servicenow®

Australia Enable AI

Last updated: 04/07/2026

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

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.

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

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

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

servicenow

Table of Contents

Enable AI experiences **7**

ServiceNow product tiers 14

ServiceNow® AI implementation 16

Now Assist organization and tools 18

AI governance 23

Data readiness 26

Application readiness 27

Solving installation and configuration issues 35

Now Assist AI assets 39

AI assets on by default 41

Now Assist skills 41

Now Assist agentic workflows 58

Platform AI agents 97

Large language models on the ServiceNow AI Platform® 103

Now LLM Service updates 105

Writing instructions for large language models 108

Long term stable models 113

Discrepancies when using different AI search tools 113

Providers and Models 114

Now Assist 114

Exploring Now Assist Admin 118

Configuring Now Assist Admin features 251

Using Now Assist Admin 258

Now Assist reference 277

Generative AI Controller 283

Now Assist AI agents 321

Explore 322

Configure 347

Create an AI agent 372

Create an agentic workflow 430

Deploy AI voice agents 446

Evaluate agentic AI assets 474

Examples of using AI agents 503

Reference 505

Model Context Protocol Client 533

Now Assist Data Kit 550

Exploring Now Assist Data Kit 551

Configuring Now Assist Data Kit 553

Using Now Assist Data Kit 553

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

Now Assist Data Kit reference....563
Now Assist Skill Kit....564
Exploring Now Assist Skill Kit....565
Configuring Now Assist Skill Kit....576
Using Now Assist Skill Kit....581
Now Assist Skill Kit reference....601
AI Control Tower....603
Explore....605
Configure....715
Use....720
Reference....755
Knowledge Graph....776
Exploring Knowledge Graph....776
Configuring Knowledge Graph....796
Using Knowledge Graph Designer....798
Using Enterprise graph schema....810
Reference for Knowledge Graph....826
MCP Server Console....827
Explore....828
Configure....830
Connect....833
Reference....839
Now Assist Readiness Evaluation....842
Explore....843
Configure....846
Use....848
Reference....857
Natural Language Understanding....857
Exploring Natural Language Understanding....861
Model management....874
Multilingual model management....968
Virtual Agent and NLU Workbench integration....983
NLU Workbench - Advanced Features....985
Natural Language Query....1029
Exploring Natural Language Query....1029
Using Natural Language Query....1031
Configuring NLQ....1033
Natural Language Query References....1039
Agentic Desktop....1041
Explore....1042
Configure....1055
Design desktop actions....1062

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

Create AI Agents....1092
Desktop action examples....1097
Execute desktop actions....1120
Reference....1125
Predictive Intelligence....1131
Explore Predictive Intelligence....1132
Install Predictive Intelligence....1136
Configure Predictive Intelligence....1137
Creating and training solutions....1173
Using Predictive Intelligence....1212
Testing and monitoring predictions....1235
Predictive Intelligence references....1248
Now Assist in Document Intelligence....1255
Explore....1257
Configure....1264
Use....1275
Reference....1277
Document Intelligence....1290
Explore....1291
Configure....1294
Integrate....1320
Use....1327
Monitor....1348
Reference....1353
Task Intelligence....1382
Exploring Task Intelligence....1383
Configure Task Intelligence....1386
Manage machine learning models with Task Intelligence....1387
Task Intelligence references....1391
ServiceNow AI Lens....1394
Explore....1395
Configure....1401
Use....1426
Reference....1448
Conversation Insights....1455
Exploring Conversation Insights....1456
Install Conversation Insights....1459
Conversation Insights reference....1460
Conversation Improvement Themes....1461
Explore....1462
Activate....1463
Use....1464

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

Reference....1466
Now Assist and Agentic AI glossary....1469
agents....1470
agentic AI....1470
agentic conversation....1470
agentic system....1470
agentic workflow....1470
AI agents....1470
AI compliance....1470
AI Control Tower....1471
AI governance....1471
AI steward....1471
Citation....1471
custom skill....1471
dynamic translation....1471
External Content Connectors....1471
generative AI....1471
Generative AI Controller....1472
guardrail....1472
Knowledge Graph....1472
large language model (LLM)....1472
multi-turn catalog ordering....1472
native translation....1472
Now Assist....1472
Now Assist Admin console....1472
Now Assist Analytics....1472
Now Assist context menu....1473
Now Assist Guardian....1473
Now Assist panel....1473
Now Assist skill....1473
Now LLM Service....1473
offensive content detection....1473
Orchestrator....1473
prompt injection detection....1473
Q&A genius result....1473
Retrieval Augmented Generation....1474
Additional resources for AI products and solutions....1474

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

Enable AI experiences

Take advantage of artificial intelligence-based tools to prioritize and automate routine tasks, #detect major incidents, and surface insights.

Get started

Choose one of these tiles to get started.

ServiceNow AI implementation



Understand the elements of Now Assist and lay the groundwork for a seamless, intelligent experience across your workflows.

Now Assist AI assets



Explore Now Assist generative AI skills, AI agents, and agentic workflows.

Large language models (LLMs)



Choose the Now LLM Service or supported third-party AI models for skills and agents.

Now Assist Admin console



Use a personalized generative AI assistant and skills to enhance your work on the ServiceNow AI Platform.

Now Assist AI agents



Build and assemble use cases that drive business outcomes.

Now Assist Data Kit



Add datasets to a data catalog.

Now Assist Skill Kit



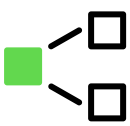
Create custom skills and prompts for your Now Assist experience.

AI Control Tower



Monitor, manage, and govern AI assets on the ServiceNow AI Platform.

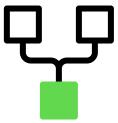
Knowledge Graph



Map entities and their relationships and add context and meaning to your data to enable intelligent search, insights, and AI-driven experiences.

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

MCP Server Console



Use Model Context Protocol (MCP) servers to allow AI applications to access data and perform actions on a ServiceNow instance.

Now Assist Readiness Evaluation



Assess processes, evaluate data readiness, and get actionable insights to help adopt Now Assist quickly.

Natural Language Understanding



Use Natural Language Understanding (NLU) to help your system work with natural human expression and intent.

Natural Language Query



Get data from your instance using plain language requests with this service on the ServiceNow AI Platform.

Predictive Intelligence



Develop ML-based solutions using Predictive Intelligence to create improved work experiences in your instance.

Document Intelligence



Automate the data extraction process for your documents and speed up your workflow with a powerful AI-based solution.

Now Assist in Document Intelligence



Use generative AI to analyze and extract information from documents

Task Intelligence



Create ML solutions with this AI tool, and track their impact on your business demands and outcomes.

Agentic Desktop



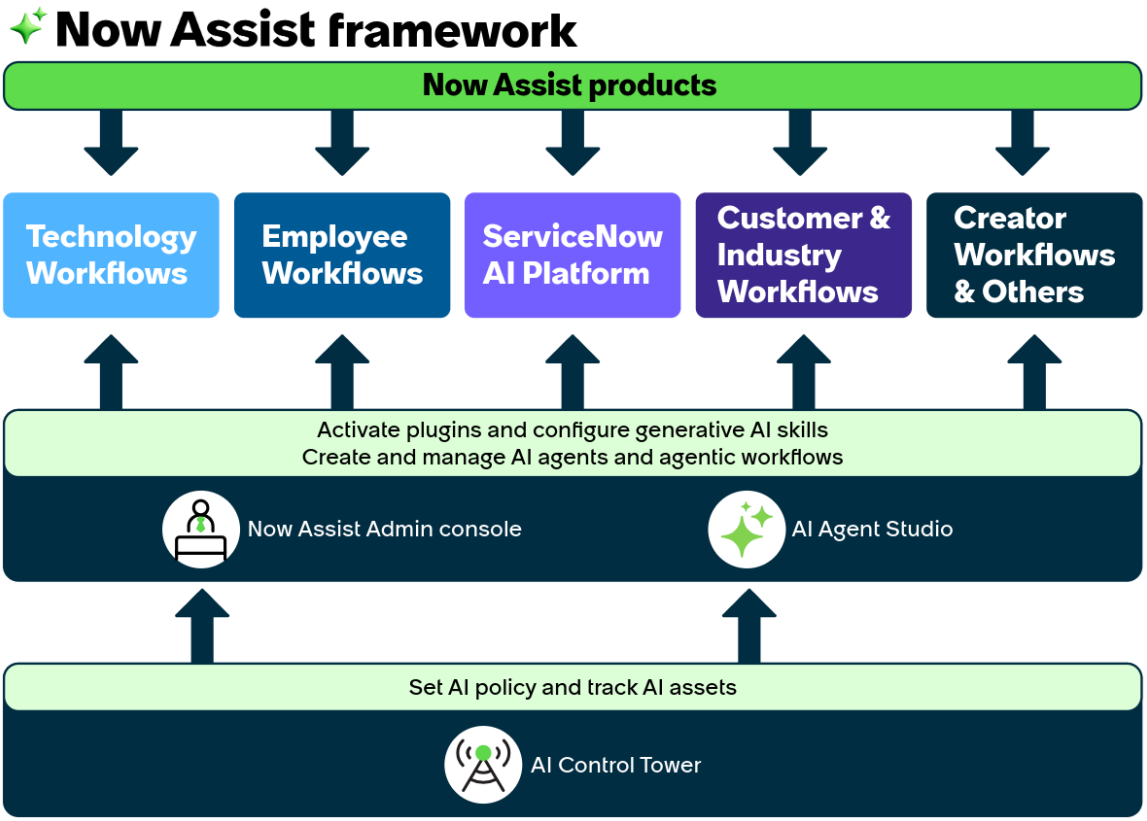
Design, configure, and manage desktop actions that automate repetitive tasks in your desktop environment.

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

ServiceNow AI Lens	Additional Resources	
Use generative AI to scan, extract, comprehend, and synthesize data and take actions to optimize your workflows.	Upgrade your AI experience to the next level	

Now Assist

Enable generative AI features on the ServiceNow AI Platform using Now Assist applications. Generative AI can help your agents and developers work faster and be more productive.



Now Assist AI agents

ServiceNow AI agents can solve essential business challenges autonomously, accelerating outcomes and empowering people with a skilled digital workforce without limits, so they can focus on their best work.

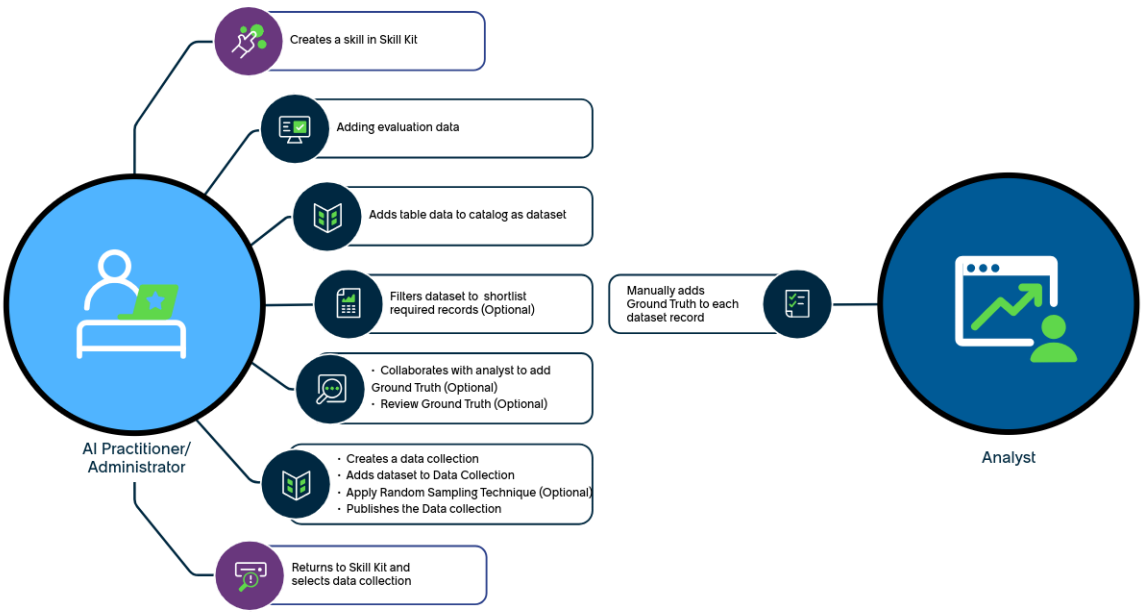
Use AI Agent Studio to create, manage, and test AI agents and use cases all in one place. The main Overview page has three sections that give you what you need to understand, begin, and continue developing AI agents and use cases.

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

Now Assist Data Kit

If the base system Now Assist skills don't fit your needs, you can use Now Assist Data Kit to create custom datasets and data collections that can be used in Now Assist Skill Kit for evaluation.

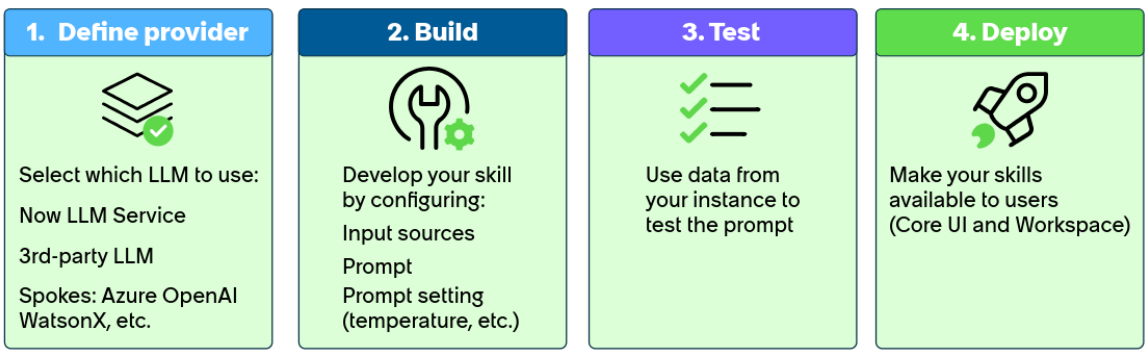
User journey for Now Assist Data Kit



Now Assist Skill Kit

Use Now Assist Skill Kit to create custom skills. These custom skills give you greater flexibility in the Now Assist environment.

User journey for Now Assist Skill Kit



AI Control Tower

AI Control Tower can help organizations monitor, manage, and govern AI assets on the ServiceNow AI Platform. An AI Control Tower framework is a set of principles, practices, and protocols that guide the development, use, and deployment of AI systems. The AI Control Tower framework helps organizations deliver AI consistent with their own responsible AI principles.

Knowledge Graph

Knowledge Graph application enhances the Now Platform by creating a semantic layer that connects data, relationships, and context across the enterprise. It structures information as

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

a graph of entities and connections, bringing context and meaning to the available raw data. By leveraging AI, it powers personalized and intelligent experiences across the Now Assist ecosystem including NAVA, AI Agents, and AI Search to deliver more relevant insights and automation.

By linking data and embedding semantic meaning, Knowledge Graph transforms raw data into actionable knowledge, fueling faster resolutions, richer insights, and more personalized AI-powered experiences.

MCP Server Console

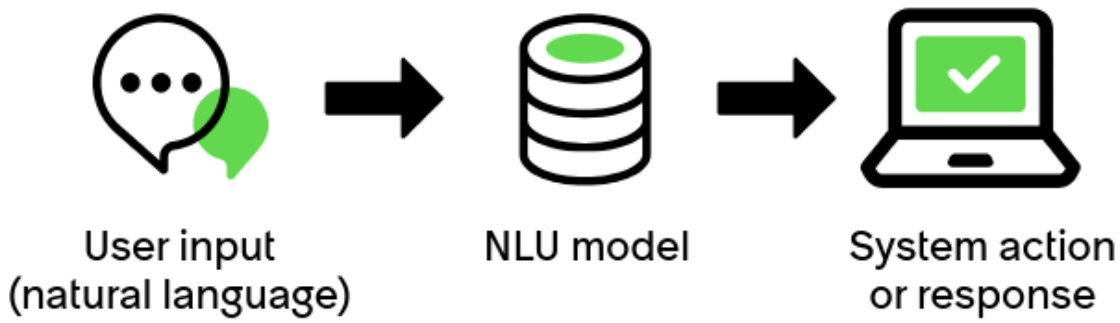
The Model Context Protocol defines a standard method of communication between large language models (LLMs) and external systems, such as a ServiceNow instance. AI applications connect to an external system from an MCP client, such as an AI agent, using an MCP server. The server tells the client which external resources it can access and how to access them. Then, from an MCP client, users can request information from the server and automate functionality using the available tools and data that the server returns.

Now Assist Readiness Evaluation

The Now Assist Readiness Evaluation app is a solution designed to simplify and automate the agentic AI and Now Assist implementation assessment process. The app helps to determine whether your organization's instance is ready to implement generative and agentic AI features in Now Assist. Previously, manually assessing your organization's instance readiness was time consuming and took significant effort. Now, you can use the Now Assist Readiness Evaluation app to automate gathering, processing, and analyzing instance data so that you can review your instance readiness more quickly. The app provides results within seconds, helping to reduce manual effort and promoting rapid and reliable assessments.

Natural Language Understanding

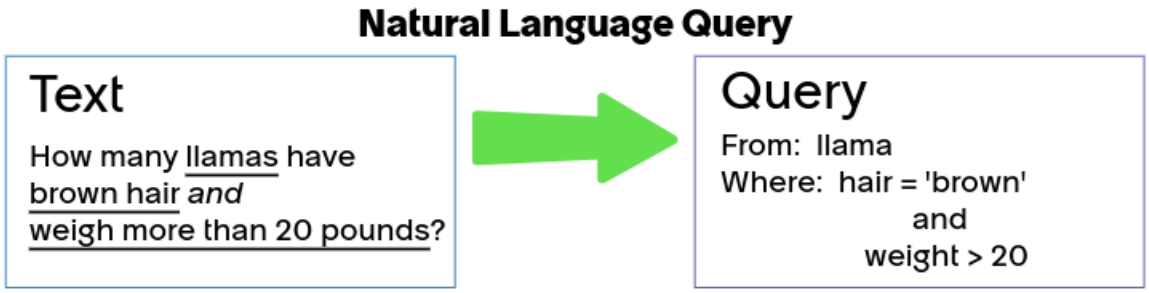
Help users communicate with your system in naturally-expressed language, using Natural Language Understanding. NLU enables your system to perform intelligent actions in response to human language input in 17 supported languages. Start from the provided pre-built models and expand them further, or build your own models from scratch.



Natural Language Query

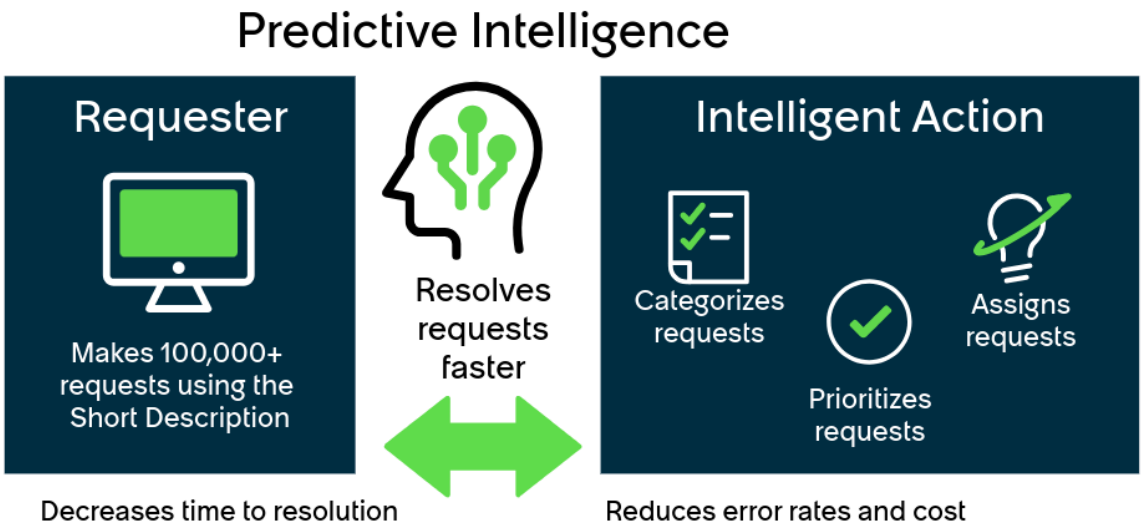
Transform natural-language questions into formal database queries with Natural Language Query (NLQ). Get data from your instance by using plain language requests in the supported languages American English, French, French Canadian, German, Japanese, and Spanish. NLQ is consumed by several other applications and features, including Analytics, Reporting, and CMDB (English is the only supported language for CMDB).

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



Predictive Intelligence

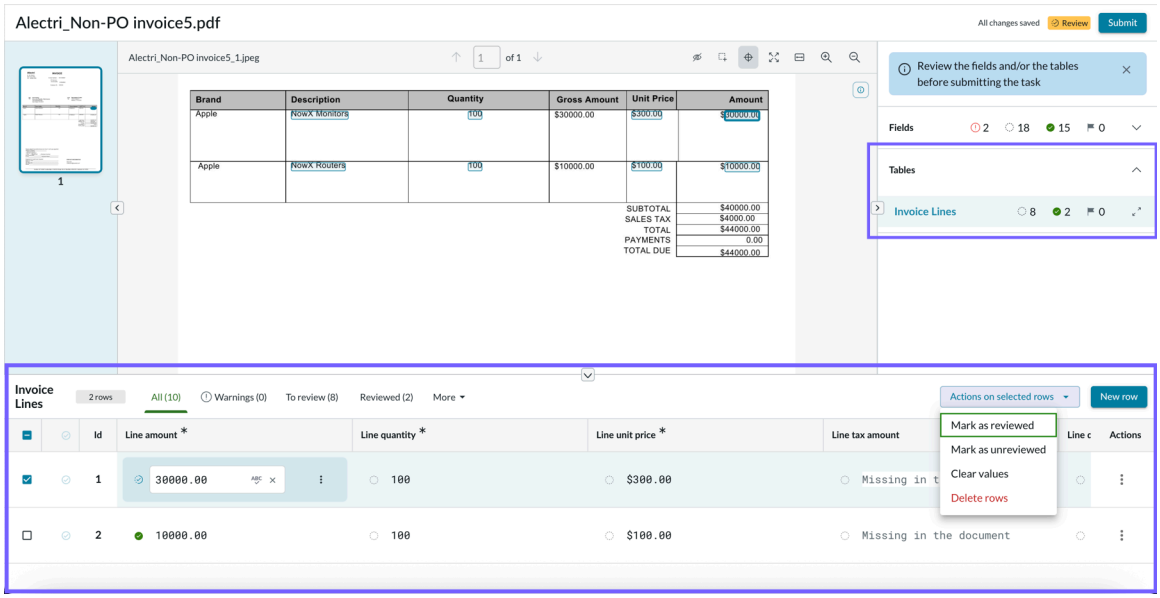
Train your system to automate request handling with Predictive Intelligence. Three machine-learning frameworks (classification, clustering, and similarity) can help you address a range of business challenges. Discover patterns in your data, reduce task-handling times and interaction or error counts, highlight issues shared between similar records, and manage records collectively. Support is available for a set of languages.



Document Intelligence

Get clear, concise data from your records quickly, and integrate it where you need it most. Document Intelligence can process single or multiple-page documents in various popular formats. This tool uses optical character recognition (OCR) and a continual-learning AI to identify and extract text and data. Document Intelligence works on a no-code basis, allowing for easy setup, use, and integration of data into workflows.

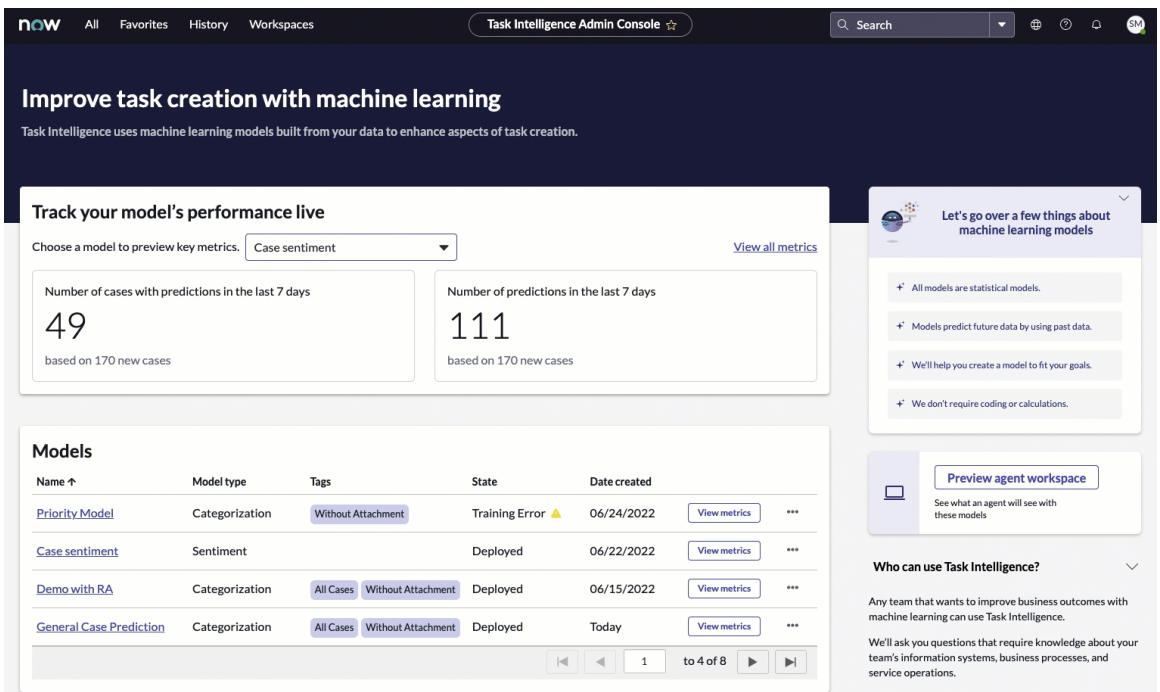
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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



You can use Now Assist in Document Intelligence skills to extract information and find answers to questions about documents.

Task Intelligence

Achieve faster response and task-solving times for your Agents, by using Task Intelligence. This tool lets you engineer ML solutions that handle data and track the efficacy of those solutions. Build, train, edit, and retrain machine-learning models in the Admin Console, and export them for use across the ServiceNow AI Platform. Follow up with the Analytics dashboard, where you can track your models' performance and effects on your business, and determine which predictions are best for future reuse.



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

ServiceNow product tiers

ServiceNow structures its products and packages in three tiers — Foundation, Advanced, and Prime. Each tier incorporates AI and builds progressively on the previous one with additional AI capabilities, agents, and governance tools.

The Foundation tier delivers AI-assisted insights and routine automation. Advanced adds agentic workflows capable of synthesizing context across complex processes. Prime unlocks a fully autonomous workforce that executes multi-step tasks end to end.

This structure makes the path to greater automation simple, predictable, and scalable across your organization.

Note: Contact your ServiceNow account team for information about availability and entitlement details for your organization. The rollout of new product tiers is independent of your organization's upgrade cycle.

Product tiers

All supported product lines offer these three tiers that include the features listed below.

Product tier feature comparison

Feature	Foundation	Advanced	Prime
AI skills and routine AI agents	Supported	Supported	Supported
Configure out-of-the-box skills and agents	Supported	Supported	Supported
Agentic workflows with contextual AI synthesis	Supported	Supported	Supported
Platform Analytics Advanced	Not supported	Supported	Supported
Create net-new custom AI skills and agents	Not supported	Not supported	Supported
MCP Server (inbound)	Not supported	Not supported	Supported
Autonomous AI workforce (AI Specialists)	Not supported	Not supported	Supported

Foundation	Advanced	Prime
AI skills, routine AI agents	Agentic workflows	Autonomous AI workforce
This solution provides task-based assistance using routine pattern recognition and categorization to accelerate understanding and help you work faster.	This solution automates entire steps in a workflow, working side-by-side with you while synthesizing new insights, understanding context, and applying deep domain knowledge.	This solution offers fully independent AI specialists to make and execute decisions autonomously. They apply deep, role-based expertise to master and run multiple agentic workflows.

AI platform enablers

Every offering has a set of platform-level AI capabilities powering the skills, agents, and governance experience across all product lines.

Now Assist

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

Now Assist is the generative AI experience embedded throughout the ServiceNow AI Platform, delivering skills such as incident summarization, sentiment analysis, reply generation, and case resolution assistance. Now Assist skills are available at every tier, across ITSM, CSM, HRSD, and industry solutions. For a full overview of the Now Assist panel and administration tools, see Now Assist panel.

Now Assist AI agents

Now Assist AI agents extend generative AI into autonomous agentic workflows. Using AI Agent Studio, administrators can configure out-of-the-box agents at Foundation and Advanced tiers, or build net-new agents from natural language instructions at the Prime tier.

AI Agent Fabric

AI Agent Fabric is the communication layer that enables ServiceNow AI agents to collaborate with each other and with third-party AI systems using open protocols — including Agent-to-Agent (A2A) and Model Context Protocol (MCP). Foundation and Advanced tiers include A2A outbound connectivity; Prime adds inbound MCP Server capability, enabling external platforms to invoke ServiceNow agents directly. For implementation details, see Enable MCP and A2A for your agentic workflows.

AI Control Tower

AI Control Tower is embedded at every tier, providing centralized governance, lifecycle management, and real-time visibility across all AI assets — whether built natively on ServiceNow or sourced from external vendors. At Foundation and Advanced, AI Control Tower discovers and manages ServiceNow AI assets; Prime extends full management and assist metering to external AI assets as well. For configuration guidance, see Exploring AI Control Tower and Configure AI Control Tower.

Workflow Data Fabric

Workflow Data Fabric grounds AI agents in real enterprise data by connecting any application, database, or system to the ServiceNow AI Platform — without requiring data to be moved or replicated. WDF Foundation, including Integration Hub, Robotic Process Automation Hub, Automation Center, and Data Catalog, is embedded in every AI native edition. WDF Advanced, adding Zero Copy Connectors and Stream Connect, is available as a paid upgrade.

RaptorDB

RaptorDB is the next-generation ServiceNow database, purpose-built to deliver the performance and scale that AI-native workloads demand. RaptorDB Standard underpins every AI native edition with improved response times and optimized query performance; RaptorDB Professional unlocks ultra-scale analytics, enhanced column-store capabilities, and advanced instance topology support for organizations running the most demanding generative AI and machine-scale data use cases.

Moveworks for ServiceNow (EmployeeWorks)

ServiceNow EmployeeWorks combines Moveworks' conversational AI and enterprise search with ServiceNow's unified portal and autonomous workflows, giving employees a single, natural-language front door to enterprise services — available across Teams, Slack, browser, and mobile. EmployeeWorks capability is bundled at a level aligned with each AI native tier, ensuring that conversational access scales with your agentic AI investment. For more information, see the Autonomous Workforce and EmployeeWorks announcement.

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

Get started

To begin implementing AI native capabilities on your instance, see the following resources.

- Now Assist overview — Learn about the applications and features that make up the Now Assist experience.
- AI governance — Learn about the importance of AI governance to ensure responsible use, regulatory compliance, and alignment with enterprise goals.
- Data readiness — Learn how to prepare your instance data for Now Assist.

Related ServiceNow applications and features

- Now Assist panel
- Exploring Now Assist
- AI Agent Studio
- AI Control Tower
- AI Control Tower Home
- Exploring AI Control Tower
- AI Control Tower dashboard
- AI Control Tower release notes
- Workflow Data Fabric
- RaptorDB
- Now Assist for ITSM release notes

ServiceNow® AI implementation

Getting ready to implement Now Assist is more than just installing plugins—it's about laying the groundwork for a seamless, intelligent experience across your workflows. Whether you're enabling conversational catalogs, automating content generation, or enhancing user interactions, a few key steps will ensure your data is ready, your applications are prepared, and your organization's AI policy is in alignment with your implementation.

Get started

Now Assist overview	AI governance	Data readiness
		
Learn about the applications and features that make up the Now Assist experience.	Learn about the importance of AI governance to ensure responsible use, regulatory compliance, and alignment with enterprise goals.	Learn how to prepare your instance data for Now Assist.

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

Application readiness	Resolve common issues	
		
Ensure that your instance is ready to take advantage of Now Assist by preparing Platform applications.	Address issues or gaps in your Now Assist configuration.	

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

Additional implementation information is available on Now Create.

For additional Now Assist resources, see Additional resources for AI products and solutions.

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.

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.

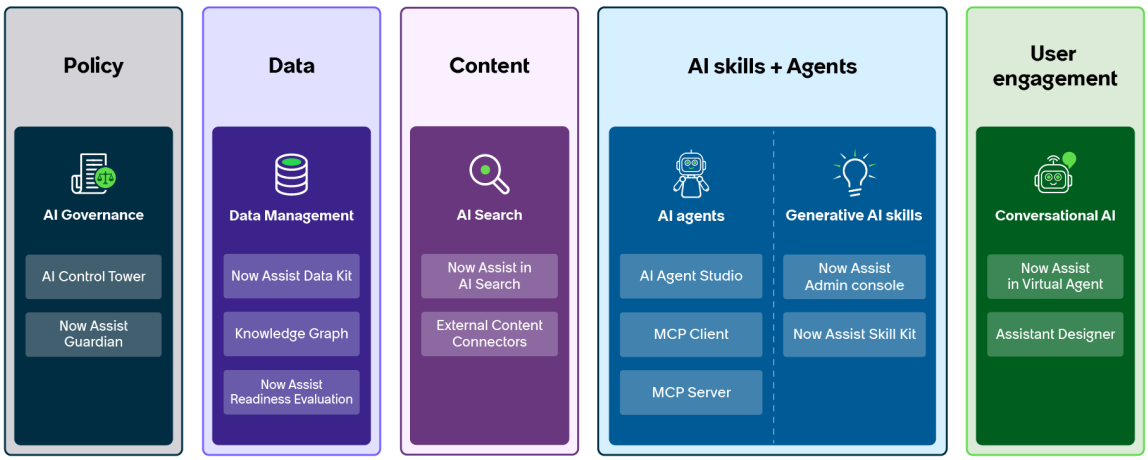
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

Now Assist organization and tools

The Now Assist experience includes generative AI skills, agentic AI, and conversational user engagement layer. A good AI experience depends on quality data and a coherent AI policy that functions as its guiding star. AI Search capabilities and other tools help connect the pieces.

The following infographic illustrates the components of the Now Assist experience. You can think of each area as a layer, with AI policy as the basis for your implementation. Data plays an important role, since poor data will lead to poor outcomes, and the search experience drives both generative and agentic AI capabilities. The conversational engagement layer allows your users to communicate with AI and perform tasks and self-serve more easily.

Components of the Now Assist experience



AI policy

AI governance is defined by your organization and then implemented on the platform using AI Control Tower. AI Control Tower is designed to scale with your organization's AI maturity. It enables end-to-end life cycle management of AI agents and models, from intake and evaluation to deployment and optimization. It also acts as a bridge between governance committees and implementation teams, ensuring that policy decisions are translated into technical configurations and operational controls. The platform includes the AI Risk and Compliance application, which comes with pre-built content aligned to the EU AI Act and NIST AI Risk Management Framework (AI RMF), helping organizations accelerate compliance and implement responsible AI practices from day one. For more information, see AI governance for Now Assist on the ServiceNow AI Platform.

For real-time guardrail enforcement for AI interactions, you can use Now Assist Guardian to monitor and mitigate risks related to offensive content, sensitive topics, and security threats in generative AI outputs. Admins can configure these guardrails directly in the Now Assist Admin console, and logs are available to track detection events and guide decisions about blocking or escalation.

Roles and responsibilities:

AI steward [sn_ai_governance_ai_steward]

The AI steward role manages the platform's data assets and AI capabilities, ensuring data quality, integrity, and ethical use. The steward also implements governance policies to maintain compliance, security, and responsible AI practices across ServiceNow applications and workflows.

AI asset/product owner [sn_ai_asset_mgmt.ai_asset_owner]

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

The AI asset (or product) owner is responsible for driving the development and delivery of AI-powered products, ensuring alignment with business goals and user needs. They define product vision, manage feature prioritization, and translate business requirements into technical deliverables.

AI Risk and Compliance Manager [sn_grc_ai_gov.ai_risk_and_compliance_manager]

The AI Risk and Compliance Manager is tasked with identifying, assessing, and mitigating risks related to AI use within the enterprise. This role ensures AI compliance with laws, regulations, and internal standards while addressing issues such as bias, privacy, and transparency.

Now Assist Admin [sn_nowassist_admin.nsa_admin]

The Now Assist admin is responsible for enabling and managing the Now Assist Guardian guardrails within the Now Assist Admin console. This includes configuring detection settings for offensiveness, prompt injection, and sensitive topics, as well as monitoring logs to assess harmful or malicious AI outputs. The admin can toggle guardrails on or off, choose actions like “Block and Log,” and tailor the configuration to match organizational risk tolerance and use cases.

For more information about these tools, see:

- AI Control Tower
- AI Risk and Compliance
- Now Assist Guardian

Data management

You know the old saying, "Garbage in, garbage out?" Poor, incomplete, or unstructured data can lead to poor AI outcomes. To prepare your data for AI implementation, see Data readiness for implementing Now Assist on the ServiceNow AI Platform.

Now Assist provides several tools to help you manage data and provide user context on your instance. The Now Assist Data Kit helps organizations manage AI training data on their ServiceNow instance. It provides a centralized workspace to curate, generate, cleanse, and evaluate datasets, including importing data from tables, adding ground truth, and creating derived datasets. It supports synthetic data generation—up to 1,000 records at a time—and includes tools to anonymize sensitive information for privacy compliance. Integrated with the Now Assist Skill Kit, it enables teams to test custom AI skills and measure performance.

Knowledge Graph makes AI in Now Assist more personal by using structured enterprise data to tailor responses to each user. It helps AI understand context, such as a user’s role, location, or assigned assets, making answers more relevant and accurate. It also supports intelligent slot-filling, automatically pulling in known information to streamline conversations and reduce repetitive questions. This results in faster, more personalized, and trustworthy AI interactions.

Roles and responsibilities:

Now Assist Data Kit admin [sn_data_kit.admin]

This role is required to create, update, and publish datasets in Now Assist Data Kit.

Knowledge Graph admin [kg_admin]

This role is required to design, manage, and audit Knowledge Graph schemas, including configuring nodes (tables), properties (columns), and relationships.

For more information about these tools, see:

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

- Now Assist Data Kit
- Knowledge Graph
- Now Assist Readiness Evaluation

Internal and external content search

Now Assist in AI Search is the primary tool used to find content on your instance, whether it's in the form of knowledge articles, Service Catalog items, system users, Knowledge Graph schemas, or content in enhanced chats using Now Assist in Virtual Agent.

Now Assist in AI Search is a plugin that enhances your existing AI Search configuration, adding support for Now Assist Genius Result answers.

You can expand the amount of content available to Now Assist in AI Search by using External Content Connectors. With this tool, AI Search applications can search content and metadata from external content repositories such as Atlassian Confluence Cloud and Microsoft SharePoint Online.

Roles and responsibilities:

AI Search admin [ais_admin]

This role is required to manage configuration settings for the AI Search application.

External Content Connectors admin [sn_ext_conn.xcc_admin]

This role is required to manage configuration and crawl settings for the External Content Connectors application.

For more information about these tools, see:

- AI Search
- Now Assist in AI Search
- External Content Connectors

Generative AI skills

Now Assist product plugins provide workflow-specific, generative AI skills that are tailored to specific use cases. For example, a skill might summarize a customer chat, suggest next steps in a workflow, or generate resolution notes for a task. For a list of available skills by product, see Now Assist skills.

Use the Now Assist Admin console to install and configure Now Assist plugins, as well as perform the following tasks:

- Set up and configure skills to align with your business processes.
- Monitor and analyze usage, including adoption rates and performance metrics across the platform.
- Manage provider policies, data handling rules, and access controls.
- Turn on and configure the Now Assist panel.
- Manage large language models (LLMs).
- Configure translation for Now Assist.

If you need to customize or adapt existing skills, you can use the Now Assist Skill Kit to build, customize, and test generative AI skills that are tailored to your organization's workflows. With features like a prompt editor, visual builder, and evaluation tools, it helps you fine-tune

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

performance and ensure quality. It also integrates with the Now Assist Data Kit and Now Assist Guardian for testing and governance, making it a flexible and secure way to extend AI across the ServiceNow platform.

Roles and responsibilities:

Now Assist Admin [sn_nowassist_admin.nsa_admin]

This role is required to install Now Assist plugins, activate skills, and configure Now Assist preferences and settings.

Skill Kit admin [sn_skill_builder.admin]

This role is required to create, update, and publish skills in Now Assist Skill Kit.

For more information about these tools and features, see:

- Overview tab in Now Assist Admin
  - Now Assist panel
  - Manage AI models
  - Analyzing Now Assist performance
  - Multilingual service for Now Assist
  - Now Assist for Mobile
- Now Assist Skill Kit

Agentic AI

In an agentic AI system, autonomous agents perform complex tasks by reasoning, planning, and executing actions across multiple steps. These agents are not just reactive; they can break down goals, make decisions, and coordinate with other agents or systems to complete workflows. In AI Agent Studio, AI agents and agentic workflows are built to handle tasks such as resolving incidents, managing requests, or navigating service catalogs using generative AI and orchestration logic. For a list of available agentic workflows by product, see Now Assist agentic workflows.

Use AI Agent Studio to perform the following tasks:

- Build or modify AI agents and agentic workflows using intuitive list views and guided setup tools.
- Simulate tasks and view execution logic to refine agent behavior.
- Access execution logs and version control to track performance and changes.
- Use the Settings page to enable Now Assist Guardian for runtime safety, including detection of offensive content and prompt injection.
- Use the AI Agent Analytics dashboard to view usage stats, efficiency gains, and execution volume to help measure value.

The Model Context Protocol (MCP) client enables you to access the Model Context Protocol tools that are hosted externally and published using an MCP Server. MCP facilitates communication between an AI host application (like AI Agent Studio) and one or more MCP servers that expose specific capabilities such as tools.

The MCP Server Console is a tool for administrators to set up and manage Model Context Protocol (MCP) servers. These servers provide AI applications with access to ServiceNow features like incident lookups, case summaries, and workflow actions. The console lets you choose which capabilities to make available, ensuring everything stays secure and compliant.

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

Then, from an MCP client, users can request information from the server and automate functionality using the available tools and data that the server returns.

Roles and responsibilities:

AI Agent Admin [sn_aia_admin]

This role is required to create, manage, update, test, and delete AI agents and agentic workflows.

MCP Client Admin [sn_mcp_client.admin]

This role is required to administer the MCP client in AI Agent Studio. It is inherited from the AI Agent Admin [sn_aia.admin] role.

MCP Server administrator [sn_mcp_server.admin]

This role is required to administer the MCP Server Console.

For more information about these tools, see:

- Now Assist AI agents
  - AI Agent Studio
  - Implement access control in Now Assist AI agents
  - AI Agent Analytics dashboard
  - Examples of using AI agents
- Model Context Protocol Client
- MCP Server Console

Conversational AI engagement

Now Assist in Virtual Agent acts as the conversational front door to your organization's AI capabilities. It enables users to interact naturally with AI-powered assistants that understand intent, recommend actions, and deliver synthesized responses—all within familiar channels like the service portal, mobile app, and Microsoft Teams. This conversational layer connects users to a wide range of AI-powered resources, including AI agents, generative AI skills, catalog items, Virtual Agent topics, Q&A Genius results from AI Search, and even custom subflows and actions.

By integrating with the AI agent framework, Virtual Agent can discover and trigger agentic workflows that dynamically solve problems based on context. For example, a user might ask about ordering a device, and the system can invoke a multi-turn catalog ordering skill, refine the request, and complete the task, all in one conversation. Similarly, Q&A Genius capabilities allow users to ask open-ended questions and receive actionable answers, links, or follow-up options, improving deflection rates and reducing the need for live agent intervention.

Now Assist in Virtual Agent also makes it easier to discover content without needing to know exact keywords. Users can search for and interact with catalog items, AI agents, and Virtual Agent topics using natural language. LLM topic discovery and Knowledge Graph slot-filling can return personalized responses and reduce repetitive questions.

To configure Now Assist in Virtual Agent, you set up an assistant. An assistant is a container for a specific conversational experience that can be delivered in portals on your instance, mobile apps, Microsoft Teams, and Slack. Use the Assistant Designer to create, configure, and manage these assistants, as well as link them to conversational assets in Virtual Agent.

To implement Now Assist in Virtual Agent, admins use guided setup that requires no deep technical expertise. Assistant configuration includes the following options:

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

- AI agent enablement and orchestration
- Search sources
- Knowledge Graph schema usage
- Conversational assets, such as Virtual Agent topics, subflows, actions, generative AI skills, and AI agents
- The portals, channels, or apps the assistant will be used in
- Branding: the look and feel of the assistant
- Chat experience, such as greeting, closing, and fallback messages

You can also customize an assistant for the Now Assist panel. The Now Assist panel is available in workspace experiences and is typically used for contextual assistance within a specific workflow or record. Now Assist panel assistants have different configuration options and skill availability.

Roles and responsibilities:

Virtual Agent [virtual_agent_admin]

This role, or the admin role, is required to install Now Assist in Virtual Agent and set up and manage assistants in Assistant Designer. It is also required for users who create and manage Virtual Agent topics and assets.

Now Assist Admin [sn_nowassist_admin.nsa_admin]

This role is required to turn on and configure the Now Assist panel in the Now Assist Admin console.

For more information about these tools, see:

- Now Assist in Virtual Agent
  - Configuring assistants overview
  - Enhanced chat
  - Integrating Now Assist in Virtual Agent with Microsoft Teams
  - Using Now Assist in Virtual Agent conversations with Slack
  - Analyzing Now Assist in Virtual Agent
- Now Assist for Mobile
- Now Assist panel

AI governance for Now Assist on the ServiceNow AI Platform

As organizations increasingly adopt AI to drive efficiency, innovation, and customer experience, AI governance becomes essential to ensure responsible use, regulatory compliance, and alignment with enterprise goals. Now Assist provides a comprehensive governance framework through key roles and applications that work together to manage AI across its life cycle.

AI policy considerations

The following policy considerations shape how AI is deployed, monitored, and maintained across the enterprise.

Data security and privacy

AI systems must comply with strict data handling protocols to protect sensitive information. This includes the following:

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

- Data classification rules for personally identifiable information, protected health information, and financial data.
- Encryption standards for data in transit and at rest.
- Data residency and sovereignty restrictions, which determine where data can be stored and processed.
- Retention and deletion policies that govern how long data is kept and when it must be purged.

Admins can configure Data Privacy for Now Assist to mask sensitive fields and control what is shared with third-party models. For details, see Data Privacy for Now Assist.

Compliance and regulations

AI deployments must adhere to a range of regulatory frameworks, including:

- HIPAA, PCI DSS, GDPR, CCPA, and FedRAMP, depending on the industry and geography.
- Third-party/vendor risk management, especially when external models or services are used.
- Logging and traceability requirements help ensure accountability in AI decisions.

Legal reviews are often required before publishing documentation or releasing features, particularly when consolidating overlapping control objectives or addressing model transparency.

Responsible AI use

To ensure ethical and effective AI, organizations should enforce the following:

- Model approval and usage guidelines, including naming conventions and branding policies for AI agents.
- Bias and fairness safeguards, with AI Stewards evaluating risks like hallucination or exposure of sensitive data.
- Human oversight requirements, ensuring AI augments rather than replaces human judgment.
- Transparency obligations, such as disclosing the use of third-party models like Azure OpenAI in product documentation.

Governance and change management

AI governance is supported by structured oversight and change control processes:

- Definition, review, and approval of enterprise-wide guardrails and new use cases.
- Change control and rollout processes ensure that AI features are deployed safely and predictably.
- Incident response and escalation plans are in place to address issues such as data breaches or model failures.

AI policy stakeholders

The following groups set and execute AI policy in an organization:

Policy setters

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

The Chief Information Officer (CIO) or Chief Technology Officer (CTO) sets the overall technology strategy, ensuring AI initiatives align with enterprise architecture and innovation goals. The Chief Information Security Officer (CISO) establishes data security and encryption standards to safeguard sensitive information across AI workflows. The Chief Data Officer (CDO) oversees data usage and governance, ensuring that AI systems handle data ethically and in accordance with organizational policies. Meanwhile, the Chief Privacy Officer and legal teams are responsible for regulatory compliance, ensuring that AI deployments meet requirements such as GDPR, HIPAA, and other jurisdictional or industry-specific standards. Together, these leaders form the foundation of AI governance, guiding implementation teams and administrators in deploying AI responsibly and securely.

Internal governance and oversight

Governance and oversight of AI in Now Assist is led by structured groups that define and enforce responsible use. An AI Governance Committee and Data Governance Council set enterprise-wide guardrails for AI, including standards for privacy, fairness, and compliance, and are responsible for reviewing and approving new AI use cases. Supporting these bodies, the AI Steward ensures that AI is used responsibly across workflows, overseeing data quality, managing risks such as bias or data exposure, and monitoring adherence to regulatory requirements. Additionally, AI Stewards monitor regulatory compliance, assess performance and user feedback, and work with admins and developers to optimize AI automation while minimizing risk.

Implementation and operations

Implementation and operations teams are responsible for securely deploying and managing AI features in alignment with governance policies. The Now Assist admin configures and manages Now Assist capabilities, ensuring that AI features are properly mapped to workflows and governed according to enterprise standards. Platform owners and ServiceNow admins oversee the deployment process, making sure that all configurations comply with established policies and technical requirements. Meanwhile, AI developers build, extend, and integrate AI features into business workflows, working closely with admins and platform teams to deliver scalable, compliant, and effective AI solutions. Together, these roles translate governance policies into secure, functional AI implementations.

For more information about AI governance user roles, see AI Risk and Compliance roles and Assign the data steward role.

For more resources about AI governance, see the following Best Practices topics:

- AI Governance White Paper
- Technical Governance Foundation Template
- Governance in operational phase: Roles and responsibilities

AI governance tools

Now Assist governance is specified in the following tools:

AI Control Tower

The AI Control Tower functions as the central hub for AI strategy, governance, and analytics within Now Assist. It offers enterprise-wide visibility into AI assets, usage patterns, and compliance status, enabling organizations to maintain oversight and accountability. Through automated discovery and inventory of approved AI assets, it streamlines asset management while embedding governance checks and compliance alerts to ensure that all AI deployments remain secure and aligned with organizational policies.

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

For more information, see AI Control Tower.

Now Assist Admin console

The Now Assist Admin console is key to managing AI governance by configuring policies, enforcing data handling rules, and ensuring compliance with security and privacy standards. Admins oversee provider policies at the skill level, track usage analytics like success rates and adoption, and collaborate with AI stewards and business SMEs to align AI with organizational goals. They also connect governance committees with technical teams to support smooth policy execution.

For more information, see Overview tab in Now Assist Admin.

Data readiness for implementing Now Assist on the ServiceNow AI Platform

High-quality data that is complete, accurate, and contextually relevant is the foundation for delivering precise, meaningful, and trustworthy AI responses.

Now Assist requires high-quality data

To unlock the full potential of Now Assist on your instance, the quality of your data is paramount. For AI to deliver accurate, context-aware, and actionable outputs, it must be trained and operate on high-quality data that is complete, consistent, and structured. When your data is well-prepared, Now Assist can interpret user queries with greater accuracy, enabling faster resolutions and more effective self-service experiences.

Whether it's summarizing complex incidents, generating resolution notes, or creating knowledge articles, Now Assist relies on detailed records that reflect the full lifecycle of a task. Short or incomplete cases often lack the depth needed for meaningful AI interpretation, which can result in vague or unhelpful responses. The importance of data readiness extends beyond technical accuracy—it directly impacts user trust and adoption.

Clean, ready data also accelerates implementation. It minimizes the need for rework, reduces deployment friction, and ensures that AI features like summarization, recommendations, and workflow guidance operate smoothly from day one. This leads to higher ticket deflection rates and improved operational efficiency, allowing teams to focus on strategic tasks rather than repetitive support.

Moreover, high-quality data fosters trust in AI outputs. When users consistently receive reliable and context-aware responses, their confidence in the system grows—driving adoption and maximizing return on investment. Ultimately, investing in data quality is an investment in user satisfaction, AI performance, and long-term success.

Follow these tips to assess your organization's data readiness:

- Audit task records for completeness and clarity, and ensure resolution fields are consistently filled out with clear summaries.
- Avoid using vague or generic language in task descriptions and updates.
- Maintain a clean and structured knowledge base and ensure that knowledge articles are linked to resolved cases.

  For details, see Knowledge Base readiness for Now Assist on the ServiceNow AI Platform.

- Audit Service Catalog items.

  For details, see Service Catalog readiness for Now Assist on the ServiceNow AI Platform.

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

- Use the Now Assist Readiness Evaluation app to assess catalog, case, and knowledge data.
- Use the Now Assist Data Kit to curate and cleanse data sources.
- Align stakeholders (data owners, product managers, engineers) around shared standards for AI-ready data.

  For details, see AI governance for Now Assist on the ServiceNow AI Platform.

For more information, see Now Assist Data Readiness Checklist.

Now Assist Readiness Evaluation app

Data readiness assessments can be time-consuming and manual, especially when evaluating whether catalog items are conversational or if knowledge articles are embedded in inaccessible formats like PDFs. The Now Assist Readiness Evaluation app helps automate this process by analyzing service catalog entries, cases, and incidents, and then providing actionable recommendations to prepare data for AI use. It also enables you to assess whether updates, installations, or customizations of your instance could affect implementation. The assessments provide direct hyperlinks to improve any issues found.

For more information, see Now Assist Readiness Evaluation.

Install Now Assist Readiness Evaluation by requesting it from the ServiceNow Store. Visit the ServiceNow Store 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.

Now Assist Data Kit

If the base system Now Assist skills don't fit your needs, use the Now Assist Data Kit to curate, cleanse, and manage data for AI evaluations. You can create custom datasets and data collections that can be used in Now Assist Skill Kit for evaluation. For more information, see Now Assist Data Kit.

Application readiness for Now Assist on the ServiceNow AI Platform

Now Assist leverages the power of the ServiceNow AI Platform to deliver AI solutions. Ensure that your instance is ready to take advantage of AI capabilities by preparing Platform applications, such as Service Catalog and AI Search.

Installing a Now Assist product such as Now Assist for IT Service Management (ITSM) gives you access to AI capabilities in Platform applications as well as capabilities specific to your product workflow. It's important to ensure that the appropriate Platform applications are ready for Now Assist. To prepare your instance for Now Assist, review the following topics:

- Knowledge Base readiness for Now Assist on the ServiceNow AI Platform
- Service Catalog readiness for Now Assist on the ServiceNow AI Platform
- AI Search readiness for Now Assist on the ServiceNow AI Platform
- Now Assist in Virtual Agent readiness on the ServiceNow AI Platform

After ensuring that your AI policy, your data, and your applications are ready, you can begin to implement Now Assist.

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

1. Ensure entitlement for at least one Now Assist application. For a list of available applications, see Exploring Now Assist Admin.

   Note: Some Now Assist products may require other entitlements.

2. In the ServiceNow Store, check version compatibilities and dependencies for Now Assist apps.

   For more information, see Evaluating version requirements and dependencies.

3. Install Now Assist products from the Now Assist Admin console.

   For more information, see Install Now Assist plugins.

Customization and Now Assist

Customizing your ServiceNow instance is often necessary to meet unique business needs, but it can introduce unintended consequences for Now Assist features. Changes to field names, UI actions, workflows, or table structures may disrupt how Now Assist skills and agents operate. For example, modifying field states or labels can interfere with skill conditions and input mapping, while altering default UI actions may prevent agents from triggering correctly. Custom resolution workflows might conflict with the logic embedded in native skills, and table-level variations can affect where and how Now Assist functions across workflows. These customizations can also create upgrade friction, requiring additional testing, rework, or reconfiguration to ensure compatibility with new platform releases.

To mitigate these risks, you can use tools like the Now Assist Readiness Evaluation app to identify high-impact customizations and plan for remediation before deploying or upgrading AI features. For more information, see Now Assist Readiness Evaluation.

To maintain compatibility with Now Assist while accommodating custom configurations, several mitigation strategies are available:

- Duplicate skills before adapting them.
- Adjust skill inputs and role conditions to reflect organizational processes.
- Add additional input sources, such as related tables or emails.

For more advanced needs, build custom generative AI skills tailored to your unique business requirements, allowing for deeper integration without compromising core functionality.

For more information about customizing AI solutions, see How to approach building custom generative AI solutions using Now Assist in the ServiceNow Community.

Knowledge Base readiness for Now Assist on the ServiceNow AI Platform

The knowledge base is the engine that enables Now Assist to deliver intelligent, accurate, and context-aware responses across AI Search, Q&A Genius Results, and other AI-powered experiences.

When users ask questions, Now Assist taps into your knowledge articles to generate answers that reflect your organization's expertise and policies. This means the quality, clarity, and structure of your knowledge content directly influence the effectiveness of AI responses. Outdated, inconsistent, or poorly written articles can lead to incorrect outputs and user frustration. That's why maintaining a clean, well-organized knowledge base is critical—not just for documentation, but for enabling reliable AI performance.

The Knowledge Management Overview dashboard is a powerful tool for assessing AI readiness. It provides key usage metrics such as article views, helpfulness ratings, and feedback trends,

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

helping content owners identify high-performing articles and prioritize updates to low-usage or outdated content. For more information, see Knowledge Management Overview dashboard.

High-level checklist

1. Review existing KB articles and knowledge bases

Get the total count of published KB articles and the number of active knowledge bases.

Why? To scope clean-up and identify high-priority content.

See: Knowledge Management Overview dashboard

2. Review the structural formatting of KB articles

Check for consistent template use (KCS-based, Q&A, Issue-Cause-Resolution, etc.). Ensure that the Knowledge Management Advanced plugin is active.

If you're using custom templates, determine whether they follow a consistent format. Also ensure that key steps are clearly defined and consistently populated.

Why? Using templates consistently improves LLM understanding and summarization.

See:

- Knowledge article templates
- Activate the Knowledge Management Advanced plugin

3. Ensure metadata completeness in KB articles

Make sure that fields such as Meta, Meta Description, Category, Author, and Validity are populated.

Why? This practice enables precise AI filtering and targeting.

See:

- Create a knowledge article
- Best practices to use your knowledge articles with Now Assist (generative AI) (ServiceNow Community)

4. Perform maintenance audits

Flag KB articles that have not been updated in a year or more, that have no metadata, are duplicates, or are about to expire.

Why? Outdated content degrades user trust and AI performance.

See:

- Knowledge Management Overview dashboard
- Retire a knowledge article
- Identify and review duplicate Knowledge articles (requires Now Assist in Knowledge Management)

5. Review top searched queries

Identify search queries and articles for optimization, and limit Now Assist Q&A functionality to high-quality, optimized KB articles.

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

Why? This prevents low-value content from surfacing in answers.

See: Knowledge Management dashboard

Service Catalog readiness for Now Assist on the ServiceNow AI Platform

A well-structured Service Catalog is essential to unlocking the full potential of Now Assist. As the backbone of many self-service workflows, the catalog enables Now Assist to interpret user requests accurately, present the right options, and minimize friction in the experience.

When catalog items are clearly defined and conversationally enabled, Now Assist in Virtual Agent can surface them more effectively, leading to higher self-service rates and better issue deflection.

However, catalogs that are cluttered, inconsistent, or overly complex make it difficult for AI to parse and respond appropriately. That's why auditing your existing catalog is a critical first step. By identifying which items are AI-ready and optimizing those that aren't, you ensure that Now Assist can deliver intelligent, context-aware responses.

High-level checklist

1. Install the Catalog Conversational Coverage (sn_catalog_con_cov) plugin

This plugin gives you access to the conversational catalog overview dashboard.

See: Install an application or plugin

2. Review your catalog inventory

Use the conversational catalog overview dashboard to review all active catalog items. From this list, determine which are most requested and the percentage that are conversational.

Why? To focus improvement efforts on the catalog items with the highest impact.

See:

- Catalog Conversational Coverage
- Conversational catalog overview dashboard

3. Determine which catalogs can be made conversational

Identify catalog items with fewer than 15 variables. (Items with more than 15 variables are better suited to be submitted as pop-up links to forms.)

> **Note:** Conversational catalog items require Now Assist in Virtual Agent. For details, see Now Assist in Virtual Agent readiness on the ServiceNow AI Platform.

Why? These are easier to convert and provide quick AI wins.

See: Service catalog variables

4. Analyze catalog item metadata

Ensure catalog items have clear, user-friendly names, tooltips, metadata, and categories. Using templates can help with this as well.

Why? This prevents LLM confusion and improves conversational flow.

See:

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

- Service Catalog properties
- Edit a catalog item in Catalog Builder
- Create a catalog item using a template

5. Review advanced configurations

Review UI policies, client scripts, and user criteria.

Why? User criteria determines who can request an item, but policies and scripts are ignored by Virtual Agent.

See:

- Service catalog UI policy
- Set security for items and categories
- Create a user criteria record in Service Catalog
- Create a Service Catalog client script

Tips

- Use Now Assist in Catalog Builder to generate catalog items.
- Limit the number of variables. Fewer is better.
- For form-based requests, select the Make item non-conversational in VA check box.
- Provide clear context for the LLM with distinct names, labels, and tooltips.
- Minimize the use of custom variable types. Use the service catalog variables included on the platform instead.
- Limit client-side scripting and use regular expressions for validation.
- Avoid complex dependencies. Simplify variable relationships to ensure a smooth conversation flow.
- Test and iterate: make changes based on feedback.

For more information about conversational catalogs in AI, see the following ServiceNow Community articles:

- Guidance for making catalog items conversational
- How to request catalog items in Now Assist in Virtual Agent

AI Search readiness for Now Assist on the ServiceNow AI Platform

Now Assist in AI Search is built directly on the robust foundation of ServiceNow® AI Search. AI Search provides the underlying infrastructure that enables Now Assist to retrieve and rank enterprise content—such as knowledge articles, records, and documentation—based on relevance and access permissions.

Because AI Search adheres to ServiceNow security models, responses are always grounded in content the user is authorized to access. Additionally, AI Search admin tools give you control over relevancy tuning, allowing you to refine how content is surfaced and ensure that Now Assist delivers the most useful and trustworthy results.

Every AI-generated answer includes citations and references, allowing users to see exactly where the information came from. This not only boosts confidence in the accuracy of responses,

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

but also reduces the risk of hallucination by grounding outputs in your organization's actual content.

To maximize its impact, Now Assist in AI Search should be activated before Now Assist is more broadly implemented. Once enabled, it powers actionable Q&A through Genius Results cards in global search, Service Portal, Employee Center, and Virtual Agent.

High-level checklist

1. Verify AI Search status

Make sure AI Search is active on your instance. To check status, navigate to **All** > **AI Search** > **AI Search Status**. If AI Search is not active, select the **Request AI Search** button to initiate activation.

See:

- Configuring AI Search
- Enable AI Search in supported ServiceNow AI Platform® applications

2. Review search sources

Configure search sources to determine what data will be searched by user queries.

See:

- Control access to searchable content using search sources
- Create a search source for AI Search

3. Configure search sources for external content

Allows Now Assist Q&A search to include external content.

See:

- Indexing and searching external content in AI Search
- Create an indexed source
- Configuring External Content Connectors

4. Enable AI Search on one or more portals

Display AI Search results in the portals on your instance.

See:

- Search application configurations
- Enabling and configuring AI Search in ServiceNow AI Platform® applications

5. Install Now Assist in AI Search

Install a Now Assist product in the Now Assist Admin console.

See: Install Now Assist in AI Search

6. Enable Now Assist Q&A Genius Results in AI Search portals

Specify the Now Assist Genius Result types you want to make available in each of your AI Search portals. Now Assist Multi-Content Response Genius Results use an LLM to generate search and chat responses that synthesize information

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

from knowledge articles, Service Catalog items, and other available content types. Answers are presented in Q&A format.

See:

- Enable Now Assist Genius Results in AI Search portals and mobile applications
- Now Assist Q&A Genius Results

7. Verify AI Search widget customizations

Customized AI Search widgets may not correctly show Q&A search results even when configured for Now Assist, since customizations get skipped during upgrades.

See:

The ServiceNow AI Platform® offers a variety of search tools, which may return different answers for the same or similar searches. This disparity in results is expected. For more information, see Search result disparities between AI Search and Now Assist search features.

Tip: For more information, see Now Assist in AI Search FAQ in ServiceNow Community.

Now Assist in Virtual Agent readiness on the ServiceNow AI Platform

Now Assist enhances Virtual Agent with AI-driven capabilities that understand natural language, guide users through complex tasks, and deliver high-confidence answers without relying on rigid keyword matching or manual configurations.

Now Assist in Virtual Agent provides the following features:

- AI asset discovery

  Say goodbye to time-consuming keyword or NLU configurations. Now Assist uses LLMs to automatically discover and match user intents to Virtual Agent topics and other AI assets, including generative AI skills, AI agents and agentic workflows, and subflows and actions.

- Simplified deployment

  Using LLM-powered Virtual Agent topics, teams can accelerate rollout and improve conversation quality. This means less effort spent on manual tuning and more time delivering value.

- AI Search Genius Results

  Users receive curated, actionable responses via Genius cards, which contain summarized knowledge with direct actions such as `Request this item`.

- Conversational catalog ordering

  Users can request Service Catalog items using natural conversation. Virtual Agent asks clarifying questions and confirms the user's intent before completing the request.

  Note: Service Catalog items must be marked as conversational to work with Virtual Agent. For details, see Catalog item conversational details page overview.

- Multi-turn Q&A

  Follow-up questions are handled seamlessly, allowing users to refine their queries and get better answers.

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

Setting up Now Assist in Virtual Agent requires customizing or creating a new LLM assistant. You can assign an assistant to one or more portals. If LLM Virtual Agent topics aren't associated with an LLM assistant, they aren't discoverable.

High-level checklist

1. Install Now Assist in Virtual Agent

You can install it from the Conversational Interfaces admin console once you have installed a Now Assist product such as Now Assist for IT Service Management (ITSM).

To set up Now Assist in Virtual Agent, you configure an assistant.

See: Configuring assistants overview

2. Review your Virtual Agent topic inventory

Review your topics and identify high-volume user intents. You can use the Conversational Analytics dashboard and Automation Discovery reports.

Why? This helps you identify the top self-solve opportunities in Virtual Agent.

See:

- Conversational Analytics dashboard in Platform Analytics experience
- Create an Automation Discovery report

Review your knowledge base

Identify KB articles that can self-serve any of the top intents you identified.

Why? This simplifies topic management and enables self-service.

See: Knowledge Base readiness for Now Assist on the ServiceNow AI Platform

4. Migrate NLU topics to LLM

Use the topic migration tool in Virtual Agent to convert NLU topics to LLM.

Why? Leverage existing Virtual Agent topics with minimal effort.

See: Migrating NLU/keyword Virtual Agent topics to LLM topics

5. Review Service Catalog items

Identify self-serve catalog items that can be replaced with existing LLM topics.

Why? To avoid redundancy and eliminate the need to create new Virtual Agent topics.

See: Service Catalog readiness for Now Assist on the ServiceNow AI Platform

Review LLM Virtual Agent topics that come with Now Assist

Use these LLM topics as a starting point for Virtual Agent topic creation.

Why? New LLM versions of older NLU Virtual Agent topics reduce rework.

See: ITSM Virtual Agent pre-built LLM topics

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

Tips

- When migrating legacy NLU topics, ensure that you optimize topic descriptions so that the topic is clearly described and aligned with the intent and expected results.

  For details, see LLM description and instruction guidelines for Virtual Agent topics.

- You can customize the look of your assistant and the chat experience during guided setup.

  For details, see Brand an assistant.

- You can choose the chat experience you want for each assistant:
  - Standard chat
  - Enhanced chat
- You can integrate Now Assist in Virtual Agent with Microsoft Teams.

  For details, see Integrating Now Assist in Virtual Agent with Microsoft Teams.

For more information about conversational catalogs in AI, see the following information from ServiceNow Community and YouTube:

- How to request catalog items in Now Assist in Virtual Agent
- Microsoft Copilot integration with Now Assist FAQ - Zurich release
- AI Academy: Enhanced chat experience with Now Assist in Virtual Agent

Solving installation and configuration issues with Now Assist

Use this checklist to address issues or gaps in your Now Assist configuration.

Common issues by Now Assist feature

Feature	Issue	Solution
AI agents	I can't find agent spokes	Install the Now Assist for Spokes application from the ServiceNow Store. For details, see Now Assist for Integration Hub.
AI Control Tower	No data on the Evaluations tab	Verify that the Conversation Evaluator [sn_na_conv_eval] plugin is active. Also verify that the Smooth Flowing Conversation Chat Eval skill is active in the Now Assist Admin console. (This skill may appear in the Platform workflow or in the Other workflow, depending on your instance.)
AI Search	External content isn't included in search results	You can use External Content Connectors to include other sources in your search results. For details, see External Content Connectors.
AI Search	Now Assist in AI Search Genius results don't appear	1. Verify that AI Search and Now Assist in AI Search are installed and configured. 2. To use Now Assist Q&A Genius Results in AI Search applications, link the Now Assist Q&A Genius Result configuration to your search profiles for those applications. For details, see Link a Genius Result configuration to a search profile.

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

Common issues by Now Assist feature (continued)

Feature	Issue	Solution
		3. Verify that all Now Assist plugins are up to date. **4.** Try repairing the plugins.
Knowledge articles	I can't generate a knowledge article	**1.** Verify that Now Assist in Knowledge Management is installed and configured. **2.** Verify that all Now Assist plugins are up to date. **3.** Verify that the knowledge generation skill is activated. For details, see Activate a Now Assist skill. **4.** Try repairing the plugins.
Knowledge articles	The KB generation popup doesn't appear in Core UI	Now Assist in Knowledge Management formatters may be missing from the form layout, possibly due to form customizations. For details, see KB1710178.
Now Assist applications	I want to view my subscriptions	See Product subscriptions overview.
Now Assist panel	Can't access skills in the Now Assist panel	• Verify that the skill is configured to display in the Now Assist panel. For details, see Edit a Now Assist skill. • If you are using Now Assist in Virtual Agent, verify that search sources were configured for the Now Assist panel. You can specify search sources for a Now Assist panel assistant when you set it up. Search sources are essential for the Now Assist panel and Virtual Agent. Without them, they cannot discover or rank skills and agentic workflows. For details, see Configuring assistants overview and Assign search sources to a chat assistant. If Now Assist in Virtual Agent is not installed, the Now Assist panel uses default search sources.
Now Assist panel	Some users don't get a response to a question	Many Now Assist skills require specific user roles. Verify that there are active skills with the user's role. For details, see Activate a Now Assist skill.
Now Assist panel	I don't want skills to be available in Now Assist panel	Disable the Now Assist panel when you configure the skill. For details, see Edit a Now Assist skill.
Now Assist panel	Options for Now Assist panel are	The sn_nowassist_admin.user role provides read-only access only. To make configuration changes, the user must have the sn_nowassist_admin.nsa_admin role.

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

Common issues by Now Assist feature (continued)

Feature	Issue	Solution
	grayed out in the Now Assist Admin console	
Now Assist setup	Errors after clone	Try repairing affected plugins. For details, see Repair a ServiceNow application.
Now Assist setup	The Q&A results skill is not in the Now Assist Admin console	Set up AI Search. For details, see Configuring AI Search.
Now Assist setup	Features/ skills are missing or generally not working	• Verify that all of your Now Assist plugins are up to date. For details, see Install an update to a ServiceNow Store application. • Verify that version and dependency requirements are met. For details, see Evaluating version requirements and dependencies. • For skills, verify that they are active in the Now Assist Admin console. For details, see Activate a Now Assist skill. • Verify that the user has the correct role for the skill. • Try clearing the cookies and cache in the web browser. • Try repairing Generative AI Controller. For details, see Repair a ServiceNow application.
Now Assist setup	Can't edit a skill	Verify that you have the Now Assist Admin role: sn_nowassist_admin.nsa_admin. You can edit a skill or make a copy of a skill to edit. For details, see Edit a Now Assist skill and Make a copy of a Now Assist skill.
Now Assist setup	Missing entries, fields, and errors	• Verify that all of your Now Assist plugins are up to date. For details, see Install an update to a ServiceNow Store application. • Try repairing the application. For details, see Repair a ServiceNow application. • Try repairing Generative AI Controller. For details, see Repair a ServiceNow application.

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

Common issues by Now Assist feature (continued)

Feature	Issue	Solution
Now Assist setup	Problems upgrading Now Assist applications	• Verify that you have a license for the application. • If the application was not previously installed, request it from the ServiceNow Store (Opt In). • If the application was previously installed, you may need to procure it from the ServiceNow Store again. For details, see Updating applications.
Now Assist setup	Skills not working in Legacy Agent Workspace	Now Assist skills are not available in Legacy Workspace. Upgrade to Service Operations Workspace to use Now Assist.
Now Assist for Code	Don't see code suggestions	• Verify that the skill is active. For details, see Activate a Now Assist skill. • Verify that the user has the appropriate role. Any authenticated builder can use the related active skill. • Verify that autocomplete is enabled. For details, see Generate code with autocomplete.
Now Assist for Creator	Skills not available or not working	• Verify that the skill is active in the Now Assist Admin console. • Verify that the user has the appropriate role. Any authenticated builder can use the related active skill.
Now Assist Skill Kit	Now Assist Skill Kit is not visible on my instance	• Verify that you have at least one Now Assist product installed. For details, see Install Now Assist plugins. • Verify that the sn_skill_builder.admin role is assigned to the user.
Now Assist Skill Kit	Skills are read-only	Verify that the sn_skill_builder.admin role is assigned to the user.
Now Assist in Virtual Agent	Topic not returning as expected	• Verify that the Virtual Agent topic is in the Active state and is published. For details, see Publish a Virtual Agent topic. • Verify that the topics are using LLM topic discovery. NLU/keyword topics cannot be used in a portal that is using Now Assist in Virtual Agent. You can migrate NLU/keyword topics to LLM, however. For details, see Migrating NLU/keyword Virtual Agent topics to LLM topics.
Now Assist in Virtual Agent	Can't add topics to the portal	Verify that the topics are using LLM topic discovery. NLU/keyword topics cannot be added to a portal that is using Now Assist in Virtual Agent. You can migrate these topics to LLM, however. For details, see Migrating NLU/keyword Virtual Agent topics to LLM topics.

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

Common issues by Now Assist feature (continued)

Feature	Issue	Solution
Now Assist in Virtual Agent	Configuration issues	1. Verify that the plugin is up to date. For details, see Install an update to a ServiceNow Store application. 2. Follow the guided setup to install and configure it. For details, see Configuring assistants overview.
Now Assist in Virtual Agent	Unable to choose fallback options in setup	If you're using a custom fallback topic, this option may be unavailable in guided setup. For details, see KB1760362: Now Assist Chat Setup Not Allowing to Set Fallback Option. You may need to log in to view the article.
Now Assist in Virtual Agent	Chat is not showing search results	• Verify that all of your Now Assist plugins are up to date. For details, see Install an update to a ServiceNow Store application. • Verify that AI Search and Now Assist in AI Search are set up and configured. For details, see Configuring AI Search and Install Now Assist in AI Search. • Verify that search sources were configured for the Virtual Agent assistant. You can specify search sources for an assistant when you set it up. Search sources are essential for Virtual Agent and the Now Assist panel. Without them, they cannot discover or rank skills and agentic workflows. For details, see Configuring assistants overview and Assign search sources to a chat assistant.
Now Assist in Virtual Agent	Error when attempting to migrate Virtual Agent NLU conversations to LLM	Verify that you have the correct role, either virtual_agent_admin or sn_vad_genai.topic_migration_admin. For more information, see Migrate NLU topics to LLM topics.

Now Assist AI assets

The Now Assist AI experience includes generative AI skills, AI agents, and AI agentic workflows. These components work alone or in combination to help achieve efficiencies and results on your instance.

The ServiceNow AI Platform now brings you an AI native experience with three licensing tiers available:

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

For more information, see ServiceNow product tiers.

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

Get started

Now Assist skills	Now Assist agentic workflows	Platform AI agents
		
Now Assist skills are predefined capabilities that use generative AI to perform tasks such as generating summaries, resolution notes, and so forth.	An agentic workflow is a structured sequence of tasks executed by one or more AI agents with minimal human intervention to fulfill a business objective. You can create and manage these workflows in AI Agent Studio.	An agent is a software entity that can perceive its environment, make decisions, and take actions to achieve specific goals. They are often used to automate tasks, interact with users, or manage complex processes.

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

For a list of AI assets that are turned on by default, see Now Assist skills, agents, and agentic workflows on by default.

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.

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.

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

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.

Now Assist skills, agents, and agentic workflows on by default

Starting with the Zurich Patch 4 release, some Now Assist skills, agents, and agentic workflows are turned on by default.

Important: Some Now Assist skills, agents, and agentic workflows are turned on by default. The default behavior works as follows:

New customers

When you install a Now Assist product, designated skills, agents, or agentic workflows are turned on automatically.

Existing customers who are upgrading (starting with Australia Patch 4)

There is no change to skills, agents, or agentic workflows that are currently enabled and customized.

An AI asset is turned on if:

- The Now Assist plugin is installed, but the asset was never turned on.
- An admin has never adjusted roles for the skill.

An AI asset is not turned on if:

- The asset was previously turned on, and then turned off again.
- An admin has adjusted roles for the asset.

For more information, see Now Assist skills, agents, and agentic workflows on by default.

Note: Some workflow skills support Now Assist functionality. Deactivating these skills may negatively impact some features.

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

Now Assist skills

Now Assist products provide generative AI skills that are tailored to meet the needs of users in different workflows.

The following sections describe the available Now Assist skills.

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

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

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.)

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

Note: Some workflow skills support Now Assist functionality. Deactivating these skills may negatively impact some features.

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

Available skills by workflow

Workflow	Product	Available skills
Technology	Now Assist for Collaborative Work Management (CWM)	• Acceptance criteria generation • Docs summarization • Doc generation • Task generation
Technology	Now Assist for Configuration Management Database (CMDB)	• Configuration item (CI) summarization • Manage duplicate CIs • Service Graph Connector diagnosis
Technology	Now Assist for Core Business Suite (CBS)	• Next Best Action • Error Resolution
Technology	Now Assist for Enterprise Architecture (EA)	• ADR DOC summarization and actions • Business application insights • Diagram change analysis • Refine text • Register a business application • Register a digital integration
Technology	Now Assist for Operational Sustainability (formerly ESG)	Extract data from utility invoices
Technology	Now Assist for Hardware Asset Management (HAM)	• Generate hardware asset insights
Technology	Now Assist for Integrated Risk Management (IRM)	• Common control objective creation • Control objective impact analyzer

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

Available skills by workflow (continued)

Workflow	Product	Available skills
		• Generate recommendation for similar control objective • Issue summarization • Recommendations for regulatory alert impacted areas • Recommendation of similar control objectives • Regulatory alert summarization • Regulatory alert impacted citations • Regulatory alert impacted control objectives • Regulatory alert impacted controls • Regulatory alert impacted policies • Risk assessment summarization • Risk event summarization Risk event summarization in the classic UI
Technology	Now Assist for IT Operations Management (ITOM)	• Alert analysis • Alert investigation • Analyze service health • Analyze service observability dashboard • LEAP installer • Service Mapping Candidate • Service mapping candidates Impact
Technology	Now Assist for IT Service Management (ITSM)	• Catalog task summarization • Change request risk explanation • Change request summarization • Chat reply recommendation • Chat summarization • Email recommendation • Incident assist • Incident sentiment analysis • Incident summarization • Investigate boot time issues • Investigate Zoom call quality issues

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

Available skills by workflow (continued)

Workflow	Product	Available skills
		• KB generation • Release notes generation • Request activity response generation • Request summarization • Requested item activity response generation • Requested item summarization • Resolution notes generation • Sidebar discussion summarization • Suggested steps generation
Technology	Now Assist for Operational Technology Manager (OTM)	Search for a related record
Technology	Now Assist for Privacy Management	• Control objective impact analyzer • Common control objective creation • Recommendation of similar control objectives • Risk assessment summary
Technology	Now Assist for Security Incident Response	• Correlation insights generation • Generate content for shift handover • Post-incident analysis • Resolution notes generation • Security incident quality assessment • Security incident recommended actions • Security incident resolution plan • Security incident summarization • Security operations metrics analysis
Technology	Now Assist for Software Asset Management (SAM)	• Error log summarization • Error resolution recommendation • Publisher compliance summarization • Product compliance summarization • Recommended actions

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

Available skills by workflow (continued)

Workflow	Product	Available skills
		• SaaS user resolution • Contract entitlement data extraction
Technology	Now Assist for Strategic Portfolio Management (SPM)	• Create a demand • EAP doc summarization • Identify similar records • Demand summarization • Goal insights • Multi feedback summarization • Planning item doc summarization • Project doc summarization • Project insights generation • Refine records • Story generation • Target generation • Write planning item
Technology	Now Assist for Third-party Risk Management (TPRM)	• TPRM issue summarization • TPRM issue management recommendation
Technology	Now Assist for Vulnerability Response	• Approval recommendation • Now Assist recommendation • SEM insights • SPC setup connector • Suggest vulnerability solutions • Vulnerable item deduplication
Customer	Now Assist for Customer Service Management (CSM)	• Activity response generation • Case summarization • Chat recommendation • Chat summarization • Email recommendation • KB generation • Resolution notes generation

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

Available skills by workflow (continued)

Workflow	Product	Available skills
		• Sentiment analysis case • Sentiment analysis dashboard • Sidebar summarization • Suggested steps generation • Trending topics dashboard
Customer	Now Assist for Field Service Management (FSM)	• KB generation • Sidebar summarization • Work order task summarization
Customer	Now Assist for Financial Services Operations (FSO)	• Case summarization • Disputes intake via Virtual Agent • Customer profile summarization • Customer interaction context summary
Customer	Now Assist for Manufacturing Commercial Operations	Enhance non conformance description
Customer	Now Assist for Order Management	Order summarization
Customer	Now Assist for Public Sector Digital Services (PSDS)	• Government case summarization • Chat summarization
Customer	Now Assist for Telecommunications, Media and Technology (TMT)	• Account onboarding case summarization • Customer play summary • Customer service summary • Engagement summary • Internal play summary • KB generation • Resolution notes generation • • Risk signal and issues summary • Service problem case summarization • Success initiative summary • Test summarization

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

Available skills by workflow (continued)

Workflow	Product	Available skills
		• Touchpoint summary • Transform mapping assist
Employee	Now Assist for Employee Experience	• Requested item summarization for approvals • Request summarization for approvals • Case summarization for approvals
Employee	Now Assist for Health and Safety	Incident summarization
Employee	Now Assist for HR Service Delivery (HRSD)	• Case summarization • Chat reply recommendation • Chat summarization • Email recommendation • KB generation • Employee information summarization • Resolution notes generation • Sentiment analysis for HR case • Sentiment analysis for HR task • Sidebar discussion summarization
Employee	Now Assist for Legal Service Delivery (LSD)	• Conversational intake for Conflict of Interest request • Get category of the legal request • Legal matter summarization • Legal request summarization • Triage legal request AI Search • Triage legal request capability
Employee	Now Assist in Contract Management	• Contract analysis • Contract metadata extraction • Contract obligation extraction • Contracts query classifier • Conversational contract search and insights

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

Available skills by workflow (continued)

Workflow	Product	Available skills
Employee	Now Assist for Workplace Service Delivery (WSD)	• Reserve Space Virtual Agent topic • Workplace Case Summarization
Creator	Now Assist for Creator	• App generation • Build Agent • Catalog item generation • App summary generation • Client script summarization • Code Assist autocomplete • Code Assist edit • Code assist summarization • Code Assist generation • Event handler generation • Experience generation • Flow generation • Flow generation with images • Flow recommendations • Flow summarization • Mobile card generation • Playbook generation • Playbook generation with images • Playbook recommendations • Process inefficiency highlights • Robotic Process Automation (RPA) bot generation • Spoke generation • Test generation • Work notes analysis
Platform	Now Assist Platform	• Article optimization • Complete record generation • Conversational Help • Document summarization • Extract information from documents • GAF skills

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

Available skills by workflow (continued)

Workflow	Product	Available skills
		• Knowledge content recommendation • Multimodal chat • Navigation • New column data generation • Potential knowledge gaps • Requester approval checklist • ServiceNow Lens • TextToResult
Data and Analytics	Now Assist skills for Analytics	AI Data Explorer skills • Analytics exploration • Exploration summarization • Refine text in exploration Query Generation skills • Analytics query generation • Analytics insight generation • Analytics follow up generation • Analytics hidden insight generation Skills installed by default with Platform: • Dashboard and visualization export • Data visualization generation
Finance & Supply Chain	Now Assist for Accounts Payable Operations (APO)	• Invoice case summarization • Purchase order summarization
Finance & Supply Chain	Now Assist for Supplier Lifecycle Operations (SLO)	• Email response • Sentiment analysis • Supplier case summarization • Supplier summarization • Supplier performance summarization
Finance & Supply Chain	Now Assist for Sourcing and Procurement Operations (SPO)	• Negotiation summarization • Procurement case summarization • Purchase requisition summarization

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

Available skills by workflow (continued)

Workflow	Product	Available skills
		• Sourcing event summarization • Sourcing request summarization
App Engine	Now Assist for App Engine	Custom app record summarization
Impact	Impact The Impact workflow contains technical accelerators that can help you get started more quickly with some Now Assist features.	• Jumpstart Your Now Assist for Creator • Jumpstart Your Now Assist in Document Intelligence • Jumpstart Your Now Assist in Virtual Agent • Jumpstart Your Now Assist Skill Kit
Vault	Now Assist for Vault	• Check role access • Generate custom data pattern • Schedule Data Discovery job
Other	Now Assist for Zero Copy Connector	• ERP data discovery • ERP data query

Now Assist skills in the Platform workflow

Most Now Assist generative AI products include skills in the Platform workflow, such as product navigation. Some Now Assist products include skills for the conversational user and platform experience, as well as knowledge article recommendations.

Auto activation of generative AI platform skills

Starting with the Zurich Patch 4 release, some generative AI platform skills for Now Assist will be active by default. This is applicable for commercial and regulated market customers.

- New users: When you install the Now Assist plugins, the related Now Assist skills will be turned on automatically.
- Existing users: When you upgrade the Now Assist plugins, any unused skill will turn on.
  1. If you have installed the plugins but never configured, activated or changed roles for a skill, any skill with 'Default On' status will be enabled individually, after the upgrade.
  2. If you have modified any role for a skill, turned it off and then back on, that skill will stay turned off after the upgrade; maintaining your last known selection.

Some skills, such as the Group Action Framework (GAF) skills, support Platform functionality and may require specific configuration.

Article optimization

The article optimization skill in ServiceNow AI Platform provides recommendations for improving the quality and health of knowledge articles, providing actionable feedback to authors and

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

managers. The recommendations for knowledge articles become available after you activate the Article Optimization skills in Now Assist Admin.

Article optimization is a Platform skill that is enabled by the admin from the Knowledge feature card. The skill allows agents and authors to review and use article optimization recommendations to improve the quality and health of their knowledge articles.

For more information, see Review and optimize articles using Article Optimization.

Catalog item form slot-fill

With the catalog item form slot-fill skill, Now Assist can automatically complete catalog item forms based on what users search for.

For instance, if requester looks up a standard laptop with 256 GB storage in silver color, the resulting form will have those specifications already filled in provided they're part of the form. Requester can quickly review these pre-filled forms, making requests faster and more efficient.

To enable this feature, activate the catalog item form slot-fill skill so that requesters see their search criteria reflected in the catalog item request forms. As soon as requesters searches for an item, relevant fields in the form are populated with their search information, minimizing repetitive entry and streamlining the request process.

Pre-filled fields show AI icons to help users validate values populated by Now Assist.

Note: Order guides, as well as the following variables, are not supported:

- MRVS (Multi-Row Variable Sets)
- Custom and custom with labels
- Attachments
- Masked
- Break
- Container
- UI page
- HTML

Now Assist Conversational Help

This skill uses Generative AI application capabilities to provide answers to the questions on the Now Assist panel.

https://player.vimeo.com/video/1088792812?h=bceb334712&badge=0&autopause=0&player_id=0&app_id=58479.

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

You can install the External Content Connectors Application Suite from the ServiceNow store

and use NAp Synthesized Search with ServiceNow Docs.

Important: This Now Assist skill is turned on by default. The skill will be automatically available to appropriate role users for the application. For more information, see Now Assist skills, agents, and agentic workflows on by default.

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

The Now Assist Conversational Help skill displays as **Get Help** on the Now Assist panel.

Note: The Get Help feature is available as a part of Now Assist entitlements and no new subscription is required. The feature is enabled by default and you can turn it off in the Now Assist Admin **Settings**. The ServiceNow Now LLM Service (Large Language Model) is the default model provider for this Now Assist skill to retrieve precise answers to the users' questions.

How Conversational Help works

You can ask your question by selecting **Get Help** skill and submit a query to use the help option in the Now Assist panel. Your query goes through the following steps to retrieve the best response.

Note: Conversational Help Skills is not available as auto-discoverable as a help skill within Now Assist panel.

1. AI search

   The query is be sent to a **central ServiceNow instance**, where **AI Search** will search the knowledge table that stores content from product doc content.

   We use Instance Data Replication (IDR) that synchronizes product documentation content from Now Support instances.

   Note: Multiple central ServiceNow instances are deployed across three regions: US East or US West, EMEA, and APJC.

2. The top matching chunks (or records) along with the user's query, are sent to the Now Assist Q&A Genius Result configuration to generate a meaningful, contextual answer.
3. Re-ranking

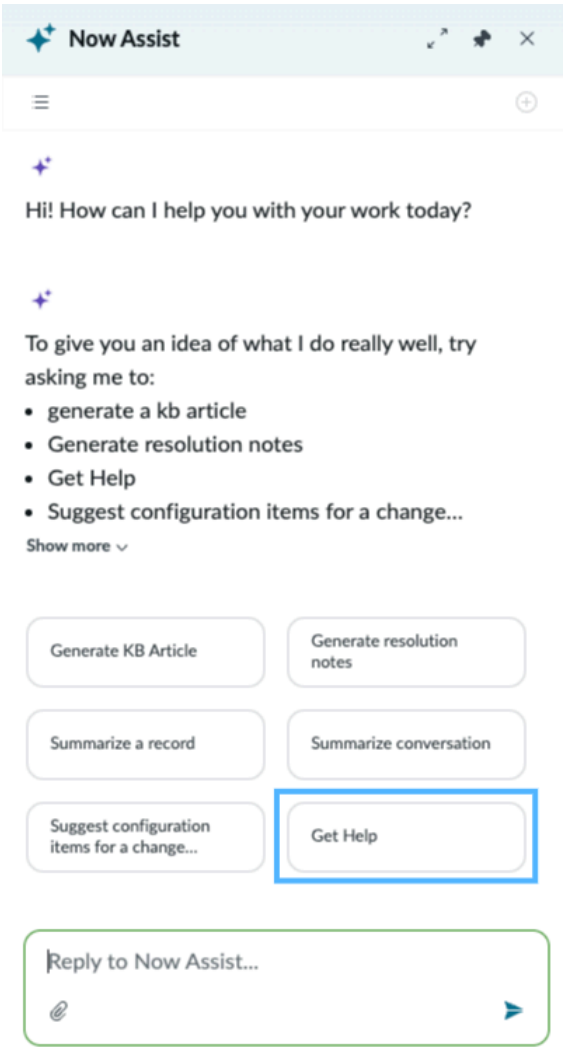
   The Now Assist Q&A, an out of box (OOB) capability, re-ranks the retrieved chunks to ensure better relevance.

4. Cache lookup
   - If the user query exists in the cache, the pre-stored answer is returned.
   - If the query is not cached, the query and re-ranked chunks are sent to Now LLM to generate a meaningful answer.

The Now LLM retrieves the most relevant result from https://www.servicenow.com/docs/ portal and displays it in the same panel.

Note: Effective from this release, the query will retrieve results based exclusively on the release version of the user's current instance. This enhancement is integrated into the query process to ensure the delivery of precise results that reflect the latest updates and features.

