure an OAuth Client credential grant.
- OAuth - JWT bearer grant: Use a JSON Web Token (JWT) grant to authenticate with your instance. For more information, see Configure an OAuth JSON web token bearer grant.

- Third party ID token issued by OIDC supporting identity provider: Configure an OAuth OIDC provider to accept identity tokens generated by a third-party OIDC provider using inbound API calls. For more information, see Configure a third party ID token.
- OAuth - Resource owner password credential grant: Authorize Access to an OAuth Endpoint using auth code flow. For more information, see Configure an OAuth resource owner password credential grant.

Note: For authorization code flow, user must complete the Authentication by local login, SSO or MFA and then provide consent.

To learn more, see New Inbound integrations experience.

Related topics

Inbound integrations

Configure an OAuth authorization code grant

Configure an OAuth JSON web token bearer grant

Configure a third party ID token

Configure an OAuth Client credential grant

Configure an OAuth resource owner password credential grant

Security findings

Provides Machine Identity security score and findings.

Machine identity security score and findings

The security score is displayed for following security findings. You can select each findings for learn more.

Security findings

Findings	Description
Accounts with no login for 100 days	Findings about the accounts that have not accessed any API in 100 days.
Accounts using Basic Authentication	Findings about the accounts that are using username and password for authentication.
Integration accounts with Web Service Access disabled	Findings about all the accounts that have WSA disabled.
Accounts performing both UI and API logins	Findings about the accounts that are used for both interactive (UI) and machine (API) logins.

Machine identity Security score and findings

The security score is based on the usage and method of machine identities. A lower score indicates a higher risk. You can view the recommendations to take preventive actions for these identities.

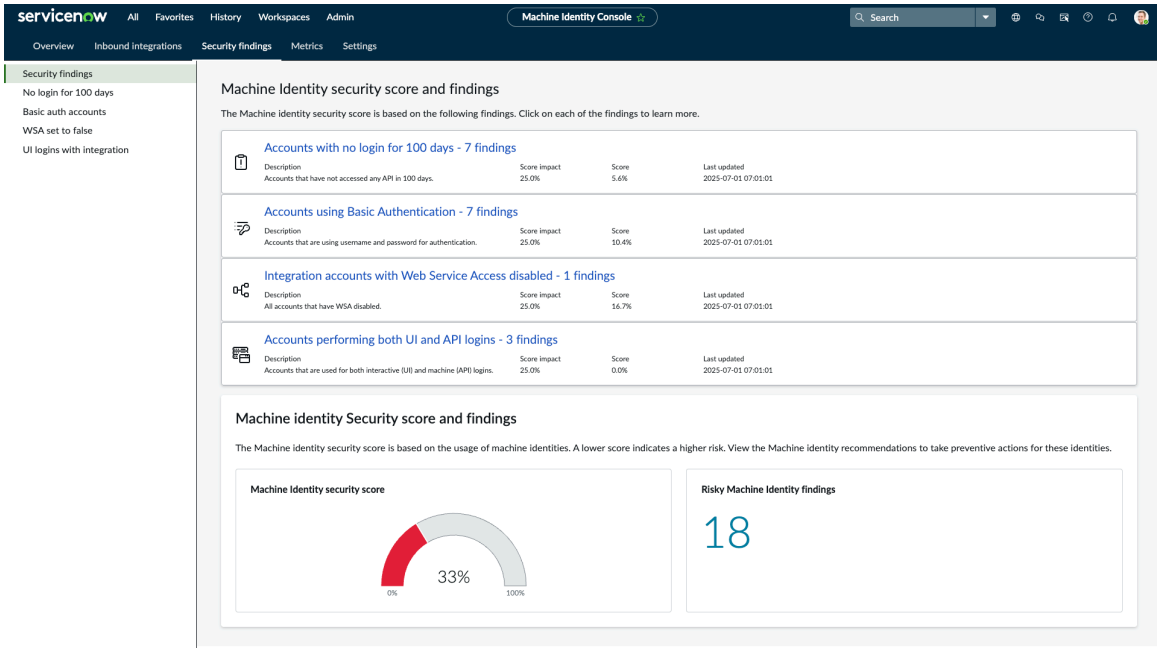

Accounts with no login for 100 days

Display the findings about the accounts that have not accessed any API in 100 days under the Security findings in the Machine Identity Console.

Accounts with no login for 100 days displays the non human identity accounts that have not accessed any API in 100 days.

Note: Any changes made to the record displayed on this page are immediately updated in the list, risk score resulting from those changes will be reflected the following day.

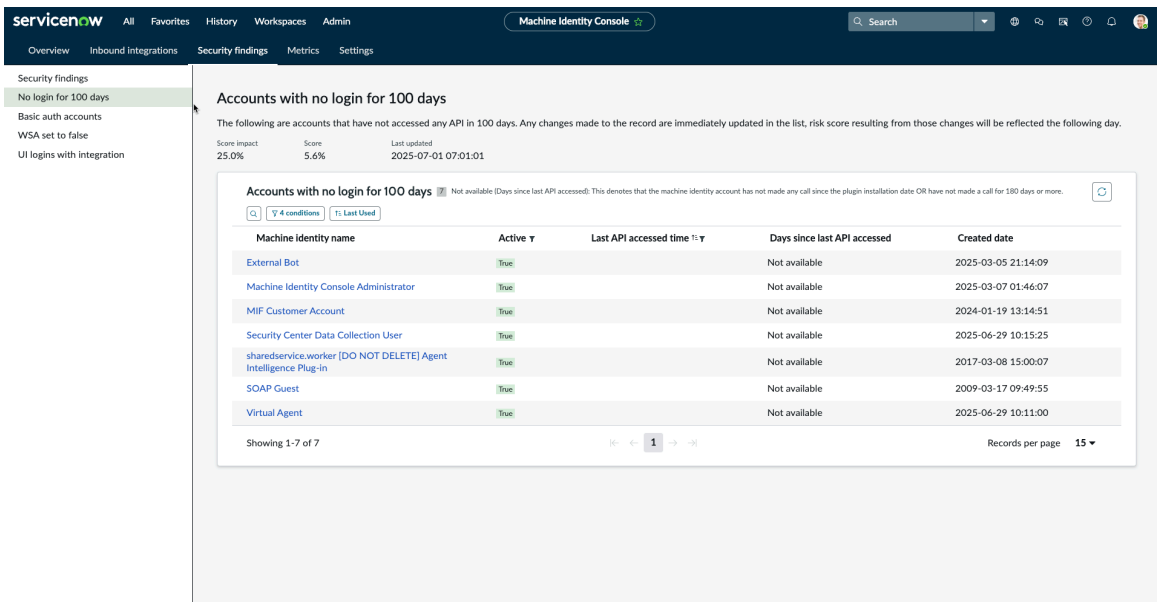

You can select the machine identity name to know more about the account and the recommendation to maintain a good security posture for the account.

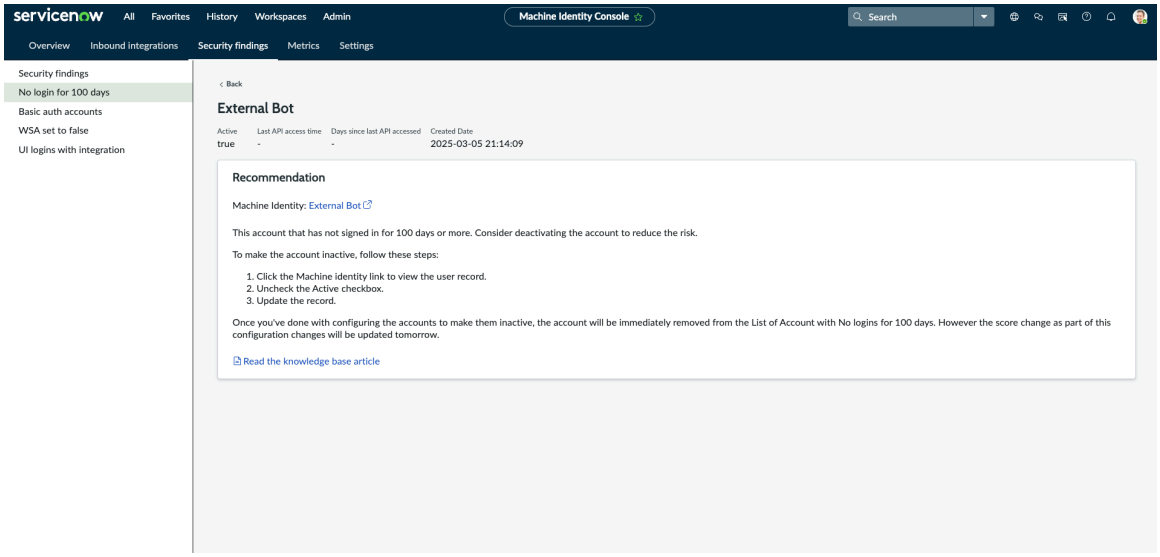

Accounts using Basic Authentication

Display the findings about the accounts that are using only basic authentication to authenticate the APIs under the Security findings in the Machine Identity Console.

Accounts using basic authentication displays the non human identity accounts that using basic authentication for the APIs to login to ServiceNow.

Note: The accounts displayed on the page are the accounts that are using username and password for authentication in the last 30 days.

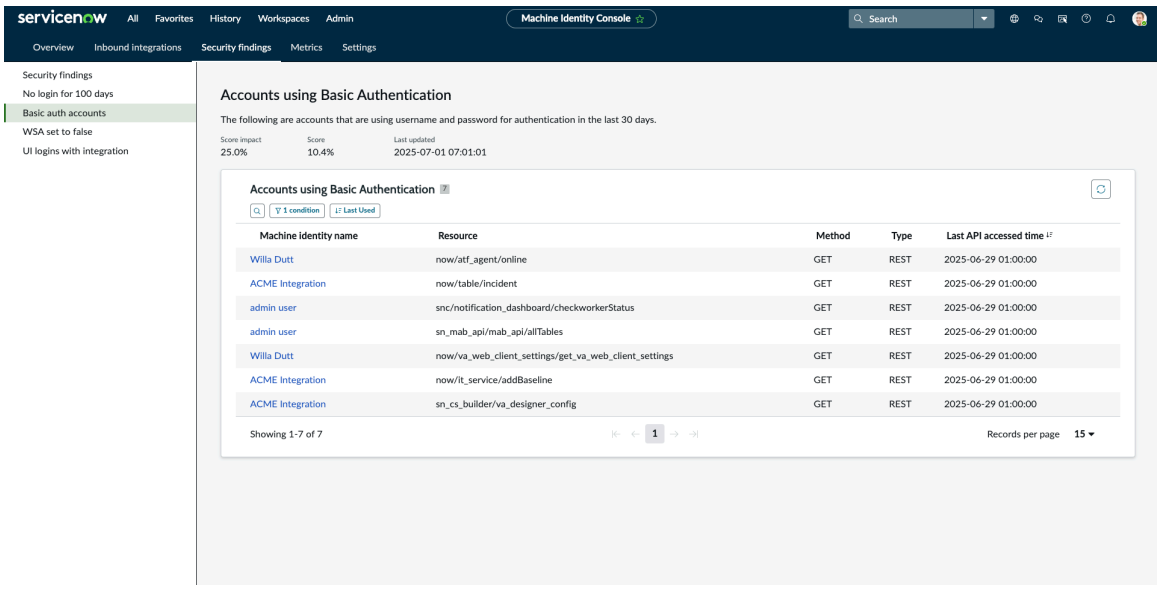

You can select the machine identity name to know more about the account and the recommendation to maintain a good security posture for the account.

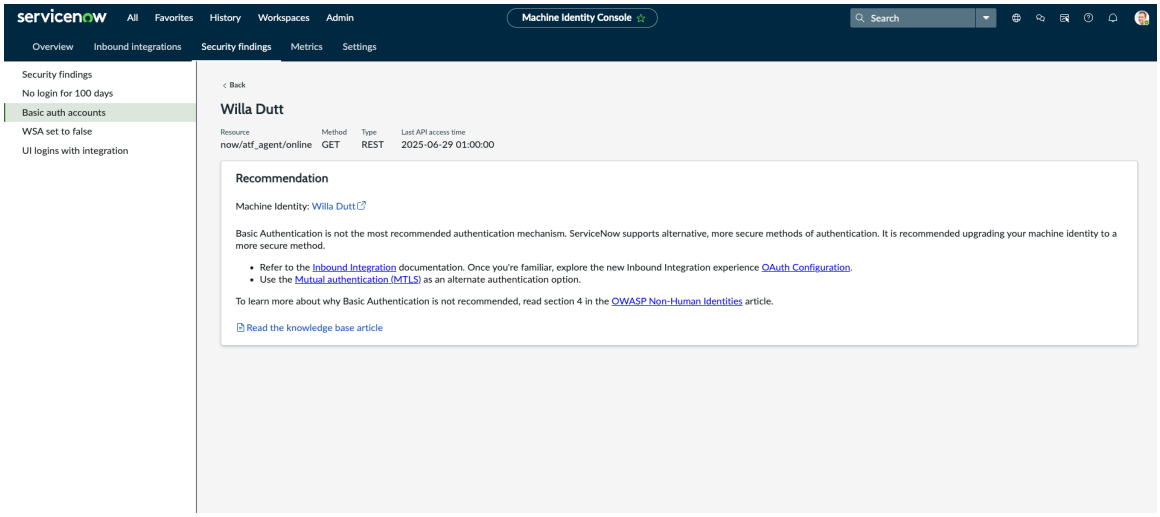

Integration accounts with Web Service Access set to false

Display the findings about the accounts that are authentication ServiceNow with the Web Service Access set to false under the Security findings in the Machine Identity Console.

Integration accounts with Web Service Access set to false displays the accounts that are using API without the Web Service Access set to false.

Note:

- Any changes made to the record displayed on this page are immediately updated in the list, risk score resulting from those changes will be reflected the following day.
- Accounts that have the **Internal Integration User** field set to `true` in their `sys_user` record will not populate data in the Machine Identity Console.

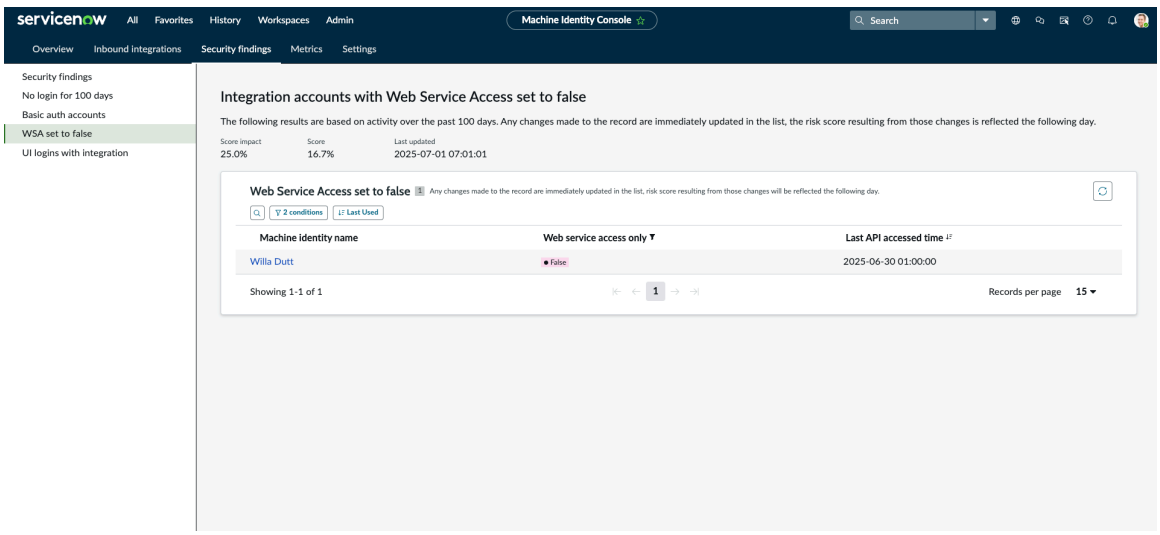

You can select the machine identity name to know more about the account and the recommendation to maintain a good security posture for the account.

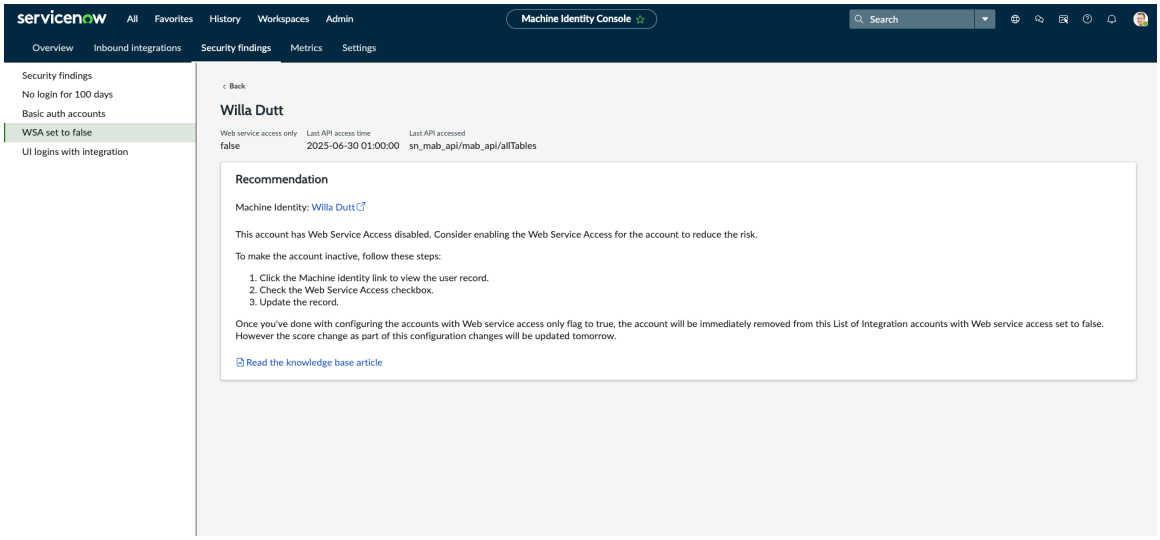

Accounts performing both UI and API login

Display the findings about the accounts that are performing UI and API login under the Security findings in the Machine Identity Console.

Accounts performing both UI and API login displays the accounts that using UI and API login to authenticate to ServiceNow.

Note: The accounts displayed on the page are based on activity over the past 30 days and any changes made today is reflected after 30 days.

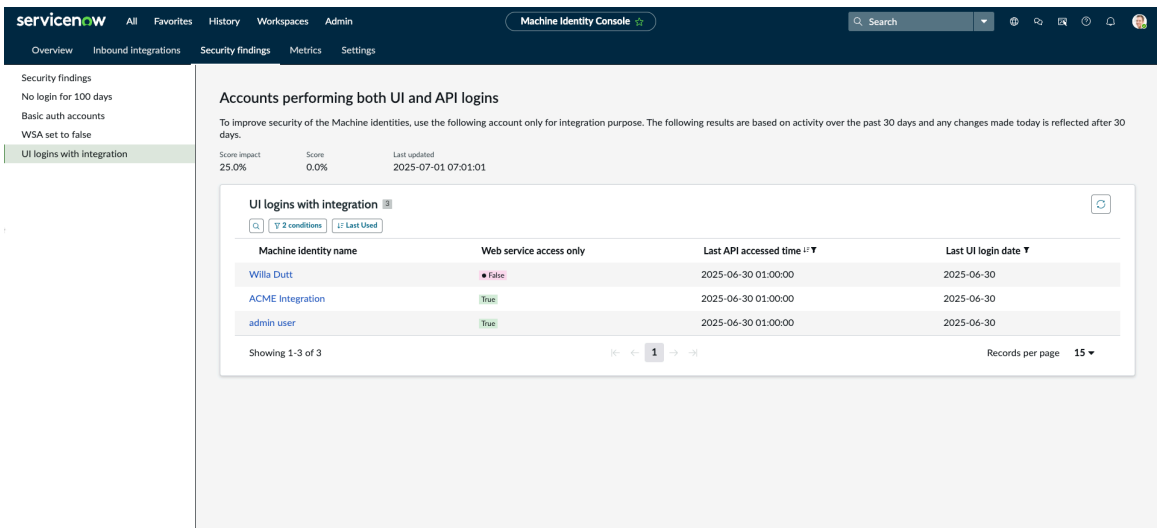

You can select the machine identity name to know more about the account and the recommendation to maintain a good security posture for the account.

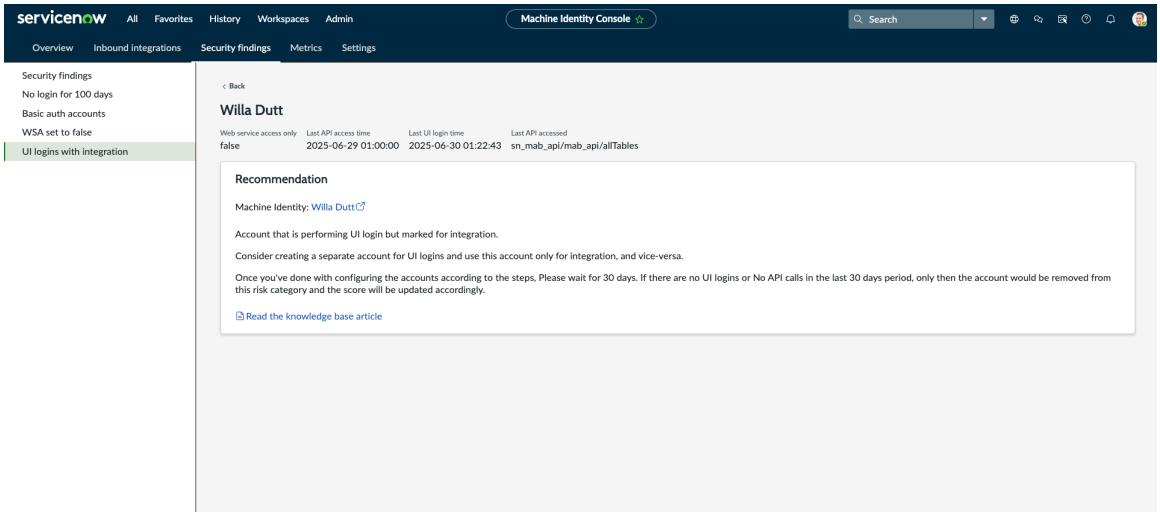

Metrics

Metrics of the machine identities.

Unique API calls in last 7 days

The following results are unique API calls in the last 7 days. Records collected in this table are in real time.

Note: Accounts that have the **Internal Integration User** field set to `true` in their `sys_user` record will not populate data in the Machine Identity Console.

Machine identity accounts

The following results are the machine integration accounts used for integrations. Records collected in this table are in real time.

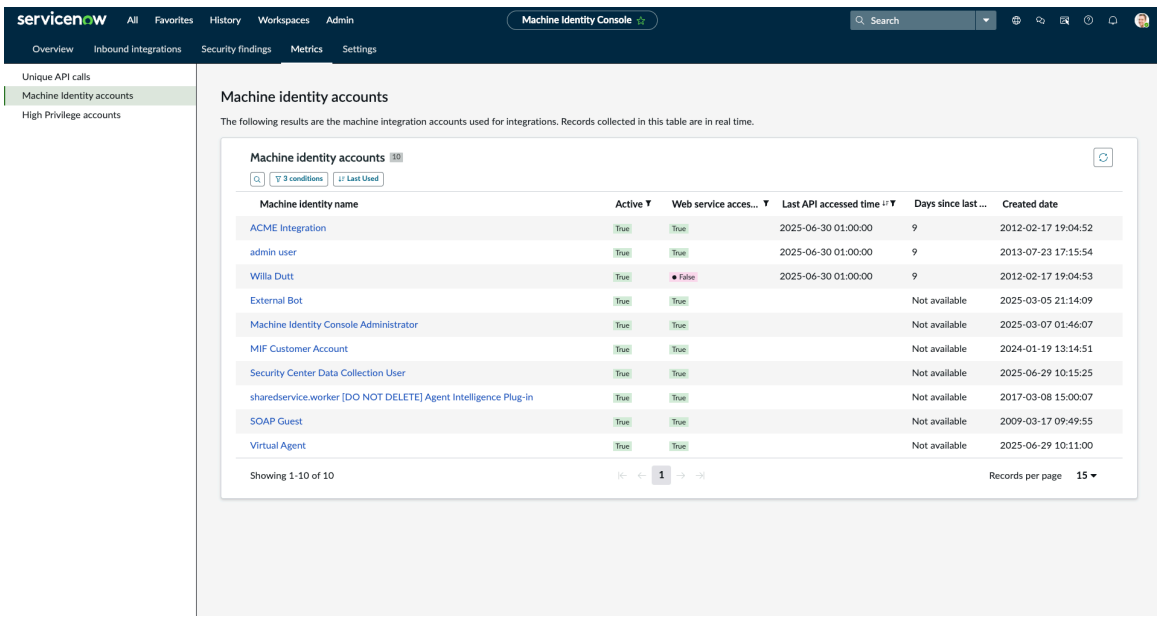

High Privilege machine accounts

The following results are the high privilege machine integration accounts used for integrations. Records collected in this table are in real time.

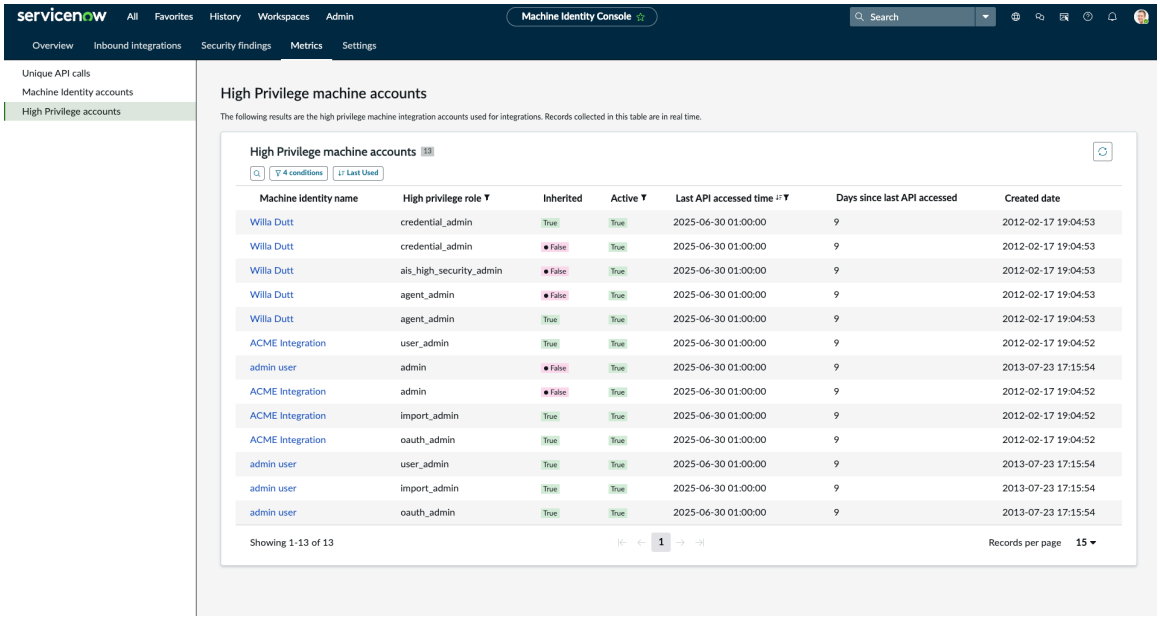

Machine Identity Console Settings

Configure the Machine Identity security score settings.

Before you begin

Role required: mi_admin

The following task is specific to your organization on how to handle the inactive accounts.

Note: If your organization does not want to report on accounts inactive for 100 days, you can deactivate this report below. This is the only report that can be set to inactive.

Procedure

1. Navigate to **All** > **Machine Identity Console** > **Settings**.
2. Select the settings based on your requirements:
   - **Accounts with no login for 100 days**
   - **Accounts using Basic Authentication**
   - **Integration accounts with Web Service Access disabled**
   - **Accounts logging into the UI despite being integration only**

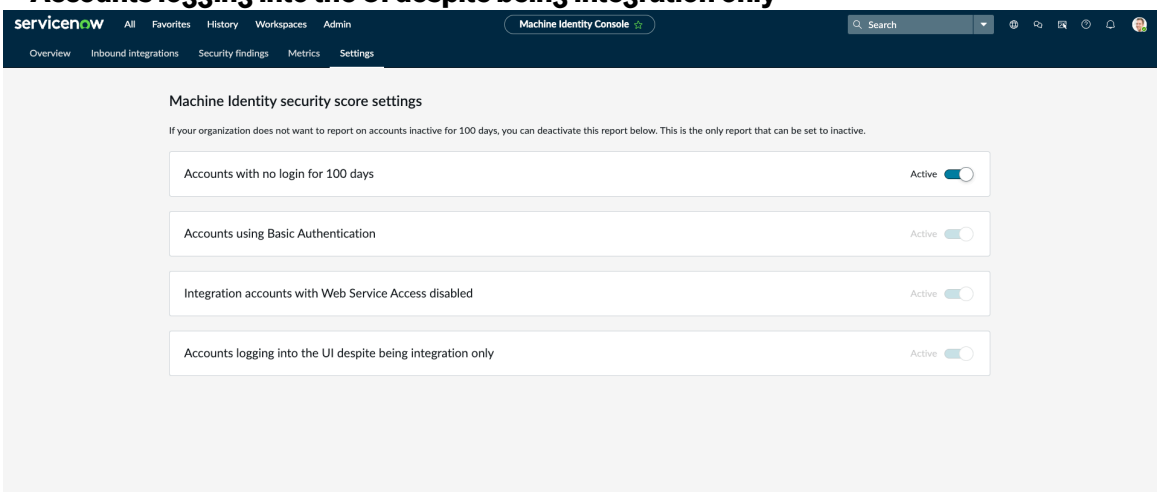

Activate Machine Identity Console

Activate the **Machine Identity Management** `com.glide.identity.machine_identity_management` to manage your service accounts which are used for integrations with ServiceNow.

Before you begin

Role required: mi_admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the **Machine Identity Management** (`com.glide.identity.machine_identity_management`) plugin using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Use Machine Identity Console

Manage your service accounts which are used for integrations with ServiceNow using the Machine Identity Console.

Before you begin

Role required: mi_admin

Activate the **Machine Identity Management** `com.glide.identity.machine_identity_management` plugin. For more information, see Activate Machine Identity Console.

Procedure

1. Navigate to **All** > **Machine Identity Console** > **Machine Identity Console**.
   The Machine Identity Console has the following tabs:
   - Overview: Displays the overall machine identity accounts and their metrics
   - Inbound integrations: Configure Inbound integrations for the APIs.
   - Security findings: Displays the security score is based on different security findings.
   - Metrics: Displays the different metrics for the machine identity accounts
   - Settings: Configure the machine identity console.
2. Select the appropriate tabs based on your security requirements.

System for Cross-domain Identity Management (SCIM)

The System for Cross-domain Identity Management (SCIM) API provides endpoints to create, read, update, and delete operations on users and groups using the SCIM protocol.