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



For more information, see Fetch end points in Now Assist Conversational Help skills.

Now Assist extract information from documents

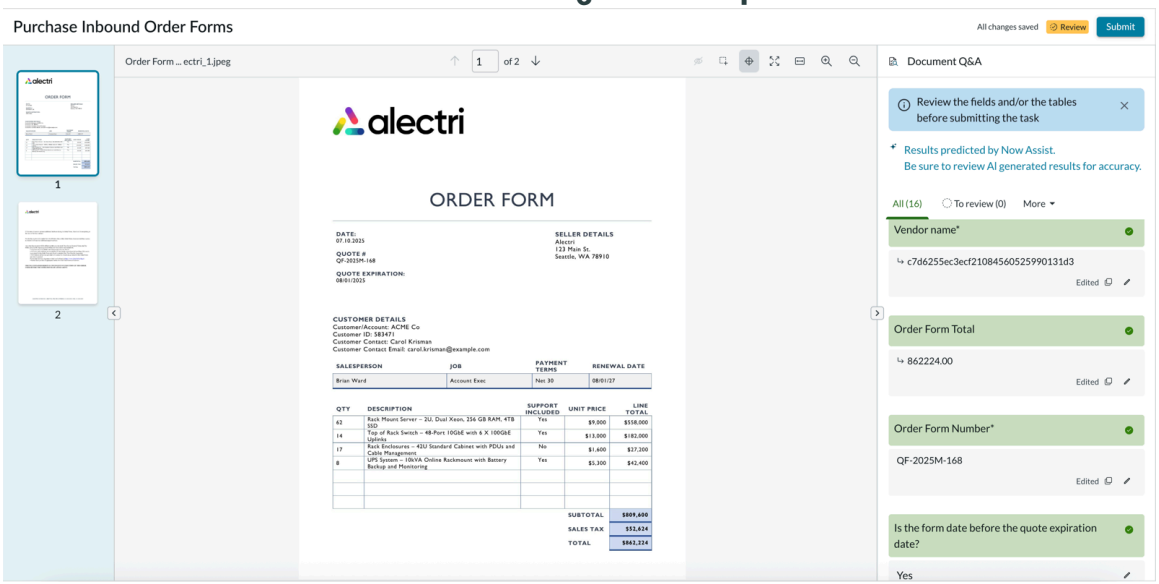
The extract information from documents skill allows you to use Now Assist predictions to extract information from document and image files.

Now Assist uses generative AI capabilities to extract values from the document based on the fields, tables, and questions defined in the skill's use case. The extracted information is turned into structured data that can be used in the platform.

Agents can use the Document Intelligence workspace to review and confirm the extracted information.

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

Extracted information in the Document Intelligence workspace



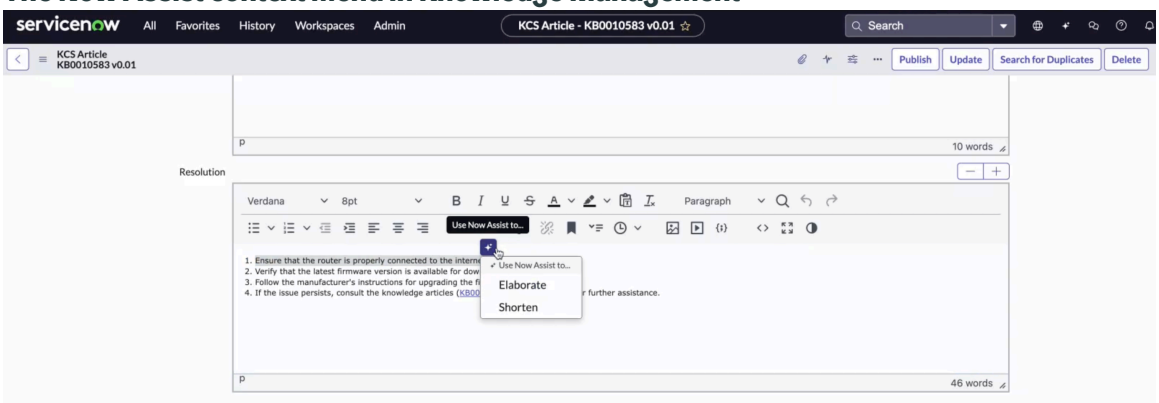
For more information, see Review extracted information in the Document Intelligence workspace.

Knowledge content recommendation

Knowledge generative AI skills on the ServiceNow AI Platform provides recommendations for editing a knowledge article. Once activated, this skill is available on the Now Assist context menu.

The Now Assist Knowledge content recommendation is a Platform skill that is enabled by the admin from the Knowledge feature card. The skill allows agents and authors to use the Now Assist context menu to elaborate and shorten content in a knowledge article.

The Now Assist context menu in Knowledge Management



For more information, see Edit an article using the Now Assist context menu.

Navigation

Use the navigation skill in Now Assist to take you where you want to go on the ServiceNow AI Platform.

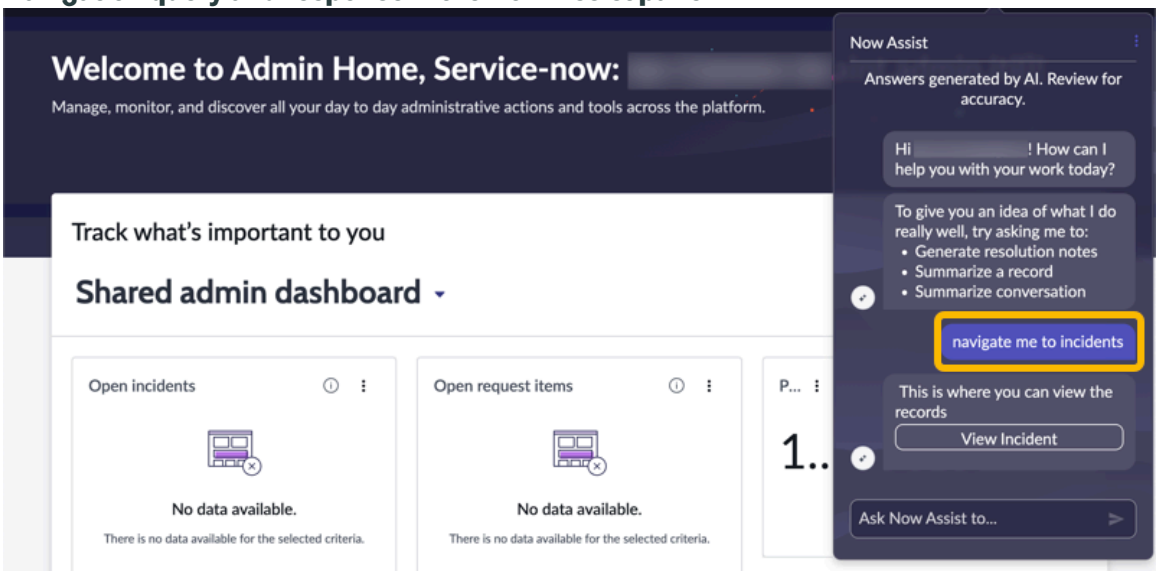
Overview of navigation

Navigation is a skill in the Now Assist panel that handles record search requests during a chat. When you ask for records or tables in plain language Now Assist shows you links in the chat to take you to the best match to your request. For example, you could enter `Show me incident records`.

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

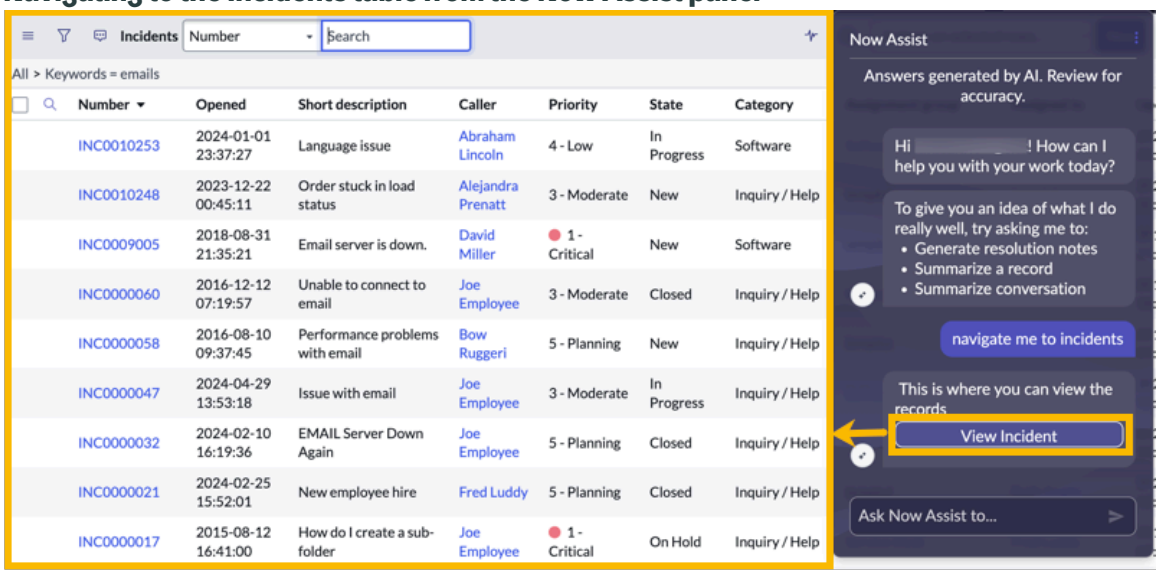
In the following figure, the user entered `navigate me to incidents` in the Now Assist panel. Now Assist responds with a link to the Incidents table.

Navigation query and response in the Now Assist panel



When you click the link, the list of all records in the Incidents table displays.

Navigating to the Incidents table from the Now Assist panel



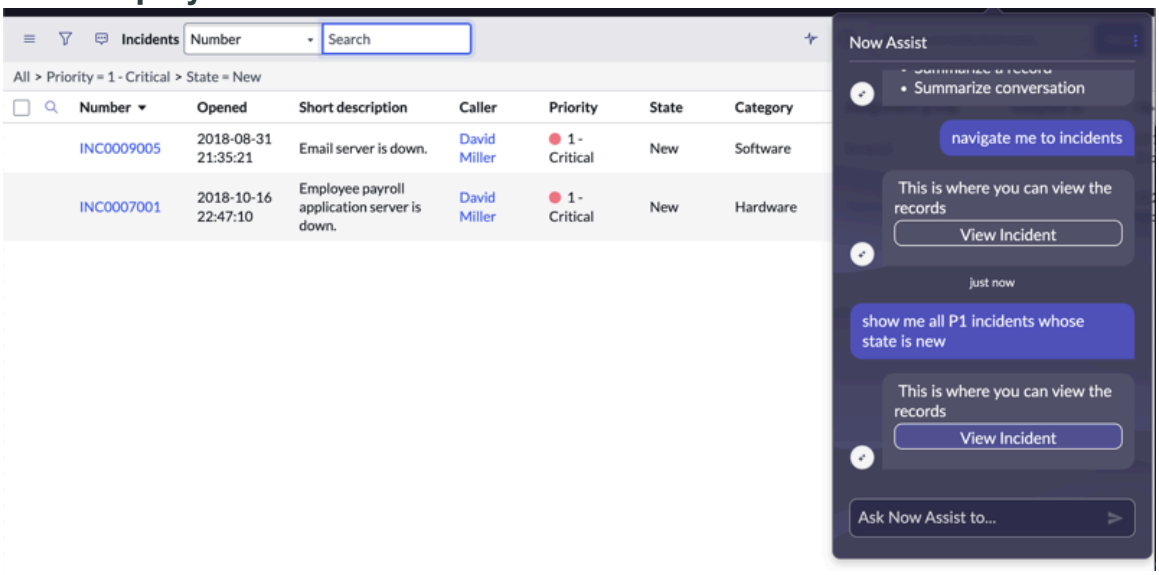
Refining your results

You can refine your results further by using more detailed requests. If you enter `Show me all incident records whose status is Complete`, Now Assist shows you only the records in the table with a Complete status.

In the following example, the user asks for all P1 incidents that are in the New state.

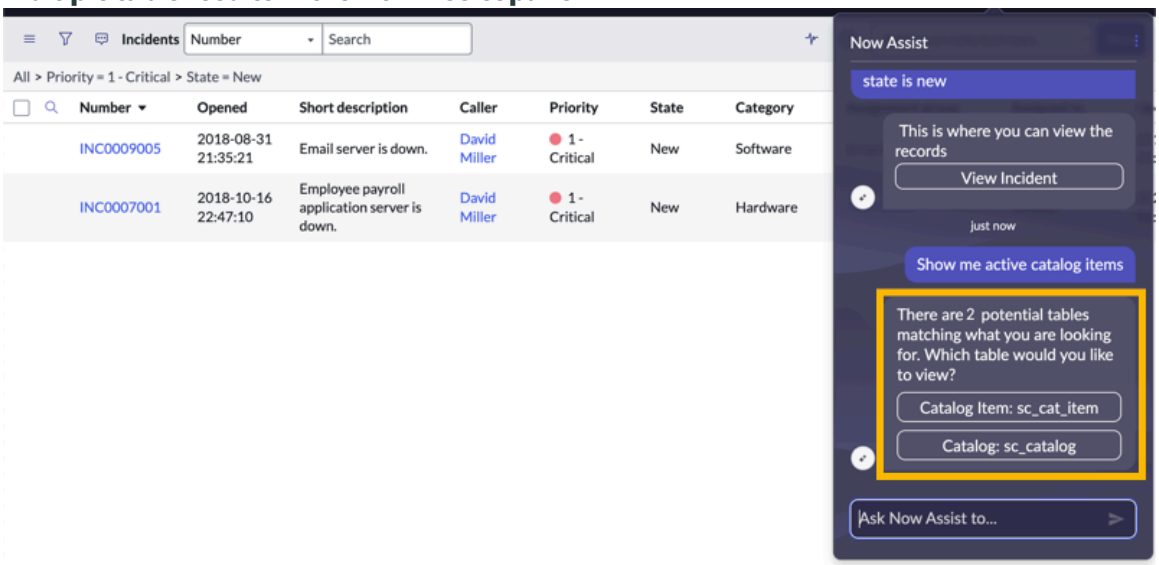
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

A refined query of the Incidents table



The number of results is based on how many potential results Now Assist finds in response to your request. If Now Assist finds more than 10 results, the list is paginated. In the following example, Now Assist finds two possible tables: the Catalog Item table and the Catalog table.

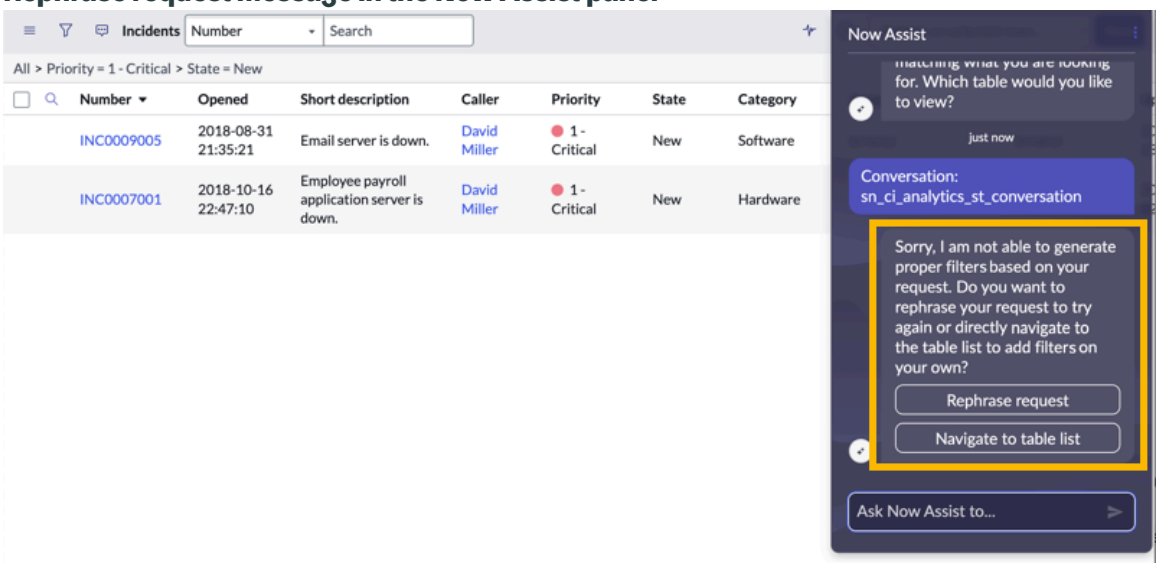
Multiple table results in the Now Assist panel



If Now Assist does not understand your request, you receive an error message asking you to rephrase your request. You can also choose to navigate to a "best guess" based on your previous request.

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

Rephrase request message in the Now Assist panel



Now Assist in Standard Ticket Page

The Now Assist in Standard Ticket Page skill allows you to display summaries generated by Now Assist, providing overviews of recent activities and ticket details.

When this skill is active, the standard ticket page on the Service Portal displays a **Summarize** button. If requesters wish to see the ticket activity and a summary of information generated by Now Assist, they can select the **Summarize** button.

Potential knowledge gaps

The potential knowledge gap skill helps find missing or incomplete knowledge coverage by analyzing service interactions. Gap recommendations become available after you activate the knowledge gaps skills in Now Assist Admin.

The knowledge gaps feature (also known as potential gaps) helps knowledge managers and service teams discover missing or insufficient knowledge articles by analyzing service case and incident patterns. When the relevant skills are activated, Now Assist generates gap recommendations to guide the creation or improvement of knowledge content.

For more information, see Manage potential knowledge gaps

Requester approval checklist skill

The requester approval checklist skill in the ServiceNow AI Platform® generates a structured checklist by mapping real-time request data against your organization's knowledge articles.

Important: This Now Assist skill is turned on by default. The skill will be automatically available to appropriate role users for the application. For more information, see Now Assist skills, agents, and agentic workflows on by default.

This skill is available within Now Assist Agents for Requester v3.1 and is designed to provide instant, data-backed decision support across in-product, agentic, and workflow experiences.

The skill uses information from the following sources to promote accuracy:

Knowledge articles

Defines the evaluation criteria based on your company policy.

System data

Fetches requester profiles and historical data from ServiceNow.

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

Approval request

Analyze the details submitted in a live request.

Requester Approval Checklist structure

The output is presented in a two-step format: Criteria (what is required) and Reason (why it passed, failed, or has missing information). Each item has one of the following statuses: Met, Not met, or Unknown.

Note: All referenced knowledge base articles are provided as links at the end of the checklist for easy verification.

Key skill considerations

Integration

The skill is built for workflows and agentic experiences, but doesn't support follow-up questions and answers.

Decision support

The skill provides information only. It doesn't perform actions like Approve, Reject, or Comment.

Flexibility and control

The Requester Approval Checklist respects the approval configuration to provide flexibility and control.

For more information, see Platform Approval assistance AI agent and Configure Service Portal Approval Configuration record.

Now Assist agentic workflows

You can customize Now Assist agentic workflows to meet the needs of your users in different workflows. For example, you can use agentic AI to help resolve tasks, execute routine but variable procedures, and investigate root causes or analytical trends.

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

Agentic workflows overview

Agentic workflows use multiple AI agents to achieve specific outcomes. Different workflows offer different available tasks. Many agentic workflow templates are available for you to activate, duplicate, or customize. For more information, see Activate an agentic workflow template and Duplicate an agentic workflow.

For more information about the in-product agentic AI experience, see In-product agentic AI. You can create UI actions for your agentic workflows in AI Agent Studio. Open the agentic workflow, navigate to the Select channels and access step in the guided setup, and create a UI action.

Workflow and product agentic workflows

The following table describes the available AI agents agentic workflows that are included in Now Assist applications.

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

Available agentic workflows by product

Product	Available agentic workflows
Now Assist for Configuration Management Database (CMDB)	• Create configuration item • Provide advice on CMDB governance • Search CMDB
Now Assist in Contract Management	• Conversational contract search and insights • Manage contract repository
	CBS setup
Now Assist for Creator	
Now Assist for Customer Service Management (CSM)	• Accelerate complaint case handling • Provide Customer 360 insights • Triage cases
Now Assist for Enterprise Architecture (EA)	Generate enterprise architecture diagram
Enterprise Asset Management	• Help repair enterprise assets • Help manage enterprise asset requests
Now Assist for Operational Sustainability Management	Generate carbon calculations for metrics
Now Assist for Field Service Management (FSM)	Create a work order
Now Assist for Finance and Supply Chain (FSC)	Conversational intake for sourcing and procurement
Now Assist for Financial Services Operations (FSO)	Help resolve friendly fraud disputes
Now Assist for Hardware Asset Management (HAM)	• Help manage hardware asset requests • Help repair hardware assets
Now Assist for Health and Safety	• Create actions for safety incidents • Help resolve health and safety cases agentic workflow • Health safety incident patterns assistant agentic workflow

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

Available agentic workflows by product (continued)

Product	Available agentic workflows
Now Assist for HR Service Delivery (HRSD)	• Help resolve tuition requests • Help resolve tuition reimbursement requests • Predict service and transfer HR cases • Resolve noncritical HR cases • Resolve critical HR cases • Generate onboarding ramp-up plan • Recommend catalog and order guide tasks in Onboarding Journey • Schedule interviews
Now Assist for Integrated Risk Management (IRM)	• Generate regulatory action plans • Get regulatory analysis • Identify risks for an entity • Optimize issue resolution
Now Assist for IT Operations Management (ITOM)	• Analyze alert impact • Agent Client Collector (ACC) diagnostic workflow • Analyze potential impact • Manage alerts autonomously • Triage and analyze alerts
Now Assist for IT Service Management (ITSM)	• Assess conflicts for a change request • Assess quality of a change request • Create outages for a change request • Create standard change request • Create standard change template proposal • DEX issue diagnosis and resolution • Explain SLA • Generate change request plans • Generate post incident reviews • Incident assist • Investigate and resolve ITSM incidents • Manage Microsoft 365 group members

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

Available agentic workflows by product (continued)

Product	Available agentic workflows
	• Notify users with Twilio • Schedule a change • Suggest configuration items for a change request • Triage and categorize ITSM incidents • Wrap-up and resolve ITSM incidents
Now Assist for Legal Service Delivery (LSD)	• Conversational contract search and insights • Triage legal requests
Now Assist for Manufacturing Commercial Operations (MCO)	Execute recall campaigns faster
Operational Technology (OT) Manager Foundation	Import the OT device spreadsheet into OT CMDB
Now Assist for Operational Technology Service Management (OTSM)	Generate OT KB articles agentic workflow
Now Assist in Notifications	Email Intent to Action
Now Assist for Public Sector Digital Services (PSDS)	Help manage public information requests
Now Assist for Sales Force Automation (SFA)	Help nurture new leads
Now Assist for Security Incident Response	• Analyze security operations metrics • Close security incident • Generate SIR shift handover report • Resolve security incident
Now Assist for Software Asset Management (SAM)	• Create software reclamation rule • Evaluate software removal candidate • Help manage software asset requests
Now Assist for Strategic Portfolio Management (SPM)	• Generate Agile stories • Monitor project tasks
Now Assist for Supplier Lifecycle Operations (SLO)	Email intent to action

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

Available agentic workflows by product (continued)

Product	Available agentic workflows
Now Assist for Telecommunications, Media and Technology (TMT)	• Help remediate bill issues • Analyze network incident • Customer voice quality issue resolutions • Monitor engagement health • Recommend risk signal solutions agentic workflow • Trigger risk mitigation touchpoint • Support renewals and expansions • Service Exchange onboarding
Now Assist for Vault	• Access observer configuration • Field encryption with Vault module • Securing custom apps with Vault agents • Summarize access observer logs
Now Assist for Workplace Service Delivery (WSD)	• Automate map updates • Help manage workplace reservations • Manage temporary space closures • Optimize cleaning activities • Process images for new tasks • Workplace Concierge
Now Assist for Zero Copy Connector	Explore ERP models
Now Assist Platform	• Analyze incident trends • Classify tasks • Generate my work plan • Generate resolution plans • Help optimize team productivity • Identify ways to improve services • Investigate IT problems • Notification agent • Process images for tasks • Suggest survey responses

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

In-product experience for agentic workflows

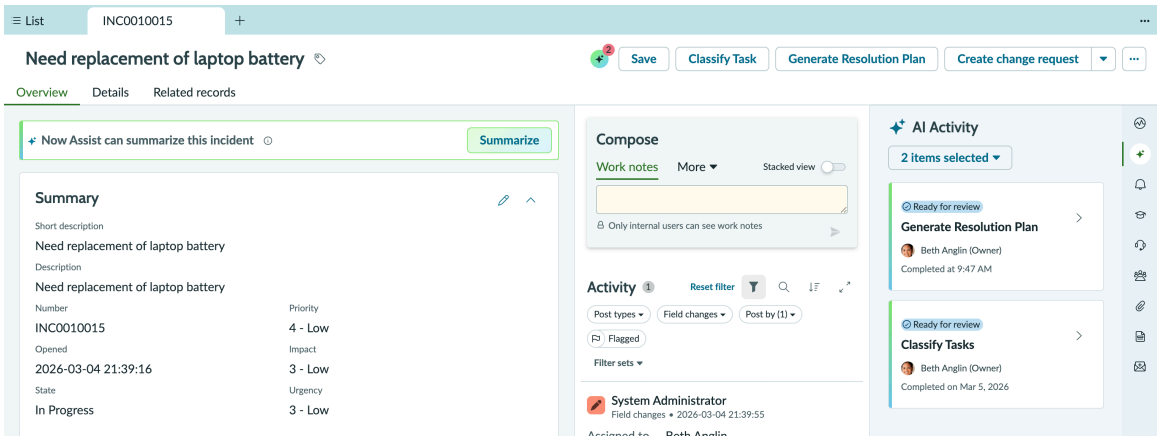
Dedicated spaces in workspaces and in the Core UI enable you to use agentic workflows directly in record forms. You can view agentic workflow progress and previous executions, and you can answer follow-up questions for ones that require human supervision.

Agentic AI in the Core UI and workspaces

Agentic workflows can help accomplish complex tasks for you, such as generating resolution notes for cases and incidents or investigating problems and root causes. You can view agentic workflows running on a record in the AI Workflows panel in the Core UI form or in workspaces. For agentic workflows that require human supervision, you can answer questions, approve next steps, or provide other input. Along with current progress, you can also review historical runs to compare results.

You can create UI actions for your agentic workflows in AI Agent Studio. Open the agentic workflow, navigate to the Select channels and access step in the guided setup, and create a UI action.

If you don't see this panel, ensure that the property *`com.glide.agentic_processes_view.enabled`* is set to `true`.



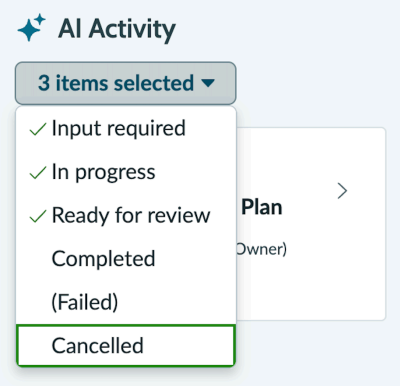
Agentic workflow execution list

The list of agentic workflow cards can be filtered by execution state. You can change which states to filter by at any time.

Note: The difference between the *`Ready for review`* state and the *`Completed`* state is that the former has generated some output. Agentic workflows without an output are marked only as *`Completed`*.

The user responsible for answering follow-up questions is identified by the *`Supervised by`* field. Only the user specified can answer questions to progress the agentic workflow, but others can see the questions asked.

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



Each card displays the current or final processing message. You can see the full list of processing messages by selecting the agentic workflow card to open the details.

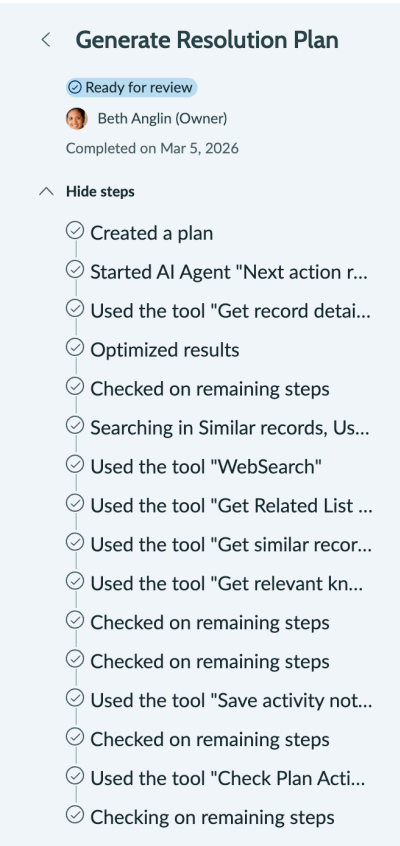
Time estimates based on previous executions are provided for currently running workflows. For completed workflows, the cards show the total time taken.

The execution list also displays the results of any workflows triggered in the Now Assist panel or triggered automatically by events once they are complete.

Agentic workflow execution details

When you select a specific agentic workflow, you can view the processing messages and history. Processing messages show you what steps the agentic workflow has taken already and which ones are still being completed.

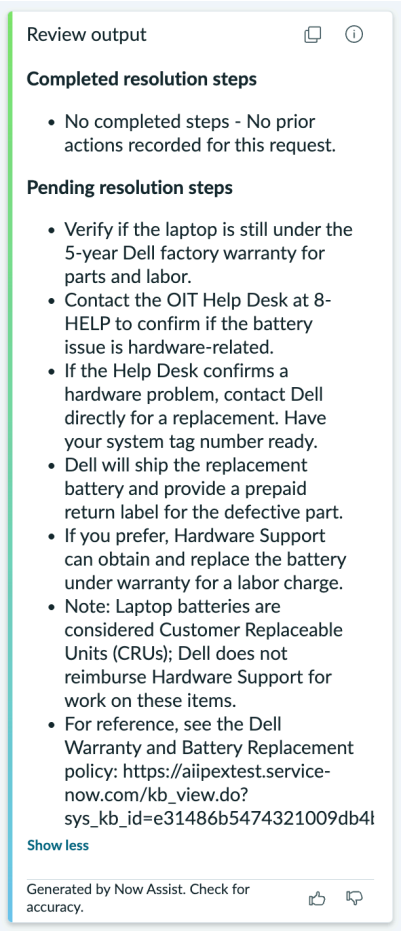
You can change the processing messages for an AI agent or tool in AI Agent Studio. For an AI agent, open the AI agent and go to the Select channels and access step. For a tool, open the AI agent, go to the Add tools and information step, and select the tool to open the form modal.



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

The history space on the card can show previous answers to follow-up questions. For completed agentic workflows, you can review the final output in the history if there is any. If there's an output, you can see the sources for the AI's reasoning by selecting the information icon. Citations can include Knowledge Base articles and lists of records.

You can copy the text of the final output of an agentic workflow by selecting the copy icon.



Fields updated with agentic AI

If an agentic workflow changes the value of a field, Now Assist displays a label under the field value stating that the value was modified by AI.

AI presence indicator

When an agentic workflow is in progress on the record, an icon is displayed at the top of the record form along with other UI actions. If selected, you can see how many agentic workflows are running on the record and the general status of each one.

If there's an agentic workflow in progress that you don't want to complete, you can cancel it. Select the more options icon on the agentic workflow card, and then select **Cancel**.

Alerts for agentic workflow status changes and required input

When an agentic workflow changes states, such as when it completes or has a follow-up question, an alert appears at the top of the screen. If the workflow generates an output, you can select the **Review** button to see the final result.

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

Enable the in-product experience for agentic workflows

Enable the AI Workflows panel and UI actions for agentic workflows on forms in the Core UI and workspaces to track agentic AI executions.

Before you begin
Role required: admin

About this task
Agentic workflows can perform work on certain records, and you can track their progress or provide input in the AI Workflows panel available for forms and workspaces. See In-product agentic AI for more details about the functions and features of the AI Workflows panel.

The following task describes the process for enabling the system property that allows you to see the AI Workflows panel and see UI actions for agentic workflows: *`com.glide.agentic_processes_view.enabled`*.

To enable users to access agentic workflows with UI actions, you can open the agentic workflow in AI Agent Studio and navigate to the **Select channels and access** step. You can select a UI action as a possible way to access the workflow. See Select channels and access for agentic workflows for more information.

Procedure
1. Navigate to **All** > **System Properties** > **All Properties**.
2. Select **New**.
3. In the *Name* field, enter `com.glide.agentic_processes_view.enabled`.
4. In the *Value* field, enter `true`.

Result
The In-product agentic experience, including the AI Workflows panel, is available.

What to do next
Monitor agentic workflow execution on forms in the Core UI and workspaces, or create UI actions to grant users access.

Platform agentic workflows

You can use the available Now Assist AI agents Platform agentic workflows to achieve business outcomes with self-executing autonomous AI agents.

Use the following agentic workflows that are available with ServiceNow AI Platform.

Available agentic workflows for Platform

Agentic workflow name	Description	Available AI agents
Analyze task trends	Detects recurring patterns, predicts disruptions, and enables proactive resolutions with actionable recommendations.	Issue trend analysis AI Agent

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

Available agentic workflows for Platform (continued)

Agentic workflow name	Description	Available AI agents
Classify tasks	Triages tasks by updating fields, evaluating sentiment, and summarizing.	Record field value prediction AI agent
Generate my work plan	Provides a personalized work plan based on current work and historical data of previous work.	Prioritize work AI agent
Generate resolution plan	Analyzes tasks, generates resolution summaries, and updates comments or work notes.	• Next best action recommendation AI agent • Resolution Action AI Agent
Help optimize team productivity	Identifies backlog items, evaluates team member performance on previous assignments, and allocates work based on similarity to those previous assignments.	Work Allocator AI Agent
Identify ways to improve service	Analyzes feedback, trends, and metrics and provides recommendations to help optimize processes.	Survey Analysis AI agent
Investigate problems	Provides insights from incident or problem details.	Problem investigation AI agent
Process images for tasks	Processes images and converts them to tasks.	• Image Processor Agent • Document and visual insights AI agent
Propose survey responses	Suggests answers for survey questions.	• Survey response suggestion AI agent • Survey filling data collection AI agent

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

Important: By default, all agentic workflows and AI agent records are read only.

To run the AI agents autonomously, you must either activate the agentic workflow template or duplicate the agentic workflow, and then proceed with the following steps:

- Activate the agentic workflow.
- Activate all agents within the agentic workflow.
- Activate the trigger to invoke the agentic workflow automatically. If you prefer to invoke it manually, activating the trigger isn't necessary.

Tools mapped to AI agents in agentic workflows

To find the tools mapped to AI agents used in the agentic workflows, you can perform the following steps:

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**
2. In the Agentic workflows tab, select the agentic workflow.
3. In the **Describe and connect** step of the guided setup, select the AI agent you want to see the tools for.
4. Go to the **Add tools and information** step.
5. Review the tools mapped to the AI agent.
6. Repeat for all AI agents in the agentic workflow.

Agentic workflows and UI actions

To enable users to access agentic workflows with UI actions, you can open the agentic workflow in AI Agent Studio and navigate to the **Select channels and access** step. You can select a UI action as a possible way to access the workflow

If you don't see your UI actions after configuring it in AI Agent Studio, ensure that the property *com.glide.agentic_processes_view.enabled* is set to `true`.

There might be AI agents installed with the Now Assist application that are not used in agentic workflows. To learn how to see all agents that are available to you, see Find AI agents.

Platform Analyze task trends agentic workflow

Use the Platform Analyze task trends AI agents agentic workflow to detect recurring task patterns so that you can flag and resolve them before they escalate.

Analyze task trends overview

The Analyze task trends agentic workflow enhances task management by detecting recurring patterns, predicting disruptions, and enabling proactive resolution to reduce downtime and improve reliability. Tasks are grouped and indexed by the large language model (LLM) to analyze common recurring issues and root patterns. The LLM then generates resolution recommendations based on the analysis and displays it to you. After the analysis is generated, you can continue the conversation to do the following:

- Get a summary of each group analysis. You have to specify which group you would like to get a summary of.
- Download the analysis as a PDF or Word document.

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

- Get more information by asking follow-up questions. For example, you can ask why certain suggestions were generated.
- Analyze the next ten groups. Each analysis is done for ten groups at a time. You can continue analyzing more groups within the same conversation, but the other actions for the previous group are no longer available.

The exact options for follow-up actions available can be configured.

The agents, tools, and triggers associated with the Analyze task trends agentic workflow are provided by Now Assist applications. You can activate the agentic workflow template by making triggers active and setting the display settings to include the Now Assist panel. If you want to change this agentic workflow's instructions, you must duplicate it, adjust the settings to suit your specific needs, and activate the duplicated version instead.

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

Prerequisites and setup

To access this workflow, you must have Now Assist for Platform installed on your instance. You can get this by installing any other Now Assist application, such as Now Assist for IT Service Management (ITSM).

For this agentic workflow to behave as expected, you should have at least 500 records on your task table.

You must also configure Group Action Framework (GAF). See Group Action Framework for more information on what GAF is and how to set it up. The Incident table is the default table for this workflow, but you can configure GAF for other task tables such as Case or HR Case.

Role masking

Required role: sn_uxc_gen_ai.platform_ai_analyze_trnds.

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with Now Assist applications have specific roles that come included with the application. If you select *`Users with specific roles`* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

In the data access settings, you must also add the necessary roles to enable reading of the tables for the records you want to access for trend analysis. For example, you can add the itil role to the agentic workflow's list of approved roles so that it can access Incident records.

Additional configuration

You can change different settings related to the agentic workflow by changing values for the Now Assist Skill Config Var Set. To access the variable set and make changes, do the following while in the Platform AI Agents and Skills scope:

- Go to the Now Assist Skill Config [sn_nowassist_skill_config] table.
- Open the record named *`Analyze Task Trends`*.
- In the Now Assist Skill Config Var Set related list, select *`Task Trends Input Config`*.

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

- Edit the variable values.
- Save or update the record.

Analyze task trends configuration

Config field	Description	Default value
Post Analysis Actions	List of possible follow-up actions a user can take before the agentic workflow completes.	• Get a summary • Get more info • Analyze next 10 groups • Download analysis
Analysis Time Frame	Range of time, in months, for the trends analyzer to look at records to identify trends. You can specify smaller ranges when running the agentic workflow, but this value defines the maximum limit.	3
GAF Record Limit	Number of records analyzed per GAF record grouping. See the previous *Prerequisites and setup* section.	8

Accessing the Analyze task trends agentic workflow

To access the agentic workflow:

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Select **Analyze task trends**.

The first step of the guided setup includes a complete list of included AI agents. Selecting the name of an AI agent opens it in a new browser tab, where you can see the full description, role, list of steps, and tools. Tools are displayed in the second step of the AI agent guided setup, Add tools and information.

In-product agentic AI and UI actions

Agentic workflows can be accessed in the Core UI and in workspaces in the AI Activity panel. From there, you can track their progress, provide or review input, and see the results of the work performed. For more information, see In-product agentic AI for more details about the AI Activity panel.

To enable users to access agentic workflows with UI actions, you can open the agentic workflow in AI Agent Studio and navigate to the **Select channels and access** step. You can select a UI action as a possible way to access the workflow

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

If you don't see your UI actions after configuring it in AI Agent Studio, ensure that the property *com.glide.agentic_processes_view.enabled* is set to `true`. See Enable the in-product experience for agentic workflows.

Testing the Analyze task trends agentic workflow

You can manually test an agentic workflow execution or access on the Testing page of AI Agent Studio if you have the sn.aia_admin role and all other roles configured in the security controls. Start a manual test, select a test type and the name of the workflow, and use utterances in the Task field like the following samples. See Test an agentic workflow execution.

If you want to evaluate the agentic workflow over many different execution logs, run an automated evaluation.

Sample utterance

After the workflow has been activated in AI Agent Studio, use similar queries to the following to run the agentic workflow in the *Now Assist* panel. You can also run this workflow on the Testing page of AI Agent Studio with the same utterance in the Task field if you have the sn.aia_admin role.

Each utterance must include the name of the table and the time frame to analyze. If the utterance doesn't include these things, the agentic workflow prompts you for more information.

The time frame specified by the user can't exceed the maximum value set by the Analysis time frame configuration.

- Analyze incident trends related to payment issues within the last two months
- Analyze case trends within the last month
- Analyze HR case trends with High Priority within the last two years
- Analyze incident and problem trends within the last two weeks

Troubleshooting

When running this agentic workflow, it's possible to see an error that states "I couldn't analyze as I didn't have the required resources." This error occurs when GAF isn't configured for the table you want to analyze. See Configure Group Action Framework for steps to configure GAF for the table. If you're still having issues after GAF is configured, reach out to Now Support.

AI agents used in the Analyze task trends agentic workflow

The following table lists the agents that are used in the Analyze task trends agentic workflow.

Important: In the Select channels and status step of each AI agent's guided setup, make sure that the Status toggle is enabled to activate the AI agent.

AI agents names and descriptions in the Analyze task trends agentic workflow

AI agent name	AI agent description	Role required
Issue trend analysis AI agent	Analyzes grouped task data to identify recurring issues and root causes. Provides detailed, actionable recommendations using structured analysis.	sn_uxc_gen_ai.platform_ai_analyze_trnds

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

Other Platform agentic workflows

For more information on other agentic workflows that are associated with the Platform workflow, see Platform agentic workflows.

Platform Classify tasks agentic workflow

Use the Platform Classify tasks AI agents agentic workflow to gather relevant information about tasks automatically and make decisions about priorities and assignments.

Classify tasks overview

The Classify tasks agentic workflow can help improve efficiency and accuracy by automatically gathering information, prioritizing tasks, assigning teams, detecting sentiment, and generating task summaries.

The agents, tools, and triggers that are associated with the Classify tasks agentic workflow are provided by Now Assist applications. You can activate the agentic workflow template by making triggers active and setting the display settings to include the Now Assist panel. If you want to change this agentic workflow's instructions, you must duplicate it, adjust the settings to suit your specific needs, and activate the duplicated version of the agentic workflow instead.

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

Prerequisites and setup

To access this workflow, you must have Now Assist for Platform installed on your instance, which you can get if you install any other Now Assist application, such as Now Assist for IT Service Management (ITSM).

For this agentic workflow to behave as expected, you must also configure Group Action Framework (GAF). See Set up AI Search for Group Action Framework and Configure Group Action Framework for more information on getting started with GAF.

Role masking

Required role: sn_uxc_gen_ai.platform_ai_classify_tasks

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with Now Assist applications have specific roles that come included with the application. If you select *`Users with specific roles`* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

In the data access settings, you must also add the necessary roles to enable reading of the tables for the records you want to classify. For example, you can add the itil role to the agentic workflow's list of approved roles so that it can access Incident records.

Additional configuration

You can change different settings related to the agentic workflow by changing values for the Now Assist Skill Config Var Set. To access the variable set and make changes, do the following while in the Platform AI Agents and Skills scope:

- Go to the Now Assist Skill Config [sn_nowassist_skill_config] table.
- Open the record named *`Task Classify Tasks Skill Config`*.

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

- In the Now Assist Skill Config Var Set related list, select *Field Predictor*.
- Edit the variable values.
- Save or update the record.

The Classify tasks configuration variable set includes the following variables. You can configure either the AIS fields or the GAF field for determining how the agentic workflow gathers what work the user has. If you configure both, GAF takes priority when running the agentic workflow. For more information about GAF, see Group Action Framework.

Classify tasks configuration

Config field	Description
Search Mode	Options include Keyword, Semantic, or Hybrid. Keyword-based searches look at individual phrases, while semantic-based searches rely on the phrasing of utterances as well. Hybrid utilizes both strategies. Default: Keyword
AIS Search Fields	Fields used by AI Search operations to find similar records for improved prediction context
Capability	OneExtend capability for enhanced field prediction functionality. See Generative AI Controller reference for more details.
GAF Config	Skill configuration for Group AI Framework (GAF).
AIS Search Fields	Fields used by AI Search to determine what work a user has
AIS Return Fields	Fields returned by AI Search to the agentic workflow to base decisions on
GAF Config	Group Action Framework grouping configuration record, which is a collection of groups of records to make searching easier
AIS Search Profile	Profile for AI search, such as Now Assist in Virtual Agent.
Capability Payload	Server-side script to generate custom payloads for OneExtend capability integration

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

Classify tasks configuration (continued)

Config field	Description
AIS Fallback Table	Table used if GAF is not configured
Return Fields	Field values returned from identified relevant records
Check for Issue Readiness	Uses the information provided by the Issue Readiness AI agent ⚠ **Warning:** Only enable if the Issue Readiness AI agent is active on your instance. Otherwise, the agentic workflow will not execute.
Fields	Field values to predict and update
Table	Table where you want to execute the agentic workflow on
Search Term Source	Options include Record-based and Utterance-based. Record-based derives search terms for generating context from the record itself. Utterance-based relies on the user-submitted utterance for search term options. Default: Record-based

Accessing the Classify tasks agentic workflow

To access the agentic workflow:

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Select **Classify tasks**.

The first step of the guided setup includes a complete list of included AI agents. Selecting the name of an AI agent opens it in a new browser tab, where you can see the full description, role, list of steps, and tools. Tools are displayed in the second step of the AI agent guided setup, Add tools and information.

In-product agentic AI and UI actions

Agentic workflows can be accessed in the Core UI and in workspaces in the AI Activity panel. From there, you can track their progress, provide or review input, and see the results of the work performed. For more information, see In-product agentic AI for more details about the AI Activity panel.

To enable users to access agentic workflows with UI actions, you can open the agentic workflow in AI Agent Studio and navigate to the **Select channels and access** step. You can select a UI action as a possible way to access the workflow

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

If you don't see your UI actions after configuring it in AI Agent Studio, ensure that the property *com.glide.agentic_processes_view.enabled* is set to `true`. See Enable the in-product experience for agentic workflows.

Testing the Classify tasks agentic workflow

You can manually test an agentic workflow execution or access on the Testing page of AI Agent Studio if you have the sn.aia_admin role and all other roles configured in the security controls. Start a manual test, select a test type and the name of the workflow, and use utterances in the Task field like the following samples. See Test an agentic workflow execution.

If you want to evaluate the agentic workflow over many different execution logs, run an automated evaluation.

Sample utterance

After the workflow has been activated in AI Agent Studio, enter these or similar phrases in the Now Assist panel to trigger the workflow. You must have the sn.now_assist_panel_user role to run the workflow. You can also run this workflow on the Testing page of AI Agent Studio with the same utterance in the Task field if you have the sn.aia_admin role.

- Populate the priority and assignment group for INC0001
- Populate the empty fields on INC0001
- Find the priority for INC001

AI agents used in the Classify tasks agentic workflow

The following table lists the agents that are used in the Classify tasks agentic workflow.

Important: In the Define availability step of each AI agent's guided setup, make sure that the Status toggle is enabled to activate the AI agent.

AI agents names and descriptions in the Classify tasks agentic workflow

AI agent name	AI agent description	Role required
Record field value prediction AI agent	Fetches, predicts, creates, and updates records with the provided details. Provides summary and justification for recommendations.	sn_uxc_gen_ai.platform_ai_field_predictor

Other Platform agentic workflows

For more information on other agentic workflows that are associated with the Platform workflow, see Platform agentic workflows.

Platform Generate my work plan agentic workflow

Use the Platform Generate my work plan agentic workflow to create personalized work plans for currently assigned work.

Generate my work plan overview

The Generate my work plan agentic workflow creates personalized work plans based on current assigned work. This reduces manual effort and guesswork in work planning. The AI agents identify and retrieve all work assigned to the user, predict effort required to complete work, and

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

generate an actionable work plan. Effort is estimated based on historical records related to the open work items. Generated work plans emphasize the highest priority work by examining the following information and can result in faster resolution and fewer missed SLAs.

- User sentiment
- Short description
- Priority
- Description
- Due date
- Escalation
- State
- Updated
- Number
- Impact
- SLA

The agents, tools, and triggers associated with the Generate my work plan agentic workflow are provided by Now Assist applications. You can activate the agentic workflow template by adding triggers and setting the display settings to include the Now Assist panel. To change this agentic workflow's instructions, duplicate it, adjust the settings to suit your specific needs, and activate the duplicated version instead.

Prerequisites and setup

To access this workflow, you must have Now Assist for Platform installed on your instance. You can get this by installing any other Now Assist application, such as Now Assist for IT Service Management (ITSM).

Now LLM is not a supported LLM provider for the Generate my work plan agentic workflow.

Role masking

Required role: sn_uxc_gen_ai.platform_ai_work_planner.

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with Now Assist applications have specific roles that come included with the application. If you select *`Users with specific roles`* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

In the data access settings, add the necessary roles to enable reading of the tables for the records you want to access for potential work plans. For example, add the itil role to the agentic workflow's list of approved roles so that it can access Incident records.

Additional configuration

Change different settings related to the agentic workflow by changing values for the Now Assist Skill Config Var Set. To access the variable set and make changes, do the following while in the Platform AI Agents and Skills scope:

- Go to the Now Assist Skill Config [sn_nowassist_skill_config] table.
- Open the record named *`Generate my work plan config`*.

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

- In the Now Assist Skill Config Var Set related list, select *New*.
- Enter a name for your Config Var Set.
- Set Config Type to be either `User schedule` or `Generate my work plan`.
- Save the Var Set record.
- Set the variables for the config type.
- Save the Var Set.

The work plan config var set includes general variables for running the agentic workflow.

Work plan config type configurations

Config field	Description
Schedule job	Creates a scheduled job that automatically runs, allowing users to see updated prioritization suggestions after making progress in their work.
Schedule job frequency	How long after a record in the original work plan is updated for the agentic workflow to run again. Default: 30 minutes
Predict estimates	Suggest estimated time to complete work per record.
Exclude task table	By default, certain fields are used for prioritizing work on the Task table. You can configure what fields to use on individual tables, such as the Incident table. Enable this to only use the table-specific configurations and not use the default fields.
Max tasks	The maximum number of tasks that can be recommended to a user for their work plan.

The Incident config var set includes the following variables. Configure either the AIS fields or the GAF field for determining how the agentic workflow gathers what work the user has. If you configure GAF, it takes priority when running the agentic workflow. For more information about GAF, see Group Action Framework.

Use this config var set as a template to change any values for other tables, such as the Case or HR Case tables.

Note: If you select *`Exclude task table`* in the work plan config var set, only the tables that have these var sets configured are included for work in the suggested work plan.

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

Incident table config type configuration

Config field	Description
Table	Table for record types that someone can work on
Conditions	Conditions for what kind of records someone can work on
Record fields	Fields used to judge how much work a task will take
Time worked	Custom or edited field for specifying time worked. By default, the agentic workflow uses the *`Time worked`* field.
AIS Search Profile	Specific AI Search profile to use to determine what work a user has
AIS Search Fields	Fields used by AI Search to determine what work a user has
AIS Return Fields	Fields returned by AI Search to the agentic workflow to base decisions on
GAF Config	Group Action Framework grouping configuration record, which is a collection of groups of records to make searching easier. By default, the agentic workflow uses the *`GAF ITSM grouping`*.
Assigned to field	The field for the user who works on the record. By default, the agentic workflow uses the *`assigned_to`* field.
Order by	JSON object containing the ordering information for how work records are selected for the work plan. For example, if there are 20 tasks assigned to an itil agent, but you only want to generate work plans with 5 tasks at a time, set the order to decide which 5 tasks are considered first. By default, the agentic workflow uses priority as its first ordering. Higher priority items are looked at first, then how soon work items are due, then how escalated the cases are, then oldest records first.

By default, the Generate my work plan agentic workflow relies on the User [sys_user] record's Schedule field. If there is no schedule defined on the User record, the workflow consults the Schedule [chm_schedule] table for the user. If neither are present, the assumed schedule is a weekday schedule from 8:00 a.m. to 5:00 p.m.

Change what tables and fields to look for schedule information with the *`User schedule`* config type.

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

Generate my work plan User schedule configuration

Config field	Description
Table	Table where user schedule information is stored
Conditions	(Optional) Conditions to determine what user schedule to associate with the user.
Schedule column	Field where specific schedule information is stored

Accessing the Generate my work plan agentic workflow

To access the agentic workflow:

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Select **Generate my work plan**.

The first step of the guided setup includes a complete list of included AI agents. Selecting the name of an AI agent opens it in AI Agent Studio, where you can see the full description, role, list of steps, and tools. Tools are displayed in the second step of the AI agent guided setup, Add tools and information.

In-product agentic AI and UI actions

Access agentic workflows in the Core UI and in workspaces in the AI Activity panel. From there, track their progress, provide or review input, and see the results of the work performed. For more information, see In-product agentic AI for more details about the AI Activity panel.

To enable users to access agentic workflows with UI actions, open the agentic workflow in AI Agent Studio and navigate to the **Select channels and access** step. Select a UI action as a possible way to access the workflow.

If you don't see your UI actions after configuring it in AI Agent Studio, verify that the property *`com.glide.agentic_processes_view.enabled`* is set to `true`. See Enable the in-product experience for agentic workflows.

Testing the Generate my work plan agentic workflow

Manually test an agentic workflow execution or access on the Testing page of AI Agent Studio if you have the sn.aia_admin role and all other roles configured in the security controls. Start a manual test, select a test type and the name of the workflow, and use utterances in the Task field like the following samples. See Test an agentic workflow execution.

To evaluate the agentic workflow over many different execution logs, run an automated evaluation.

Sample utterance

After the workflow is activated in AI Agent Studio, enter `What should I work on today?` or similar phrases in the Now Assist panel to trigger the workflow. You must have the

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

sn.now_assist_panel_user role to run the workflow. Also run this workflow on the Testing page of AI Agent Studio with the same utterance in the Task field if you have the sn.aia_admin role.

Users must have the sn_uxc_gen_ai.platform_ai_work_planner role to execute the agentic workflow.

AI agents used in the Generate my work plan agentic workflow

The following table lists the agents used in the Generate my work plan agentic workflow.

Important: In the Select channels and status step of each AI agent's guided setup, verify that the Status toggle is enabled to activate the AI agent.

AI agents names and descriptions in the Generate my work plan agentic workflow

AI agent name	AI agent description	Role required
Prioritize work AI Agent	Dynamically orders tasks based on parameters like urgency, priority, SLAs, sentiment, and impact.	sn_uxc_gen_ai.platform_ai_work_planner

Other Platform agentic workflows

For more information on other agentic workflows associated with the Platform workflow, see Platform agentic workflows.

Platform Generate resolution plan agentic workflow

Use the Platform Generate resolution plan AI agents agentic workflow to fetch task record details, generate resolution summary steps, and update comments or work notes.

Generate resolution plan overview

The Generate resolution plan agentic workflow can help resolve tasks by collecting record details and generating resolution summaries that can be added to comments or work notes. Due to the dynamic nature of AI agents, this agentic workflow can be used for tasks that require complex logic even when provided with minimal details.

The agents, tools, and triggers that are associated with the Generate resolution plan agentic workflow are provided by Now Assist applications. You can activate the agentic workflow template by making triggers active and setting the display settings to include the Now Assist panel. If you want to change this agentic workflow's instructions, you must duplicate it, adjust the settings to suit your specific needs, and activate the duplicated version of the agentic workflow instead.

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

Prerequisites and setup

To access this workflow, you must have Now Assist for Platform installed on your instance, which you can get if you install any other Now Assist application, such as Now Assist for IT Service Management (ITSM).

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

For this agentic workflow to behave as expected, you must also configure Group Action Framework (GAF). See Set up AI Search for Group Action Framework and Configure Group Action Framework for more information on getting started with GAF.

Role masking

Required role: sn_uxc_gen_ai.platform_ai_grp_workflow.

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with Now Assist applications have specific roles that come included with the application. If you select *`Users with specific roles`* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

In the data access settings, you must also add the necessary roles to enable reading of the tables for the records you want to access for potential resolution plans. For example, you can add the itil role to the agentic workflow's list of approved roles so that it can access Incident records.

Additional configuration

You can change different settings related to the agentic workflow by changing values for the Now Assist Skill Config Var Set. To access the variable set and make changes, do the following while in the Platform AI Agents and Skills scope:

- Go to the Now Assist Skill Config [sn_nowassist_skill_config] table.
- Open the record named *`Generate Resolution Plans Skill Config`*.
- In the Now Assist Skill Config Var Set related list, select the configuration variable set you want to edit.
- Set the variables for the configuration type.
- Save the Var Set.

The Generate resolution plan configuration variable set includes the following variables. You can configure either the AIS fields or the GAF field for determining how the agentic workflow gathers what work the user has. If you configure both, GAF takes priority when running the agentic workflow. For more information about GAF, see Group Action Framework.

Review button params configuration

Config field	Description	Default value
Show button	Display a UI action for reviewing a generated resolution plan	Enabled
Button label	Label for the UI action button to review a generated resolution plan	Review

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

Next Best Action Recommender configuration

Config field	Description
Search Mode	Options include Keyword, Semantic, or Hybrid. Keyword-based searches look at individual phrases, while semantic-based searches rely on the phrasing of utterances as well. Hybrid utilizes both strategies. Default: Keyword
Include Decomposition Agent in the workflow	Determines whether the Decomposition Agent, one of the agents in the workflow, is used. If this is unselected, the Next best action recommendation AI agent still runs.
use_websearch	Enables the agentic workflow to use web searches to help collect relevant information for generating the resolution plan
AIS Semantic Indexed Names	???A
script_to_run	Optional script to run when the Next best action recommendation AI agent executes
save_activity	Determines whether a tool creates an activity record on the sys_aia_agent_execution_activity table when saving resolution notes
AIS Search Profile	Profile for AI search, such as Now Assist in Virtual Agent.
Return Fields	Field values returned from identified relevant records
GAF Config	Group Action Framework grouping configuration record, which is a collection of groups of records to make searching easier
AIS Search Fields	Fields used by AI Search to determine what work a user has
AI Search Matching Threshold	Threshold value for whether AI Search considers a record to be related and relevant
Table	Table which has the records you want to generate resolution plans for

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

Next Best Action Recommender configuration (continued)

Config field	Description
Related List Tables	Tables which contain related records that you want the AI agent to explore before generating resolution plans

Decomposition Agent configuration

Config field	Description
agent	Name of the AI agent to use for generating resolution actions
tablesinfo	Table information for providing context for generated action steps

Accessing the Generate resolution plan agentic workflow

To access the agentic workflow:

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Select **Generate Resolution Plan**.

The first step of the guided setup includes a complete list of included AI agents. Selecting the name of an AI agent opens it in a new browser tab, where you can see the full description, role, list of steps, and tools. Tools are displayed in the second step of the AI agent guided setup, Add tools and information.

In-product agentic AI and UI actions

Agentic workflows can be accessed in the Core UI and in workspaces in the AI Activity panel. From there, you can track their progress, provide or review input, and see the results of the work performed. For more information, see In-product agentic AI for more details about the AI Activity panel.

To enable users to access agentic workflows with UI actions, you can open the agentic workflow in AI Agent Studio and navigate to the **Select channels and access** step. You can select a UI action as a possible way to access the workflow

If you don't see your UI actions after configuring it in AI Agent Studio, ensure that the property *com.glide.agentic_processes_view.enabled* is set to `true`. See Enable the in-product experience for agentic workflows.

Testing the Generate Resolution Plan agentic workflow

You can manually test an agentic workflow execution or access on the Testing page of AI Agent Studio if you have the sn.aia_admin role and all other roles configured in the security controls. Start a manual test, select a test type and the name of the workflow, and use utterances in the Task field like the following samples. See Test an agentic workflow execution.

If you want to evaluate the agentic workflow over many different execution logs, run an automated evaluation.

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

Sample utterance

After the workflow has been activated in AI Agent Studio, enter these or similar phrases in the Now Assist panel to trigger the workflow. You must have the sn.now_assist_panel_user role to run the workflow. You can also run this workflow on the Testing page of AI Agent Studio with the same utterance in the Task field if you have the sn.aia_admin role.

- Generate resolution plan for INC0001
- Create detailed resolution steps for INC0001
- Resolve INC0001

AI agents used in the Generate resolution plan agentic workflow

The following table lists the agents that are used in the Generate resolution plan agentic workflow.

Important: In the Select channels and status step of each AI agent's guided setup, make sure that the Status toggle is enabled to activate the AI agent.

AI agents names and descriptions in the Generate resolution plan agentic workflow

AI agent name	AI agent description	Role required
Next best action recommendation AI agent	Identifies the steps for resolving tasks by referencing the similar task details and reviewing knowledge articles.	sn_uxc_gen_ai.platform_ai_next_best_action
Decomposition Agent	Analyzes and breaks down each previous resolution step into smaller, actionable substeps, then creates records	sn_uxc_gen_ai.platform_ai_resolution_action_ai_agent

Other Platform agentic workflows

For more information on other agentic workflows associated with the Platform workflow, see Platform agentic workflows.

Platform Help optimize team productivity agentic workflow

Use the Platform Help optimize team productivity AI agents agentic workflow to gather relevant information about tasks automatically and make decisions about priorities and assignments.

Help optimize team productivity overview

The Help optimize team productivity agentic workflow offers multiple capabilities that can enhance your team's performance. Improving assignments based on previous work and balancing workloads, the agentic workflow can address problems that traditional work assignment rules may have.

The agentic workflow performs the following tasks:

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

- Analyzes historical performance of team members
- Calculates workloads relative to each team member's typical capacity
- Enables proactive team management with data-driven insights

The agents, tools, and triggers that are associated with the Help optimize team productivity agentic workflow are provided by Now Assist applications. You can activate the agentic workflow template by making triggers active and setting the display settings to include the Now Assist panel. If you want to change this agentic workflow's instructions, you must duplicate it, adjust the settings to suit your specific needs, and activate the duplicated version of the agentic workflow instead.

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

Prerequisites and setup

To access this workflow, you must have Now Assist for Platform installed on your instance, which you can get if you install any other Now Assist application, such as Now Assist for IT Service Management (ITSM).

Role masking

Required role: sn_uxc_gen_ai.platform_ai_help_allocate_work.

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with Now Assist applications have specific roles that come included with the application. If you select *`Users with specific roles`* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

In the data access settings, you must also add the necessary roles to enable reading of the tables for records assigned to team members. For example, you can add the itil role to the agentic workflow's list of approved roles so that it can access Incident records.

Additional configuration

You can change different settings related to the agentic workflow by changing values for the Now Assist Skill Config Var Set. To access the variable set and make changes, do the following while in the Platform AI Agents and Skills scope:

- Go to the Now Assist Skill Config [sn_nowassist_skill_config] table.
- Open the record named *`Optimal Ticket Assignment`*.
- In the Now Assist Skill Config Var Set related list, select *`Work Allocator`*.
- Edit the variable values.
- Save or update the record.

The Help optimize team productivity configuration variable sets include the following variables.

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

Help optimize team productivity Work Allocator configuration

Config field	Description
Tickets per batch	The number of tickets processed by the agentic workflow for evaluation. Default: 25
Metrics Calculation Window	How many days back to include in metric calculations. Default: 600
Assignment Group	The assignment group whose tickets are used for evaluation, analysis, and recommendations.
Table for Query	The table used to gather information about the work assigned to team members.
Backlog Ticket Query	Filter conditions for determining what counts as "backlog" work for team members to complete.
Similarity Score Table	The table used for gathering information about other work that has been previously assigned to team members.
Similarity Score Fields	Fields to consider for calculating which tasks are similar to the one that needs assigning. For example, selecting category and subcategory as similarity score fields includes those fields in calculating whether a team member has worked on similar assignments within the metric calculation window.

Help optimize team productivity Trigger frequency configuration

Config field	Description
Enable scheduled job	Enable the agentic workflow to run on its own using a scheduled job.
Assignment Group	The assignment group whose tickets are used for evaluation, analysis, and recommendations.
Scheduled job	The name of the scheduled job that runs to trigger the agentic workflow.
Scheduler frequency	How often the scheduled job runs and triggers the agentic workflow.

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

Accessing the Help optimize team productivity agentic workflow

To access the agentic workflow:

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Select **Help optimize team productivity**.

The first step of the guided setup includes a complete list of included AI agents. Selecting the name of an AI agent opens it in a new browser tab, where you can see the full description, role, list of steps, and tools. Tools are displayed in the second step of the AI agent guided setup, Add tools and information.

In-product agentic AI and UI actions

Agentic workflows can be accessed in the Core UI and in workspaces in the AI Activity panel. From there, you can track their progress, provide or review input, and see the results of the work performed. For more information, see In-product agentic AI for more details about the AI Activity panel.

To enable users to access agentic workflows with UI actions, you can open the agentic workflow in AI Agent Studio and navigate to the **Select channels and access** step. You can select a UI action as a possible way to access the workflow

If you don't see your UI actions after configuring it in AI Agent Studio, ensure that the property *com.glide.agentic_processes_view.enabled* is set to `true`. See Enable the in-product experience for agentic workflows.

Testing the Help optimize team productivity agentic workflow

You can manually test an agentic workflow execution or access on the Testing page of AI Agent Studio if you have the sn.aia_admin role and all other roles configured in the security controls. Start a manual test, select a test type and the name of the workflow, and use utterances in the Task field like the following samples. See Test an agentic workflow execution.

If you want to evaluate the agentic workflow over many different execution logs, run an automated evaluation.

Sample utterance

After the workflow has been activated in AI Agent Studio, enter "Give me optimized assignment evaluations for the Software assignment group" or similar phrases in the Now Assist panel to trigger the workflow. You must name the specific assignment group you're allocating work for.

You must have the sn.now_assist_panel_user role to run the workflow. You can also run this workflow on the Testing page of AI Agent Studio with the same utterance in the Task field if you have the sn.aia_admin role.

AI agents used in the Help optimize team productivity agentic workflow

The following table lists the agents that are used in the Help optimize team productivity agentic workflow.

Important: In the Define availability step of each AI agent's guided setup, make sure that the Status toggle is enabled to activate the AI agent.

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

AI agents names and descriptions in the Help optimize team productivity agentic workflow

AI agent name	AI agent description	Role required
Work Allocator AI Agent	Streamlines ticket management by automatically fetching backlog tickets, evaluating agent workloads and performance metrics, and assigning tickets to agents in a way that optimizes distribution and operational efficiency.	sn_uxc_gen_ai.platform_ai_help_allocate_work

Other Platform agentic workflows

For more information on other agentic workflows that are associated with the Platform workflow, see Platform agentic workflows.

Platform Identify ways to improve service agentic workflow

Use the Platform Identify ways to improve service agentic workflow to analyze feedback, performance metrics, and historical trends that identify areas for service improvement.

Identify ways to improve service overview

The Identify ways to improve service agentic workflow optimizes service delivery and customer satisfaction by analyzing feedback, metrics, and trends to provide actionable process improvement recommendations. After the analysis is generated, you can continue the conversation to ask follow-up questions or download the analysis as a PDF or Word document.

The Now Assist applications provide the agents, tools, and triggers for the Identify ways to improve service agentic workflow. You can activate the agentic workflow template by making triggers active and setting the display settings to include the Now Assist panel. To change this agentic workflow's instructions, duplicate it, adjust the settings to suit your specific needs, and activate the duplicated version instead.

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

Prerequisites and setup

To access this workflow, you must have Now Assist for Platform installed on your instance. You can get this by installing any other Now Assist application, such as Now Assist for IT Service Management (ITSM).

Because this agentic workflow relies on survey data, you must have Assessment records associated with task tables to analyze.

Role masking

Required role: sn_uxc_gen_ai.platform_ai_improve_services.

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with Now Assist applications have specific roles that come included with the application. If you select *`Users with specific roles`* for user access, you must

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

configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

In the data access settings, add the necessary roles to enable reading of the Assessment table and other required tables. For example, add the itil role to the agentic workflow's list of approved roles so it can access Incident records.

Additional configuration

You can change agentic workflow settings by changing values for the Now Assist Skill Config Var Set. To access the variable set and make changes, do the following while in the Platform AI Agents and Skills scope:

- Go to the Now Assist Skill Config [sn_nowassist_skill_config] table.
- Open the record named *`Identify ways to improve service`*.
- In the Now Assist Skill Config Var Set related list, select *`Survey analysis input config`*.
- Edit the variable values.
- Save or update the record.

Identify ways to improve service configuration

Config field	Description	Default value
Maximum number of records	Maximum number of records included in analysis	500
Analysis Time Frame	Time range, in months, for the survey analyzer to examine records to identify trends. Users can specify smaller ranges when running the agentic workflow, but this value defines the maximum limit.	3

Accessing the Identify ways to improve service agentic workflow

To access the agentic workflow:

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Select **Identify ways to improve service**.

The first step of the guided setup includes a complete list of included AI agents. Selecting an AI agent name opens it in a new browser tab, where you can see the full description, role, list of steps, and tools. Tools are displayed in the second step of the AI agent guided setup, Add tools and information.

In-product agentic AI and UI actions

You can access agentic workflows in the Core UI and in workspaces in the AI Activity panel. From there, you can track their progress, provide or review input, and see the results of the work performed. For more information, see In-product agentic AI for more details about the AI Activity panel.

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

To enable users to access agentic workflows with UI actions, open the agentic workflow in AI Agent Studio and navigate to the **Select channels and access** step. You can select a UI action as a possible way to access the workflow

If you don't see your UI actions after configuring it in AI Agent Studio, verify that the property *com.glide.agentic_processes_view.enabled* is set to `true`. See Enable the in-product experience for agentic workflows.

Testing the Identify ways to improve service agentic workflow

You can manually test an agentic workflow execution or access on the Testing page of AI Agent Studio if you have the sn.aia_admin role and all other roles configured in the security controls. Start a manual test, select a test type and the workflow name, and use utterances in the Task field like the following samples. See Test an agentic workflow execution.

To evaluate the agentic workflow over many different execution logs, run an automated evaluation.

Sample utterance

After the workflow is activated in AI Agent Studio, enter similar phrases to the following in the Now Assist panel to trigger the workflow. You can also run this workflow on the Testing page of AI Agent Studio with the same utterance in the Task field if you have the sn.aia_admin role.

- Make suggestions for continuous service improvement based on surveys within the last 3 months
- Summarize surveys to improve services within the last 2 years
- Give ideas for improving case resolution based on Post-Case survey results over the last month

AI agents used in the Identify ways to improve service agentic workflow

The following table lists the agents used in the Identify ways to improve service agentic workflow.

Important: In the Select channels and access of each AI agent's guided setup, verify that the Status toggle is enabled to activate the AI agent.

AI agents names and descriptions in the Identify ways to improve service agentic workflow

AI agent name	AI agent description	Role required
Survey Analysis AI agent	Acquires survey data within the specified time range and analyzes it for ways to improve services	sn_uxc_gen_ai.platform_ai_improve_services

Other Platform agentic workflows

For more information on other agentic workflows associated with the Platform workflow, see Platform agentic workflows.

Platform Investigate problems agentic workflow

Use the Platform Investigate problems AI agents agentic workflow to perform root cause and risk assessments so that you can create an actionable resolution plan for a problem.

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

Investigate problems overview

The Investigate problems agentic workflow can help to assist agents and subject matter experts (SMEs) in investigating problems in their IT landscape. A problem can be associated with many incidents, and any investigator must be aware of a large number of details when looking at a problem. The agentic workflow can help provide insights from the incident and problem details and suggest plans or possible solutions.

The agents, tools, and triggers that are associated with the investigate problems agentic workflow are provided by Now Assist applications. You can activate the agentic workflow template by making triggers active and setting the display settings to include the Now Assist panel. If you want to change this agentic workflow's instructions, you must duplicate it, adjust the settings to suit your specific needs, and activate the duplicated version of the agentic workflow instead.

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

Prerequisites and setup

To access this workflow, you must have Now Assist for Platform installed on your instance, which you can get if you install any other Now Assist application, such as Now Assist for IT Service Management (ITSM).

Because this agentic workflow analyzes problems and incidents related to those problems, you must have records on the Problem and Incident table.

Role masking

Required role: sn_uxc_gen_ai.platform_ai_problem_investigator.

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with Now Assist applications have specific roles that come included with the application. If you select *`Users with specific roles`* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

In the data access settings, you must also add the necessary roles to enable reading of the tables for Problems and related records. For example, you can add the itil role to the agentic workflow's list of approved roles so that it can access Incident records.

Accessing the Investigate problems agentic workflow

To access the agentic workflow:

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Select **Investigate problems**.

The first step of the guided setup includes a complete list of included AI agents. Selecting the name of an AI agent opens it in a new browser tab, where you can see the full description, role, list of steps, and tools. Tools are displayed in the second step of the AI agent guided setup, Add tools and information.

In-product agentic AI and UI actions

Agentic workflows can be accessed in the Core UI and in workspaces in the AI Activity panel. From there, you can track their progress, provide or review input, and see the results of the work

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

performed. For more information, see In-product agentic AI for more details about the AI Activity panel.

To enable users to access agentic workflows with UI actions, you can open the agentic workflow in AI Agent Studio and navigate to the **Select channels and access** step. You can select a UI action as a possible way to access the workflow

If you don't see your UI actions after configuring it in AI Agent Studio, ensure that the property *com.glide.agentic_processes_view.enabled* is set to `true`. See Enable the in-product experience for agentic workflows.

Testing the Investigate problems agentic workflow

You can manually test an agentic workflow execution or access on the Testing page of AI Agent Studio if you have the sn.aia_admin role and all other roles configured in the security controls. Start a manual test, select a test type and the name of the workflow, and use utterances in the Task field like the following samples. See Test an agentic workflow execution.

If you want to evaluate the agentic workflow over many different execution logs, run an automated evaluation.

Sample utterance

After the workflow has been activated in AI Agent Studio, enter `investigate the problem PRB001` or similar phrases in the Now Assist panel to trigger the workflow. You can also run this workflow on the Testing page of AI Agent Studio with the same utterance in the Task field if you have the sn.aia_admin role.

Troubleshooting

If a Problem has a large number of related incidents, you may run into an error that states "This model's maximum context length is 128000 tokens. However, your messages resulted in [X] tokens. Please reduce the length of the messages." This maximum token count is in place for all Now Assist skills, so there is currently no way to work around this error. You can try the agentic workflow again using a different Problem.

AI agents used in the Investigate problems agentic workflow

The following table lists the agents that are used in the Investigate problems agentic workflow.

Important: In the Define availability step of each AI agent's guided setup, make sure that the Status toggle is enabled to activate the AI agent.

AI agents names and descriptions in the Investigate problems agentic workflow

AI agent name	AI agent description
Problem investigation AI agent	Identifies the root causes, performs an impact assessment, and plans resolutions.

Other Platform agentic workflows

For more information on other agentic workflows associated with the Platform workflow, see Platform agentic workflows.

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

Platform Process images for new tasks agentic workflow

Use the Platform Process images for new tasks agentic workflow to convert images to actionable tasks.

Process images for new tasks overview

The process images for new tasks agentic workflow can help decrease manual data entry and improve task organization by automatically converting images into task records. First, the agentic workflow extracts information from the image, such as error messages, and presents the analysis to the user. Then, the user is presented with the details of the task, such as short description, category, and priority, before it is submitted so that they can make any changes. Once the information is confirmed, the agentic workflow creates an incident record and attaches the image.

The agents, tools, and triggers that are associated with the process images for new tasks agentic workflow are provided by Now Assist applications. You can activate the agentic workflow template and set the display settings to include the Now Assist panel. If you want to change this agentic workflow's instructions, you must duplicate it to create a custom agentic workflow, adjust the settings to suit your specific needs, and activate the duplicated version of the agentic workflow instead.

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

Prerequisites and setup

To access this workflow, you must have Now Assist for Platform installed on your instance, which you can get if you install any other Now Assist application, such as Now Assist for IT Service Management (ITSM).

Users must have the *sn_uxc_gen_ai.platform_ai_image_processor* role to invoke the agentic workflow.

If you want the ability for users to create tasks from images using Now Assist for Virtual Agent, you must activate the Image Processor Agent, Record management AI agent, and Document and visual insights AI agent and set the display to include Virtual Agent. This agentic workflow cannot be discovered in Virtual Agent, so you must enable the individual AI agents that comprise it.

Role masking

Required role: sn_uxc_gen_ai.platform_ai_image_processor.

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with Now Assist applications have specific roles that come included with the application. If you select *Users with specific roles* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

In the data access settings, you must also add the necessary roles to enable reading of the tables for the records you want to be able to make tasks on. For example, you can add the itil role to the agentic workflow's list of approved roles so that it can access Incident records.

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

Additional configuration

You can change different settings related to the agentic workflow by changing values for the Now Assist Skill Config Var Set. To access the variable set and make changes, do the following while in the Platform AI Agents and Skills scope:

- Go to the Now Assist Skill Config [sn_nowassist_skill_config] table.
- Open the record named *`Image to task config`*.
- In the Now Assist Skill Config Var Set related list, select the *`Tables Mapping`* configuration variable set.
- Set the variables for the configuration type.
- Save the Var Set.

Tables Mapping configuration

Config field	Description
role_and_table_map	When a user triggers this agentic workflow, it checks the role of the invoking user and uses the map to determine what tables a user can make a record on. For example, if you map itil to Incident table, then users with the itil role can make Incident records with images.
get_location_details	Determines whether information about the image's location is included in metadata extraction.
tableslist	List of tables where you can make tasks from images

Accessing the Process images for new tasks agentic workflow

To access the agentic workflow:

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Select **Process images for new tasks**.

The first step of the guided setup includes a complete list of included AI agents. Selecting the name of an AI agent opens it in a new browser tab, where you can see the full description, role, list of steps, and tools. Tools are displayed in the second step of the AI agent guided setup, Add tools and information.

In-product agentic AI and UI actions

Agentic workflows can be accessed in the Core UI and in workspaces in the AI Activity panel. From there, you can track their progress, provide or review input, and see the results of the work performed. For more information, see In-product agentic AI for more details about the AI Activity panel.

To enable users to access agentic workflows with UI actions, you can open the agentic workflow in AI Agent Studio and navigate to the **Select channels and access** step. You can select a UI action as a possible way to access the workflow

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

If you don't see your UI actions after configuring it in AI Agent Studio, ensure that the property *com.glide.agentic_processes_view.enabled* is set to `true`. See Enable the in-product experience for agentic workflows.

Sample utterance

After the workflow has been activated in AI Agent Studio, enter `Convert image to new task` or similar phrases in the Now Assist panel to trigger the workflow. You can also run this workflow on the Testing page of AI Agent Studio with the same utterance in the Task field if you have the sn.aia_admin and sn_uxc_gen_ai.platform_ai_image_processor roles.

AI agents used in the Process images for new tasks agentic workflow

The following table lists the agents that are used in the Process images for new tasks agentic workflow.

Important: In the Define availability step of each AI agent's guided setup, make sure that the Status toggle is enabled to activate the AI agent.

AI agents names and descriptions in the Process images for new tasks agentic workflow

AI agent name	AI agent description	Role required
Document and visual insights AI agent	Extracts information from images.	platform_ml_di.creation_agent

Other Platform agentic workflows

For more information on other agentic workflows associated with the Platform workflow, see Platform agentic workflows.

Platform Propose survey responses agentic workflow

Use the Platform Propose survey responses AI agents agentic workflow to assist requesters in completing surveys.

Propose survey responses overview

The Propose survey responses agentic workflow can help simplify and increase survey responses. Many callers end up not filling out surveys after requests have been fulfilled, and this agentic workflow helps them to make informed decisions to answer survey questions quickly.

When this agentic workflow and trigger are activated, the assignee receives an email with AI-suggested answers to their survey based on the associated Incident or Request. They then have the option to accept the AI-generated answers with a link at the bottom of the email. They can also choose to fill out the survey manually.

The agents, tools, and triggers that are associated with the Propose survey responses agentic workflow are provided by Now Assist applications. You can activate the agentic workflow template by making the trigger active. If you want to change this agentic workflow's instructions, you must duplicate it, adjust the settings to suit your needs, and activate the duplicated version of the agentic workflow instead.

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

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

Prerequisites and setup

To access this workflow, you must have Now Assist for Platform installed on your instance, which you can get if you install any other Now Assist application, such as Now Assist for IT Service Management (ITSM).

You can use this workflow for any survey triggered on the Incident or Request table. For example, this agentic workflow can be used by the Short Customer Satisfaction Survey with Smiley Face when an Incident is closed. You can change existing surveys, and the agentic workflow can still propose answers. The agentic workflow can also be used for custom surveys triggered by the Incident or Request table as long as there's a trigger associated with the survey. See Configure a trigger condition for a survey for instructions on adding a trigger to a survey.

Propose survey responses isn't available for Now Assist panel.

Role masking

Required role: sn_uxc_gen_ai.platform_ai_survey_response.

Agentic workflows and their AI agents use role masking to determine which users can access them. Ones installed with Now Assist applications have specific roles that come included with the application. If you select *Users with specific roles* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an agentic workflow.

In the data access settings, you must also add the necessary roles to enable reading of survey tables and other related tables.

Accessing the Propose survey responses agentic workflow

To access the agentic workflow:

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Select **Propose survey responses**.

The first step of the guided setup includes a complete list of included AI agents. Selecting the name of an AI agent opens it in a new browser tab, where you can see the full description, role, list of steps, and tools. Tools are displayed in the second step of the AI agent guided setup, Add tools and information.

In-product agentic AI and UI actions

Agentic workflows can be accessed in the Core UI and in workspaces in the AI Activity panel. From there, you can track their progress, provide or review input, and see the results of the work performed. For more information, see In-product agentic AI for more details about the AI Activity panel.

To enable users to access agentic workflows with UI actions, you can open the agentic workflow in AI Agent Studio and navigate to the **Select channels and access** step. You can select a UI action as a possible way to access the workflow

If you don't see your UI actions after configuring it in AI Agent Studio, ensure that the property *com.glide.agentic_processes_view.enabled* is set to `true`. See Enable the in-product experience for agentic workflows.

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

Sample utterance

After the workflow has been activated in AI Agent Studio, you need to set a trigger for the workflow to run. This workflow cannot be triggered from the Now Assist panel.

If you want to test the agentic workflow in AI Agent Studio and you have the sn.aia_admin role, enter the following phrase in the Testing page in the Task field: `Help me with survey AINST00XXXX. List of fields to extract inc/case/req details: number, short_description, description, calendar_stc, escalation, reopen_count, close_code, close_notes, state, priority, caller_id.` You must include the list of fields to extract.

AI agents used in the Propose survey responses agentic workflow

The following table lists the agents that are used in the Propose survey responses agentic workflow.

Important: In the Define availability step of each AI agent's guided setup, make sure that the Status toggle is enabled to activate the AI agent.

AI agents names and descriptions in the Propose survey responses agentic workflow

AI agent name	AI agent description	Role required
Survey response suggestion AI agent	Suggests answers to questions based on record details.	sn_uxc_gen_ai.platform_ai_survey_response
Survey filling data collection AI agent	Collects the data related to the record and the survey questions to gather feedback.	sn_uxc_gen_ai.platform_ai_survey_response

Other Platform agentic workflows

For more information on other agentic workflows that are associated with the Platform workflow, see Platform agentic workflows.

Platform AI agents

Explore different AI agents available in platform Now Assist applications.

There might be AI agents installed with the Now Assist application that are not used in agentic workflows. To learn how to see all agents that are available to you, see Find AI agents.

AI agents are also provided with some of the Integration Hub spokes or plugins that are preconfigured and are callable from interfaces such as AI Agent Studio and Virtual Agent. For more information, see Available AI agents for Integration Hub.

Platform Approval assistance AI agent

The Approval assistance AI agent is an AI agent that enables you to see your list of pending approvals, as well as see the details about your pending approvals. You can then approve or reject requests and tickets from Now Assist in Virtual Agent.

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

Approval assistance AI agent overview

Roles required: approver_user, sn_request_approver_read

Note: Add the necessary roles to enable reading of other tables whose records go through the approval process. For example, you can add the sn_change_read role to read Change Request records.

The benefit of using the Approval assistance AI agent is that you don't need to navigate to a specific page to approve your tickets. You can ask the AI agent about your pending approval requests and then tell the AI agent to approve or reject those approvals. The AI agent will ask follow-up questions and offer context-aware responses to simplify your experience.

Ask you administrator to configure the display fields and the knowledge base (KB) search fields to generate a Gen AI checklist to assist the Approval assistance AI agent in making targeted decisions. The checklist uses KB articles and policies to assist the Approval assistance AI agent in decision making. The checklist fetches information from knowledge base articles about specific requester approval tickets. An approval_admin and admin role are required to configure the Approval assistance AI agent. For more information, see Configure Service Portal Approval Configuration record.

Note: Provide cross-scope privileges to the Now Assist Agents for Requester plugin for tables whose records are restricted within the scope of an application. For more information, see Define cross-scope access to an application resource.

Prerequisites and setup

To access this AI agent, you must have Now Assist Agent for Requester (version 3 of the standard plugin) installed on your instance, which is installed with Now Assist for Platform. You can get Now Assist for Platform when you install any other Now Assist application, such as Now Assist for IT Service Management (ITSM).

The Approval assistance AI agent displays data from the Approvals [sysapproval_approver] table. If the user has been assigned to approve a request, the approval record is shown.

Role masking

Required role:

AI agents use role masking to determine which users can access them. Ones installed with Now Assist applications have specific roles that come included with the application. If you select *`Users with specific roles`* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an AI agent.

In the data access settings, you must also add the necessary roles to enable reading of the tables for the records you want to evaluate for readiness. For example, you can add the itil role to the AI agent's list of approved roles so that it can access Incident records.

Accessing the Approval assistance AI agent

To access the AI agent:

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Go to the **AI Agents** tab.
3. Select **Approval assistance AI agent**.

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

Sample utterances

After the agent has been activated in AI Agent Studio, enter phrases such as the following or similar queries to run the AI Agent in Virtual Agent, the *Now Assist* panel, or Microsoft Teams.

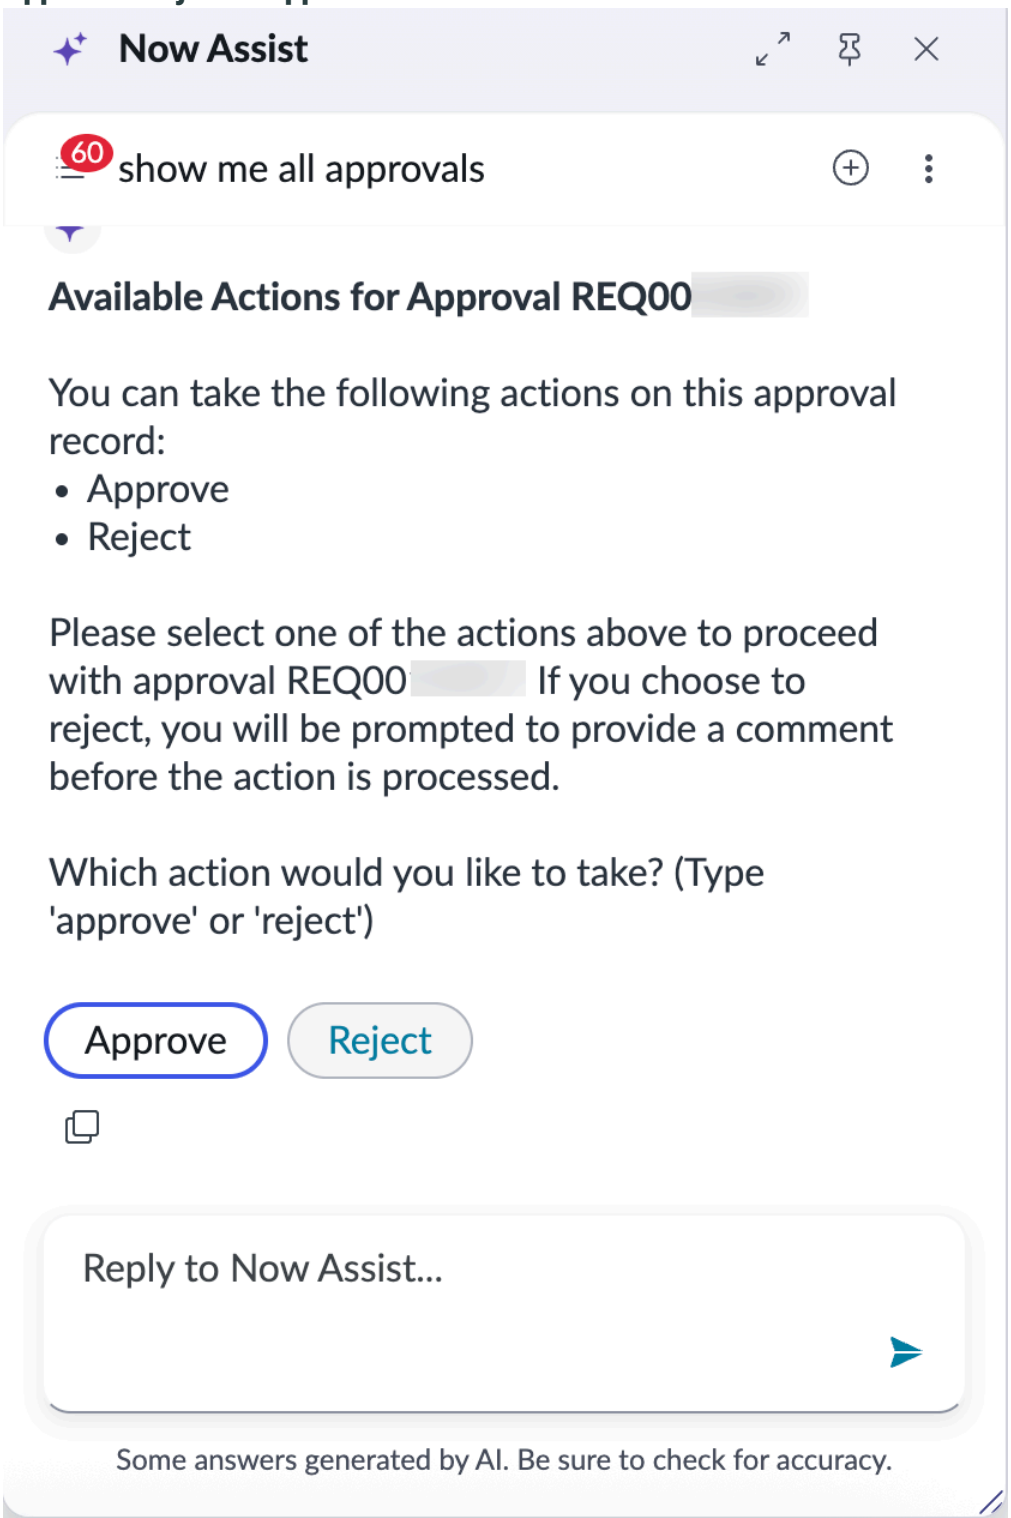
- Can you give me a list of pending approvals?
- What are the pending approvals for time off requests?
- Show me a list of pending approval requests by priority.
- Give me details about my approval request?

Performing approval action on an approval record

The Approval assistance AI agent shows approval requests that require an action on your part. After the approval information is provided by the agent, select **Approve** or **Reject** for a specific approval record, or type `Approve` or `Reject`. If you reject an approval request, then add a comment about the rejection reason.

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

Approve or reject an approval record



Platform Issue Readiness AI agent

The Issue Readiness AI agent analyzes the current state of a task to assess and determines whether it is ready to be worked on.

Issue Readiness AI agent overview

The Issue Readiness AI agent takes the information on a record, analyzes it, determines whether all relevant information is present on the record, and makes a decision about whether the task

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

is ready for work. Using the AI agent can help fulfillers focus on actionable work by reducing manual checks for missing supporting details or back-and-forth conversations. Instead of having to check whether tasks are waiting on pending approvals or unanswered questions, fulfillers can see what actionable work is available and prioritize it.

Prerequisites and setup

To access this AI agent, you must have Now Assist Agent for Platform installed on your instance, which is installed with any other Now Assist application, such as Now Assist for IT Service Management (ITSM).

Role masking

Required role:

AI agents use role masking to determine which users can access them. Ones installed with Now Assist applications have specific roles that come included with the application. If you select *`Users with specific roles`* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an AI agent.

In the data access settings, you must also add the necessary roles to enable reading of the tables for the records you want to evaluate for readiness. For example, you can add the itil role to the AI agent's list of approved roles so that it can access Incident records.

Accessing the Issue Readiness AI agent

To access the AI agent:

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Go to the **AI Agents** tab.
3. Select **Issue Readiness**.