SCIM Provider

The SCIM provider synchronizes the changes made to identities in the IdP, including creating, updating, or deleting records.

SCIM Client

The SCIM Client is used for creating, updating, and deleting identity resources in a system that supports SCIM compliant REST requests.

The SCIM protocol is an application-level HTTP-based protocol based on the HTTP (RFC7230) standard. Use this API for provisioning and managing identity data, such as users and groups. Use the API on the web and in cross-domain environments, such as enterprise-to-cloud service providers or inter-cloud scenarios.

To access this API, you must activate the SCIM v2 - ServiceNow® Cross-domain Identity Management (com.snc.integration.scim2) plugin. To know more about the SCIM API, see System for Cross-domain Identity Management (SCIM) API.

SCIM Provider

The Service Provider provisions users and groups using the SCIM API.

Explore

Learn about SCIM Provider.

Activate

Activate SCIM.

SCIM customization

Get details about how to customize SCIM.

Source definition

Learn more about source definition for SCIM.

Explore SCIM Provider

The Service Provider provisions users and groups using the SCIM API.

As a SCIM provider, the ServiceNow schemas support SCIM APIs to provision users and groups.

The SCIM provider synchronizes the changes made to identities in the IdP, including creating, updating, or deleting records. These changes are automatically synchronized to the provider according to the SCIM protocol. Also, the IdP can read identities from the provider to add to the IdP directory. The IdP can then detect incorrect values in the provider that could create security vulnerabilities. The synchronization enables end users to have seamless access to applications for which they're assigned, with up-to-date profiles and permissions.

Configurations for SCIM Provider

To configure the SCIM Provider, perform the following tasks:

- Activate the **SCIM v2 - ServiceNow Cross-domain Identity Management** plugin. To learn more, see Activate the SCIM plugin.
- Activate the other plugins that other plugins that are required for SCIM:
  - OAuth 2.0
  - REST API Provider
  - REST API Access Policy
- Add the scim_admin role as part of the SCIM service.

⚠ **Warning:** Grant this role carefully. The scim_admin role is equivalent to giving the user the admin role, where the scmin_admin can add or update Personally Identifiable Information (PII).

Tables

Two tables, sys_user and sys_group, contain the SCIM attributes that do not map to existing ServiceNow tables. To know more about the tables, see the SCIM-specific tables.

Activate the SCIM plugin

For SCIM activation, install the SCIM v2 - ServiceNow Cross-domain Identity Management (com.snc.integration.scim2) plugin.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the SCIM v2 - ServiceNow Cross-domain Identity Management (com.snc.integration.scim2) plugin using the filter criteria and search bar.

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

3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

Tutorial: Configure SCIM for user provisioning with a Provider

Configuring SCIM automatically provisions and de-provisions users and groups to ServiceNow by using the providers provisioning service.

Before you begin

Install SCIM plugin.

Role required: scim_admin

Warning: Grant this role carefully. The scim_admin role is equivalent to giving the user the admin role, where the scmin_admin can add or update Personally Identifiable Information (PII).

About this task

You can provision users using SCIM by the following authentication methods:

- Provision user using Basic Authentication
- Provision user using OAuth

Provision user using Basic Authentication

Configuring SCIM automatically provisions and de-provisions users and groups to ServiceNow by using the providers' provisioning service with Basic Authentication.

Before you begin

Role required: scim_admin

Warning: Grant this role carefully. The scim_admin role is equivalent to giving the user the admin role, where the scmin_admin can add or update Personally Identifiable Information (PII).

You must activate the SCIM plugin.

Procedure

1. Navigate to **All** > **System Web Services** > **REST API Access Policies** to check the details on the REST API Access Policies.
2. In the API Access Policy page, click the **SCIM API Policy** record.
3. Verify the **SCIM API Basic Auth** record is available in the Authentication Profiles sections.
4. Select **Basic Auth** as the Type.
5. Create the required configurations at the provider's end.

6. Test the connection to ensure that the provider can connect to ServiceNow.

> **Note:** If the connection fails, ensure that your ServiceNow account has admin permissions and try again.

Provision user using OAuth

Configure the provider for SCIM automatically provisions and de-provisioning of users and groups to ServiceNow by using the providers provisioning service with OAuth.

Before you begin

Role required: scim_admin

> **Warning:** Grant this role carefully. The scim_admin role is equivalent to giving the user the admin role, where the scmin_admin can add or update Personally Identifiable Information (PII).

You must activate the SCIM plugin.

Procedure

1. Navigate to **All** > **System OAuth** > **Application Registry**.
2. In the Application Registries page, click the **SCIM API** record.
3. Verify the **SCIM API** record details.

   These details should be provided while configuring the ServiceNow application on Azure AD.
4. Navigate to **All** > **System Web Services** > **REST API Access Policies** to check the details on the REST API Access Policies.
5. In the API Access Policy page, click the **SCIM API Policy** record.
6. Verify the **SCIMAPIOAuthOnly** record is available in the Authentication Profiles sections.
7. Check if the **OAuth Entity** field is specified with **SCIM API** record that was earlier configured or verified as application registry.
8. Create the required configurations at the provider's end.
9. Test the connection to ensure that the provider can connect to ServiceNow.

> **Note:** If the connection fails, ensure that your ServiceNow account has admin permissions and try again.

SCIM Troubleshooting

Common error scenarios integrating with SCIM.

Invalid Rest API URL

Action: Enter valid API URL. Would be able to cross check the REST API URL in the **REST API Explorer**.

No Redirect URL is set in ServiceNow instance

Action: Enter valid Redirect URL for SCIM OAuth entity in ServiceNow. Enter redirect URL while configure OAuth entity in ServiceNow.

When Redirect URL is different than the request

Action: '**redirect_url**' provided in '**Authorization Endpoint'** should be same as the OAuth entity configured in ServiceNow.

Note:

This error occurs when there is a mismatch between Azure 'Authorization Endpoint' and ServiceNow 'Redirect URL'

When an invalid client secret is passed

Action: Value entered in 'Client Secret' should be same as the OAuth entity configured in ServiceNow.

When an invalid 'ClientId' passed in Azure 'Authorization EndPoint'

Action: Value entered for 'client_id' parameter in 'Authorization Endpoint' has to be same as OAuth entity configured in ServiceNow.

SCIM customization

Customize SCIM protocols for your identity management.

SCIM customization enables you to do the following:

- Support custom fields on the sys_user and sys_user_group tables through dynamic extension schema generation.
- Provide an ability to override the default SCIM mappings.

A SCIM admin can define custom extension schemas for user and group resources. The attributes defined in the custom extension schema can be mapped to fields in the sys_user or sys_user_group tables.

Configurations for SCIM customization

For SCIM customization, you should perform the following tasks:

- Define a custom extension schema for users and groups in the SCIM Extension schema table. #For more information see, Create a SCIM Extension schema.
- Create entities in an ETL definition for custom schema attributes. The entities are created for the target table that is mapped with either the sys_user or sys_user_group attributes. For more information see, Create a SCIM ETL definition.
- Create an RTE mapping between these two entities.#For more information see, step 5 in the Create a SCIM ETL definition.
- Send custom schema attributes with data in the SCIM API request payload.

The SCIM API calls the RTE engine with the defined mapping. The data is stored in the respective fields in the target table as defined in the mapping.

SCIM customization properties and schemas

The SCIM customization includes the following properties, supported schemas, and unsupported schemas.

Properties

SCIM customization adds the following system properties.

Properties

Name	Description
com.snc.integration.scim2.max.member	The SCIM maximum member count.
com.snc.integration.scim2.resolve	Resolve SCIM resources based on the source definition of the requesting client if multiple resources are found with an external ID SCIM filter. **Note:** As a prerequisite, if provisioning is done from multiple sources, then the sources must be defined in the SCIM Source Definition table. If sources are not defined or if this property is inactive, then all matching resources are returned with an external Id filter response.
com.snc.integration.scim2.user.etl	The SCIM User ETL definition ID.
com.snc.integration.scim2.group.etl	The SCIM Group ETL definition ID.
com.snc.integration.scim2.rte.verbose	Enables verbose logging for SCIM User and Group RTE definitions.
com.snc.integration.scim2.string.field.validation	The length validation for fields. This property enables validation instead of truncating and saving the field.
com.snc.integration.scim2.provider	The ID of the script include for customizing the SCIM responses.

Supported schemas

SCIM customization adds the following supported schemas.

Supported Schemas

Schemas	Description	Prefix	Example
urn:ietf:params:scim:schemas:core:2.0:User	Includes core attributes for the resources.		name.middleName
urn:ietf:params:scim:schemas:extension:servicenow:2.0:User	Includes attributes that are related to ServiceNow.	*servicenow*	servicenow.manager.value
urn:ietf:params:scim:schemas:custom:2.0:User	Includes custom attributes that are not mapped as part of the core extension or the ServiceNow extension schema.	*custom*	custom.social

Unsupported schema

urn:ietf:params:scim:schemas:extension:enterprise:2.0:User: Includes attributes commonly used in representing users that belong to or act on behalf of a business or an enterprise.

Note: Enterprise schema is a valid schema but its attributes are mapped to any table. Because database persistence is not supported, there will be no error displayed if an enterprise schema is included in the request body.

Create a SCIM Extension schema

Create custom attributes to map to fields that are not mapped as part of either the core schema or the ServiceNow extension schema.

Before you begin
Role required: scim_admin

Warning: Grant this role carefully. The scim_admin role is equivalent to giving the user the admin role, where the scmin_admin role can insert new records into the tables that can bypass business logic or ACL protection.

Procedure

1. Navigate to **All** > **SCIM** > **SCIM Extension schemas**.
2. Click **New**.
3. On the form, fill in the fields.

Note: Only one extension schema can be mapped to the **Resource Type** field. For example, User as a resource type can be mapped to a user extension schema.

SCIM Extension schema

Field	Description
Name	Name of the extension schema.
Active	Option to activate the schema. Select this field if the record must be considered as custom extension schema. Only one custom extension schema record can be active at a time for a specific resource type, whether for the User or Group type.
Resource Type	Resource type that has to be mapped to the extension schema. The following are the options: ○ **User** ○ **Group**
Application	Application scope for this record.
Schema JSON	Details within the JSON schemas. For more information about defining the extension schema with attributes, see Datatracker.

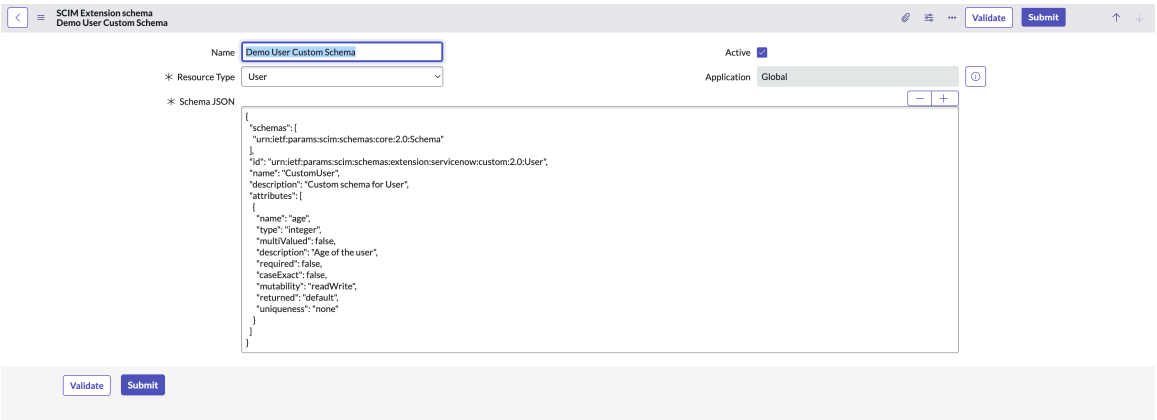

4. Validate the attributes by clicking **Validate**.
5. Click **Submit**.

Result

The extension schema with custom attributes related to User or Group resource type is created. Use the SCIM ETL Definitions to map the resources based on the extension schema on the sys_user and sys_user_group table. For more information, see Create a SCIM ETL definition.

Create a SCIM ETL definition

Use the SCIM ETL definitions to map the custom attributes with the sys_user or sys_user_group tables.

Before you begin

Role required: scim_admin

Warning: Grant this role carefully. The scim_admin role is equivalent to giving the user the admin role, where the scmin_admin role can insert new records into the tables that can bypass business logic or ACL protection.

Note:

- SCIM Group and SCIM User ETL definitions are part of the base system for resource mapping. You can use the same resource mappings and change the criteria as required, or you can create new resource mappings.
- There is no support for [*] fields through RTE in SCIM mapping.

Procedure

1. Follow the instructions in Create Extract Transform Load (ETL) definitions.
2. Open the newly created record and view the details.
3. In the ETL Entities section, create an entity by clicking **New**.
   You have to create entities for the following users:
   - scim-user: For the fields that are from SCIM.
   - user (sys_user) or group (sys_user_group) table: For the fields that you want to map from the database table with SCIM. For example, for customization of user details through SCIM, you can use the sys_user table.

a. On the form, fill in the fields.

ETL Entity form

Field	Description
Name	Name of the ETL entity.
Application	Application scope for this record.
Table	Target table for the ETL entity.
Definition	Selected ETL entity.
Path	Unique path for this entity. Do not specify any path for the entity representing the import set table. When an entity represents a collection, the path must end with an asterisk [*]. This requirement applies to intermediate entries and the target table entity.
Synchronize Inserts	Option to guarantee only one record with unique coalesced field values by synchronizing record inserts.
Run business rules	Option to run business rules.

b. Click **Submit**.

The ETL entities are created for scim-user and user tables. You have to create ETL entity fields within these ETL entities, and map both entities by creating an RTE Entity Mapping.

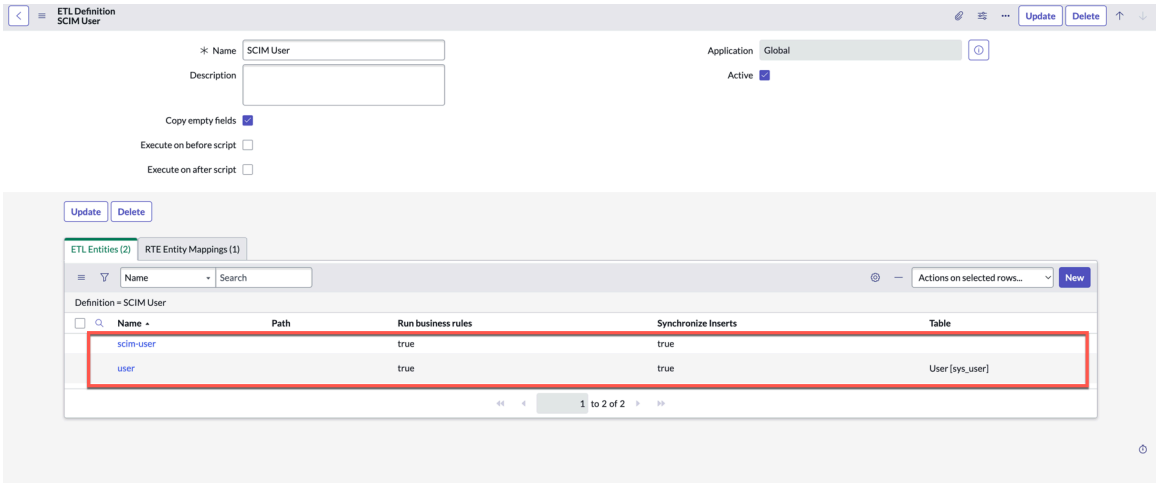

4. Create the entities and map them.

a. Open the submitted record (scim-user and user).

b. In the ETL Entity fields, add the fields by clicking **New**.

c. On the form, fill in the fields.

ETL Entity form

Field	Description
Name	Name of the ETL Entity field definition.
Application	Selected ETL entity that this field definition belongs to.
Field/Path	This field is either a column or path. ▪ The field is a column name when the entity is the Import or Target table. ▪ The field is a path when the field has nested structures.
Entity	Entity that this operation applies to. Choose the entity using the look-up icon.
Coercion action	What the system should do if a reference or choice could not be found. Options are as follows: ▪ **Create**: Create a new reference or choice. Assign the reference or choice to the current record. ▪ **Reject**: Do not save the whole record to the database. ▪ **Ignore**: Set the current value as empty.
Definition	Selected ETL entity that this field definition belongs to.
Coalesce	Option to query the existing records.

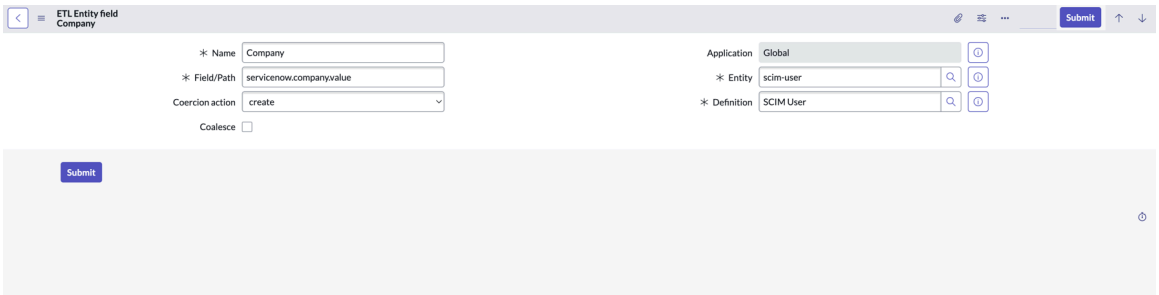

d. Submit the entries by clicking **Submit**.

You can add multiple entries as your ETL Entity field based on your requirement.

The scim-user ETL entity can have entities with the core extension (user), the ServiceNow extension, or custom extension.

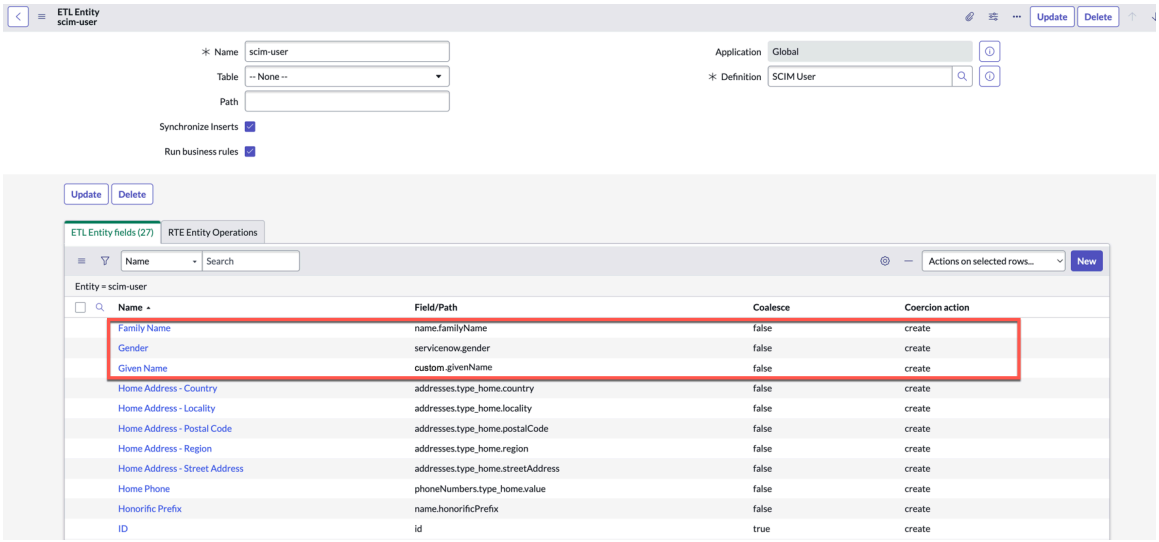

The user ETL entity can have entities from the database table. For example, sys_user table.

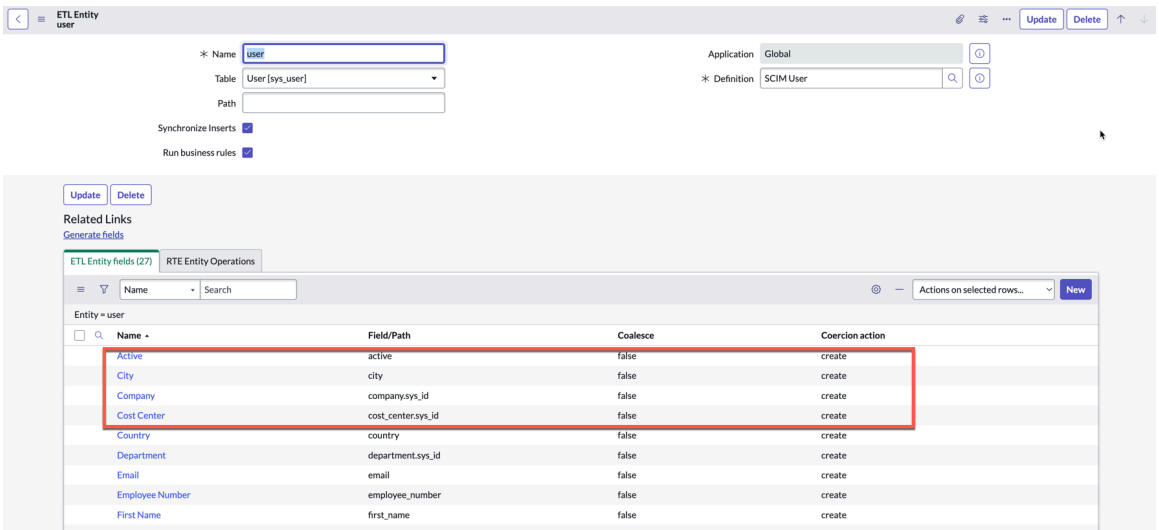

Note: To add a filter in the incoming SCIM fields, use an underscore (_). This underscore translates to an EQ filter. For example, the attribute *email.type_work.value* applies the SCIM filter of email *[type eq "work"].value*.

After the fields are created in both the scim-user and user ETL Entity records, you have to create an RTE Entity mapping record. You must then specify the source and target definition to map both the fields.

5. In the RTE Entity Mappings section, create an entity mapping by clicking **New**
6. On the form, fill the fields.

Field	Description
Name	Name of the mapping.
Source Entity	Source entity for the mapping.

Field	Description
Target Entity	Target entity for the mapping.
Order	Order in which the mapping should be processed.
Is Conditional	Option to designate the mapping as conditional.
Condition Script	Script that defines the conditions that must be met for the mapping.
Application	Application scope for this record.
Definition	Selected ETL entity that this mapping belongs to.
Ignore	Option to designate if this ETL entity mapping should be ignored when performing data integrations using Robust Import Set Transformation (RTE).

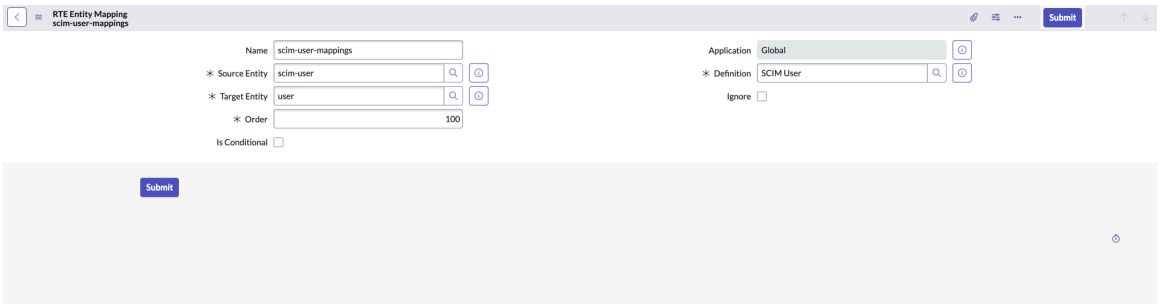

7. Click **Submit**.

   The following example shows a record created for mapping both the scim-user and user ETL Entity records.

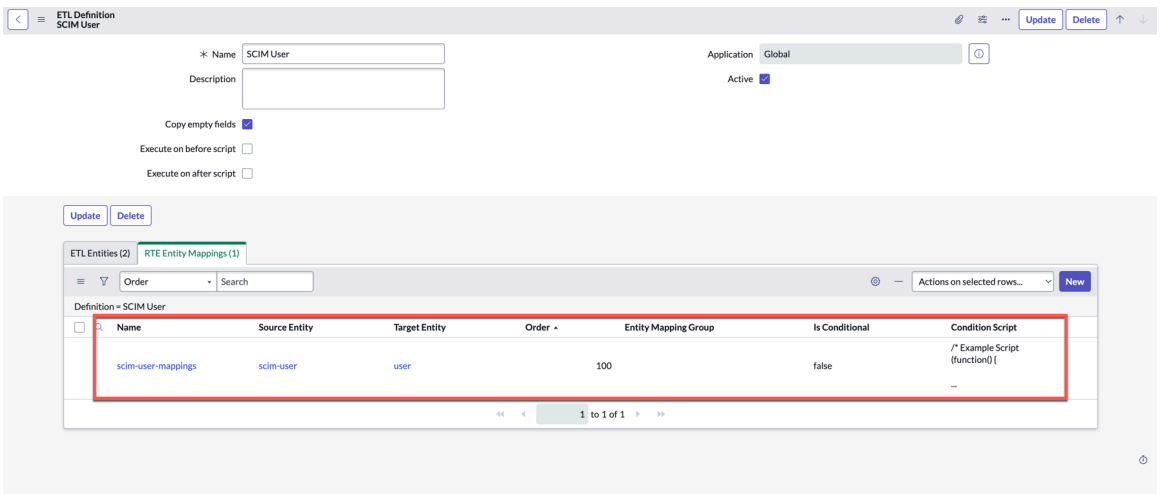

8. Open the submitted record (scim-user-mappings) and create a mapping between the scim-user and user ETL Entity records.

   a. In the RTE Field Mappings section, click **New**.

   b. On the form, fill in the fields.

ETL Entity fields

Field	Description
Source Field	Application scope for this record.
Application	Selected ETL entity that this field definition belongs to.
Target Field	Output ETL field for the operation if the operation takes a single output.
Entity Mapping	Entity mapping that this operation applies to.
Referenced Entity	Entity that is referenced and the operation that it applies to.
Definition	Selected ETL entity that this field definition belongs to.
Order	Order in which the operation runs on the entity.

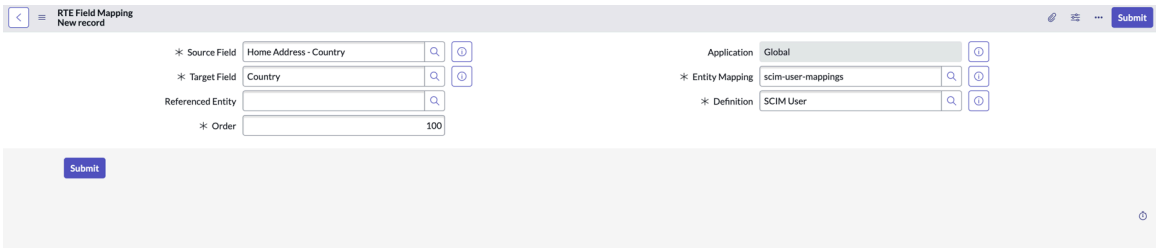

The source field **Home Address - Country** (scim-user ETL Entity) maps the target field as **Country** (user ETL Entity).

c. Submit the entries by clicking **Submit**.

You can add multiple entries as your RTE Entity mappings based on your requirement.

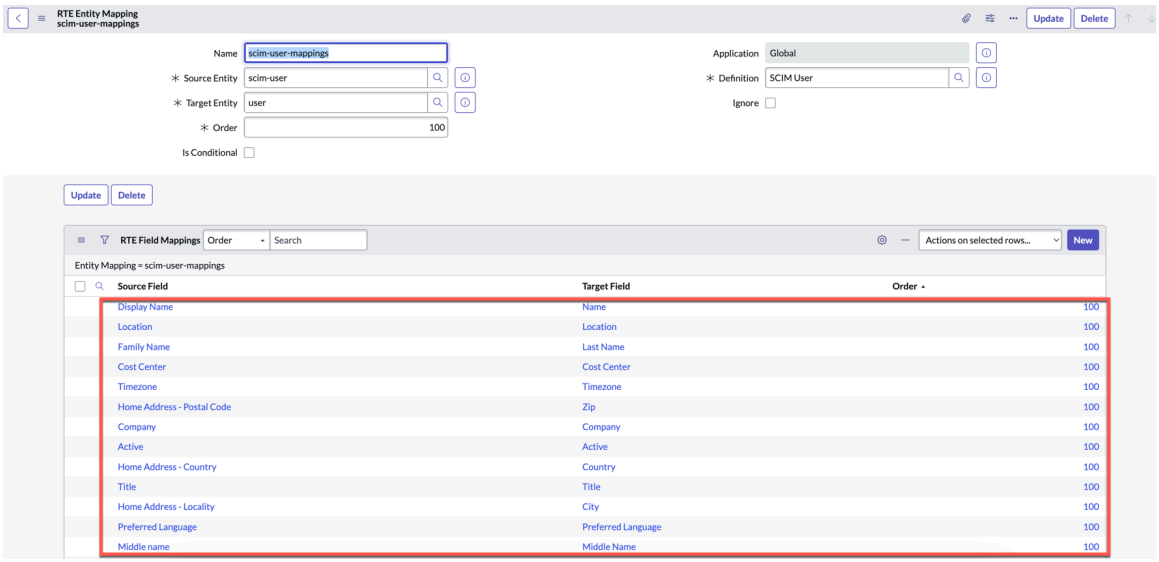

The source fields and targets fields are mapped as configured. When you perform CRUD operations using SCIM, the customized values are updated in the respective table.

Result

These ETL definitions and mappings enable you to extract data from a source table, transform the data as desired, and load the data into a target table.

Related topics

Create Extract Transform Load (ETL) definitions

Handle unmapped fields

You can handle unmapped fields in SCIM customization in different ways.

During SCIM customization, the fields that are not part of the sys_user and sys_user_group tables can be mapped by performing the following functions.

Customize SCIM (Create or Update)

You can create or update the SCIM Client.

- The SCIM admin can add custom scripts in the *onBefore* and *onAfter* scripts for fields that are not mapped in ETL Definition or RTE.
- The SCIM admin can override RTE Mappings by adding custom scripts in the *onBefore* and *onAfter* scripts.
- You can invoke a scriptable API in the RTE *onBefore* or *onAfter* scripts to access incoming request and perform transformations on other tables, lists, and unmapped attributes.
- You can use the *sn_auth.SCIM2Util.getScimProviderCustomizationContext()* method to provide the SCIM request context that contains the *scimResource* object. The *scimResource* in context represents the following in each operation:
  - **POST**: The SCIM resource sent in the request payload.
  - **PUT**: The current SCIM resource from database replaced with the SCIM resource sent in thee request payload.
  - **PATCH**: The current SCIM resource from the database after performing the patch operations.