Testing the Issue Readiness AI agent

You can manually test the AI agent execution on the Testing page of AI Agent Studio if you have the sn.aia_admin role and all other roles configured in the security controls. Select the AI agent, start a manual test, and use utterances in the Task field like the ones below.

Sample utterances

After the agent has been activated in AI Agent Studio, enter phrases such as the following or similar queries to run the AI Agent in Virtual Agent, the *`Now Assist`* panel, or Microsoft Teams.

- Is INC001234 ready for work?
- Do I have enough details on CSM0001234 to start?

Platform Request status AI agent

The Request status AI agent is an AI agent that enables you to view tickets and make updates to them from Now Assist in Virtual Agent, the Now Assist panel, or Microsoft Teams.

Request status AI agent overview

The Request status AI agent enables you to view your open tickets, check the status of tickets, and add comments through Now Assist in Virtual Agent, the Now Assist panel, or Microsoft Teams.

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

When you ask for the details of a request, you can perform any other ticket tasks configured by your administrator in the Standard Ticket configuration, such as reopening an incident, and resolving a ticket. You don't need to navigate to a specific page to view your tickets. The AI agent can ask follow-up questions and offer context-aware responses to simplify your experience.

The tools and triggers that are associated with the Request status AI agent are provided by Now Assist applications. You can activate the AI agent by making triggers active and setting the display settings to include Virtual Agent. If you want to change this AI agent's instructions, you mustIf you want to change this AI agent's instructions, you must duplicate it, adjust the settings to suit your specific needs, and activate the duplicated version of the AI agent instead.

Request status AI agent overview

The Request status AI agent enables you to view your open tickets, check the status of tickets, view a summary of the previous comments from the AI agent or others working on the ticket, and add comments through Now Assist in Virtual Agent, the Now Assist panel, or Microsoft Teams.

You can upload a file as an attachment to an open ticket or incident to support a request action. For example, if you lose your identity card you may request a replacement in the Request status AI agent. You may be asked to upload an email or document that has your manager's approval to get the replacement ID card. In the Request status AI agent, you just type, I need to attach a document to this incident or ticket. The Request status AI agent then provides the **Click here to upload a file** option within the AI agent chat for you to upload an attachment to the ticket. You upload the manager approval and the service agent can then approve your request for a new ID.

To find more information about an open ticket, you can ask the Request status AI agent follow-up questions based on previous answers from the agent. For example:

- Can you check the latest progress on my most recent request?
- When did I submit this ticket?
- Who is working on it?
- I want to add a comment to my ticket.
- I also want to add an attachment to my ticket.

When you ask for the details of a request, you can perform any other ticket tasks configured by your administrator in the Standard Ticket configuration, such as reopening an incident, and resolving a ticket. You don't need to navigate to a specific page to view your tickets. The AI agent can ask follow-up questions and offer context-aware responses to simplify your experience.

The tools and triggers that are associated with the Request status AI agent are provided by Now Assist applications. You can activate the AI agent by making triggers active and setting the display settings to include Virtual Agent. If you want to change this AI agent's instructions, you must duplicate it, adjust the settings to suit your specific needs, and activate the duplicated version of the AI agent instead.

Prerequisites and setup

To access this AI agent, you must have Now Assist Agent for Requester installed on your instance, which is installed with Now Assist for Platform. You can get Now Assist for Platform when you install any other Now Assist application, such as Now Assist for IT Service Management (ITSM).

To configure which actions are available, a user with the admin or sp_admin role can configure the Standard Ticket configuration for a table. In the Standard Ticket actions related list, you can add, change, or remove actions. All actions available from the Standard Ticket configuration can

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

be used by the Request status AI agent. See Configure actions for standard ticket page for more information.

To make the Request status agent available for users, you must navigate to the **Toggle display** step of the guided setup in AI Agent Studio. Toggle Virtual Agent to `true` and select an assistant.

To make the Request status AI agent available in Microsoft Teams, you must configure an assistant for Now Assist in Virtual Agent to use a Teams channel. See Display your assistant on a portal or channel for steps to enable Teams for your assistant. Then, in the Request status AI agent guided setup in AI Agent Studio, select the assistant you configured for Microsoft Teams in the **Toggle display** step.

Role masking

Required role:

AI agents use role masking to determine which users can access them. Ones installed with Now Assist applications have specific roles that come included with the application. If you select *`Users with specific roles`* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an AI agent.

In the data access settings, you must also add the necessary roles to enable reading of the tables for the records you want to evaluate for readiness. For example, you can add the itil role to the AI agent's list of approved roles so that it can access Incident records.

Accessing the Request status AI agent

To access the AI agent:

1. Navigate to **All** > **AI Agent Studio** > **Create and manage**.
2. Go to the **AI Agents** tab.
3. Select **Request status AI agent**.

Testing the Request Status AI agent

You can test the Request Status AI agent on the Testing page of AI Agent Studio if you have the sn.aia_admin role. Select the AI agent and use utterances in the **Task** field like the following examples:

Sample utterances

After the agent has been activated in AI Agent Studio, enter phrases such as the following or similar queries to run the AI Agent in Virtual Agent, the *`Now Assist`* panel, or Microsoft Teams.

- Can you list all open tickets I've created?
- What's the current status of my incident INC001?
- Add comment to INC001 "I need your help in fixing this request ASAP. Please prioritize this."
- Can you check the latest progress on my most recent request?

Large language models on the ServiceNow AI Platform®

ServiceNow AI Platform® utilises large language models (LLM) for generative AI and agentic AI functionality. You can choose the ServiceNow® LLM or supported third-party models. LLMs

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

bring about the Generative AI application capabilities by understanding and generating human language, and processing vast amount of data.

Third-party model strategy

ServiceNow® supports third-party AI model providers, which signifies that the admin persona can select a third-party large language or an AI model provider for Now Assist skills that are available out of box. The Now LLM Service is the default LLM. Find the supported model providers, the LLM services and the version.

Third party AI model providers

Model provider	Large language model	Version	Deployment policy
ServiceNow®	Now LLM Service	Latest update for your instance	Default
Microsoft Azure	Azure OpenAI	OpenAI GPT-4.1 and GPT-4.1-mini	https://learn.microsoft.com/en-us/azure/ai-services/openai/how-to/deployment-types#global-standard
Google	Google Gemini	Gemini 2.5 Flash and 2.5 Pro	https://cloud.google.com/vertex-ai/generative-ai/docs/learn/locations
AWS	AWS Claude	Claude 3.7 Sonnet	https://docs.aws.amazon.com/bedrock/latest/userguide/models-regions.html **Note:** Some out of box skills, Agents, and Agentic Workflows do not support Now LLM Service or third party model providers. To learn more about the exceptions to LLM service and third party model providers support, see https://support.servicenow.com/kb?id=kb_article_view&sysparm_articl

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

Note: Some out of box skills, Agents, and Agentic Workflows do not support Now LLM Service or third party model providers. To learn more about the exceptions to LLM service and third party model providers support, see https://support.servicenow.com/kb?id=kb_article_view&sysparm_article=KB2222333

Large language model selection

This selection is supported at various levels such as, skill, skill group and instance levels. However, the controls that define LLMs mapping to different skills, allowed at various global locations, are configured and approved in AI Control Tower by the AI steward. The model provider options for custom skills created in Now Assist skill kit and AI Agent Studio are also configured in AI Control Tower. See Configure third-party LLMs using AI Control Tower to learn more.

The **Manage large language models** feature enables the admin persona to perform additional actions like:

- Edit the model provider at the instance, skill group and skill levels
- Deactivate active skills which are non-compliant with fallback as **No**. See Explore AI model providers to learn more.
- View audit history on policy updates by AI steward in AI Control Tower
- Select allowed model providers across domain separated instances.

See Manage AI models to learn about managing.

For more information on selecting and updating model providers, see Manage model providers.

Now LLM Service updates

The Now LLM Service provides access to specialized large language models (LLMs) that are developed by ServiceNow. It also provides access to open-source LLMs that are selected, configured, or enhanced by ServiceNow, from the ServiceNow community and partners. Review these reference materials and model cards for additional information about the Now LLM Service and about the models used.

Model cards

Large language models (LLMs) are complex machine-learning models that are trained on large datasets like websites and documentation to perform language-related tasks, such as text generation for case summaries and resolution notes.

Model cards explain the specific model's context, intended use, training data, limitations, and other important information.

These model cards are for skills that use the Now LLM Service. There are certain skills, such as Now Assist Multi-Turn Catalog Ordering, that use Azure OpenAI instead. To see what LLM a skill is using, you can check the skill list in the Now Assist Admin console and review the LLM service column.

Model card for ServiceNow large language model

Model used for AI-driven solutions to support natural language understanding, automation, and decision support.

This model card is available in Yokohama patch 1 and later.

Model card for ServiceNow large language model (V2)

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

Model for enterprise AI that enhances text-based automation and content generation in ServiceNow workflows, including requester OOTB skills, custom skills, and agentic use cases.

This model is for Generative AI Controller application 11.2 or higher.

Model card for ServiceNow small language model

Model used for enterprise AI applications by enhancing text-based automation and content generation within ServiceNow workflows.

This model card is available in Yokohama patch 1 and later.

Model card for ServiceNow small language model (V2)

Model for enterprise AI that enhances text-based automation and content generation in ServiceNow workflows, including creator and fulfiller OOTB skills as well as custom skills.

This model is for Generative AI Controller application 11.2 or higher.

Model card for ServiceNow third party large language model

Model used for AI-driven solutions for text generation, summarization, and conversational AI.

Model card for ServiceNow Voice AI Speech-to-Text and Text-to-Speech models

Models used within ServiceNow AI Voice Agents for converting spoken user input to text and generating natural-sounding speech from AI responses.

Model card for ServiceNow Now Assist Guardian model

This model provides content moderation and helps identify different kinds of prompt injection attacks and offensive content.

Model card for ServiceNow Inferred CSAT and Factors large language model

This model is designed to ingest a conversation and predict a CSAT score as well as factors that explain the predicted score.

May 2025

An advanced 12B general-purpose small language model (SLM) with a singular, high-performance architecture that supports a wide range of tasks in ServiceNow's context was released. Fine-tuned on Mistral-Nemo-12B-Instruct, this model is designed and optimized for tasks like Agent Assist, Text-to-Flow, Text-to-Cypher, Safety & Content Moderation and Text-to-Code.

Key Enhancements:

- Enhanced instruction adherence: Improved the model's capability to accurately interpret and follow user instructions, ensuring that the model can better understand and execute complex commands. Leading to more precise and reliable outcomes than previous releases.
- Increased context window: increased context window from 16K to 32K, enabling the model to better understand long-form inputs, maintain coherence over extended interactions, and support more complex tasks with richer contextual awareness.
- Improved multilingual proficiency: Boosted performance across languages compared to previous releases, with notable enhancements in Japanese processing.

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

- Optimized for ServiceNow workflow related capabilities: Extended support coverage for Text-to-Flow, and improved the performance of Text-to-Code, Text-to-Cypher etc.
- Continuously enhanced model deployment consolidation: Integrates ServiceNow-related tasks into a single model, reducing system complexity at the same time while elevating overall performance.

March 2025

A powerful 12B general-purpose small language model (SLM) designed to enhance a wide range of applications, including text-to-code and agent use cases was released. Fine-tuned on Mistral-Nemo-12B, it streamlines deployment and consolidates multiple functionalities into a singular, architecture.

Key Enhancements:

- Optimized to fulfill use cases: Enhances case summarization, chat summarization, resolution notes, and knowledge base generation across supported languages, including improvements in Japanese quality.
- Superior text-to-code and text-to-cypher performance: Delivers major advancements in Glide JavaScript and generic JavaScript editing and generation, along with improved accuracy in query generation and execution for structured databases.
- Robust content moderation and safety: Provides stronger protection against adversarial prompts, jail-breaking attempts, and harmful content generation, ensuring safer deployment with built-in content filtering.
- Unified model deployment:integrates ServiceNow-related tasks into a single model, thereby reducing system complexity while elevating overall performance.
- Improved instruction adherence: Delivers better instruction following and consistency across varying levels of prompt and instruction strictness than the current text-to-text NowLLM.

November 2024

Several key improvements were added to the Now LLM Service that are aimed at enhancing performance and quality.

- Multilingual support: Now LLM Service supports 8 additional languages, enabling global teams to use the model in their native languages.

  The supported languages are: English, German, French, Japanese, Dutch, French Canadian, Spanish, Brazilian Portuguese, and Italian.

- JSON format support: The model now provides output in JSON format, making it easier for developers to integrate with various applications and automate workflows seamlessly.
  - Deterministic responses: JSON mode ensures structured, consistent output, which improves predictability and reliability when integrating with applications.
  - Error reduction: Unlike free-form text mode, JSON responses are less prone to format errors or stray characters, minimizing integration issues.
  - Lower token consumption: The fixed structure of JSON can reduce token usage, making it more efficient and cost-effective for applications with high response frequency.
- Improvements in instruction following: The model has been fine-tuned to understand and follow instructions more precisely. This enables the model to deliver more to-the-point and actionable responses, helping users get the information they need faster and more efficiently.

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

General guidelines for writing instructions for generative AI large language models (LLMs)

When using Now Assist products and skills, you may have the option to give specific instructions or other guidance to the LLM. Writing generative AI instructions is different from conducting a keyword search. Use the following general guidelines when crafting your instructions.

Writing instructions for generative AI is very different from using search keywords. Keywords are the words that you might expect to appear in your results. For example, if you search for "gray bobtail cats," then you can reasonably expect your search results to return with topics or media that is about gray cats, bobtail cats, or even just cats in general. But with generative AI, you are asking the LLM to perform a task for you. The phrase "gray bobtail cats" does not include a verb to tell the LLM what to do. What about these gray cats? Should it locate all there is to know about them? Should it find gray bobtail cats to adopt? Should it create a picture of a gray bobtail cat? Should it be a realistic picture, or more of a line drawing? Generative AI needs more than just keywords.

General LLM instructions

Use instructions or questions to tell the LLM what you want. They can include four parts:

Goal

What kind of result do you want from the LLM?

Example: `I want a list of 3-5 bullet points to prepare me...`

Context

Why do you need it, and who is involved?

Example: `...for an upcoming meeting with [client], focusing on their current state and what they're looking to achieve with their "Phase 3+" brand campaign.`

Expectations

How should the LLM best fulfill your request?

Example: `Please use simple language so I can get up to speed quickly.`

Source

What information or other resources would you like the LLM to use?

Example: `Focus on email and Teams chats with [people] since June.`

Sample LLM instructions



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

Continually test and refine your instructions. Creating good LLM instructions is an iterative process, and as the LLM model learns, you may want to modify your instructions over time.

Stating your goal

When constructing an LLM description or instruction, consider these basic guidelines.

- Lead with action verbs. Use the imperative form or direct commands.
- Be direct and use simple sentences rather than complex ones.
- Be specific.
- Don't use jargon or slang terms.
- Avoid references to third parties or pronouns. Removing the subject or any identifiers generally prevents the LLM from personifying or otherwise misidentifying the end user.
- Your words instruct the logic that generative AI will use. Detailed, chain-of-thought instructions work well for this.

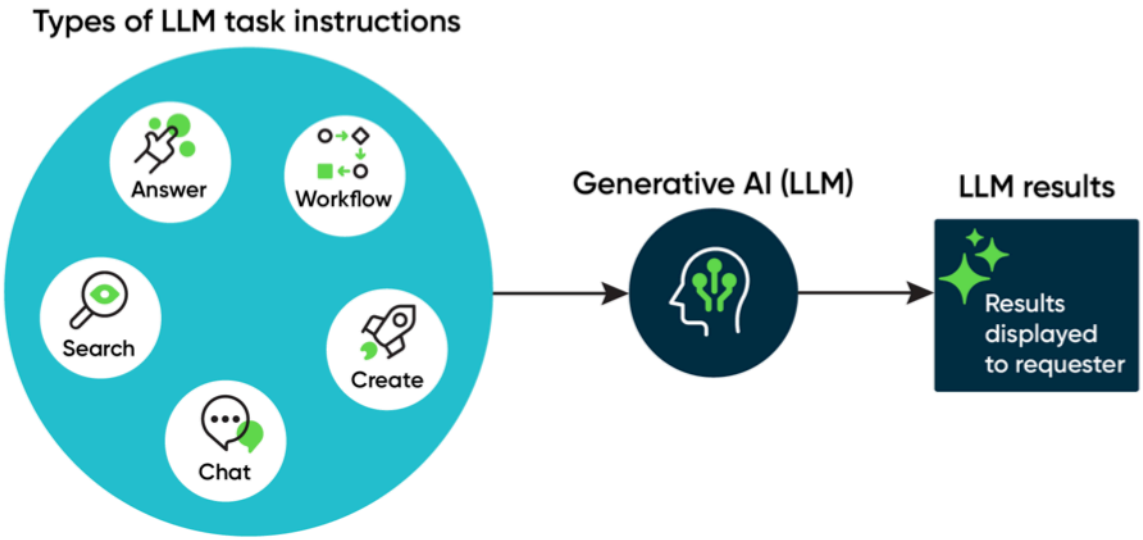
Instructions should also be tailored to the type of task. The following table describes the different kinds of tasks and the sort of instructions you might write for each circumstance.

Types of generative AI tasks and example instructions

Task type	Description	Example instruction
Simple search	Simple search for an answer.	When is the next company holiday?
Answer	Gather information from multiple sources and provide a summarized answer.	What were the major customer support issues in the past 30 days?
Chat	A back-and-forth conversation in which the LLM is getting additional information from the requester.	I have a new phone and now I can’t access Okta.
Create	Create a new ServiceNow component.	Write a new KB on common reasons for slow query execution and how to fix it, based on problems created in the last 12 months.
Workflow	Leverage existing workflows and create conversations from them.	Reset my Okta password.

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

LLM instruction workflow



Providing context in your instructions

Providing context to the LLM may feel like you are stating the obvious. For example, you may need to explain why your user would want to perform the task, or explain more about what the task is about. If you're using language that could have alternative meanings, you may want to define your terms. For example, if your instructions are about Microsoft Teams, you may need to say something like, "Teams refers to Microsoft Teams, an application that employees in a company can use to communicate with one another individually or in groups."

When providing context, think about the target audience for the task. This will help you to write better instructions. Mentioning whom the task is for also helps the LLM carry out the task.

Types of users to consider:

Admins

Admins configure skills in the Now Assist Admin console. They work with platform owners and product owners for tasks and requirements. Subject-matter experts check the accuracy of generative AI results. Governance boards or committees may oversee final sign-off on the skill.

Builders

Builders create assets such as applications and workflows. Their skill level may vary from no-code, low-code, mid-skill, or high-skill. They mostly interact with each other and admins.

Agents

Agents provide technical assistance, customer support, or other problem-solving help for users of a product, service, or organization. Agents work with their peers and support many kinds of requesters.

Requesters

Requesters may include your organization's employees, partners, or customers. Requesters encounter generative AI mostly in a self-service context. Generative AI provides the opportunity to requesters to solve the problem themselves before they resort to asking an agent for help.

In addition, you may need to "instruct" the LLM how to arrive at the kind of answer you want. Detailed chain-of-thought instructions are more likely to produce the results you're looking for. If you make assumptions in your instructions, your results may vary wildly.

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

Example of chain-of-thought instructions

Standard instructions	Chain-of-thought instructions
Q: Start with 5 tennis balls. Buy 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls are there now? A: The answer is 11. Q: The cafeteria had 23 apples. If the cafeteria used 20 to make lunch and bought 6 more, how many apples does the cafeteria have?	Q: Start with 5 tennis balls. Buy 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls are there now? A: There are 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. 5+6=11. The answer is 11. Q: The cafeteria had 23 apples. If the cafeteria used 20 to make lunch and bought 6 more, how many apples does the cafeteria have?
LLM results (incorrect!)	LLM results (correct!)
A: The answer is 27.	A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had 23-20=3. They bought 6 more apples, so they have 3+6=9. The answer is 9.

Expectations

You can specify how you want your results to sound and look. For example, you can tell the LLM what kind of tone you want it to use.

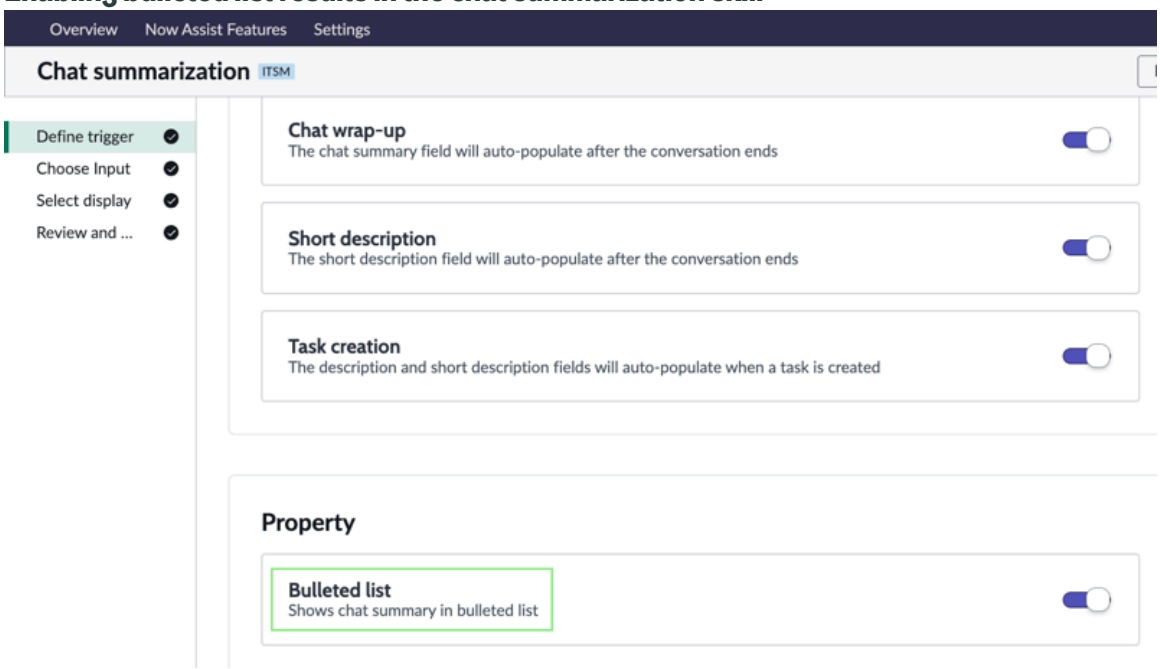
Modifying the tone of LLM results

Instruction	LLM result	Evaluation of result
Greet the user with a casual, laid-back tone.	Hey there! (: Welcome to May 22, 2024, a Wednesday. Chill and have a great day!	Too casual!
Greet the user with a professional, business-like tone.	Good day, welcome to our business. How may I assist you today?	A little too formal.
Greet the user with a professional, but friendly tone.	Hello there! I'm glad to assist you today, May 22, 2024. How can I help you?	Great, but I don't want the date.
Greet the user with a professional, but friendly tone. Do not include the date.	Hello! I'm here to help. How can I assist you today?	Perfect.
Instruction	LLM result	Evaluation of result

In some Now Assist applications, you may have the option to specify formatting. In those cases, you can tell the LLM to provide answers in a bulleted list, for example. Bulleted lists are often easier to read.

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

Enabling bulleted list results in the chat summarization skill



Source

You can suggest a variety of sources that the LLM should use to find answers, including Microsoft Teams conversations, Microsoft SharePoint Online sites, incidents and cases, and internal knowledge articles. In order for an agent or skill to access all of these sources, an admin must configure access on the instance.

For developers, the sources that a skill or AI agent can access may vary depending on the desired outcome. For example, the incident summarization skill uses the Incident table as its source. AI agents use different tools and knowledge sources, customized for the task they perform.

For more information about configuring sources for skill or AI agent use, see the following topic areas:

- AI Search
- Now Assist in AI Search
- External Content Connectors
- Overview tab in Now Assist Admin
- Now Assist AI agents

Additional guidelines for Now Assist skills and tools

Resources for writing LLM instructions for Now Assist skills

Skill	Reference
App generation	General guidelines for using app generation
Analytics generation	Guidelines and example questions
Catalog item generation	Suggestions to describe catalog items
Code generation	General guidelines for code generation

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

Resources for writing LLM instructions for Now Assist skills (continued)

Skill	Reference
Flow generation	
LLM topic skill for Virtual Agent	LLM description and instruction guidelines for Virtual Agent topics
Now Assist Skill Kit	General guidelines for Now Assist Skill Kit
RPA bot generation	General guidelines for RPA bot generation
Test generation	Design considerations for prompting
UI generation	General guidelines UI generation

Long term stable models

Long term stable (LTS) models support regulated industries, such as financial institutions, with stronger AI lifecycle management, governance, transparency, and compliance tools.

Long term stable models enable you to adopt AI confidently while supporting customers in their operational and compliance needs. These models also integrate with tools like AI Control Tower to support customers with governance, monitoring, and compliance controls along with guidance for implementing AI responsibly.

LTS models provide predictable AI updates, longer model stability windows, and early access testing for smoother adoption.

You can configure long terms stable models in AI Control Tower. You can also navigate to **Manage model providers** section in the Now Assist settings and select Now LLM Service-LTS as a provider for specific Now Assist skills.

Discrepancies when using different AI search tools

Different AI search tools may return different answers for the same or similar searches. This difference in results is expected. It occurs because each large language model (LLM) uses a different approach to find results and generate answers that match your search.

Overview of how LLMs create responses

Large Language Models don't retrieve answers from a single source. Instead, they generate responses by predicting the next most likely word, one at a time, based on patterns they've learned from the data they were trained on.

Because of that, several factors can make the answers vary each time. The factors include:

- Probability
- Multiple good answers
- Context sensitivity
- LLM variation

Answer generation involves probability, context, and variation. Therefore, responses to the same question can differ slightly each time it is asked, and from provider to provider.

For more information on differences in ServiceNow AI results, see Search result disparities between AI Search and Now Assist search features.

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

Example: Now LLM Service response vs. Azure response

The difference in responses by Now LLM Service and Azure for information about tuition reimbursement

Providers and Models

You can bring your own large language model (LLM) provider and API to use with Now Assist.

Bringing your own large language model enables:

- Data Privacy & Compliance: Keep sensitive data within your own environment and meet regulatory requirements.
- Customization: Fine-tune models for domain-specific needs and control behavior.
- Integration Flexibility: Connect easily with existing infrastructure and preferred cloud providers.
- Cost Control: Use existing contracts and optimize resource usage to avoid extra charges.
- Strategic Autonomy: Maintain vendor neutrality and avoid lock-in for future-proofing.

Now Assist

ServiceNow® Now Assist uses generative AI to enhance user productivity and efficiency through conversation and proactive experiences.

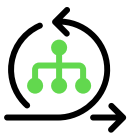
Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows and AI agents. For more information see ServiceNow product tiers.

See https://www.servicenow.com/community/now-assist-articles/now-assist-quick-start-guide/ta-p/2685044 to read about the foundation of generative AI at ServiceNow®.

Get started

Explore	Configure	Analyze
Learn more about Now Assist on the ServiceNow AI Platform.	Explore, configure, use, and analyze generative artificial intelligence (AI) features and skills.	Analyze and monitor Now Assist performance.

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

Now LLM Service model updates	Reference	Generative AI Controller
		
Learn about the Now LLM Service models.	Learn about user roles, data usage policy, and domain separation in Now Assist.	Customize your generative AI experience on the ServiceNow AI Platform with the Generative AI Controller.

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.

Now Assist products and skills

The Now Assist framework is supported across the ServiceNow AI Platform. To use Now Assist skills, activate one or more of the following Now Assist products.

Now Assist products

Workflow	Business areas	Available products
Technology	The Technology workflow includes IT applications, such as IT services and operations, managing your strategy to deliver products and services, and platform security.	• Now Assist for Collaborative Work Management (CWM) • Now Assist for Configuration Management Database (CMDB) • Now Assist for Enterprise Architecture (EA) • Now Assist for Operational Sustainability (formerly ESG) • Now Assist for Hardware Asset Management (HAM)

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

Now Assist products (continued)

Workflow	Business areas	Available products
		• Now Assist for Integrated Risk Management (IRM) • Now Assist for IT Operations Management (ITOM) • Now Assist for IT Service Management (ITSM) • Now Assist for Operational Technology Manager (OTM) • Now Assist for Operational Technology Service Management (OTSM) • Now Assist for Security Incident Response • Now Assist for Software Asset Management (SAM) • Now Assist for Strategic Portfolio Management (SPM) • Now Assist for Third-party Risk Management (TPRM) • Now Assist for Vulnerability Response
Customer	The Customer workflow includes applications that support customer service, including field service, financial services, telecommunications and media, and the public service sector.	• Now Assist for Customer Service Management (CSM) • Now Assist for Field Service Management (FSM) • Now Assist for Financial Services Operations (FSO) • Now Assist for Order Management • Now Assist for Public Sector Digital Services (PSDS) • Now Assist for Sales Force Automation (SFA) • Now Assist for Telecommunications, Media and Technology (TMT)
Employee	The Employee workflow supports HR Service Delivery and Employee Experience features.	• Now Assist for Employee Experience • Now Assist for Health and Safety

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

Now Assist products (continued)

Workflow	Business areas	Available products
		• Now Assist for HR Service Delivery (HRSD) • Now Assist for Legal Service Delivery (LSD) • Now Assist for Workplace Service Delivery (WSD)
Creator	The Creator workflow supports a variety of Platform tools and builders, including the following: • App Engine Studio • ServiceNow AI Platform scripting • Platform Analytics • Service Catalog • Workflow Studio • RPA Hub • Process Mining	
Finance & Supply Chain	The Finance & Supply Chain workflow supports purchase requisitions, sourcing requests, and request for products or services.	• Now Assist for Accounts Payable Operations (APO) • Now Assist for Supplier Lifecycle Operations (SLO) • Now Assist for Sourcing and Procurement Operations (SPO)
App Engine	The App Engine workflow supports App Engine products and offerings so that you can enhance your custom applications with AI.	Now Assist for App Engine
Vault	The Vault workflow supports generating custom data patterns from text descriptions to streamline your workload, checking role access for an encrypted column to monitor your instance's encryption access posture, and scheduling Data Discovery jobs to detect sensitive data.	Now Assist for Vault
Other	The Other workflow accommodates additional plugins and skills that don't fit into any other workflow.	Now Assist for Zero Copy Connector

For a complete list of Now Assist skills and their workflows, see Now Assist skills.

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

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 Admin

Learn how Now Assist Admin brings generative AI capabilities to the ServiceNow AI Platform. With Now Assist Admin, you can improve the productivity and efficiency in your organization, deliver better self-service, recommend actions and provide answers, and empower your users to search more effectively.

Now Assist Admin overview

Now Assist is a growing cross-platform family of generative AI features, which are tasks that a large language model (LLM) can perform. Generative AI features are based on the initial training and architecture.

A skill delivers a feature plus the use case to the user. An example of a skill is chat summarization within a customer workspace. Now Assist products provide features and skills that are tailored to meet the needs of users in different workflows.

Now Assist framework

The Now Assist framework is supported across the ServiceNow AI Platform. To use Now Assist skills, activate one or more of the following Now Assist products.

Now Assist products

Workflow	Business areas	Available products
Technology	The Technology workflow includes IT applications, such as IT services and operations, managing your strategy to deliver products and services, and platform security.	• Now Assist for Collaborative Work Management (CWM) • Now Assist for Configuration Management Database (CMDB) • Now Assist for Enterprise Architecture (EA) • Now Assist for Operational Sustainability (formerly ESG) • Now Assist for Hardware Asset Management (HAM)

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

Now Assist products (continued)

Workflow	Business areas	Available products
		• Now Assist for Integrated Risk Management (IRM) • Now Assist for IT Operations Management (ITOM) • Now Assist for IT Service Management (ITSM) • Now Assist for Operational Technology Manager (OTM) • Now Assist for Operational Technology Service Management (OTSM) • Now Assist for Security Incident Response • Now Assist for Software Asset Management (SAM) • Now Assist for Strategic Portfolio Management (SPM) • Now Assist for Third-party Risk Management (TPRM) • Now Assist for Vulnerability Response
Customer	The Customer workflow includes applications that support customer service, including field service, financial services, telecommunications and media, and the public service sector.	• Now Assist for Customer Service Management (CSM) • Now Assist for Field Service Management (FSM) • Now Assist for Financial Services Operations (FSO) • Now Assist for Order Management • Now Assist for Public Sector Digital Services (PSDS) • Now Assist for Sales Force Automation (SFA) • Now Assist for Telecommunications, Media and Technology (TMT)
Employee	The Employee workflow supports HR Service Delivery and Employee Experience features.	• Now Assist for Employee Experience • Now Assist for Health and Safety

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

Now Assist products (continued)

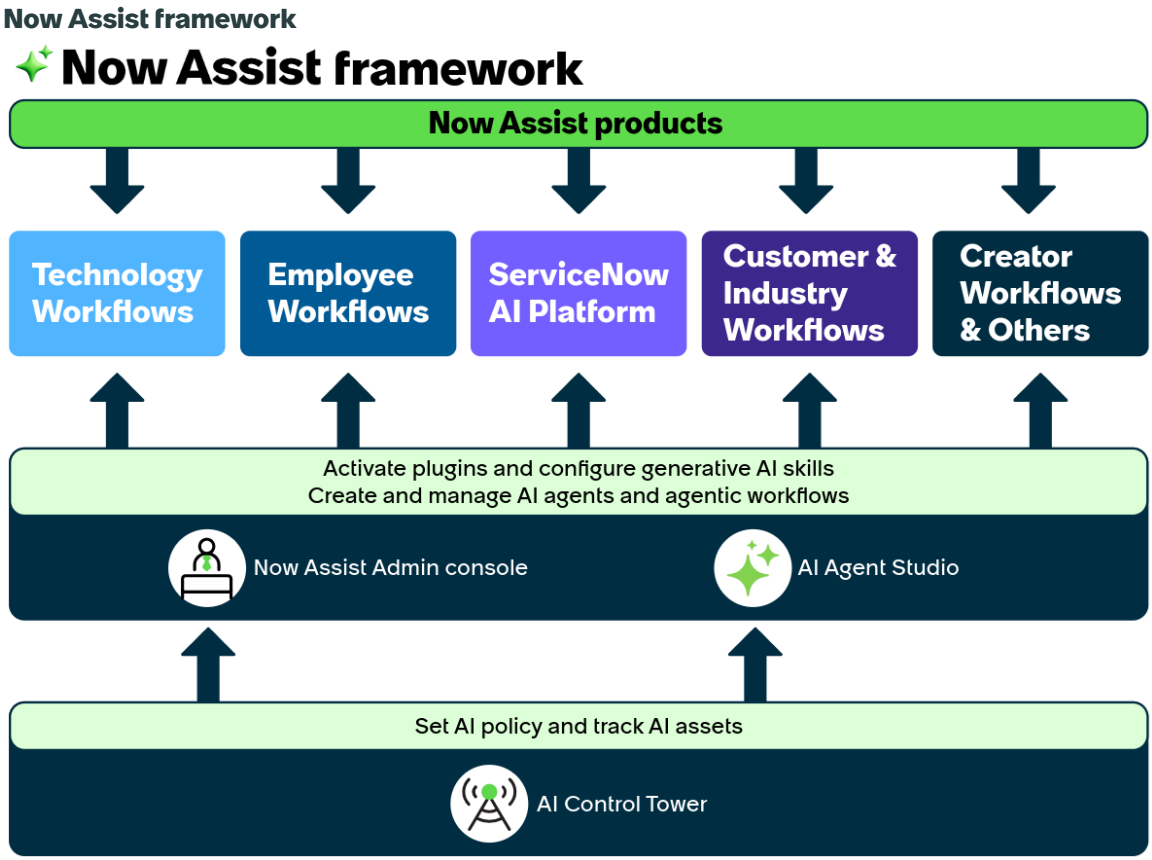
Workflow	Business areas	Available products
		• Now Assist for HR Service Delivery (HRSD) • Now Assist for Legal Service Delivery (LSD) • Now Assist for Workplace Service Delivery (WSD)
Creator	The Creator workflow supports a variety of Platform tools and builders, including the following: • App Engine Studio • ServiceNow AI Platform scripting • Platform Analytics • Service Catalog • Workflow Studio • RPA Hub • Process Mining	
Finance & Supply Chain	The Finance & Supply Chain workflow supports purchase requisitions, sourcing requests, and request for products or services.	• Now Assist for Accounts Payable Operations (APO) • Now Assist for Supplier Lifecycle Operations (SLO) • Now Assist for Sourcing and Procurement Operations (SPO)
App Engine	The App Engine workflow supports App Engine products and offerings so that you can enhance your custom applications with AI.	Now Assist for App Engine
Vault	The Vault workflow supports generating custom data patterns from text descriptions to streamline your workload, checking role access for an encrypted column to monitor your instance's encryption access posture, and scheduling Data Discovery jobs to detect sensitive data.	Now Assist for Vault
Other	The Other workflow accommodates additional plugins and skills that don't fit into any other workflow.	Now Assist for Zero Copy Connector

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

Now Assist products include some or all of the following foundational platform tools for Now Assist. For more information, see Now Assist skills in the Platform workflow.

- Administrators install plugins, manage skills, and analyze usage and performance with the Overview tab in Now Assist Admin.
- Users can take advantage of Now Assist skills by using the Now Assist panel on the instance.
- Use Now Assist in AI Search to generate answers for AI Search.
- Use Now Assist for Mobile to run generative AI skills in a mobile environment.
- Use Now Assist in Virtual Agent to create conversational catalog experiences and author topics that use LLM topic discovery.
- Developers can use the Generative AI Controller to integrate generative AI features in custom flows and conversations by using your own third-party large language model (LLM) licenses.

The following diagram shows what's available in the Now Assist framework.



Now Assist benefits

Benefit	Feature	Users
Leverage the power of search with the Now LLM generative AI model to answer questions in user searches with actionable AI-generated summaries of relevant knowledge articles.	Now Assist in AI Search	Everyone
Install and configure Now Assist applications and the skills they provide.	Now Assist Admin console	Administrators

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

Benefit	Feature	Users
Choose which skills to turn on, and which users can access them.	Now Assist Admin console	Administrators
Monitor the usage and performance of generative AI features and capabilities offered under Now Assist.	Now Assist Analytics	Administrators
Access generative AI skills in context through a user-friendly interface.	Now Assist panel	Everyone
Use Now Assist skills on mobile devices.	Now Assist for Mobile	Everyone
Customize your workflows and use your own third-party LLM license.	Generative AI Controller	Administrators or developers
Use Now Assist in other platform features.	Now Assist skills in the Platform workflow	Administrators or developers
Monitor Now Assist consumption on your instance.	Monitoring Now Assist usage in Subscription Management	Administrators

Overview tab in Now Assist Admin

The Now Assist Admin console provides quick and effortless access to the important information that you need to set up, configure, and monitor Now Assist applications and features.

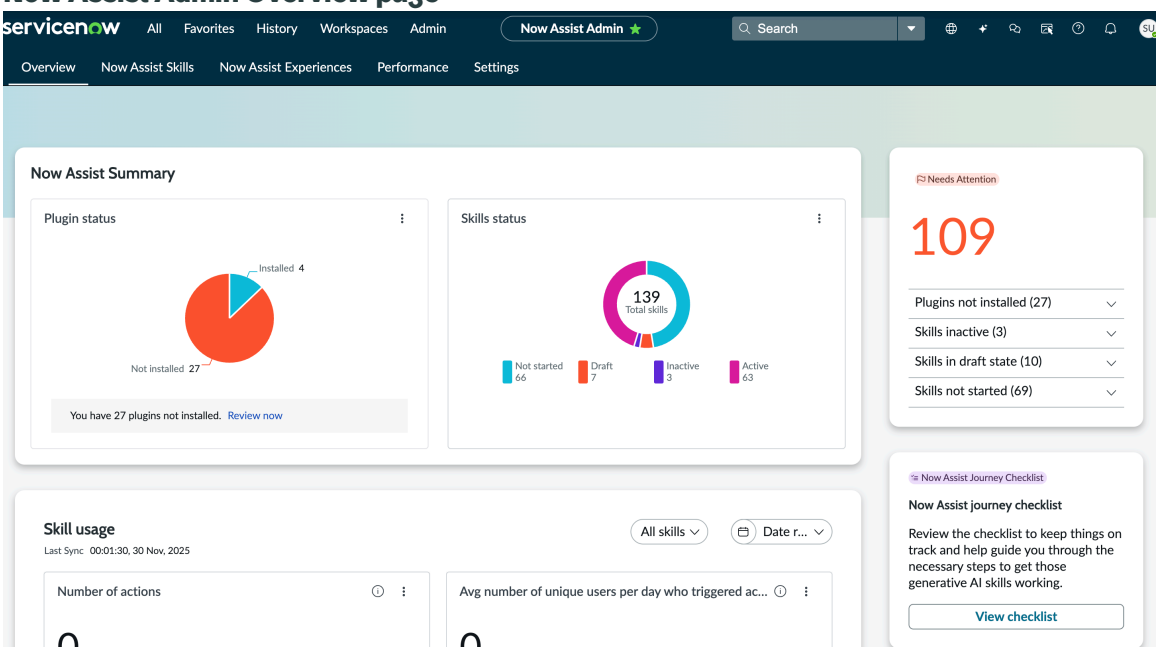
x Now Assist Admin overview

Now Assist Admin overview tab

Begin your exploration of the Now Assist features and skills in the Now Assist Admin console. This console contains everything that you need to install, configure, and learn about the different generative AI features on the ServiceNow AI Platform.

The following example shows the Now Assist Admin Overview page.

Now Assist Admin Overview page



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

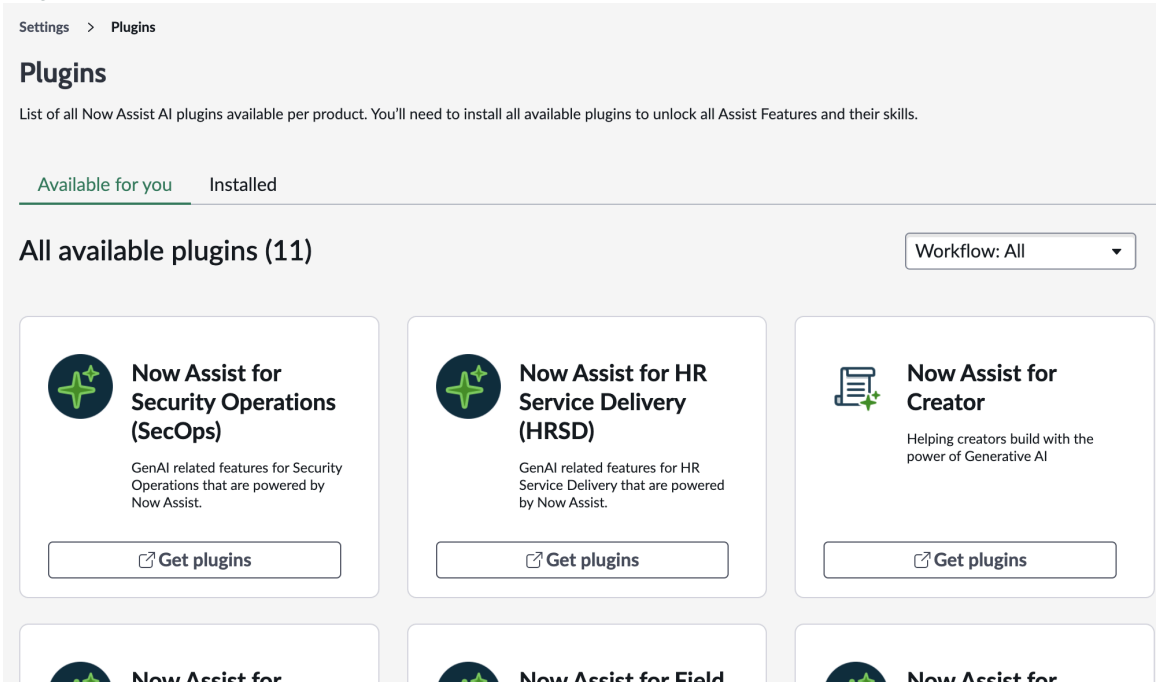
Now Assist Admin workflow

Take five steps to begin using the Now Assist Admin console.

1. **Install plugins.**

   On the **Available for you** tab of the Settings page, you can review the available plugins and install the ones that are relevant to your business needs. Each plugin contains the skills that you can activate to enable generative AI features on your instance. The following diagram shows the **Available for you** tab.

   Available for you tab on the Now Assist Admin Settings > Plugins page



2. **(Optional) Turn on the Now Assist panel.**

   The Now Assist panel integrates the Now Assist skills into the Next Experience UI. By turning on the Now Assist panel directly from the Now Assist Admin console, you enable agents to access skills from anywhere on the ServiceNow AI Platform.

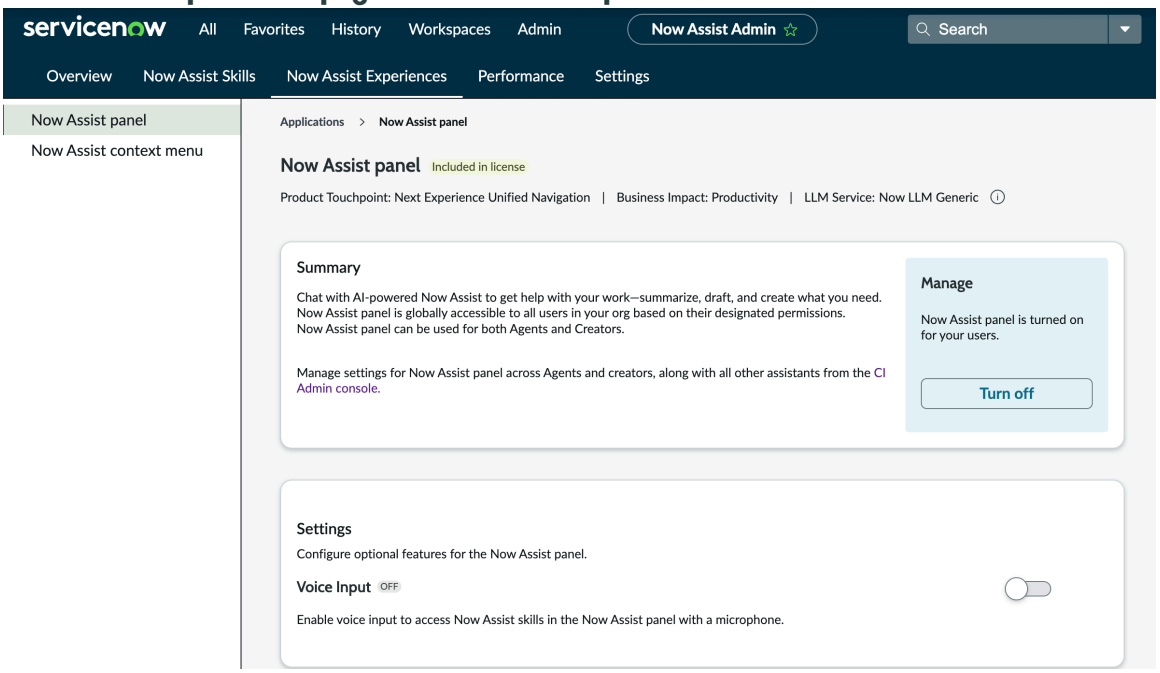
   For more information about the Now Assist panel, see Now Assist panel overview.

   This step is optional because the skills can also display in-product in the Core UI and in Workspaces.

   You can turn the Now Assist panel on from the Now Assist Experiences page of the Admin console.

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

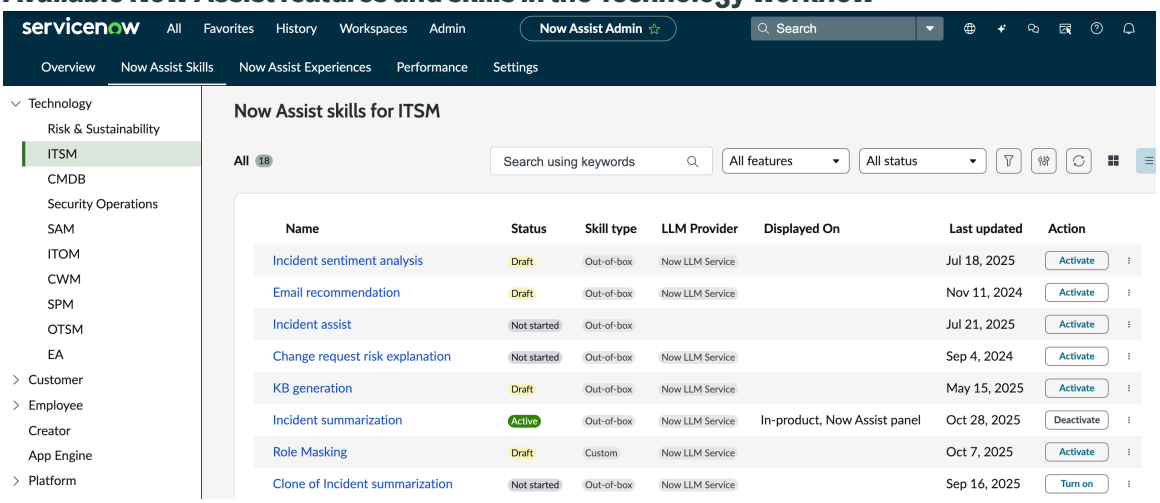
Now Assist Experiences page with Now Assist panel



3. Activate skills.

Skills are features that are created for a specific use case in a Now Assist application. Use the Now Assist Skills page to explore the skills that are available with your installed plugins. The following diagram shows the available features and skills in the Technology workflow.

Available Now Assist features and skills in the Technology workflow

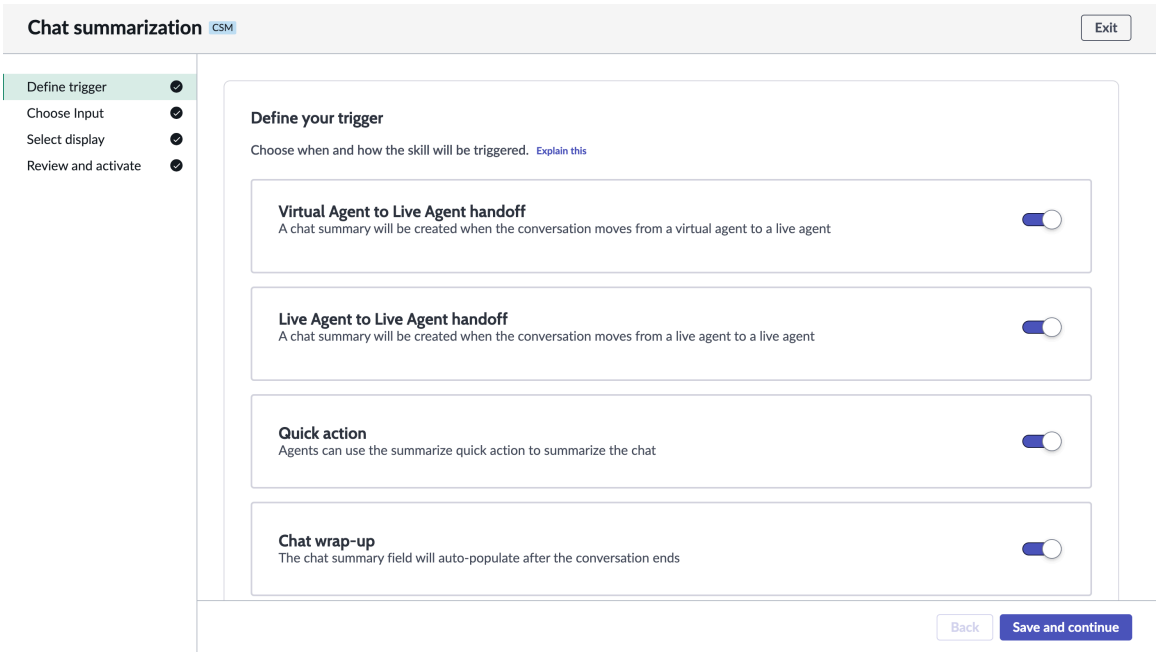


After deciding which skills best fit your business needs, you can activate them from the console. Some skills require configuration so that you can customize the skill to your needs, such as determining the skill inputs and triggers. You can select the skills that you want to configure in the Now Assist Admin Features page.

The following example shows a step in the guided setup of the skill activation process, which is choosing where to display the skill. Skills can be displayed in-product, including the Core UI and Workspaces, or in the Now Assist panel. You can choose either or both locations. The following diagram shows the process for chat summarization activation in Now Assist Admin.

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

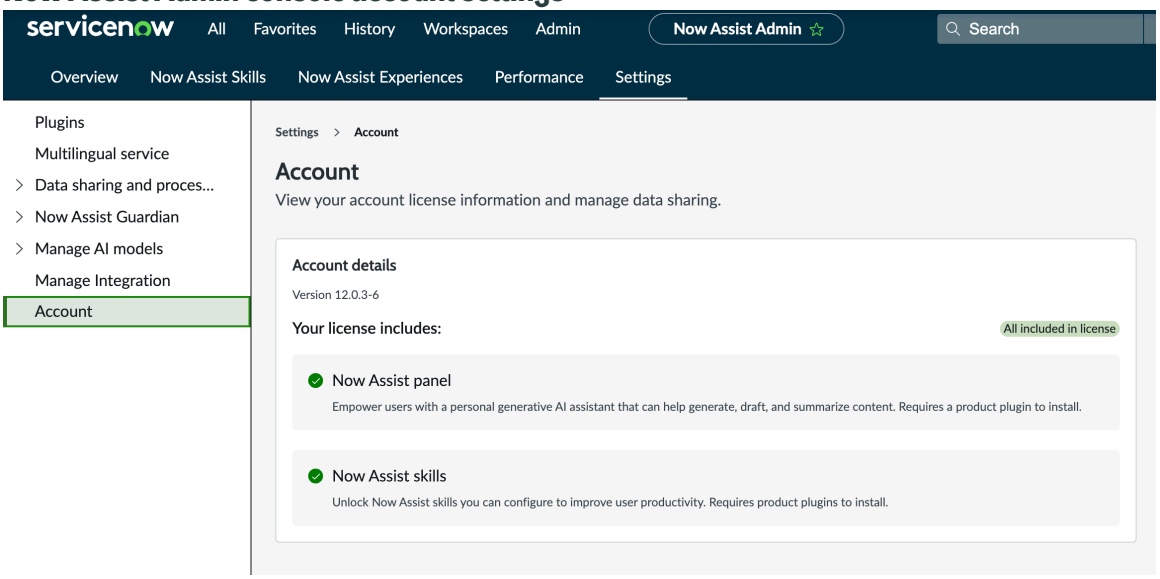
Chat summarization activation in Now Assist Admin



4. Review your Now Assist account settings.

The Now Assist Admin console Settings page enables you to set up language support, if you have Dynamic Translation enabled on your instance, and review your account details. Get up-to-date information about what plugins are available to you and the status of data sharing on your instance.

Now Assist Admin console account settings

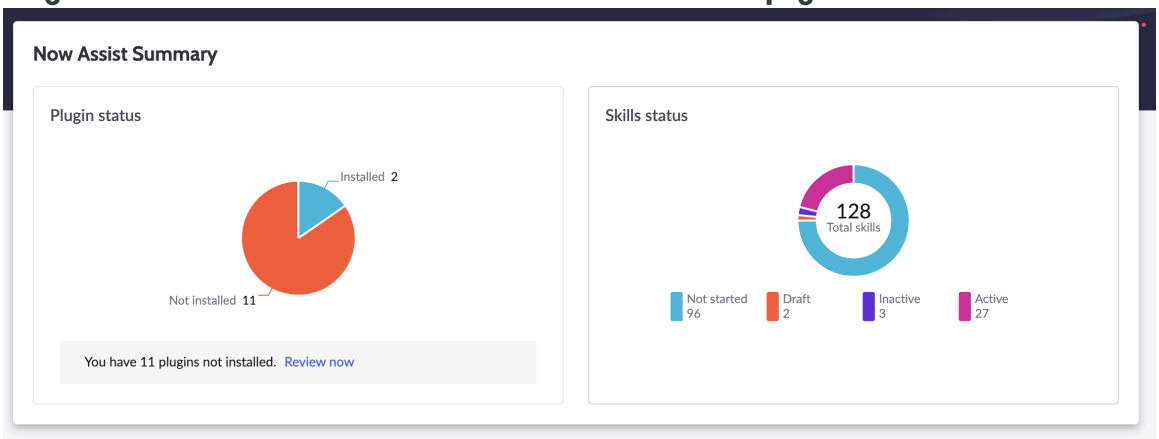


5. Monitor and analyze skill performance.

Use the metrics available on the Overview page to review the summaries, performance information, and issues that need your attention. The following example shows the current plugin status, as well as the number of active skills.

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

Plugin and Skills status on the Now Assist Admin Overview page



The Now Assist Admin console contains the **Now Assist journey checklist** with additional instructions for implementing Now Assist on your instance.

Now Assist journey checklist

You can consult the following checklist from the Now Assist Admin console Overview page at any time to guide your implementation of Now Assist applications, features, and skills.

Now Assist checklist

The following steps make up your Now Assist journey and will allow you to get the most out of generative AI skills.

1. **Install Now Assist plugins**
   Visit the store to install the plugin specific to the product workflows. Browse all available plugins.
2. **Turn on Now Assist panel**
   After a plugin is installed, return to this console to turn on the Now Assist panel which will allow the skills associated with those plugins to work.
3. **Activate Now Assist skills**
   Once the panel is turned on, configure and customize information sources that define how the skills work.
4. **Review account settings**
   Review account information and manage data sharing.
5. **Analyze Now Assist skills**
   Track and monitor the progress of your Now Assist skills.

Done

Learn more about the Now Assist journey

- Get an overview of the Now Assist framework.
- Discover the Now Assist panel.
- Explore the Now Assist skills available.
- Activate and configure a Now Assist skill using the Now Assist Admin console.
- Analyze and monitor Now Assist skill usage and performance.

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

Now Assist Experiences

Explore Now Assist panel and Now Assist context menu under the *Now Assist* **Experiences** tab of Now Assist Admin.

Now Assist panel

With the Now Assist panel, you can get assistance from generative AI experiences to solve customer issues faster. Use this conversational interface to summarize a chat, case, or incident, get help, or generate resolution notes so that you can get the context of this information more quickly.

You must fulfill the following conditions to use the Now Assist panel:

- Next Experience must be enabled. For more information, see Considerations for activating Next Experience.
- You must have the now_assist_panel_user role.
- Your role must be applied to at least one active Now Assist skill to use skills in the Now Assist panel.

There are two versions of the Now Assist panel: standard chat and enhanced chat. Select one of the versions to learn more:

- Standard chat - provides conversational support assistance within a standard, static chat window. To activate the Now Assist panel standard chat, see Activate the Now Assist panel standard chat.
- Enhanced chat - provides conversational support assistance within a dynamic window that gives a more intuitive and personalized experience. You must activate the Now Assist panel enhanced chat and configure your Now Assist instance to use the enhanced chat. See Activate Now Assist panel enhanced chat for more information.

If you want to use assistants, you must activate them. See Activate Now Assist panel assistants for information on activating assistants.

Note: Voice input is automatically activated when you activate the Now Assist panel. As of the Zurich Patch 4 release, voice input is configured in the Configure Next Experience accessibility preferences.

Standard chat

With the Now Assist panel standard chat, you can get assistance from generative AI experiences to solve customer issues faster. Use this conversational interface to summarize a chat, case, or incident, get help, or generate resolution notes so that you can get the context of this information more quickly.

Note: Next Experience must be enabled to use the Now Assist panel. For more information, see Considerations for activating Next Experience.

Now Assist panel overview

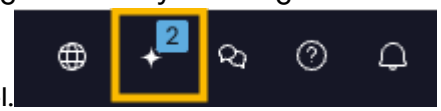
Agents can use the Now Assist panel to interact with and get assistance from generative AI. On the Now Assist panel, you can increase your productivity and efficiency by using the generative AI experience to summarize a chat, case, incident, or generate resolution notes.

Conversational aspects of the Now Assist panel, such as skill detection, are powered by Now LLM Service.

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

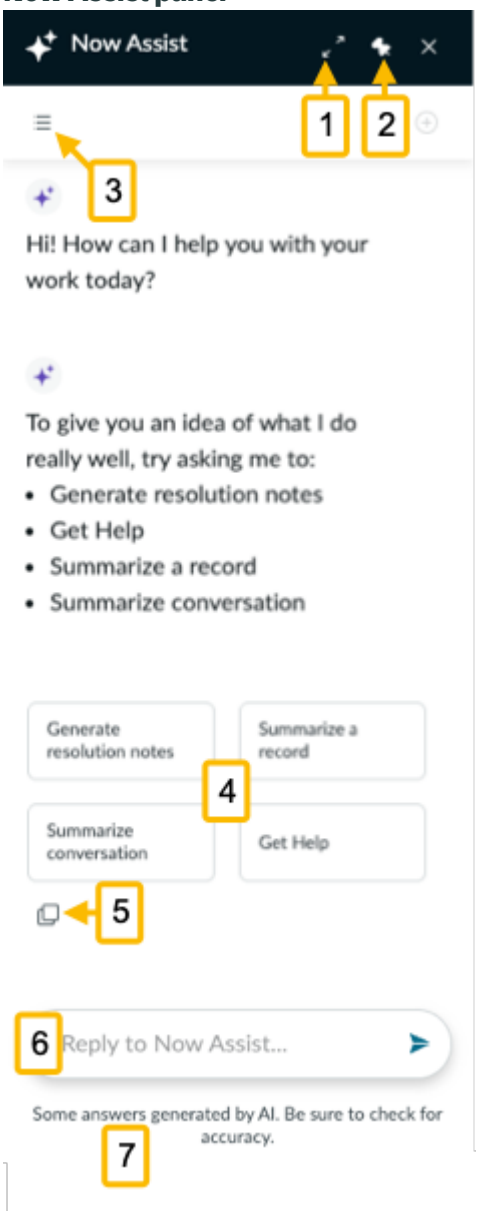
Note: Now Assist skills must be enabled to appear on the Now Assist panel. For more information, see Now Assist skills.

Let's get started by selecting the Now Assist icon to display the Now Assist panel.

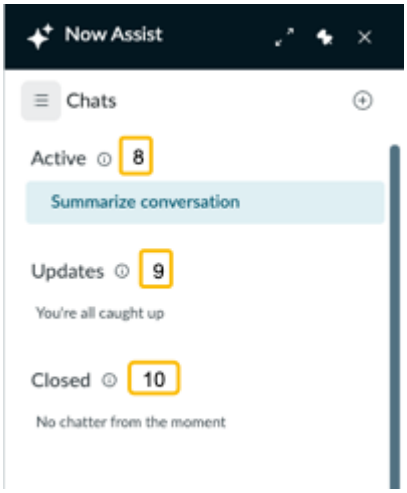


If a number in a square appears, it indicates how many messages you missed when the Now Assist panel was closed.

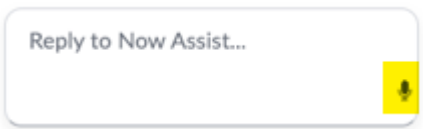
Now Assist panel



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



The Now Assist panel includes:

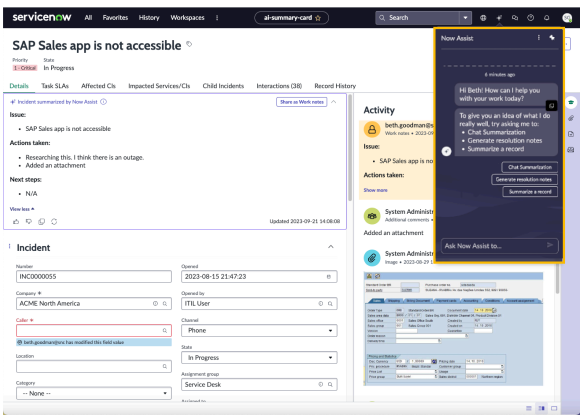
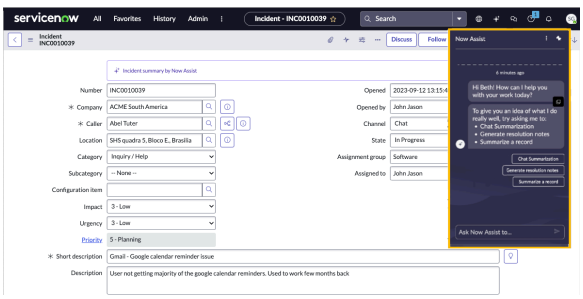
Item number	Description
1 - Enter modal button	Expands the chat into a 90% screen-size window.
2 -	Positions, or pins, the Now Assist panel to the screen.
3 - Display chats button	Displays the Chats window.
4 - Option buttons	Displays the available options.
5 - Copy message button	Copies the Now Assist panel message.
6 - Reply to Now Assist... field	Enter actions.
Voice Input 	If Voice Input is activated, select the microphone icon or the keyboard shortcut to use your voice to interact with the Now Assist panel. After you speak, there's a pause while the system transcribes the text and then displays it on the screen. See Enable voice input for Now Assist panel for information on enabling Voice Input. See Next Experience keyboard shortcuts for the Now Assist menu (Voice Input mode) shortcuts for Microsoft and macOS.
8 - Active chats	All active chats display in the Active chats section. in the Active chats section of the Now Assist. You can create additional chats by selecting the + icon in the heading.

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

Item number	Description
9 - Updates	Displays important updates and reminders.
10 - Closed chats	Displays all closed chats. If you select one of the closed chats, you can see that chat's history.

Now Assist panel is available on Next Experience and ServiceNow Studio. The following screenshots show the Now Assist panel in a workspace and on Core UI screens under Next Experience.

Now Assist panel

Next Experience	Core UI
	

Navigating from the Now Assist panel

Navigate from the Now Assist panel without leaving the current conversation by entering a navigation request in the **Ask Now Assist to...** field. If you enter "navigate me to active incidents," Now Assist displays a button that enables you to view the active incidents.

Chat summarization

Quickly learn the details of a chat by reading a chat summarization. The chat summarization gives you enough details about the chat so that your requester doesn't have to repeat the same information to you.

To generate a chat summarization from the Now Assist panel, select **Chat Summarization** or enter `summarize chat` in the **Ask Now Assist to** field.

Note: You can also generate a chat summarization by using the `/summarize` quick action in Agent Chat.

Case or incident summarization

Quickly learn the details of a case or incident by reading a case summarization. The summarization gives you enough details about the interaction so that your requester doesn't have to repeat the same information to you.

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

You can generate a case or incident summarization from the Now Assist panel for Now Assist for CSM, Now Assist for HRSD, or Now Assist for ITSM:

- For Now Assist for CSM, select **Summarize record** or enter `summarize a record` in the **Ask Now Assist to** field.
- For Now Assist for HRSD, select **Summarize record** or enter `summarize a record` in the **Ask Now Assist to** field.
- For Now Assist for ITSM, select **Summarize incident** or enter `summarize an incident` in the **Ask Now Assist to** field.

Conversation Help

Get specific and accurate answers to your queries by using the Get Help skill option on the Now Assist panel. This skill is available to everyone entitled to Now Assist capabilities.

For more information about the Now Assist Conversational Help skill which represents as Get Help on the Now Assist panel, see Now Assist Conversational Help.

Resolution notes generation

Quickly learn the details of how an interaction was resolved by generating and reading resolution notes.

To generate resolution notes from the Now Assist panel, select **Generate resolutions notes** or enter `generate resolutions notes` in the **Ask Now Assist to** field.

Streaming responses

After you enter a question or request on the Now Assist panel, Now Assist gathers information from Knowledge Base articles, external content, product documentation, catalog items, and workflows and combines them into a synthesized, comprehensive answer. Instead of waiting for the entire message to render, the synthesized response streams in real time and stops streaming after the entire message has been delivered. An animated sparkle icon (✨) appears while the response is generated and changes to the static sparkle icon after the response has fully loaded.

Enhanced chat

Now Assist panel enhanced chat is a conversational support experience within a dynamic window that also includes the ability to have multiple active conversations and superior search capabilities. Use Now Assist panel enhanced chat to improve your productivity and efficiency by leveraging generative AI to perform tasks such as summarize a chat, case, or incident, request help, generate resolution notes, among others.

Agents can use the Now Assist panel enhanced chat to interact with and get assistance from generative AI. You can move the chat window by selecting the header and dragging the chat window to the desired location. You can resize the chat window by using the window's edges to resize to your desired size or use the toolbar icons.

Conversational aspects of the Now Assist panel, such as skill detection, are powered by Now LLM Service.

Note: To use the full capabilities of Now Assist panel enhanced chat, AI Search must be enabled for your portal. Without it, Now Assist panel enhanced chat functions in a limited capacity. Basic conversational interactions such as predefined topic flows and simple questions and answers are available, but knowledge article retrieval, AI responses grounded in instance content, and semantic search capabilities require AI Search. For more information, see Enable and configure AI Search in Service Portal.

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

Prerequisites

You must fulfill the following conditions before using the Now Assist panel enhanced chat:

- Next Experience must be enabled. For more information, see Considerations for activating Next Experience.
- You must have the now_assist_panel_user role.
- Your role must be applied to at least one active Now Assist skill to use skills in the Now Assist panel.
- Now Assist panel enhanced chat must be activated before you can use it. See Activate Now Assist panel enhanced chat for more information.
- Now Assist skills must be enabled to appear on the Now Assist panel. For more information, see Now Assist skills.
- If you want to use assistants, you must activate them. See Activate Now Assist panel assistants for information on activating assistants.
- If you want to use voice input, you have to enable voice input Enable voice input for Now Assist panel.

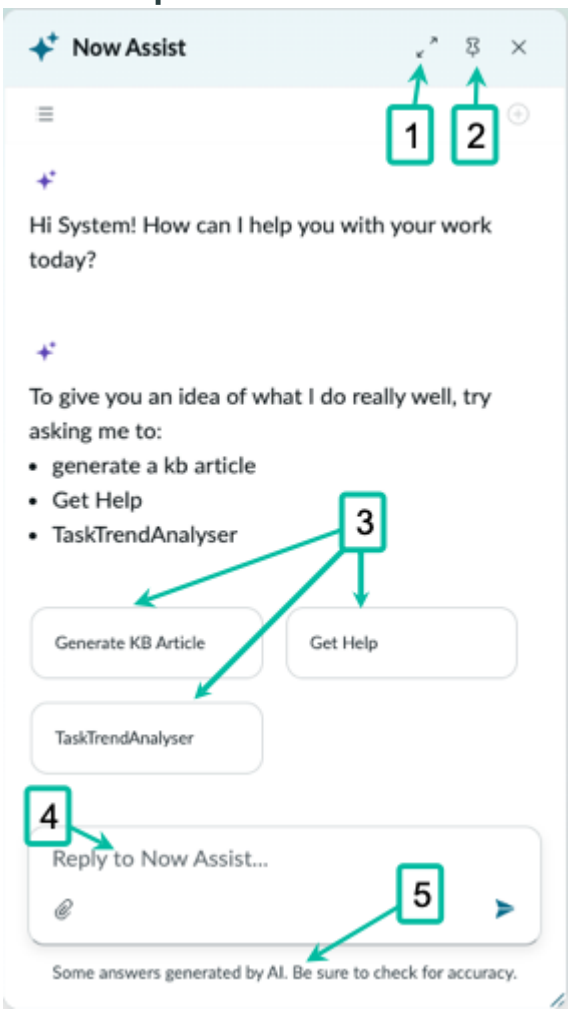
> **Note:** Voice input is automatically activated when you activate the Now Assist panel. As of the Zurich Patch 4 release, voice input is configured in the Configure Next Experience accessibility preferences.

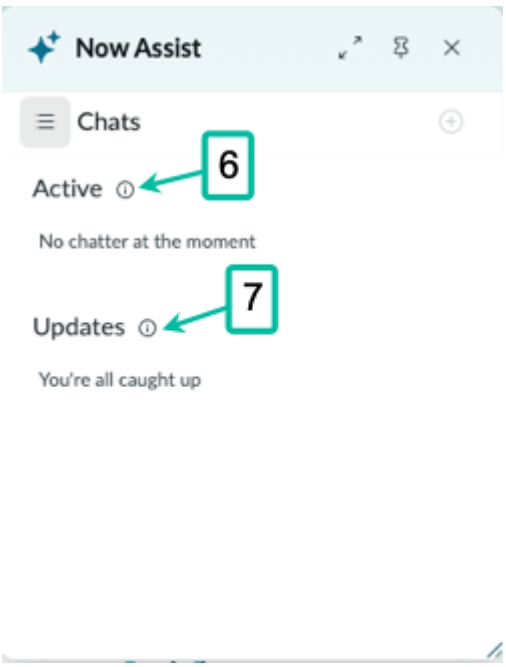
Now Assist panel enhanced chat overview

Let's get started by selecting the Now Assist icon to display the Now Assist panel.

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

Now Assist panel





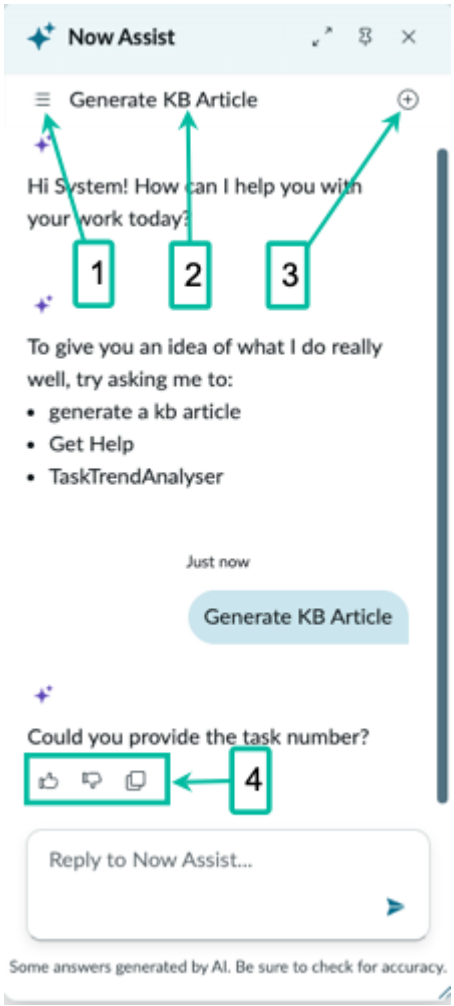
The Now Assist panel includes:

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

Item number	Description
1 - 	Expands the chat into a 90% screen-size window. The 90% screen-size window can't be resized or moved. Selecting the icon again resizes the chat back into the floating window.
2 - 	Positions, or pins, the Now Assist panel to the screen.
3 - Option buttons	Displays the available options.
4 - Reply to Now Assist...	Enter actions.
5 - Now Assist message	Indicates that the answers are generated by AI.
6 - Active chats	All active chats display in the Active chats section. in the Active chats section of the Now Assist. You can create additional chats by selecting the + icon in the heading.
7 - Closed chats	Displays all closed chats. If you select a closed chat, the chat's history displays.

The Now Assist subheader consists of these elements:

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



Now Assist subheader elements

Element	Description
1. Chats (☰)	All chats appear. Chats are organized with the most recent conversations at the top. Selecting a chat opens the chat in the conversation area. If there are unread chats or notifications, a badge number appears on the Chats icon (☰). Any unread chat or notification appears with a red dot next to it and the chat title appears in bold. Additionally, if you switch to a new chat while another active chat is ongoing, a pop-up message on the Chats icon (☰) appears: `Your previous chat was saved. You can revisit all of your past chats and continue ones that are still active.` The following list includes the chat sections that you may see in the chats area.

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

Now Assist subheader elements (continued)

Element	Description
	• Active: Chats where you can continue the conversation. If applicable, active chats move to the Closed chats section after two hours of inactivity. This 2 hour time limit can be configured within the Messaging Channels {sys_cs_channel.list} table. To change the inactivity time limit, from the Messaging Channels {sys_cs_channel.list} table, select the **NASS** record and populate the **Conversation Idle Timeout** field with your preferred active chat time limit. If you have no active chats, `No chatter at the moment` is displayed. If more than 12 active chats are running, a **Show more** link appears to view more chats. Selecting **Show more** displays an additional 10 chats. • Updates: Updates for important notifications and reminders can be configured to display. When *`sn_nowassist_va.nass_notification_enabled`* is set to `true` (the default), simple (nonactionable) Now Assist panel notifications appear on your portal home page. After selecting a notification, you're redirected to this Updates section. If you have no updates, `You're all caught up` is displayed. If more than four updates are available, a **Show more** link appears to view more updates. Selecting **Show more** displays an additional 10 updates. **Note:** If *`sn_nowassist_va.nass_notification_enabl`* is set to `false`, the Updates section doesn't appear. • Closed: Closed chats can be configured to display. A message closes when the designated time has passed (2 hours of inactivity) or you receive the following response in the chat: `It looks like you're finished with this chat, so I'll go ahead and close it.` Turn on closed chats by selecting the **Show closed chats** check box within **Conversational Interfaces** > **Assistants** > **[Selected Assistant Name]** > **Chat experience** > **Closed chats**. After being turned on, closed chats are displayed for as long as they're available within the Conversations (sys_cs_conversation) table. Closed chats appear in a read-only

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

Now Assist subheader elements (continued)

Element	Description
	mode and can't become active again. If more than four closed chats are available, a **Show more** link appears to view more closed chats. Selecting **Show more** displays an additional 10 closed chats. After a conversation has closed, you can't reopen it. Hovering over a closed chat displays the delete icon (). Confirm the chat deletion on the Delete this chat? modal to permanently delete the chat from the interface.
2. [Chat name]	The name of the conversation. If you select a promoted asset or query, that asset's title appears as the chat name. If instead you enter an utterance into the **Reply to Now Assist** field, your initial utterance becomes the chat name. The chat name appears in both the Now Assist subheader and **Chats list** > **Active** section.
3. New chat ( )	A new conversation begins. You may be prompted with a greeting message along with any promoted conversational assets such as topics, subflows, and/or actions, or suggested queries.
4. Feedback icons panel	You can indicate if the response was helpful by selecting the like thumbs up icon or if the response wasn't helpful, you can select the dislike thumbs down icon . You can also copy the response by selecting the copy message icon .

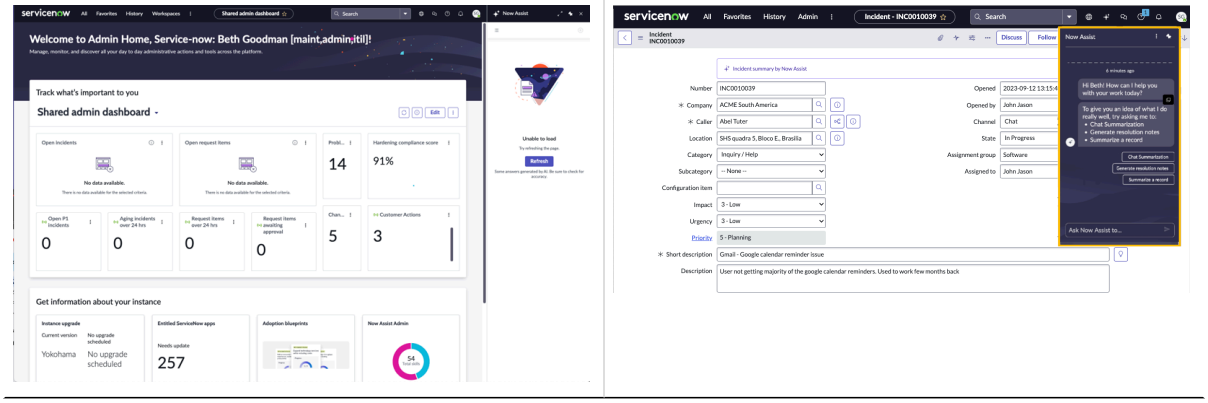
Now Assist panel is available on Next Experience and ServiceNow Studio. The following screenshots show the Now Assist panel in a workspace and on Core UI screens under Next Experience.

Now Assist panel

Next Experience	Core UI

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

Now Assist panel (continued)



Agentic conversations

Note: Admins must first enable AI agents before end users can experience agentic conversations. Now Assist panel discovers and executes agentic workflows. For more information on agentic workflows, see Now Assist agentic workflows and Multiple conversations in Now Assist AI agents.

When you ask a question to the Now Assist panel enhanced chat, the agent understands the query and begins a flow. When you submit a message with multiple questions or requests, Now Assist panel enhanced chat answers the multiple questions consecutively in its response. It can reason, plan, and execute across AI agents, Now Assist panel topics, conversational actions and subflows, catalogs, Knowledge Base articles, custom skills, and any Now Assist supported skills to help you. You receive on-screen messages to let you know where the agent is in the agentic processing flow prior to receiving the response. After the processing has completed its flow, a View AI Steps section header appears, where the processing flow steps can be expanded and viewed. You can stop the agentic processing flow at any time by selecting the End flow icon ( ). After an action starts, it can't be stopped. Selecting the End flow icon only stops the proceeding processing steps.

Navigating from the Now Assist panel

You can navigate from the Now Assist panel without leaving the current conversation by entering a navigation request in the **Ask Now Assist to...** field. If you enter "navigate me to active incidents," Now Assist displays a button that enables you to view the active incidents.

Chat summarization

Quickly learn the details of a chat by reading a chat summarization. The chat summarization gives you enough details about the chat so that your requester doesn't have to repeat the same information to you.

To generate a chat summarization from the Now Assist panel, select **Chat Summarization** or enter `summarize chat` in the **Ask Now Assist to** field.

Note: You can also generate a chat summarization by using the `/summarize` quick action in Agent Chat.

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

Case or incident summarization

Quickly learn the details of a case or incident by reading a case summarization. The summarization gives you enough details about the interaction so that your requester doesn't have to repeat the same information to you.

You can generate a case or incident summarization from the Now Assist panel for Now Assist for CSM, Now Assist for HRSD, or Now Assist for ITSM:

- For Now Assist for CSM, select **Summarize record** or enter `summarize a record` in the **Ask Now Assist to** field.
- For Now Assist for HRSD, select **Summarize record** or enter `summarize a record` in the **Ask Now Assist to** field.
- For Now Assist for ITSM, select **Summarize incident** or enter `summarize an incident` in the **Ask Now Assist to** field.

Conversation Help

Get specific and accurate answers to your queries by using the Get Help skill option on the Now Assist panel. This skill is available to everyone entitled to Now Assist capabilities.

For more information about the Now Assist Conversational Help skill that represents as Get Help on the Now Assist panel, see Now Assist Conversational Help.

Resolution notes generation

Quickly learn the details of how an interaction was resolved by generating and reading resolution notes.

To generate resolution notes from the Now Assist panel, select **Generate resolutions notes** or enter `generate resolutions notes` in the **Ask Now Assist to** field.

Streaming responses

After you enter a question or request on the Now Assist panel, Now Assist gathers information from Knowledge Base articles, external content, product documentation, catalog items, and workflows and combines them into a synthesized, comprehensive answer. Instead of waiting for the entire message to render, the synthesized response streams in real time and stops streaming after the entire message has been delivered. An animated sparkle icon (✨) appears while the response is generated and changes to the static sparkle icon after the response has fully loaded.

Fallback options

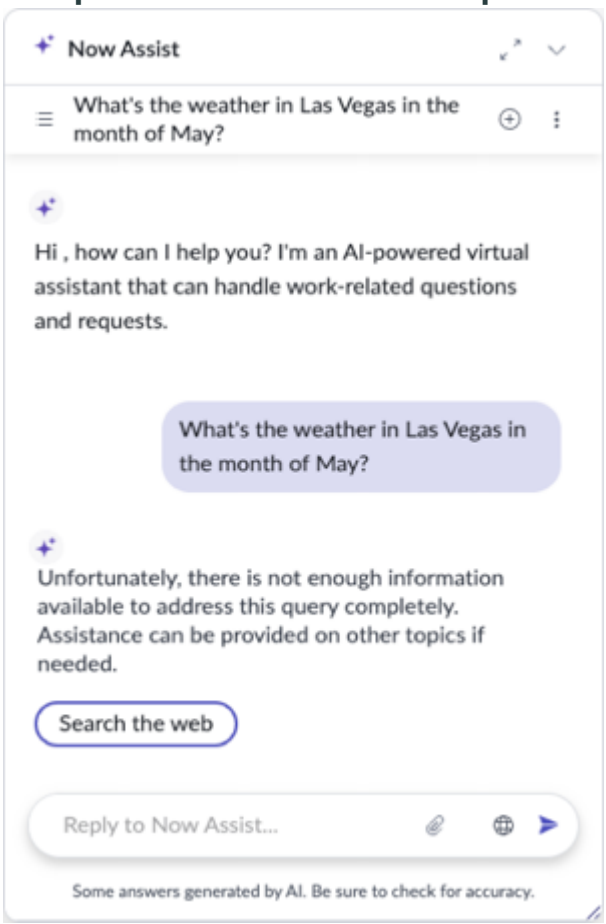
Note: For more information about where and how to enable fallback options, see Manage an assistant chat experience.

A fallback state can occur whenever search results are unavailable. Scenarios where search results are unavailable include when Now Assist didn't understand the query, complaint small talk was found, or an error occurred. When search results are unavailable, the **Search the web** fallback option may appear. If you select the **Search the web** fallback option, the web search mode is triggered and uses the internet to search for the results.

Note: Only the last query entered into the conversation is considered when entering web search mode via this **Search the web** fallback option.

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

Example of Search the web fallback option



Premium chat

Now Assist panel premium chat provides a full-featured conversational AI experience that combines generative AI chat skills with advanced capabilities to speed up issue resolution. Premium chat is the default experience for new installations and upgrades.

Note: Please note these important considerations:

- Next Experience must be enabled to use the Now Assist panel. For more information, see Considerations for activating Next Experience.
- To use the full capabilities of Now Assist panel premium chat, AI Search must be enabled for your portal. Without it, Now Assist panel enhanced chat functions in a limited capacity. Basic conversational interactions such as predefined topic flows and simple questions and answers are available, but knowledge article retrieval, AI responses grounded in instance content, and semantic search capabilities require AI Search. For more information, see Enable and configure AI Search in Service Portal.

Agents can use the Now Assist panel enhanced chat to interact with and get assistance from generative AI. You can move the chat window by selecting the header and dragging the chat window to the desired location. You can resize the chat window by using the window's edges to resize to your desired size or use the toolbar icons.

Now Assist panel enhanced chat must be activated before you can use it. See Activate Now Assist panel enhanced chat for more information.

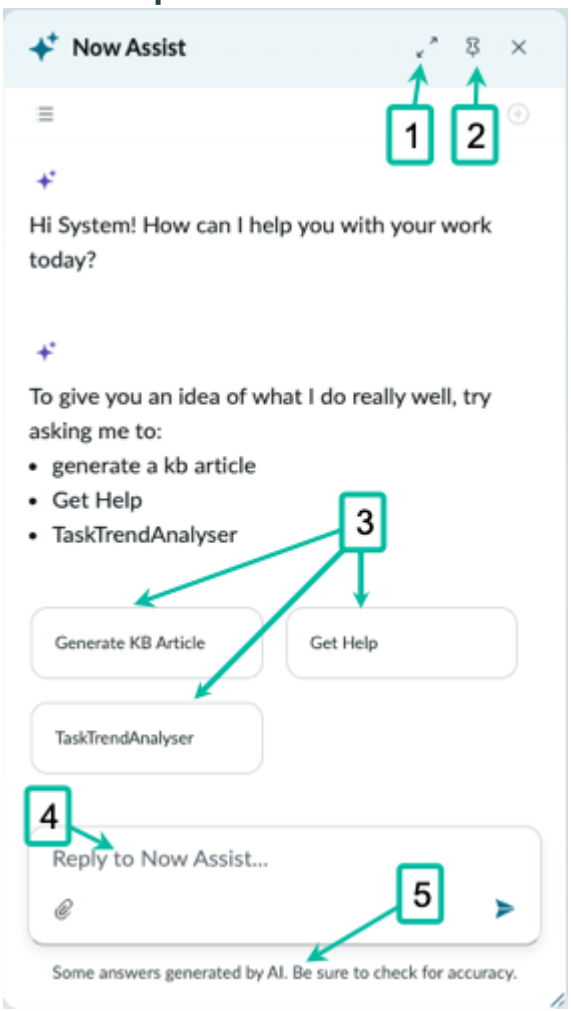
Conversational aspects of the Now Assist panel, such as skill detection, are powered by Now LLM Service.

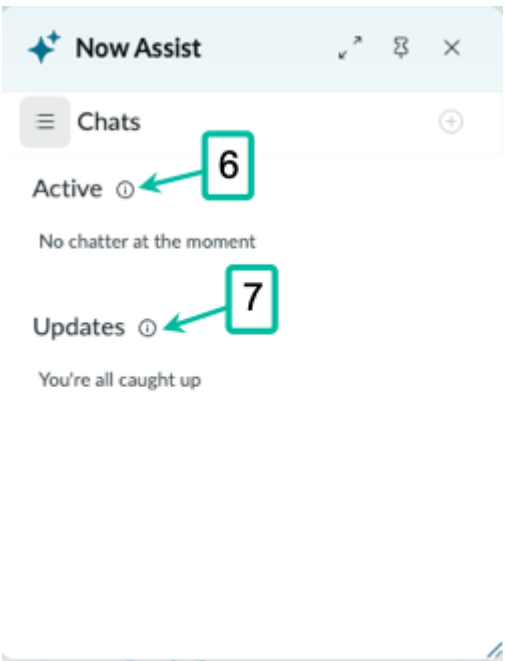
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

Note: Now Assist skills must be enabled to appear on the Now Assist panel. For more information, see Now Assist skills.

Let's get started by selecting the Now Assist icon to display the Now Assist panel.

Now Assist panel





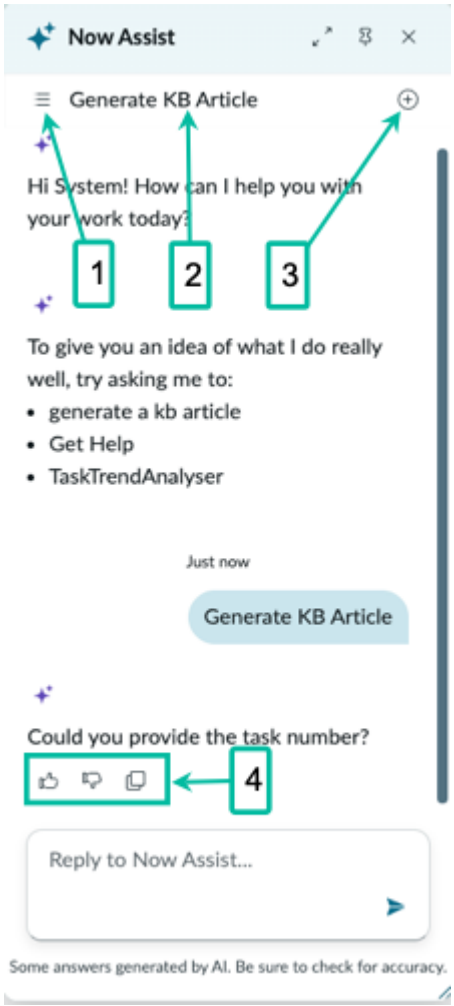
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

The Now Assist panel includes:

Item number	Description
1 - 	Expands the chat into a 90% screen-size window. The 90% screen-size window can't be resized or moved. Selecting the icon again resizes the chat back into the floating window.
2 - 	Positions, or pins, the Now Assist panel to the screen.
3 - Option buttons	Displays the available options.
4 - Reply to Now Assist...	Enter actions.
5 - Now Assist message	Indicates that the answers are generated by AI.
6 - Active chats	All active chats display in the Active chats section. in the Active chats section of the Now Assist. You can create additional chats by selecting the + icon in the heading.
7 - Closed chats	Displays all closed chats. If you select a closed chat, the chat's history displays.