The following is an example of an *onAfter* script.

```
(function onAfter(source, target, importLog) {

    var ctx =
sn_auth.SCIM2Util.getScimProviderCustomizationContext();
    gs.info("scim context ee: " +
JSON.stringify(ctx.scimResource));

    var roles = ctx.scimResource.roles;
    if(roles) {
        var removingRolesGR = new
GlideRecord("sys_user_has_role");
        removingRolesGR.addQuery("user",
target.sys_user[0].sys_id);
        removingRolesGR.query();
        removingRolesGR.deleteMultiple();

    for (var i = 0; i < roles.length; i++) {
        var addingRolesGR = new
GlideRecord("sys_user_has_role");
```

```
        addingRolesGR.setValue("user",
 target.sys_user[0].sys_id);
        addingRolesGR.setValue("role", roles[i].value);
        addingRolesGR.setValue("state", "active");
        addingRolesGR.insert();
        }
        }
        var customUserExtn = new
 global.SCIMProviderCustomization().getCustomExtensionUrn('User
');
        var salary = ctx.scimResource[customUserExtn].salary;
        if (salary) {
            var gr = new GlideRecord("u_user_salary");
            gr.addQuery("user", target.sys_user[0].sys_id);
            gr.query();
            if (gr.next()) {
                gr.setValue("salary", salary);
                gs.info("scim update: " + gr.update());
            } else {
            gr.setValue("salary", salary);
            gr.setValue("user", target.sys_user[0].sys_id);
            gr.insert();
            }
           }

})(source, target, importLog);
```

Customize SCIM response

For the GET API calls, any response back to the SCIM client can be customized using the script by extending the *SCIMProviderCustomization* script.

While extending the script, the author can override the *customizeUserResponse* and *customizeGroupResponse* methods to modify the responses for User and Group resources.

The *com.snc.integration.scim2.provider.customization.script.id* property enables the SCIM plugin to use the script that should be used for response customization.

The following is an example of extending the base script.

```
var SCIMCustomizationScript = Class.create();
SCIMCustomizationScript.prototype =
 Object.extendsObject(SCIMProviderCustomization, {
    initialize: function() {

 SCIMProviderCustomization.prototype.initialize.call(this);
    },
    customizeUserResponse: function(context) {
        try {
            var rolesGR = new GlideRecord("sys_user_has_role");
            rolesGR.addQuery("user", context.scimResource.id);
            rolesGR.query();
            var i = 0;
            context.scimResource.roles = [];
            while (rolesGR.next()) {
                context.scimResource.roles[i] = {
```

```
                    display:
 rolesGR.getElement('role.name').getValue(),
                    value:
 rolesGR.getElement('role.sys_id').getValue()
                };
                i++;
            }
            var userGR = new GlideRecord("u_user_salary");
            userGR.addQuery("user", context.scimResource.id);
            userGR.query();
            if (userGR.next()) {
                var salary = userGR.getValue("salary");
                if (salary) {
                    var customExtensionValue =
 SCIMProviderCustomization.prototype.getCustomExtensionNodeValue
.call(this, "User", context);
                    customExtensionValue.salary = salary;

 SCIMProviderCustomization.prototype.setCustomExtensionNodeValue
.call(this, "User", context, customExtensionValue);
                }
            }
        } catch (ex) {
            gs.error("err: " + ex);
        }
        return context;
    },
    customizeGroupResponse: function(context) {
        return context;
    },
    type: 'SCIMCustomizationScript'
});
```

Note:

- The parameter that the *customizeUserResponse* and *customizeGroupResponse* methods contain is a context object with one attribute called *scimResource*. This attribute has all the attributes of a user or group resource object.
- A customized script include can only be created and viewed by the admin.
- If a user or group resource is modified, then you must return the context back.
- If there are no modification of any attribute in the resource object, then set the *com.snc.integration.scim2.provider.customization.script.id* to empty or return as null.
- If certain attributes are persisted through the *onAfter* script, they must be populated with database values in the *scimResource* object inside the customized script. This action is required so that the system can do the following:
  - To get the correct *scimResource* object in *onAfter* scripts during the PUT and PATCH operation.
  - To include the attributes that persisted through the *onAfter* script in the response back to the client.

Helper functions

The following are some of the helper functions for SCIM customization. These functions enable you to fetch or set different types of information.

Helper functions

Function	Purpose
`SCIMProviderCustomization.prototype.getCustomExtensionUrn.call(this, "User");`	Fetch the value of the custom extension schema.
`SCIMProviderCustomization.prototype.getServiceNowExtensionUrn.call(this, "User");`	Fetch the value of the ServiceNow extension schema.
`SCIMProviderCustomization.prototype.getCustomExtensionNodeValue.call(this, "User", context);`	Fetch the custom extension node value from the response.
`SCIMProviderCustomization.prototype.getServiceNowExtensionNodeValue.call(this, "User", context);`	Fetch the ServiceNow extension node value from the response
`SCIMProviderCustomization.prototype.setCustomExtensionNodeValue.call(this, "User", context, customExtensionValue);`	Set the custom extension node value in the response.

The following is an example of using the helper function:

```
var customExtensionUrn =
SCIMProviderCustomization.prototype.getCustomExtensionUrn.call(t
his, "User");
var customExtensionValue =
SCIMProviderCustomization.prototype.getCustomExtensionNodeValue.
call(this, "User", context);
customExtensionValue.age = "18";
SCIMProviderCustomization.prototype.setCustomExtensionNodeValue
 call(this, "User", context, customExtensionValue);
```

Note: RTE supports setting data in tables other than the sys_user and sys_user_group tables.

Creating a source definition

Create a source definition to capture information about which identity source a resource is provisioned from.

Before you begin

Role required: scim_admin

Warning: Grant this role carefully. The scim_admin role is equivalent to giving the user the admin role, where the scmin_admin can add or update Personally Identifiable Information (PII).

About this task

Using a source definition, the provisioning identity source can be mapped to an OAuth entity using which it authenticates while provisioning.

After a source definition is created, all resources getting provisioned from that identity source is mapped to its corresponding source definition ID.

The source definition captures the required source information, such as by doing the following:

- Identifies the SCIM Client from which the resource is provisioned.
- Resolves duplicate information provided by the external ID:
  - If multiple identity sources are provisioning resources, there can be two or more resources have the same external ID value because the external Id is only unique to the identity source.
  - If multiple resources are returned with an external ID SCIM filter then the resources can be resolved based on the source definition of the requesting identity source.

Note: A source definition can be created only for the identity source, which, uses an OAuth authentication method.

Procedure

1. Navigate to **All** > **SCIM** > **Source Definition**.
2. On the SCIM Source Definitions page, click **New**.
3. On the form, fill in the fields.

SCIM Source Definition

Field	Description
Name	Name of the source definition.
Application	Application scope for this record.
OAuth Entity	The OAuth entity of the integration user. The entity is used for provisioning the user by the identity source provider.
Identity Source	The name of the identity source provider. For example, Azure AD, Okta, and so on

Note: Enable the *com.snc.integration.scim2.resolve.externalid.conflict* property to return only SCIM resources created by the requesting identity source. By default, all the matching resources with an external ID filter are returned.

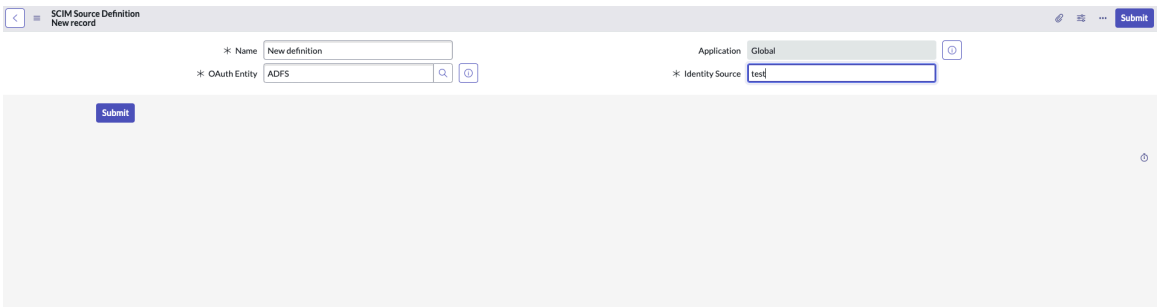

4. Click **Submit**.

Result

The SCIM source definition is created. Use the SCIM ETL Definitions to map the resources based on the extension schema on the sys_user and sys_user_group table. For more information, see Create a SCIM ETL definition.

SCIM Client

The SCIM Client facilitates provisioning and updates on identity resources through CRUD operations exposed by SCIM endpoint on an external system.

Explore  Learn about SCIM Client.	Activate  Activate SCIM Client.
SCIM Client Properties  Get details about how to customize SCIM.	Troubleshoot  Learn more about source definition for SCIM.

Explore SCIM Client

The SCIM Client facilitates provisioning and updates on identity resources through CRUD operations exposed by SCIM endpoint on an external system.

The SCIM Client is used for creating, updating, and deleting identity resources in a system that supports SCIM compliant REST requests. The client is used for identity life-cycle management and for identity attribute synchronization across ServiceNow instances or between ServiceNow and other SCIM providers.

Because the APIs are exposed by the client, you can automate the process of creating, updating, or deleting any resources on a single or multiple SCIM providers. For example, if a developer joins the organization, you have to provide access to Git, workplace, and so on.

The SCIM client enables you to perform the following actions:

- Provision identities and access for user or group membership.
- Synchronize identity and related resources with SCIM compliant systems.
- Integrate any SCIM Provider on ServiceNow®.
- De-provision identities and access.

The SCIM client provides scriptable APIs that the integration developer can use to build workflows or automations to do specified jobs. To know more about the scriptable API, see SCIM2Client API.

Configurations for SCIM Client

To configure the SCIM Client, perform the following tasks:

- Create a REST message for all outbound calls for a particular SCIM Provider. For more information, see Create a REST message.
- Create a SCIM Provider to fetch resource types and schemas information from the SCIM Provider with the REST message. Enable the configuration of the HTTP Method (PUT or PATCH) to update a resource in the SCIM Provider. For more information, see Create a SCIM Provider.
- Create the mappings of SCIM attributes to ServiceNow attributes for a particular resource type and SCIM Provider. To know more, see Create a SCIM Provider Resource Mapping.
- Perform the mapping of SCIM field with the database table and field name. Pass the default value or write a script to fetch the value. For more information know, see Create a SCIM attribute mapping.

Activate the SCIM Client plugin

For SCIM Client activation, install the SCIM v2 - ServiceNow Cross-domain Identity Management Client (com.snc.integration.scim2.client) plugin.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the SCIM v2 - ServiceNow Cross-domain Identity Management Client (com.snc.integration.scim2.client) plugin using the filter criteria and search bar.

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

   Note: When domain separation and delegated Admin are enabled in an instance, the administrative user must be in the **global** domain. Otherwise, the following error appears: `Application installation is unavailable because another operation is running: Plugin Activation for <plugin name>.`

   You will see a message after installation is completed. For information about the components installed with a plugin, see Find components installed with an application.

SCIM Client properties, tables, scriptable APIs, and logs

The SCIM v2 - ServiceNow Cross-domain Identity Management Client (com.snc.integration.scim2.client) plugin includes the following system properties, tables, scriptable APIs, and logs.

Properties

SCIM Client adds the following system properties.

Properties

Name	Description
com.snc.integration.scim2.client.log.property	This property determines whether to write all the log records or just the error log records. The possible values are **FAILURE** or **ALL**. Default value: **ALL**
com.snc.integration.scim2.client.log.purge.duration	This property determines the number of days for clearing the logs. Default value: **180**

To set the properties, navigate to **All** > **SCIM** > **SCIM Client Properties**.

Tables

The SCIM Client adds the following tables.

Tables

Name	Description
SCIM Provider (sys_scim_provider)	Stores data for each SCIM provider, such as the name, REST message resource definitions, and so on.
SCIM Provider Resource Mapping (sys_scim_provider_resource_mapping)	Stores the primary table information for each provider and resource name.
SCIM Attribute Mapping (sys_scim_attribute_mapping)	Stores the source details where each SCIM attribute value should come from, such as the table field, script, and so on.
SCIM Client Log (sys_scim_client_logs)	Stores the statuses of each call triggered to SCIM Provider.

Scriptable API

The SCIM2Client API calls the System for Cross-domain Identity Management (SCIM) Provider (server role) to create, update, or delete data in a service provider (SP). The scriptable API of the SCIM Client should be used in the scripts that are running in the system context or by a system admin user.

For example, you can use the script while running the integration hub workflow as a system user, while running the scheduled jobs, and so on.

The following are some of the use cases for using the scriptable APIs:

- As an admin, provision identity information from background scripts, business rules, script include calls, workflows, and so on.
- As an admin, run a scheduled job or an on-demand job for identity provisioning.
- Run a workflow or sub-workflow with the Script step using the provision scriptable API call.
- Add the provision script directly in a business rule or script include. The script can be triggered by non-admin users. This use-case works in the following situations:

- The user has access to the token, meaning that the user has the role to generate the token from the REST template.
- The user has access to retrieve the SCIM attribute values from the mapped tables.

To know more about the scriptable API, see SCIM2Client API.

SCIM Client Logs

The SCIM Client Logs display the provisioning status about the SCIM APIs. To view the provisioning status, navigate to **All** > **SCIM** > **SCIM Client Logs**.

Create a REST message

Configure a REST message for all outbound calls for a particular SCIM Provider.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Web Services** > **Outbound** > **REST Message**.
2. Click **New**.
3. On the form, fill in the fields.

SCIM Provider

Field	Description
Name	Descriptive name for this message.
Endpoint	Base URL of SCIM Provider. For example, `https://example.service-now.com/api/now/scim`.
Authentication Type	Type of authentication to be used to connect to the external SCIM Provider. For more information see, Outbound REST authentication.
HTTP Headers	The content type that is expected from the external SCIM Provider. For example, Header Name is the content type in the following API request body: `Content-type: application/scim+json`

4. In the#HTTP Methods#elated list, click#**New**.
5. Configure the following HTTP Methods.

Note: The URLs have variables that you must replace.

HTTP Methods

Methods	Sample URL
GET	`https://example.com/api/now/scim/${resourceName}`
PATCH	`https://example.com/api/now/scim/${resourceName}/${resourceId}`
PUT	`https://example.com/api/now/scim/${resourceName}/${resourceId}`
DELETE	`https://example.com/api/now/scim/${resourceName}/${resourceId}`
POST	`https://example.com/api/now/scim/${resourceName}/${resourceId}`

Note:

- You must create all the HTTP methods for the operation of the SCIM Client.
- A sample REST message is shipped from the base system.

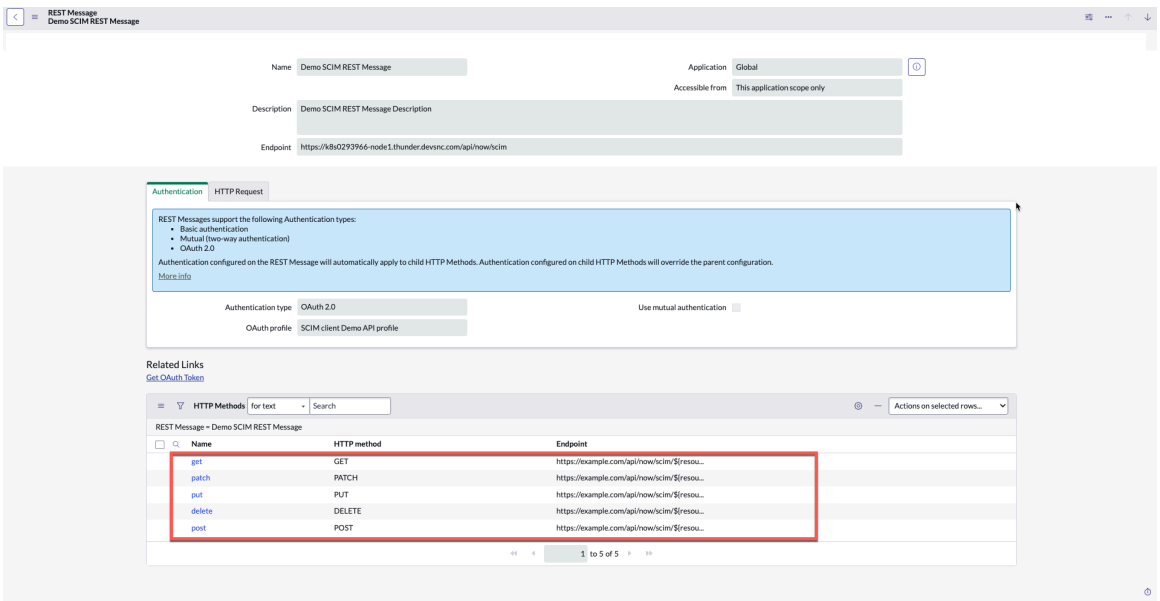

6. Click **Submit**.

Result
The REST message record is created.

What to do next

Use the REST message to create a SCIM Provider. For more information, see Create a SCIM Provider.

To learn more about how to create a REST message, see Create a REST message.

Create a SCIM Provider

Create a SCIM Provider to fetch resource types and schemas information from the SCIM Provider with the REST message. Enable the configuration of the HTTP Method (PUT or PATCH) to update a resource in the SCIM Provider.

Before you begin

Note: A sample SCIM Provider is part of the base system. You can use and configure based on your requirement, or create a new record.

Roles required: admin

Procedure

1. Navigate to **All** > **SCIM Client** > **SCIM Provider**.
2. In the SCIM Providers page, click **New**.
3. On the form, fill in the fields.

SCIM Provider form

Field	Description
Name	Name of the SCIM Provider
Outbound REST Message	Message that is used to call the API of the SCIM Provider. For more information, see Create a REST message.
HTTP Method for Update	Type of HTTPS method that is used for updating the resource mapping. The PATCH or PUT method can be used by the Client while updating the identity resource during the provision of an already existing resource.

Note:

- The **Resource Definition** and **Schema Field Definition** fields are fetched from the */ResourceTypes* and */Schemas* public endpoints of the SCIM Provider. These fields are auto-populated after the rest message is selected.
- If the REST message, Schemas, or the Resource Types are updated on the SCIM Provider, then click **Refresh** and then **Update** to update the **Resource Definition** and **Schema Field Definition** fields.

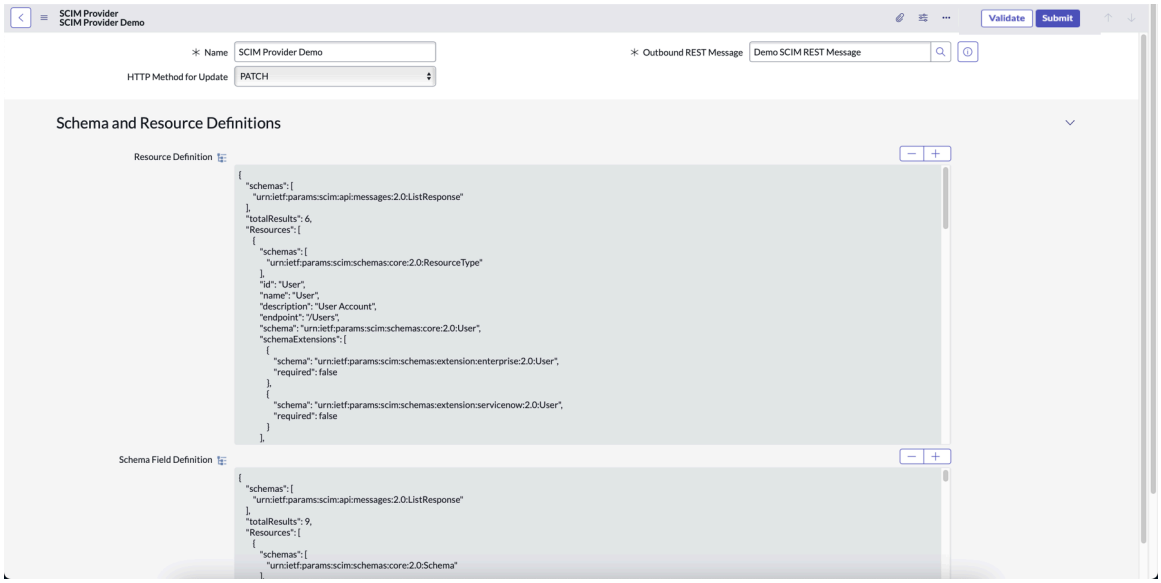

4. Click **Submit**.

Result

The SCIM Provider details are created successfully. Use the SCIM Provider Resource Mapping to map the SCIM Provider details to the resources such as users or groups. For more information, see Create a SCIM Provider Resource Mapping.

Create a SCIM Provider Resource Mapping

Define the mappings of SCIM attributes to ServiceNow attributes for a particular resource type and SCIM Provider.

Before you begin

Roles required: admin

Procedure

1. Navigate to **All** > **SCIM Client** > **SCIM Provider Resource Mapping**.

   The SCIM Provider Resource Mapping is shipped with the User and Group mapping by default.

   Note: The User or Group mappings contains sample mappings, which you can use as a reference. You can also create mapping based on the user or group resources.

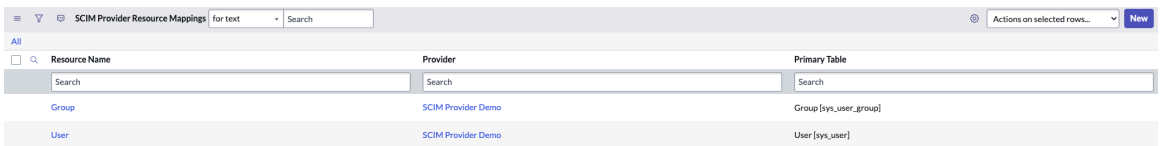

2. Create a resource mapping by clicking **New**.
3. On the form, fill in the fields.

Resource Mapping

Fields	Description
Provider	Name of the SCIM Provider. Refer to the configured Provider name when creating a SCIM Provider.
Resource Name	Resource for which the mapping must be defined.
Primary Table	The table that contains the sys_id of the resource being mapped.

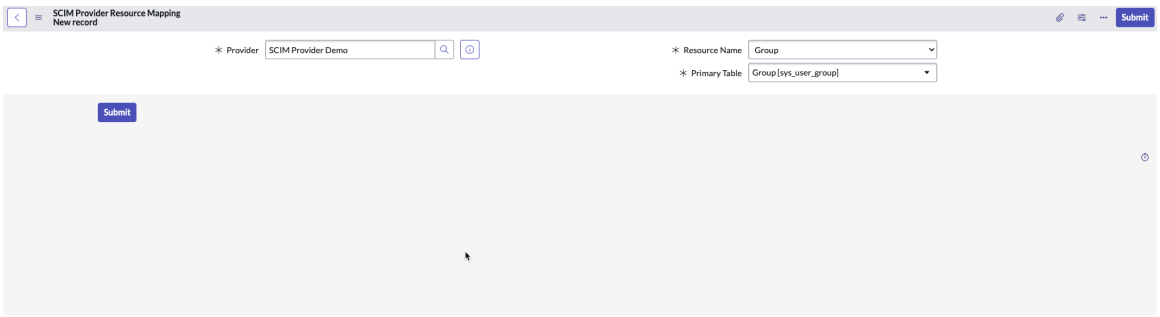

4. Click **Submit**.

Result
The record is created and displayed in the SCIM Provider Resource Mapping page. Use the SCIM Attribute Mappings to further map the attributes from schemas. For more information, see Create a SCIM attribute mapping.

Create a SCIM attribute mapping

Create a SCIM attribute mapping and use it as a single source of resource to the ServiceNow table fields.

Before you begin
Roles required: admin

About this task
The following are the attribute mapping types and their descriptions.

Attribute mapping types

Attribute mapping type	Description
Direct	The SCIM attribute is populated with the help of the **Database Table Name** and **Database Field Name** fields.
Constant	The SCIM attribute is populated with the help of a specified default value.
Script/Custom	The SCIM attribute is populated with the help of the return value of a script. This attribute requires enabling the run script option.

Note:

- The password attribute is not supported
- A sample attribute mapping for the User and Group resources is part of the base system. You can use and configure the mappings based on your requirements, or you can create a new record.

Procedure

1. Navigate to **All** > **SCIM Client** > **SCIM Provider Resource Mapping**.
2. Selected the SCIM Resource Mapping that is created for the configuration.
3. In the#SCIM Attribute Mappings#elated list, click#**New**.

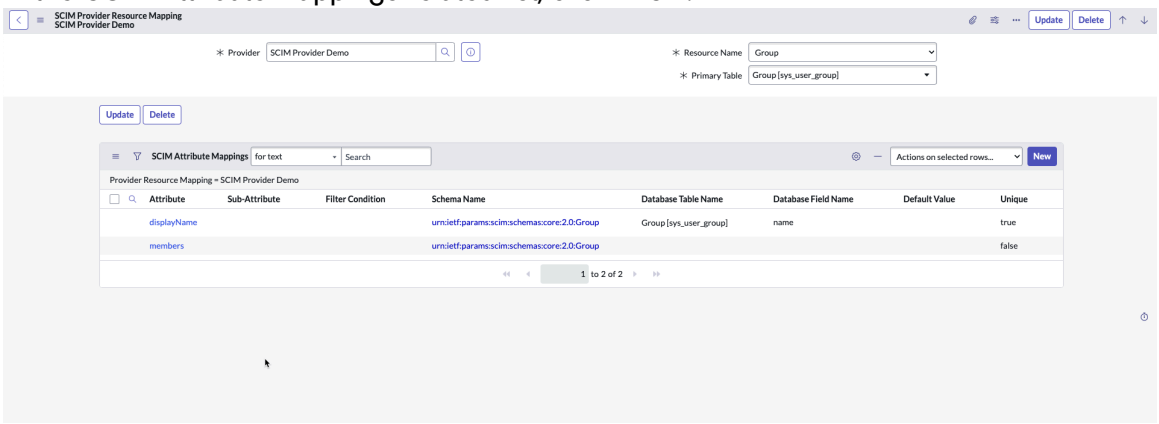

4. On the form, fill in the fields.

SCIM Attribute Mappings form

Fields	Description
Provider Resource Mapping	Attribute mapping for a Provider-resource combination. If this field is not populated automatically, then use the search option. Select the Provider Resource Mapping record.
Schema Name	Schema name of the SCIM attribute for which the mapping must be defined. For example, *urn:ietf:params:scim:schemas:core:2.0:User*.
Attribute	Attribute for which the mapping must be defined. For example, *userName*.
Sub-Attribute	Sub-Attribute (if any) for which the mapping must be defined. For example, if there is a complex type attribute like *name.familyName*, then the

Fields	Description
	attribute is **name** and the sub-attribute is **familyName**.
Field Type	The SCIM Attribute data type. This field is auto-populated from the schemas defined by the SCIM Provider. For example, **boolean**.
Multi-Value	Multiple values that are assigned to the attribute. An attribute can have multiple values like work email, home email, or another email. For example, *`emails`*. This field is populated using the schemas defined by SCIM Provider.
Filter Condition	Some multi-valued attributes can have additional information that can be specified by using a Filter Condition. The choices of filter condition are populated using the schemas defined by the SCIM Provider. For example, the *`phoneNumbers`* attribute has multiple types like work, mobile, home, and so on
Unique	Option to uniquely identify a resource between the SCIM Client and SCIM Provider systems. Multi-valued attributes cannot be marked as unique. For example, for a User Resource, the *`username`* attribute can be made unique.
Database Table Name	Use this field to map the attribute table name to the ServiceNow table name. If you choose direct attribute mapping, you must define this field. For example, the *`username`* SCIM Attribute can be mapped to the **User** field in the **Database Table Name** field.
Database Field Name	TheDatabase Field Name field maps the SCIM attribute to the ServiceNow field name to be mapped with the SCIM Attribute. If you choose direct attribute mapping, you must define this field.

Fields	Description
	For example, the *username* SCIM Attribute can be mapped to the **User ID** field in the **Database Field Name** field.
Default Value	Details about the default value that is passed to the SCIM Provider. Can be used if the direct attribute mapping of the field returns null, or if the default can be used to return a hard-coded value. In the case of a hard-coded value, the database table name and field name should be **None**. For example, the primary sub-attribute value of work email can be hard-coded as **true**.
Run Script	Option to fetch the value of the attribute via script. This option is required for the multi-valued attributes that don't contain a filter condition. For a complex type of attribute, a script can supply the value at an attribute or sub-attribute level. For example, the **Members** attribute of group resource has no filter condition. So, the script option should be defined on the parent attribute level of the **Members** attribute.
Script	Script that is used to fetch the attribute value. The return type of script should be a string, or a JSON converted as a string. The output of the script should be in the proper format as expected by the provider for that attribute.

5. Click **Submit**.

Attribute Mapping references

The attribute mappings enables you to use the attributes as a single source of resource to the ServiceNow table fields.

Attribute

The attribute for which mapping needs to be defined. For example, *userName*.

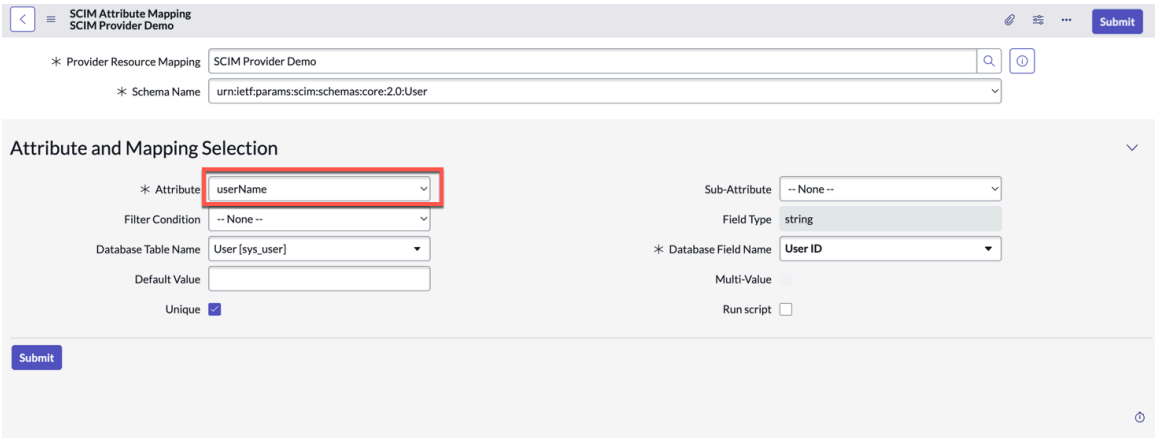

Sub-Attribute

Select the sub-attribute, if any, for which a mapping needs to be defined.

For example, if there is a complex type attribute like *`name.familyName`*, then the attribute is *`name`* and the sub-attribute is *`familyName`*.

For simple attributes like user name, the **Sub-Attribute** value would be **None**.