The Now Assist subheader consists of these elements:

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



Now Assist subheader elements

Element	Description
1. Chats (☰)	All chats appear. Chats are organized with the most recent conversations at the top. Selecting a chat opens the chat in the conversation area. If there are unread chats or notifications, a badge number appears on the Chats icon (☰). Any unread chat or notification appears with a red dot next to it and the chat title appears in bold. Additionally, if you switch to a new chat while another active chat is ongoing, a pop-up message on the Chats icon (☰) appears: `Your previous chat was saved. You can revisit all of your past chats and continue ones that are still active.` The following list includes the chat sections that you may see in the chats area.

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

Now Assist subheader elements (continued)

Element	Description
	• Active: Chats where you can continue the conversation. If applicable, active chats move to the Closed chats section after two hours of inactivity. This 2 hour time limit can be configured within the Messaging Channels {sys_cs_channel.list} table. To change the inactivity time limit, from the Messaging Channels {sys_cs_channel.list} table, select the **NASS** record and populate the **Conversation Idle Timeout** field with your preferred active chat time limit. If you have no active chats, `No chatter at the moment` is displayed. If more than 12 active chats are running, a **Show more** link appears to view more chats. Selecting **Show more** displays an additional 10 chats. • Updates: Updates for important notifications and reminders can be configured to display. When *`sn_nowassist_va.nass_notification_enabled`* is set to `true` (the default), simple (nonactionable) Now Assist panel notifications appear on your portal home page. After selecting a notification, you're redirected to this Updates section. If you have no updates, `You're all caught up` is displayed. If more than four updates are available, a **Show more** link appears to view more updates. Selecting **Show more** displays an additional 10 updates. **Note:** If *`sn_nowassist_va.nass_notification_enabl`* is set to `false`, the Updates section doesn't appear. • Closed: Closed chats can be configured to display. A message closes when the designated time has passed (2 hours of inactivity) or you receive the following response in the chat: `It looks like you're finished with this chat, so I'll go ahead and close it.` Turn on closed chats by selecting the **Show closed chats** check box within **Conversational Interfaces** > **Assistants** > **[Selected Assistant Name]** > **Chat experience** > **Closed chats**. After being turned on, closed chats are displayed for as long as they're available within the Conversations (sys_cs_conversation) table. Closed chats appear in a read-only

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

Now Assist subheader elements (continued)

Element	Description
	mode and can't become active again. If more than four closed chats are available, a **Show more** link appears to view more closed chats. Selecting **Show more** displays an additional 10 closed chats. After a conversation has closed, you can't reopen it. Hovering over a closed chat displays the delete icon (). Confirm the chat deletion on the Delete this chat? modal to permanently delete the chat from the interface.
2. [Chat name]	The name of the conversation. If you select a promoted asset or query, that asset's title appears as the chat name. If instead you enter an utterance into the **Reply to Now Assist** field, your initial utterance becomes the chat name. The chat name appears in both the Now Assist subheader and **Chats list** > **Active** section.
3. New chat ( )	A new conversation begins. You may be prompted with a greeting message along with any promoted conversational assets such as topics, subflows, and/or actions, or suggested queries.
4. Feedback icons panel	You can indicate if the response was helpful by selecting the like thumbs up icon or if the response wasn't helpful, you can select the dislike thumbs down icon . You can also copy the response by selecting the copy message icon .

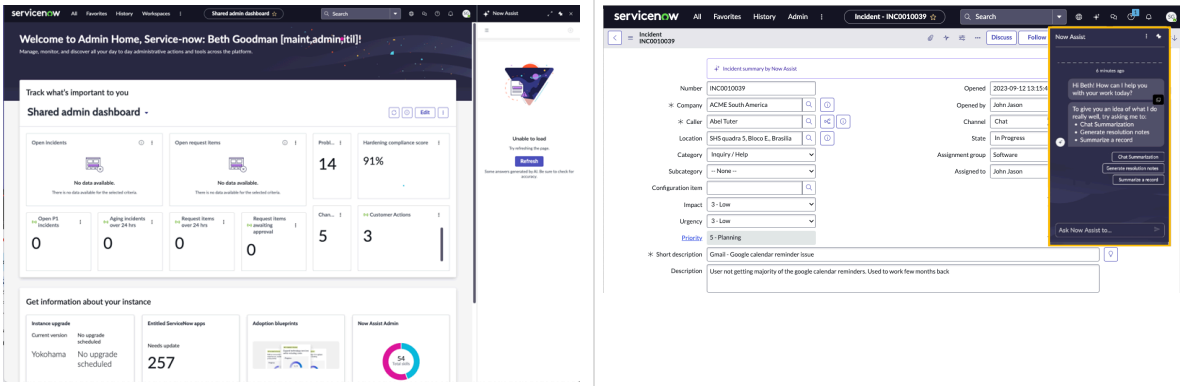
Now Assist panel is available on Next Experience and ServiceNow Studio. The following screenshots show the Now Assist panel in a workspace and on Core UI screens under Next Experience.

Now Assist panel

Next Experience	Core UI

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

Now Assist panel (continued)



Agentic conversations

Note: Admins must first enable AI agents before end users can experience agentic conversations. Now Assist panel discovers and executes agentic workflows. For more information on agentic workflows, see Now Assist agentic workflows and Multiple conversations in Now Assist AI agents.

When you ask a question to the Now Assist panel enhanced chat, the agent understands the query and begins a flow. When you submit a message with multiple questions or requests, Now Assist panel enhanced chat answers the multiple questions consecutively in its response. It can reason, plan, and execute across AI agents, Now Assist panel topics, conversational actions and subflows, catalogs, Knowledge Base articles, custom skills, and any Now Assist supported skills to help you. You receive on-screen messages to let you know where the agent is in the agentic processing flow prior to receiving the response. After the processing has completed its flow, a View AI Steps section header appears, where the processing flow steps can be expanded and viewed. You can stop the agentic processing flow at any time by selecting the End flow icon ( ). After an action starts, it can't be stopped. Selecting the End flow icon only stops the proceeding processing steps.

Navigating from the Now Assist panel

You can navigate from the Now Assist panel without leaving the current conversation by entering a navigation request in the **Ask Now Assist to...** field. If you enter "navigate me to active incidents," Now Assist displays a button that enables you to view the active incidents.

Chat summarization

Quickly learn the details of a chat by reading a chat summarization. The chat summarization gives you enough details about the chat so that your requester doesn't have to repeat the same information to you.

To generate a chat summarization from the Now Assist panel, select **Chat Summarization** or enter `summarize chat` in the **Ask Now Assist to** field.

Note: You can also generate a chat summarization by using the `/summarize` quick action in Agent Chat.

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

Case or incident summarization

Quickly learn the details of a case or incident by reading a case summarization. The summarization gives you enough details about the interaction so that your requester doesn't have to repeat the same information to you.

You can generate a case or incident summarization from the Now Assist panel for Now Assist for CSM, Now Assist for HRSD, or Now Assist for ITSM:

- For Now Assist for CSM, select **Summarize record** or enter `summarize a record` in the **Ask Now Assist to** field.
- For Now Assist for HRSD, select **Summarize record** or enter `summarize a record` in the **Ask Now Assist to** field.
- For Now Assist for ITSM, select **Summarize incident** or enter `summarize an incident` in the **Ask Now Assist to** field.

Conversation Help

Get specific and accurate answers to your queries by using the Get Help skill option on the Now Assist panel. This skill is available to everyone entitled to Now Assist capabilities.

For more information about the Now Assist Conversational Help skill that represents as Get Help on the Now Assist panel, see Now Assist Conversational Help.

Resolution notes generation

Quickly learn the details of how an interaction was resolved by generating and reading resolution notes.

To generate resolution notes from the Now Assist panel, select **Generate resolutions notes** or enter `generate resolutions notes` in the **Ask Now Assist to** field.

Streaming responses

After you enter a question or request on the Now Assist panel, Now Assist gathers information from Knowledge Base articles, external content, product documentation, catalog items, and workflows and combines them into a synthesized, comprehensive answer. Instead of waiting for the entire message to render, the synthesized response streams in real time and stops streaming after the entire message has been delivered. An animated sparkle icon (✨) appears while the response is generated and changes to the static sparkle icon after the response has fully loaded.

Fallback options

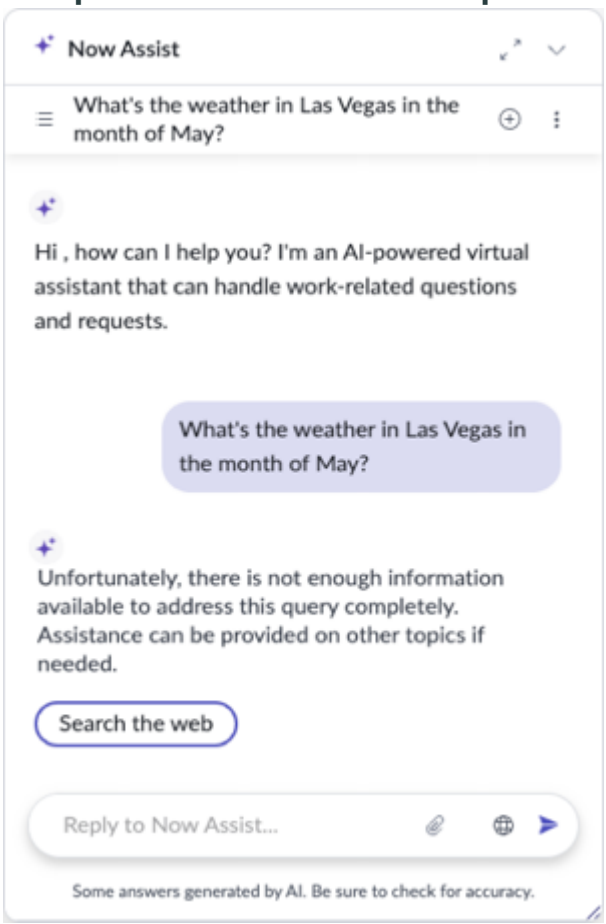
Note: For more information about where and how to enable fallback options, see Manage an assistant chat experience.

A fallback state can occur whenever search results are unavailable. Scenarios where search results are unavailable include when Now Assist didn't understand the query, complaint small talk was found, or an error occurred. When search results are unavailable, the **Search the web** fallback option may appear. If you select the **Search the web** fallback option, the web search mode is triggered and uses the internet to search for the results.

Note: Only the last query entered into the conversation is considered when entering web search mode via this **Search the web** fallback option.

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

Example of Search the web fallback option



Now Assist context menu

The Now Assist context menu uses generative AI to help agents summarize, create, and edit written content, thus streamlining their writing tasks.

Agents produce various types of content, including emails and chat replies. The Now Assist context menu uses generative AI to assist agents in summarizing, creating, and editing emails and chat replies. Agents can preview AI generated content, scroll through previously created content, and refine the text using the AI. The Now Assist context menu can be triggered from any application or field in a ServiceNow workspace where the Now Assist context menu is enabled.

The Now Assist context menu unlocks the power of generative AI and is available in Next Experience for:

- Customer Service Management (CSM)
- Human Resources (HR)
- IT Service Management (ITSM)
- Strategic Portfolio Management (SPM)
- IT Operations Management (ITOM)

The Now Assist context menu isn't available with Core UI.

Using the Now Assist context menu

The Now Assist context menu is available on any field where the floating Now Assist button ( ) appears. If you start typing in the field, a menu appears with the available Now Assist context

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

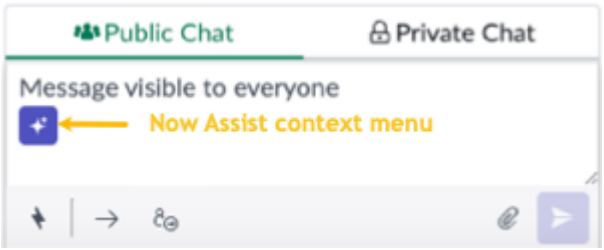
menu actions. The Now Assist context menu helps you summarize, create or modify existing documentation.

You can configure the Advanced filter to hide or show the Now Assist Context Menu quick actions option using the wwna_quick_actions table. To view more variable set configuration options, see #unique_131.

Chat window using the Now Assist context menu

You can use the Now Assist context menu when communicating with users in Agent Chat:

If it takes too long to generate text or the Now LLM Service isn't available, an error message appears.

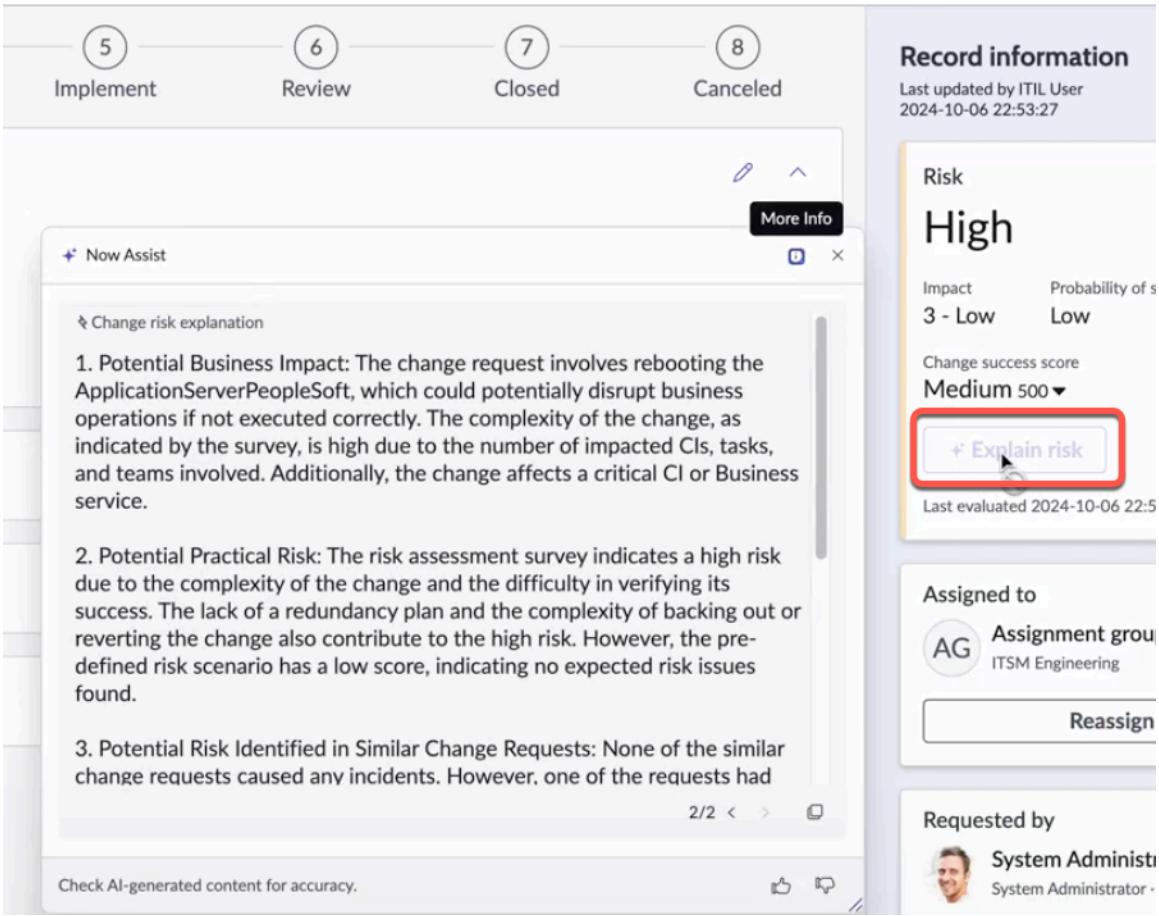


Change request risk explanation using the Now Assist context menu

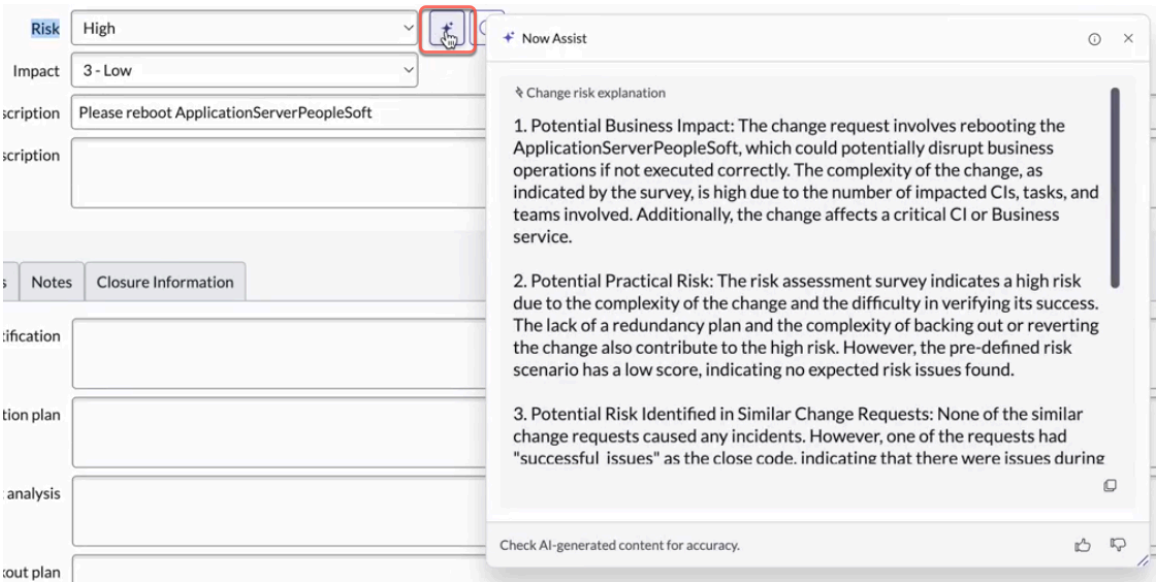
The Now Assist context menu makes the change request risk explanation available on the workspace and on UI16 after assessment and calculation.

When the risk is assessed and calculated, you'll see the **Explain risk** button with the Now Assist icon on the workspace showing the risk explanation in a dialogue box in the Record information section.

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



When the risk is assessed and calculated, you'll see the Now Assist icon against the **Risk** field on UI16, showing the risk explanation in a dialogue box.



Note: The risk explanation that is presented in the dialogue box is assessed and calculated on the change request form.

For more information about risk assessment and calculation, see Risk assessment.

Content editing in Knowledge Base articles using the Now Assist content menu

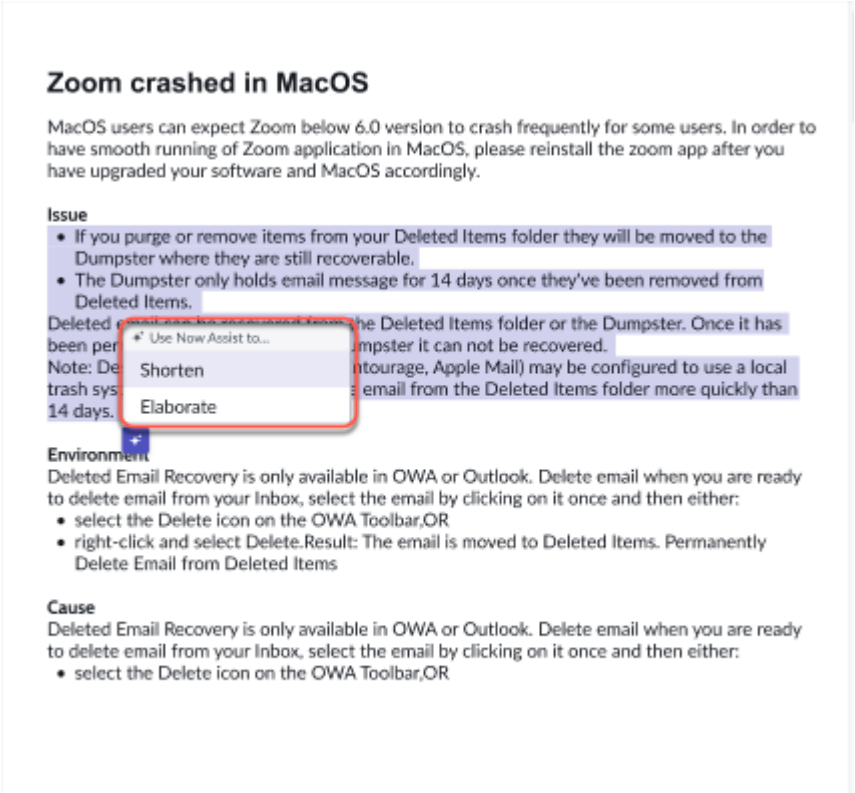
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

The Now Assist context menu enables generative AI assisted content editing capabilities for Knowledge Base authoring and to provide resolution notes in workspaces and UI16.

When you open a knowledge article and select the content in it, you'll see the Now Assist icon pop-up and float along with your mouse device.

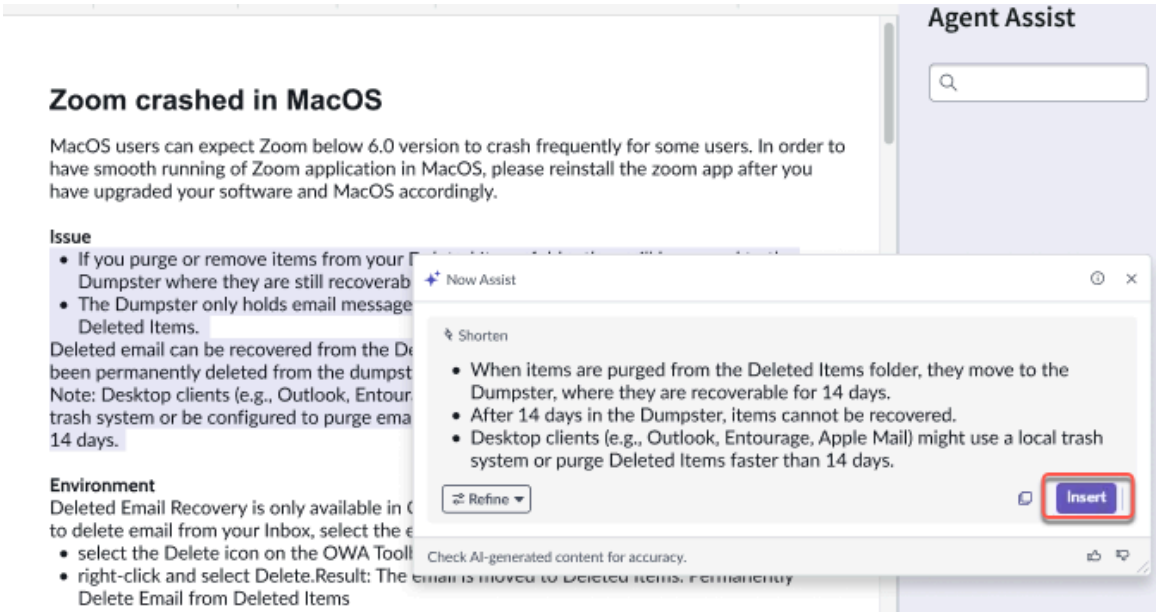
When you hover over the Now Assist icon, you'll see the following menu options to help you edit the content:

- **Elaborate**: Generative AI details the selected text.
- **Shorten**: Generative AI shortens the selected text.



You can insert the generative AI elaborated or shortened content into the Knowledge Base articles, using the **Insert** button and update or publish them.

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



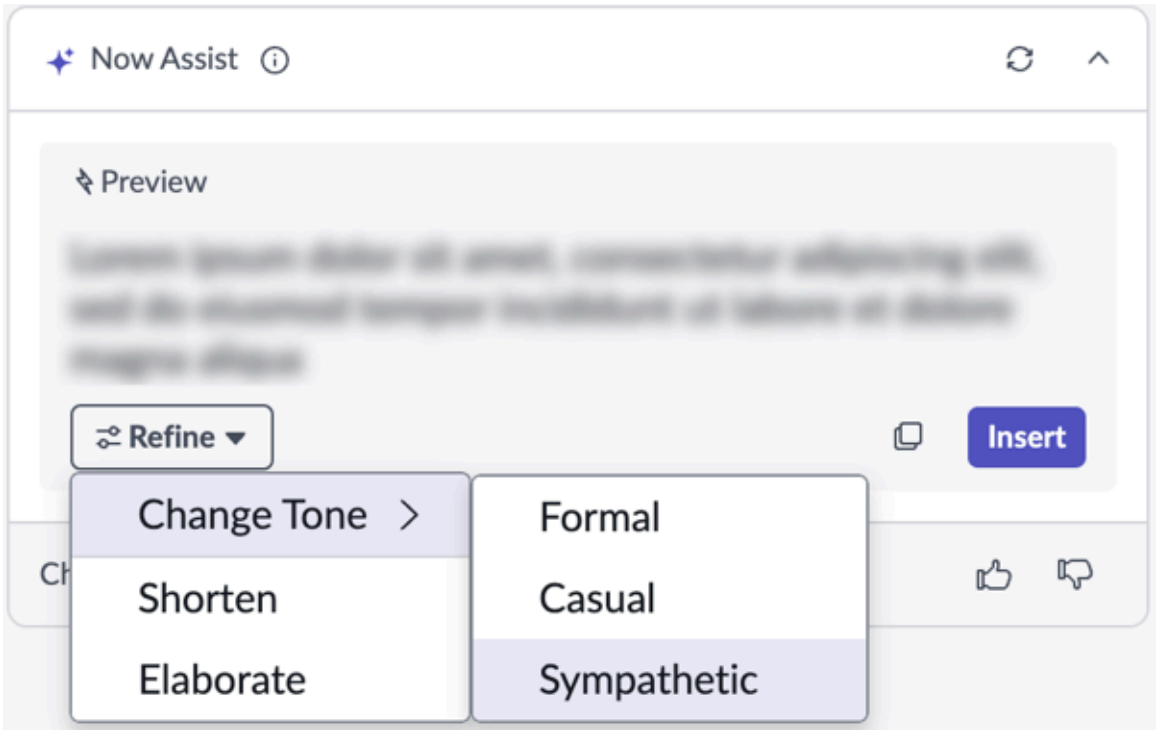
For more information about generating Knowledge Base articles using the Now Assist context menu, see Edit an article using the Now Assist context menu.

Change Tone using Now Assist context menu

The Now Assist context menu enables users to change the tone of their content. Users can choose between casual, formal, and sympathetic tone to enhance their content further, using the generative AI capabilities.

When you hover over the Now Assist icon or select **Refine** menu, you see **Change Tone** menu option.

You can further choose a preferred tone and select **Formal**, **Casual**, and **Sympathetic**. Review the changes and select **Insert** to finalize the new text.



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

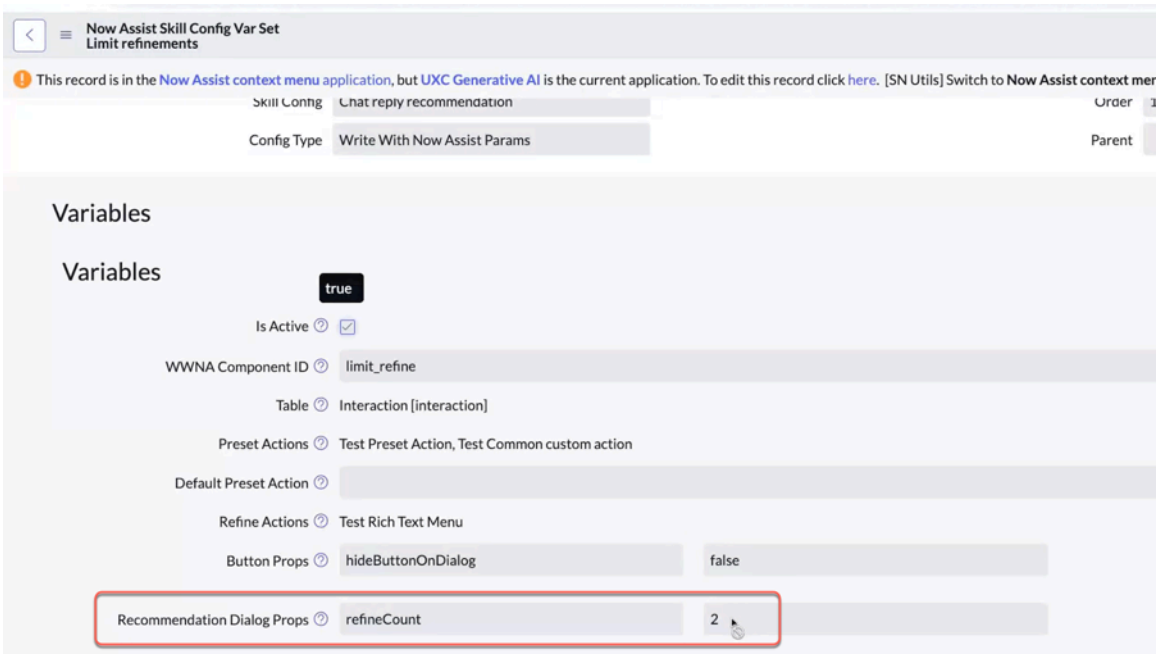
Note: If you do not see the change tone option for your application or product, reach out to ServiceNow.

Limit the number of content refinement calls using the Now Assist context menu

Configure and limit the number of content refinements to the skill per session using the Now Assist context menu. By default, the maximum number of refinements isn't set and you can configure to limit it.

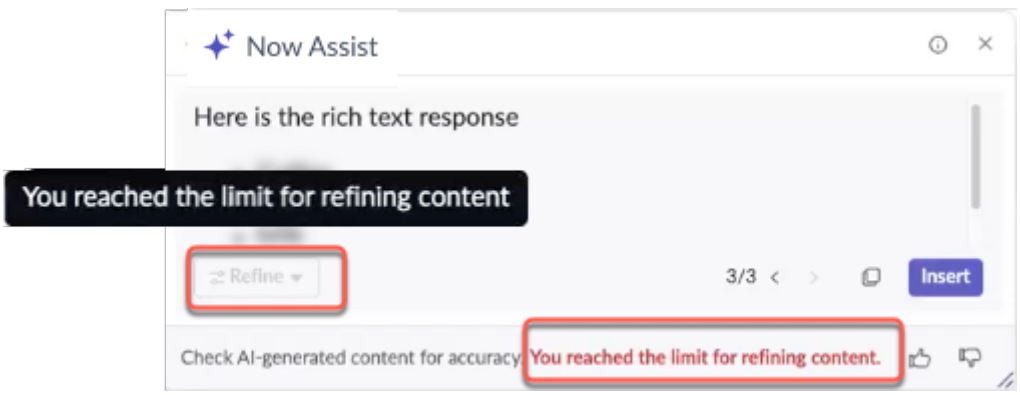
You can configure the refinement calls limitation at the **Recommendations Dialog Props** field in the **Limit refinements** record using the *`refineCount`* property.

Note: The **Limit refinements** record is available in the Now Assist config var set record in the Now Assist Skill context menu application.



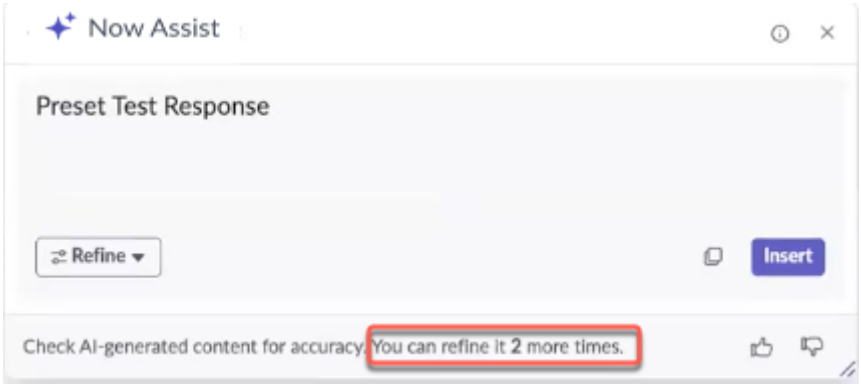
You can set the refineCount as follows:

- The default value of the property is **-1**. If the value is less than 0, then the number of refinements to the skill are unlimited.
- If you provide 0 as the value, then the refine button will be disabled and the `You reached the limit for refining content.` message is displayed.



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

- If you configure it with a value greater than 0, you'll be able to refine the content according to the set value. For example, if you set the refineCount to 2, then you will be able to refine the content only twice.



Create knowledge article with open prompt

Use NAcm to create inline knowledge article using open prompts. See Generate KB article with Now Assist context menu.

Set the minimum word count for the Now Assist icon

You can set a minimum word count required for the Now Assist icon to appear, to control when the icon is displayed based on content length. To set the minimum word count, update the value for the `minSelectedWordCount` property.

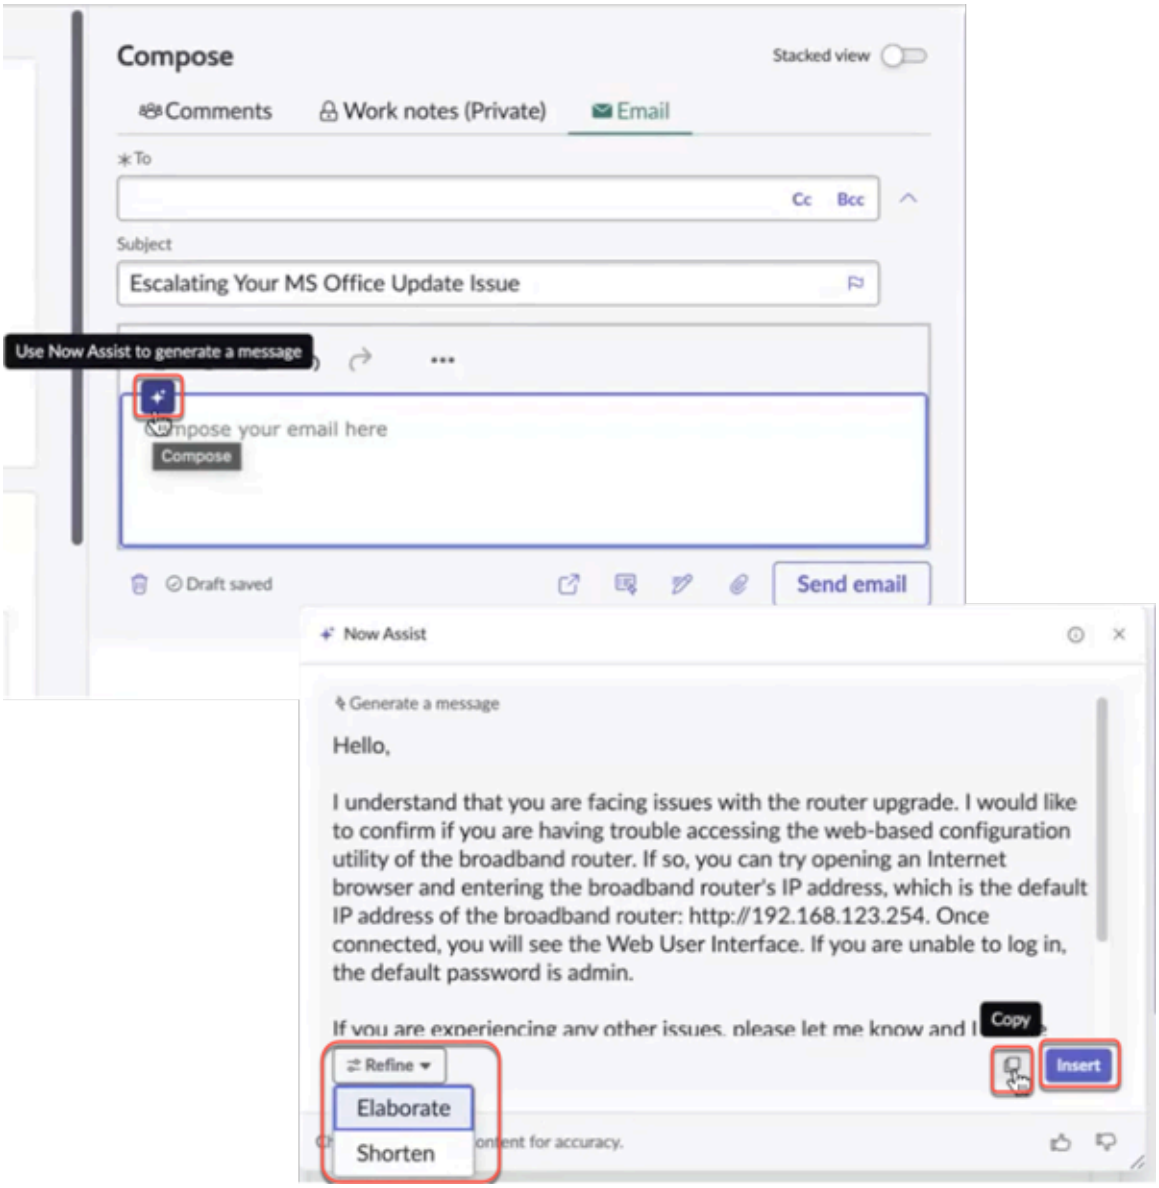
Email recommendations using the Now Assist context menu

Use the Now Assist context menu to compose or respond to emails with recommendations from Now Assist with generative AI template suggestions. The Now Assist context menu enables users to generate email response recommendations in new, forward, reply, or reply all scenarios.

Compose email

While you compose a new email, you see the Now Assist context menu icon ( ) that displays the message `Use Now Assist to generate a message`. When you select the icon, generative AI generates a message recommendation.

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

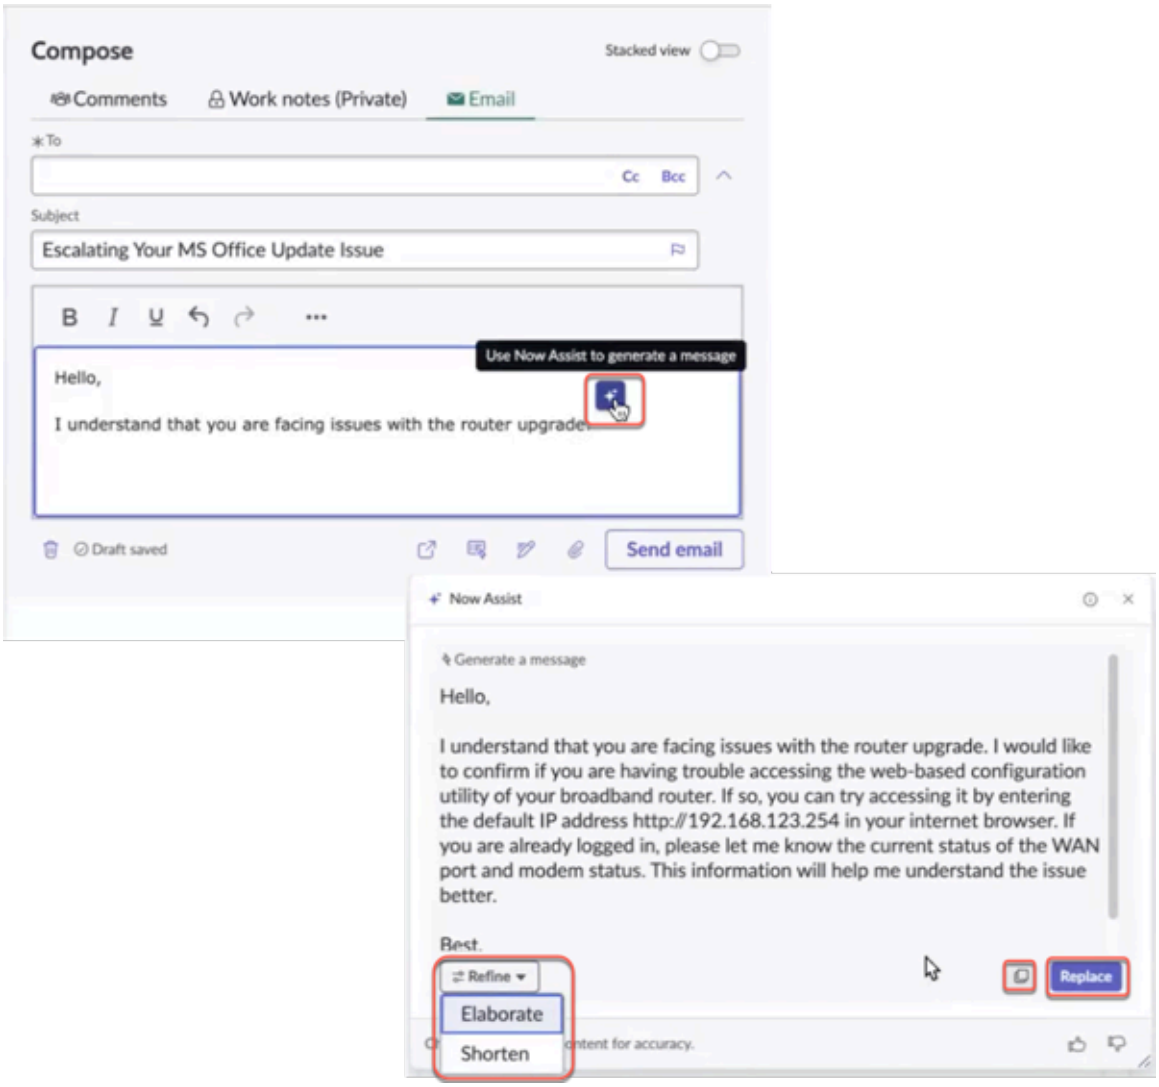


You can refine the generated text either by selecting the **Elaborate** or **Shorten** context menu options on the Now Assist Admin model. You can also copy the new text or **Insert** the generated text into the email body.

Complete draft emails

You can use the Now Assist context menu to finish your drafted emails. You can enter some text and use generative AI to help you with complete your drafts by selecting the Now Assist context menu icon ( ) that displays the message `Use Now Assist to generate a message`. When you select the icon, you see that generative AI is generating the message for you.

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



You can further refine the generated text by selecting **Elaborate** or **Shorten** context menu options on the Now Assist admin model. You can also copy the new text or select **Replace** to replace the generated text into the email body.

View inline citations

You can select the LLM response and view the all the related information sources like KB articles, emails, related records, etc.

Select the numbered citations to view details along with the link to the source. You can also select the **Show sources** option to view the list of all the related resource links.

Select **Replace** to add the source information in the email body as plain text. The numbered citation and all the source links will not be copied.

You can also select a record from the Now Assist Context Menu inline citation to navigate directly to the record page, where the referenced section will be highlighted for easy identification.

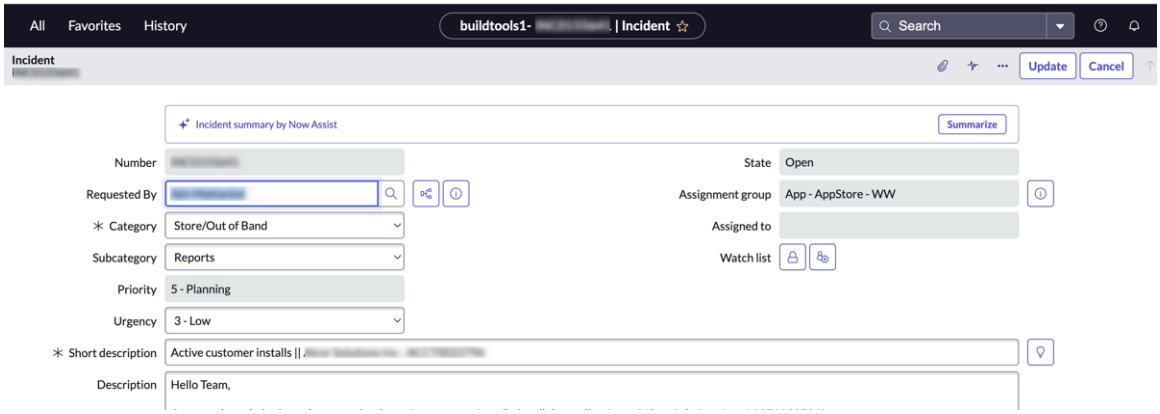
Summarize records with the Now Assist context menu

Use the Now Assist context menu to generate a record summary for the page, using Generative AI application assisted summarization capabilities in workspaces and UI16. The Now Assist

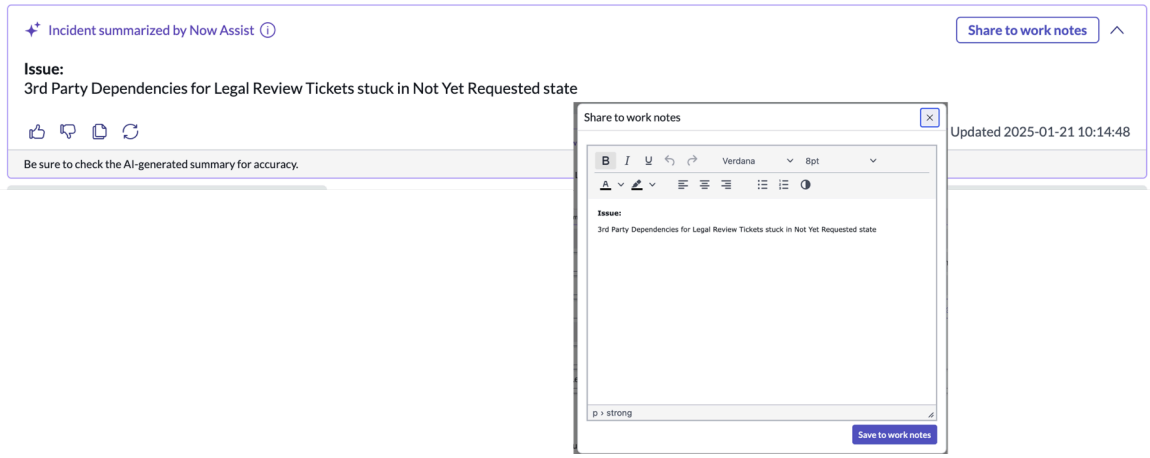
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

context menu can generate a new summary, expand or collapse the summary card, share the summary to work notes, regenerate, or copy the summary.

Go to any record page to access the Now Assist context menu icon ( ) with the message `Incident Summary by Now Assist`. Select **Summarize** to use generative AI to generate a record summary.



You can use the **Share to work notes** button to share the summary to work notes. You can also expand or collapse the summary card as required.



To provide feedback, use the feedback icons ( ) or ( ) If you're not satisfied with the provided summary, use the regenerate icon ( ) to regenerate the summary. You can also copy the summary to the clipboard using the copy icon ( )

Improve Docs content in Strategic Portfolio Management with Now Assist context menu

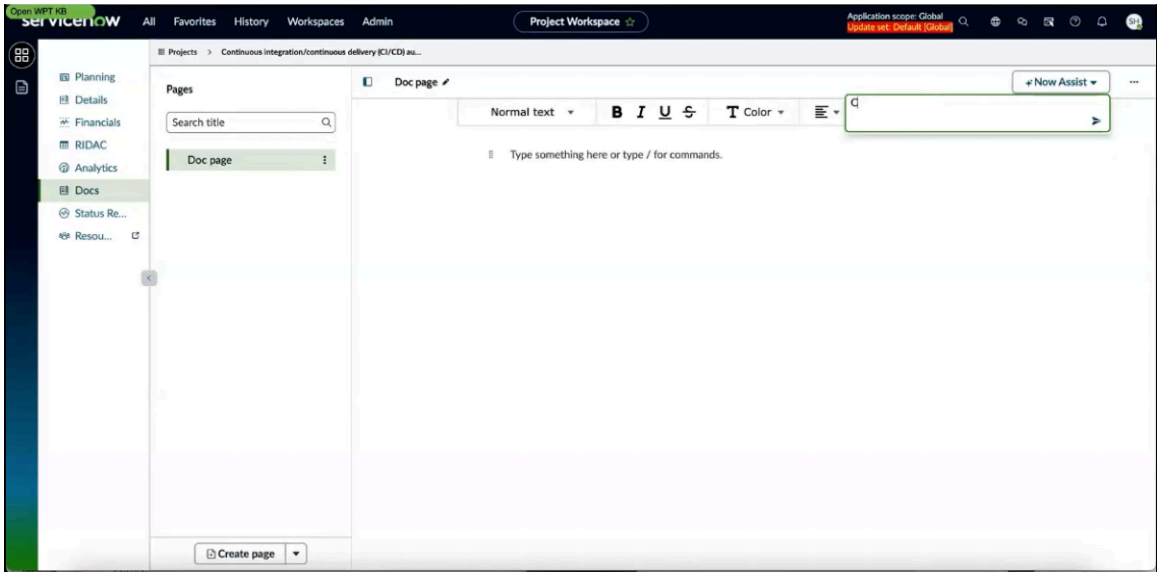
Use the Now Assist context menu to start conversations and raise queries, using generative AI capabilities and custom built in skills in Strategic Portfolio Management.

You can now interact with Now Assist directly from the docs section in Strategic Portfolio Management, Strategic Planning Workspace, Portfolio Planning Workspace, and Project Workspace to ask open-ended queries, create new content, add context, or improve existing sections. This helps you draft faster, refine ideas, and keep your work relevant without leaving the page.

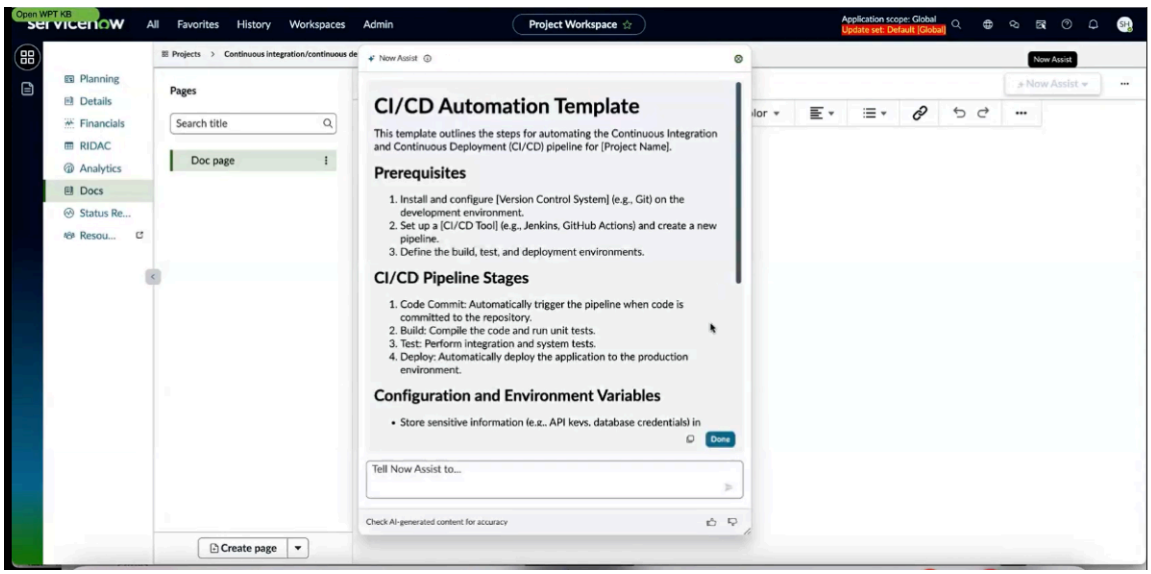
To access this feature, go to any doc page within Strategic Portfolio Management, Portfolio Planning Workspace or Project Workspace.

Select the Now Assist icon ( ) and enter your query.

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



The query response is generated in a new Now Assist window. You can use the input field to provide follow-up query responses and instructions.



You can also select the Now Assist icon ( ) and use the **Summarize** option from the drop-down menu to use generative AI for generating a record summary.

The Now Assist icon ( ) along with the text input field is also displayed when you select any text from the docs page.

You can also see the option to **Summarize**, **shorten**, **elaborate**, or **change tone** for the selected content.

Related topics

Generate and improve Docs content with Now Assist for Strategic Portfolio Management (SPM)

Generate KB article with Now Assist context menu

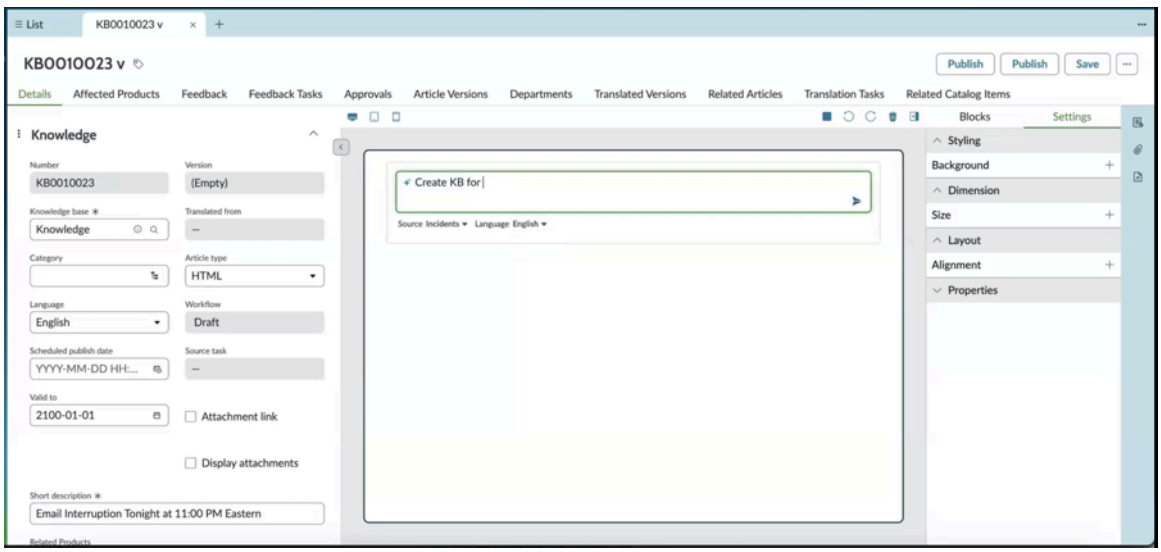
Use the Now Assist context menu to generate Knowledge Base articles in Knowledge Management.

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

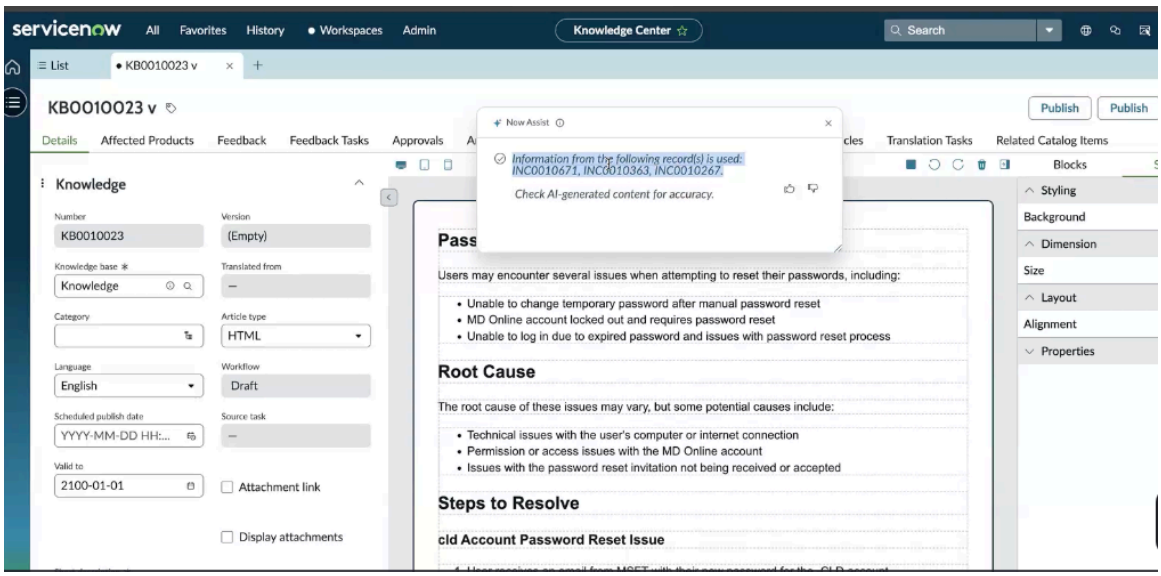
Before you begin
Role required: admin

Procedure

1. Navigate to **All** > **Knowledge** > **Knowledge Center**.
2. Select the list menu icon, ☰ to view all the KB articles.
3. Select a KB article from the list.
   The KB article details page is displayed with an open prompt inline option for KB generation.



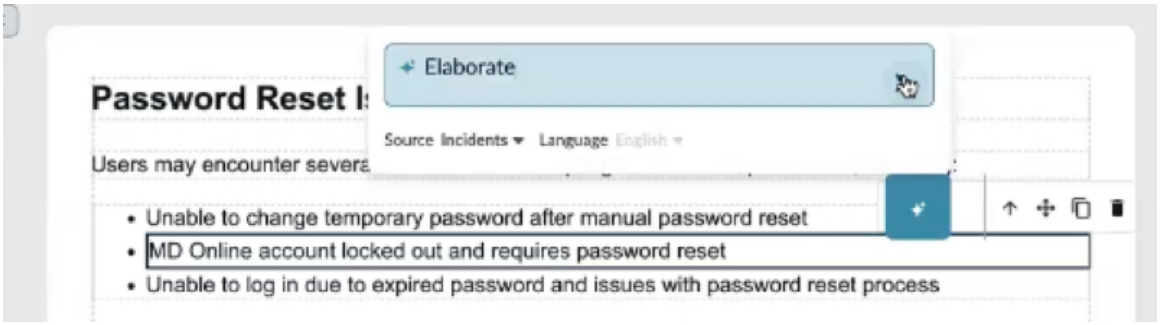
4. Enter the instructions for generating content, such as `Generate an article for password reset.`
5. Select the **Source** type and **Language**.
6. Press enter to start the KB article generation.
   The KB article preview is displayed.



   Review the generated article. You can edit the content, check grammar, and make the article easy to read. Select **Regenerate** to recreate the article based on your prompt.

7. Select the text to further refine the content with Now Assist Context Menu pop-over variant.

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



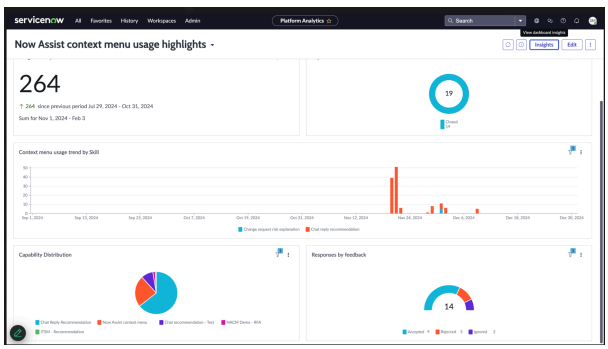
8. Type or select a menu option and press enter.
   Using the Now Assist Context Menu open prompt inline and pop-over variant you can quickly produce high-quality, accurate, and lucid knowledge articles, saving time and improving support in Knowledge Management.

Now Assist context menu usage dashboard

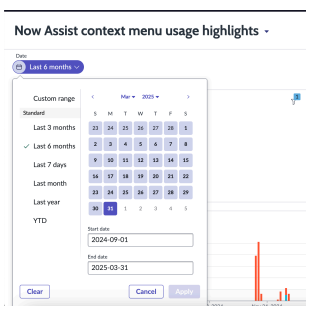
Use the Now Assist Context Menu dashboard to monitor the use of Now Assist context menu across the different applications.

The Now Assist context menu usage dashboard provides insights into usage patterns, frequency, and effectiveness of the context menu actions for the users and helps you refine the functionalities accordingly. The dashboard contains indicators that reveal:

- Usage metrics
- Usage based on application
- Usage based on skills
- Capacity distribution
- Responses by feedback

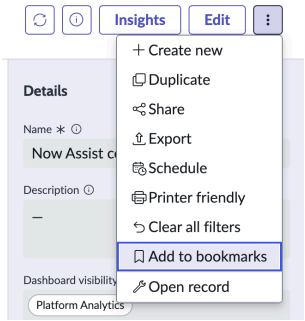


The Now Assist context menu usage dashboard enables you to filter the data based on the usage date. You can choose from the standard options or provide a custom range and select apply to filter the data.

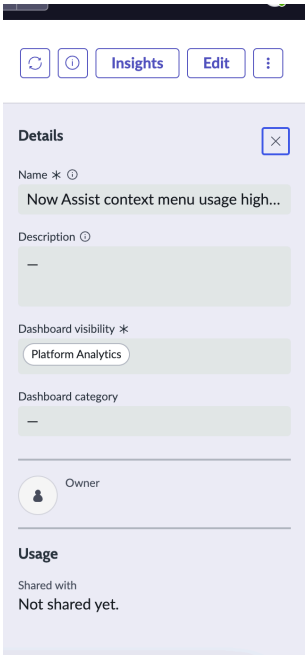


© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

You have an option to refresh, duplicate, edit, create, and export a dashboard.



You can also select the info icon () to view the dashboard-related information.



Key features

Usage matrix: The count or number of times, Now Assist context menu has been used during the selected duration. You have the option to refresh this matrix, view insights, or suggestion available and filter this data further based on the date.

Implicit feedback duration: The breakdown of feedback based on whether the response is inserted and closed during the selected time range. You have the option to refresh this matrix, view insights, or suggestion available and filter this data further based on the date.

Usage trend by skill: The total usage distribution based on the skills that use the Now Assist context menu. You have the option to refresh this matrix, view insights, or suggestion available and filter this data further based on the date.

Capacity Distribution: The capacity distribution based on the different applications that use the Now Assist context menu. You have the option to refresh this matrix, view insights, or suggestion available and filter this data further based on the date.

Response by feedback: The feedback field based on the Generative AI logs. The options are accepted, rejected, or ignored. You can also refresh this matrix, view insights, or suggestion available and filter this data further based on the date.

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

Insights: View insights and suggestions for Now Assist usage.

`

Use Now Assist context menu for custom skill deployment

Use the Now Assist context menu to deploy the custom skills created using Now Assist skill kit.

AI practitioners can use Now Assist skill kit to create a custom skill to provide custom solutions, with Now Assist's generative AI capabilities.

Now Assist allows the administrators to choose the custom skills for deployment. Use Now Assist context menu as a preferred channel to deploy a custom skill, when you configure deployment setting in Now Assist skill kit. For more information, see Configure skill deployment settings.

You must also select Now Assist context menu as a display option when you activate the skill in Now Assist Admin. For more information, see Activate a Now Assist skill.

To complete the activation process, create a new Now Assist context menu configuration. For more information, see Create Now Assist context Menu configuration

Create Now Assist context Menu configuration

Create a new Now Assist context Menu configuration to deploy and activate a custom skill.

Before you begin

To configure custom skills in action, ensure that the skill is activated in Now Assist.

Role required: sn_skill_builder.admin

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Now Assist Experiences** > **Now Assist Context Menu**.
2. Select **Configurations** tab in the Now Assist context menu home page.
3. Select **Create New**.
4. Enter values for the following fields on the Configure Now Assist context menu form.
   - Workflow
   - Product
   - Name
5. Select **Start Configuration**.
6. Provide the following details in the **General details** page, on the configuration form:
   - Name
   - Description
   - Table name
7. Select a location where you want to add the trigger for the Now Assist context menu.
   The options are:
   - Record form field: Select this option to add the trigger on UI16 form
   - Custom location: Select this option to add the trigger on any other desired location.
8. If you select Record form fields, add the form fields where you want the Now Assist context menu icon.

Note: Now Assist context menu can only be set up at the table level. Currently, it doesn't support filtering or configuring options based on specific record conditions.

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

9. If you select Custom location, choose the context menu display type.
   Choose between the following options of display:
   - Modeless window: A draggable and resizable dialog box.
   - Embedded card: A fixed window displayed on the page.

   If you select **Custom location**, ensure that you have completed the UIB Now Assist context menu component configurations. For setup information see Now Assist context menu UIB Setup.
10. Select **Save and continue**.
    You will land on the Configure experience page.
11. Select options for the following fields on the Configure experience page:
    - Actions
    - Refinement actions
    - Turn on to prevent access refinement action
    - Maximum refinements
    - Actions for generated content
    - Enable users to provide feedback on the recommendations

    Provide the same information for each of the form fields you have selected.

    If you proceed without providing the required values for each form field, system will prompt you to provide configuration.
12. **Optional:** When prompted, select **Yes, apply** on the Apply Configurations prompt if you want to apply the same configurations.
13. Select **Save and continue** once you review and edit the values, if required.
    You will land on the Review and activate page.
14. Select an option from the **Select a record to test the configurations** drop-down menu.
15. Select **Preview** and **Done**.

Create multiple Now Assist context menu skill configurations

Create multiple Now Assist context menu configuration for the same field and table.

Before you begin

You can now have multiple Now Assist context menu configurations on the same record form and field. You just have to create the required configuration and add to the form or field.

Role required: admin

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Now Assist Experiences** > **Now Assist Context Menu**.
2. Select **Configurations** tab in the Now Assist context menu home page.
3. Select **Create New**.
4. Enter values for the following fields on the Configure Now Assist context menu form.
   - Workflow
   - Product
   - Name

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

5. Select **Start Configuration**.
6. Provide the following details in the **General details** page, on the configuration form:
   - Name
   - Description
   - Table name

   The UI now allows you to add multiple configurations to the same field without any errors. This is applicable only if you select Record form field.
7. Select Record form field as the location, to add the trigger for the Now Assist context menu on the UI16 form.
8. Add and select one or multiple form fields in the **Form Fields** field.
9. Select **Save and continue**.
   You will land on the Configure experience page. You will see a table that lists the skill configurations along the order value, for the selected record field. The precedence is always given to the configuration with the lowest order value.

   Note:
   - In case there is a conflict due multiple configurations to the same field in a table,the precedence is give to the configuration with the lowest order value.
   - In case there are multiple configurations, but the order value is same, the configuration with the most recently modified date is applied to field.

10. Select options for the following fields on the Configure experience page:
    - Default Action
    - Actions
    - Refinement actions
    - Turn on to prevent access refinement action
    - Maximum refinements
    - Actions for generated content
    - Enable support for extended tables.

    Provide the same information for each of the form fields you have selected.

    Note:
    - To disable the inheritance model for child tables, toggle off the **Enable support for extended tables** button. This model applies the parent table's configurations to a child table if the child table does not have its own configuration.
    - Do not merge configurations. Use only the child table configuration and retrieve configurations from the child table.
    - Apply the merging logic to use the lowest value configuration for the fields.

    If you proceed without providing the required values for each form field, system will prompt you to provide configuration.
11. Select **Save and continue** once you review and edit the values, if required.
    You will land on the Review and activate page.

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

12. Select an option from the **Select a record to test the configurations** drop-down menu.
13. Select **Preview** and **Done**.

Analyzing Now Assist performance

Use the Now Assist Analytics dashboard to monitor the usage and performance of generative AI features and capabilities offered under Now Assist.

https://player.vimeo.com/video/1063716199?h=247ef99aa4&badge=0&autopause=0&player_id=0&app_id=58479

Get started

Explore Learn more about Now Assist Analytics	**Configure** Configure Now Assist Analytics
Use Use Now Assist Analytics	**Reference** Learn about user roles in Now Assist Analytics

Troubleshoot and get help

- Ask questions and explore other resources for in the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

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.

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

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 Analytics

Learn how Now Assist Analytics enables users with the Now Assist Analytics Viewer or Now Assist Analytics Admin role to monitor the usage, value, and performance of generative AI features and capabilities offered under Now Assist.

Now Assist Analytics overview

The Now Assist Analytics dashboard is built on the Platform Analytics experience. The dashboard contains indicators and breakdowns that provide actionable insights into the usage, value, and performance of your Now Assist implementation, including Now Assist Guardian, Now Assist context menu, and search.

Now Assist Analytics users

Users

User	Description
Now Assist Analytics Admin	Now Assist Analytics admins monitor usage, value, and performance of the Now Assist implementation, including Now Assist Guardian, Now Assist context menu, and search. They also configure skill-dashboard mappings to view individual skill usage and performance indicators.
Now Assist Analytics Viewer	View the dashboard pages to monitor usage and performance indicators.

Now Assist Analytics benefits

Now Assist Analytics benefits

Benefit	Feature	Users
Monitor usage, value, and adoption of Now Assist, skill performance, Now Assist Guardian, Now Assist context menu, and search.	• Usage and adoption • • Skills performance • Now Assist Guardian analytics • Now Assist context menu analytics • User search analyzer	Now Assist Analytics Admin or Now Assist Analytics Viewer

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

What to explore next

To learn more about configuring and using Now Assist Analytics, see:

- Configuring Now Assist Analytics
- Using Now Assist Analytics
- Now Assist Analytics reference

Configuring Now Assist Analytics

Configure the Now Assist Analytics dashboard to view the usage, value, and performance indicators of Now Assist.

Configuration overview

Now Assist Analytics requires at least one Now Assist application, for example, Now Assist for Customer Service Management (CSM), to be installed and configured on your instance. See Installing Now Assist Analytics for more information.

The following is an optional configuration task used to map a Now Assist skill to a dashboard.

Map a skill to a dashboard to view skill usage and performance indicators.

Domain Separation

Now Assist Analytics supports domain separation only for indicators using the following data collection jobs.

- [GenAI Analytics] Daily Data Collection
- [GenAI Analytics] Historical Data Collection
- [Now Assist Analytics] Daily Data Collection
- [Now Assist Analytics] Historical Data Collection

See Approaches to Performance Analytics with domain separation for more information on applying domain separation configuration.

Note: Be sure to check the Run as field in the data collection job records has a valid user.

Installing Now Assist Analytics

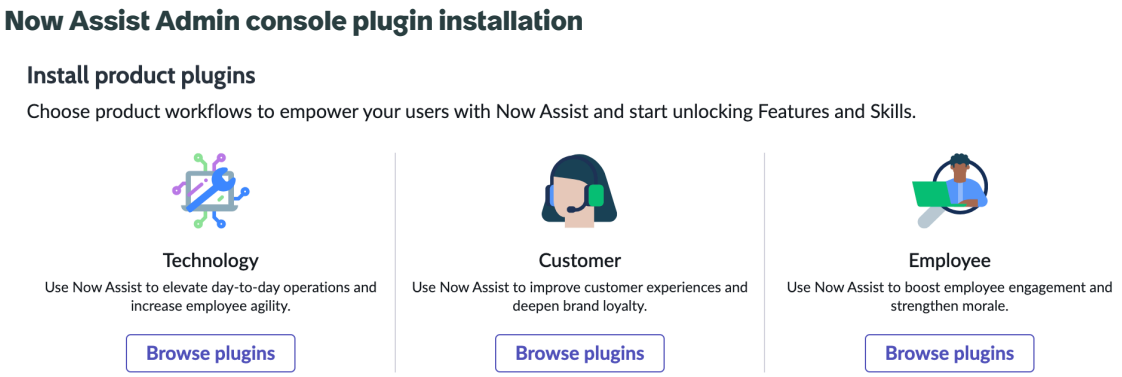
You can install the Now Assist Analytics application (sn_na_analytics) with Now Assist applications if you have the admin role.

Installation requirements

You must be on Australia Patch 0 or later.

Now Assist Analytics is included as a dependency for all Now Assist applications. It is not recommended to install the application by itself. Instead, you can install Now Assist applications from the Now Assist Admin console or directly from the ServiceNow Store. For details, see Install Now Assist plugins.

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



Compatibility matrix

Now Assist Analytics has a dependency on Now Assist Admin. Be sure to have the compatible version of the Now Assist Admin console based on the following compatibility matrix.

Now Assist Analytics version	Now Assist Admin console version	Release
Now Assist Analytics 2.0.14	Now Assist Admin 5.0.7	Australia Patch 1
Now Assist Analytics 1.1.11	Now Assist Admin 4.1.16	Australia Patch 0

Map a skill to a dashboard

Map a Now Assist skill to a dashboard to view skill performance indicators and skill details.

Before you begin

Be sure to map a dashboard with a skill in the same domain.

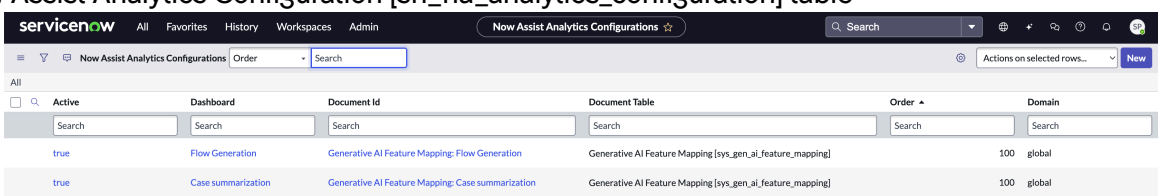
Note: You can only map a skill to a dashboard. Mapping a feature (that consists of multiple skills) is currently not supported.

If you're mapping a skill to a custom dashboard, be sure to share appropriate access to the dashboard. See Share a Platform Analytics dashboard for more information.

Roles required: sn_na_analytics.admin and sn_nowassist_admin.nsa_admin

Procedure

1. Navigate to the **All** menu and enter `sn_na_analytics_configuration.list`.
   The Now Assist Analytics Configuration [sn_na_analytics_configuration] table



   appears.
2. Create a new mapping by selecting **New**
3. On the form, fill in the fields.

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

Now Assist Analytics Configuration record form

Field	Description
Dashboard	The Dashboard that you want to map to a skill. Use the lookup icon () to search for and select the dashboard.
Application	The application that contains the record.
Document Table	Table that contains configured skills. Use the lookup icon () to search for and select the Now Assist Skill Config [sn_nowassist_skill_config] table.
Document Id	The Skill that you want to map to the dashboard. Use the lookup icon () to search for and select the skill that you want to map to the dashboard.
Active	Check box used to enable or disable the mapping.
Order	Order to set that determines the priority of the mapping in cases where multiple skills are mapped to the same dashboard.

4. Select **Submit**.

What to do next

After you've completed the mapping, go to the Skill details page and select the skill from the drop-down list to verify that the mapped dashboard is displayed.

Using Now Assist Analytics

The Now Assist Analytics dashboard provides indicators and breakdowns that help monitor the performance of generative AI features, capabilities, and skills active on your instance.

Access the dashboard by navigating to **All** > **Now Assist Admin** > **Analytics**. You must have Now Assist Analytics Admin [sn_na_analytics_admin] or Now Assist Analytics Viewer [sn_na_analytics_viewer] role to view the dashboard. The following sections explain the dashboard pages in more detail.

Usage and adoption

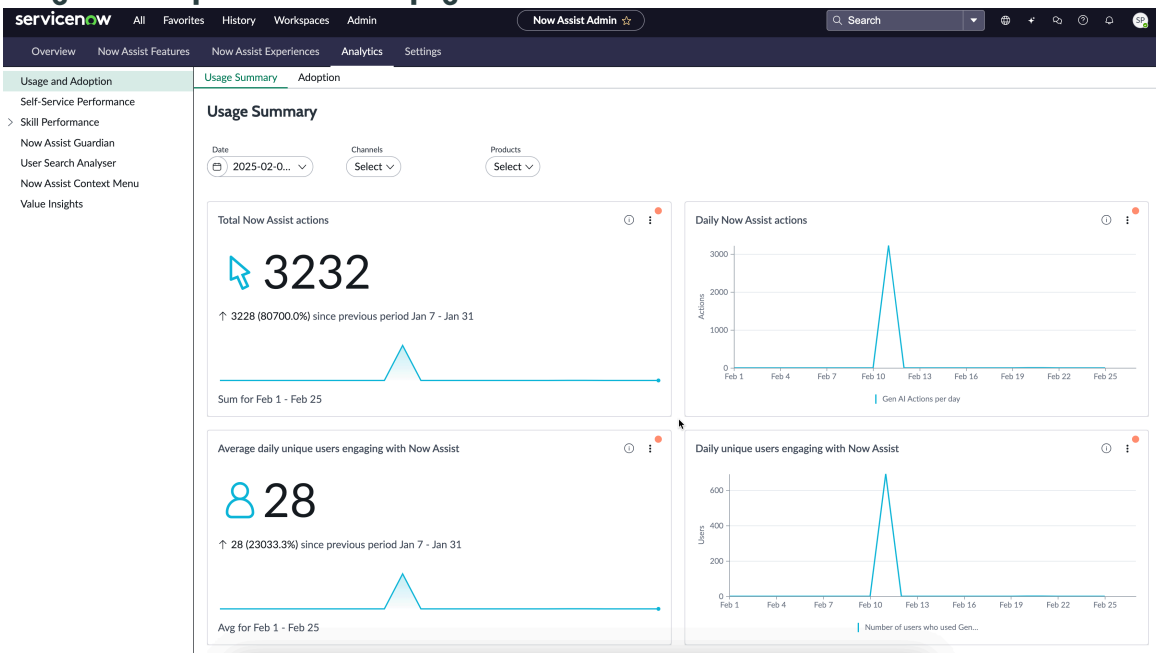
The Usage and adoption dashboard page contains key usage and performance indicators that help you evaluate the adoption of Now Assist in your organization.

The Usage and adoption dashboard page is the landing page for the dashboard. The Usage and adoption dashboard page contains two tabs that provide insights on usage summary and adoption.

The Usage summary page includes indicators on total and daily Now Assist actions, skill distribution and engagement trend, and daily unique users who have engaged with Now Assist. The Adoption page tracks the departments with the highest Now Assist usage, comparison of actions by department, and feedback and error details. Use the date range, channel, and product filters to view usage and adoption indicators for the selected period, Now Assist channel, and product, respectively. The default date range is 30 days.

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

Usage and adoption dashboard page



The indicators on the Usage and adoption dashboard page provide the following insights. See Now Assist Analytics dashboard indicator details for information on data source and calculations behind each indicator on the page.

- Skills engagement trend for a selected period can reveal skills that have been used more frequently or less frequently.
- Total and daily actions for a selected period can reveal the scale of Now Assist actions executed. The trend line in the visualization shows periods of increased or declining engagement.
- Average and daily unique users can reveal user engagement with Now Assist.

Usage summary indicators

Total Now Assist actions

This area of the dashboard shows the total number of Now Assist actions in the selected date range. A single use of a Now Assist skill represents an action. Select a filter combination to view the number of actions by products or Now Assist channels. For example, you can view the number of Now Assist actions executed through the Now Assist in Virtual Agent channel for Now Assist for Customer Service Management (CSM) product.

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

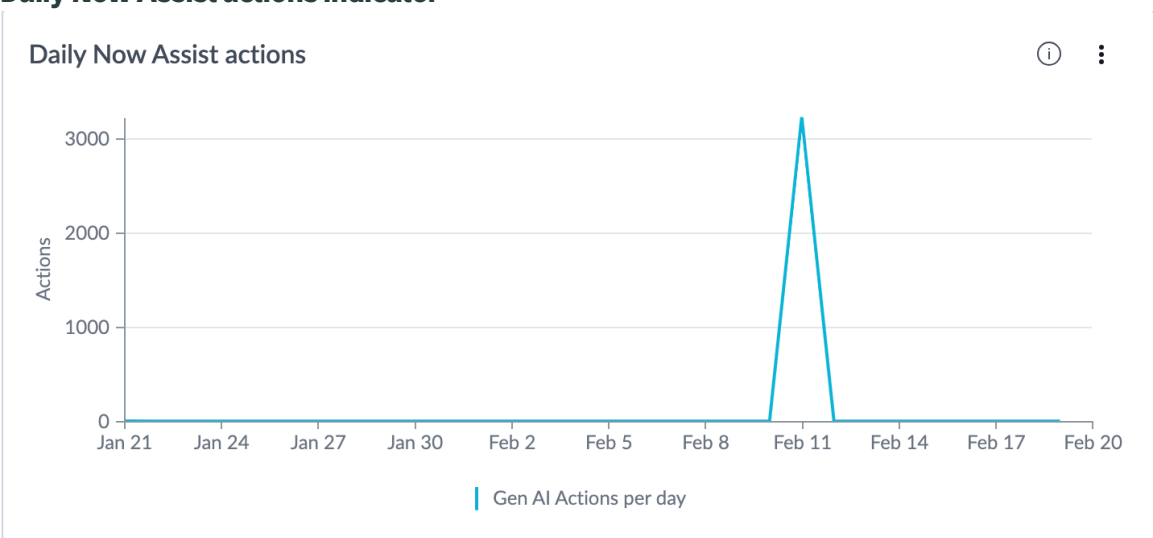
Total Now Assist actions indicator



Daily Now Assist actions

This area of the dashboard shows the number of Now Assist actions executed per day in the selected date range.

Daily Now Assist actions indicator

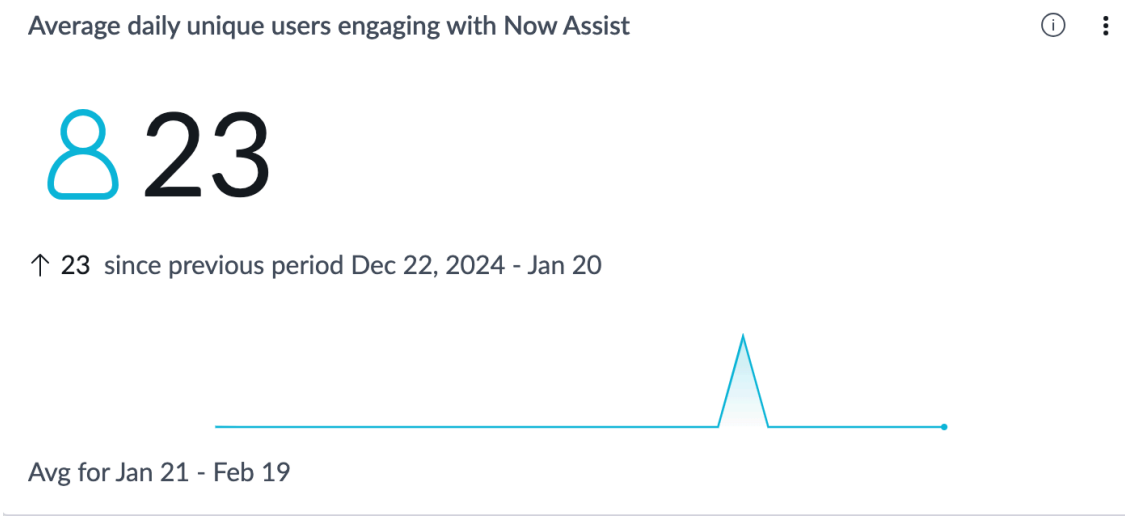


Average daily unique users engaging with Now Assist

This area of the dashboard shows the average number of daily unique users who engaged with Now Assist actions in the selected date range. Average number of daily unique users is an indicator of the level of Now Assist engagement across selected Now Assist products and channels.

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

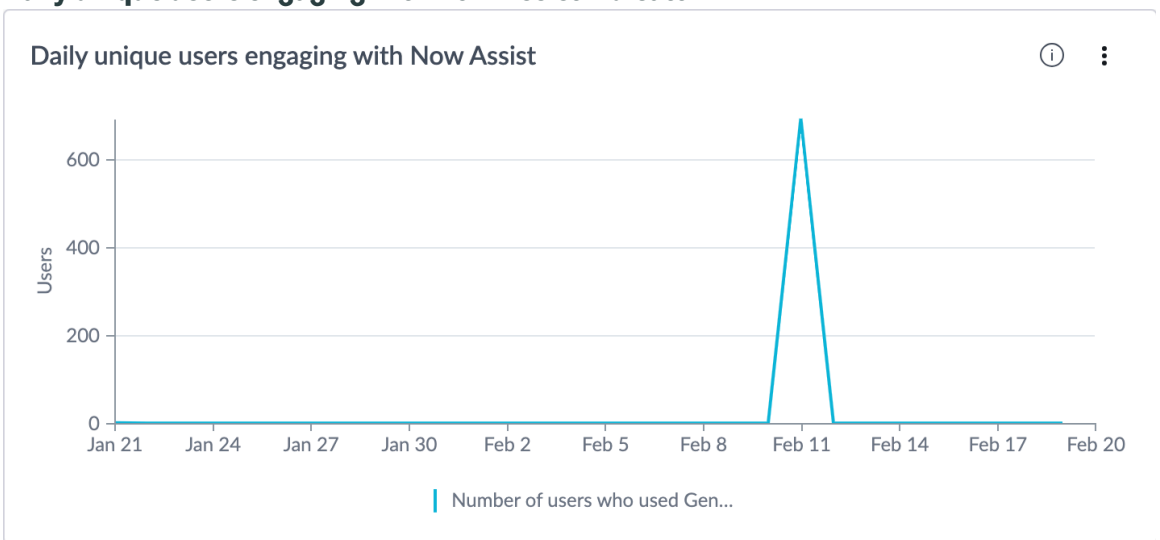
Average number of unique users per day who triggered actions indicator



Daily unique users engaging with Now Assist

This area of the dashboard shows the daily unique users who engaged with Now Assist actions in the selected date range.

Daily unique users engaging with Now Assist indicator

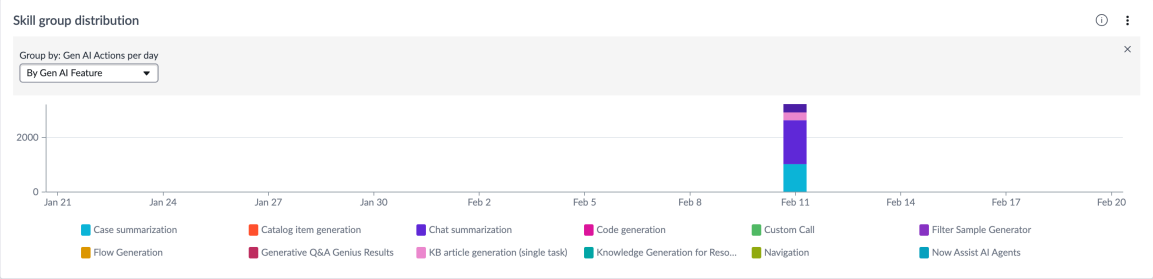


Skill group distribution

This area of the dashboard shows the skill usage in a trend chart for the selected filters. The visualization is interactive. Use the Group by filter to apply predefined breakdowns. Hover over the trend lines to see the number of times that each skill was used. Selecting one or more items in the legend removes the trend line for those items. You can compare and contrast skill usage distribution in a date range to understand in-demand skills, and skills that are low on usage and need further analysis.

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

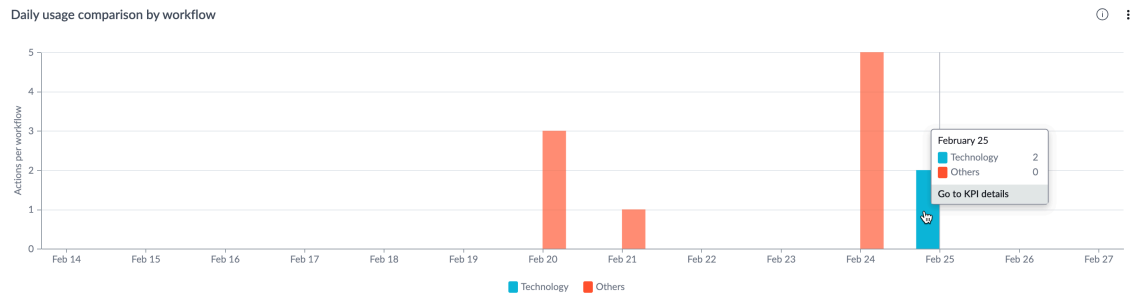
Skill group distribution indicator



Daily usage comparison by workflow

This area of the dashboard shows the number of Now Assist actions executed per day against each workflow in the selected date range. Hover over a bar to view the count of Now Assist actions executed for a workflow on that day.

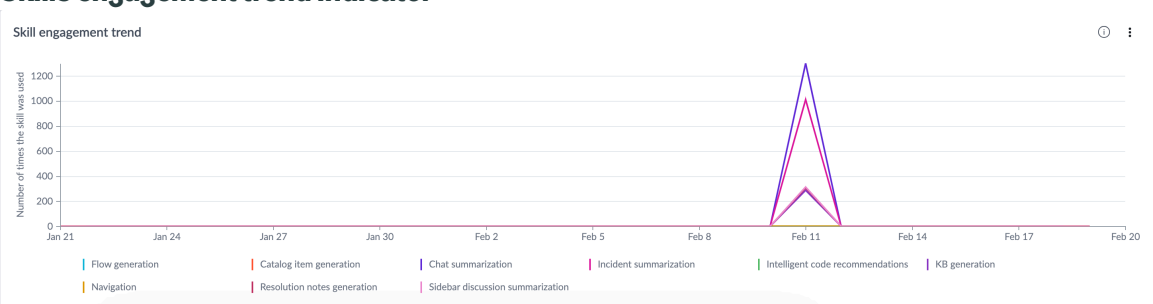
Daily usage comparison by workflow indicator



Skill engagement trend

This area of the dashboard shows the skill usage in a trend chart for the selected filters. The visualization is interactive. Hover over the trend lines to see the number of times that each skill was executed. Selecting one or more skills in the legend removes the trend line for those skills. You can compare and contrast skill usage levels in a date range to understand in-demand skills, and skills that are low on usage and need further analysis.

Skills engagement trend indicator



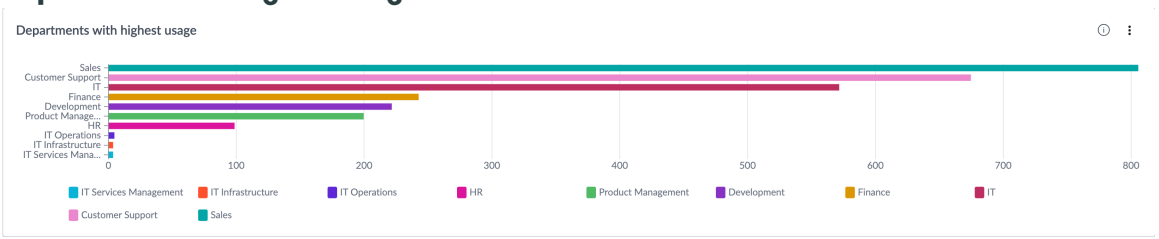
Adoption indicators

Departments with highest usage

This area of the dashboard shows the departments with the highest usage of Now Assist actions in the selected date range. Hover over the horizontal bars to view the count of actions for a particular department.

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

Departments with highest usage indicator



Now Assist actions comparison by user department

This area of the dashboard shows the number and type of Now Assist actions executed against each department in the selected date range.

Now Assist actions comparison by user department indicator

Now Assist actions comparison by user department

By Department / By Skills Config	Sales	Customer Support	IT	Finance	Development
Chat summarization	344	250	218	62	87
Incident summarization	238	225	190	61	71
Sidebar discussion summarization	75	58	53	21	24
Resolution notes generation	80	77	53	14	22

Feedback details

This area of the dashboard shows the number of Now Assist actions, number of feedback provided by users of Now Assist actions, and number of Now Assist actions with positive feedback in the selected date range.

Feedback details indicator

Feedback details

By Skill Family	Number of Now Assist actions	Number of Now Assist actions with feedback	Number of Now Assist actions with positive feedback
Others	5K	12	6
ITSM	3	0	0
Creator	3	1	1
Total	**5K**	**13**	**7**

Error details

This area of the dashboard shows the number of Now Assist actions and the number of Now Assist actions resulting in errors.