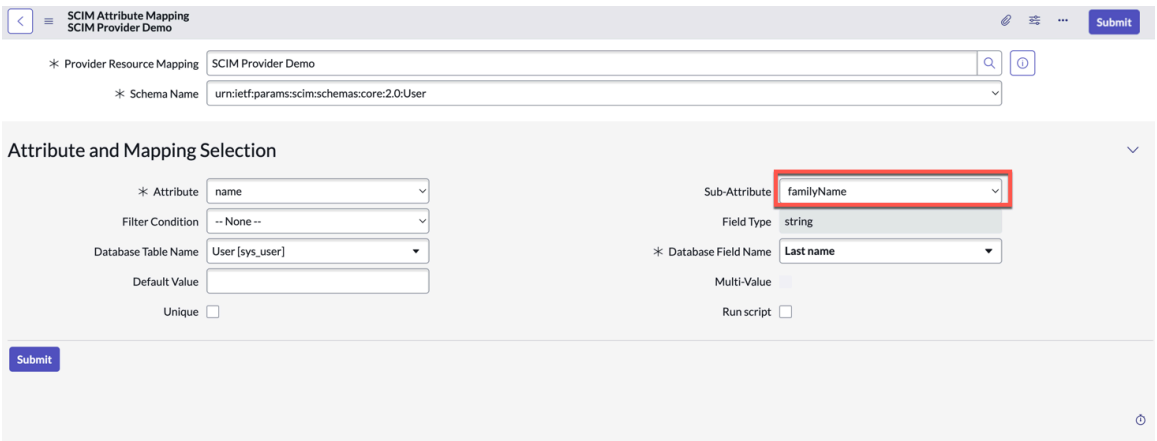

Filter Condition

A multi-valued attribute can have additional information that can be specified by using a Filter Condition. The choices for the filter condition are populated using the schemas defined by the SCIM Provider.

For example, the *`phoneNumbers`* attribute has multiple types like work, mobile, home, and so on.

You can specify a Filter Condition from a set of possible values. For example, the *`phoneNumber`* attribute can have the Filter Condition as **type eq "mobile"**.

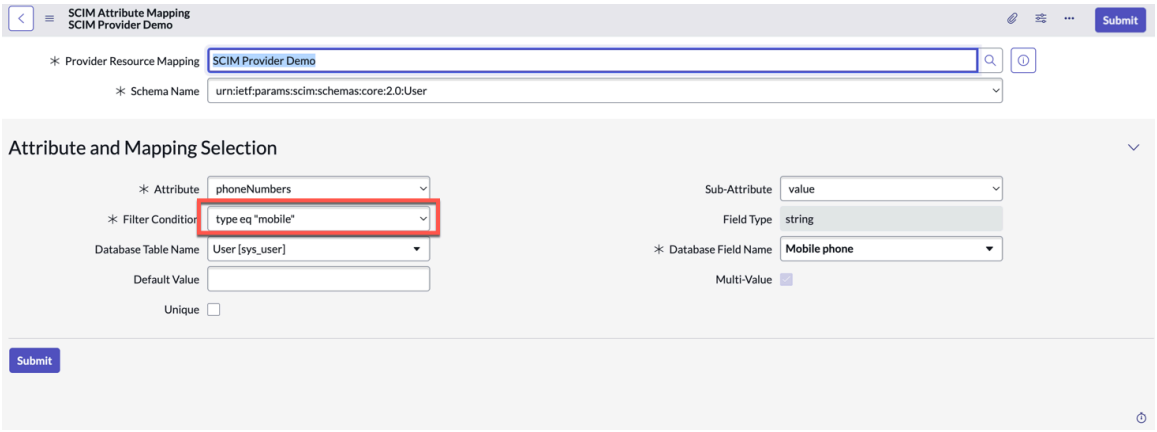

The *phoneNumber* attribute can instead have a Filter Condition as **type eq "work"**.

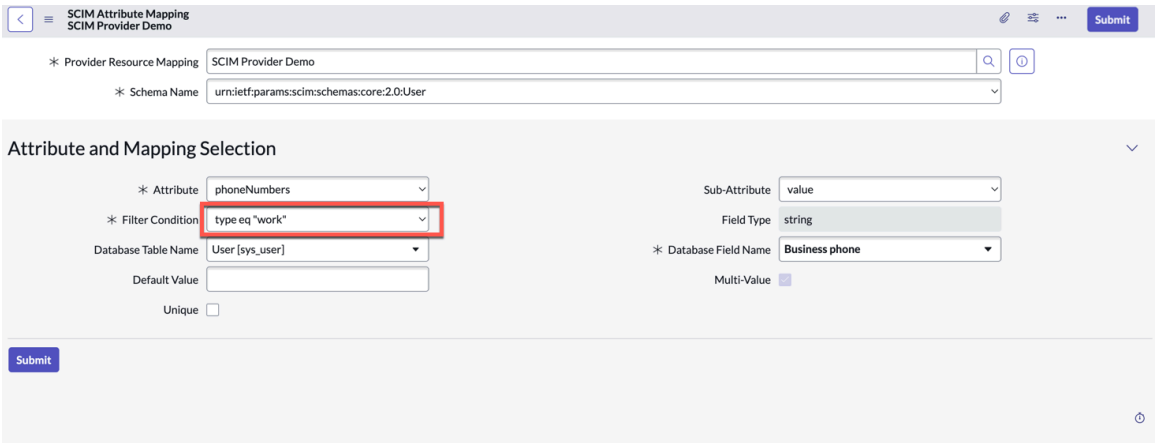

Database Field Name

If the direct attribute mapping option is chosen, then this attribute needs to be defined. The **Database Field Name** field represents the ServiceNow field name that is mapped with the SCIM Attribute.

For example, the *username* SCIM Attribute can be mapped to a user as the **Database Table Name** field, and to the user ID field as the **Database Field Name** field.

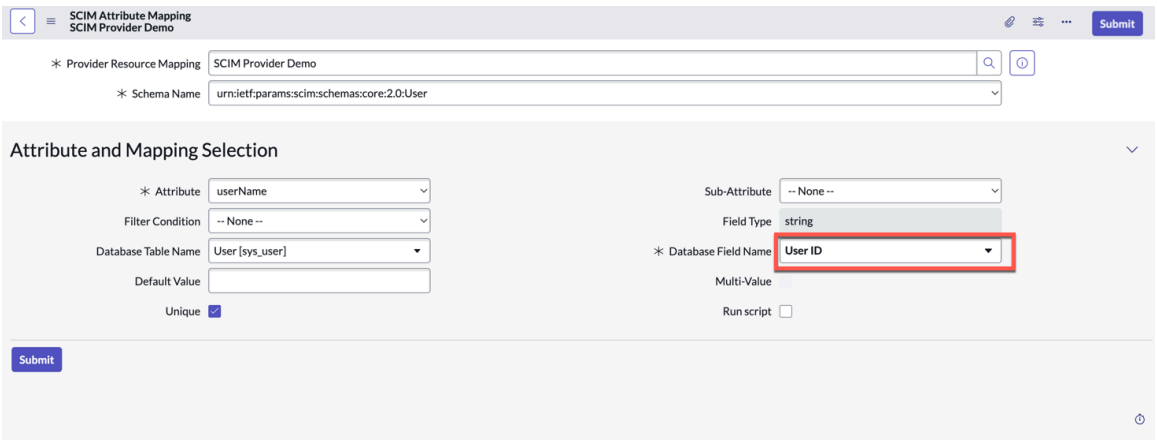

You can also dot-walk using the **Database Field Name**. For example, the **department** SCIM Attribute can be mapped to the **Department Name** field.

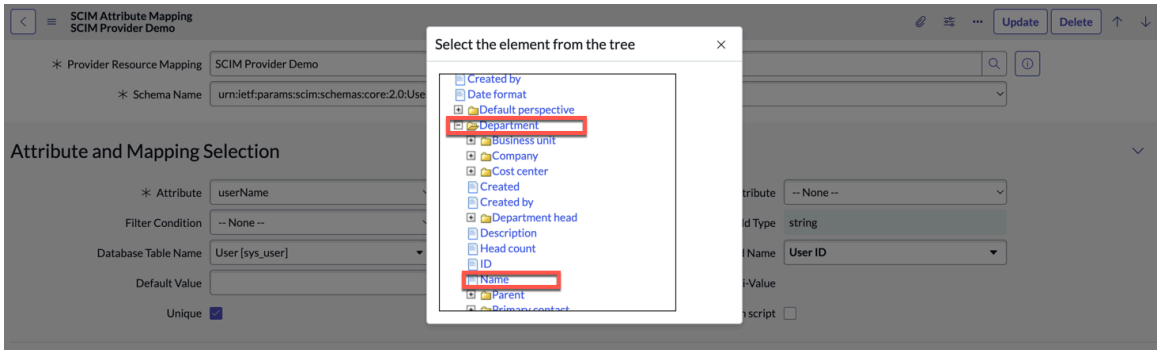

Here the Database Table is **User** and the Database field Name is **Department Name**.

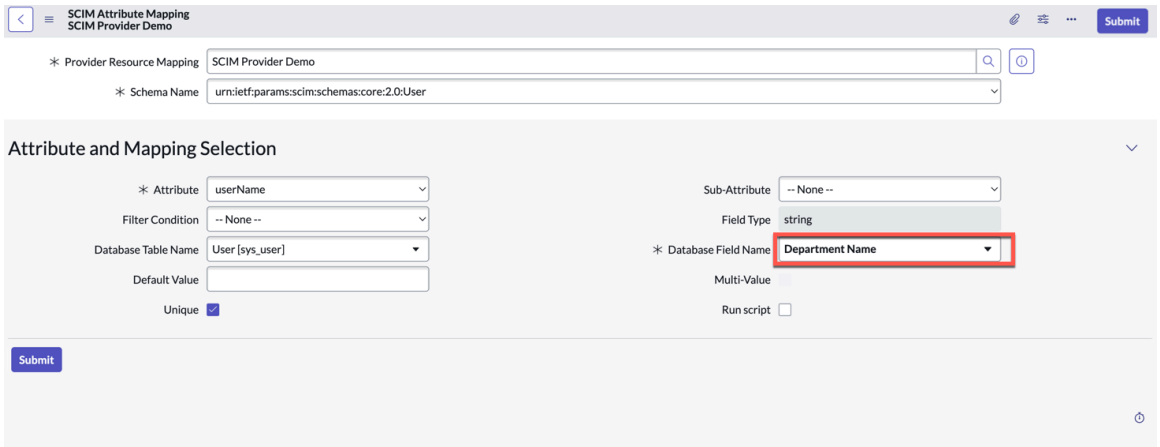

Default Value

The default value is passed to the SCIM Provider if direct attribute mapping of that field returns null. The default value can also be used to return a hard-coded value.

In the case of a hard-coded value, the database table name and field name should be **None**.

For example, the primary sub-attribute value of work email can be hard coded as **true**.

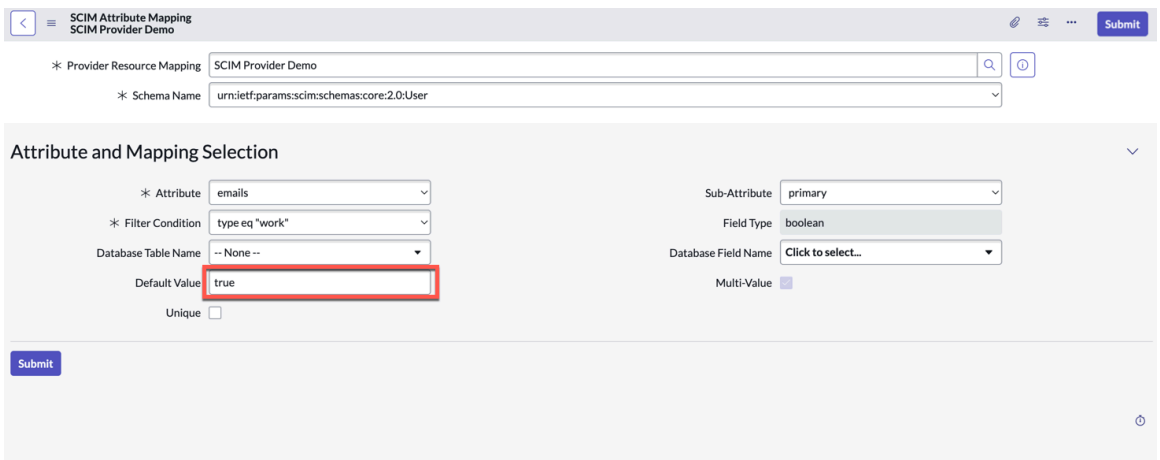

Script

The script is used to fetch the attribute value. The return type of the script should be always a string, or a JSON converted as a string. The output of the script should be in the proper format as expected by the provider for that attribute.

The following is a sample script for a multi-valued attribute.

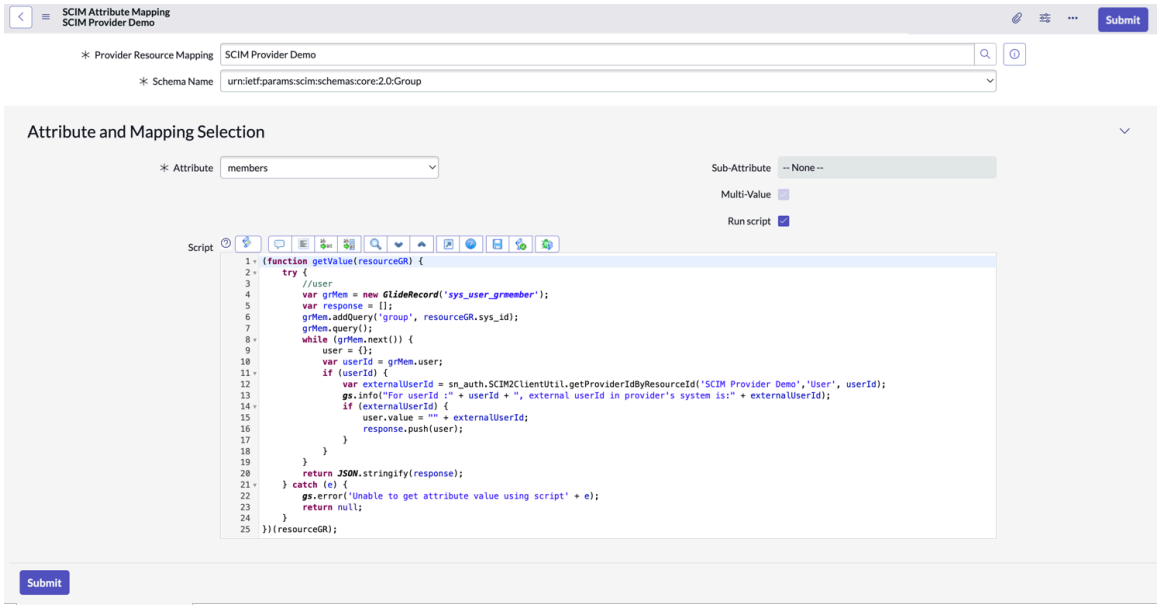

The output of the script should have a stringified JSON Array.

The following is a sample Script of a simple-valued Attribute.

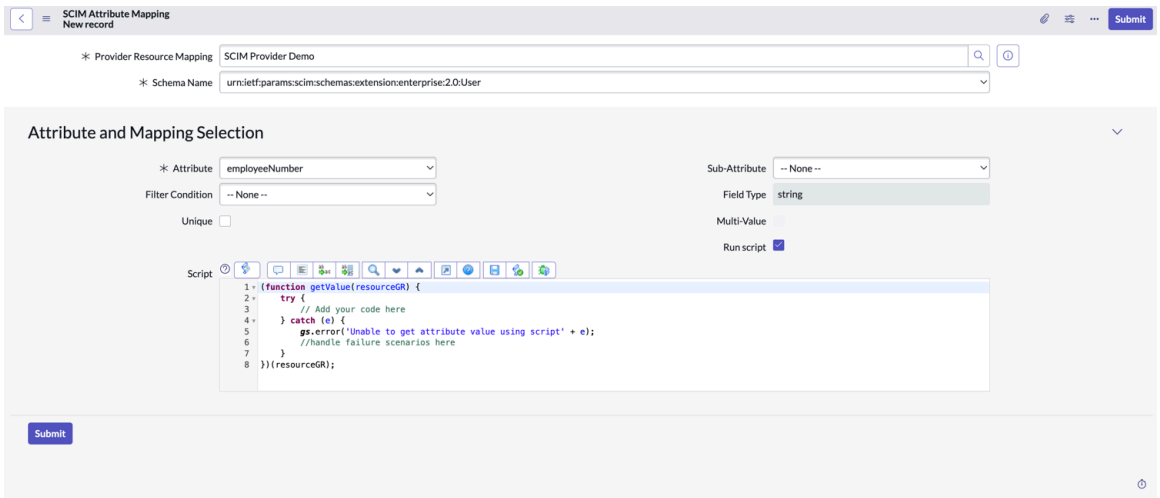

The output of the script should be a string.