Error details indicator

Error details

By Skill Family	Number of now Assist Actions	Number of now Assist Actions with error status
Creator	3	1
ITSM	3	2
Others	5K	23
Total	**5K**	**26**

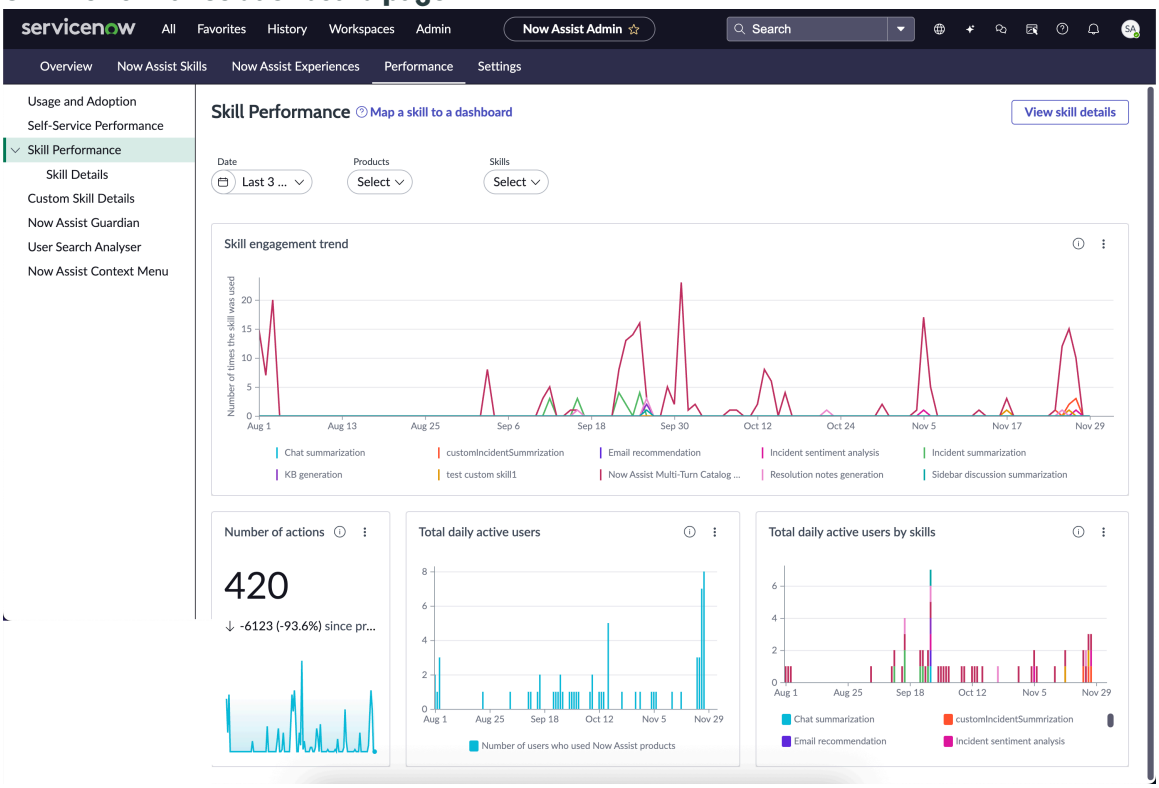
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective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

Skills performance

Use the Skills performance dashboard page to view usage and performance indicators of one or more Now Assist skills that are active.

The Skills performance dashboard page contains indicators that help you analyze the usage and performance of active skills. Use the **Date range**, **Product**, and **Skills** filters to break down by date range, Now Assist product, and skill, respectively. The filter selection applies to all visualizations on the page. See Now Assist Analytics dashboard indicator details for information on the data and calculations behind each indicator.

Skill Performance dashboard page



The indicators on the Skills performance dashboard page provide the following insights. See Now Assist Analytics dashboard indicator details for information on the data and calculations behind each indicator.

- Skills usage trend visualization for a selected period can reveal skills that have been used more frequently or less frequently.
- The Number of actions visualization for a selected period can reveal the scale of Now Assist skill executions. The trend line comparison shows the increasing or decreasing trend from the previous period.
- The Total daily active users indicator shows a breakdown of daily active users by product and skill.

Select the **View skill details** button to view the usage and performance indicators of skills mapped to a dashboard.

Skills performance indicators

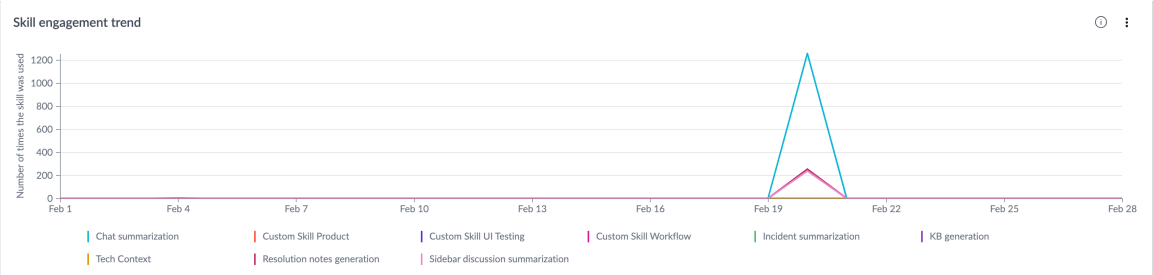
Skills engagement trend

This area of the dashboard shows the skill usage in a trend chart for the selected filters. The visualization is interactive. Hover over the trend lines to see the number

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

of times that each skill was used. Selecting one or more skills in the legend removes the trend line for those skills. You can compare and contrast skill usage levels in a date range to understand in-demand skills, and skills that are low on usage and need further analysis.

Skills usage trend indicator

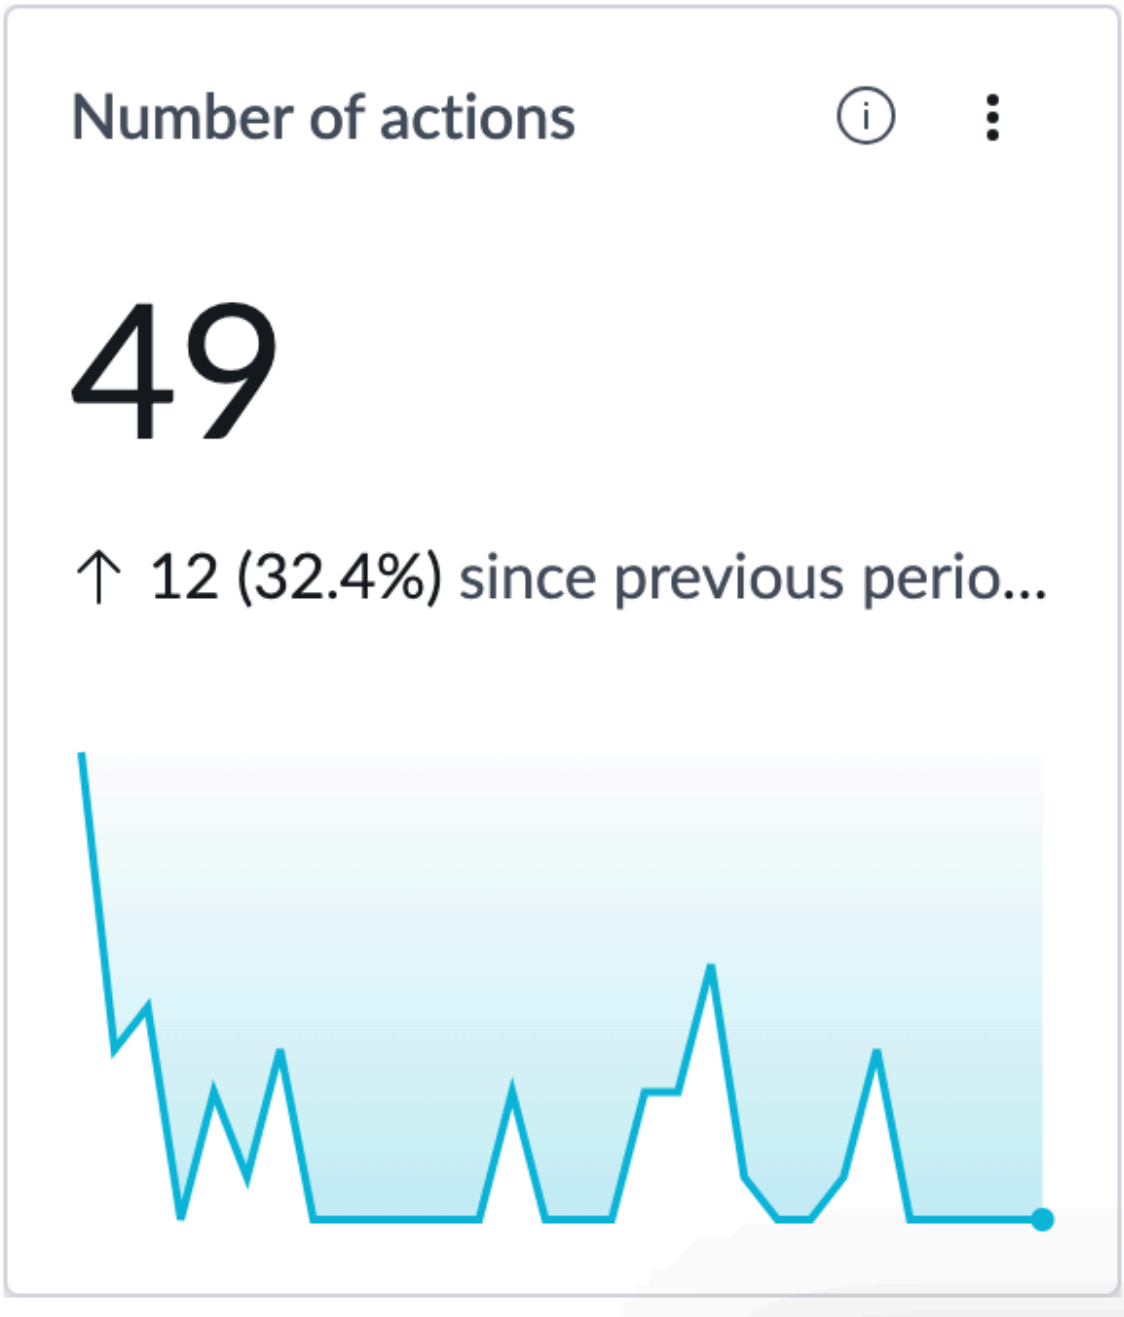


Number of actions

This area of the dashboard shows the number of Now Assist actions for the selected date range, product, and skill. A single use of a skill represents an action. The headline number is an indicator of the scale of skill usage across products. The trend line comparison shows the increase or decrease in the number of actions from the previous period.

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

Number of actions indicator

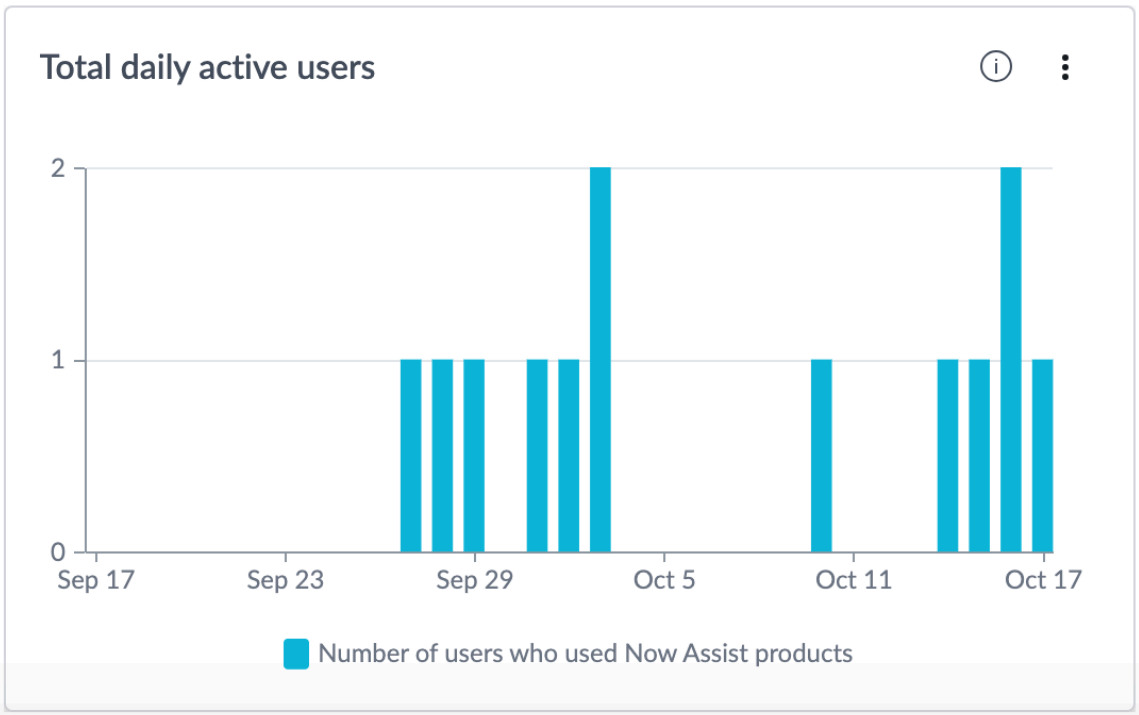


Total daily active users

This area of the dashboard shows the cumulative number of daily active users who have used Now Assist products with one or more active skills. Select a filter combination to see skill usage patterns. For example, you can visualize the number of users who have used the Chat Summarization skill across all products in the last month.

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

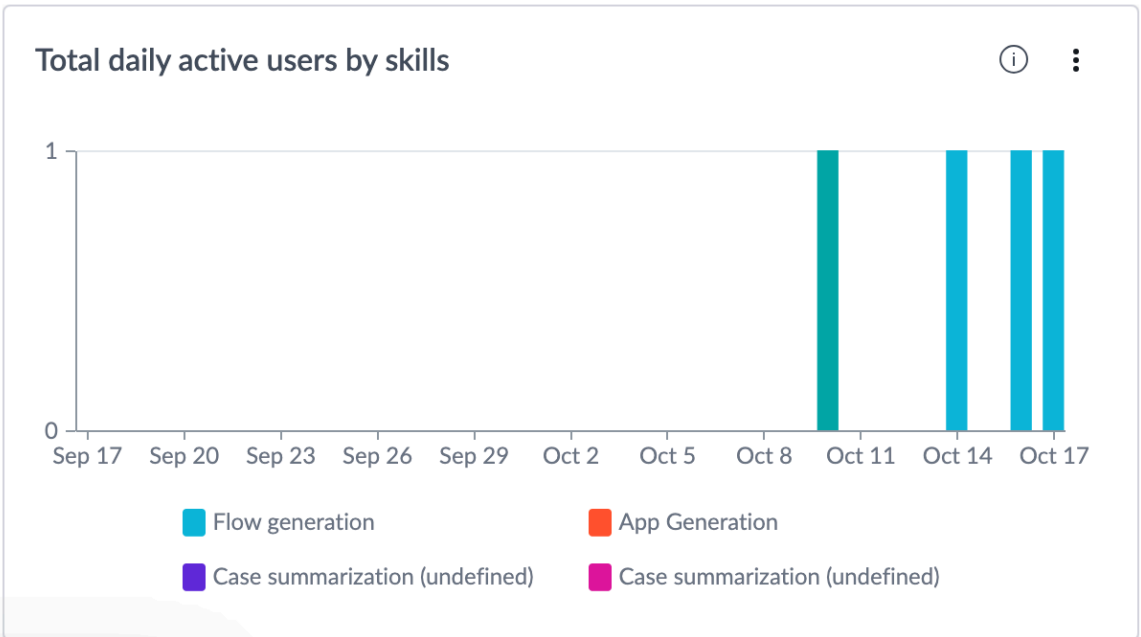
Total daily active users indicator



Total daily active users by skills

This area of the dashboard shows the number of daily active users who have used Now Assist skills in the selected date range. When you hover over a date in the visualization, a pop-over shows all active skills with the count of users against each skill. Select a filter combination to visualize patterns in user activity.

Total daily active users by skills indicator



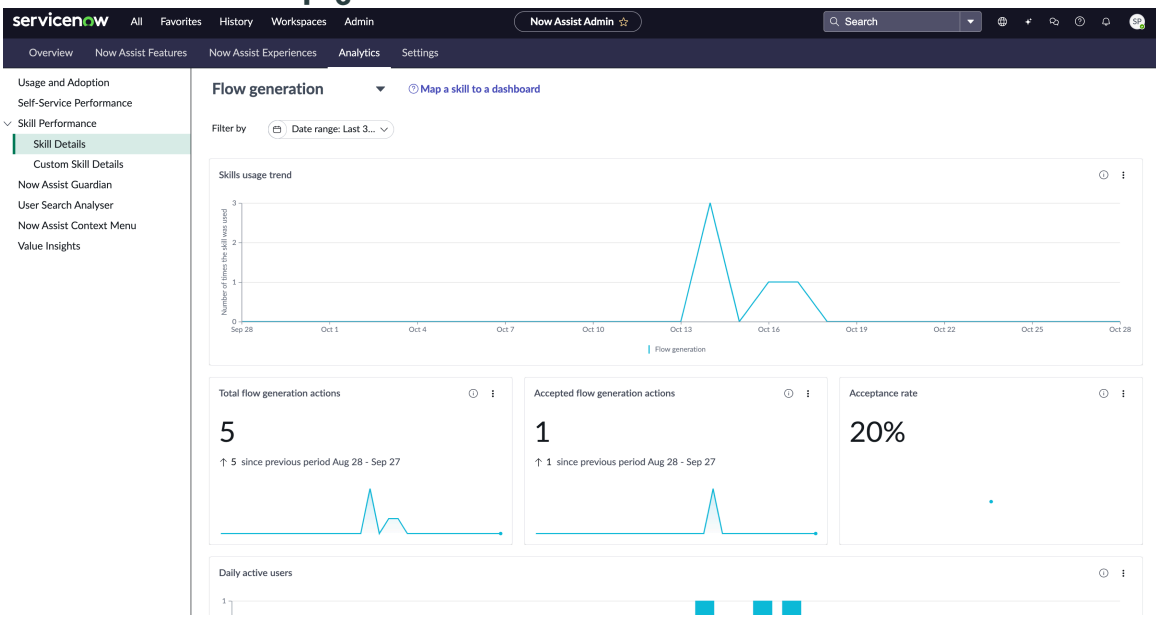
Skill details

Use the Skill details dashboard page to view usage and performance indicators of a skill.

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

The Skill details dashboard page contains indicators pertaining to a specific skill. The indicators provide insight into skill usage and performance. Select a skill from the drop-down list to view the indicators. The drop-down lists both active and inactive skills. Each skill has a subtitle that identifies the skill family that it belongs to, for example, ITSM, HR, and so on. Use the date range filter to view skill usage and performance over a certain period. The date range filter selection applies to all visualizations on the page. See Now Assist Analytics dashboard indicator details for information on the data and calculations behind each indicator.

Skill details dashboard page



The indicators on the Skill details dashboard page provide the following insights. See Now Assist Analytics dashboard indicator details for information on the data and calculations behind each indicator.

- Skill engagement trend visualization for a selected period can reveal patterns in skill usage.
- Acceptance rate visualization shows how well the skill met the requirements of users who used the skill. A high acceptance rate for a skill is an indicator of good performance. A low acceptance rate among skill users indicates that the skill doesn't meet the requirements either fully or partially.
- The number of actions visualization for a selected period can reveal the scale of the skill executions. The trend line comparison shows the increasing or decreasing trend from the previous period.
- The daily active users visualizations show a breakdown of daily active users by skill to help you see user activity on the skill.

Skill details indicators

The indicators on skill details pages might differ based on the skill selected. For example, summarization skills might have different set of indicators compared to generation skills because each skill is mapped to its own dashboard that contains a set of indicators related to the skill.

The Now Assist Analytics dashboard comes with some default skill-to-dashboard mappings to get you started. The default dashboards are visible to users with the Now Assist Analytics Viewer [sn_na_analytics.viewer] role. You can create your own dashboards and map them to skills. See Create a dashboard with the in-line editor and Map a skill to a dashboard for more information on creating custom dashboards and mapping them to skills, respectively.

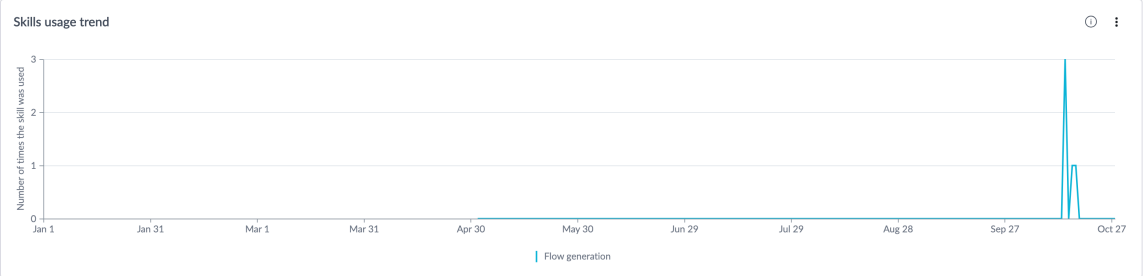
The following indicators are for the Flow Generation skill.

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

Skill usage trend indicator

This area of the dashboard shows the Flow Generation skill usage in a trend chart for the selected filters. The visualization is interactive. Hover over the trend lines to see the number of times the Flow Generation skill was used.

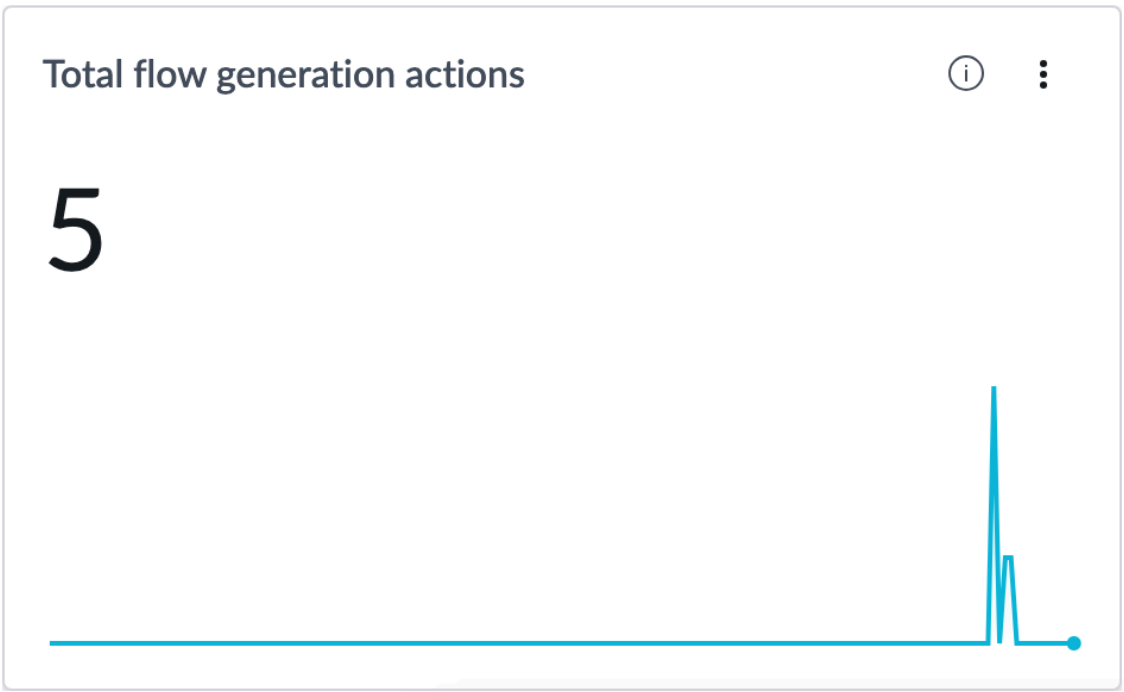
Skill usage trend indicator



Total Flow Generation actions indicator

This area of the dashboard shows the number of flow generation actions for the selected date range. A single use of the Flow Generation skill represents an action. The headline number is an indicator of the scale of flow generation skill usage across products. The trend line comparison shows the increase or decrease in the number of actions from the previous period.

Total flow generation actions indicator

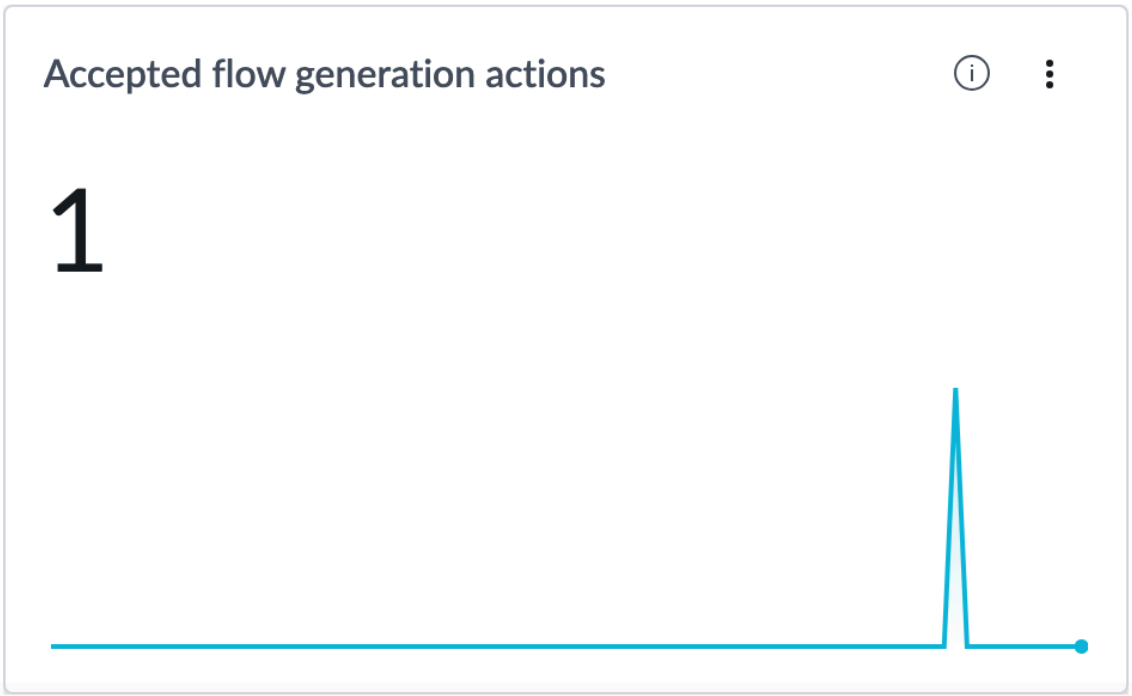


Accepted flow generation actions indicator

This area of the dashboard shows the number of flow generation actions, that is, the number of flow generation skill executions that resulted in flows that were accepted by the users.

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

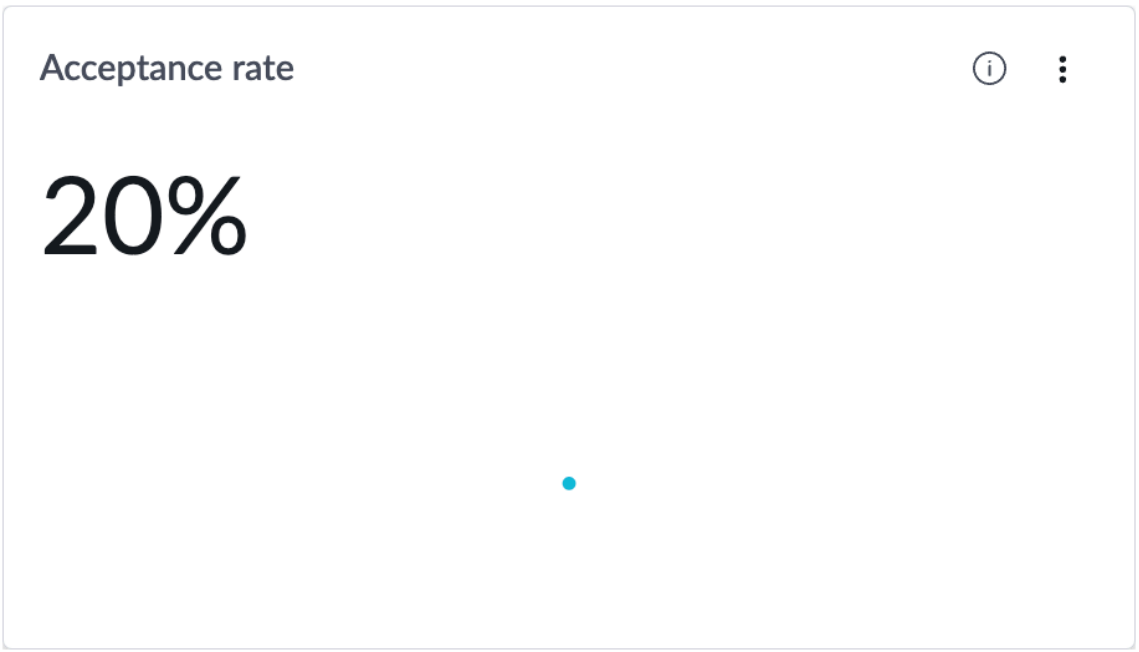
Accepted flow generation actions indicator



Acceptance rate indicator

This area of the dashboard shows the acceptance rate of Flow Generation skill based on user acceptance of the flow. The percentage is calculated using the formula: (Total number of accepted flow generation actions/Total number of flow generation actions) x 100.

Acceptance rate indicator



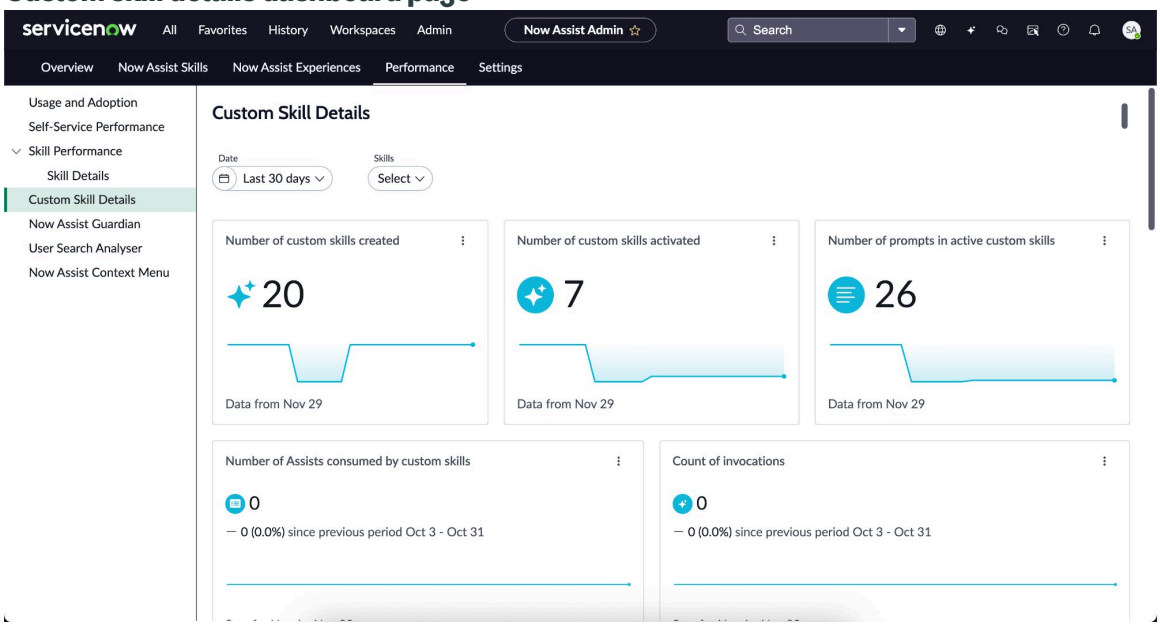
Custom skill details

Use the Custom skill details dashboard page to view usage and performance indicators of custom skills.

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

The Custom skill details dashboard page contains indicators pertaining to a custom skill. The indicators provide insight into skill usage and performance. Select a skill from the Skills drop-down list to view the indicators. The drop-down lists both active and inactive skills. Each skill has a subtitle that identifies the skill family that it belongs to, for example, ITSM, HR, and so on. Use the date range filter to view skill usage and performance over a certain period. The date range filter selection applies to all visualizations on the page. See Now Assist Analytics dashboard indicator details for information on the data and calculations behind each indicator.

Custom skill details dashboard page



The indicators on the Custom skill details dashboard page provide the following insights. See Now Assist Analytics dashboard indicator details for information on the data and calculations behind each indicator.

- Skill engagement trend visualizations for a selected period can reveal patterns in skill usage across workflows, products, and features.
- The daily unique users visualization shows a breakdown of daily unique users by the selected skill to help you see user activity and engagement with the skill.
- Acceptance rate visualization shows how well the skill met the requirements of users who used the skill. A high acceptance rate for a skill is an indicator of good performance. A low acceptance rate among skill users indicates that the skill doesn't meet the requirements either fully or partially.
- Skills feedback visualization shows that the feedback recorded based on user response to each skill execution in the selected date range.

Skill details indicators

The Custom skill details page contains a common set of indicators that give you insights into custom skills. Only those skills where the skill family is Other or the category is Custom are shown in the skills filter on the page.

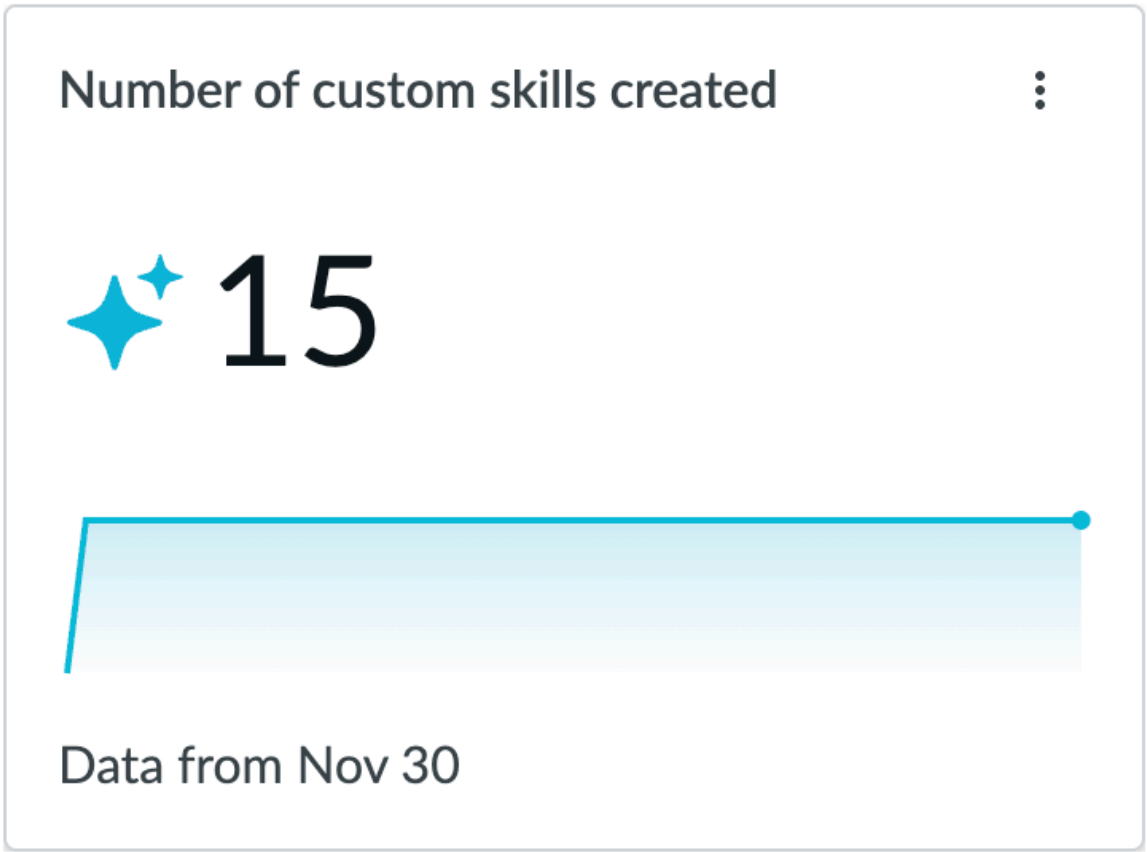
The following indicators are common across all custom skills.

Number of custom skills created

This area of the dashboard shows the number of custom skills created by the users in the selected date range.

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

Number of custom skills created indicator

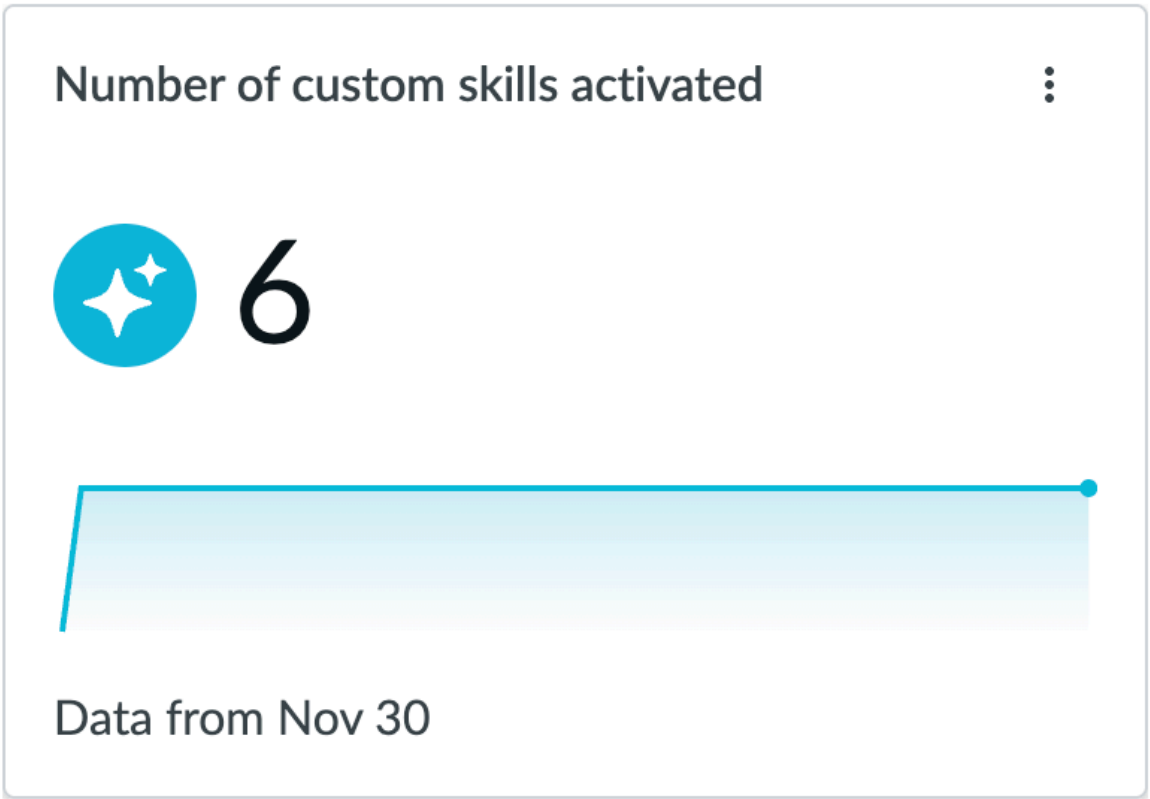


Number of custom skills activated

This area of the dashboard shows the number of custom skills activated by the users in the selected date range.

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

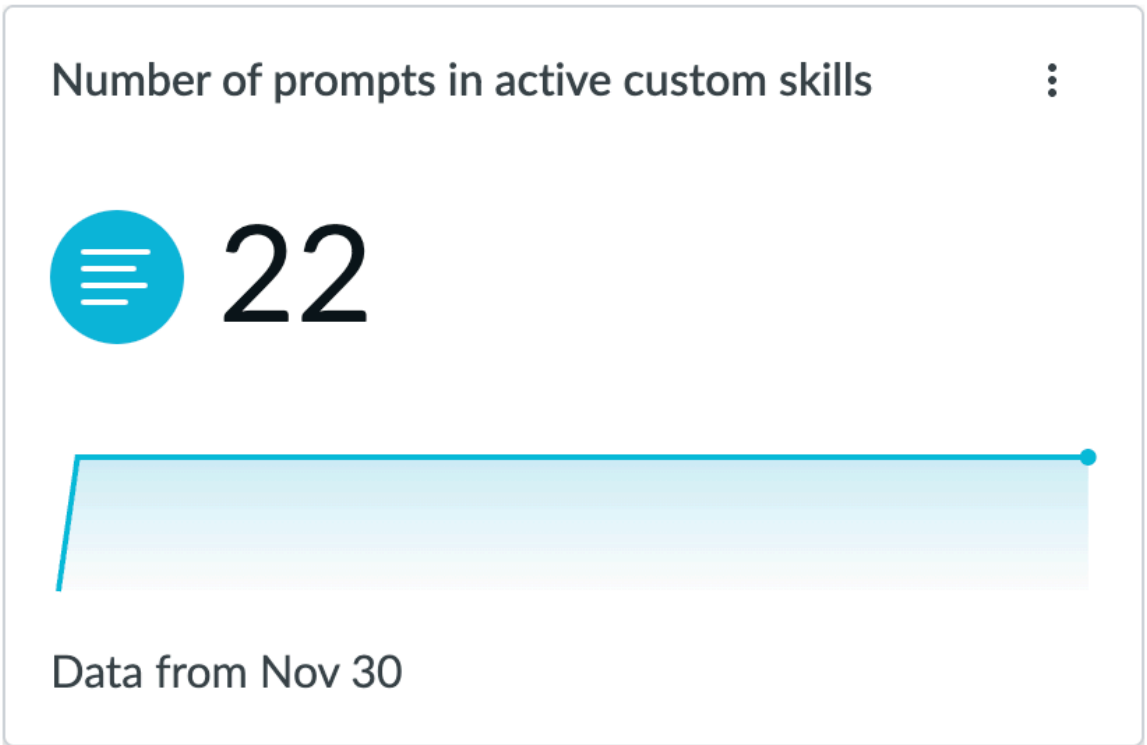
Number of custom skills activated indicator



Number of prompts in active custom skills

This area of the dashboard shows the number of prompts in active custom skills in the selected date range.

Number of prompts in active custom skills indicator

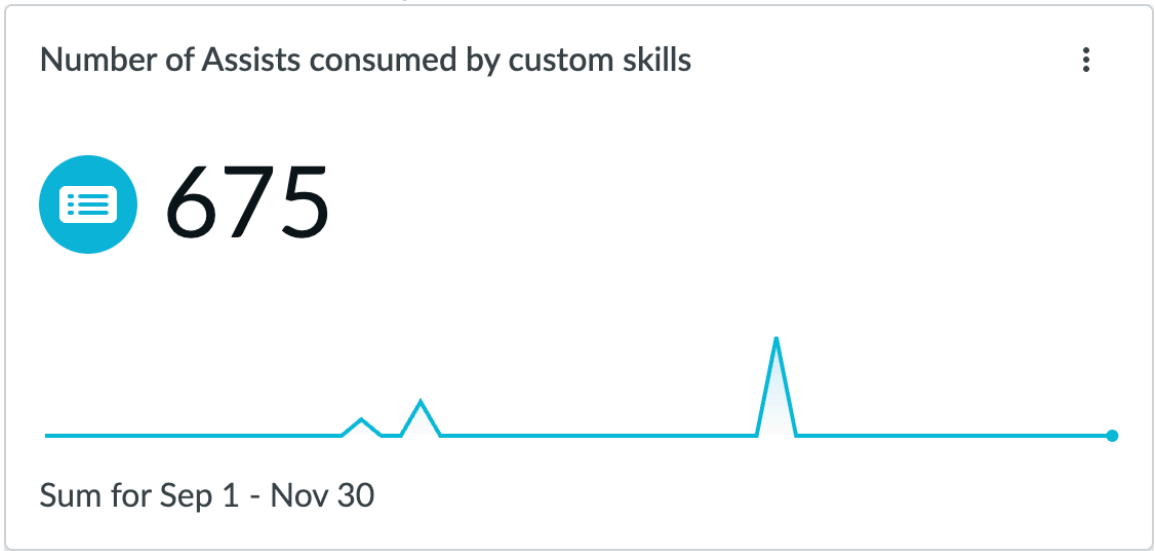


Number of Assists consumed by custom skills

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

This area of the dashboard shows the number of Assists consumed by custom skills in the selected date range.

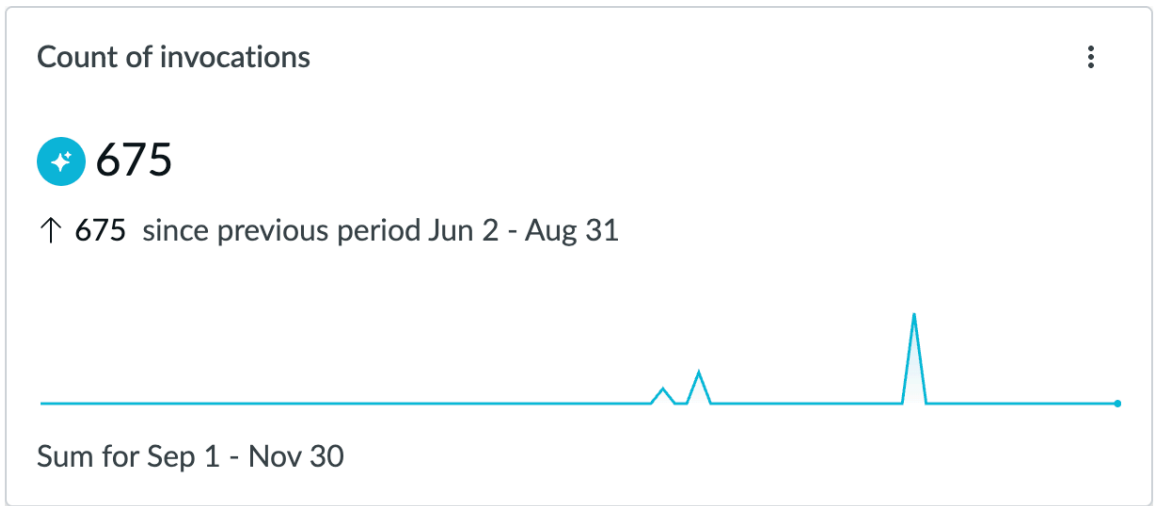
Number of Assists consumed by custom skills indicator



Count of invocations

This area of the dashboard shows the count of invocations for the custom skills in the selected date range.

Count of invocations indicator

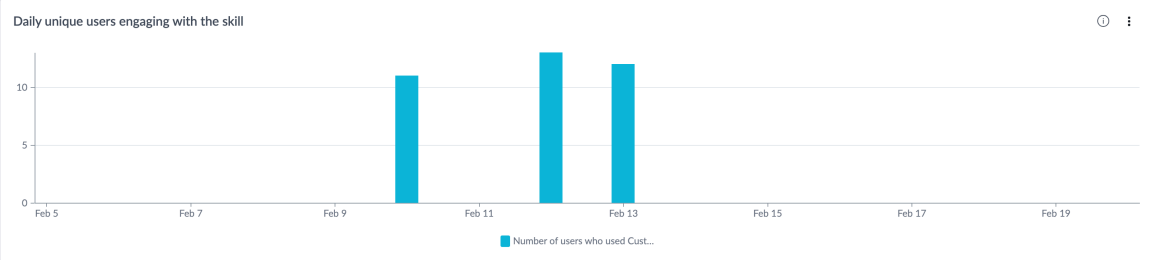


Daily unique users engaging with the skill

This area of the dashboard shows the number of unique users per day who engaged with the skill in the selected date range. The bar chart shows a trend of increase or decrease in the number of unique users to help you understand periods of high and low skill engagement.

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

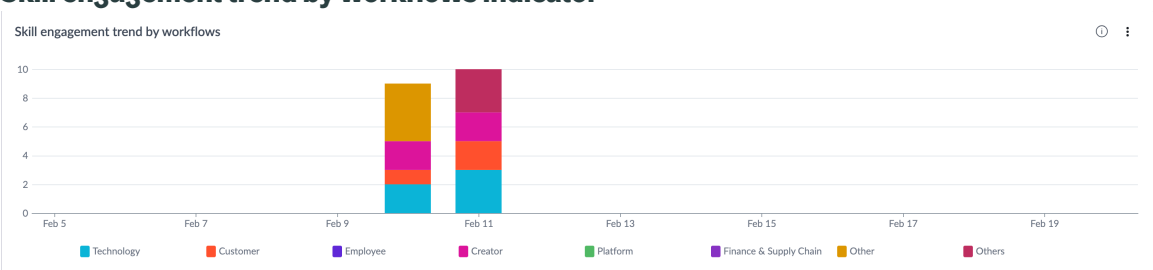
Daily unique users engaging with the skill indicator



Skill engagement trend by workflows

This area of the dashboard shows the skill usage across workflows in a bar chart for the selected date range. The visualization is interactive. Hover over the bars to see the number of times the skill was used in each of the workflows.

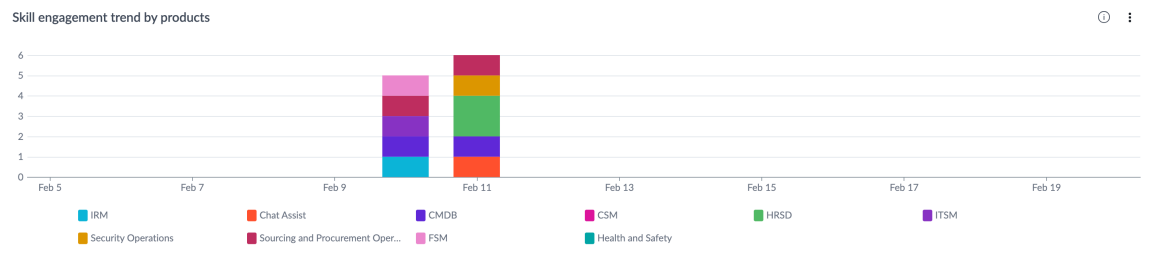
Skill engagement trend by workflows indicator



Skill engagement trend by products

This area of the dashboard shows the skill usage across Now Assist products in a bar chart for the selected date range. The visualization is interactive. Hover over the bars to see the number of times the skill was used in each of the products.

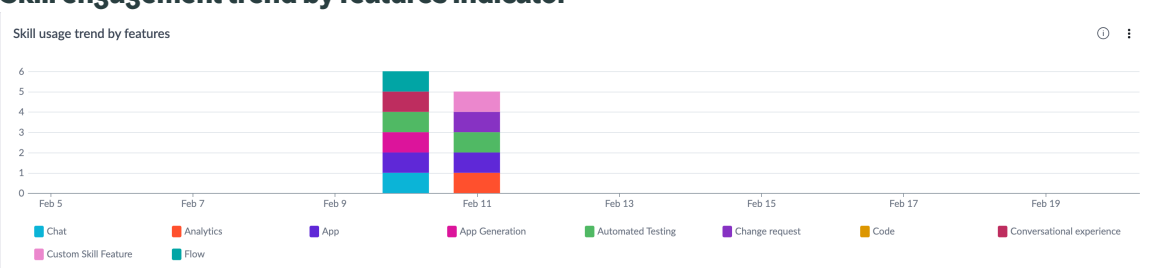
Skill engagement trend by products indicator



Skill engagement trend by features

This area of the dashboard shows the skill usage across Now Assist features in a bar chart for the selected date range. The visualization is interactive. Hover over the bars to see the number of times the skill was used in each of the features.

Skill engagement trend by features indicator

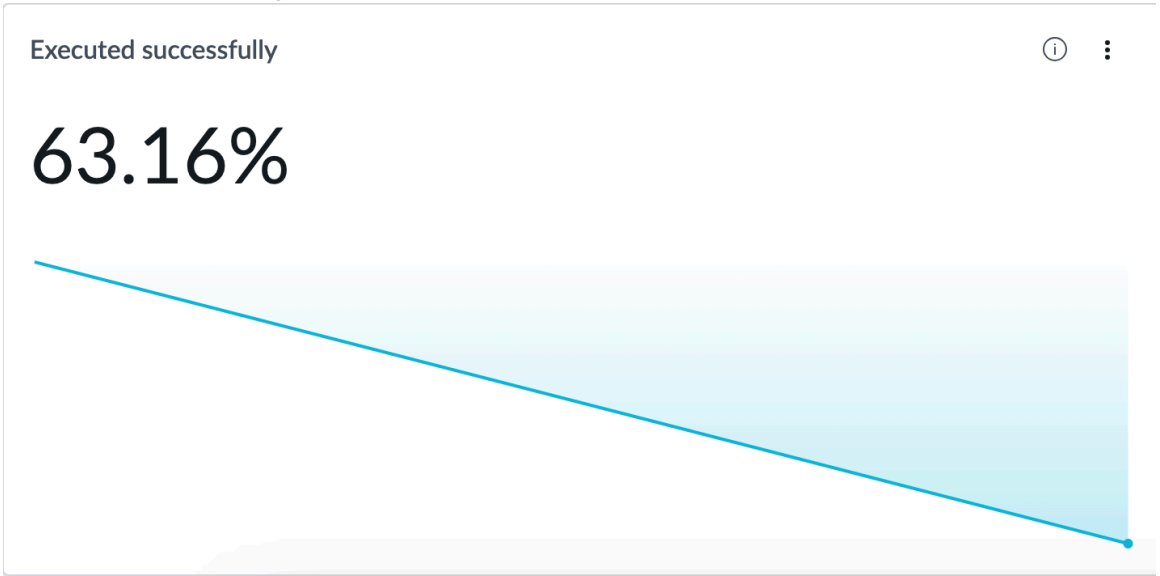


© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

Executed successfully

This area of the dashboard shows the acceptance rate of the selected skill based on user feedback. The percentage is calculated using the formula: (Total number of accepted skill executions/Total number of skill executions) x 100.

Executed successfully indicator

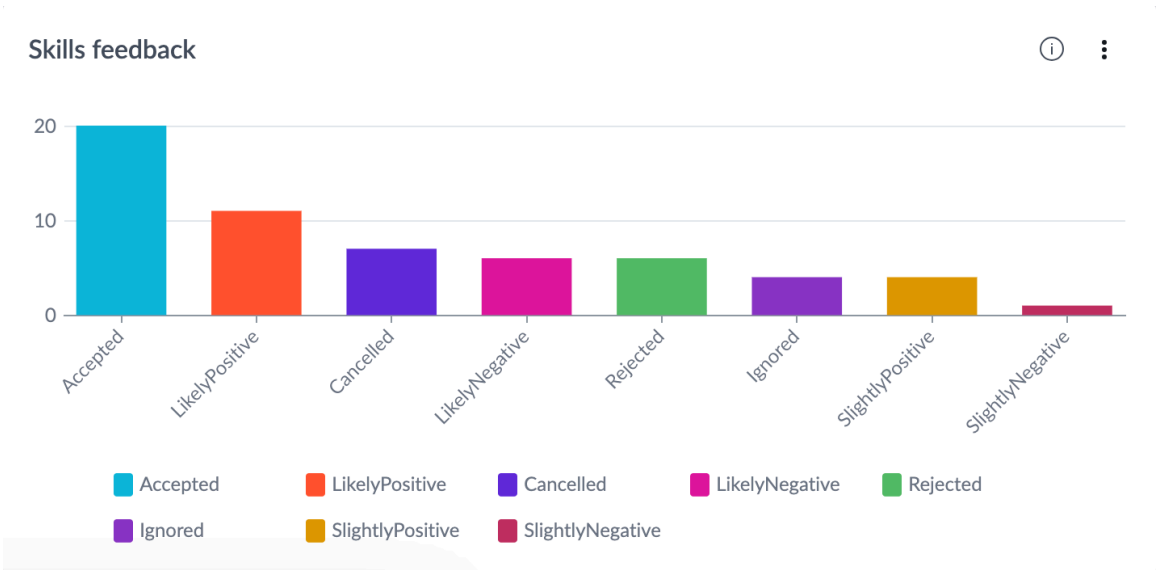


Skills feedback

This area of the dashboard shows the feedback recorded for each execution of the selected skill. User feedback is categorized into the following:

- Accepted (edited & non-edited): The user accepted the skill output with or without further edits to the output.
- Rejected: The user rejected the skill output.
- Canceled: The user canceled the skill execution.
- Ignored: The user didn't take any action based on the skill output.

Skills feedback indicator



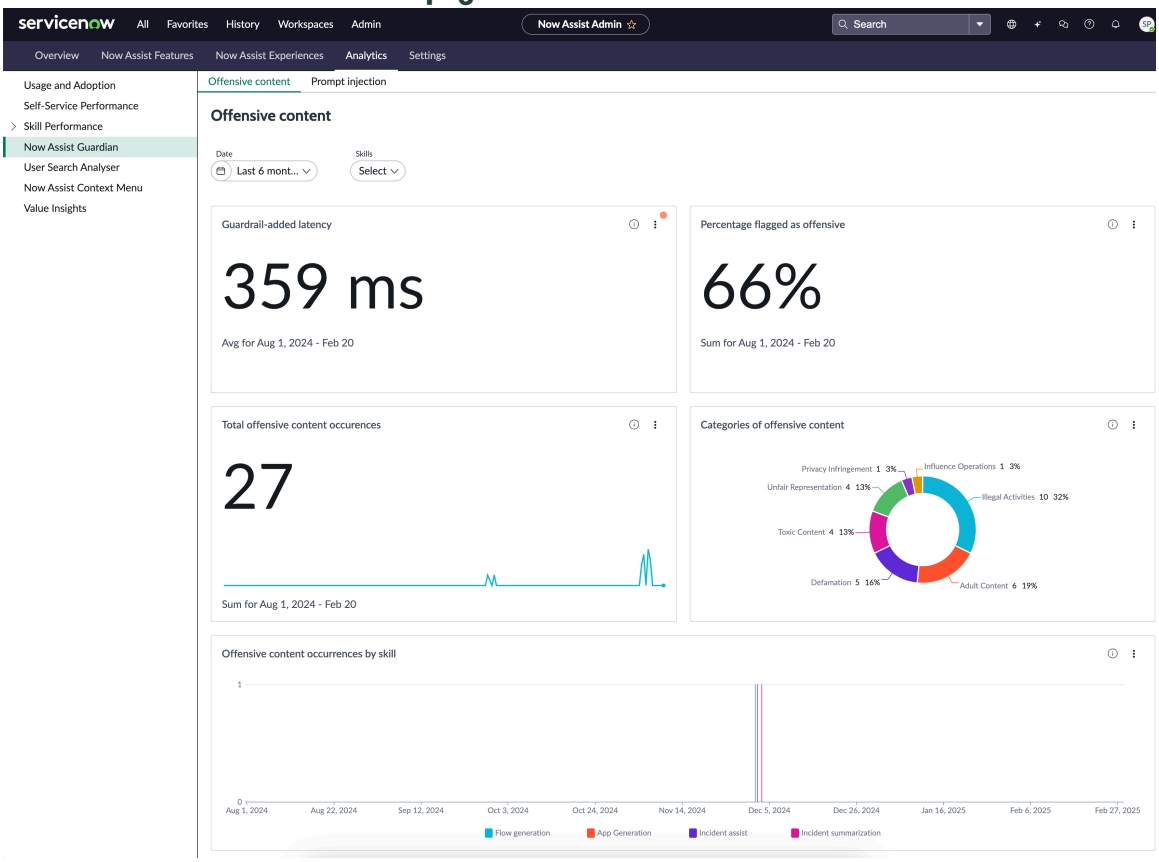
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

Now Assist Guardian analytics

Monitor the performance of guardrails enabled through Now Assist Guardian.

The Now Assist Guardian analytics dashboard helps admins monitor and evaluate the effectiveness of offensive content and prompt injection guardrails in tracking and analyzing requests sent to large language models (LLM) and their responses.

Now Assist Guardian dashboard page



The indicators on the Now Assist Guardian dashboard page provide the following insights.

- Average latency as a result of active offensive content and prompt injection guardrails. High latency could mean increased guardrail activity in the period.
- Count and percentage of offensive content and prompt injection occurrences.
- Skills where offensive content and prompt injection occurrences were detected.

Apply the filters on the dashboard to view guardrail activity for skills in a date range. See Now Assist Analytics dashboard indicator details for information on the data and calculations behind each indicator.

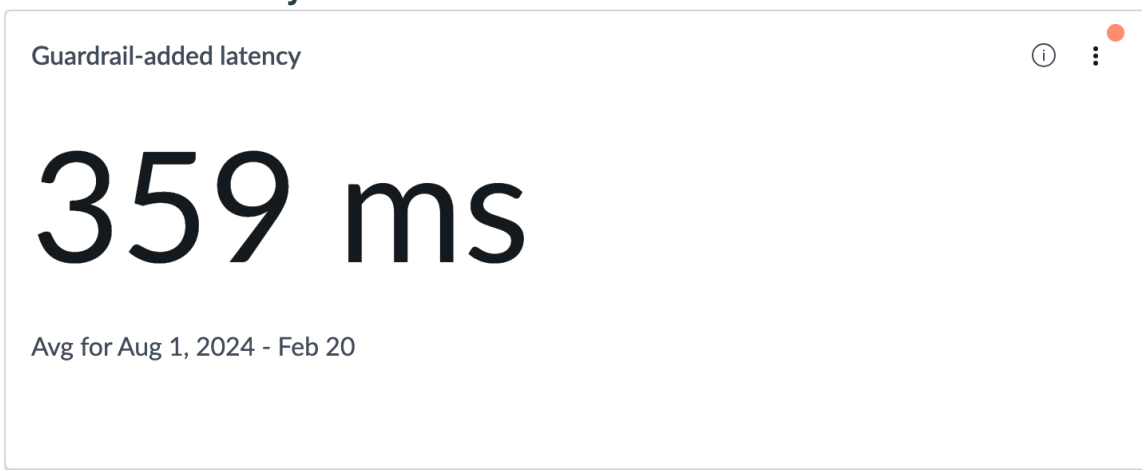
Offensive content indicators

Guardrail-added latency

This area of the dashboard shows the average latency as a result of the active offensive content guardrail for the selected skills and date range.

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

Guardrail-added latency indicator



Percentage flagged as offensive

This area of the dashboard shows the percentage of requests and responses to and from the LLM service that are flagged for offensive content.

Percentage flagged as offensive indicator

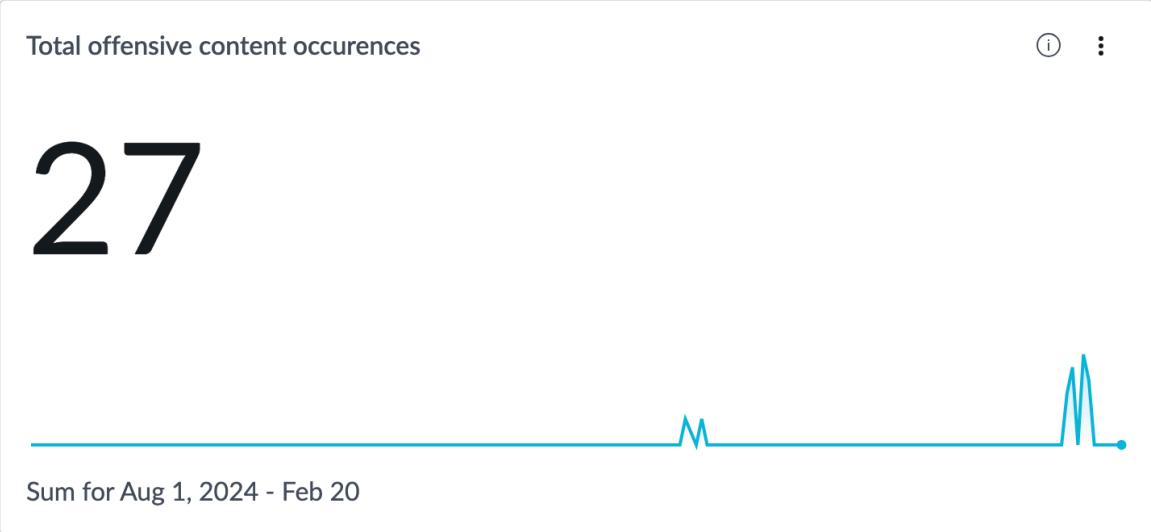


Total offensive content occurrences

This area of the dashboard shows the total number of offensive content occurrences for the selected skills and date range.

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

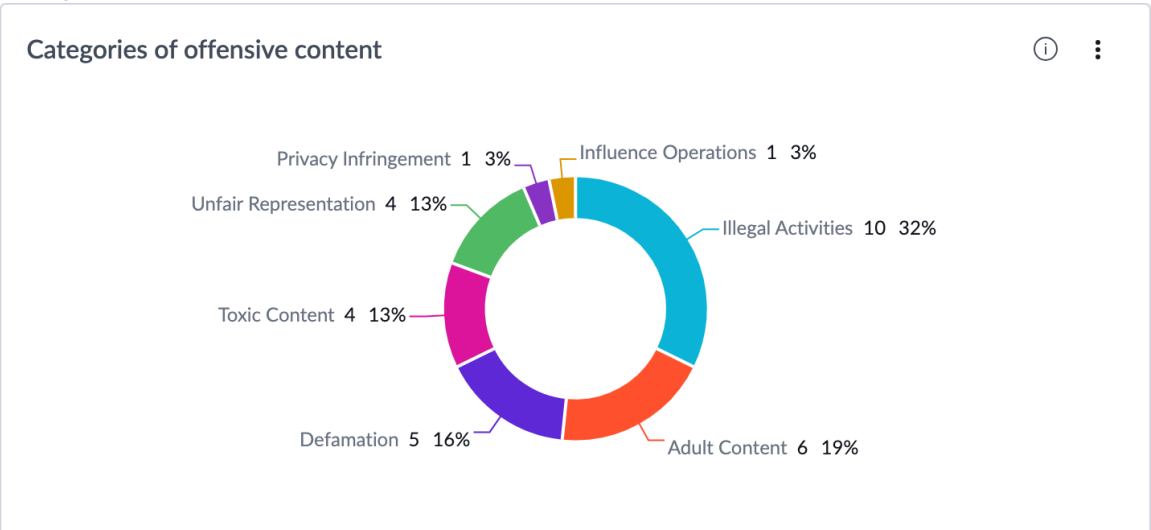
Total offensive content occurrences indicator



Categories of offensive content

This area of the dashboard shows a breakdown of offensive content occurrences by the categories. If content is deemed to be offensive under more than one category, for example, toxic and defamatory, the occurrence is counted individually toward both the categories. For more information on offensive content categories, see Now Assist Guardian.

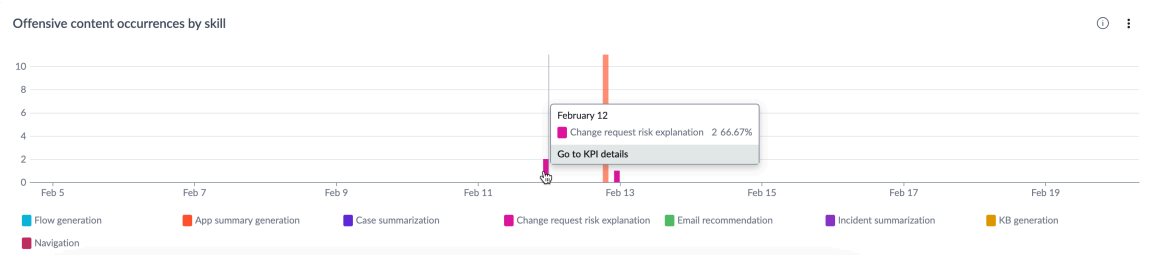
Categories of offensive content indicator



Offensive content occurrences by skill

This area of the dashboard shows the number of offensive content occurrences over time by the skills in which the content is detected.

Offensive content occurrences by skill indicator



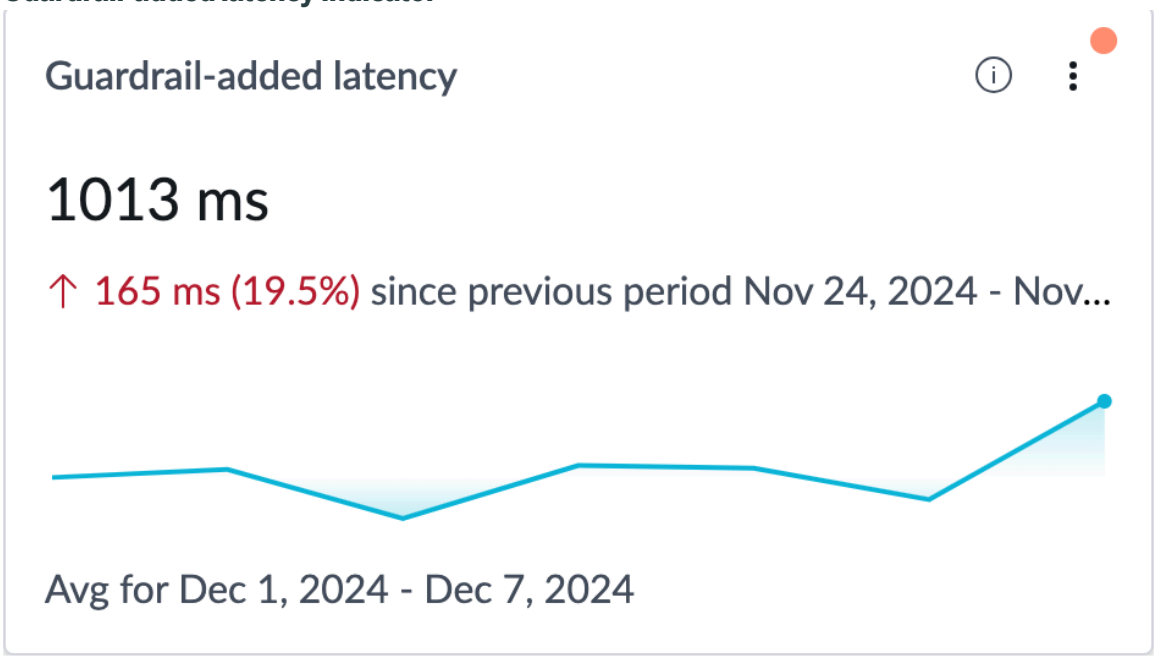
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

Prompt injection indicators

Guardrail-added latency

This area of the dashboard shows the average latency as a result of the active prompt injection guardrail for the selected skills and date range.

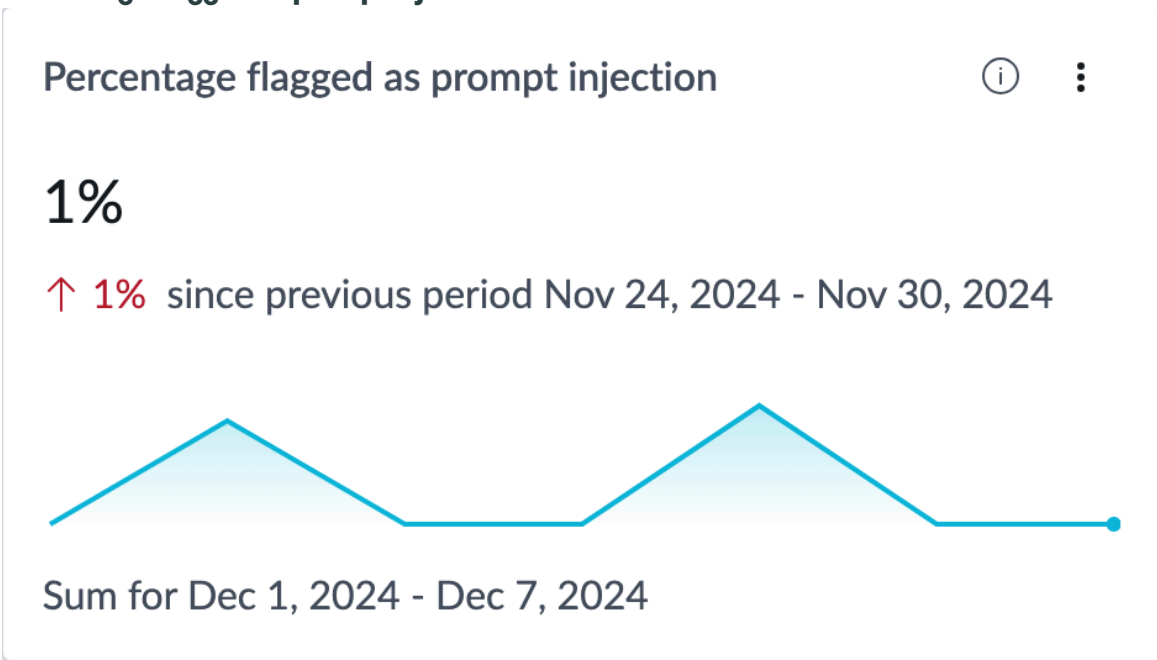
Guardrail-added latency indicator



Percentage flagged as prompt injection

This area of the dashboard shows the percentage of requests and responses to and from the LLM service that are flagged for offensive content.

Percentage flagged as prompt injection indicator



Total prompt injection occurrences

This area of the dashboard shows the total number of offensive content occurrences for the selected skills and date range.

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

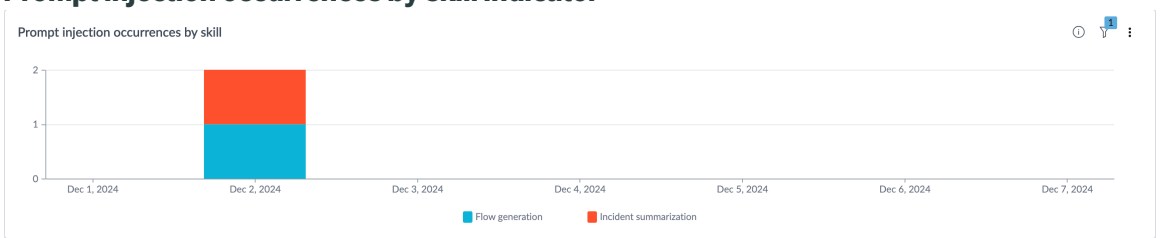
Total prompt injection occurrences indicator



Prompt injection occurrences by skill

This area of the dashboard shows the number of prompt injection occurrences over time by the skills where prompt injection attempts were detected.

Prompt injection occurrences by skill indicator



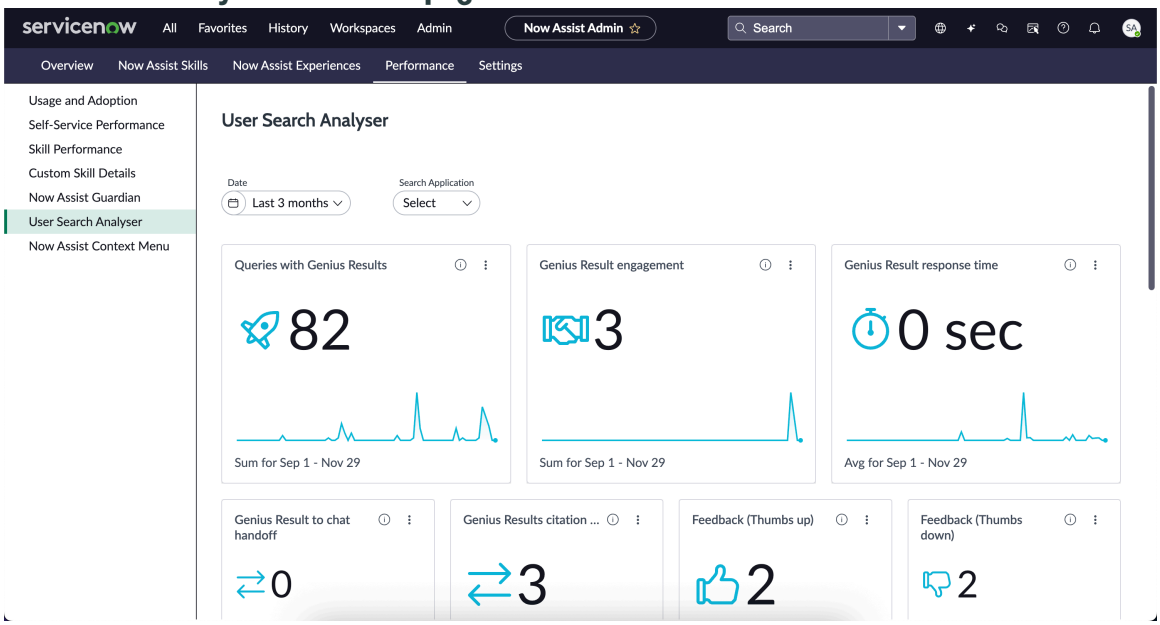
User search analyzer

Gain insights into user search queries and results provided by Now Assist.

The User search analyzer dashboard page contains indicators that help admins understand the effectiveness of search in enhancing the self-service experience. Equipped with insights from the dashboard, Knowledge admins can improve Knowledge content and availability for search.

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

User search analyzer dashboard page



The indicators on the User search analyzer dashboard page provide the following insights. See Now Assist Analytics dashboard indicator details for information on the data and calculations behind each indicator.

- Search queries that yielded Knowledge Base articles and catalog items as Genius Result.
- Distribution of search queries by the source that they originated from, for example, Now Assist in Virtual Agent, Service Portal.
- List of the most common queries that did or didn't yield genius results.

User search analyzer indicators

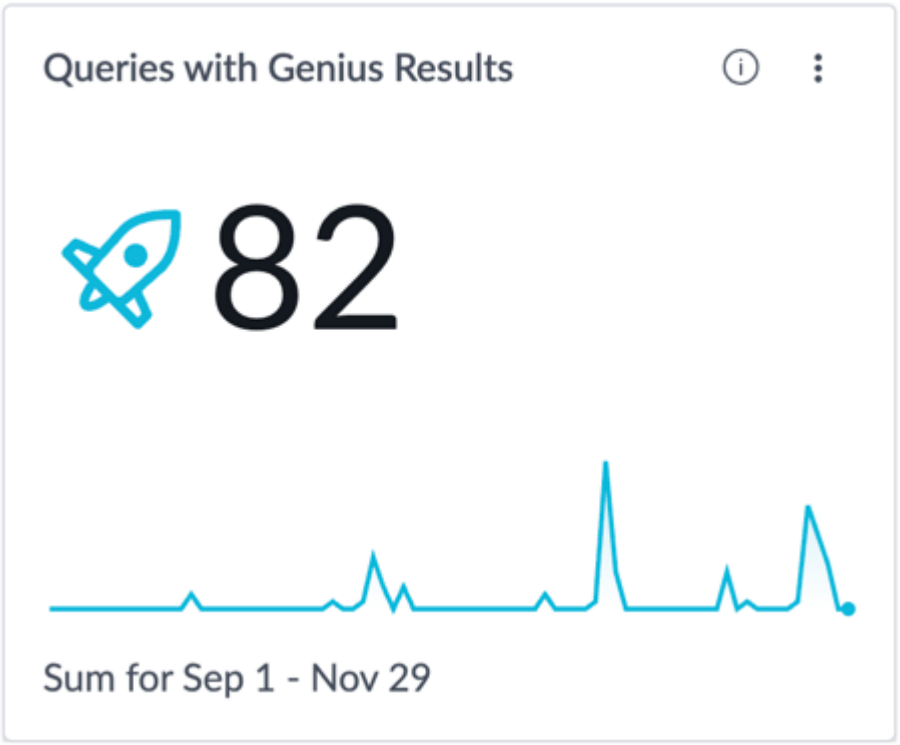
Queries with Genius Results

This area of the dashboard total number of search queries that had a Genius Result returned by Now Assist.

Note: A query can have more than one genius result.

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

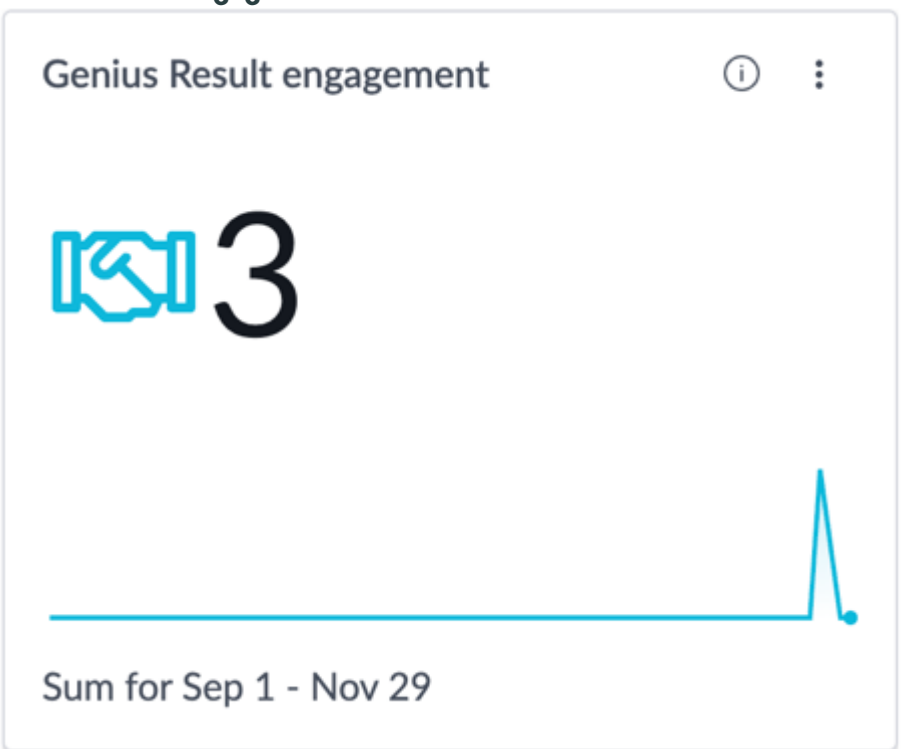
Queries with Genius Results indicator



Genius Result Engagement

This area of the dashboard shows the number of genius results returned by Now Assist that covers the interactions like "Ask a follow up" or "Show full answer", and clicking on citations within the result card.

Genius Result engagement indicator

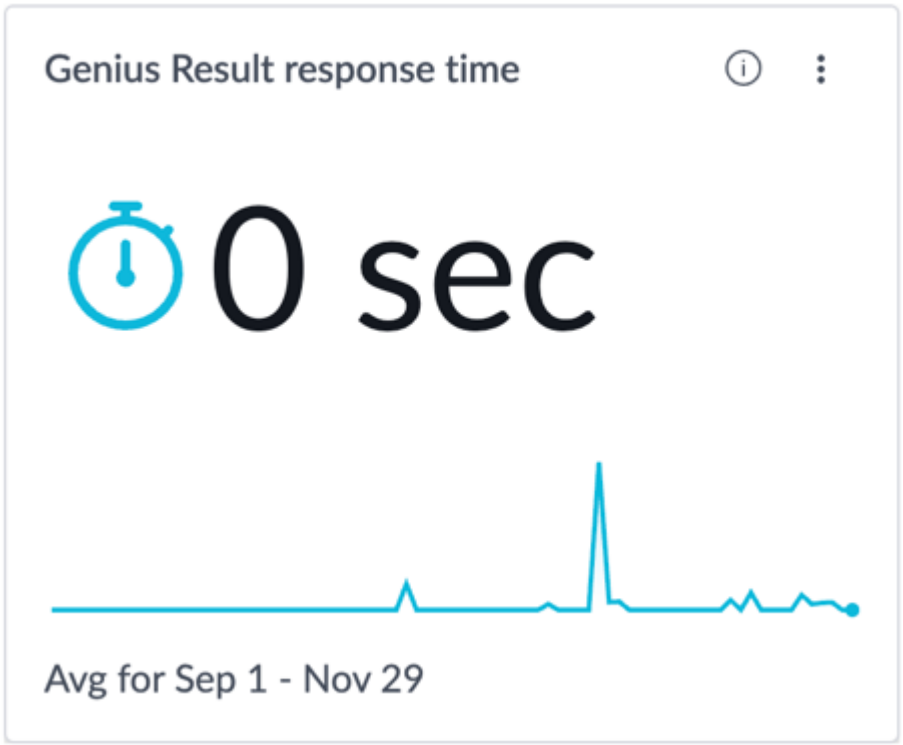


Genius Result response time

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

This area of the dashboard shows the average response time taken for a Genius Result to load.

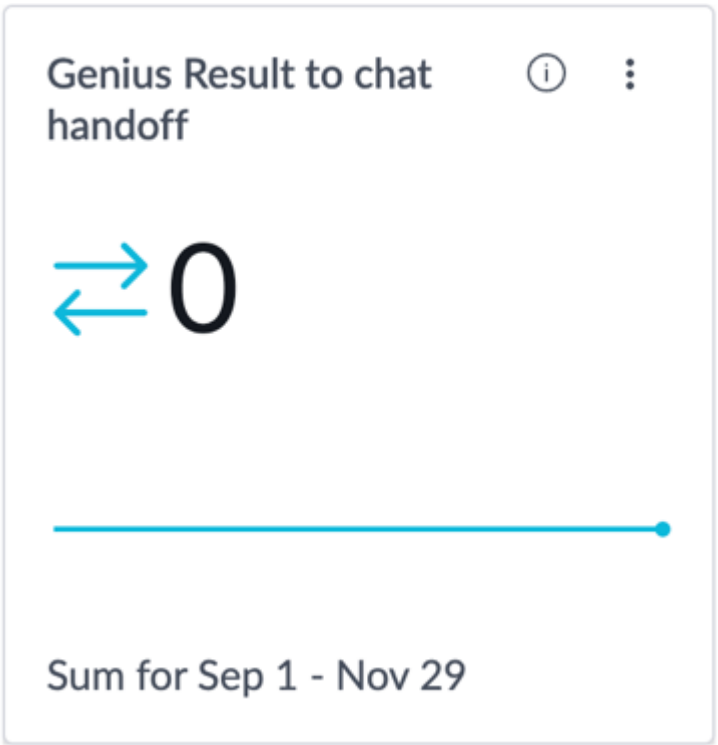
Genius Result response time indicator



Genius Result to chat handoff

This area of the dashboard shows the number of times users clicked the **Ask a follow-up** button within the synthesized Genius Result.

Genius Result to chat handoff indicator

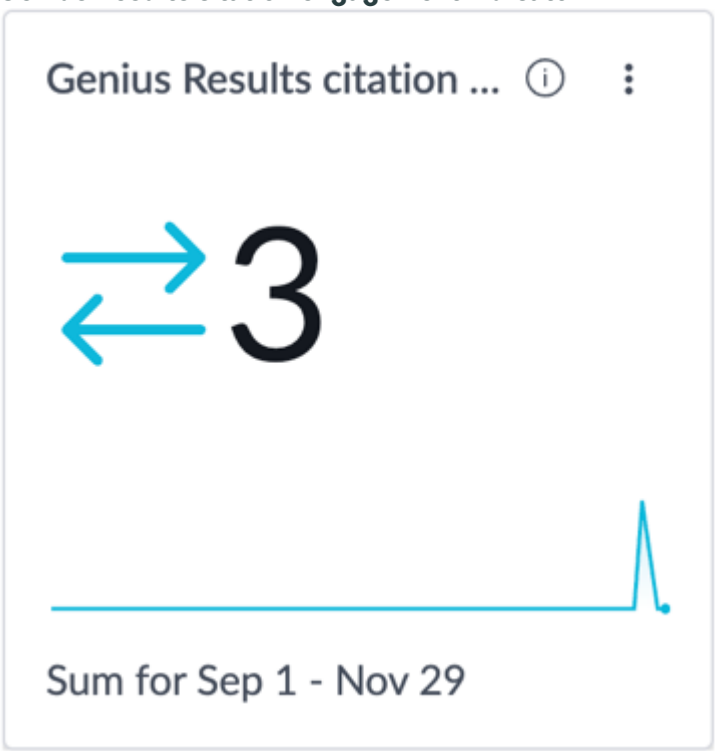


Genius Results citation engagement

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

This area of the dashboard shows the number of interactions with citations within Genius Result returned by Now Assist.

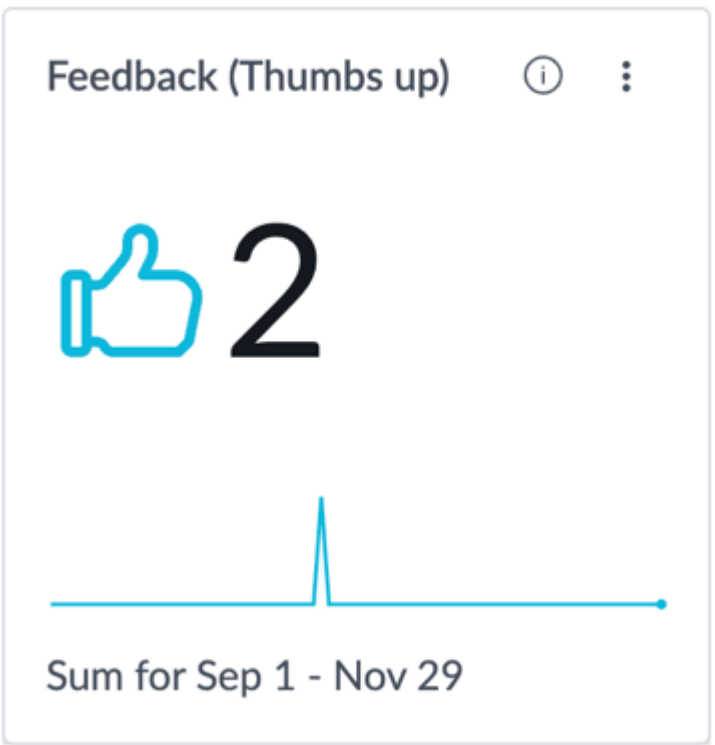
Genius Results citation engagement indicator



Feedback (Thumbs up)

This area of the dashboard shows the count for positive feedback for Genius Result returned by Now Assist.

Feedback (Thumbs up) indicator

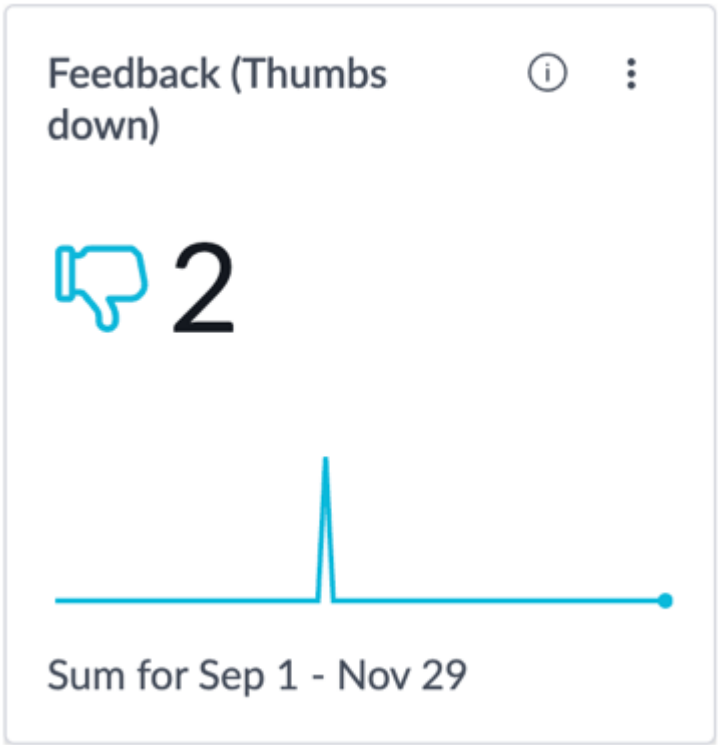


Feedback (Thumbs down)

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

This area of the dashboard shows the count for negative feedback for Genius Result returned by Now Assist.

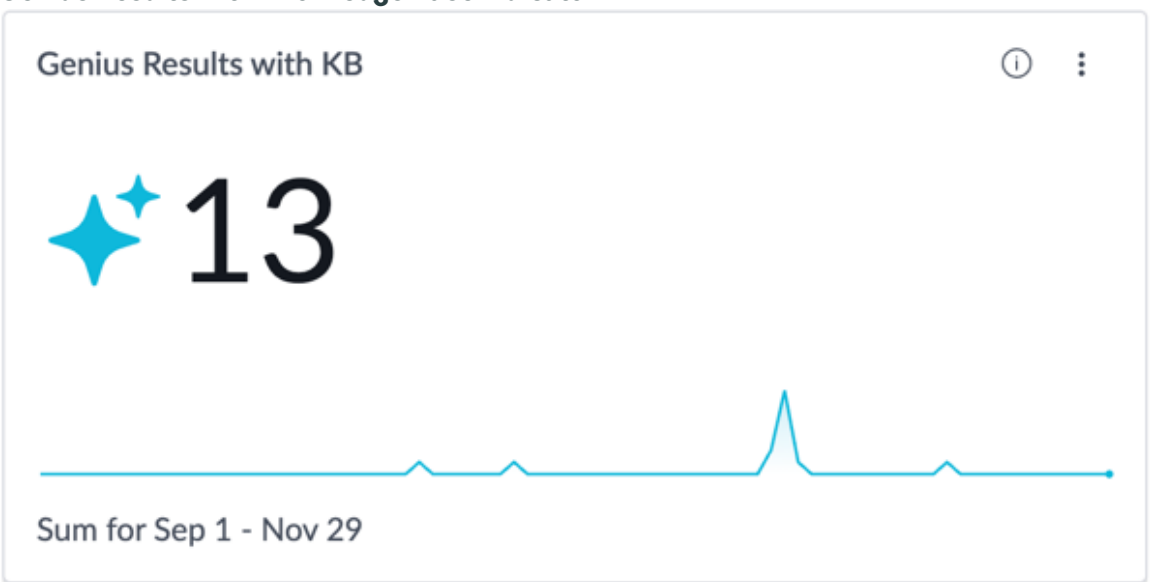
Feedback (Thumbs down) indicator



Genius Results with Knowledge Base

This area of the dashboard shows the number of search queries where Now Assist returned Knowledge Base articles as Genius Result for the users to review.