Troubleshoot SCIM Client

Troubleshooting actions can help resolve common issues when setting up or running the SCIM Client.

Troubleshooting

Issue	Action
Response contains the following message: `“message":"Unable to access the table core_company with id: 0c441abbc6112275000025157c651c89,`	The message is displayed if the API is called in a user context and the user does not have access to the table. You must make sure that the scriptable API is called in the system context.

Troubleshooting (continued)

Issue	Action
`Please cross check the Access control rules"`	
Response contains the following message: `{"message":"User Not Authenticated","detail":"Required to provide Auth information”}`	• Make sure that the token is generated through the corresponding REST message and it is valid. • Make sure that the scriptable API is called in the system context.
Response contains the following message: `Script execution failed, the reason is: Cannot cast java.lang.Integer to java.lang.String`	In the SCIM attribute mapping, if the field is defined to fetch from a script with this issue, then make sure that the return type should always be a string.
Response contains the following message: `"status": "400", "scimType": "invalidValue", "detail": "Manager id : 02826bf03710200044e0bfc8bcbe5ds8 doesn't exist"`	For any SCIM attribute that expects the ID, then that ID is always of the provider system. Make sure the ID that is passed in the payload is valid in the provider system.

Areas to check for troubleshooting

The following are some of the areas that can be checked for troubleshooting errors when using the SCIM Client:

- If any issue is found while running any of the scriptable APIs, see the SCIM Client logs section.

Logs Fields

Field	Description
Request ID	Unique ID that represents a scriptable API called.
SCIM Provider	Provider name for which the API is called.
Resource	Name of the resource for which the API is called.
Resource ID	The ID for which the API is called. For deletion, the ID represents the resource ID in the provider system and the ID is in the client system.
Action	API that is called
Status	Status of the log as a success or failure

Field	Description
Message	Success or error message. The error message can be from the SCIM provider or because of configuration issues in the SCIM Client.

- Check the request body by configuring and viewing outbound calls. To learn more, see Outbound web services logging.
- Update log levels by adding the content-type, testing the sample, and navigating to corresponding the HTTP Method form in the REST message of the corresponding provider.
- If the request body is truncated, then increase the limit by using system property *glide.outbound_http.content.max_limit*.

Access observer

Use Access Observer to understand people and processes access data on your instance.

With Access Observer security administrators can:

- Understand which entities like users, roles, scoped apps, and scripts that access your data.
- Use this foreknowledge to know how best to apply security to limit unnecessary access to your data, while confirming those requiring access and still operate normally.
- Avoid broken automation by seeing a clear view of how your data is accessed before making security changes.
- Address the need to provide information regarding how encryption is applied on your instance.

Configure Access Observer by creating Access Observer configuration records. Within these records, you define a specific table and column to be observed and the time in which the column is observed.

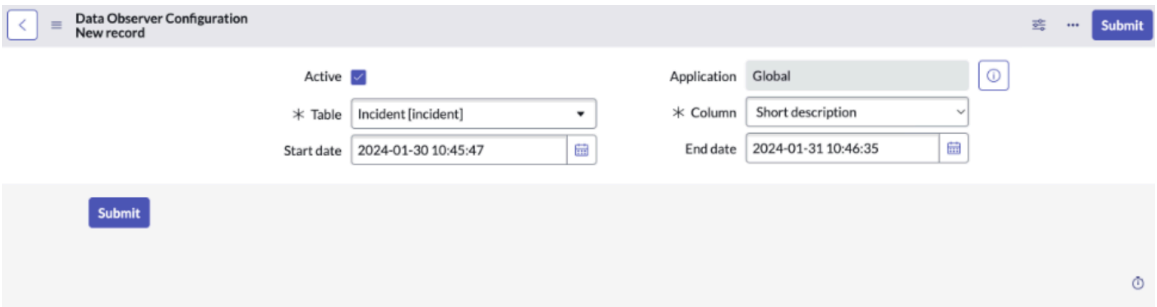

Find the results of your observations on the Access Observer log record table. On this table you can see a record detailing each time the specified column is accessed.

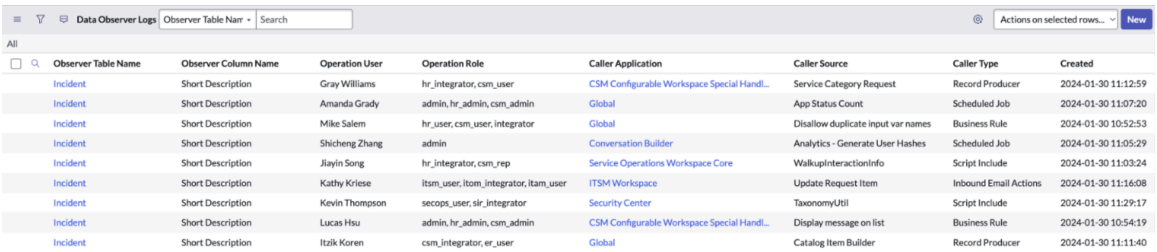

Observer Table Name	Observer Column Name	Operation User	Operation Role	Caller Application	Caller Source	Caller Type	Created
Incident	Short Description	Gray Williams	hr_integrator, csm_user	CSM Configurable Workspace Special Handl...	Service Category Request	Record Producer	2024-01-30 11:12:59
Incident	Short Description	Amanda Grady	admin, hr_admin, csm_admin	Global	App Status Count	Scheduled Job	2024-01-30 11:07:20
Incident	Short Description	Mike Salem	hr_user, csm_user, integrator	Global	Disallow duplicate input var names	Business Rule	2024-01-30 10:52:53
Incident	Short Description	Shicheng Zhang	admin	Conversation Builder	Analytics - Generate User Hashes	Scheduled Job	2024-01-30 11:05:29
Incident	Short Description	Jiayin Song	hr_integrator, csm_rep	Service Operations Workspace Core	WalkupInteractionInfo	Script Include	2024-01-30 11:03:24
Incident	Short Description	Kathy Kriese	itsm_user, itom_integrator, itam_user	ITSM Workspace	Update Request Item	Inbound Email Actions	2024-01-30 11:16:08
Incident	Short Description	Kevin Thompson	secops_user, sir_integrator	Security Center	TaxonomyUtil	Script Include	2024-01-30 11:29:17
Incident	Short Description	Lucas Hsu	admin, hr_admin, csm_admin	CSM Configurable Workspace Special Handl...	Display message on list	Business Rule	2024-01-30 10:54:19
Incident	Short Description	Itzik Koren	csm_integrator, er_user	Global	Catalog Item Builder	Record Producer	2024-01-30 11:11:40

Configure access observation

Create an access observation record to review access to a data column during a specified time window.

Before you begin

Role required: security_admin

Procedure

1. Navigate to **All** > **Access Observer** > **Access Observer Configuration**.
2. Select **New** to create a record.
3. On the form, fill in the fields.

Access observer configuration fields

Field	Description
Active	Whether the record is active. Your record is automatically marked active during the observation period and is inactive otherwise.
Application	The scoped application for the record. This field is read only.
Table	Table that contains the column to be observed
Column	The column within the selected table to be observed. The following field types are supported: ○ Email ○ Date ○ Date/Time ○ Journal ○ Journal Input ○ Phone number (E164) ○ String ○ String (UTF-8) ○ Translated ○ Translated Field ○ Translated HTML ○ Translated Text ○ URL
Start job immediately	If selected, observation begins as soon as the record is created.
End date and time	The time at which observation ends.
Start date and time	The time at which observation begins. This field is only visible if **Start job immediately** isn't selected.

4. Select **Submit** to save the record.

What to do next

Once the observation window your defined has started, you can see each instance where the column was access detailed in the records on the Access Observer log [sys_data_ob_log] table.

Review Access Observer logs

Use information in the Access Observer log records for insights on how your data is accessed.

Once you've configured one or more access observation records, your instance begins creating Access Observer log records with details about access to the columns you have chosen. You can find these records on your instance at **All** > **Access Observer** > **Access Observer Log**.

Access observer log results

Use the following table to understand the information presented in the logs.

Access observer log fields

Field	Description
Observer Table Name	The table selected in the access observation record that generated this record.
Observer Column Name	The column selected in the access observation record that generated this record.
Operation User	User that accessed the column
Operation Role	Roles for the user that accessed the column
Caller Application	The scoped application from which the data was accessed.
Caller Type	The type of element which has accessed the column, such as record producer, scheduled job, or business rule.
Caller Source Document ID	
Caller Source	Element which accessed the column. Used along the **Caller Type** field to see what specifically has accessed the column. For example if the caller type is `Business Rule`, the caller source is the name of the business rule which accessed the column.
Primary Hash	
Repeat Count	
Java Stack	
JavaScript Stack	
Session ID	ID of the session in which the column was accessed.

Granular admin roles required to secure your instance

Verify proper access management by assigning roles that define user permissions and responsibilities. By doing so, organizations can maintain security, enforce conformance, and optimize their operations effectively.

Roles are a fundamental part of managing access and maintaining security within your instance. They define what you can see and do, verifying that you have the appropriate level of access based on your responsibilities. By assigning the correct roles to the users, organizations can safeguard sensitive data, enforce compliance, and streamline operations.

To optimize access management within the ServiceNow AI Platform, consider adopting granular admin roles. This approach enables you to assign specific permissions to developers or users who perform minor administrative tasks, without granting them unrestricted access to the full admin role.

By adopting granular admin roles, you can create a more secure and efficient access management system that aligns with your organization's needs.

Note:

- Each product within the ServiceNow AI Platform has its own set of granular admin roles. To determine the appropriate roles for your developers, refer to the specific product documentation.
- Granular admin roles are separate from the existing admin role and must be assigned independently.

servicenow

Additional resources for Platform Security products and solutions

If you're looking for Platform Security best practices, troubleshooting, or other implementation guidelines, select a feature or resource type to discover ServiceNow resources on other relevant websites.

Note: Many resources in this table require you to log in to sites like ServiceNow University, Now Create, or the ServiceNow Community. If the expected resource does not load, please log in and try to access the resource again.

Resource links

Platform Security features or products	Resource type	Resources
Security Center	Getting Started	ServiceNow Security Center
Security Center	Best practices	• Best Practices - Security Center • Security Center Hardening best practices
Security Center	FAQs	ServiceNow Security Hardening - Security Center
ServiceNow Vault	Getting Started	• What is ServiceNow Vault? • ServiceNow Vault Console Overview
Code Signing	Getting Started	• Code Signing for enhanced defense-in-depth • Code Signing and Circle of Trust (CoT): Introduction
Data Privacy	Tips and examples	[Washington Release] Mitigating the risk of inadvertent sensitive data exposure with Data Privacy
Field Encryption	Tips and examples	Field Level Encryption in ServiceNow
Field Encryption Enterprise	Getting Started	Field Level Encryption vs Platform Level Encryption differences
Field Encryption	Troubleshooting	• Why is Field Level Encryption not applied for fields of the Phone Number data type • We are encrypting the Field Level on employee relations case • Encrypting attachments using "Field Level Encryption" • Unable to delete an Encrypted Field Configuration

Resource links (continued)

Platform Security features or products	Resource type	Resources
		• Case table is not available in Encrypted Field Configuration • Querying encrypted field not returning consistent results in Scheduled Jobs
Field Encryption Enterprise	Troubleshooting	• Field Level Encryption Enterprise: Initial Setup • Field Level Encryption Enterprise: Limitations and Considerations • Field Level Encryption Enterprise: Administration
Edge Encryption	Troubleshooting	• Edge Encryption on ServiceNow Overview • Edge Encryption: Initial Setup • Edge Encryption Rules • Edge Encryption: Limitations • Edge Encryption: Administration
Edge Encryption	Tips and examples	• Edge Encryption Resources • Edge Encryption Proxy logs • Edge Encryption Certificate Expiry Notification
Edge Encryption	Troubleshooting	• Edge Encryption Performance Graphs • Edge Encryption Extends Field Length
Database Encryption	Tips and examples	Database Encryption in ServiceNow: TSE/DBE using CCS
Access Management	Tips and examples	Identity And Access Management With ServiceNow
Zero Trust Access	Troubleshooting	Help with Access Analyzer, Zero Trust Access, and Adaptive Authentication on ServiceNow
Domain separation for service providers	Troubleshooting	ServiceNow Domain Separation for Service Providers
Data filtration	Getting started	Understanding the "Data Filtration" plugin
Security Roles	Tips and examples	How do security roles work?
Connections and Credentials	Getting started	Connections and Credentials Overview - Learn Integrations on the ServiceNow AI Platform

Resource links (continued)

Platform Security features or products	Resource type	Resources
ServiceNow access controls	Tips and examples	Access Controls - The Easy Way
Security Center	Best practices	Best Practices - Security Center
Platform Security	Best practices	• ServiceNow Security Best Practices Guide • Securing the ServiceNow AI Platform

Virtual infrastructure security

Use virtualization with the flexibility to install multiple MID Servers in virtualized operating systems and networks in shared physical hardware.

Access to the hypervisor and virtual infrastructure management consoles must be protected to prevent unauthorized cloning of the virtualized MID Server:

- Limit the access to the hypervisor and virtual infrastructure management consoles by the trusted few admins.
- Only allow connections to the internal trusted network for management console such as ESX Server and VirtualCenter.
- Use VLANs to prevent network attacks.
- Follow the security guidelines published in the virtualization hardening guides from the vendors.

Operating system security

Learn how the MID Server stores its ServiceNow AI Platform username and password in its configuration file, named `config.xml`, for secure authentication to the instance.

The MID Server must be deployed in a secure and hardened operating system for protection against unauthorized access to the credentials:

- Limit the access to the operating system to a few, trusted administrators.
- Monitor the operating system logs to detect any unauthorized access, particularly attempts to access the config.xml file, as this file contains important MID Server configuration information.
- Install operating system security patches regularly with the latest anti-virus software and update AV definitions regularly.
- The MID Server requires a current Java framework to run. Keep Java updated regularly.
- Remove or disable unnecessary services and applications.
- Install an OS firewall to limit access to unauthorized ports.
- Follow the security guidelines published in the operating system hardening guides from the vendors.

Network security

The MID Server communicates on port 443 using SSL to the instance and requires no inbound connections.

To properly secure your MID server in a network, do the following:

- Install the MID Server on a secured server behind a corporate firewall for protection against unauthorized access from the internet.
- Configure the firewall on the MID Server to accept no inbound connections other than the ones required for corporate management of the operating system and hardware.
- The system that hosts the MID Server must be able to access the ServiceNow download site at **install.servicenow.com**.

  Note: This URL is to the ServiceNow download site, which is not accessible from this topic.

  The MID Serve host machine must be able to access that site to download the installer package. It contacts **install.servicenow.com** every 60 minutes to see if a newer version is available and if so, it performs an automatic upgrade. A MID Server upgrade also takes place when you upgrade an instance.