Genius Results with Knowledge Base indicator

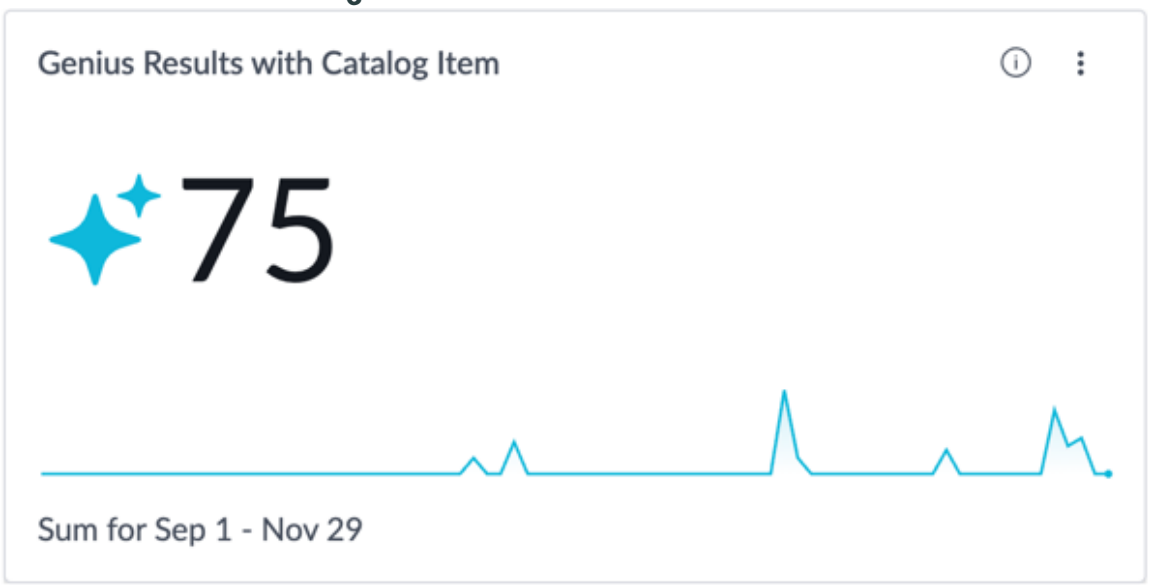


Genius Results with Catalog Item

This area of the dashboard shows the number of search queries where Now Assist returned catalog items as Genius Result for the users to review.

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

Genius Results with Catalog Item indicator

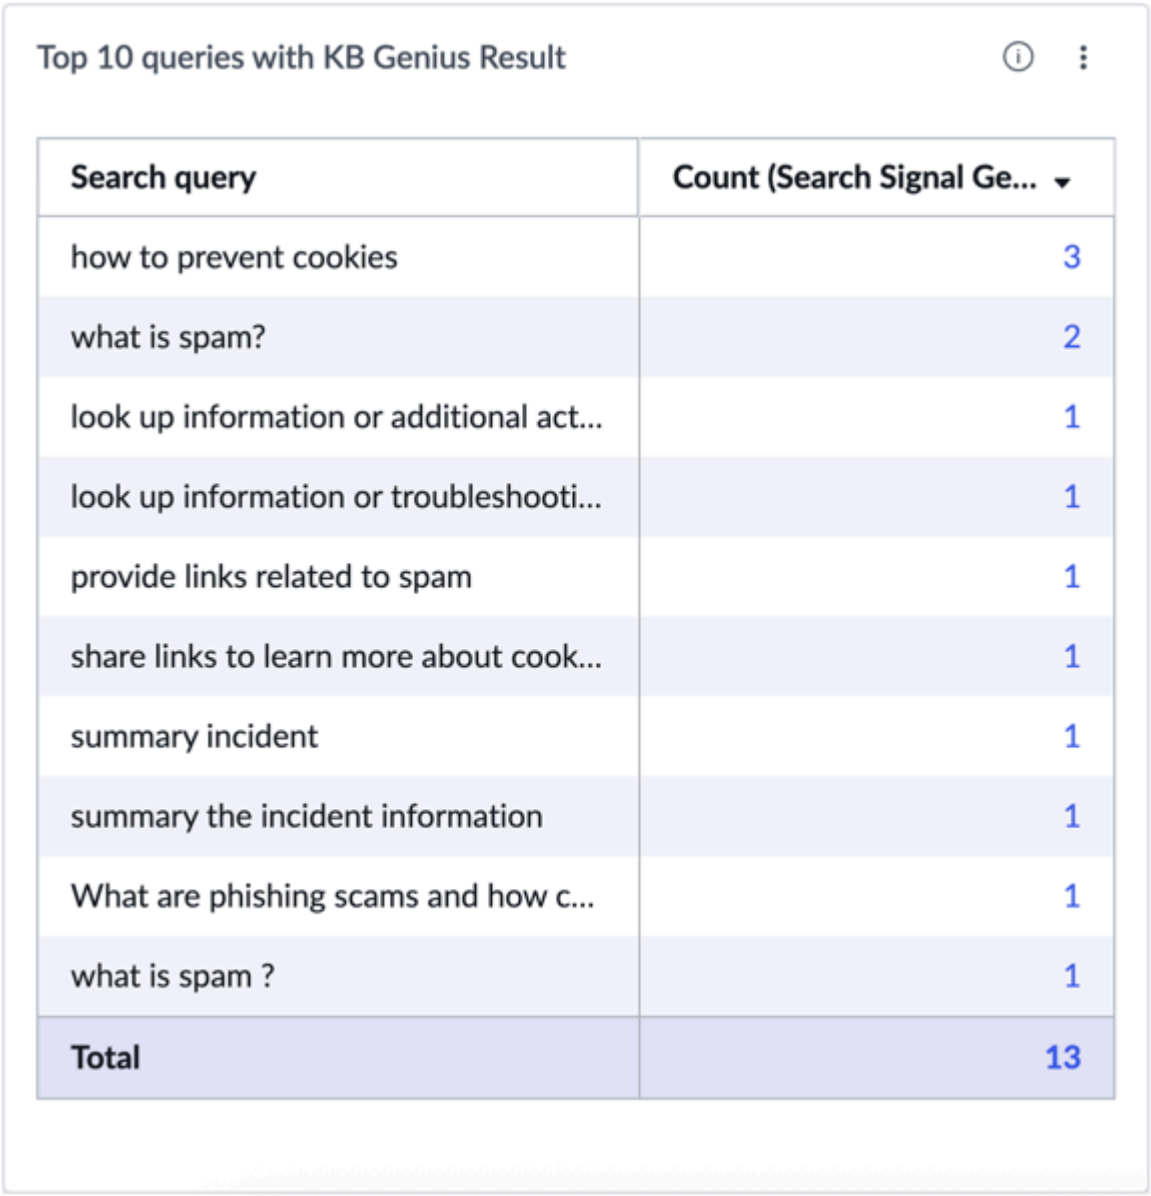


Top 10 queries with Knowledge Base Genius Result

This area of the dashboard shows the top ten search queries where Now Assist returned Knowledge Base as Genius Result, with respective query count.

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

Top 10 queries with Knowledge Base Genius Result



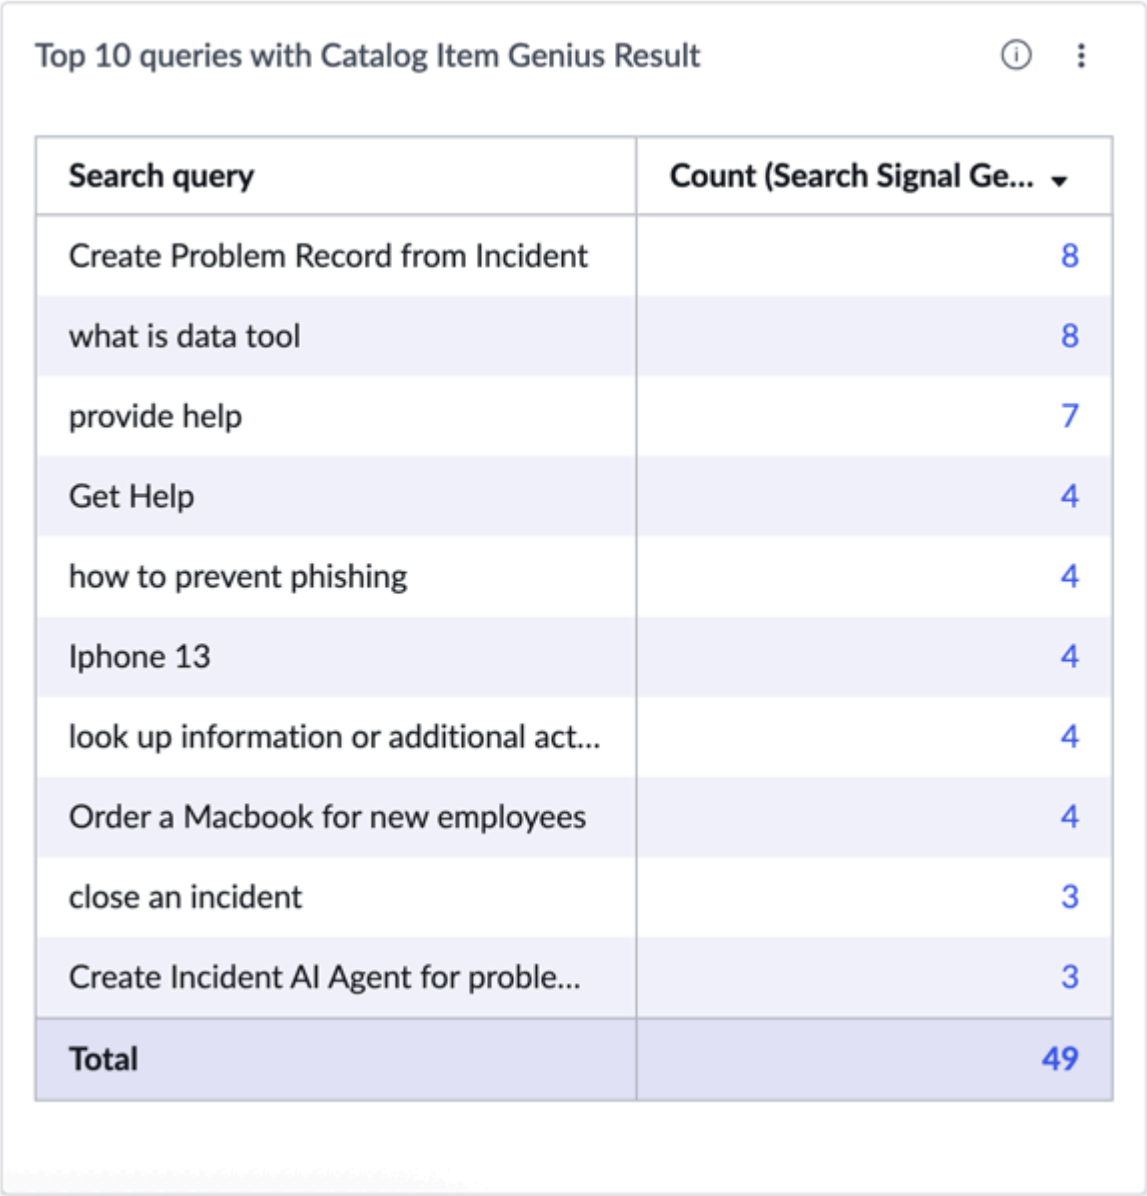
Search query	Count (Search Signal Ge...
how to prevent cookies	3
what is spam?	2
look up information or additional act...	1
look up information or troubleshooti...	1
provide links related to spam	1
share links to learn more about cook...	1
summary incident	1
summary the incident information	1
What are phishing scams and how c...	1
what is spam ?	1
Total	**13**

Top 10 queries with Catalog Item Genius Result

This area of the dashboard shows the top ten search queries where Now Assist returned catalog items as Genius Result, with respective query count.

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

Top 10 queries with Catalog Item Genius Result indicator



Search query	Count (Search Signal Ge...
Create Problem Record from Incident	8
what is data tool	8
provide help	7
Get Help	4
how to prevent phishing	4
Iphone 13	4
look up information or additional act...	4
Order a Macbook for new employees	4
close an incident	3
Create Incident AI Agent for proble...	3
Total	**49**

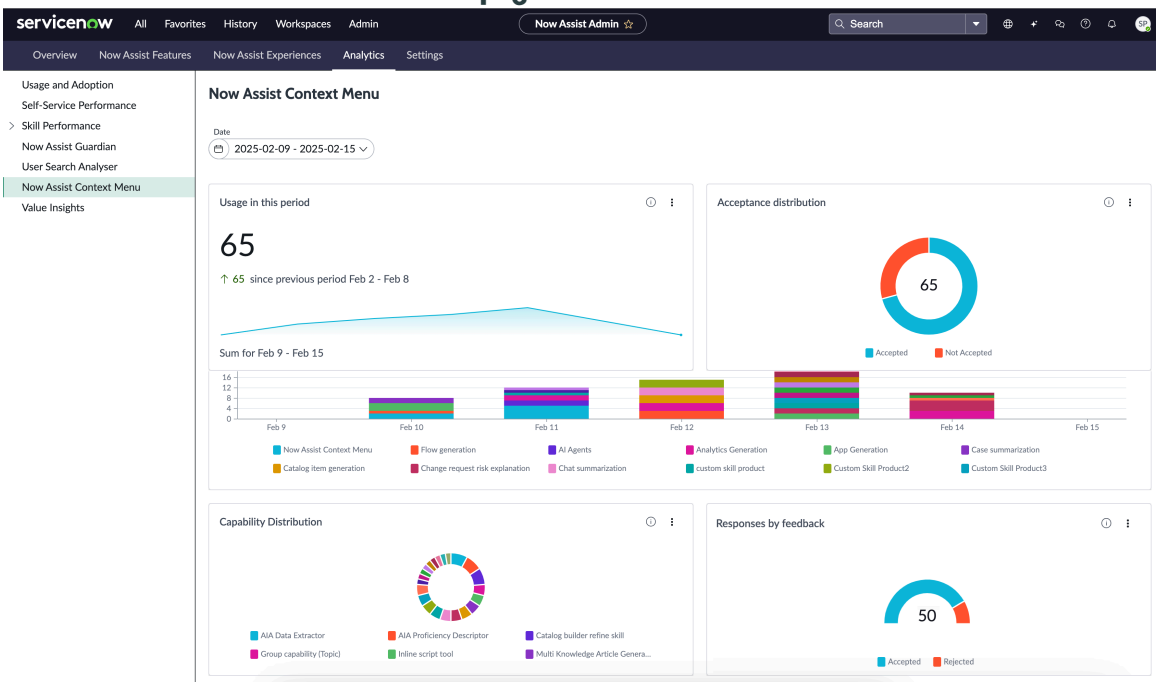
Now Assist context menu analytics

Monitor the usage and performance of the Now Assist context menu.

The Now Assist context menu dashboard helps admins to evaluate the effectiveness of context menu actions in assisting agents with summarizing, creating, and editing emails and chat replies.

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

Now Assist context menu dashboard page



The indicators on the Now Assist context menu dashboard page provide the following insights. See Now Assist Analytics dashboard indicator details for information on the data and calculations behind each indicator.

- Count and distribution of context menu actions in a date range. This indicator reveals the scale of usage, and the most and least used actions.
- Acceptance and feedback on context menu actions. This indicator reveals actions that were accepted or rejected by the users with a positive or negative feedback.
- Usage trend chart reveals patterns in usage of context menu actions over a date range.

Now Assist context menu indicators

Usage in this period

This area of the dashboard shows the total number of Now Assist context menu actions used in the selected date range.

Usage in this period indicator



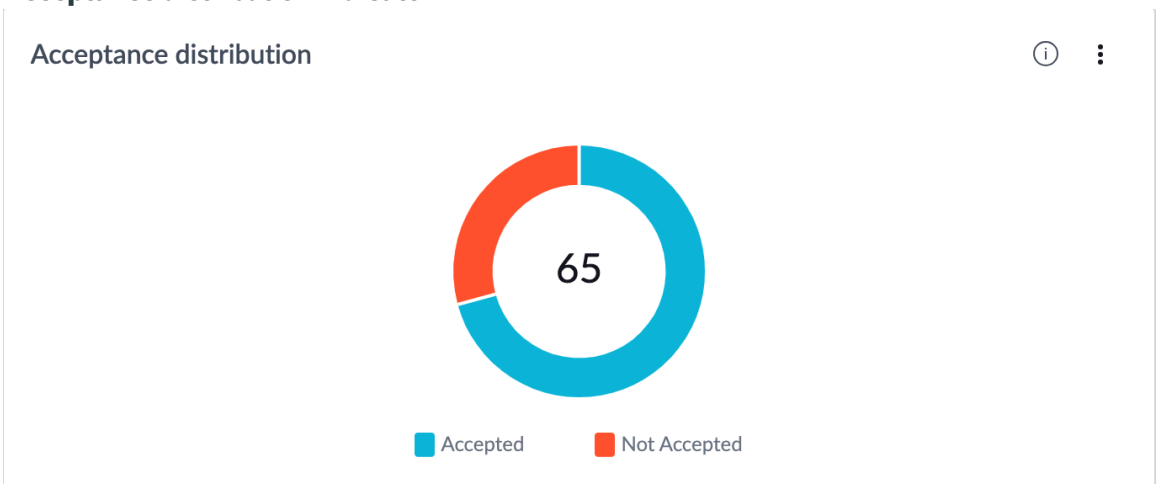
Acceptance distribution

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

This area of the dashboard shows the distribution of Now Assist context menu actions by their acceptance.

- Accepted: The user has inserted the context menu response into their workspace.
- Not Accepted: The user has closed the window or dialog box containing the context menu response.

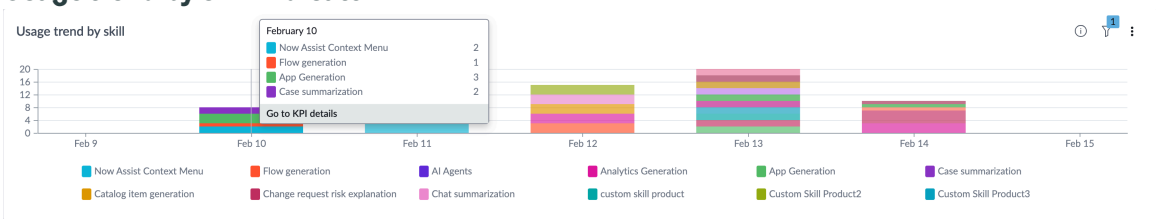
Acceptance distribution indicator



Usage trend by skill

This area of the dashboard shows the usage trend of Now Assist context menu actions in the selected date range.

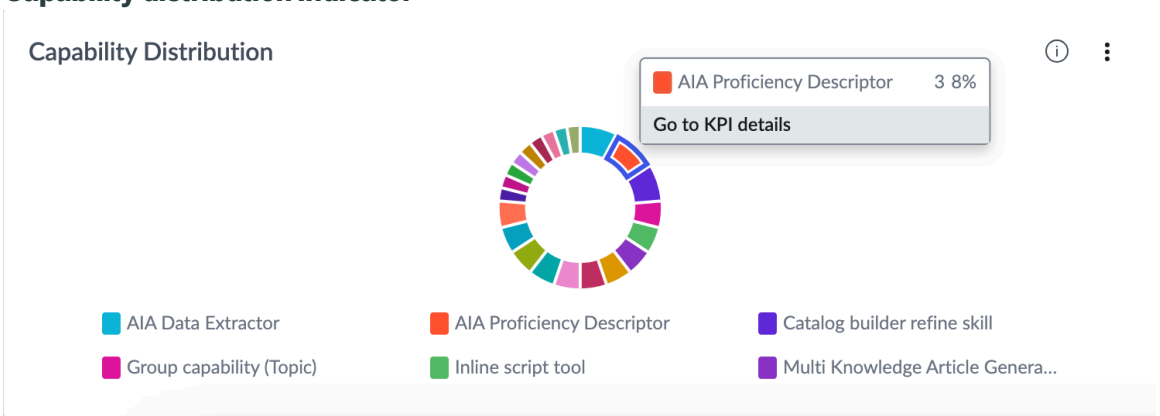
Usage trend by skill indicator



Capability distribution

This area of the dashboard shows the usage distribution of various Now Assist context menu actions in the selected date range.

Capability distribution indicator



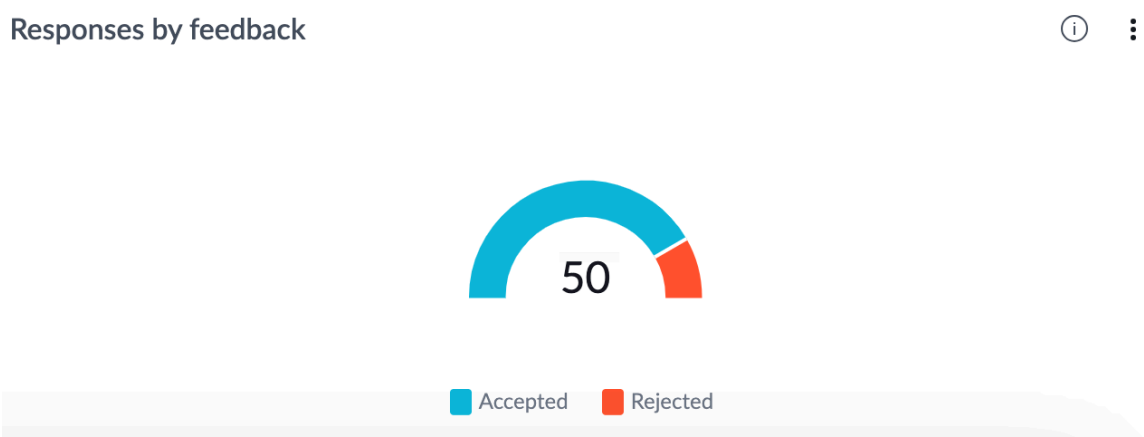
Responses by feedback

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

This area of the dashboard shows a breakdown of Now Assist context menu responses by the user feedback received.

- Accepted: The user gave a thumbs up on the context menu response.
- Rejected: The user gave a thumbs down on the context menu response.

Responses by feedback indicator



Now Assist Analytics reference

Now Assist Analytics reference topics include information about user roles and details of the indicators on the dashboard.

Now Assist Analytics roles

Now Assist Analytics requires the following roles to view and manage the dashboard functionality.

Now Assist Analytics Viewer [sn_na_analytics.viewer]

Users with the Now Assist Analytics Viewer role can view the Now Assist Analytics dashboard in the Now Assist Admin console, and have read access to data source tables.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups that this role is assigned to by default.

None.

Special considerations

None.

Now Assist Analytics Admin [sn_na_analytics.admin]

Users with Now Assist Analytics Admin role can view the Now Assist Analytics dashboard in the Now Assist Admin console, and read and write to some data source tables.

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

Contains Roles

List of roles contained within the role.

Now Assist Analytics Viewer [sn_na_analytics.viewer].

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.

Now Assist Analytics dashboard indicator details

Indicator details help you understand the data and calculations behind an indicator that is presented in the form of a visualization on the dashboard.

Now Assist Analytics indicators contain the following details: indicator type, data source, calculation, available breakdowns, unit, and so on.

To access these indicators, navigate to **Platform Analytics Administration** > **Indicators**. You must have the Now Assist Analytics Admin [sn_na_analytics_admin] role to access the indicators.

These indicators collect data at a daily frequency. Data is only available for dates before the current date. If you want to see results from the current day, you must wait until the next day.

Note: The Generative AI Usage Log [sys_gen_ai_usage_log] table is maint-only.

Usage and adoption indicator details

Visualizati	Indicator type	Indicator source table	Calculatic	Available breakdowr	Frequency	Unit	Precision
Total Now Assist actions	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of all actions	By Gen AI Feature, By Generative AI Skill Execution Modality, By Skill Family, By Skills Config, By Skills Config, By Workflow	Daily	#	0
Daily Now Assist actions	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of daily actions	By Gen AI Feature, By	Daily	#	0

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

Visualizati	Indicator type	Indicator source table	Calculatic	Available breakdowr	Frequency	Unit	Precision
				Generative AI Skill Execution Modality, By Skill Family, By Skills Config, By Skills Config, By Workflow			
Average daily unique users engaging with Now Assist	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Average of daily unique users	By Gen AI Feature, By Skills Config, By Skill Family, By Generative AI Skill Execution Modality	Daily	#	0
Daily unique users engaging with Now Assist	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of daily unique users	By Gen AI Feature, By Skills Config, By Skill Family, By Generative AI Skill Execution Modality	Daily	#	0
Skill group distribution	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of daily skill executions grouped By Gen AI Feature, By Generative AI Skill Execution Modality, By Skill Family, By	By Gen AI Feature, By Generative AI Skill Execution Modality, By Skill Family, By Skills Config, By Skills Config, By Workflow	Daily	#	0

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

Visualizati	Indicator type	Indicator source table	Calculatic	Available breakdowr	Frequency	Unit	Precision
			Skills Config, By Skills Config				
Daily usage comparison by workflow	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of actions grouped by workflows	By Gen AI Feature, By Generative AI Skill Execution Modality, By Skill Family, By Skills Config, By Skills Config, By Workflow	Daily	#	0
Skill engagement trend	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of actions grouped by skill	By Gen AI Feature, By Generative AI Skill Execution Modality, By Skill Family, By Skills Config, By Skills Config, By Workflow	Daily	#	0
Departments with highest usage	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of actions grouped by user department and sorted by descending order	By Department, By Skills Config	Daily	#	0
Now Assist actions comparison	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of actions grouped	By Department, By Skills Config	Daily	#	0

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

Visualizati	Indicator type	Indicator source table	Calculatic	Available breakdowr	Frequency	Unit	Precision
by user department			by user department				
Feedback details	Automated	Gen AI Log Metadata[sys_gen_ai_log_metadata]	Count of actions grouped by Skill Family	By Skill Family, By Skills Config	Daily	#	0
	Automated	Gen AI Log Metadata[sys_gen_ai_log_metadata]	Count of actions with feedback grouped by Skill Family	By Feedback, By Skill Family, By Skills Config	Daily	#	0
	Formula	Gen AI Log Metadata[sys_gen_ai_log_metadata]	(Count of actions with feedback / Total count of actions) x 100 grouped by Skill Family	By Skill Family, By Skills Config	Daily	%	2
	Automated	Gen AI Log Metadata[sys_gen_ai_log_metadata]	Count of actions where By Feedback is Accepted grouped by Skill Family	By Skill Family, By Skills Config, By Feedback	Daily	#	0
	Formula	Gen AI Log Metadata[sys_gen_ai_log_metadata]	(Count of actions where By Feedback is Accepted / Total count of actions with	By Skills Config, By Skill Family	Daily	%	2

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

Visualizati	Indicator type	Indicator source table	Calculatic	Available breakdowr	Frequency	Unit	Precision
			feedback) x 100 grouped by Skill Family				
Error details	Automated	Gen AI Log Metadata[sys_gen_ai_log_metadata]	Count of actions grouped by Skill Family	By Skill Family, By Skills Config	Daily	#	0
	Automated	Gen AI Log Metadata[sys_gen_ai_log_metadata]	Count of actions with error status grouped by Skill Family	By Gen AI Log Metadata Status , By Skill Family, By Skills Config	Daily	#	0
	Formula	Gen AI Log Metadata[sys_gen_ai_log_metadata]	(Count of actions with error status/ Total count of generative AI metadata records) x 100 grouped by Skill Family	By Skill Family, By Skills Config	Daily	%	2

Skill performance indicator details

Visualizatio	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
Skill engagement trend	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of skill executions grouped by skill	By Gen AI Feature, By Generative AI Skill Execution Modality, By Skill	Daily	#	0

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

Visualizatio	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
				Family, By Skills Config, By Workflow			
Number of actions	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of all skill executions	By Gen AI Feature, By Generative AI Skill Execution Modality, By Skill Family, By Skills Config, By Workflow	Daily	#	0
Total daily active users	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of daily unique users	By Gen AI Feature, By Skills Config, By Skill Family, By Generative AI Skill Execution Modality	Daily	#	0
Total daily active users by skills	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of daily unique users grouped by Skills Config	By Gen AI Feature, By Skills Config, By Skill Family, By Generative AI Skill Execution Modality	Daily	#	0

Skill details indicator details

Visualizatio	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
Skill engagement trend	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of skill executions grouped by Skills Config	By Gen AI Feature, By Generative AI Skill Execution Modality,	Daily	#	0

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

Visualizatio	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
				By Skill Family, By Skills Config, By Workflow			
Total skill actions	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of skill executions	By Gen AI Feature, By Generative AI Skill Execution Modality, By Skill Family, By Skills Config, By Workflow	Daily	#	0
Accepted skill actions	Automated	Gen AI Log Metadata[sys_gen_ai_log_metadata]	Count of skill executions where Feedback is Accepted	By Feedback, By metadata response accepted without edits, By Skills Config, By Skill Family	Daily	#	0
Acceptance rate	Formula	Gen AI Log Metadata[sys_gen_ai_log_metadata]	(Count of skill executions where Feedback is Accepted/ Count of skill executions) x100	By Skills Config	Daily	%	0
Daily active users	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of daily unique users	By Gen AI Feature, By Skills Config, By Skill Family, By Generative AI Skill Execution Modality	Daily	#	0

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

Custom skills indicator details

Visualization	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
Skill engagement trend by workflows	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of custom skill executions grouped by workflows	By Custom Skill (Workflow Type), By Custom Skill (Feature Type), By Custom Skill (Product Type), By Custom Skill Completion Status, By Skills Config	Daily		0
Skill engagement trend by products	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of custom skill executions grouped by products	By Custom Skill (Workflow Type), By Custom Skill (Feature Type), By Custom Skill (Product Type), By Custom Skill Completion Status, By Skills Config	Daily		0
Skill engagement trend by features	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of custom skill executions grouped by features	By Custom Skill (Workflow Type), By Custom Skill (Feature Type), By Custom Skill	Daily		0

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

Visualizatio	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
				(Product Type), By Custom Skill Completion Status, By Skills Config			
Daily unique users engaging with the skill	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count of unique user_id executed custom skills	By Skills Config	Daily	#	0
Executed successfully	Formula	Generative AI Usage Log[sys_gen_ai_usage_log]	(Count of custom skill executions with status Completed/ Count of custom skill executions)x100	By Skills Config	Daily	%	2
Skills feedback	Automated	Generative AI Usage Log[sys_gen_ai_usage_log]	Count	By Feedback, By Skills Config	Daily	#	0

Now Assist Guardian offensive content guardrail indicator details

Visualizatio	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
Guardrail-added latency	Automated	Generative AI Metric[sys_generative_ai_metric]	Average time taken by the guardrail to evaluate offensive content	By Skills Config	Daily	Milliseconds	0
Percentage flagged as offensive	Formula	Generative AI Metric[sys_generative_ai_metric]	(Count of offensive content occurrences/ Total number	By Skills Config	Daily	%	0

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

Visualizatio	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
			of LLM calls for which the guardrail was enabled)x100				
Total offensive content occurrences	Automated	Generative AI Metric[sys_generative_ai_metric]	Count of offensive content occurrences	By Gen AI metric value, By Skills Config	Daily	#	0
Categories of offensive content	Automated	Generative AI Metric[sys_generative_ai_metric]	Count of offensive content occurrences grouped by categories	By Offensiveness Type, By Skills Config	Daily	#	0
Offensive content occurrences by skill	Automated	Generative AI Metric[sys_generative_ai_metric]	Count of offensive content occurrences grouped by skill	By Gen AI metric value, By Skills Config	Daily	#	0

Now Assist Guardian prompt injection guardrail indicator details

Visualizatio	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
Guardrail-added latency	Automated	Generative AI Metric[sys_generative_ai_metric]	Average time taken by the guardrail to evaluate prompt injection	By Skills Config	Daily	Milliseconds	0
Percentage flagged as prompt injection	Formula	Generative AI Metric[sys_generative_ai_metric]	(Count of prompt injection occurrences/ Total number of LLM calls for which the guardrail	By Skills Config	Daily	%	0

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

Visualizatio	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
			was enabled)x100				
Total prompt injection occurrences	Automated	Generative AI Metric[sys_gen_ai_metric]	Count of prompt injection occurrences	By Gen AI metric value, By Skills Config	Daily	#	0
Prompt injection occurrences by skill	Automated	Generative AI Metric[sys_gen_ai_metric]	Count of prompt injection occurrences grouped by skill	By Gen AI metric value, By Skills Config	Daily	#	0

User search analyzer indicator details

Visualizatio	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
Queries with Genius Results	Automated	Search Event[sys_search_event]	Count of search queries where Genius Result Displayed = true	AIS Search Application	Daily	#	0
Genius Result engagement	Automated	Genius Result Event Action[sys_search_genius_result_event_action]	Count of genius results where Card Type = Now Assist Multi-Content Response and Action Type = Show Full Answer or Action Type = Ask a Follow Up or Action Type = uxf_client_action	AIS Search Application	Daily	#	0

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

Visualization	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
Genius Result response time	Automated	Search Signal Event[sys_search_result_signal_event]	Average (Genius Result Response Time in Seconds)	AIS Search Application	Daily	Seconds	0
Genius Result to chat handoff	Automated	Genius Result Event Action[sys_search_genius_result_event_action]	Count of queries where Card Type = Now Assist Multi-Content Response and Action Type = Ask a Follow Up	AIS Search Application	Daily	#	0
Genius Results citation engagement	Automated	Genius Result Event Action[sys_search_genius_result_event_action]	Count of queries where Card Type = Now Assist Multi-Content Response and Action Type = uxf_client_action	AIS Search Application	Daily	#	0
Feedback (Thumbs up)	Automated	Genius Result Event Action[sys_search_genius_result_event_action]	Count of queries where Card Type = Now Assist Multi-Content Response and Action Type = Feedback and Action Value = Helpful	AIS Search Application	Daily	#	0

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

Visualizatio	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
Feedback (Thumbs down)	Automated	Genius Result Event Action[sys_search_genius_result_event_action]	Count of queries where Card Type = Now Assist Multi-Content Response and Action Type = Feedback and Action Value = Not Helpful	AIS Search Application	Daily	#	0
Genius Results with KB	Automated	Search Signal Genius Result Triggered Event[sys_search_genius_result_triggered_event]	Count of queries where Source Table = kb_knowledge	AIS Search Application	Daily	#	0
Genius Results with Catalog Item	Automated	Search Signal Genius Result Triggered Event[sys_search_genius_result_triggered_event]	Count of queries where Source Table = sc_cat_item	AIS Search Application	Daily	#	0
Top 10 queries with KB Genius Result	Automated	Search Signal Genius Result Triggered Event [sys_search_genius_result_triggered_event]	Top 10 queries with count of genius results where Source Table = kb_knowledge	Not Applicable	Daily	Not Applicable	Not Applicable
Top 10 queries with Catalog Item Genius Result	Automated	Search Signal Genius Result Triggered Event [sys_search_genius_result_triggered_event]	Top 10 queries with count of genius results where Source Table = sc_cat_item	Not Applicable	Daily	Not Applicable	Not Applicable

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

Now Assist context menu analytics indicator details

Visualization	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
Usage in this period	Automated	Generative AI Log[sys_generative_ai_log]	Count of actions where Source is Now Assist context menu	By NAcm implicit feedback, By Skills Config, By Feedback, By Status, By Skill Capability	Daily	#	0
Acceptance distribution	Automated	Generative AI Log[sys_generative_ai_log]	Count of actions where content generated by Now Assist context menu grouped by Accepted or Not Accepted	By NAcm implicit feedback, By Skills Config, By Feedback, By Status, By Skill Capability	Daily	#	0
Usage trend by skill	Automated	Generative AI Log[sys_generative_ai_log]	Count of context menu actions grouped by Now Assist skills	By NAcm implicit feedback, By Skills Config, By Feedback, By Status, By Skill Capability	Daily	#	0
Capability distribution	Automated	Generative AI Log[sys_generative_ai_log]	Count of context menu actions grouped by capability	By NAcm implicit feedback, By Skills Config, By Feedback, By Status, By Skill Capability	Daily	#	0
Responses by feedback	Automated	Generative AI Log[sys_generative_ai_log]	Count of actions where feedback for the content generated	By NAcm implicit feedback, By Skills Config, By Feedback, By Status,	Daily	#	0

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

Visualizatio	Indicator type	Indicator source table	Calculation	Available breakdowns	Frequency	Unit	Precision
			by Now Assist context menu grouped by Accepted or Rejected.	By Skill Capability			

Now Assist Admin Settings

Configure general settings for all Now Assist applications from the **Settings** tab in the Now Assist Admin console.

As you begin to explore Now Assist, use **Settings** in the Now Assist Admin console to activate the plugins. See Install Now Assist plugins.

Turn on the Now Assist panel, review multi-language support, data sharing and processing, and Now Assist Guardian features, and view the Now Assist account settings. You can also manage the large language model integrations under **Settings**.

Configure multilingual service for Now Assist applications

Turn on multilingual service for user-entered text with native translation or Dynamic Translation in Now Assist applications.

Before you begin

To see a list of all available languages supported in Now Assist, see Multilingual service for Now Assist.

To use Dynamic Translation, you must install and activate the application and install at least one language pack. For more information, see Activate Dynamic Translation and Activate a language.

Note: Dynamic Translation is not supported in regulated markets.

If you don't see a language and region section of your Now Assist Admin console, make sure that you have installed at least one Now Assist application. This installs or updates the required dependencies to the latest version.

Role required: sn_generative_ai.nsa_admin

About this task

There are two translation services available to translate user-generated content in Now Assist applications. For more information on the differences between the two, see Multilingual service for Now Assist. In that documentation, you can also find the steps for adding another language to your service if it is not preconfigured.

After you have enabled either Dynamic Translation or native translation, translations will be available for in-product experiences, Virtual Agent, and the Now Assist panel.

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

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Settings**.
2. On the left-hand panel, select the **Multilingual service**.
3. In the native translation or Dynamic Translation tab, toggle the switch to **On**.
   A modal displays, confirming your choice and stating that Now Assist streaming of LLM (large language model) responses will be unavailable when Dynamic Translation is active.

   Note:

   When Dynamic Translation is inactive, Now Assistant streaming is available. The modal message displayed will change according to the status of the Dynamic Translation.

4. You can enable both Dynamic Translation and native translation.
   Native translation is applied first. If the language is not supported through native translation, then Dynamic Translation will be applied.
5. **Optional:** To enable the preferred language from the available languages supported in the model providers selected in AI Control Tower, select **Edit** and then select the check-box next to the languages you want to use for translation.

   (Optional) See Multilingual service for Now Assist for more information about adding new languages for your preferred model to support.

   To update your choice, select **Save**.

Result
Multilingual service is enabled for Now Assist applications.

Multilingual service for Now Assist

Now Assist applications use the multilingual capabilities of large language models (LLMs) to translate user-generated content.

Multilingual service overview

You can configure the multilingual services for Now Assist applications to communicate with your users in their preferred languages. The LLM detects the language used based on the user's language choice in their Language & Region preferences or the contents of the text, such as a conversation message.

Multilingual support is enabled for multiple LLM providers, such as Now LLM Service, Microsoft Azure OpenAI, Anthropic Claude on AWS, and Google Gemini. Azure OpenAI offers the most language support on the ServiceNow AI Platform.

How the different multilingual services work

The inherent multi-language capabilities of the LLM is called native translation. In native translation, text is sent as-is to the model and the model responds in the user's language. For example, if a user sends a message in Spanish, the multilingual LLM supports Spanish, so it receives the Spanish input directly from the user. It then generates a response in Spanish, regardless of the language used in the source material it uses to create that response.

For Dynamic Translation, once non-English content is detected, the translation service translates the content to English before sending the request to the large language model (LLM). After the LLM returns a response, the response is translated back into the user's preferred language for them to see. For example, a user enters a message in Slovenian. Dynamic Translation translates

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

the Slovenian to English before giving the message to the LLM. The LLM generates a response in English, and then Dynamic Translation translates that response to Slovenian to show to the user.

If you enable both native translation and Dynamic Translation, native translation has preference and is used first.

If you have native translation enabled, the LLM checks the user's language to see if it is supported. If the language is supported, only native translation is used and no Dynamic Translation occurs. A small number of languages are not yet natively supported by all models. Now Assist Dynamic Translation is automatically applied to address this for unsupported languages. Additionally, if you have native translation enabled, then LLM requests that return outputs in the fallback language will not be translated by Dynamic Translation. Otherwise, they'll be returned as-is to the user. This process helps increase skill performance and decrease latency.

Note: If you want to add a language that is not part of the regularly configured languages, you can do so. See the section below, Adding languages beyond the P1, P2, and P3 language groups.

Native translation

Supported languages are categorized into three basic groups.

P1 languages have the strongest native support across most model providers, meaning high-quality translation and generative capabilities with fewer errors. Nuance, idioms, and context are usually handled well because these languages are deeply integrated into the models.

P2 languages are supported natively by smaller ServiceNow models and most external providers, but not always by every model (e.g., some exceptions for Canadian French and Dutch). Accuracy is good, but there may be occasional gaps in nuanced understanding compared to P1 languages.

P3 languages include a broader set of languages supported by major providers like Microsoft Azure OpenAI, Google Gemini, and Anthropic Claude on AWS, but not natively by ServiceNow's own large models. While coverage is comprehensive, these languages generally have less training depth, so subtle tone or idiomatic expressions may be harder to capture perfectly.

- P1
  - English
  - French
  - German
  - Italian
  - Spanish
  - Brazilian Portuguese
- P2
  - Canadian French
  - Japanese
  - Dutch
- P3
  - Swedish
  - Finnish
  - Czech

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

- Hebrew
- Hungarian
- Korean
- Norwegian
- Polish
- Portuguese
- Russian
- Simplified Chinese
- Traditional Chinese
- Thai
- Turkish
- Arabic
- Danish

Specific models vary in their support.

- ServiceNow Large Language Model v1 (Llama 3.3 70B) only supports English natively.
- ServiceNow Large Language Model v2 (GPT-OSS) supports English, French, German, Spanish, and Italian natively.
- ServiceNow Small Language Model v1 (Mistral Nemo 12B) and v2 (Apriel) supports all P1 and P2 languages.
- Microsoft Azure OpenAI GPT-4.1 and GPT-4.1-mini support more than 80 languages, including P1, P2, and P3.
- Google Gemini 2.5 Flash/2.5 Pro currently does not support French Canadian natively.
- Anthropic Claude 4.0 Sonnet currently does not support French Canadian or Dutch natively.

Adding languages beyond the P1, P2, or P3 language groups

Some providers, such as Azure OpenAI, offer additional language support beyond the P1, P2, and P3 language groups previously described. You can configure your instance to include additional languages.

Note: Before adding a language that is not in the P1, P2, or P3 groups, you should confirm whether the language is supported by the provider. If you add a language that isn't supported, you may get unexpected results. The LLM could fail to detect the language, and/or it could not be able to translate effectively.

Ensure the language is on the sys_language table.

Create a record for the language if necessary. The record should include the language name and code. The *Active* field should be set to true.

Enable the language for the Now Assist model.

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

- Go to the sys_generative_ai_model_config table and open the record for the model you want to add the language to, such as gpt_small by Azure OpenAI.

  Note: The model for AI Search is named E5FT. If you want to make changes to the language support for AI Search, you can configure that record. After, you need to reindex all sources in the new language.

- In the *Supported languages* list, add the language and save the record.

Activate the language for Now Assist applications.

- Go to **All** > **Now Assist Admin** > **Settings** and open the **Multilingual service** tab.
- Scroll to the provider you edited, then select the edit button.
- Select the language you added and save.

  Note: If you don't see your language as an option, ensure that the language is on the sys_language table and *Active* is set to true.

Note: If semantic search is not working even after installing a language and enabling it for AI Search, you may need to clear the embedding model cache. To do so, run the following: `GlideCacheManager.flush("ais_semantic_embedding_model_cache");`.

Dynamic Translation

Dynamic Translation for Now Assist uses the Microsoft Azure OEM translation service through ServiceNow. Every language available as a language pack on the ServiceNow platform is supported by Microsoft Azure OEM. These languages are:

- Swedish
- Finnish
- Czech
- Hebrew
- Hungarian
- Korean
- Norwegian
- Polish
- Portuguese
- Russian
- Simplified Chinese
- Traditional Chinese
- Thai
- Turkish
- Arabic
- Danish

For more information, see Microsoft Azure OEM for Dynamic Translation in Now Assist.

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

Choosing a translation service

There are different factors that can go into your decision to use one translation service over the other, such as latency and quality of translation. Your entitlements determine what is available to you, but there are no additional costs associated with native translation service.

Note: The native translation service is included in Pro Plus and Enterprise licenses, by default.

For more information on tracking Now Assist usage, see Monitoring Now Assist usage.

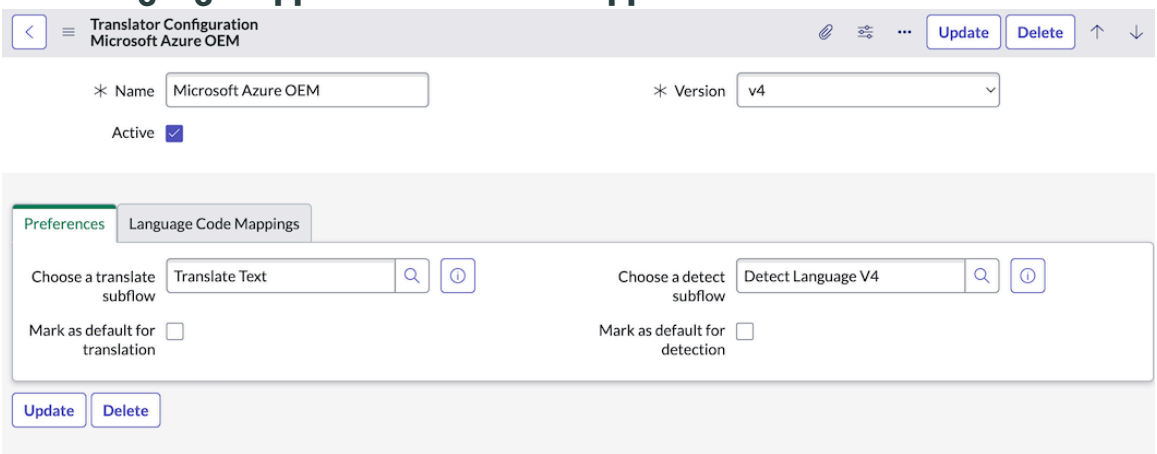
Enabling translation for Now Assist

For more information on turning on multilingual services for Now Assist applications, see Enable translation for Now Assist applications.

Microsoft Azure OEM for Dynamic Translation in Now Assist

Support multi-language input in Now Assist applications with the Microsoft Azure OEM translator service included in Generative AI Controller.

Multi-language support for Now Assist applications



When dynamic translation in the Generative AI Controller is enabled, your users' input text is handled by Microsoft Azure OEM (com.snc.microsoft_oem_translation_spoke) as follows:

1. The language of the text is detected.
2. The text is dynamically translated to English, then provided to Now Assist.
3. Now Assist returns a response in English.
4. The English response is translated to the user's preferred language and displayed in the UI.

For more information about enabling dynamic translation in this context, see Configure multilingual service for Now Assist applications and Enable Dynamic Translation for capabilities in Generative AI Controller.

From the Zurich release, this spoke has available support for specific regulated markets, subject to conditions. For information see KB0743854 on the Now Support portal.

Activation

The Microsoft Azure OEM translator service is included with Generative AI Controller. Microsoft Azure OEM requires Dynamic Translation tables, so Generative AI Controller activates Dynamic Translation if it isn't already active.

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

Microsoft Azure OEM appears as a row in the Translator Configurations (sn_dt_translator_configuration) table. This record is not editable. The connection credentials of this translator service are already configured. For more information about this integration, see https://www.servicenow.com/company/media/press-room/gen-ai-now-platform.html.

If you try to update this record in the Translator Configurations table, the following error is displayed: "The Microsoft Azure OEM translator configuration is used in Generative AI flows. It cannot be edited".

With the Xanadu Patch 3 release, the Microsoft Azure OEM translator includes support for Exclusion Framework in Dynamic Translation. Also the API used by flows and subflows is automatically updated to v4.

The application scope for Microsoft Azure OEM is sn_ms_oem.

For more information, see Dynamic Translation.

Related topics

- Dynamic Translation
- Generative AI Controller

Data sharing and processes

Data sharing and processes topic description.

Configure Now Assist privacy policies

Configure privacy policies to anonymize data in Now Assist.

Before you begin

You must have the following applications installed on your instance:

If you do not have a data steward, see Assign the data steward role documentation.

Role required: sn_generative_ai.data_steward

About this task

The Now Assist **Privacy Policies** enables you to review and configure how personally identifiable information (PII) is de-identified in Now Assist.

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Settings**.
2. In the navigation pane, go to **Data sharing and processing** > **Privacy policies**.
3. Select **View detail** of the privacy polic you want to review.
   Review the policy details and current active patterns in the pop-up window.
4. **Optional:** Select **Edit in Data Privacy**.
   See Configuring Data Privacy for Now Assist for more information on configuring Now Assist policies.

Opt out of data sharing for Now Assist

Data sharing improves ServiceNow AI products. You can opt out of data sharing from the Now Assist Admin console Settings page.

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

Before you begin

Important: Data sharing is not available for GCC or self-hosted instances. You don't need to opt out because data sharing is never enabled. If you have any questions, reach out to your account representative.

If you do not have a data steward, see Assign the data steward role documentation.

Role required: sn_generative_ai.data_steward

Note: The **Opt out** button appears only if you've installed at least one Now Assist application or plugin. For a list of all Now Assist applications, see Exploring Now Assist Admin.

About this task

By opting out of the ServiceNow customer data sharing program, you can no longer provide data to improve ServiceNow AI products. By sharing data with the ServiceNow AI development program, you provide relevant data to help improve prediction accuracy, user experience, tailor products to your business needs, and reduce hallucinations for your activated Now Assist skills.

You can choose to opt out a ServiceNow instance from data sharing from the Now Assist Admin console if you don't want to participate. Repeat the opt-out process for all instances that use the Now Assist functionality.

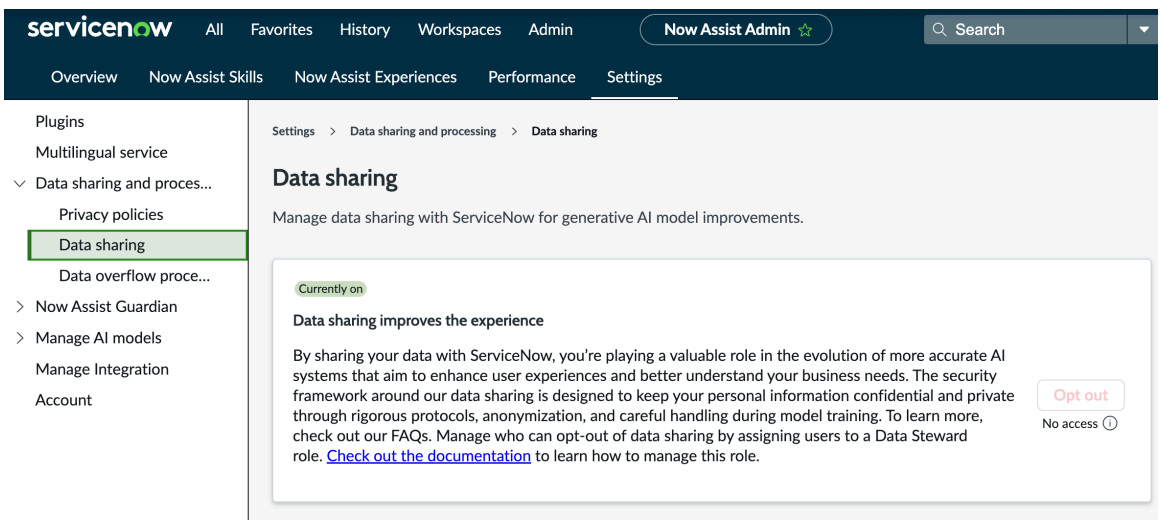
Opting out can take up to five business days to process.

Note: You must opt out of data sharing to turn off ServiceNow ability to gather Customer data for AI and data products.

Procedure

1. Change the current session scope by selecting the Globe icon in the top right next to the search bar and setting Application Scope to Generative AI Controller.
2. Navigate to **All** > **Now Assist Admin** > **Settings**.

   If you're already in the Now Assist Admin console, select the **Settings** tab.
3. In the Settings panel, select the **Data sharing and processing** > **Data sharing** tab.



4. Select **Opt Out**.
5. In the confirmation window, select **Opt Out**.

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

Result
Your data sharing preference is saved on the instance. If you want to opt back in to data sharing, you must consult with your account executive.

Configure Now Assist data overflow processing

Configure where Now Assist data is processed during periods of high traffic.

Before you begin
Role required: sn_generative_ai.data_steward

About this task
By default, Now Assist data is handled in ServiceNow datacenters. During periods of high network traffic, data can be "burst" to Microsoft Azure datacenters to help keep consistent load times and avoid capacity overflow errors. Users with the data steward role can opt out of or into this behavior from the Now Assist Admin console.

Procedure
1. Navigate to **All** > **Now Assist Admin** > **Settings**.
2. In the navigation pane, go to **Data sharing and processing** > **Now overflow processing**.
3. Select either **Opt out** or **Opt in**.
   Instances are opted in by default. You can change this setting at any time.

Assign the data steward role

Select a data steward to make decisions about data sharing with ServiceNow in Now Assist applications.

Before you begin
Role required: sys_admin

About this task
A data steward is responsible for determining whether prompts, inputs, and responses with generative AI in Now Assist applications on your instance are shared with ServiceNow.

The ServiceNow customer data sharing program allows you to provide data to improve ServiceNow AI products. This data helps improve prediction accuracy, enhance user experience, tailor products to your business needs, and reduce hallucinations for your activated Now Assist skills.

The following steps explain how to assign a role to a specific user. Another option is to make a group, such as "AI Data Stewards," assign the role to the group, and then put the user in that group. See Assign a role to a group for more information.

Procedure
1. Navigate to **All** > **User Administration** > **Users** and then open a user record.
2. In the Roles related list, select **Edit**.
3. In the collection of roles, select the sn_generative_ai.data_steward role, and then select the **Add** ( ) icon.
4. Select **Save**.

Result
Once the data steward has the correct role, they can review the data sharing information on the Now Assist Admin console in the **Settings** tab under **Data privacy and sharing** > **Data sharing**.

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

Now Assist Guardian

Now Assist Guardian is built on the ServiceNow Small Language Model (SLM) and monitors generative AI interactions to detect offensive content, prompt injection attacks, and sensitive topics.

Now Assist Guardian Overview

Generative AI is an emerging technology. Human interactions are unpredictable, and outputs generated by large language model (LLM) are probabilistic, which means that they're based on probabilities. Running the same input twice may generate two different outputs. Managing this risk is an important consideration when implementing generative AI on your instance. Now Assist Guardian evaluates requests sent to LLMs and their responses in real-time to reduce that risk.

Guardrails

Now Assist Guardian provides three guardrails. Each guardrail has a different scope of applicability:

Guardrail	What it detects	Scope
Offensiveness detection	Offensive or harmful content in AI inputs and outputs.	Specific Now Assist skills and workflows.
Prompt injection detection	Attempts to override LLM instructions or expose restricted information.	All generative AI applications and features.
Sensitive topic filters	Subjects not suited for AI responses, such as workplace safety or employee compensation.	Virtual Agent conversational skills only (available for HR Service Delivery and Customer Service Management).

Note: The scope of each guardrail differs. Prompt injection detection applies to all generative AI applications and features. Offensiveness detection applies only to supported Now Assist skills and workflows. Sensitive topic filters apply only to Virtual Agent conversations and require HR Service Delivery.

Offensive content

Due to the probabilistic nature of generative AI, it's possible for an LLM to generate offensive content. If there's offensive content in the input of the request, offensive content can also occur in the response. Examples of offensive content include language that is toxic, defamatory, or fraudulent.

When offensive content is detected, Now Assist Guardian logs the event. You can also configure it to block the content. This guardrail applies to specific Now Assist skills and workflows.

Prompt injection

Prompt injection is a type of security attack where someone tries to override the normal instructions of an LLM to access restricted information or cause unintended behaviors. Now Assist Guardian detects prompt injection attempts by using an LLM trained on various types of prompt injection techniques, such as role playing, paraphrasing, repetition, instructions to ignore other instructions, and persuasion.

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

Note: Due to the probabilistic nature of the model and evolving attack techniques, Now Assist Guardian may not identify every prompt injection attempt in some cases.

Prompt injection protection applies to all generative AI applications and features on your instance. It is not limited to specific skills or workflows.

Filtered subjects

Certain subjects, such as workplace safety employee compensation, or personal well-being may not be best suited for generative AI responses. You can activate filters that detect these kinds of subjects in Virtual Agent conversations and redirect users to the Sensitivity Detection: Fallback Virtual Agent topic instead of generating an AI response.

Note: Sensitive topic filters apply only to Virtual Agent conversational skills. These filters are available only with HR Service Delivery and Customer Service Management.

Logging and blocking

Now Assist Guardian logs detected events for offensiveness and prompt injection. You can access logs from **Now Assist Admin** > **Settings** > **Now Assist Guardian**. Log data includes information about the request, the conversation that contains the offensive content, and any user feedback.

In addition to logging, you can configure Now Assist Guardian to block offensive content or prompt injection attempts. When blocking is enabled and content is detected, you see a standard error message instead of the generated response. The standard error message displays that the request couldn't be completed, and you don't see what the AI generated. Before enabling blocking, review logs for a period of time to understand how frequently these issues occur in your environment.

Redirection for sensitive filtered topics

After a filter detects a sensitive topic, Now Assist Guardian redirects you to the Sensitivity Detection: Fallback topic in Virtual Agent. This topic can redirect you to a live agent or help you create an HR case.

You can override the redirection by selecting **Proceed, not sensitive**. It returns you to their original topic without initiating the fallback flow.

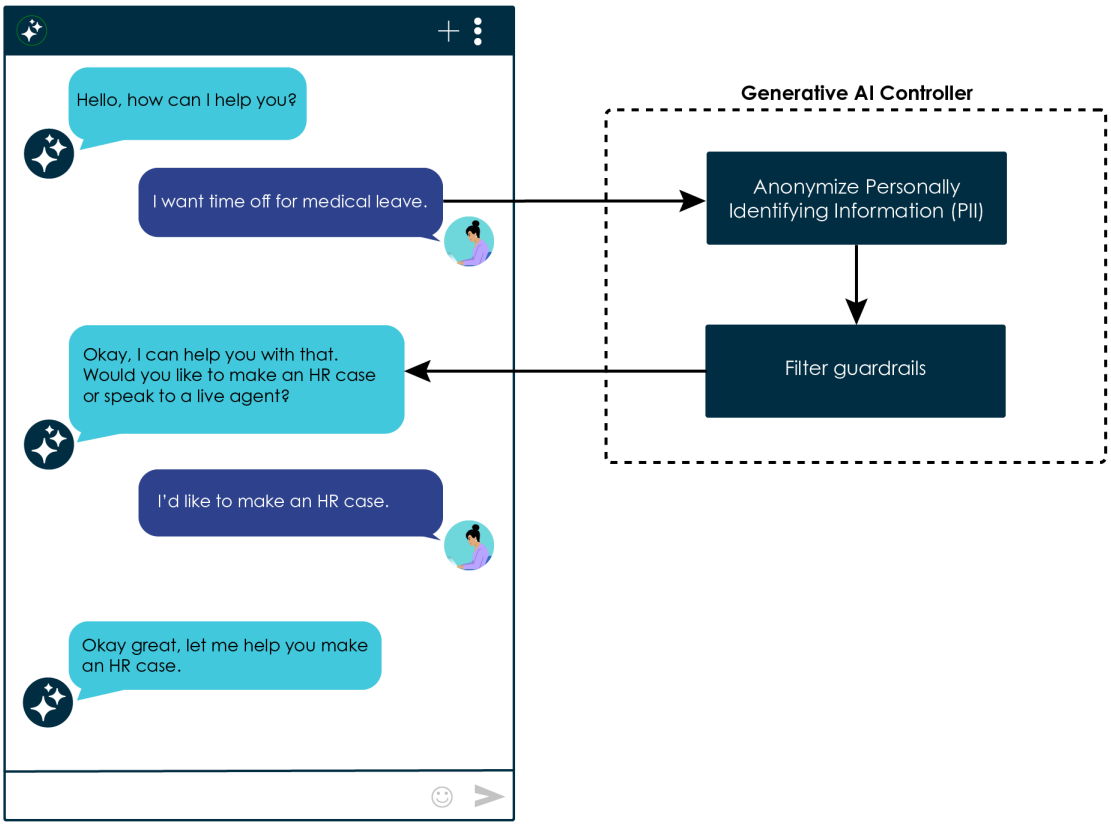
Note: After you continue with the fallback topic, for example, by starting the flow to create an HR case, Virtual Agent does not continue detecting sensitive topics in that conversation.

Now Assist Guardian at runtime

All skills that use Now Assist Guardian remove personally identifiable information (PII) before the request reaches the LLM. You can configure what type of data is anonymized. For more information see, Configuring Now Assist for Data Privacy.

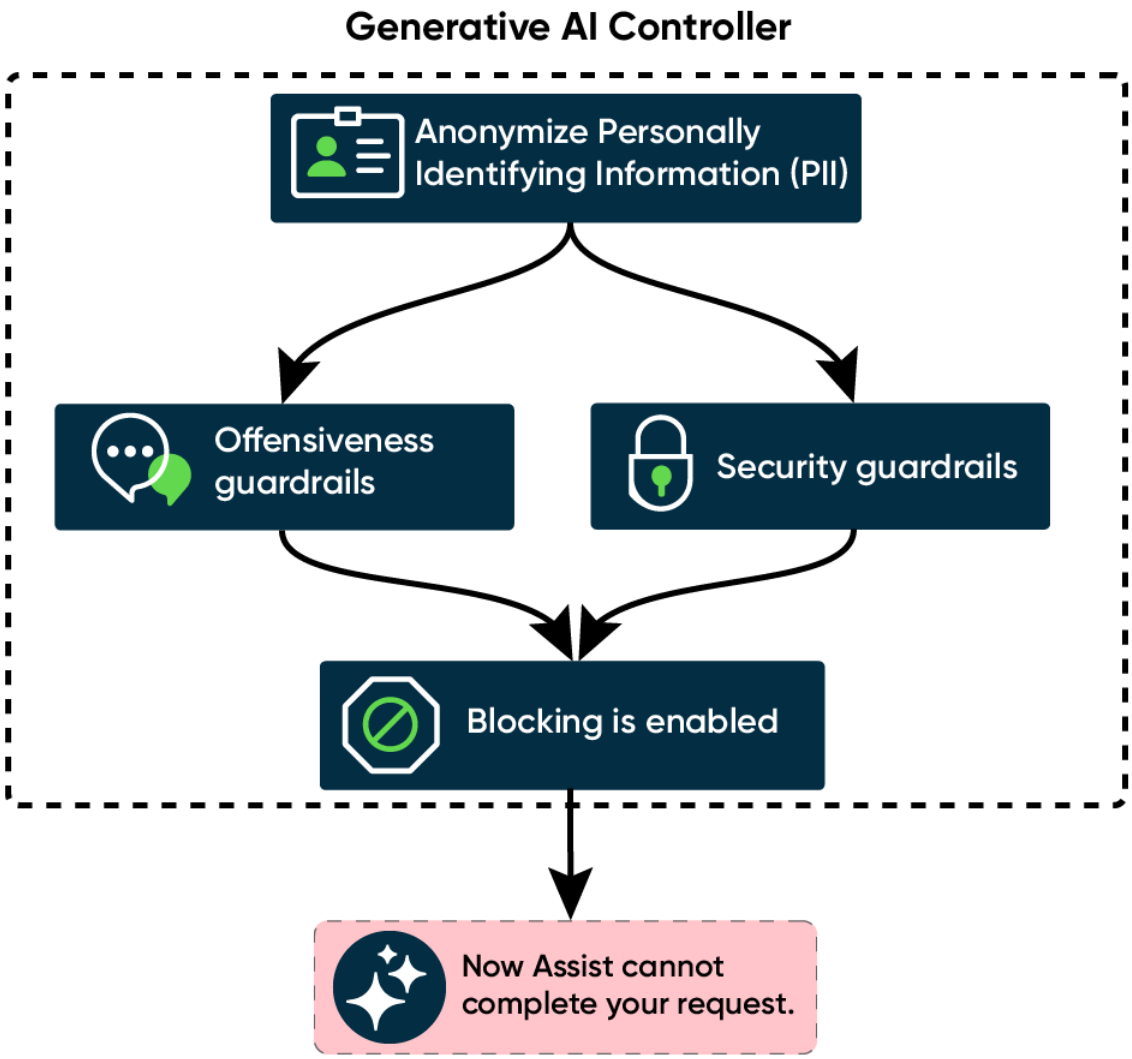
For conversational skills, semantic search processes requests to determine whether a filter has been detected. If so, the user is redirected to a Virtual Agent topic that asks if they want to make an HR case or speak to a live agent.

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



For catalog item generation and agent skills, such as summarization and resolution note generation, offensiveness and prompt injection guardrails run on inputs and outputs of requests. If either is detected, Now Assist Guardian logs the request. If you've chosen to block this content, then a standard error message appears, and the user doesn't see the result of the request.

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



Safety categories captured in Now Assist Guardian logs

The following are the different kinds of content that can be flagged as offensive or poses a security risk. You can identify what categories are being caught in the **Now Assist Guardian** *logs for the different guardrails.*

Category	Category ID	Description
Toxic Content	O1	Should not: Promote hate speech, incite violence like child abuse or bullying or engage in personal attacks or harassment. Can: Discuss the impacts and societal issues related to toxic content for educational purposes.
Unfair Representation	O2	Should not: Assist in perpetuating stereotypes or encourage discrimination or bias.

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

Category	Category ID	Description
		Can: Discuss issues of representation and promote awareness and inclusivity.
Adult Content	O3	Should not: Display sexually explicit material or content inappropriate for minors including pornography, erotic chat and sexual activity. Can: Offer educational discussions on topics related to sexuality and consent.
Erosion of Trust in Public Information	O4	Should not: Spread misinformation or disinformation undermining public trust such as misleading medical advice, illegal drugs and advice. Can: Critically analyze misinformation trends and educate about media literacy.
Propagating Misconceptions/ False Beliefs	O5	Should not: Disseminate unverified, misleading information, conspiracy theories, science denialism, historical revisionism, or content that could lead to radicalization. Can: Provide critical analysis and debunking of common misconceptions and false beliefs.
Risky Financial Practices	O6	Should not: Promote unethical trade practices such as multi-level marketing, deceptive lending practices, or gambling. Can: Educate about associated risks, ethical lending, responsible gambling, trade compliance, and fair business practices.
Trade and Compliance	O7	Should not: Involve labor issues, engage in unethical trade affairs, contribute to unfair labor practices or promote fraudulent trade practices. Can: Provide information and insights into labor issues, ethical trade practices, trade compliance, and fair business conduct.
Dissemination of Dangerous Information	O8	Should not: Spread fake news, propagate social media rumors, disclose government secrets or leak financial data and other contents that contributing to public panic or national security. Can: Educate about the dangers of such information and promote safety.
Privacy Infringement	O9	Should not: Involve doxxing, unauthorized data sharing, or surveillance like private information inferring or leaking. Can: Discuss the importance of privacy rights and digital security.

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

Category	Category ID	Description
Security Threats	O10	Should not: Spread malware or encourage terrorism or other security threats such as weapon generation, cyberattack or chemical harms. Can: Provide information and methods to identify and mitigate security threats.
Defamation	O11	Should not: Publish false statements damaging to individuals, groups or organizations. Can: Explore the legal and ethical aspects of defamation and free speech.
Fraud or Deceptive Action	O12	Should not: Engage in academic dishonesty, perpetrate scams, promote pseudo-pharmaceuticals, or disseminate fake online content. Content should not deceive or defraud individuals. Can: Educate about recognizing and preventing fraudulent activities.
Influence Operations	O13	Should not: Manipulate public opinion or disrupt societal harmony for political or ideological purposes like surveillance or campaign. Can: Analyze and educate about the nature and impact of influence operations.
Illegal Activities	O14	Should not: Promote or provide guidance on engaging in illegal actions, such as crimes or terrorism. Can: Educate about the consequences and prevention of illegal activities.
Persuasion and Manipulation	O15	Should not: Involve persuasion or manipulation techniques that promote self-harm, psychological manipulation, or any form of harm to individuals' well-being. Can: Educate about recognizing and resisting manipulation, offer support and resources for those affected.
Violation of Personal Property	O16	Should not: Enagge in violating an individual's personal property,malicious mimicking, plagiarism, copyright infringement or identity forgery. Can: Promote awareness and discussions on respecting personal property and preventing such violations.

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

Skills that support offensive detection

Supported skills by workflow

Workflow	Application	Supported skills
Technology	Now Assist for Configuration Management Database (CMDB)	• Configuration item (CI) summarization • Manage duplicate CIs • Service Graph Connector diagnosis
Technology	Now Assist for IT Operations Management (ITOM)	• Alert analysis • Alert investigation
Technology	Now Assist for IT Service Management (ITSM)	• Change request risk explanation • Change request summarization • Chat reply recommendation • Chat summarization • Incident assist • Incident summarization • KB generation • Resolution notes generation • Sidebar discussion summarization
Technology	Now Assist for Security Incident Response	• Post-incident analysis • Resolution notes generation • Security incident recommended actions • Security incident summarization
Technology	Now Assist for Strategic Portfolio Management (SPM)	• Multi feedback summarization • Planning item doc summarization • Project doc summarization • Project summary emails • Story generation • Write planning item
Customer	Now Assist for Customer Service Management (CSM)	• Case summarization • Chat recommendation • Chat summarization

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

Supported skills by workflow (continued)

Workflow	Application	Supported skills
		• Email recommendation • KB generation • Resolution notes generation • Sidebar summarization
Customer	Now Assist for Field Service Management (FSM)	• KB generation • Sidebar summarization • Work order task summarization
Customer	Now Assist for Financial Services Operations (FSO)	• Case summarization • Disputes intake via Virtual Agent
Customer	Now Assist for Public Sector Digital Services (PSDS)	• Government case summarization • Chat summarization
Employee	Now Assist for Health and Safety	Incident summarization
Employee	Now Assist for HR Service Delivery (HRSD)	• Case summarization • Chat summarization • KB generation • Resolution notes generation
Employee	Now Assist for Legal Service Delivery (LSD)	Legal request summarization
Employee	Now Assist in Contract Management	• Contract analysis • Contract metadata extraction
Creator		Catalog item generation
Finance & Supply Chain	Now Assist for Accounts Payable Operations (APO)	Record summarization
Finance & Supply Chain	Now Assist for Supplier Lifecycle Operations (SLO)	Supplier case summarization
Finance & Supply Chain	Now Assist for Sourcing and Procurement Operations (SPO)	Record summarization

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

Activate offensiveness protection for generative AI

Activate offensiveness detection to log or block offensive content generated by Now Assist skills and workflows.

Before you begin

Role required: sn_generative_ai.nsa_admin

About this task

Generative AI output is probabilistic, which means that the same input can produce different outputs. Some of the AI generated content may be offensive, which includes toxic, sexist, or other harmful language. Now Assist Guardian enables you to detect offensive content in both inputs and outputs, and logs the event when it is detected. You can also cofigure it to block offensive material so that users see a standard error message instead of the generated response.

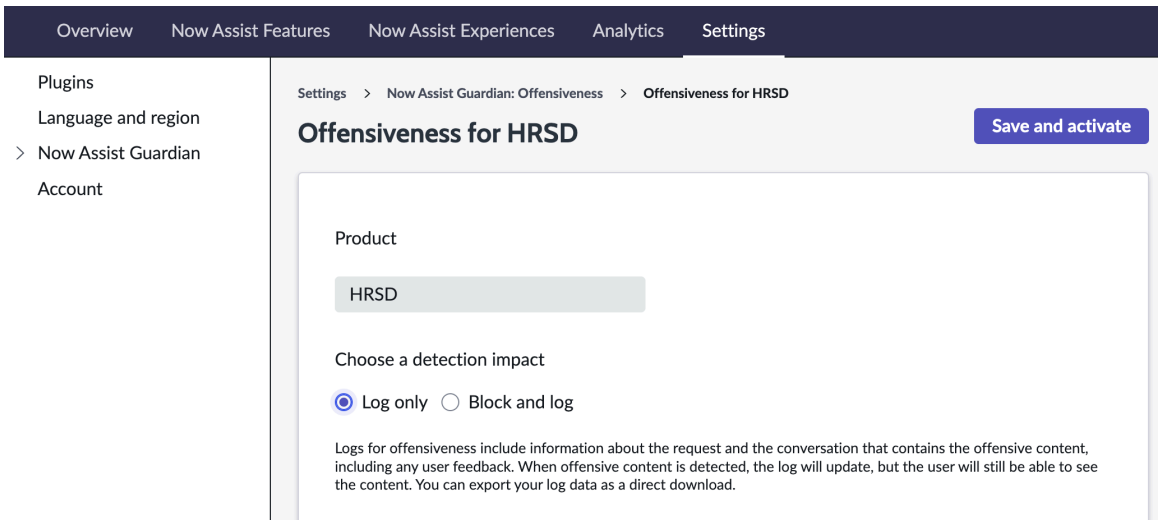
Note: Offensiveness detection applies only to specific Now Assist skills and workflows. It is not available for all Now Assist applications. For more information about the list of skills that support offensiveness detection, see Now Assist Guardian.

You can export logs for review. For more information, see Export Now Assist Guardian logs.

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Settings**.
2. In the side panel, select the **Now Assist Guardian** > **Offensiveness** tab.
3. Go to the **Available for you** tab to see which workflows you can choose from.

   Offensiveness guardrails that are already activated appear in the **Active** tab.
4. Select **Activate** for the workflow that you want to enable offensiveness detection.
5. Select your impact detection.
   - Select **Log only** to record the events when offensive content is detected. The content is still shown to the user.
   - Select **Block and log** to record the event and prevents the content from being shown to the user. The user sees a standard error message instead.



6. Select **Save**.

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

Result
Offensiveness detection guardrail is enabled on your instance for the selected workflow. Events are logged when offensive content is detected or generated.

What to do next
You can enable offensiveness detection for separately for each supported Now Assist application and workflow. Repeat this task for each workflow on which you want offensiveness protection enabled.

To change the detection impact for an active workflow, select more options (⋮) icon in the list of active workflows and then select **Edit**.

To deactivate offensiveness protection for a workflow, select more options (⋮) icon in the list of active workflows and then select **Deactivate** .

Configure prompt injection attack protection

Activate or deactivate prompt injection attack detection to protect all generative AI applications and AI-generated text and conversations.

Before you begin
Role required: sn_generative_ai.nsa_admin

About this task
Prompt injection attacks are a type of cybersecurity attack where someone tries to override the initial instructions of an LLM to cause unintended behaviors. Now Assist Guardian detect and log these prompt injection attack attempts across all generative AI applications and features. You can also configure the prompt injection detection guardrail to block the AI-generated response when an attack is detected in addition to logging it.

You can export logs for review. For more information, see Export Now Assist Guardian logs.

Procedure
1. Navigate to **All** > **Now Assist Admin** > **Settings**.
2. In the side panel, go to **Now Assist Guardian** > **Prompt Injection**.
3. Select the toggle to activate or deactivate prompt injection attack detection.
4. **Optional:** Under **Detection impact**, select the options icon (⋮) and then select **Edit** to change how detected attacks are handled.

   (Optional) You can choose whether prompt injection attacks are blocked as well as logged.

Result
Prompt injection detection is configured on your instance. When enabled, you see a standard error message when an attack is detected.

Configure sensitive topic filters

Set up filters in Now Assist Guardian to redirect Virtual Agent users to a live agent or HR case when a sensitive subject is detected in a conversation.

Before you begin
Role required: sn_nowassist_admin.nsa_admin

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

About this task

Now Assist Guardian detects sensitive subjects in Virtual Agent conversations and redirect users to the Sensitivity Detection: Fallback Virtual Agent topic instead of continuing the conversation with generative AI. You can activate and configure the subjects that trigger this redirection, and add sample phrases to improve detection accuracy.

Note: The sensitive topic filter guardrail applies only to Virtual Agent conversational skills and is available for HR Service Delivery and Customer Service Management (CSM).

See Now Assist Guardian for more information.

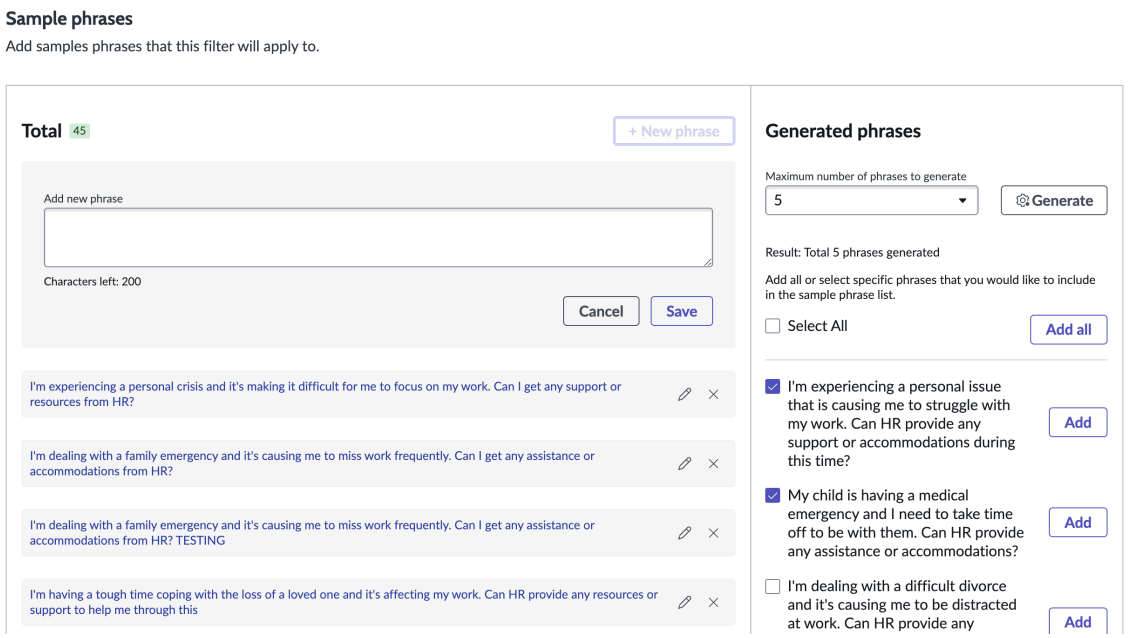
Procedure

1. Navigate to **All** > **Now Assist Admin** > **Settings**.
2. Navigate to **Now Assist Guardian** > **Filters**.
3. Go to the **Available for you** tab to see which filters are available.

   Note: When the CSM plugin is installed on your instance, the default sensitive filters for CSM appear in the **Available for you** tab.

4. Select **Activate** on the filter you want to enable.
5. Review and generate sample phrases.

   Sample phrases tell the LLM what kind of text the filter should detect. The more sample phrases you provide, the more accurately the LLM detects these topics.



Option	Description
Add a phrase	**a.** Select **New phrase** to add a phrase manual ly. **b.** Enter text in the Add new phrase text box. **c.** Select **Save**.

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

Option	Description
Generate a phrase using generative AI	**a.** In the **Generative samples** panel, select the maximum number of phrases to gener ate from the list. **b.** Select **Generate**. **c.** Select the check box next to the phrase and select**Add** to add individual sample phrase. **d.** Select the check box next to each phrase and then select**Add all** to add multiple phrases together.
Edit a phrase	**a.** Select the pencil icon and start editing the phrase. **b.** select **Save**.
Delete a phrase	Select the delete icon to remove the phrase.

6. Select **Save and Continue**.
7. Choose which Virtual Agents you want to enable the sensitive topic filters for by selecting the checkbox next to the Virtual Agent name.
8. Select **Save and Continue**.
9. Review the fallback topic for when a filter is detected.

   By default, when a sensitive topic is detected, the user is redirected to the Sensitivity Detection: Fallback topic, which prompts the user to create an HR case or speak to a live agent. You can preview the topic from the guided setup.
10. Select **Save and continue**.
11. Review your filter configuration in the final step.

    If you want to make any changes, use the step navigation panel to return to a previous step.
12. Select **Activate** to activate the filter.

Result
The sensitive topic filter is activated and appears in the **Active** tab of the **Filters** section of Now Assist Guardian settings. When the filter detects a sensitive topic in a Virtual Agent conversation, the user is redirected to the Sensitivity Detection: Fallback Virtual Agent topic.

What to do next
To modify the filter configuration after activation, go to the **Active** tab of the Now Assist Guardian settings options and select the more actions icon (⋮) in the **Actions** column.

Export Now Assist Guardian logs

Export logs from Now Assist Guardian to get insights into how often different guardrails are being detected and used.

Before you begin
Role required: sn_generative_ai.nsa_admin

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

About this task

Now Assist Guardian logs all three types of guardrails available. Reviewing the logs can help you determine how often offensive content is generated, prompt injection attack attempts occur, or sensitive topics are detected.

See Now Assist Guardian for more information.

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Settings**.
2. Select the more options menu icon (⋮) and select **Export**.
   a. For the offensiveness guardrail, go to **Now Assist Guardian** > **Offensiveness**.

      Use the list in the **Active** tab to select which workflow you want to export logs for.

   b. For the prompt injection guardrail, go to **Now Assist Guardian** > **Prompt Injection**.

   c. For the sensitive topics, go to **Now Assist Guardian** > **Filters**.

Result

The log is exported as a .csv file to your computer.

What to do next

If you do not see any log data, then it is most likely that the guardrail has not been triggered yet. If you believe you should be seeing data but aren't, reach out to Now Support.

Configuring a Guardrail Service Provider

Now Assist Guardian supports default guardrails and third-party or custom guardrail service providers to extend AI content monitoring with your organization's AI governance policies.

Guardrail Service Provider overview

Now Assist Guardian includes default guardrails to detect offensive content, block prompt injection attacks, and filter sensitive topics. In addition to these guardrails, Now Assist Guardian supports third-party guardrail service providers. You can also create your own custom guardian using the Bring Your Own Guardrail (BYOG) feature to monitor and detect malicious content.

Custom guardian evaluates prompts and responses in real time, just like the default guardrails. You can set each custom guardrail to log events, block content, or both based on your specific security or content policies.

ServiceNow supported third-party guardrail service providers

Now Assist Guardian supports the following cloud AI service providers:

- ServiceNow guardrail: The default provider built into the platform. No additional configuration is required.
- Third-party providers: Integrate with supported cloud AI safety services. Following are the supported providers:

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

- Microsoft Azure Content Safety
- Amazon Bedrock
- Google Model Armor

- Custom (BYOG): Create your own guardrail to apply your organization's specific compliance, security, or content policies.

Only one guardrail service provider can be active at a time. The selected provider monitors and evaluates all Now Assist interactions for harmful, offensive, or prompt injection content.

Custom guardrails for Now Assist Guardian

The Bring Your Own Guardrail (BYOG) feature lets you create and connect a custom guardrail in Now Assist Guardian. A custom guardrail connects to an external endpoint that you control, allowing you to apply policies specific to your organization. It evaluate prompts and responses in real time, the same as default guardrails. You can configure each custom guardrail to log events, block content, or both.

Setting up a custom guardrail involves three steps:

1. Create a connection and credential alias for the custom guardian and complete its setup. The alias supports Basic Auth, OAuth 2.0, and API Key credential types.

   For more information, see Create a Connection & Credential alias.

2. Create a custom guardian in the Generative AI Custom Guardian Transformer [sys_generative_ai_custom_guardian_transformer_list.do] table.

   For more information, see Create a custom guardian.

3. Select your preferred guardrail service providers to monitor content.

   For more information, see Setup a Guardrail Service Provider.

Create a custom guardian

Create your own custom guardian to monitor and detect requests sent to LLM.

Before you begin

Configure a connection and credential alias for your preferred provider. For more information, see Create a Connection & Credential alias.

Role required: admin

Procedure

1. Navigate to **All**, and then enter `sys_generative_ai_custom_guardian_transformer_list.do` in the filter to go to the Generative AI Custom Guardian Transformer [sys_generative_ai_custom_guardian_transformer_list.do] table.
2. Select **New**.
3. In the **Name** field, enter the name of the custom guardian.
4. Create a request transformer script.

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

a. In the **Request Transformer** filed, enable the **Turn on ECMAScript 2021 (ES12) mode** toggle to create script.

The following example shows the request transformer script for the custom guardian request structure.

```
((function(inputs) {
 /*
 inputs structure: {
  prompt: object, The textual content that needs to be
 evaluated with guardian
  threshold_setting: object
  connection_attributes: object
 }
 write code here to construct the request body and any custom
 headers needed using the inputs object.
 */
 // construct body using the inputs like input.prompt and so
 on.
 var body = {};
 //construct headers using the inputs
 var headers = {};
 return {
  body: body,
  headers: headers
 };
})(inputs);
```

5. Select the **Active** option.
6. Create a response transformer script.

a. In the **Response Transformer** filed, enable the **Turn on ECMAScript 2021 (ES12) mode** toggle to create script.

The following example shows the response transformer script for the custom guardian response structure.

```
(function(inputs) {
 /*
 write code here to transform the custom guardian response
 into standard guardian response format expected by platform
 inputs structure: {
  response_body: object,
  threshold_setting: object
 }
 */
 // write code here to populate the response object.
 var guardianResponse = {};
 //Set flagged to true, if the guardian has flagged the text.
 Categories are optional and are for logging purpose only
 // if transformer setting is Prompt Injection, expected
 response format: {security:{flagged:false,categories:[]}}
```

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

```
 // if transformer setting is Offensiveness, expected
 response format: {safety:{flagged:false,categories:[]}}
 // if transformer setting is ALL, expected response format:
 {safety:{flagged:false,categories:[]},security:{flagged:fals
e,categories:[]}}
 return guardianResponse;
})(inputs);
```

7. In the **Guardian Setting** field, select a guardrail of your choice.
   - Select **All** to handle the detection of both offensiveness and prompt injection together.
   - Select **Offensiveness** to enable the detection of offensive or harmful content.
   - Select **Prompt Injection** to enable the detection of prompt attacks.

   Note: If your LLM can evaluate both offensiveness and prompt injection in a single request, select All. If not, create separate guardrail records for each offensiveness and prompt injection.

8. In the **Connection And Credential Alias** field, select the alias that you want to integrate with your custom guardian.
9. Set up the custom guardian token limit.

   a. Navigate to **All** > **System Definition** > **Tables**.

   b. Select the **Gen AI Control Setting Data [sys_gen_ai_control_setting_data]** table.

   c. In the Related Links section, select **Show List**.

   d. Search for the **Custom Guardian Token limit** record in the **Gen AI Control Data** column.

   e. Update the **Value** field of the Custom Guardian Token limit record with the maximum number of tokens per request.
   The token limit defines the maximum number of tokens the guardrail can process in a single evaluation. If the request content exceeds this limit, the content splits into smaller parts. Each part is evaluated separately against the guardrail rules.

   For example, if the token limit is 1000 and the request content is 2500 tokens, the system splits it into three parts and evaluates each one individually.

Setup a Guardrail Service Provider

Select a guardrail service provider of your choice in Now Assist Guardian to monitor and detect Now Assist interactions for harmful, offensive, and prompt injection content.

Before you begin
Role required: sn_nowassist_admin.nsa_admin

About this task
Now Assist Guardian supports multiple guardrail service provider types to monitor and detect harmful content. Select a provider based on your organization's requirements. Only one provider can be active at a time.

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

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Settings**.
2. Navigate to **Now Assist Guardian** > **Guardrail service providers**.
3. In the **Set default provider** section, select a service provider from the list of guardrail service providers.
4. Select **Save and activate**.

Result

The selected guardrail service provider monitors and evaluates all Now Assist interactions for harmful, offensive, or prompt injection content.

Manage AI models

Access and select the LLM (large language model) provider used for various Now Assist skills. The selection impacts all the skills within the capability.

Before you begin

Role required: admin, sn_nowassist_admin.nsa_admin

This feature is available in Now Assist Admin version 6.2 and Yokohama patch 6 onwards.

To enable the model provider selection, ensure that the skill is available in that region. The configuration controls for the approved model providers and corresponding compliant skills for specific regions are approved in AI Control Tower by the AI steward. For example, to update the model provider for Now Assist Q&A Genius Results, ensure that all the skills under Conversational skills are available and active in the region.

As per the AI Control Tower settings, the model provider selection is available at multiple levels, including skill, skill group, and instance levels in Now Assist Admin. Therefore, in a situation where a skill is not compliant or a model provider is unavailable in a particular region, Now Assist Admin console presents the admin persona with alternate region scope options seeking approval permissions to proceed with model provider selection. For example: You opt to switch to another model provider to use a particular Now Assist skill. Consider, the region scope must be changed to global location for the skill to work. Here, you can proceed only after consenting to the region scope change.

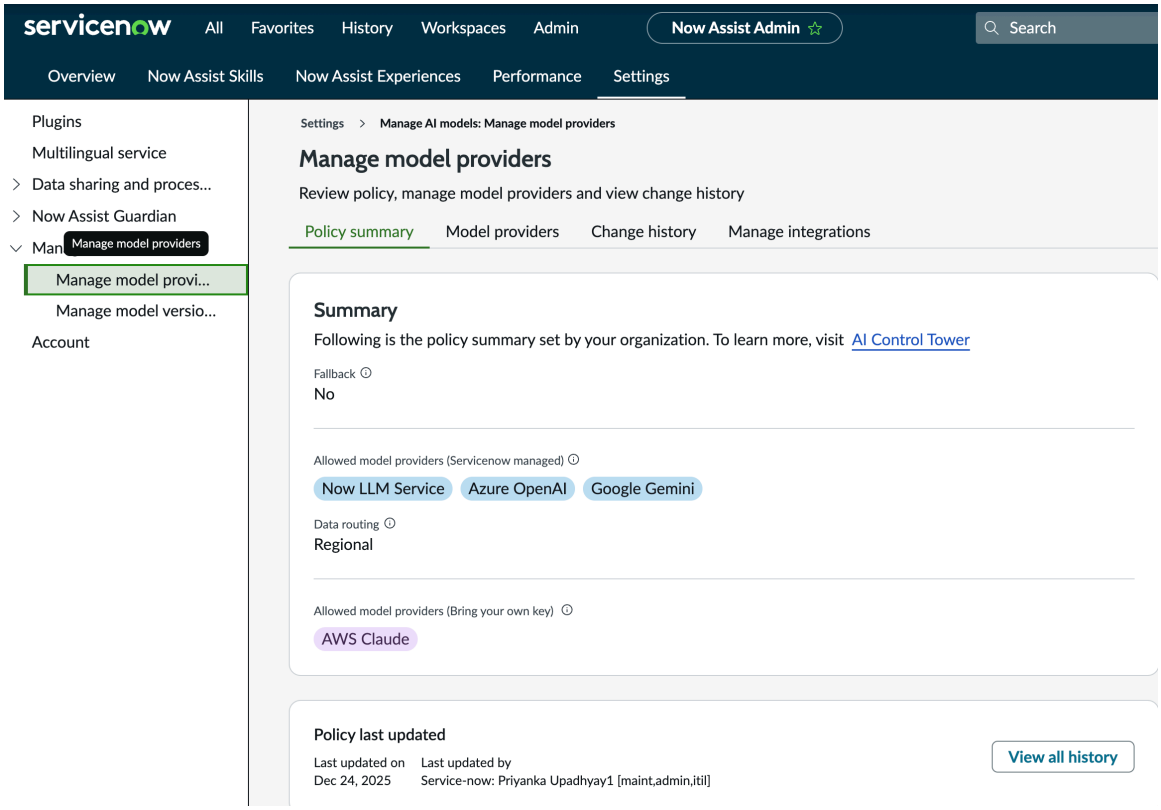
Note: For regulated markets:

- The **Manage LLM** option will not be available to customers on the Now Assist Admin user interface.
- Only Now LLM Service skills will be available to customers. Any skill that leverages external providers (e.g., Azure) will be hidden on the Now Assist Admin console. With the exception of NSC region, skills leveraging Azure OpenAI is available to the users.

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Settings**.
2. Navigate to **Settings** > **Manage model providers**.

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



3. Review these on the **Manage model providers** page:
   - Policy Summary set by your organization about skills using non-compliant model providers.

     Note: Note the routing and fallback exception details. These details are also derived from the AI Control Tower configurations.

   - Model providers assigned to the Now Assist skills and skill groups.
   - Policy and skill updates by the AI steward in AI Control Tower under the **Change History** tab.

Manage model providers

Edit or customise the model provider for a skill or skill group at the instance level from the list of supported third party model providers, including the default Now LLM Service. You can also review the model policy set by your organisation, and view the change history here.

Before you begin

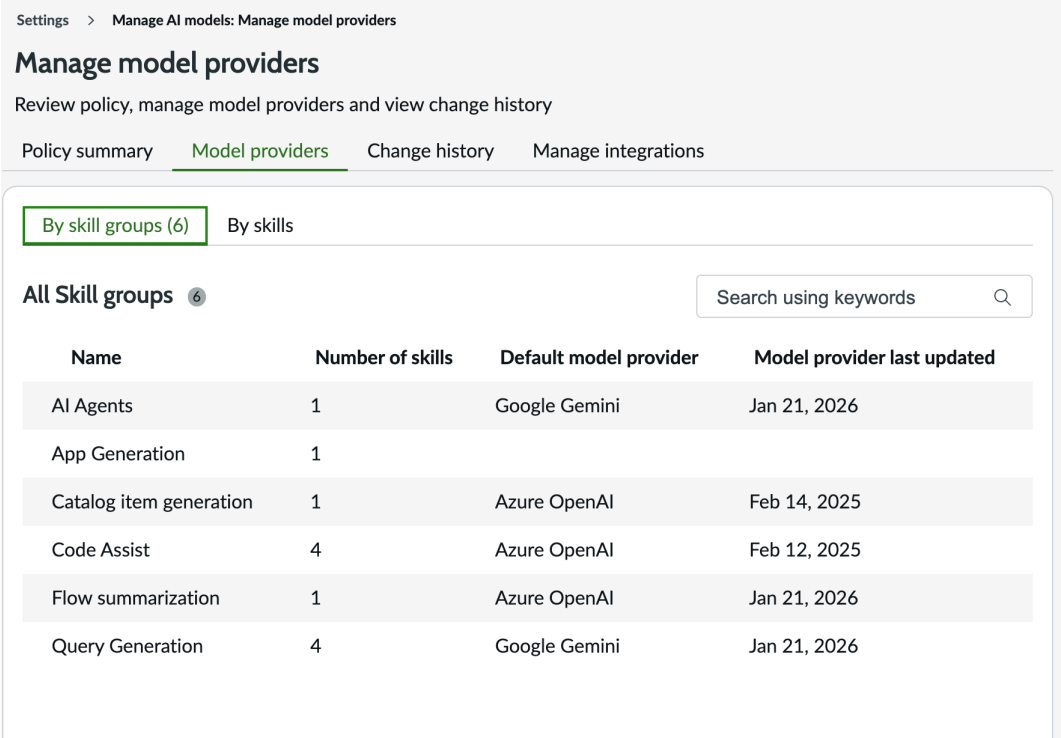
Role required: admin

Let's take an example of editing the model provider for skill group.

Procedure

1. Navigate to **All** > **Now Assist Admin**.
2. Navigate to **Settings** > **Manage AI models**.
3. Select **Manage model providers**.

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



4. Select **Edit model provider** to edit the model provider for the instance or customize it to update it at the skill or skill group level.
5. Select **Instance** scope to update the model provider for all the skills and skill groups in the instance.

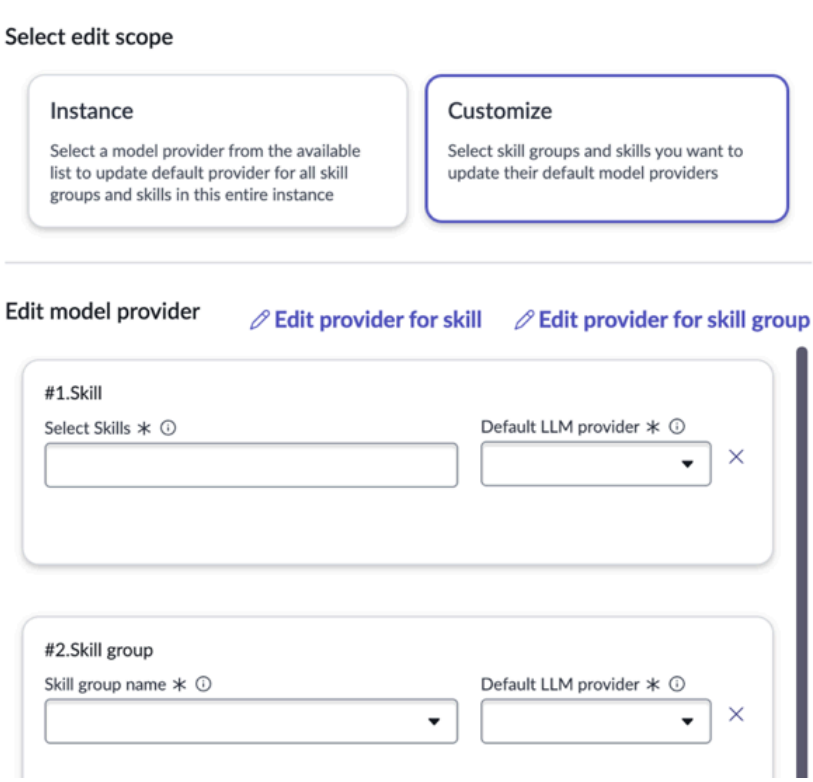
   a. Choose the **Default model provider** and select **Save and activate** to update your selection. The impacted skills and skill groups will display alongside.

   b. Select **Save** to update the model provider.
   You will receive a confirmation message if you have reviewed the change and the impact.

   Note: The agreement acceptance message is displayed only for the first time user.

6. Edit the model provider for a skill or skill group, after selecting the instance level model provider.
   This customization will override the instance level selection and update all capabilities associated with that group.
7. Select **Customize** to select the model provider for a particular skill or skill group.

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



a. Select the **Default model provider** for specific skill or skill group.
Select the skill group as **AI agents** to change the model provider for **Now Assist in Virtual Agent**.

b. Select **Save and activate** to update your selection.

Note: The Websearch AI search answers skill is only compatible with Google Gemini as the LLM provider. If you want to change the LLM provider to Azure OpenAI for this skill, see Configure AI search answers capability for web search for more information.

8. Select **Cancel, go back** to revert or **Yes, activate** to save the updates.
The status summary of skills for which the model provider change updated successfully or failed, displays.
9. Review and accept the consent modal message to accept the terms for changing the region scope of the selected model provider.
There are certain market regions restricted for particular model providers. The region scope must be changed to global location for the skills to be able to use the selected model provider.
10. Select **Yes, activate** to accept and apply the selection.
You receive a message confirming model provider selection. The selection will also notify if any of the skills within that capability are not updated with the selection.

Note: The skill and skill level customizations are overwritten once you update the model providers at the instance level.

You can also select **Cancel, go back** to review or reject the changes.

11. Follow similar steps for editing the model provider for a particular skill.

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

Manage Integration

Choose the preferred integration type for configuring the available model providers. There are two ways to configure a model provider in Now Assist Admin. You can either select Original Equipment Manufacturer (OEM) or Bring Your Own Key (BYOK).

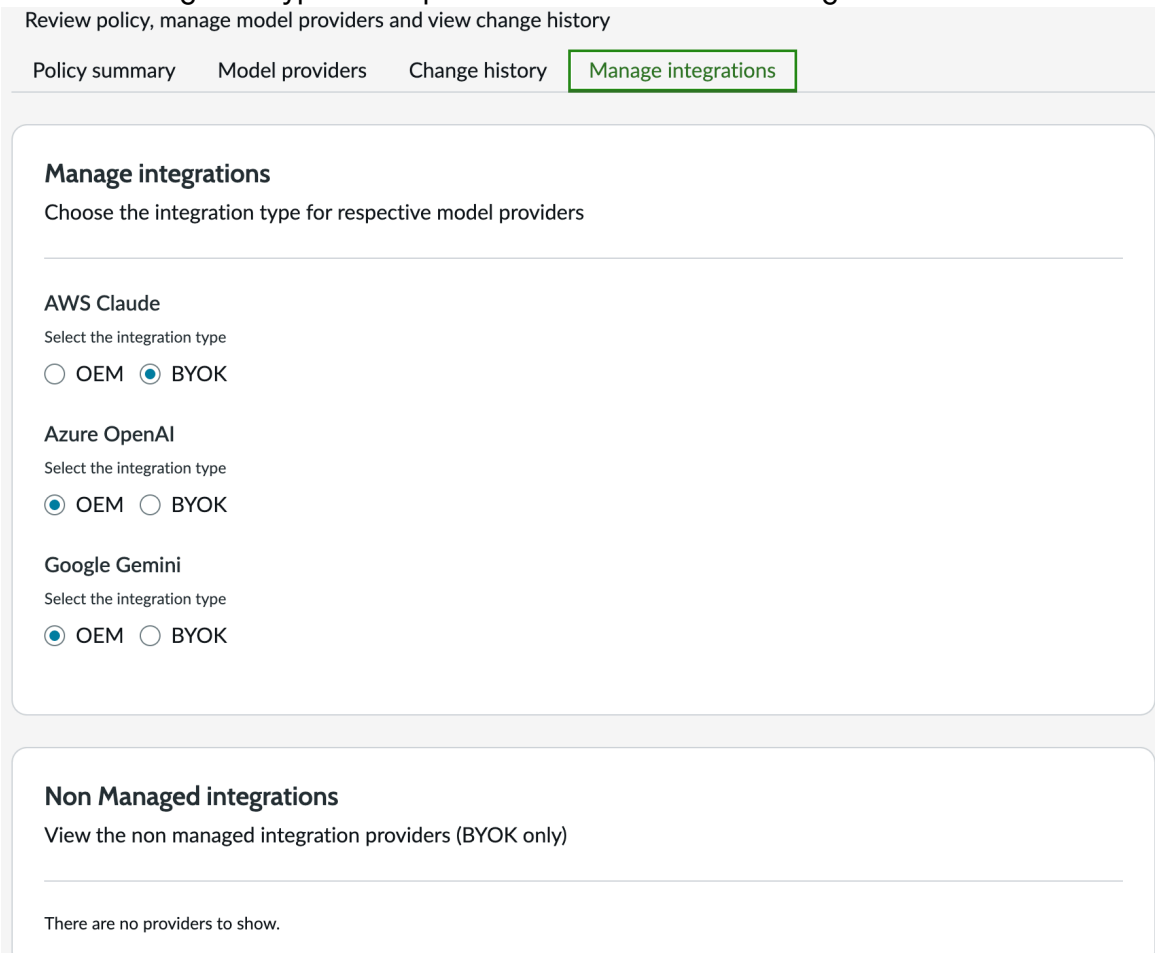
Before you begin

Role required: admin

With **Bring Your Own Key (BYOK)** option, you can apply your own license key to use the Generative AI application capabilities from the given provider. The alternate option is **OEM** where, ServiceNow® uses license from the provider to access the Generative AI application capabilities.

Procedure

1. Navigate to **All** > **Now Assist Admin**.
2. Navigate to **Settings** > **Manage integration**.
3. Change the model provider to use either the ServiceNow® key or **Bring Your Own Key (BYOK)** across all the third party model providers.
   This enables **Bring Your Own Key (BYOK)** selection for third party model providers across all features except AI Agents.
4. Select the integration type for the providers that allow such a configuration.



You can find the providers that allow BYOK type integration only, under **Non Managed integration** section.

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

Manage version

Manage the version of the model providers across skills and instance levels. You can change and update versions for the out-of-box and custom skills.

Before you begin

Role required: admin

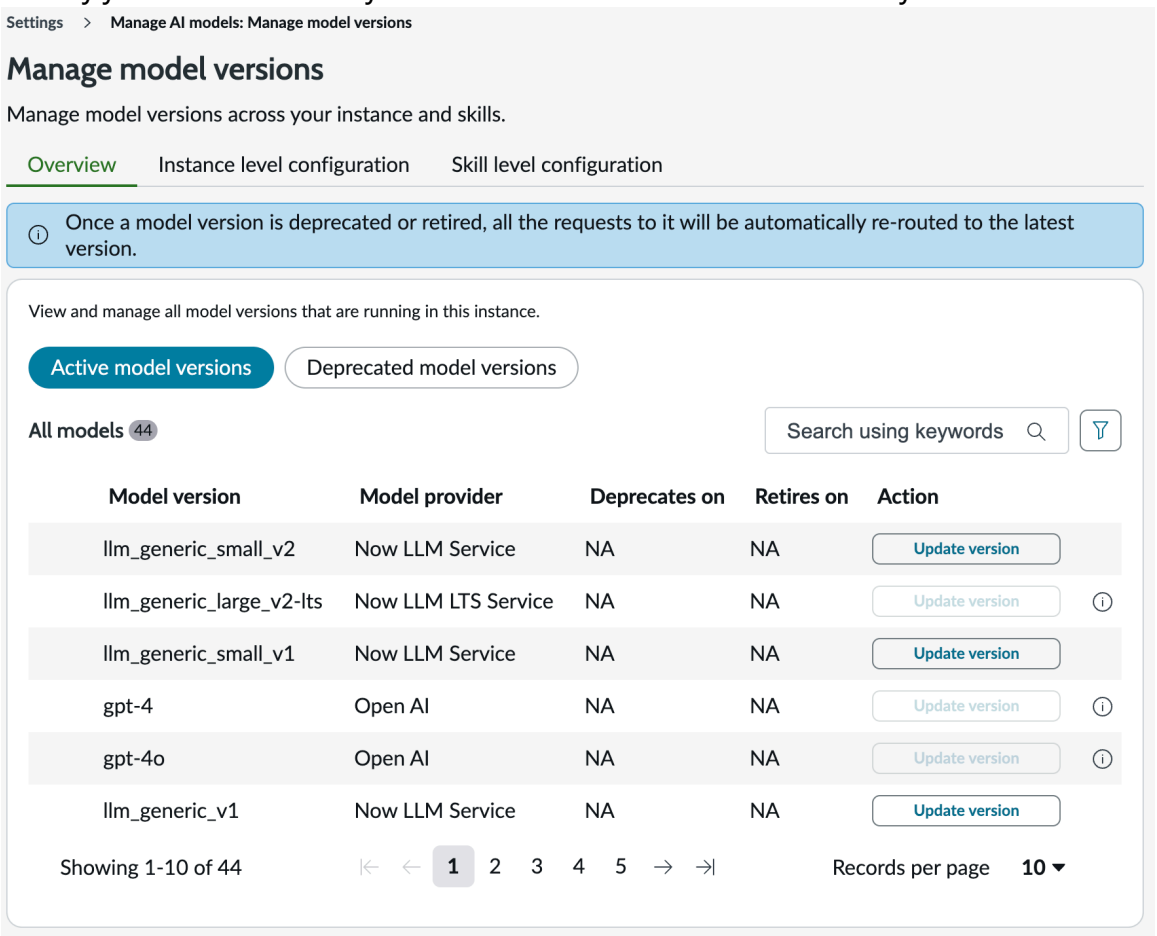
See Default and target model version to know more about default and target model versions.

Procedure

1. Navigate to **All** > **Now Assist Admin**.
2. Navigate to **Settings** > **Manage AI models**.
3. Select **Manage model version**.
4. View the summary of the active, deprecated, retiring model versions, and recommended actions, providing effective version management system to the admins.

   Note: A model version will not be available for selection as target model version, if it is in deprecated, retired, in review or rejected state.

5. Identify your use case and key factors to choose the model that best fits your needs.

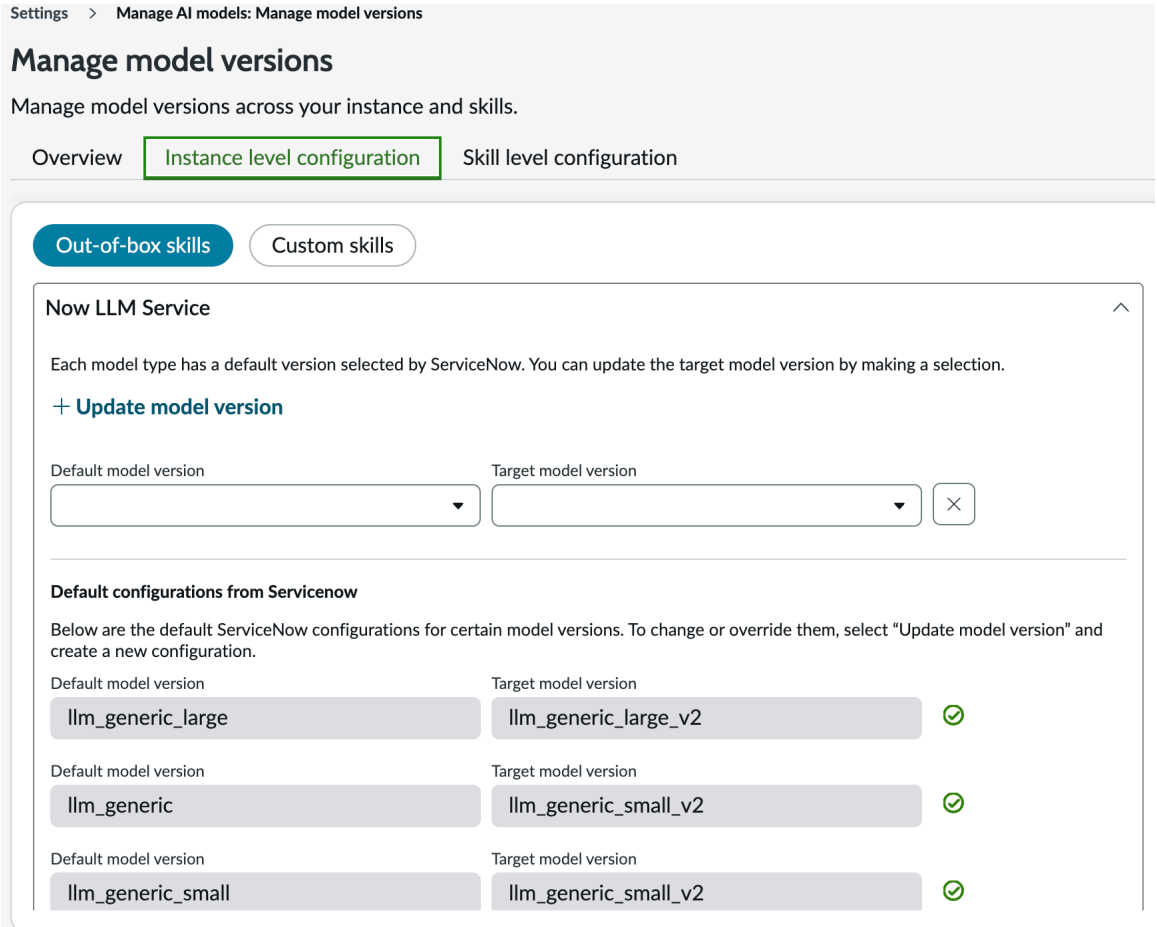


6. View or update the model version instantly under the **Overview** tab.
7. You can edit the model provider version for out-of-box or custom skills for the instance or customize it to update it at the skill or skill group level, under the respective tabs.

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

Customizing the model version for skills replaces the instance-level model version currently assigned to each provider. This action is typically reserved for specific situations.

8. Select **Instance level configuration** to manage the model provider version for the skill across the instance.



   a. Select **Out-of-box skills** to set the target model version against the current version of the applicable model providers, for the default skills.

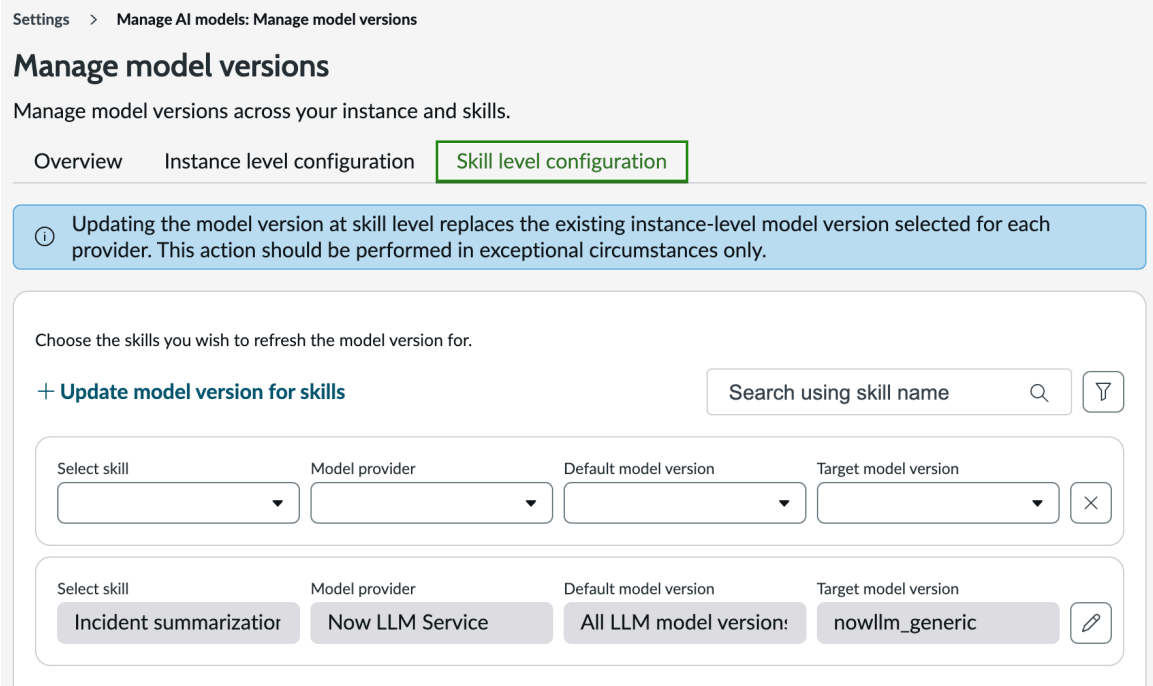
   b. Select **Custom skills** to set the target model version against the current version of the applicable model providers, for the custom skills.

   c. Select **Save and activate** to update your changes.

   d. Select **Update model version** to override ServiceNow shipped mappings or create a new default and target model versions mapping.

9. Select **Skills** to manage the model provider version at the skill level.

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



a. Search or select the skills you wish to refresh the model version for, and the applicable model provider.

b. Select the default and the target model provider version.

c. Select **Save and activate** to update your changes.

d. Select **Update model version for skills** to override ServiceNow shipped mappings or create a new default and target model versions mapping, at the skill level.

Review Now Assist account

Review your Now Assist license details on the Account page of the Now Assist Admin console to make sure that you're up to date on what's available to you.

Before you begin

Role required: sn_generative_ai.nsa_admin

About this task

From the Account page, you can review the licensing information to verify which plugins and features you're entitled to and what their status is.

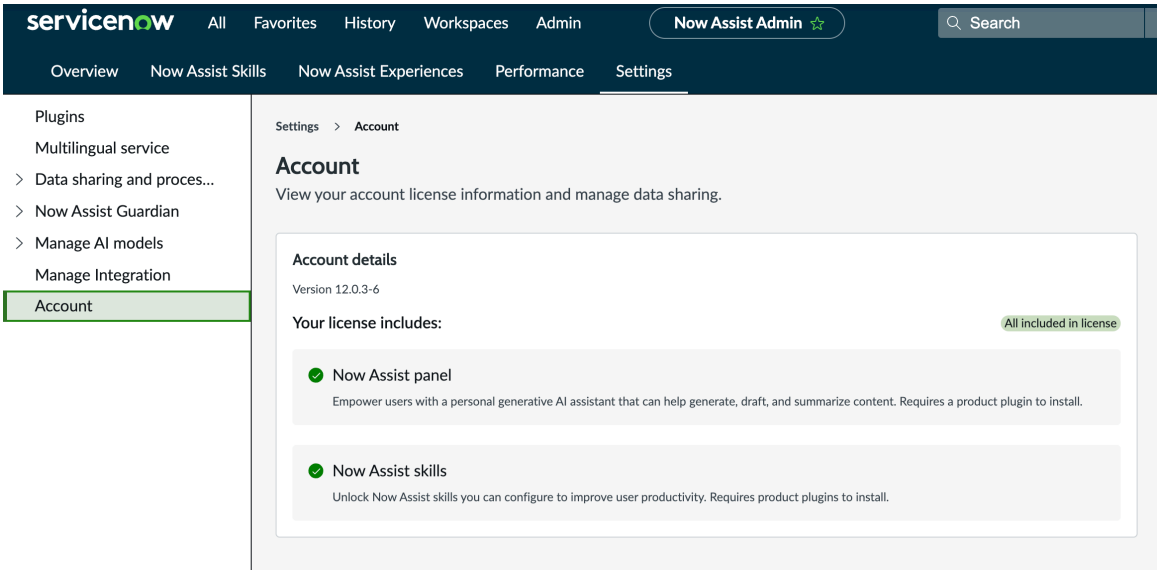
Procedure

1. Navigate to **All** > **Now Assist Admin** > **Settings**.

If you're already in the Now Assist Admin console, select the **Settings** tab.

2. In the Settings page, select **Account**.

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



3. In the Account details section of the panel, review what is included with your Now Assist license.

What to do next

Turn on the skills that you want for your Now Assist workflow applications so that you can use generative AI capabilities across the ServiceNow AI Platform. For more information, see Configuring Now Assist Admin features.

Configuring Now Assist Admin features

Use the Now Assist Admin console to activate the various Now Assist applications and skills that you're entitled to.

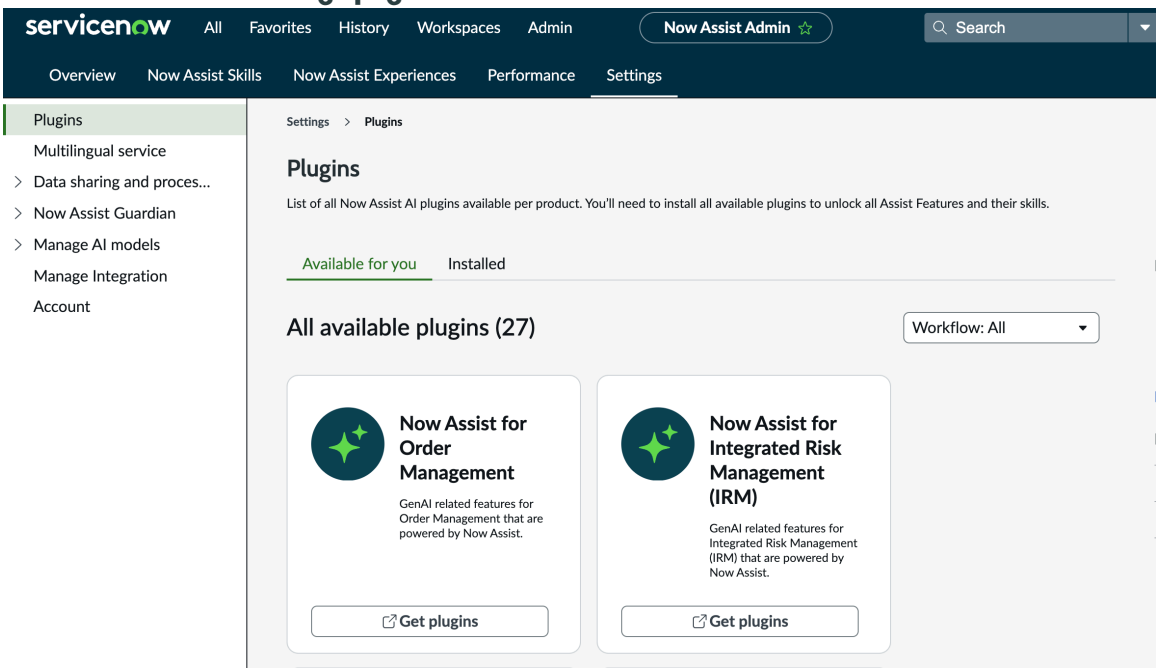
Configuration overview

Install and activate the plugins, turn on the Now Assist panel, and view the account settings. You must install at least one Now Assist application before you can configure any skills.

The following example shows the Settings page with four available plugins to install, including Now Assist for HRSD and Now Assist for Creator.

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

Now Assist Admin settings page



Install Now Assist plugins

Install Now Assist plugins to enable generative AI on your instance.

Before you begin

Role required: admin

Follow these instructions to get started with Now Assist Admin:

1. To get started with Now Assist, you must install at least one Now Assist application on your instance.
2. Activate any Now Assist plugin from the ServiceNow Store to access Now Assist Admin.
3. Install additional plugins, from the Now Assist Admin console.
4. The Now Assist Admin console guides your implementation, starting with installation.
5. Check out the Now Assist Journey Checklist for more information.

About this task

Then, activate any Now Assist plugin from the ServiceNow Store to access Now Assist Admin. to install additional plugins, you must first activate any Now Assist plugin from the ServiceNow StoreThe Now Assist Admin console can guide your implementation, starting with installation. Check out the Now Assist Journey Checklist for more information.

Procedure

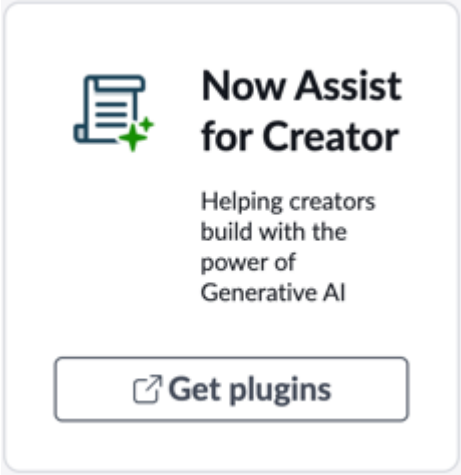
1. Navigate to **All** > **Now Assist Admin** > **Settings**.

   If you're already in Now Assist Admin, select the **Settings** tab.

2. On the **Settings** page, select **Plugins**.

   Plugins appear as cards. Review all Now Assist plugins on the **Available for you** tab. Plugins that you have already installed appear on the **Installed** tab.

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



3. Select **Get plugins** on the card for the plugin that you want to install.
4. In the confirmation window, select **Install Plugin** to open the ServiceNow Store page for the plugin in a new browser tab.
5. Install the plugin from the ServiceNow Store page.

   Some applications may require you to request the app from the ServiceNow® Store first. After you've requested the application from the ServiceNow® Store page, navigate to **All** > **System Applications** > **All Available Applications** > **All** to finish the installation.
6. Return to the Now Assist Admin console.
7. In the dialog box, select **Refresh**.
8. Either close the dialog box to review all of your plugins or select **View all (Plugin) Assists and Skills** to review the features of your new plugin.

Result
Your plugin is successfully installed.

If you encounter issues installing or updating applications, see this knowledge article for steps that may address your issue. Otherwise, you can make a Support case.

What to do next
Activate the Now Assist panel standard chat or Activate a Now Assist skill

Activate the Now Assist panel standard chat

Activate the Now Assist panel standard chat to enable your agents to use Now Assist skills, such as task summarization or navigation, in a side panel on the user interface.

Before you begin
You must install at least one Now Assist application before you can turn on the Now Assist panel.

Role required: sn_generative_ai.nsa_admin

About this task
You must have the now_assist_panel_user role to have access to the Now Assist panel once you turn it on.

To learn more about the Now Assist panel, and how it can assist your agents, see Now Assist panel.

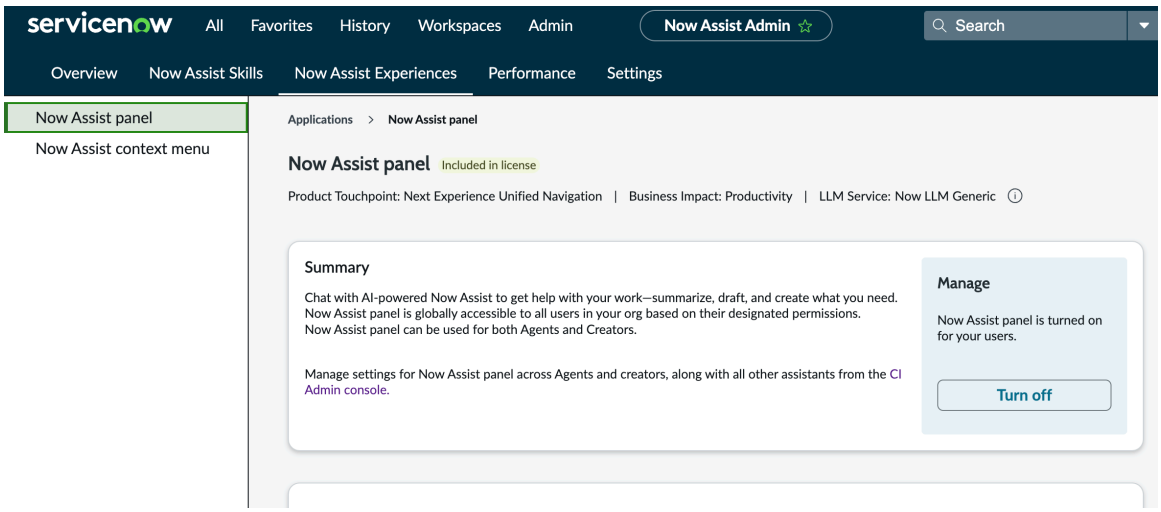
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

Procedure

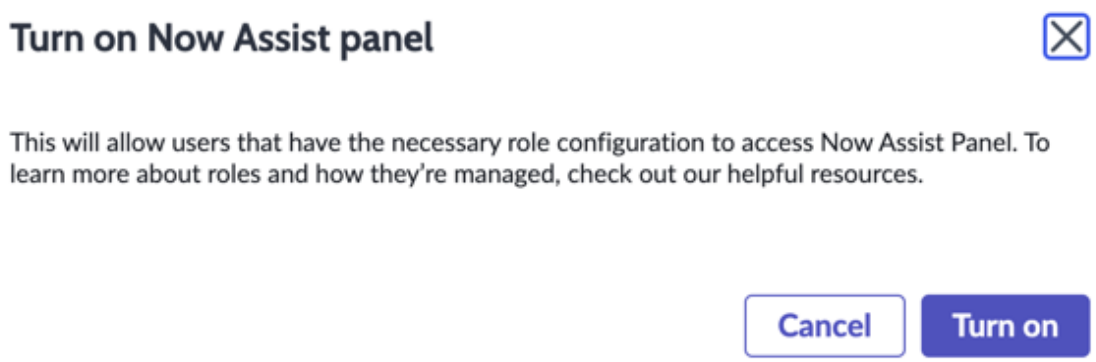
1. Navigate to **All** > **Now Assist Admin** > **Now Assist Experiences**.

   If you're already in Now Assist Admin, select the **Now Assist Experiences** tab.

2. Select **Now Assist panel**.
3. Enable the Now Assist panel on your instance by selecting **Turn On**.



4. In the Turn on Now Assist panel dialog box, select **Turn on**.



5. If you want to use assistants, you must activate them.
   See Activate Now Assist panel assistants for information on activating assistants.

What to do next

For examples of the Now Assist panel in action for Now Assist applications, see the following topics:

- Summarize a chat conversation by using Now Assist for Customer Service Management (CSM)
- Summarize a Sidebar discussion by using Now Assist for IT Service Management (ITSM)
- Summarize an issue using Now Assist for Integrated Risk Management (IRM)
- Generate a knowledge article from HR Agent Workspace with Now Assist for HRSD

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

Activate Now Assist panel enhanced chat

Activate the Now Assist panel enhanced chat to enable your agents to use Now Assist skills, such as task summarization or navigation, in a side panel on the user interface.

Before you begin

You must install at least one Now Assist application before you can turn on the Now Assist panel.

Role required: sn_nowassist_admin.nsa_admin

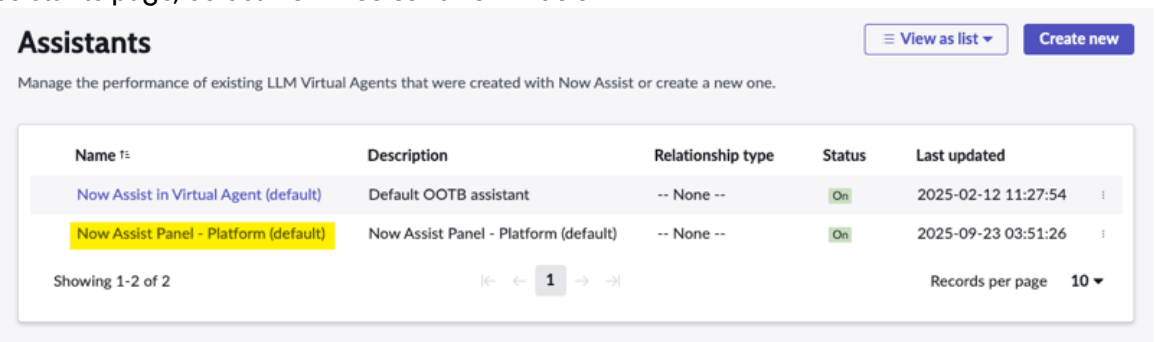
About this task

You must have the now_assist_panel_user role to have access to the Now Assist panel once you turn it on.

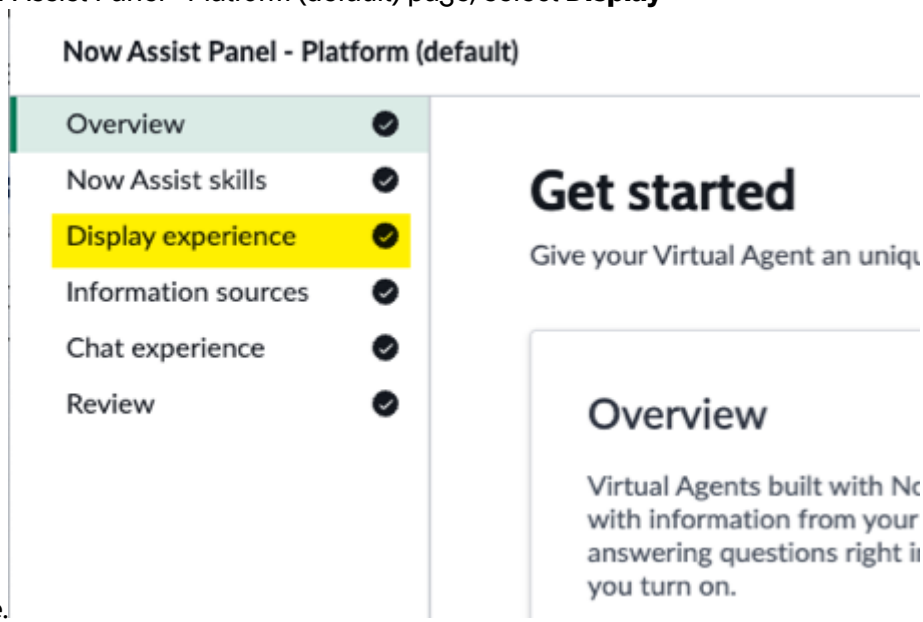
To learn more about the Now Assist panel, and how it can assist your agents, see Now Assist panel.

Procedure

1. Navigate to **All** > **Conversational Interfaces** > **Assistants**.
2. On the Assistants page, select **Now Assist Panel - Platform (default)**.



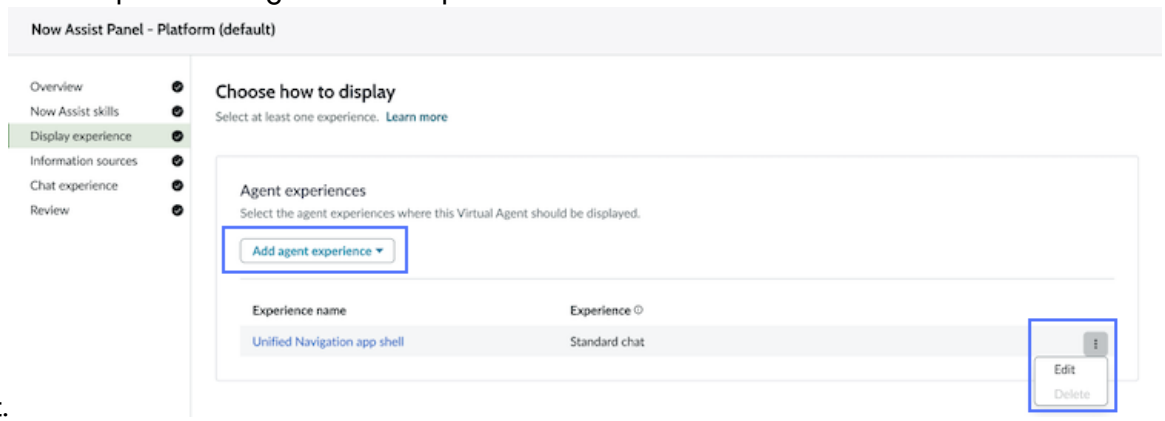
3. On the Now Assist Panel - Platform (default) page, select **Display experience**.



4. Select the **Add agent experiences** drop-down menu and then select **Unified Navigation app shell**.

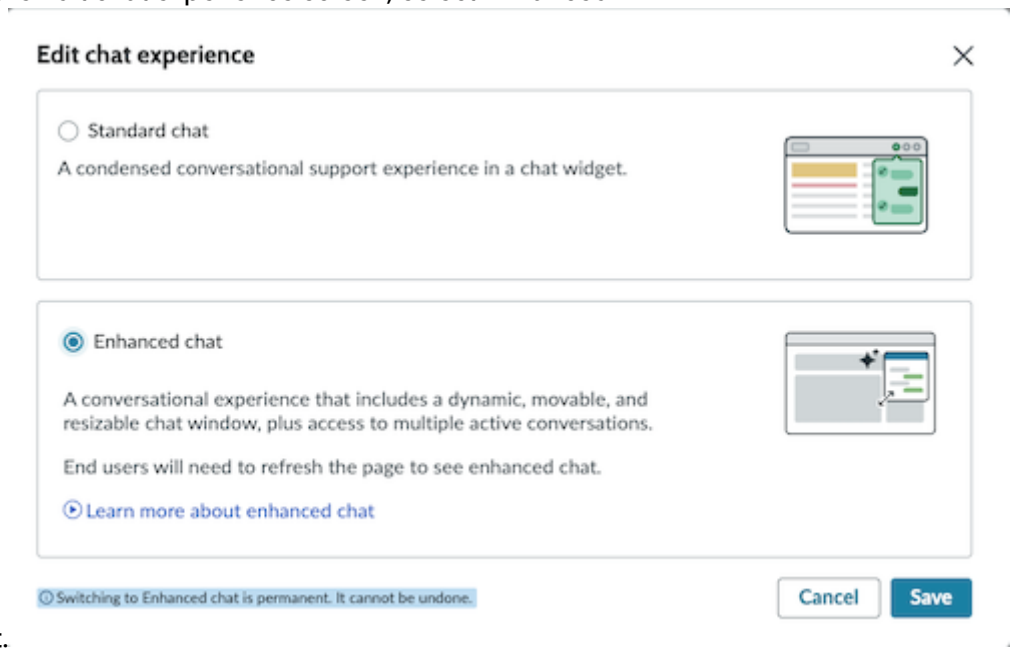
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

5. Select the ellipsis to change the chat experience from standard chat to enhanced chat.



Important: Switching to enhanced chat is permanent and can't be undone after you have selected **Save** and then **Save and continue**.

6. On the Edit chat experience screen, select Enhanced chat.



7. Select **Save**.
8. If you want to use assistants, you must activate them.
   See Activate Now Assist panel assistants for information on activating assistants.

What to do next

For examples of the Now Assist panel in action for Now Assist applications, see the following topics:

- Summarize a chat conversation by using Now Assist for Customer Service Management (CSM)
- Summarize a Sidebar discussion by using Now Assist for IT Service Management (ITSM)
- Summarize an issue using Now Assist for Integrated Risk Management (IRM)
- Generate a knowledge article from HR Agent Workspace with Now Assist for HRSD

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

Activate Now Assist panel assistants

You must activate Now Assist panel assistants before using them.

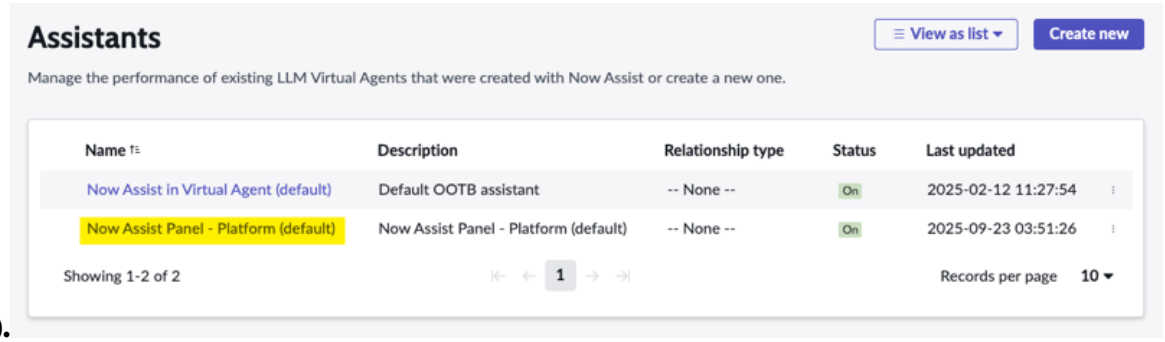
Before you begin

Role required: admin

Note: Voice input is automatically activated when you activate the Now Assist panel.

Procedure

1. Navigate to **All** > **Conversational Interfaces** > **Assistants**.
2. On the Assistants screen, select **Now Assist Panel - Platform (default)**.



   The Now Assist Panel - Platform (default) page displays.
3. If the Review tab does not display on the Now Assist Panel - Platform (default) page, select **Review**.
4. On the Review tab, activate the Now Assist panel assistant.

Enable voice input for Now Assist panel

Give users the option to use their voice when interacting with the Now Assist panel to make the panel more accessible. Voice input enables you to use the panel without needing to use a keyboard.

Before you begin

Note: Voice input is automatically activated when the Now Assist panel is activated. As of the Zurich Patch 4 release, voice input is configured in Enable additional chat features and not with this option.

You must have installed at least one Now Assist application with a skill that uses the Now Assist panel. See Now Assist panel for more information about supported skills.

Role required: sn_generative_ai.nsa_admin

About this task

You can give users the option to use voice input in the Now Assist panel. This feature provides an additional input method to interact with Now Assist skills in English. Once it's enabled, users can choose to activate this feature in their personal accessibility preferences by toggling on **Enable voice input for the Now Assist panel**. See Configure Next Experience accessibility preferences for more information about setting personal accessibility preferences.

Voice-to-text input can help users with mobility impairments access generative AI skills without using a keyboard. This feature can also be useful to blind or low-vision users, neurodivergent users, non-native language speakers, and mobile users on the go, such as field service agents.

The voice input feature is not supported in regulated markets.

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

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Now Assist Experiences**.

   If you're already in the Now Assist Admin console, navigate to the Now Assist Experiences page.

2. Go to the Now Assist panel tab.
3. In the Settings section, turn on the toggle for **Voice Input**.

Result

Users can choose whether they can use their voice to interact with the Now Assist panel in their Next Experience accessibility preferences.

Using Now Assist Admin

Use Now Assist Admin to explore the various Now Assist plugins, skills and associated Generative AI application features you are entitled to.

Activate a Now Assist skill

Configure the triggers, settings, and display locations for Now Assist skills to enable GenAI capabilities across the ServiceNow AI Platform.

Before you begin

Role required: sn_generative_ai.nsa_admin

About this task

Activate the skills that are most relevant to your use cases and business needs. For a full list of available skills, see Now Assist skills. After the skills have been activated, they're accessible across the ServiceNow AI Platform based on the availability and display settings you choose.

Procedure

1. Navigate to **All** > **Now Assist Admin Console** > **Features**.

   If you're already in the Now Assist Admin console, select the **Now Assist Features** tab.

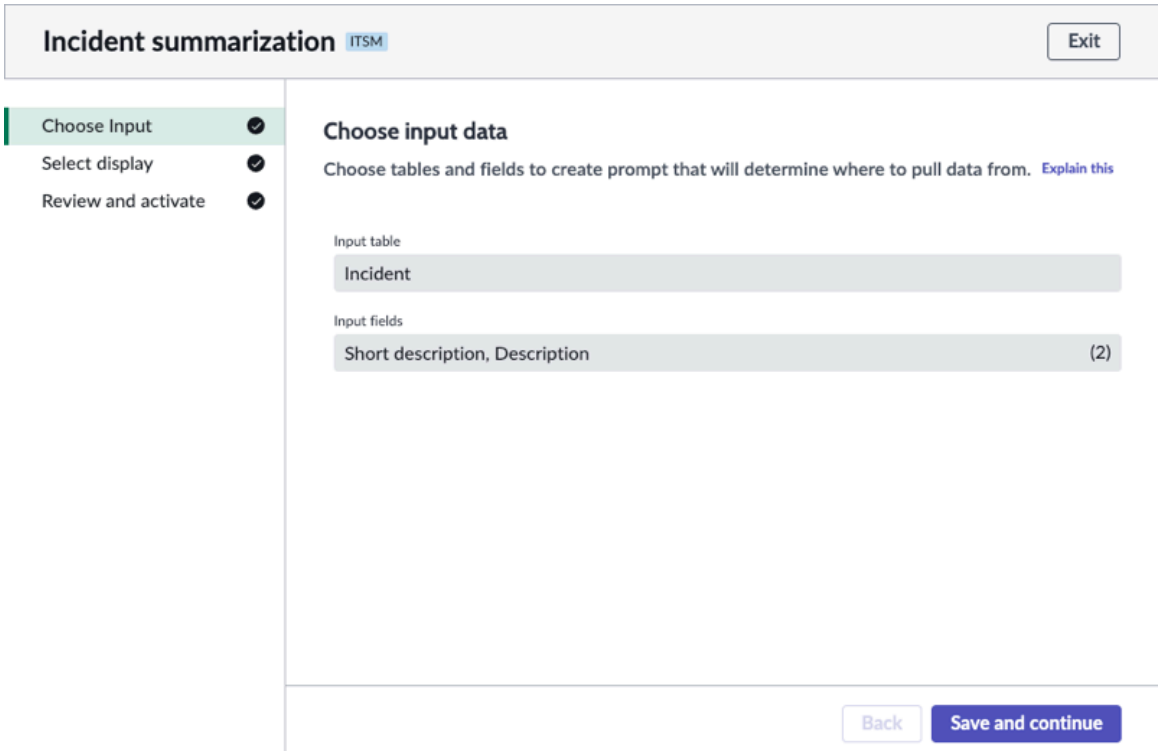
2. On the navigation panel, select a workflow, such as **Technology**.

   Each workflow contains feature sets.

3. On the feature card that is associated with the skill you'd like to activate, select **View details**.
4. In the All available skills section, select **Activate Skill**.
5. In the first step of the skill configuration, determine which inputs or triggers that you want to associate with the skill.

   Each skill configuration has steps that are shown in the guided setup. The exact steps vary from skill to skill. A symbol next to each step indicates whether the step is completed, partially completed, or not completed. After configuring a step, select **Save and continue** to go to the next step. Return to a previous step by selecting **Back**.

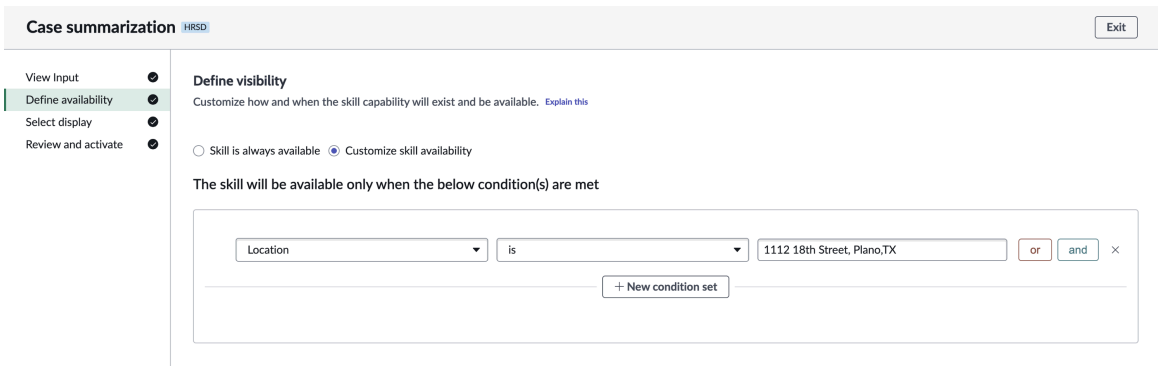
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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



Note: Some configuration options are read only.

6. After you've configured the current step, select **Save and continue** to go to the next step.
7. **Optional:** For some skills, the next step is to define the availability.

   (Optional) You can select **Skill is always available** if you do not want to place any restrictions on when the skill is available for use. If you want to add conditions, select **Customize skill availability**. Selecting this option opens up a condition builder for you to select fields and values that determines whether someone can use the skill.



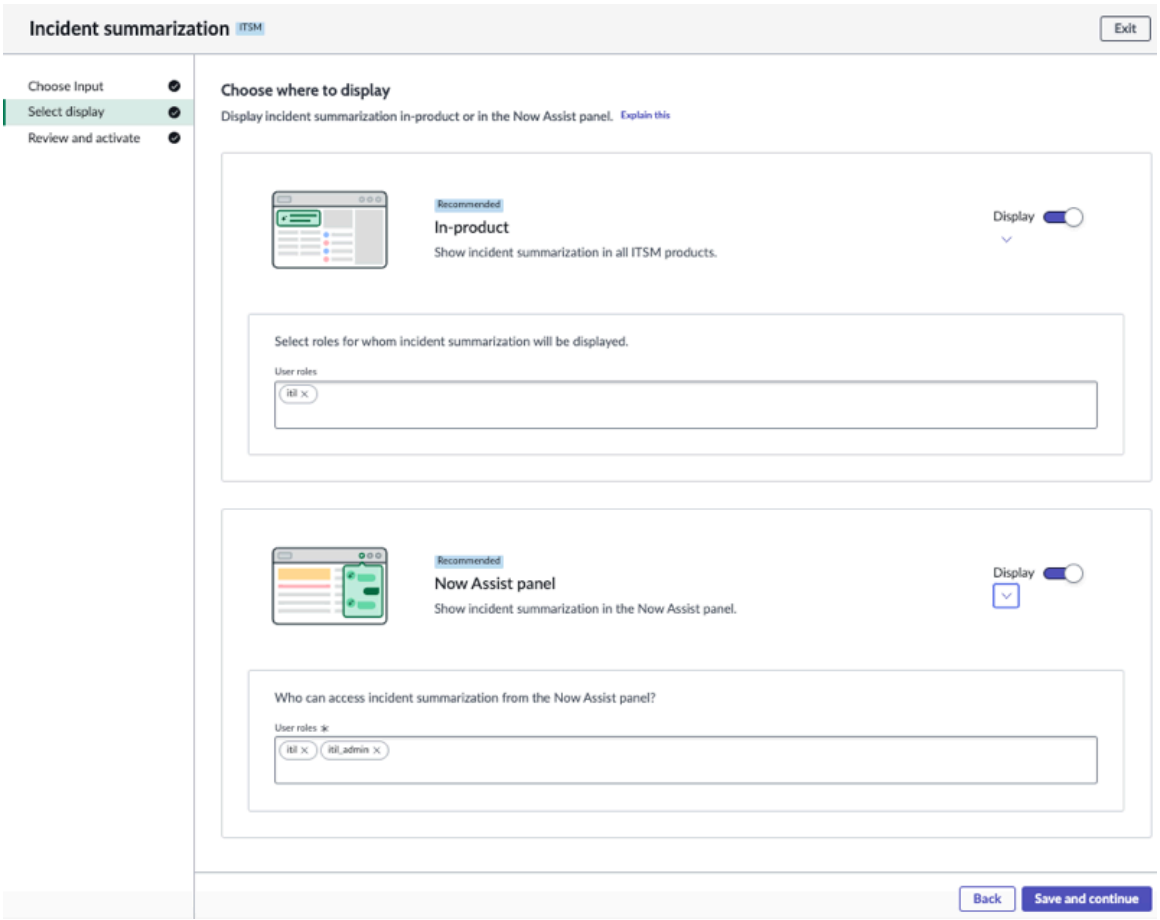
8. In the next step of the skill configuration, select where you'd like to display the skill.

   Options vary from skill to skill. Some options are only available for certain skills.

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

- **In-product desktop**: When selected, Now Assist skills are displayed on forms and workspaces.
- **Now Assist panel**: When selected, Now Assist skills are available in the Now Assist panel. If you don't see this option, you must activate the Now Assist panel. For more information, see Activate the Now Assist panel standard chat.
- **Core UI**: When selected, the Now Assist skill will display as a UI action in the Core UI.

Select the down arrow next to the Display toggle to select the roles that can use the skill. Roles can be added by entering the name of the role in the User roles field. Existing roles can be removed by selecting the X icon in the role bubble. You must have at least one role specified, but you can add as many as you like.



9. Review your choices and select **Activate** to complete the configuration.

What to do next

Use the Now Assist applications and skills that you've activated.

Configure chat summarization and chat reply recommendation skills in the Now Assist Admin console

Define the triggers, inputs, and display location for chat summarization and chat reply recommendation by using the guided setup in the Now Assist Admin console. The activation steps are conceptually same for both the skills.

Before you begin

Role required: sn_generative_ai.nsa_admin

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

Procedure

1. Navigate to **All** > **Now Assist Admin Console** > **Features**.

   If you're already in the Now Assist Admin console, select the **Now Assist Features** tab.

2. On the navigation panel, select a workflow that has chat recommendation, either **Technology** or **Customer**.

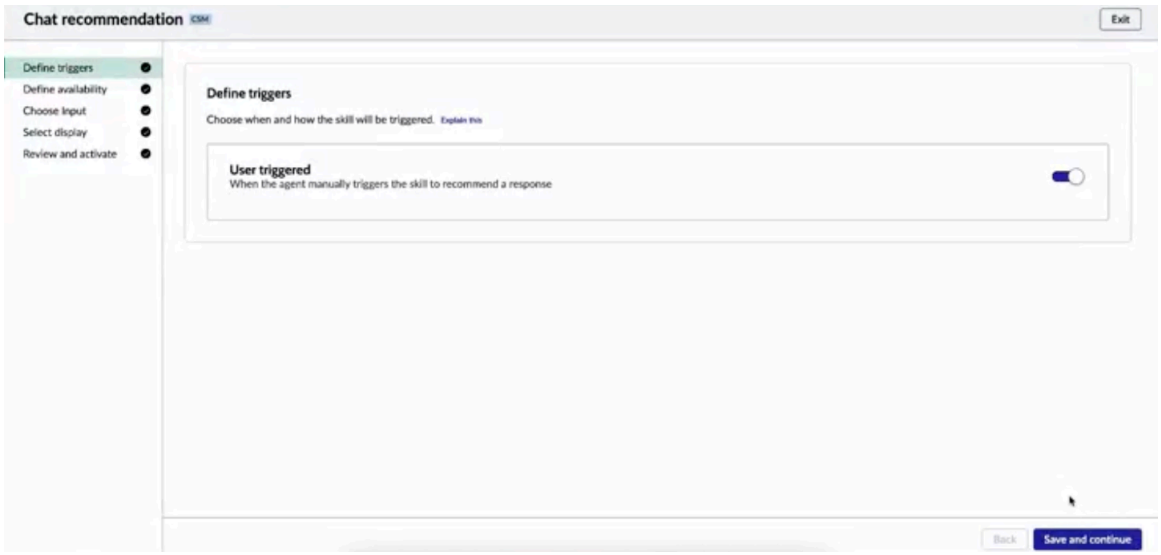
   Each workflow contains feature sets.

3. On the Chat feature card, select **View details**.
4. In the All available Chat skills section of the chat recommendation card, select **Activate Skill**.
5. Go to **Define Trigger**, the first step in the guided setup.

   By default, many of the options in the setup are configured for the most common use cases. You might need to select the step in the guided setup navigation to go back and change the configurations in previous steps. You can also use **Back** to navigate through the steps.

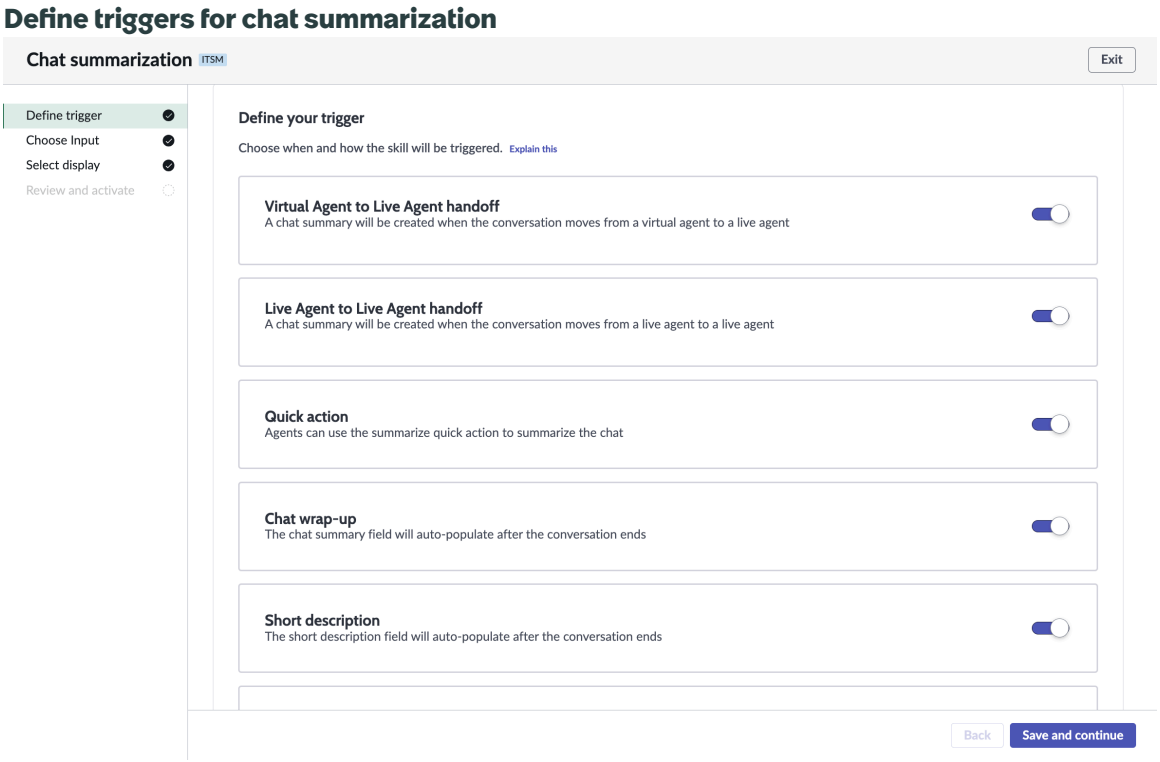
6. Using the toggles, select the actions trigger the chat recommendation skill.

Define triggers for chat recommendation



7. Using the toggles, select the actions trigger the chat summarization skill.

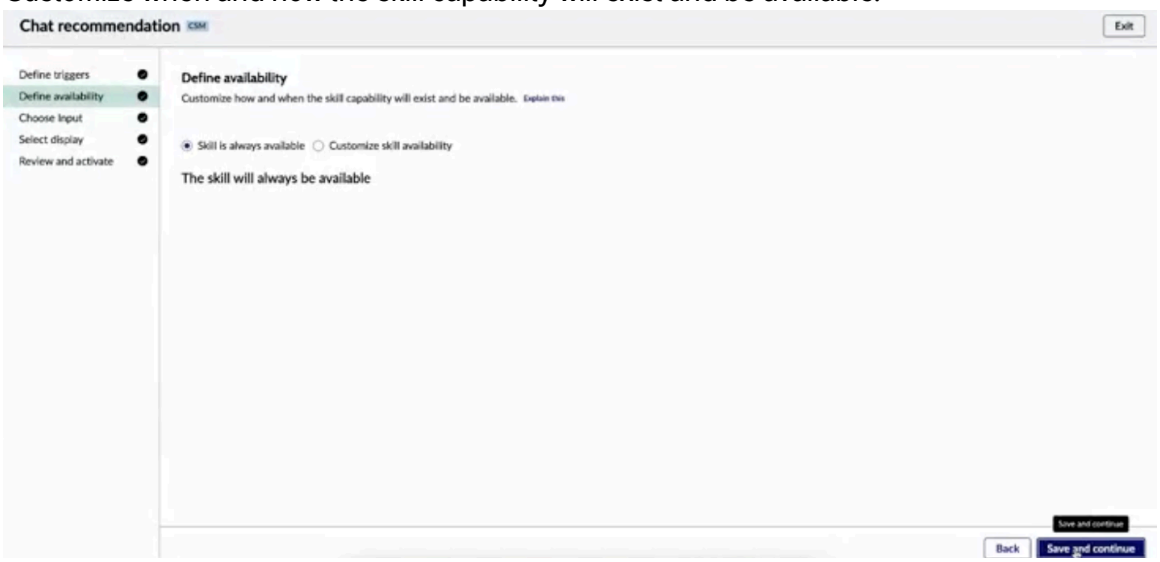
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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



8. Select whether you want the summary to be formatted with bullet points.

   By default, the summaries are written with bullet points, but you can turn off this format so that the generated summary uses paragraphs instead.

9. Go to **Define availability**, the next step, by selecting **Save and continue**.
10. Customize when and how the skill capability will exist and be available.

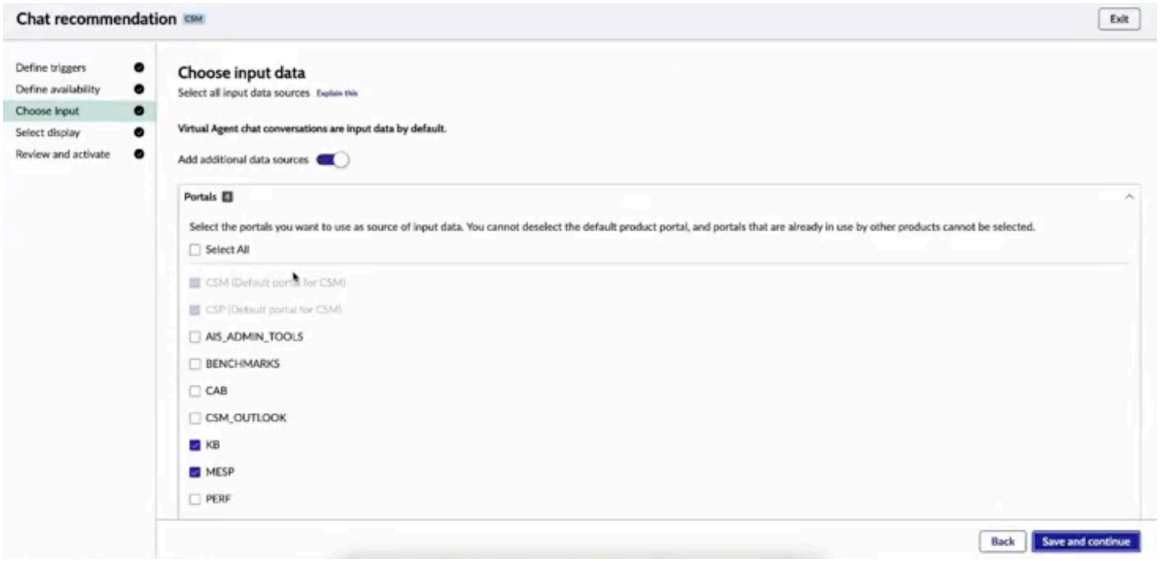


11. Select **Customize skill availability** if you want to define the skill to be available for a certain domain.
12. Go to **Choose Input**, the next step, by selecting **Save and continue**.
13. Select any additional data sources that you want the Large language model (LLM) to take into account when generating a recommendation.

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

14. Select a portal for the data source to allow chat summarization/recommendation to be generated for the conversation occurring on that portal.

    This is a mandatory step. The admin must specify a portal and enable a specific channel on **Choose Input** page, to enable the skill for chats sent in the selected portal/channel. Else the agent will receive an error message, "Chat summaries won't appear until your IT administrator completes all the required steps involved in the setup".



15. Select **Save and continue**.
16. Go to **Select display**, the last step, and select where you would like to display the skill.

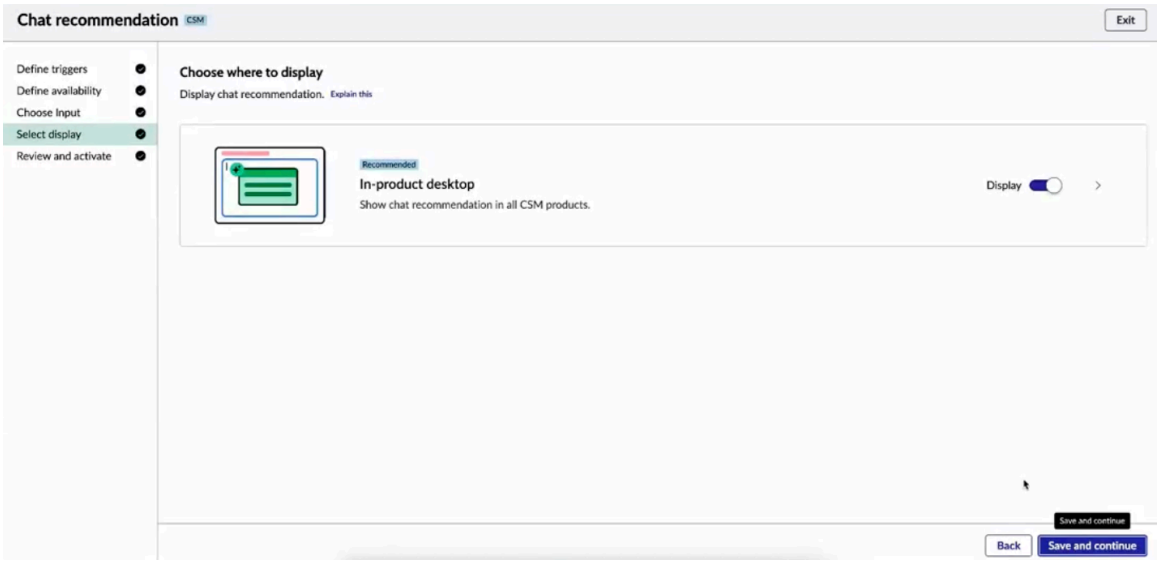
    You can select both in-product, Now Assist panel, or both.

    Note: Chat recommendation is not available in the Now Assist panel.

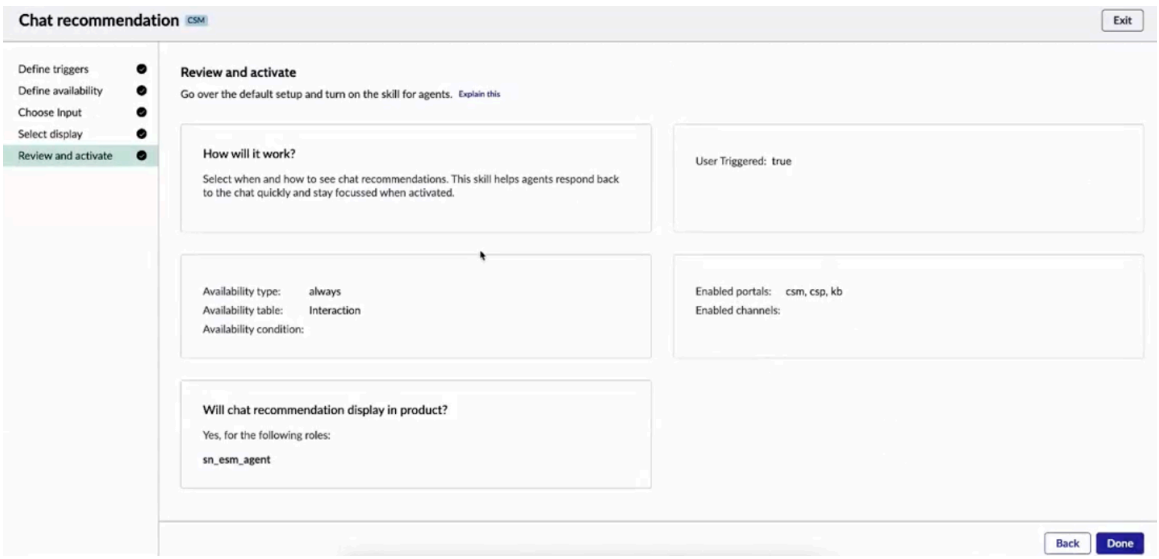
    - **In-product desktop**: When selected, Now Assist skills are displayed on forms and workspaces.
    - **Now Assist panel**: When selected, Now Assist skills are available in the Now Assist panel. Select the down arrow to identify the roles that can use the skill. Select the arrow next to toggle, to select roles who can access the skill. You can add roles by entering the name of the role in the **User roles** field. You can remove existing roles by selecting the X icon in the role bubble. You must have at least one role specified, but you can add as many roles as you like.

    Note: You can use different roles for chat recommendation in different workflows. You can see which workflow you're configuring by checking the label next to the skill name at the top of the guided setup, such as "ITSM" or "HRSD."

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



17. Review your choices and complete the configuration by selecting **Activate**.



Result

Chat recommendation or reply recommendation for the workflow is active on the instance.

What to do next

Analyze your skill performance and usage on the Now Assist Admin console to help determine the success of the skill. Learn more about tracking your Now Assist usage at Monitoring Now Assist usage in Subscription Management.

Configure email reply recommendation in the Now Assist Admin console

Configure the email recommendation Now Assist skill to enable agents to draft email replies based on contextual information.

Before you begin

Role required: sn_generative_ai.nsa_admin

About this task

The email recommendation skill is available in multiple workflows. The exact steps and the order in which they appear for the guided setup for the skill vary depending on the workflow.

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

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Now Assist Features**.

   If you're already in the Now Assist Admin console, select the **Now Assist Features** tab.

2. On the navigation panel, select a workflow, such as **Technology**.

   Each workflow contains feature sets. The available workflows for email recommendation are Customer (CSM) and Employee (HRSD).

3. On the feature card that is associated with the skill you'd like to activate, select **View details**.
4. In the All available skills section, select **Activate Skill**.

   If you have already activated the skill and want to edit the configuration, in the Active skills section, select the more options menu item (⋮). Then select **Edit**.

5. In the **Write with Now Assist** step, choose whether you want the Now Assist context menu to be active and then select the actions you want to be available.

   After each step, select **Save and continue**. For more information about these choices, see Skill inputs and triggers for Now Assist for HRSD.



6. Choose when you want the skill to be available.

   You can select **Skill is always available** if you don't want to place any restrictions on when the skill is available for use. If you want to add conditions, select **Customize skill availability**. Selecting this option opens up a condition builder for you to select fields and values that determine whether someone can use the skill.

7. Choose the inputs or email parameters, including input data and whether you want to display template suggestions.

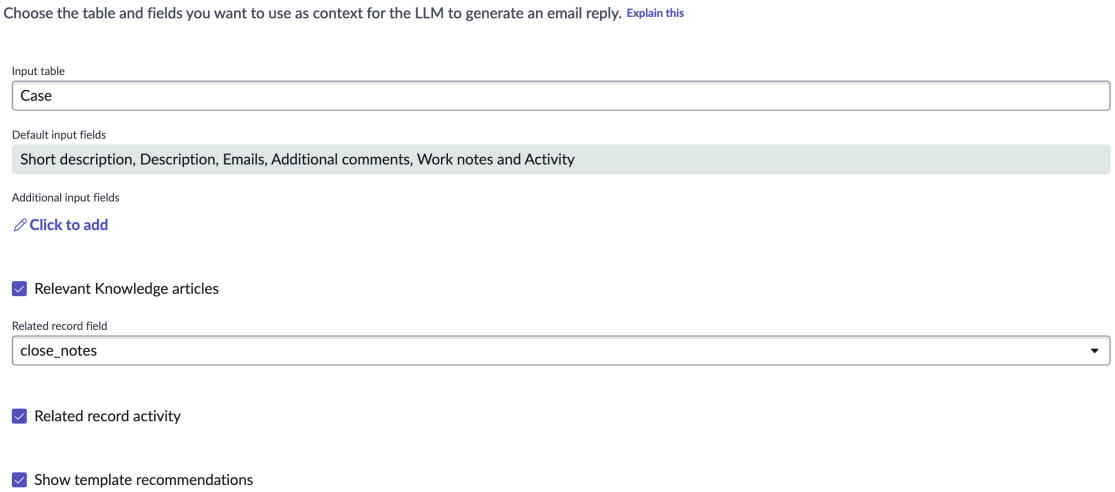
   Some options may be read-only in certain workflows.

   Additional context can improve the quality of the email recommendation. In the **Additional input fields** field, you can choose any fields not already selected in the **Default input fields**. You can include Knowledge articles by selecting the check box. You can choose a field from

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

related records, such as resolution notes, in the **Related record field**. If there's no value for the related record field or you don't select one, you can also opt to include activity from related records.

If you select **Show template recommendations**, then agents see a list of template suggestions generated by Now Assist when they go to select a template.



8. In the next step of the skill configuration, select where you'd like to display the skill.

   You can select both in-product, Now Assist panel, or both.

   - **In-product**: When selected, Now Assist skills are displayed on forms and workspaces.

     For the skills that appear in-product, select the down arrow to identify the roles that can use the skill.

   - **Now Assist panel**: When selected, Now Assist skills are available in the Now Assist panel. If you don't see this option, you must activate the Now Assist panel. For more information, see Activate the Now Assist panel standard chat.

     For the skills that appear in the Now Assist panel, select the down arrow to identify the roles that can use the skill.

   Roles can be added by entering the name of the role in the User roles field. Existing roles can be removed by selecting the X icon in the role bubble. You must have at least one role specified, but you can add as many as you like.

9. Review your choices and select **Activate** to complete the configuration.

Result

Agents can generate email drafts with generative AI.

Edit a Now Assist skill

Edit the configuration of a Now Assist skill to choose the inputs or triggers and the display location of the skill output.

Before you begin

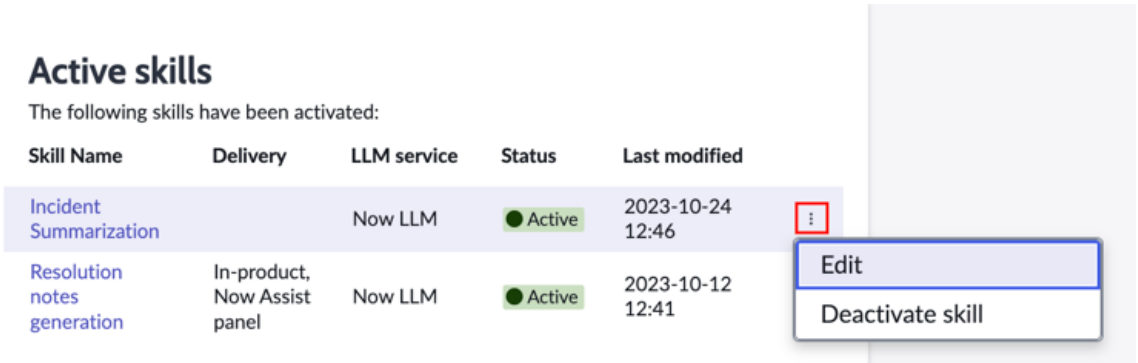
Role required: sn_generative_ai.nsa_admin

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

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Features**.
   If you're already in the Now Assist Admin console, select the **Now Assist Admin Features** tab.
2. Select a workflow on the navigation panel, such as **Technology**.

   Each workflow contains feature sets.
3. On the feature card associated with the skill you'd like to edit, select **View details**.
4. Under **Active skills**, select the more options icon ⋮ next to the skill that you want to configure, then select **Edit**.



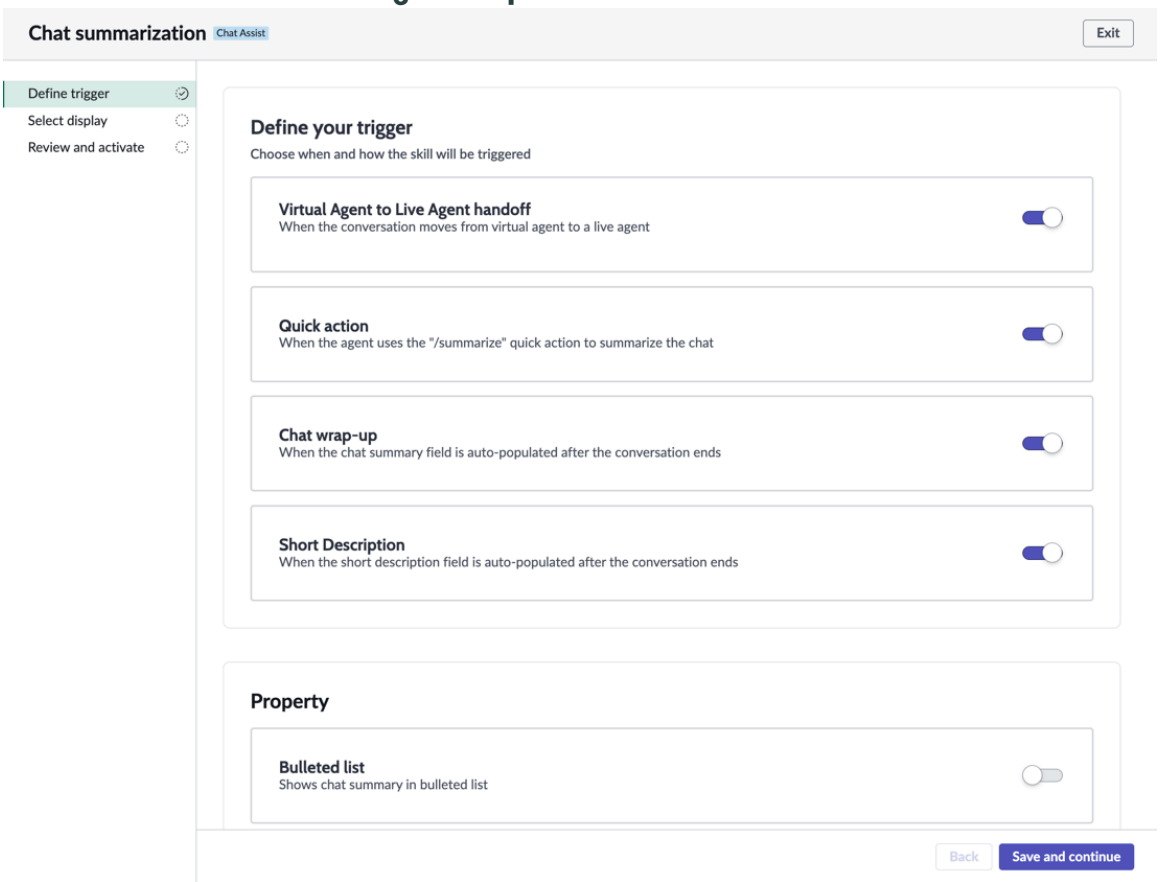
The first step in the guided setup for the skill is displayed.

Each skill configuration has a number of steps shown in the guided setup. The exact steps vary from skill to skill. A symbol next to each step indicates whether the step is completed, partially completed, or not completed.

Note: Some configuration options are read-only.

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

Chat summarization skill configuration panel



5. Proceed to the next step when you've finished configuring the current step by selecting **Save and continue**.

   You can return to a previous step by selecting **Back**.

6. Apply the new settings after reviewing your changes by selecting **Done**.

Result

The skill is activated with your preferred settings. You can now install other plugins or activate other skills.

Make a copy of a Now Assist skill

The 'Make a copy' feature enables you to create a copy of a Now Assist skill so that you can experiment with skill settings and configure the skill to fit your business needs.

Before you begin

Role required: sn_generative_ai.nsa_admin

About this task

The skills that come with the Now Assist applications have default configurations that are optimized to serve the most common use cases. If you want to change the skill settings, you can edit a skill with the Now Assist Admin console or you can create a copy of the skill. Creating a copy leaves the original skill configuration intact in case you want to use it later or want to create another copy from the original. You can activate and configure the copies of the skills by using the same guided setup as the default skills.

Note: The 'Make a copy' feature is not available for all Now Assist skills.

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

Note: In a default scenario, only one version of a skill can be active at a time. If you create and activate a copy of the skill, any previously activated version of the skill is deactivated.

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Features**.

   If you're already in the Now Assist Admin console, select the **Now Assist Features** tab.

2. In the navigation pane, select the workflow of the skill that you want to copy, such as Technology or Customer.
3. On the feature card that contains the default skill, select **View details**.
4. In the All available skills or Active skills section, select the more options icon next to the skill that you want to make a copy of and select **Make a copy**.

   Note: Only one version of a skill can be active at a time. If you create and activate a copy of the skill, any previously activated version of the skill is deactivated.

5. In the modal, select **Make a copy**.

Result
A copy of the skill is generated and you're taken to the guided setup.

What to do next
Continue the steps in the guided setup to activate the skill. For more information, see Activate a Now Assist skill.

If you're making a copy of the case or incident summarization skill and would like to learn more about your options, see the documentation for configuring record summarization.

Configure case or incident summarization in the Now Assist Admin console

Configure case or incident summarization by using the guided setup in the Now Assist Admin console. You can choose the input tables and fields as well as customize the prompt output for copies of the record summarization skills.

https://player.vimeo.com/video/996395898?h=e609c55303&badge=0&autopause=0&player_id=0&app_id=58479

Before you begin
You can only customize the input data and prompt output for a copy of a record summarization skill. To learn more about making a skill copy, see Make a copy of a Now Assist skill. After you create a skill copy, you can learn the steps to complete the skill setup here.

Role required: nsa_admin

About this task
By default, many settings for Now Assist record summarization are optimized for general use cases. Review your goals for incorporating generative AI on your instance to determine whether you want to make changes and what those changes are. After you have made a plan, you can create a copy of a skill and modify the input sources and prompt output.

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

Procedure

1. In the **Name of the skill** field, enter the skill name.

   The default name adds (copy) to the end of the original skill name. For example, Case summarization (copy). Changing the name to be more specific makes it easier to identify later if you want to make changes.

2. **Optional:** Add a description for the skill.
3. Go to the next step by selecting **Save and continue**.
4. Choose the input fields and data sources for each input template.

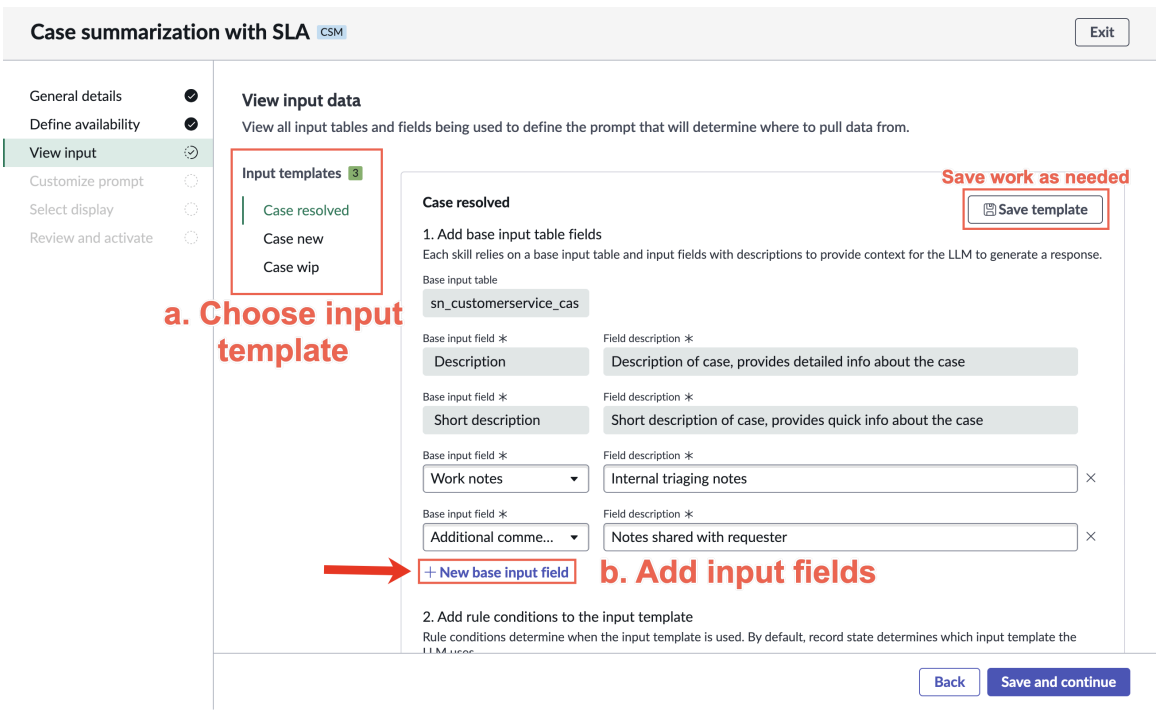
   The default options are selected for you. Some options are read only. After you make any changes to the input template, you can save your work by selecting **Save template**.

   a. Select an input template.

   The three default input templates are for new records, records that are in progress, and closed records.

   b. Add the base input table fields by selecting **New base input field**, choosing a field, and entering a field description.

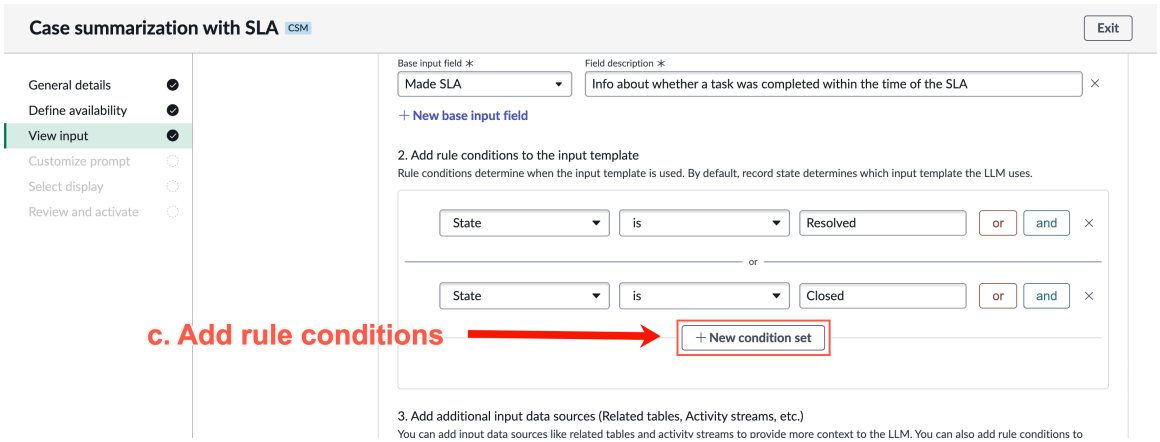
   Each base input table field requires a description. The description informs the large language model (LLM) what the field is for and how the information should be interpreted. The more information that you put in the description means that the model has more context for the data.



   c. Add or modify the rule conditions for the base table.

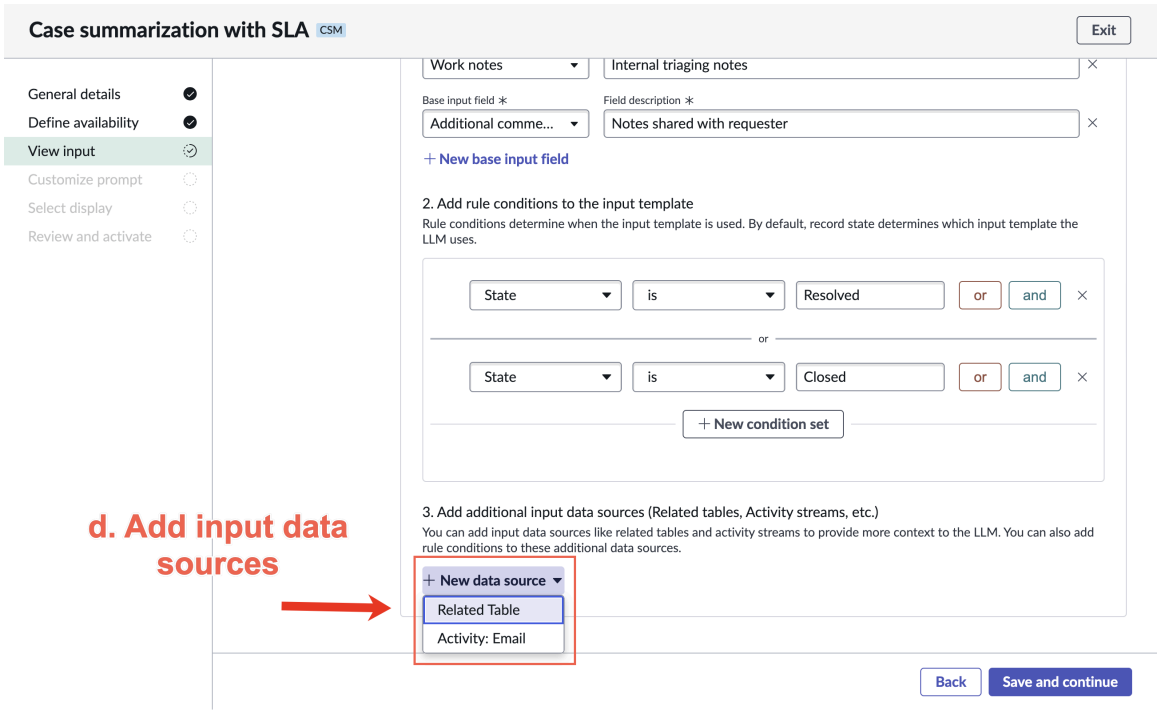
   The rule conditions determine when the input template is used. Record summarization is only available to the records that match the rule conditions of an input template.

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



d. Optional: Add additional input data sources by selecting **New data source** and choosing either **Related Table** or **Activity: Email**.

(Optional) Each related table is configured with input fields and descriptions. More specific descriptions for related tables help provide more context to the LLM. Activity fields, such as Email, don't have input fields that you can configure.

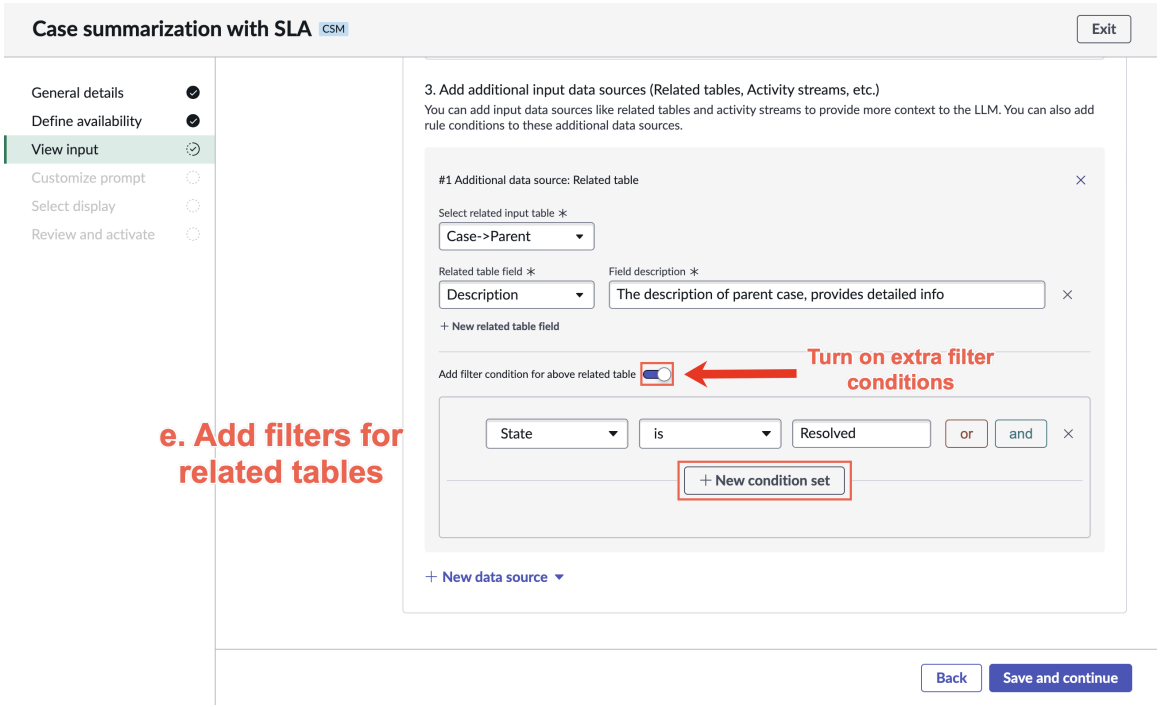


e. Optional: Add a filter condition to the related table.

(Optional) You can add more rule conditions to the related table. These rule conditions determine whether the data from the additional data source is incorporated into the

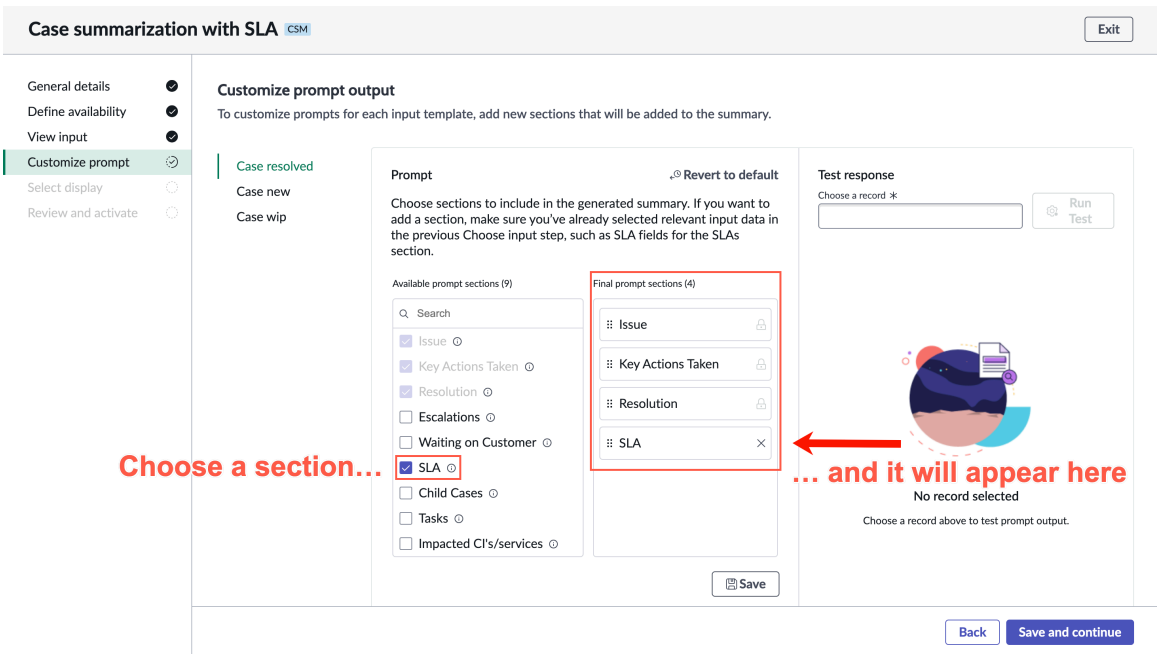
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

summary. You can generate summaries on cases that don't match additional data source rule conditions as long as the base table rule conditions are met.



5. Select **Save and continue**.
6. Choose prompt output sections to appear in summaries by moving a prompt section in the Available prompt sections list to the Final prompt sections list.

   You can reorder sections by dragging the boxes in the Final prompt sections list. Some input templates have sections that are marked with the lock icon ( ). These sections must appear in the final summary, but you can still reorder them with any sections you have added.



7. In the Test response panel, select a record from the **Choose a record** field.
8. Generate a summary for the chosen record by selecting **Run Test**.

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

Important: Each time that you test your prompt output, the operation counts as an assist that is tracked by your Now Assist subscription. To track your Now Assist usage, Monitoring Now Assist usage in Subscription Management.

Running multiple tests with different records can help ensure that you're satisfied with the results.

9. Select **Save and continue**.
10. Choose when the skill is available by selecting either **Skill is always available** or **Customize skill availability**.

    If you choose **Customize skill availability**, you can use the condition builder to add conditions that restrict when the skill is available. For example, you could make the skill only available for certain assignment groups.

11. Select **Save and continue**.
12. Choose where you want record summarization to be available by selecting the toggle next to your preferred display option.

    You can select both in-product, Now Assist panel, or both.

    - **In-product**: When selected, Now Assist skills are displayed on forms and workspaces.

      For the skills that appear in-product, select the down arrow to identify the roles that can use the skill.

    - **Now Assist panel**: When selected, Now Assist skills are available in the Now Assist panel. If you don't see this option, you must activate the Now Assist panel. For more information, see Activate the Now Assist panel standard chat.

      For the skills that appear in the Now Assist panel, select the down arrow to identify the roles that can use the skill.

    You can add roles by entering the name of the role in the **User roles** field. You can remove existing roles by selecting the X icon in the role bubble. You must specify at least one role, but you can add as many roles as you like.

13. Select **Save and continue**.
14. Review your choices and select **Activate** to complete the configuration.

Result

Your customized version of case or incident summarization is active on the instance.

What to do next

Analyze your skill performance on the Now Assist Admin console to help determine the success of the new version of the skill. Learn more about tracking Now Assist usage at Monitoring Now Assist usage in Subscription Management.

Analyzing Now Assist usage

Use the Now Assist analytics and monitoring tools in the Overview page to review the summaries, skill usage information, and issues that need your attention.

Now Assist Admin Overview page

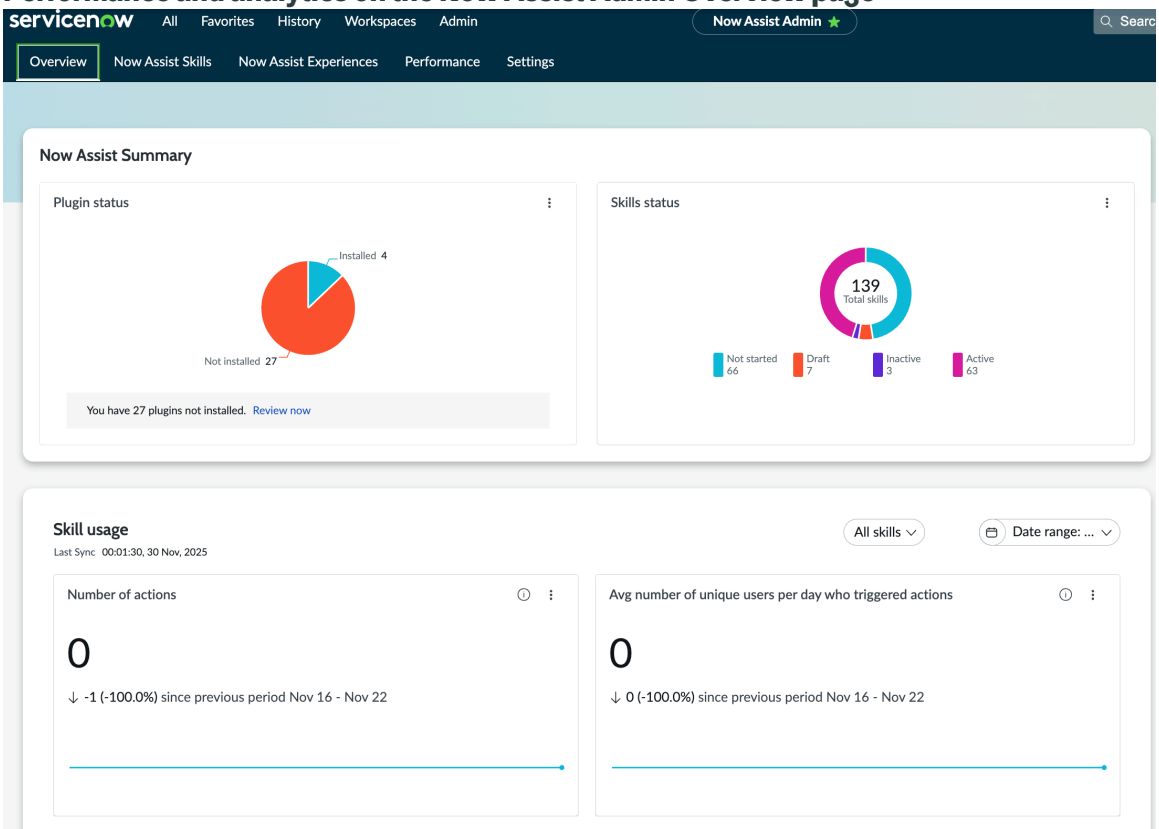
After you activate and begin using the skills, you can measure their usage over time with graphs that display important metrics. The console Overview page contains the information about skill

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

performance, including the task completion over time, and a summary of which skills are active or available.

The following example shows the Now Assist Admin Overview page.

Performance and analytics on the Now Assist Admin Overview page



Now Assist Summary section

The Now Assist Summary section displays information about the status of your skills and plugins in a graphical format.

Plugin status

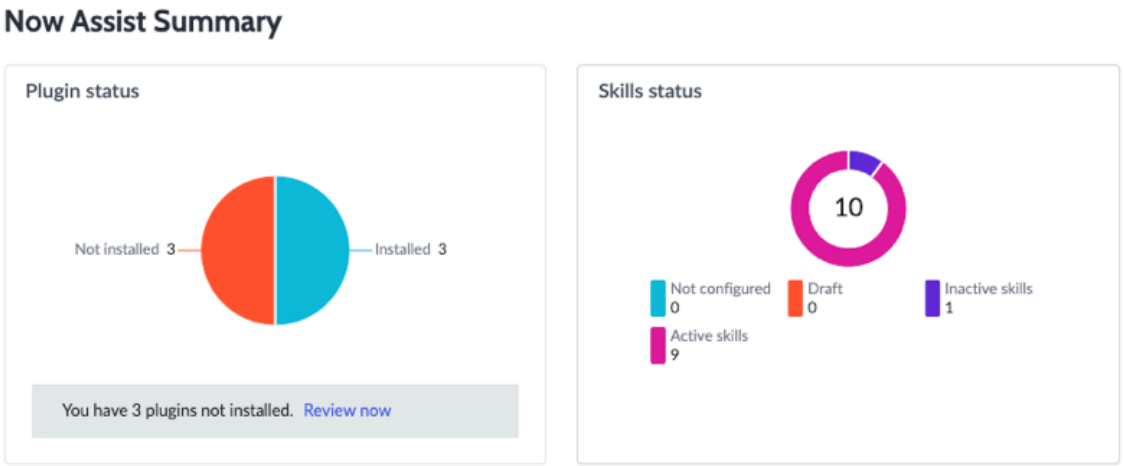
Displays a pie chart that shows the number of plugins that are installed and the number of uninstalled plugins that are available. Select the **Review now** link to review and install additional plugins.

Skills status

Displays a donut chart that shows your skills according to the state of their configuration: active, inactive, draft, or not configured.

The following diagram shows the plugin status that displays as a pie chart, and the Skills status that displays as a donut chart.

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



Skills usage section

The Skills usage section displays the metrics on your installed active skills. Select the info icon to see the information about what each card's metrics represents, or select the option icon to refresh the card.

Use the configuration controls to configure the charts.

All skills

View all skills, or select one or more skills individually to change the view.

Date range

Change the date range for the analytics.

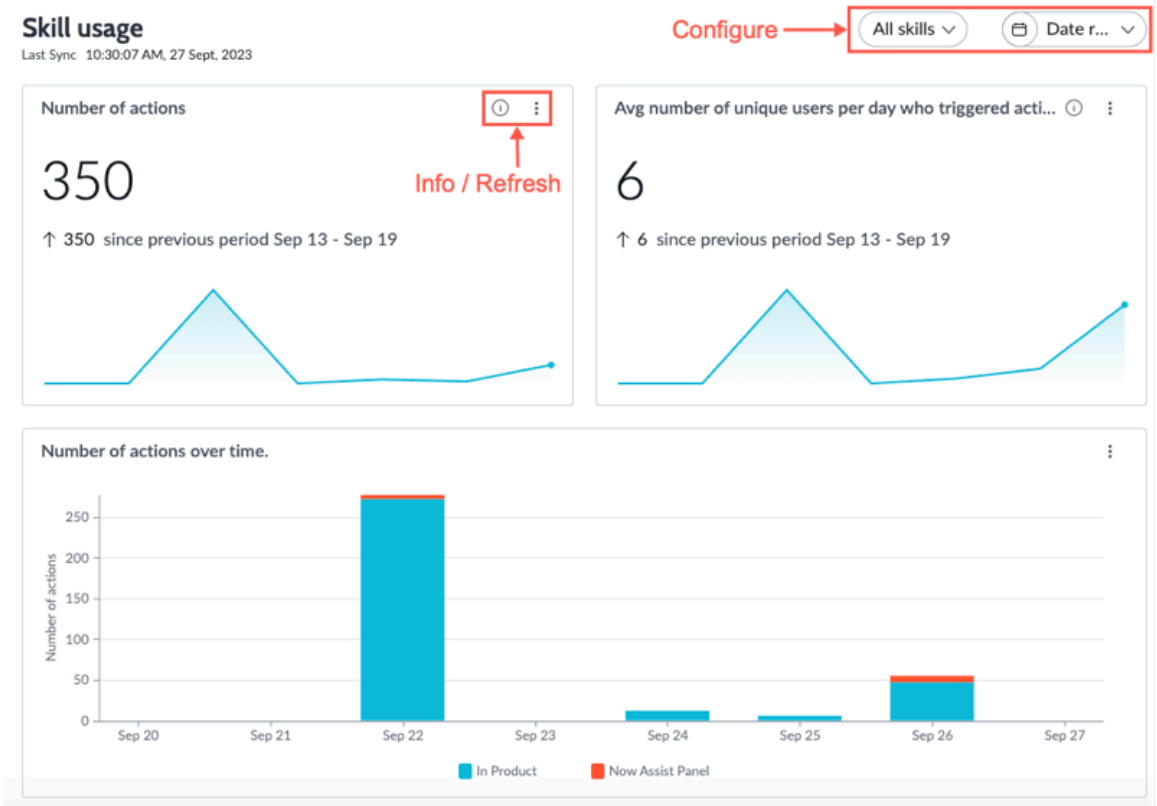
The following diagram shows the skills usage graphs and controls, which include the number of actions, average unique users per day, and the number of actions over time.

Data is collected once a day and uses Performance Analytics and Reporting (PAR) for data collection and visualization.

For more detailed information about performance, see Now Assist Analytics and

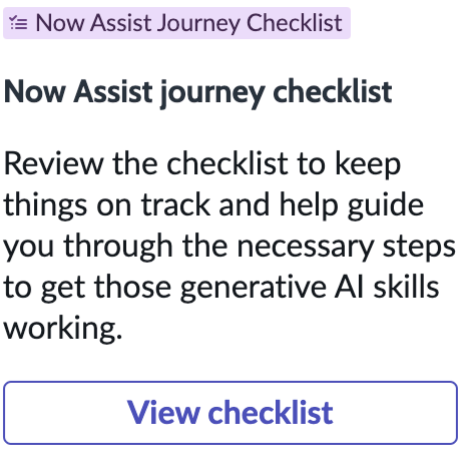
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

Skills usage graphs and controls



Now Assist journey checklist

The Now Assist journey checklist provides you with a guide of the Now Assist workflow. If you are confused as to what your next step should be when implementing Now Assist on your instance, access the checklist at any time by selecting **View checklist**.



Additional resources on the Overview page

The Helpful resources and FAQs sections provide links to documentation and answer common questions.

The Needs Attention section displays a count of items that need your attention. Below the count are categorized links to those items.

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



Now Assist reference

Reference topics include information about user roles, data usage, and domain separation for Now Assist.

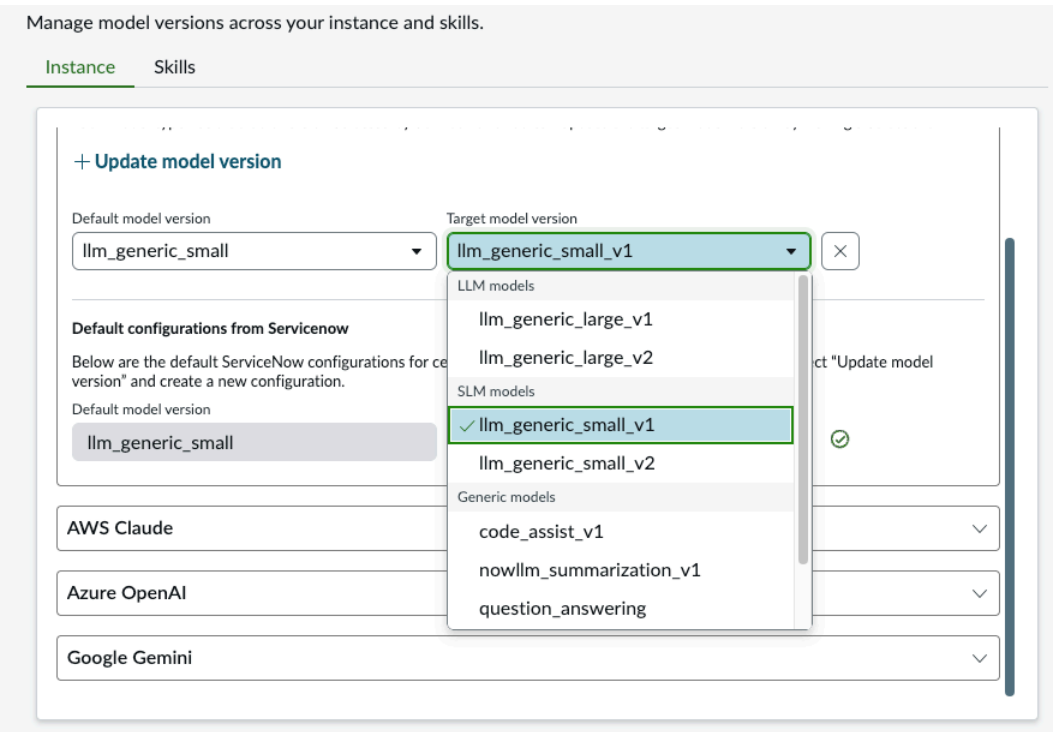
Default and target model version

Model version is the large language model version a skill uses to route requests to process users' queries. Default model version is where all the requests route to by default. This is pre-set by ServiceNow®. A target model version is chosen to route your requests to a different version at run-time, rather than using the default version.

Updating the target model version at the instance level

Note: A model version will not be available for selection as target model version, if it is in deprecated, retired, in review or rejected state.

A default mapping of default and target model version is pre-configured. If you update the target model version for a selected default model version, all the associated skills with this version mapping at the current instance level, get impacted.



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

Updating the target model version at the skill level

If you update the target model version for a selected default model version at the skill level, the mapping is updated for that skill only. Customizing the model version for skills overrides the instance-level model version currently assigned to each provider. This action is typically reserved for specific situations.

Domain separation in the Now Assist Admin console

Domain separation is supported for the Now Assist Admin console. 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: Basic

- Business logic: Ensure that data goes into the proper domain for the application's service provider use cases.
- The application supports domain separation at run time. The domain separation includes separation from the user interface, cache keys, reporting, rollups, and aggregations.
- The owner of the instance must set up the application to function across multiple tenants.

Sample use case: When a service provider (SP) uses chat to respond to a tenant-customer's message, the customer must be able to see the SP's response.

For more information on support levels, see Application support for domain separation.

In the Now Assist Admin console, generative AI capabilities are organized into skills. Each skill can be configured differently for each domain or you can create a variant of a skill for a domain. By default, all skills exist in the global domain.

How domain separation works in the Now Assist Admin console

You must enable domain separation on your instance first before you can use it for Now Assist skills.

Now Assist works with domain separation. When you use Now Assist in a domain-separated environment, users are only able to access data within their domain. For example, if a user uses the summarization skill, Now Assist only uses material that exists within the user's domain when generating that summary. When a skill is domain separated, only users who are in that domain can use the skill that you have configured for that scope.

If you're a service provider that hosts multiple clients in the same instance, you can set up domain separation to separate tenant data, processes, and administrative tasks. However, Assist consumption is tracked according to instance without differentiating between tenants. You can track your Now Assist usage in the Subscription Management dashboard.

If you want a domain to have a different version of an existing skill, you can reconfigure and activate the skill or create a variant in the preferred domain. See the section on granting access to Now Assist skills to a domain.

Use cases

You can configure the inputs, roles, triggers, and prompts when you're activating or editing a skill or a later variant of the skill.

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

Some use cases include the following examples:

- Use the **Activity** field as an input in the incident summarization in one domain but only use the **short description** and **description** fields in another domain.
- Grant certain roles access to the Now Assist panel in one domain while another domain has no role restrictions.
- Trigger the generative AI capabilities by using quick actions in Agent Chat in only one domain.
- Create a variant of a skill to test one prompt in one domain while another domain uses the default prompt for the skill.

Granting a domain access to Now Assist skills

Domain separation is possible at the skill level and at the individual configuration level. When using the guided setup in the Now Assist Admin console, each configuration option has its own record that you can separate by domain. To create a record in a different domain, you must set up the skill while in the scope of your preferred domain.

1. Navigate to the Now Assist Skill Config (sn_nowassist_skill_config) table.
2. Add the Domain field to the list. If it isn't present, select the gear icon at the top of the list and add the Domain field into the Selected column, then select **OK**.
3. Find the skill that you want to enable on a domain-by-domain basis. Set Active to false on the skill that is in the global scope. You might need to change the scope to edit the record.
4. Change your current domain to the domain that you want to enable the skill in.
5. Navigate to **All** > **Now Assist Admin Console** > **Features**.
6. Navigate to the skill that you want to activate according to domain and select **Activate skill**.
7. Configure the skill as usual. For more information, see Activate a Now Assist skill.
8. Return to the Now Assist Skill Config (sn_nowassist_skill_config) table. There should be a new record in the current domain. Open the new record.
9. In a different browser tab, return to the Now Assist Skill Config table and open the deactivated skill record in the global domain.
10. Compare the global skill record to the one created within your domain. Records on the related list may not be present in the domain-specific skill. If they are not there, you must recreate those records in your domain and attach them to the related list in your domain-specific skill.
11. Repeat the process for each skill and each domain where you want to have the skill available.

Related topics

Domain separation for service providers

Fetch end points in Now Assist Conversational Help skills

The Now Assist Conversational Help skills architecture solves latency by fetching answers hosted at the nearest location, which is best suited to the user.

Fetching solutions hosted on multiple geographical location

The possible solutions to users queries posted on Now Assist Conversational Help skills are hosted on three main locations: Japan, EMEA and the US.

For optimal performance, the central instance should be situated in the same geographic area as the user's instance. To achieve this alignment, we utilize DISH service, a tool within the user's

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

instance, that helps identify the correct endpoint. The Now Assist Conversational Help skills uses the Mimir lookup service feature, inside the different data centers.

The DISH service communicates with the Mimir lookup table to determine the end points matching the user's location. Once we obtain this endpoint, it gets securely stored in the sys-service table within the customer instance. Moving forward, all the users' queries are routed through the specific endpoint, ensuring consistency and reliability. The Get Help skill uses the endpoint present in the sys-property `com.snc.get_help.endpoint`.

Note: The Now Assist Conversational Help skill version is stored in sn_ads_now_help.com.snc_now_help_skill.version, ensuring backward compatibility within the conversational shared services.

Now Assist Admin roles

Certain roles are required to use Now Assist Admin functionality.

Now Assist Admin [sn_nowassist_admin.nsa_admin]

This user can create and update the Now Assist Admin experience by editing and configuring skills.

Contains Roles

List of roles contained within the role.

ACE User [ace_user].

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.

Now Assist Admin console user [sn_nowassist_admin.user]

This user can access the console and view skills and their configurations, but cannot make edits.

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.

Now Assist panel user [now_assist_panel_user]

Users who have access to the Now Assist panel.

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

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.

User data usage policy for Now Assist

Now Assist is designed to keep user data safe and secure. You can also mask sensitive data or opt-out of sharing data for model improvements.

How your data is sent and stored

Your AI workloads are securely sent using Transport Layer Security (TLS) 1.2 from your ServiceNow instance to one of three centralized ServiceNow compute hubs (datacenters with GPUs for AI workloads), where the AI prediction processing takes place. The data used to generate the response is deleted from the compute hubs after the response has been generated. The result is then returned to the ServiceNow instance.

The input and output data isn't cached or stored on the compute hub and is transient.

Your data isn't commingled with other customer data when using Now LLM Service for generative AI.

Also, there's no commingling of data for domain-separated instances when you use Generative AI services.

When appropriate, ServiceNow might leverage third-party endpoint services (for example Azure OpenAI Service) to augment Now LLM Service to power Now Assist capabilities. Azure OpenAI is a third-party model, but policies same as the company's, apply, as it's hosted in our compute hubs.

Further, to confirm quality of service, ServiceNow might use Azure-hosted GPUs for Now LLM Service capacity bursting in case of high customer demand. Data processed by third-party endpoints isn't subject to use or access by third-party providers and are operated within the ServiceNow network boundary.

Mask sensitive data

Sensitive data can be masked before sending it to LLMs using Now Assist for Data Privacy. To learn more, see Configuring Data Privacy for Now Assist.

After you enable the plugin, it's designed to mask sensitive data before it's sent to the LLM, but could result in less accurate results because the specific data isn't included within the prompt.

Note: This plugin, within the context of generative AI products, doesn't mask the sensitive data that exists in records within your instance, nor does it help prevent new sensitive data from being stored on the instance itself.

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

ServiceNow might use Retrieval Augmented Generation (RAG) for selected AI features (for example NowAssist for AI Search) and passes information to the LLM based on what the requester can access within the system. If a user searches for something in the portal using Now Assist in AI Search, AI Search finds the article and then sends that to the LLM as a part of the prompt. Because AI Search knows what the user has access to, it won't send an article that the user isn't able to access.

Note: For some features, such as case summarization, the agent generating the summary might have more permissions than other people who have access to the record. If they choose to paste that summary to the work notes, the agent should check to confirm that the data they're sharing in the work notes is appropriate to share with others who have access to that record.

Opt out of data sharing

Data Sharing helps ServiceNow to continuously advance and improve its Now LLMs, based on the latest customer usage. If you no longer want to participate in the customer data-sharing program, you're able to opt out.

To opt out, follow the instructions in Opt out of data sharing for Now Assist.

Troubleshoot a Now Assist skill

Run diagnostics for a skill on the Now Assist Admin console and get information about the status of your skill configuration.

Before you begin

Role required: sn_nowassist_admin.nsa_admin

About this task

Certain skills have diagnostic scripts that you can run from the Now Assist Admin console. These diagnostic scripts check for successful skill execution and setup of the underlying capability definitions. If you've made a copy of a skill, you will not be able to run diagnostics on the skill copy.

Procedure

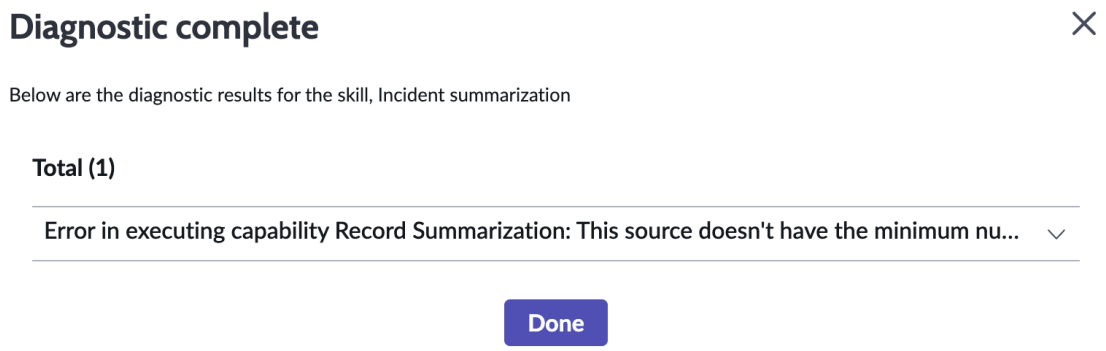
1. Navigate to **All** > **Now Assist Admin** > **Features**.

   If you're already in the Now Assist Admin console, select the **Now Assist Features** tab.

2. In the navigation pane, select the workflow of the skill that you want to troubleshoot, such as **Technology** or **Customer**.
3. On the feature card that contains the skill you want to troubleshoot, select **View details**.
4. In the All available skills or Active skills section, select the more options icon next to the skill that you want to make a copy of and select **Run diagnostics**.
   Run diagnostics is greyed out for skills where this option is not available or applicable.

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

5. After the diagnostics are complete, review the results of each test.



What to do next
If you have identified any problems with your skill configuration, you can edit the skill from the Now Assist Admin console.

If editing the skill does not solve the issue, you can contact ServiceNow Support for additional help.

Now Assist panel system properties

Use system properties to customize Now Assist panel. Some properties are available on a system properties form, but some lesser-used properties are available only from the System Property [sys_properties] table.

Property	Description
sn_nowassist_va_enable_nap_aix_experience	Enabled and integrated chat is available for customers to choose from.
sn_nowassist_va_enable_nap_aix_experience	Disabled, integrated chat is not available for customers to choose from.

Generative AI Controller

Use Generative AI Controller to integrate third-party large language models (LLMs) with your workflows.

Get started

Get started with Generative AI Controller to integrate directly with external LLMs. With Workflow Studio and Virtual Agent Designer, you can create your own use cases for AI-generated text and sentiment analysis, including advanced workflows and custom scripts.

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

Explore

Explore Generative AI concepts and terminology.

Configure

Configure Generative AI Controller provider capabilities.

Use

Use Generative AI Controller provider capabilities.

Reference

Reference for Generative AI Controller.

Troubleshoot and get help

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

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 Generative AI Controller

Learn more about generative AI concepts and how to integrate third-party generative AI into the ServiceNow AI Platform to create content, summarize task records, and analyze user sentiment.

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

Generative AI Controller overview

Complex algorithms and deep learning models learn patterns and use that knowledge to generate new outputs. With Generative AI Controller, you can generate content directly within the ServiceNow AI Platform.

Generative AI Controller integrates with external LLMs, including ones by OpenAI, Azure OpenAI, Google Cloud (AI Studio and Vertex), Aleph Alpha, IBM watsonx, and Amazon Bedrock. These capabilities are available in Workflow Studio flows, Virtual Agent topics, and scripting like background scripts and business rules.

Generative AI Controller benefits

Benefit	Feature
Generate text to summarize complex information.	Workflow Studio Actions to Generate Content and Summarize
Analyze user sentiment to identify and alleviate end-user concerns.	Sentiment Analysis Workflow Studio Action
Query a large language model (LLM) directly.	Generic Prompt Workflow Studio Action
Write scripts for AI model capabilities that are designed to increase the accuracy and scalability of your custom content.	Script with generative AI
Integrate with third-party AI service providers to customize your AI experience	OpenAI, Azure OpenAI, Google AI, Aleph Alpha, IBM watsonx, Amazon Bedrock

Get started with Generative AI Controller

- The Generative AI Controller application is installed with any Now Assist application.
- Sign up and create an account with a generative AI provider.
  - To sign up with OpenAI, go to their official platform website.
  - To get started with Azure OpenAI, go to their documentation.
  - To start using AI Studio with Gemini API, go to the AI Studio homepage.
  - To use Vertex AI on the Google Cloud, go to the Vertex AI homepage.
  - To get started with Aleph Alpha, go to their website and create an Aleph Alpha API account.
  - To get started with IBM watsonx, go to Getting started with IBM watsonx as a Service.
  - To get started with Amazon Bedrock, set up an IAM user with the correct permissions and then explore the Converse API.
- Configure credentials for your preferred AI service provider for the Generative AI Controller capabilities.
- Add Generative AI Controller capabilities to Virtual Agent Designer to start creating Virtual Agent topics with generative AI with your AI service provider.
- Use Generative AI Controller with Virtual Agent Designer, Flow Designer, and scripts.

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

Domain separation and Generative AI Controller

Domain separation is supported for Generative AI Controller. 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.

Domain separation enables you to create partitions in the application data and administrative processes. Because the generative AI tables are domain separated, Generative AI Controller supports domain separation for OneExtend capabilities. The capabilities are the basic building blocks for Virtual Agent Designer topics, components, flows, and scripts that use generative AI. With domain separation, you can isolate the data and control access so that users in one domain don't have access to the capabilities of another domain.

For more details on domain separation and Virtual Agent, check out the documentation.

How domain separation works in Generative AI Controller

Domain separation is possible at the generative AI OneExtend capability level. Records that are related to the execution and configuration of OneExtend capabilities, such as log tables that are accessible to ServiceNow personnel, are also separated according to the capability's domain.

If you want to create a copy of an existing generative AI capability in a different domain, you must create a record in the OneExtend Capabilities (sys_one_extend_capability) table. See the reference for Generative AI Controller for more information about the OneExtend Capabilities table.

You set the domain when the record is created. The domain is based on the domain that you're in at the time that you create the record. When you're creating a capability record, you can use an existing OneExtend Capability record as a blueprint to help confirm that the capability works as intended.

After you create the OneExtend Capability record, you must create records for the following attribute and config records in the new domain:

- OneExtend Capability Attribute records with the same values as the capability in the global domain.
- A OneExtend Capability Definition that corresponds to the new capability.
- A OneExtend Definition Config definition record that includes the OneExtend Capability Definition for the new domain.

For more information on setting up OneExtend Definition Configs to use in Virtual Agent Designer, see enabling Generative AI Controller for Virtual Agent.

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

You can also create these records by using the related lists in the OneExtend Capability record default view.

Note: The OneExtend Capability Definition record that you add must be the same as the capability that you want in the new domain. For example, if you're creating a capability in a new domain for sentiment analysis, you could add the Sentiment Analysis (OpenAI Completion) record. Adding the Summarize (OpenAI Completion) Config could result in unexpected behavior. The OneExtend Definition Config record that you select should include the OneExtend Capability Definition record that you added.

Use cases

With domain-separated capabilities, you can build different Virtual Agent topics, flows with Workflow Studio, or different background and business rule scripts that are also domain separated.

Related topics

Domain separation for service providers

Configuring Generative AI Controller

Configure Generative AI Controller providers and capabilities.

Configuration overview

Generative AI Controller enables you to choose between several AI providers or a generic connector.

- Microsoft Azure OpenAI
- OpenAI
- Google AI Studio
- Google Vertex AI
- Aleph Alpha
- IBM watsonx
- Amazon Bedrock

Configure API credentials

Configure your large language model (LLM) credentials to use third-party integrations with OpenAI, Azure OpenAI, Google Cloud, Aleph Alpha, IBM watsonx, Amazon Bedrock, and generic models to control the third-party integration.

Configure a generic LLM connector

Set up an external LLM to use generative AI capabilities to add finer control over prompts and transformer scripts.

Configure an AI service provider for a capability

Set your preferred AI service provider for each capability and tailor generative AI content to suit your business needs.

Enable generative AI capabilities in Virtual Agent Designer

Enable generative AI capabilities in Virtual Agent Designer to build custom topics and topic blocks with content created by AI.

Enable recursive summarization for large inputs

Enable recursive summarization to retain context for large inputs in LLM calls.

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

Enable Dynamic Translation for generative AI capabilities

Use Dynamic Translation for in-product generative AI capabilities to support users who speak different languages.

Installing Generative AI Controller

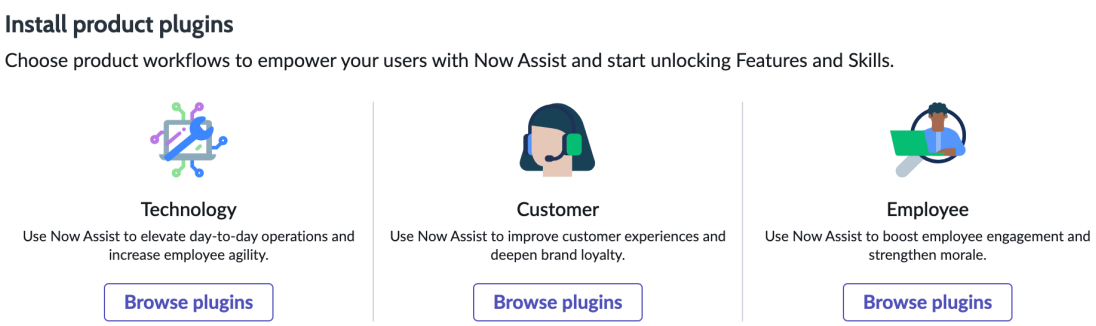
You can install the Generative AI Controller application (sn.generative.ai) with Now Assist applications if you have the admin role.

Installation requirements

You must be on Vancouver patch 2 or later.

Generative AI Controller is included as a dependency for all Now Assist applications. It is not recommended to install the application by itself. Instead, you can install Now Assist applications from the Now Assist Admin console or directly from the ServiceNow Store. For details, see Install Now Assist plugins.

Now Assist Admin console plugin installation



Configuring API credentials for generative AI capabilities

Configure the credentials and connections for your preferred generative AI service provider to integrate third-party generative AI on the ServiceNow AI Platform.

You can use several different models and generative AI service providers to integrate into your custom flows, scripts, and topics.

- Aleph Alpha
- Amazon Bedrock
- Azure OpenAI
- Google - AI Studio
- Google - Vertex AI
- IBM watsonx
- OpenAI
- Generic third-party LLM

Configure API Credentials for Amazon Bedrock

Configure your API credentials to use AWS Bedrock in custom workflows and Virtual Agent Designer topics.

Before you begin

You must have an AWS account and the IAM user needs permission to access the bedrock:InvokeModel action.

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

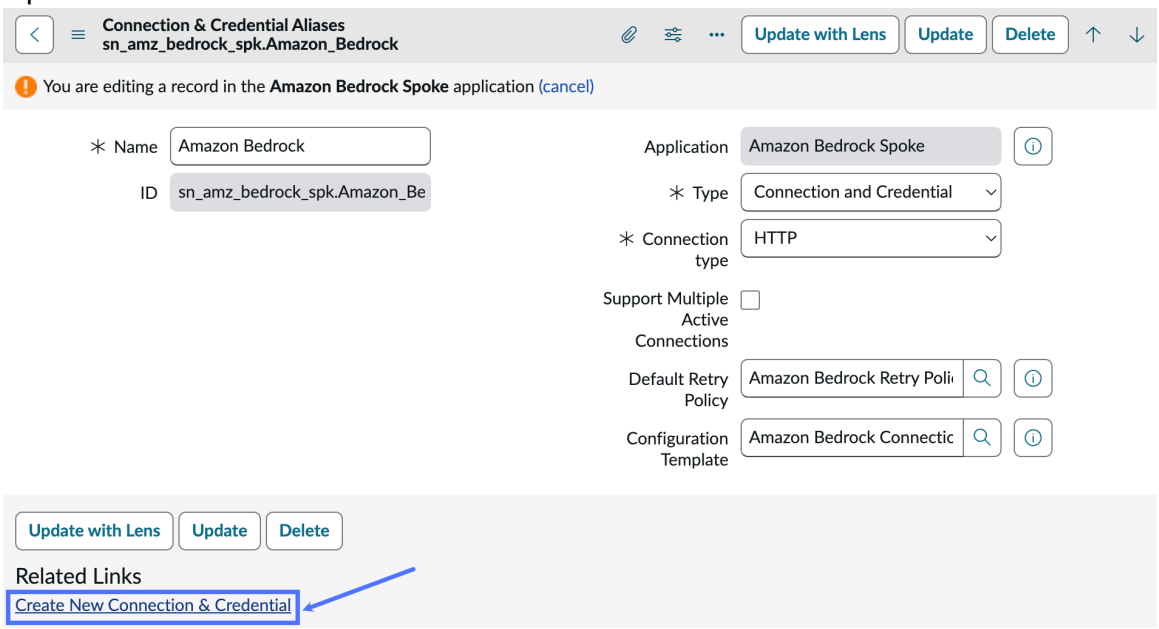
Role required: admin

About this task

To use Amazon Bedrock as your LLM provider for Generative AI Controller capabilities, you must have an active connection configured.

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connections & Credential Aliases**.
2. Open the record for Amazon Bedrock.

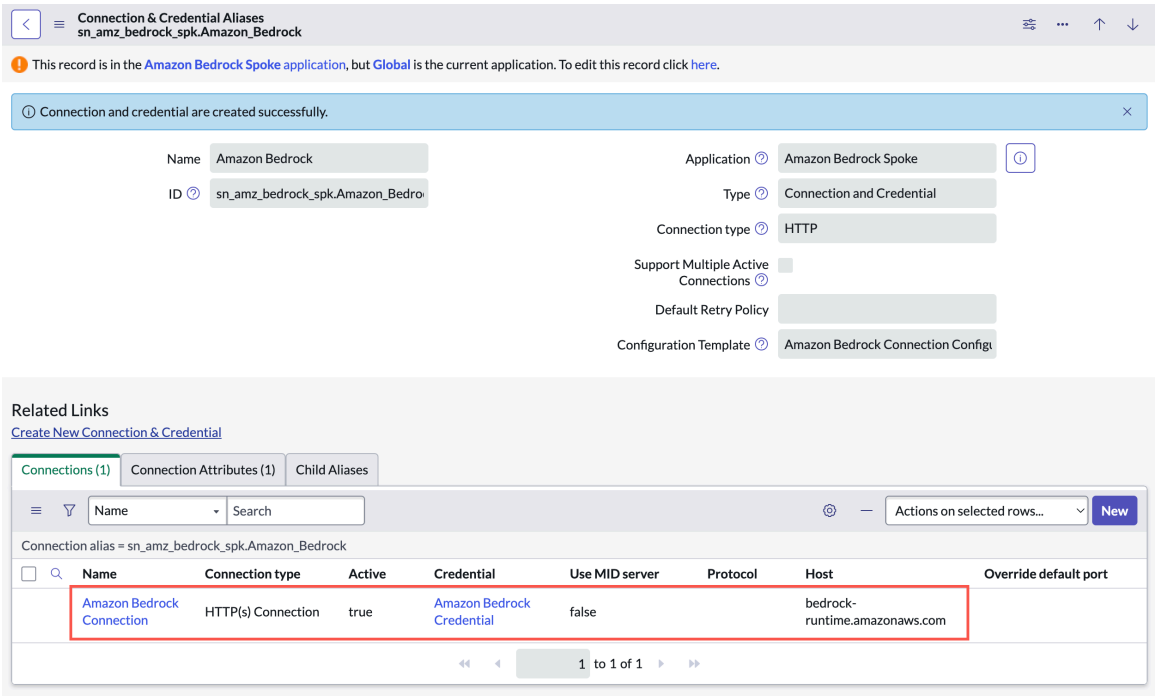


3. Select the **Create New Connection & Credential** related link.
4. Enter your **Region**, such as `us-east-1`.
5. Enter your API key.
   For more information, see AWS documentation on generating API keys.
6. Select **Create**.

Result

You can use Amazon Bedrock as your provider for Generative AI Controller capabilities in Flow Designer, Virtual Agent Designer, and scripts to create custom experiences with generative AI.

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



Configure API credentials for Azure OpenAI

Configure your API credentials to use Azure OpenAI in custom workflows and Virtual Agent Designer topics.

Before you begin

In order to use generative AI capabilities with Azure OpenAI, you must have an Azure resource with an API key.

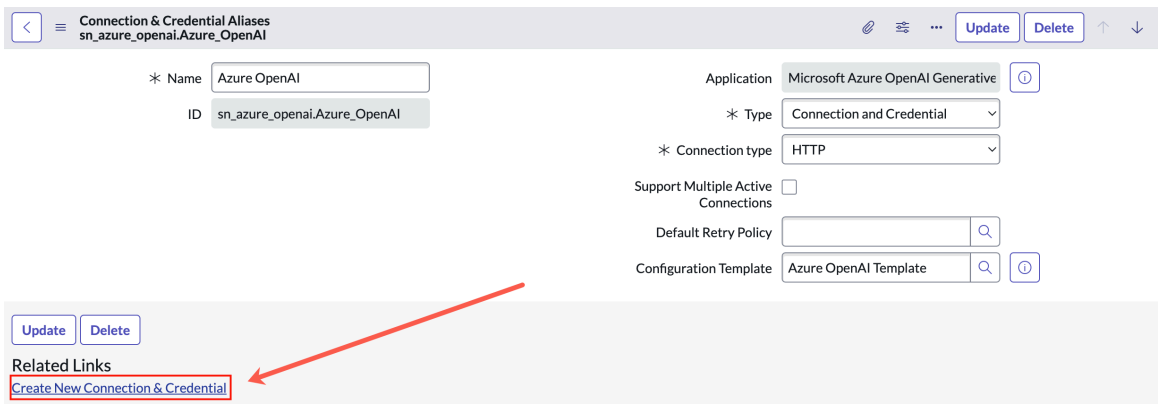
Role required: admin

About this task

In order to use models with Azure OpenAI as your LLM provider for Generative AI Controller capabilities, you must have an active connection configured.

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connections & Credential Aliases**.
2. Open the Generative AI provider record for Azure OpenAI.
3. Select the **Create New Connection & Credential** related link.



4. Edit the Connection URL to include your resource name.

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

For Azure OpenAI, your Connection URL is in the form `https://{your-resource-name}.openai.azure.com`. See the Azure OpenAI documentation for more information.

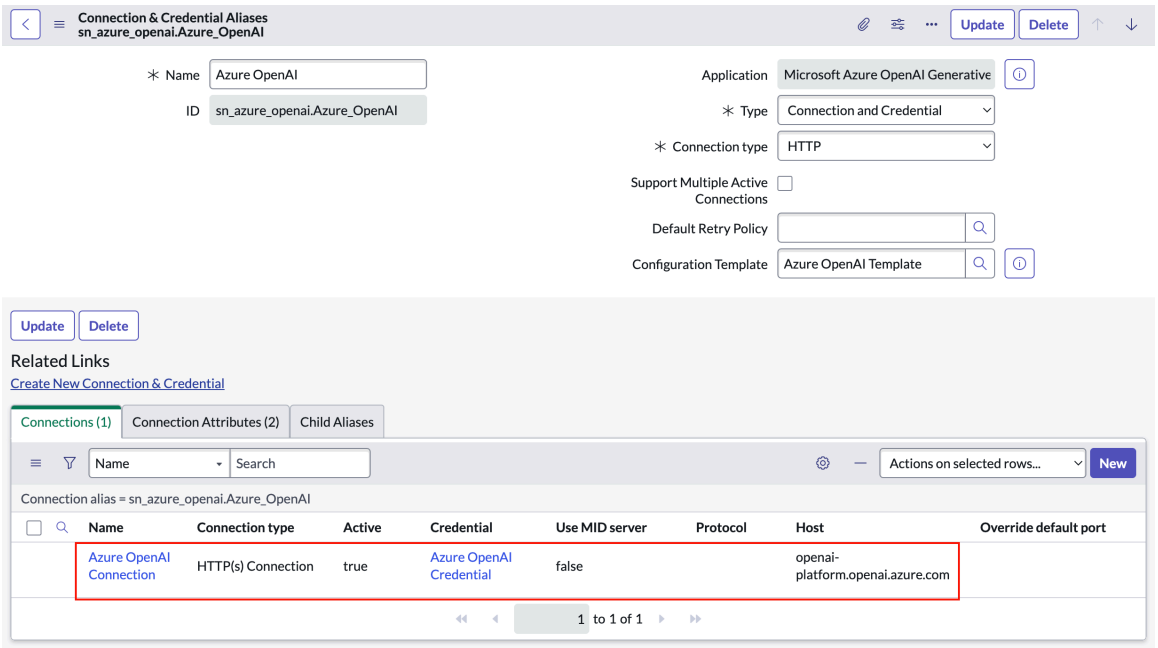
5. In the API key field, enter the API key for the provider.

> **Note:** The characters in the API key field are masked in the user interface.

6. Create a connection by selecting **Create**.

Result

You can now use capabilities labeled with Azure OpenAI in Flow Designer, Virtual Agent Designer, and scripts like background scripts and business rules to create custom experiences with generative AI.



What to do next

If you want to use generative AI capabilities through your MID Server, open the new Connection record, select the **Use MID server** check box, and save the record.

Related topics

- Bring your own key for third-party AI provider integration
- Configure a custom resource path for BYOK models

Configure API credentials for Google AI Studio

Configure your API credentials to use Google AI Studio in custom workflows and Virtual Agent Designer topics.

Before you begin

You must have a Google account in order to use AI Studio.

Role required: admin

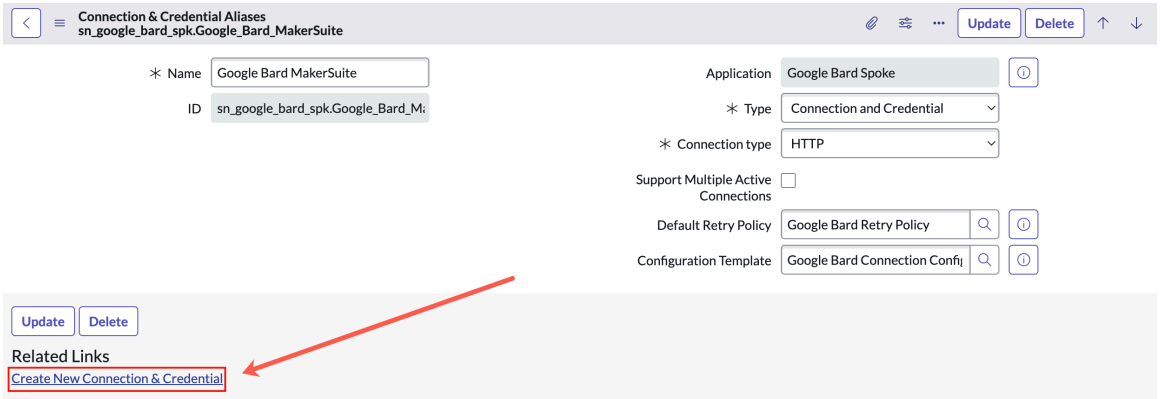
About this task

In order to use Google AI Studio as your LLM provider for Generative AI Controller capabilities, you must have an active connection configured.

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

Procedure

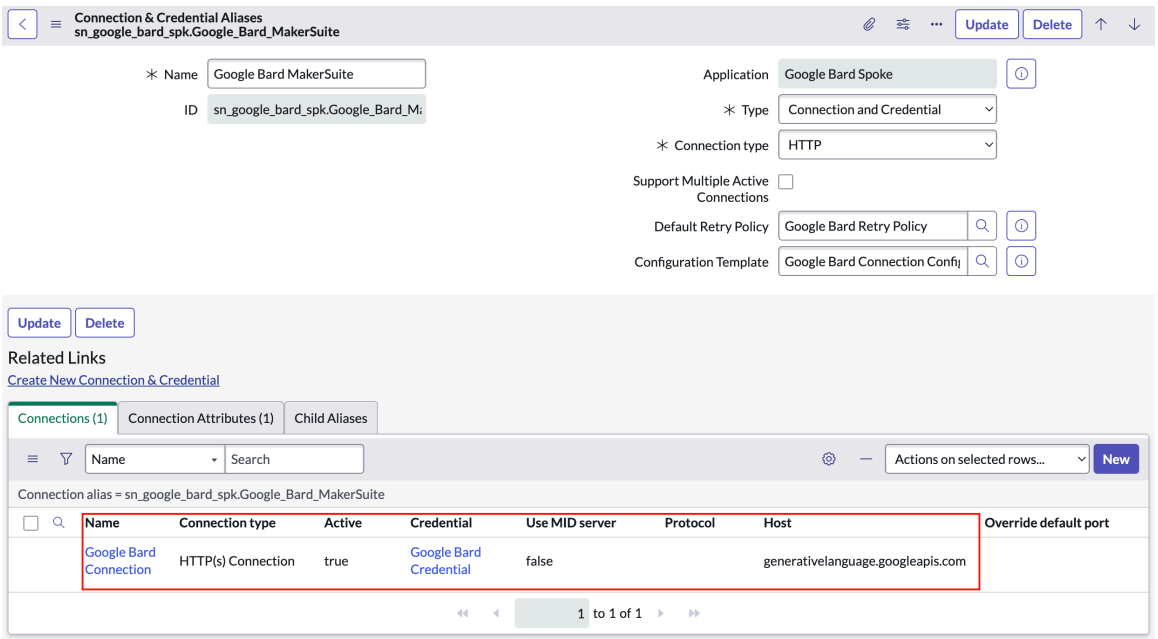
1. Navigate to **All** > **Connections & Credentials** > **Connections & Credentials Aliases**.
2. Open the record for Google Bard MakerSuite.
3. Select the **Create New Connection & Credential** related link.



4. Enter your API key.
   You can find your API key in AI Studio by selecting **Get API Key** from the navigation menu.
5. Select **Create**.

Result

You can now use AI Studio (Google Cloud Chat Completion) and AI Studio (Google Cloud Chat Completion) in Flow Designer, Virtual Agent Designer, and scripts to create custom experiences with generative AI.



What to do next

Use your LLM provider to create flows with Flow Designer, topics with Virtual Agent Designer, or scripts to provide the benefits of generative AI to your users.

Create API credentials for Google Vertex AI

Configure your API credentials to use Google Vertex AI in custom workflows and Virtual Agent Designer topics.

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

Before you begin

You must have a Google Cloud project and the permissions to generate new OAuth credentials.

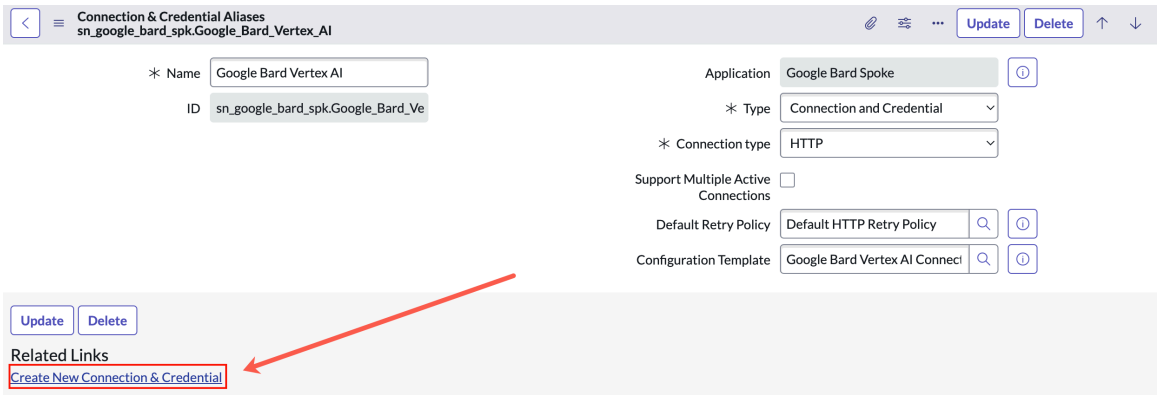
Role required: admin

About this task

In order to use Google Vertex AI as your LLM provider for Generative AI Controller capabilities, you must have an active connection configured.

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connections & Credentials Aliases**.
2. Open the record for Google Bard Vertex AI.
3. Select the **Create New Connection & Credential** related link.



4. Fill in the required fields.

Google OAuth Connection

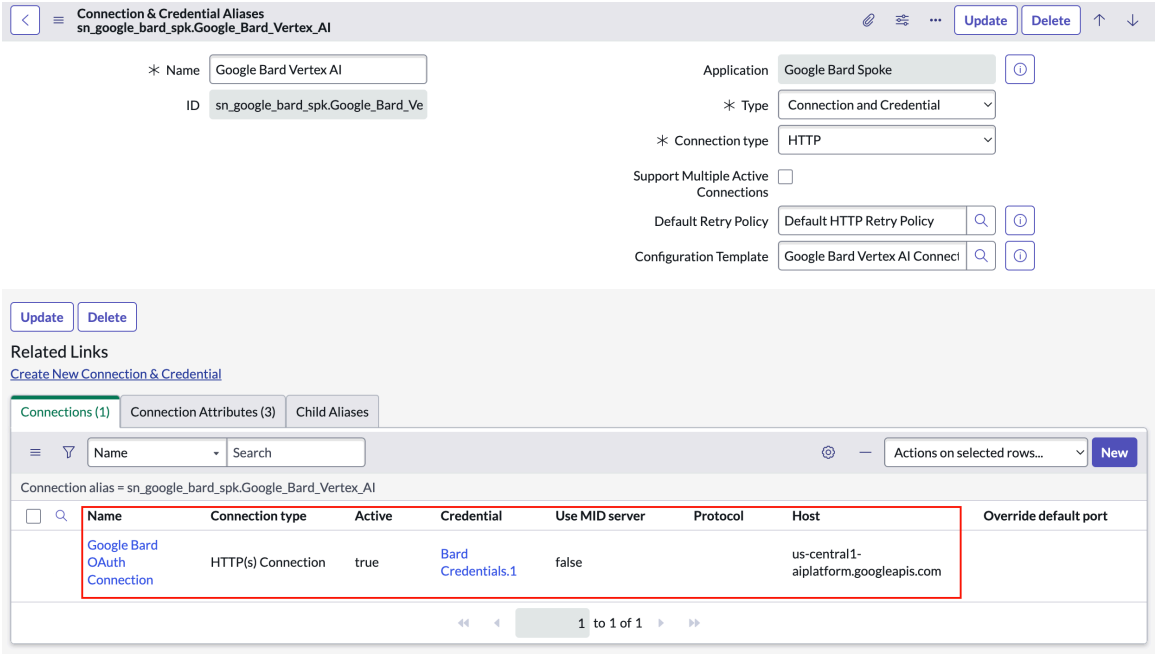
Field	Value
Project ID	The Project ID found in the Google Cloud console
Credential Name	The name of your credential, such as `Google OAuth Credential`
OAuth Name	The name of your OAuth authentication, such as `Google Registry`
OAuth Client ID	To get the OAuth Client ID, create a new OAuth Client ID with the Google Cloud console with the following attributes: **a.** Application type: `Web application` **b.** Authorized redirect URI: URL in the OAuth Redirect URL field, usually `<instance>.service-now.com/oauth_redirect.do` For more information, see the Google documentation for creating OAuth client IDs. Once you have created the OAuth client, a pop-up window will have the Client ID and Client secret for you to copy into your clipboard.
OAuth Client Secret	Client secret from your OAuth Client ID found in the Google Cloud console

5. In the pop-up window, log in to a Google Account with access to the project.
6. When prompted for Google Cloud access for gsuite spokes, select **Allow**.

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

Result

You can now use Completions – Vertex AI and Chat Completions – Vertex AI in Flow Designer, Virtual Agent Designer, and scripts to create custom experiences with generative AI.



What to do next

Use your LLM provider to create flows with Flow Designer, topics with Virtual Agent Designer, or scripts to provide the benefits of generative AI to your users.

Configure API credentials for IBM watsonx

Configure your API credentials to use IBM watsonx Granite models in custom workflows and Virtual Agent Designer topics.

Before you begin

To use IBM watsonx large language models (LLMs), you must configure your credentials to use the Generative AI Controller capabilities.

Role required: admin

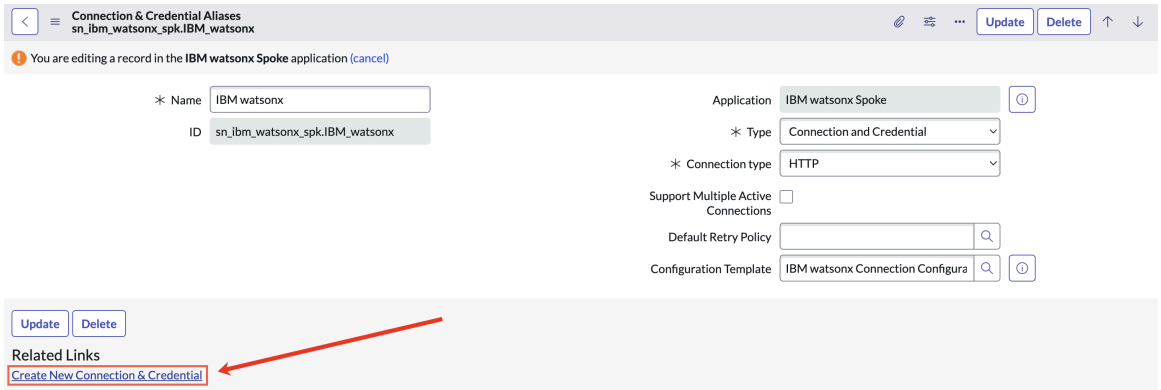
About this task

In order to use models with IBM watsonx as your LLM provider for Generative AI Controller capabilities, you must have an active connection configured.

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connections & Credential Aliases**.
2. Open the record named IBM watsonx.
3. Select the **Create New Connection & Credential** related link.

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



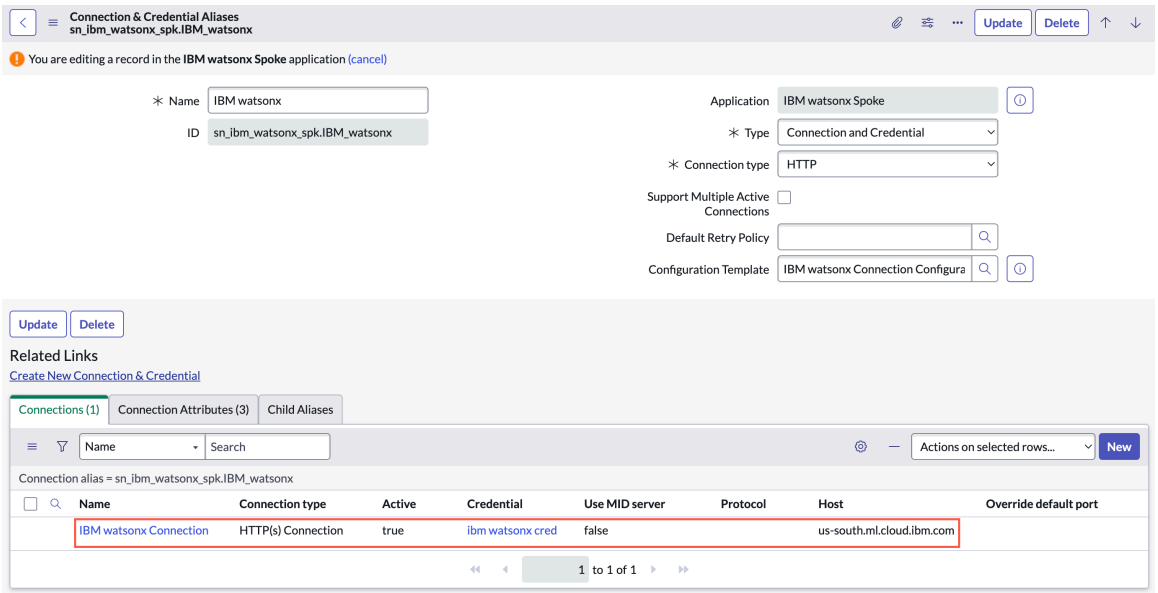
4. In the API key field, enter your API key.

   For more information about generating an API key, see IBM's documentation for generating API keys for authentication.

5. Create a connection by selecting **Create**.

Result

You can now use capabilities labeled with IBM watson in Flow Designer, Virtual Agent Designer, and scripts like background scripts and business rules to create custom experiences with generative AI.



What to do next

If you want to use generative AI capabilities through your MID Server, open the new Connection record, select the **Use MID server** check box, and save the record.

Configure API credentials for OpenAI

Configure your API credentials to use OpenAI in custom workflows and Virtual Agent Designer topics.

Before you begin

To use the OpenAI and Azure OpenAI large language models (LLMs), you must configure your credentials to use the Generative AI Controller capabilities.

Role required: admin

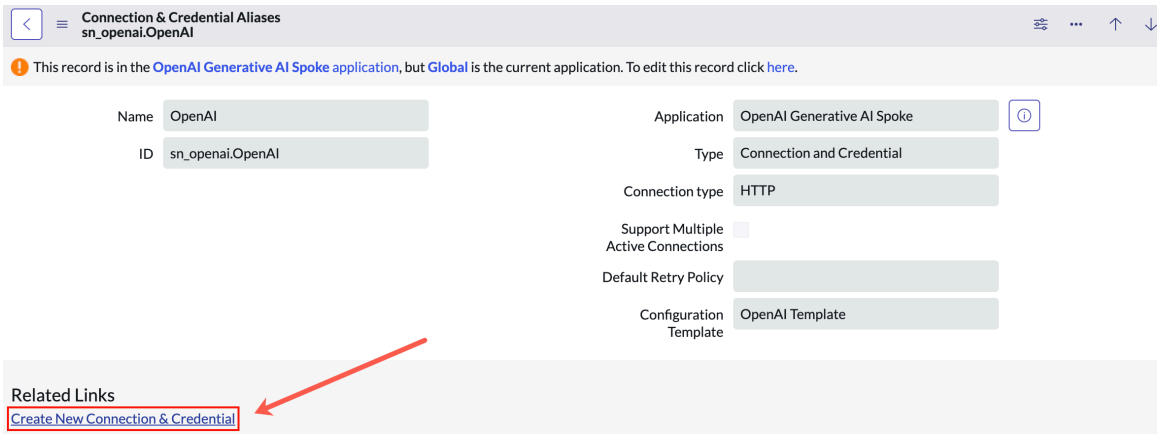
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

About this task

In order to use models with Azure OpenAI as your LLM provider for Generative AI Controller capabilities, you must have an active connection configured.

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connections & Credential Aliases**.
2. Open the Generative AI provider record for OpenAI.
3. Select the **Create New Connection & Credential** related link.



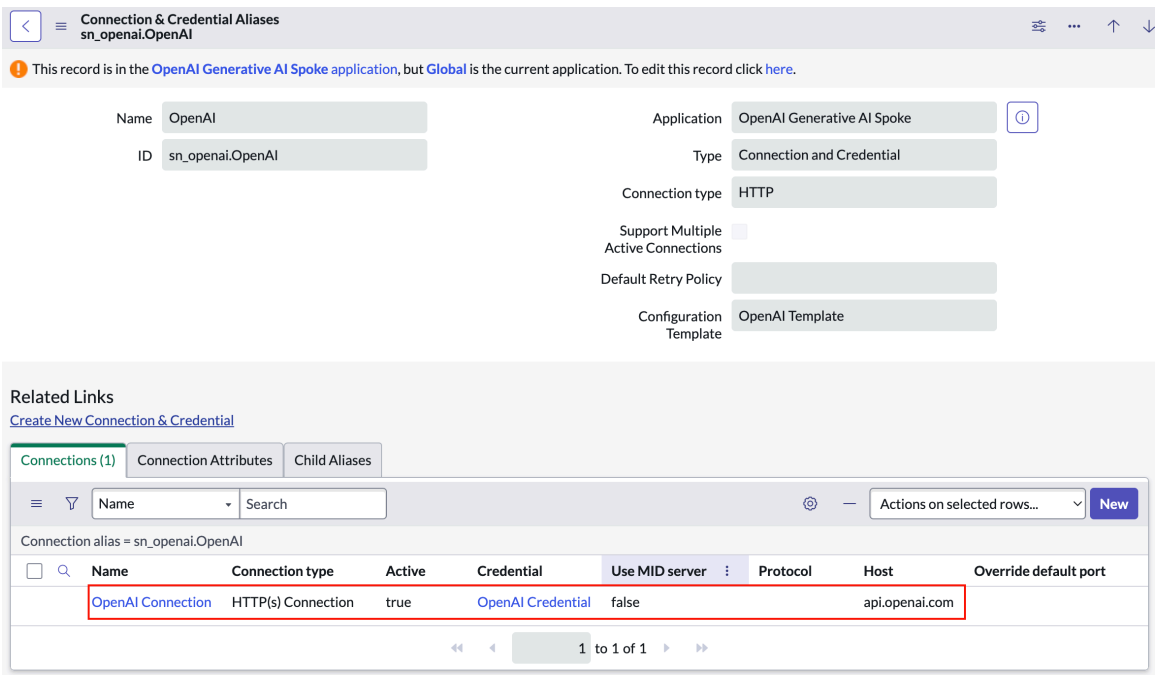
4. In the API key field, enter the API key.

   For OpenAI, see your OpenAI API keys for the credentials information.

5. Create a connection by selecting **Create**.

Result

You can now use capabilities labeled with OpenAI in Flow Designer, Virtual Agent Designer, and scripts like background scripts and business rules to create custom experiences with generative AI.



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

What to do next

If you want to use generative AI capabilities through your MID Server, open the new Connection record, select the **Use MID server** check box, and save the record.

Configure API credentials for a generic large language model (LLM) connector

Use a generic LLM connector to connect the ServiceNow AI Platform with an external AI provider to use generative AI capabilities in custom Virtual Agent topics, Flows, or scripts, like background and business rule scripts.

Before you begin

You must have access to the API of an external LLM in order to configure the credentials.

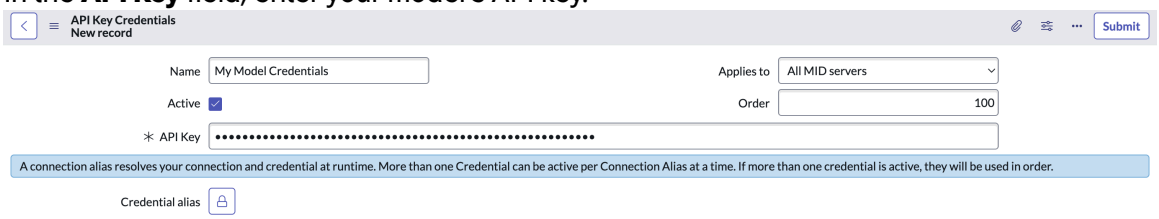
Role required: admin

About this task

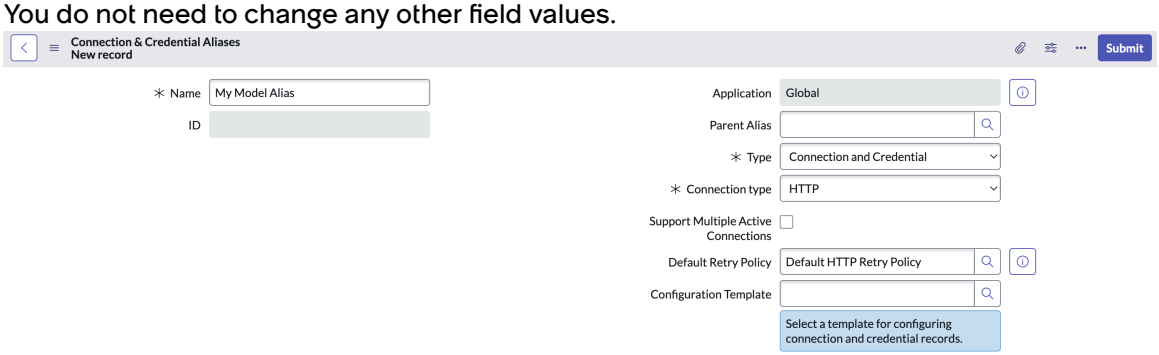
You can connect an external LLM to the ServiceNow AI Platform by creating a connection alias to the model, create a model record that points to the alias, create a prompt record for the model, and then create a transformer record to transform the request or response from the LLM.

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Credentials**.
2. Select **API Key Credentials**
3. In the **Name** field, give your API credential a name.
   Including the name of the model in the API Key Credential makes it easier to identify later.
4. In the **API Key** field, enter your model's API key.



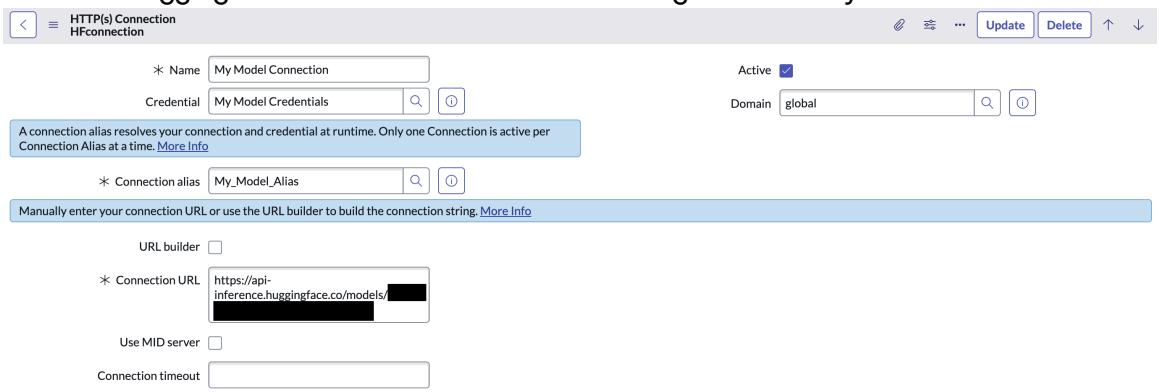
5. Select **Submit** to create the Credential.
   At this step, you have the API key information for the external LLM configured.
6. Navigate to **All** > **Connections & Credentials** > **Connections & Credential Aliases**.
7. Select **New**.
8. In the **Name** field, enter a name for the credential alias.
   You do not need to change any other field values.



9. Select **Submit.**
10. After you are redirected to the List view, find and open the record you just created.
11. Select **New** in the Connections related list to create a new Connection record.

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

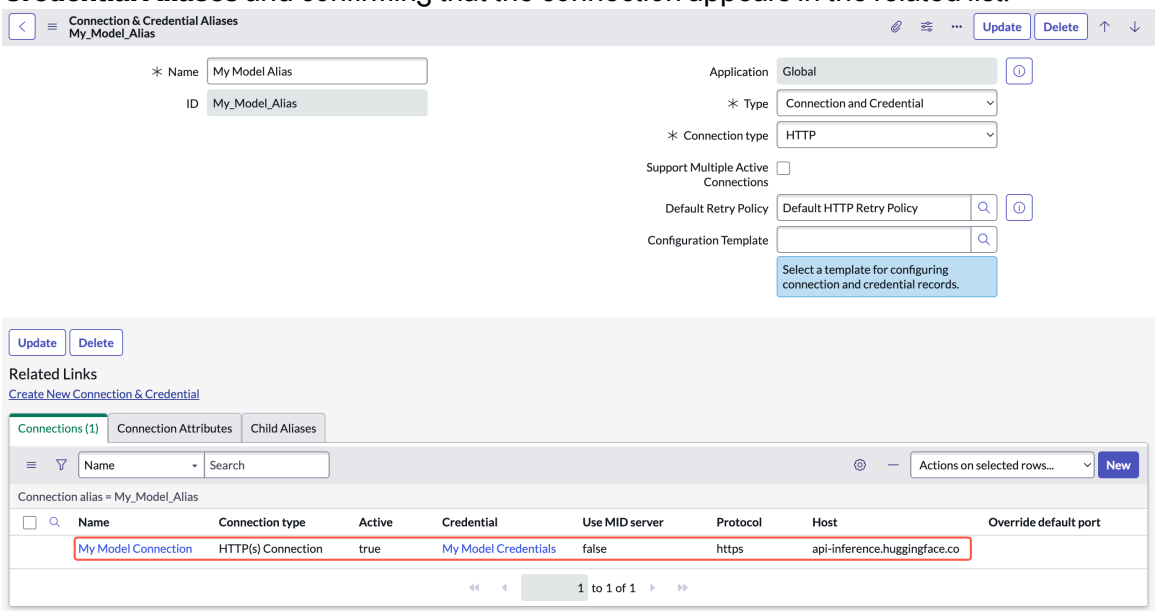
12. Enter the name of the connection in the Name field.
13. In the **Credential** field, select the Credential record you created with the model's API key. You can search the list of Credentials by typing the name of the Credential in the field. You can also select the lookup icon () to open a modal with the full list.
14. In the **Connection alias** field, select the alias record you created.
15. For the Connection URL, enter the endpoint URL for the model. For example, a model from Hugging Face might start with "https://api-inference.huggingface.co". You can leave the remaining fields as they are.



16. Select **Submit**.

Result

You have the connection and credential alias to use for connecting a generic LLM to use generative AI capabilities on the ServiceNow AI Platform. You can confirm this by navigating to **All** > **Connections & Credentials** > **Connections & Credential Aliases** and confirming that the connection appears in the related list.



What to do next

For more information on configuring a generic LLM, see configure a generic LLM connector

Configure a generic large language model (LLM) connector

Connect an external LLM to the ServiceNow AI Platform by using a generic LLM connector. With a connector, you can write your own prompts to send to the LLM and create your own generative AI capabilities.

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

Before you begin

You must configure API credentials for your LLM before setting up the connector. For more information, see configure API credentials for a generic LLM.

Role required: admin

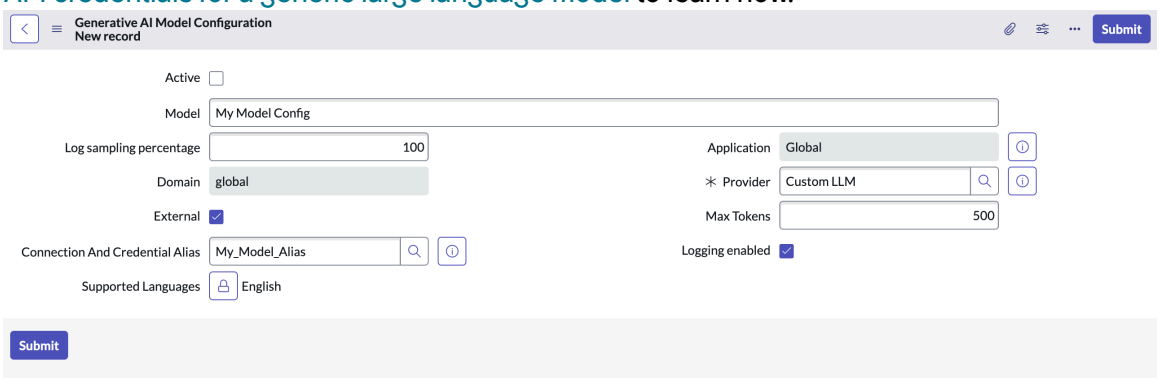
About this task

Generative AI Controller offers several base system connections to popular AI service providers such as Azure OpenAI and Google Vertex. However, you might want to incorporate an industry LLM, a case-specific LLM into your custom workflows, or an LLM that adheres to your organization's data handling policies.

Note: Generative AI Controller only supports text generation.

Procedure

1. In the navigation filter, go to the Generative AI Model Configuration [sys_generative_ai_model_config] table by entering `sys_generative_ai_model_config.list`.
2. Select **New**.
3. In the **Model** field, enter the name of the model.
   If you are using Azure OpenAI, then the model name is the deployment name of your resource.
4. In the **Provider** field, select **Custom LLM**.
5. In the **Connection and Credential Alias** field, select the alias that you created when you set up your credentials.
   If you haven't already configured your API credentials, see configure API credentials for a generic large language model to learn how.

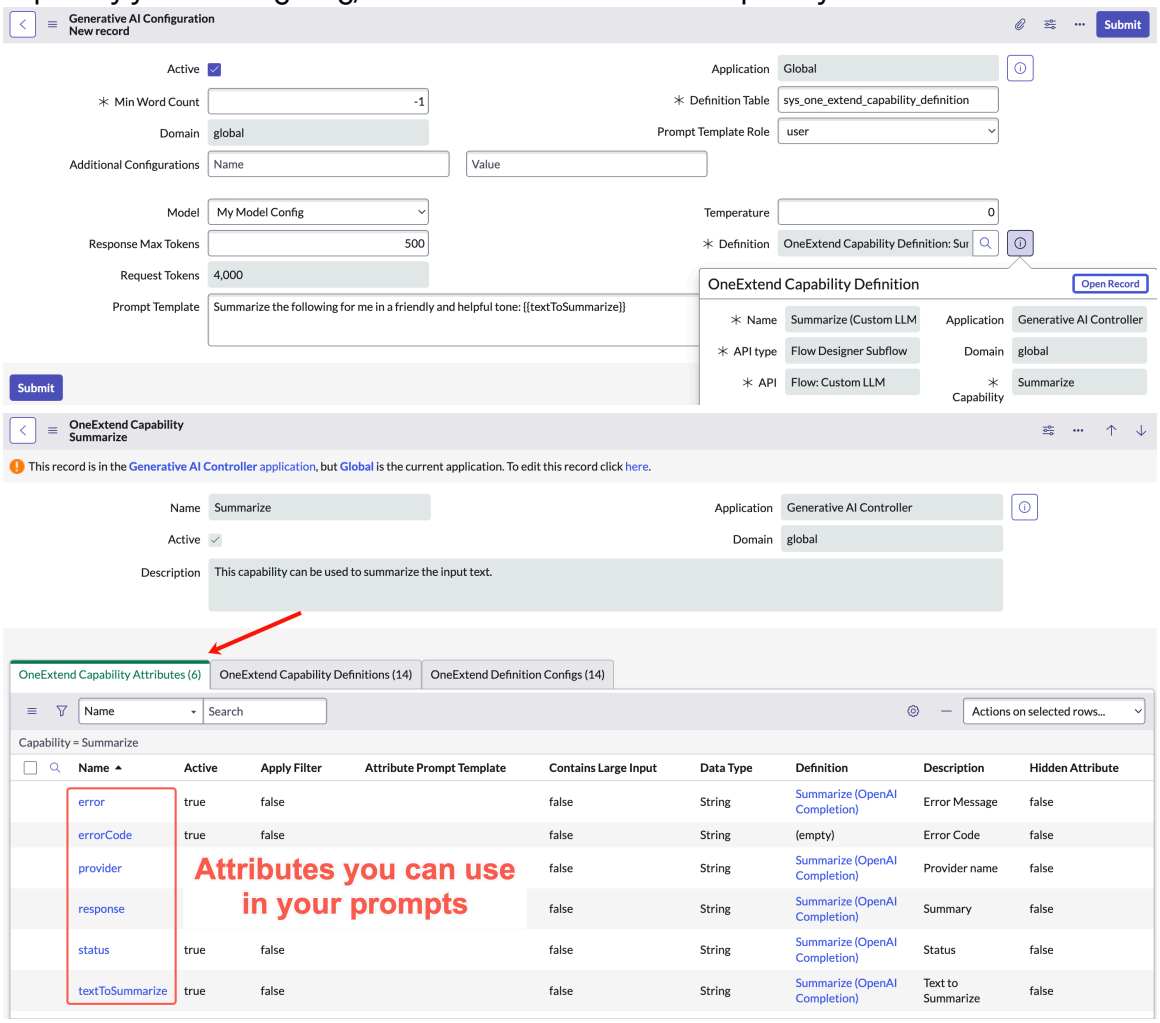


6. Create the model configuration by selecting **Submit**.
   The model configuration is associated with the API key of the external model.
7. In the navigation filter, go to the Generative AI Configuration [sys_generative_ai_config] table by entering `sys_generative_ai_config.list`.
8. Select **New**.
9. Select the list lookup icon () next to the **Definition** field.
10. In the Select the document modal, select a capability that you want to configure with the Custom LLM provider.
    For example, if you're configuring a Summarize capability, select **Summarize (Custom LLM)**.
11. Select **OK** to save your selection and close the modal.
12. In the **Model** field, choose the model configuration that you created in step 6.
    After you select a **Definition**, the **Model** field changes into a drop-down menu with options that use the Custom LLM provider. The model configuration that you made should appear in

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

this list. If it doesn't, go back to step 1 and make sure that your model configuration has the **Provider** field set to **Custom LLM**.

13. In the **Prompt template** field, enter the prompt for the capability.
    The prompt template is the instruction that is sent to the LLM to execute a capability. Use two braces around the capability attribute names to incorporate them into the prompt template. For example, if you're configuring a Summarize capability, your prompt template could be `Summarize the following for me in a friendly and helpful tone: {{textToSummarize}}`. To learn which attributes are available to include in the prompt, go to the OneExtend Capability [sys_one_extend_capability] table, find the record for the capability you're configuring, and look at the OneExtend Capability Attributes related list.



    You might need to experiment with different prompts to determine what works best for your use cases.

14. Create the new Generative AI Configuration by selecting **Submit**.

15. In the navigation filter, go to the Generative AI Custom LLM Transformer [sys_generative_ai_custom_llm_transformer] table by entering `sys_generative_ai_custom_llm_transformer`.

16. Select **New**.

17. Write transformer scripts.
    To have Generative AI Controller understand the format of the inputs and outputs of your custom LLM, you must write transformer scripts. When you create a transformer record, code and comments are provided for you to use as a guide while you edit. These scripts depend on the expected request and response objects that are interpreted by your model.

    For example, the Azure OpenAI request structure looks like the following script:

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

```
{"messages": [{"role":"user", "content":"Summarize the
 following text: <<content>>"}], "max_tokens": 800,
 "temperature": 0.7}
```

The request transformer script for that request structure is the following script:

```
(function(inputs) {
    /* write code here to construct the request body and any
 custom headers needed using the inputs object.
    inputs structure: {
        prompt_data: object,
        request_data: object
    } */
    var requestData = inputs.request_data;
    var promptData = inputs.prompt_data;
    var prompt = promptData.prompt;
    var model = promptData.model;

    // construct body using the inputs
    var body = {
        messages: [{
        "role": "user",
        "content": prompt
        }],
        max_tokens: parseInt(promptData.max_tokens),
        temperature: parseInt(promptData.temperature)
    };

    //construct headers using the inputs
    var headers = {};
    return {
        body: body,
        headers: headers
    };
})(inputs);
```

The response structure from Azure OpenAI looks like this script:

```
{
    "choices": [{
    "finish_reason": "stop",
    "index": 0,
    "message": {
        "content": "<<response>>",
        "role": "assistant"
        }
    }],
    "created": 1714994995,
    "id": "chatcmpl-9LqpXeLVXDAi6kciPfLeIDjmALeea",
    "model": "gpt-35-turbo-16k",
    "object": "chat.completion",
    "usage": {
        "completion_tokens": 47,
        "prompt_tokens": 70,
        "total_tokens": 117
    }
}
```

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

Because of that response structure, the response transformer script looks like this script:

```
(function(inputs) {
    /* write code here to transform the llm response into an
 array of text responses, using the inputs object
    inputs structure: {
        prompt_data: object,
        request_data: object,
        response_body: string,
        response_headers: string
    } */
    var requestData = inputs.request_data;
    var promptData = inputs.prompt_data;
    var responseBody = JSON.parse(inputs.response_body);
    gs.info("response : " + inputs.response_body);
    var responseTexts = [];

    // write code here to populate the responseTexts array.

 responseTexts.push(responseBody.choices[0].message.content);

    return responseTexts;

})(inputs);
```

18. Create the transformer scripts by selecting **Submit**.

Result

Your external LLM is connected to Generative AI Controller. You can use the AI service provider for generative AI capabilities on the ServiceNow AI Platform.

What to do next

After you connect the LLM, you can set a provider for the generative AI capabilities of Generative AI Controller. The capabilities are Generic Prompt, Generate Content, Sentiment Analysis, and Summarize. For more information on choosing a provider, see Set a provider for a generative AI capability.

Related topics

Bring your own key for third-party AI provider integration

Configure a custom resource path for BYOK models

Set a provider for a generative AI capability

Determine which AI provider to use for each of the generative AI capabilities of Generative AI Controller.

Before you begin

Configure your credentials for your preferred provider. See Configuring API credentials for generative AI capabilities for more details.

Role required: admin

About this task

Generative AI Controller comes with four capabilities by default. You can only configure a provider for the capabilities listed in the table in step 2. If you have installed other Now Assist applications, you may see other capabilities on the OneExtend Capability table. Those other capabilities use the Now LLM Service as their provider and can't be reconfigured.

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

Procedure

1. In the navigation filter, search for the OneExtend Capability table by entering `sys_one_extend_capability.list`.
2. Open the record for the capability that you would like to configure, such as Sentiment Analysis.

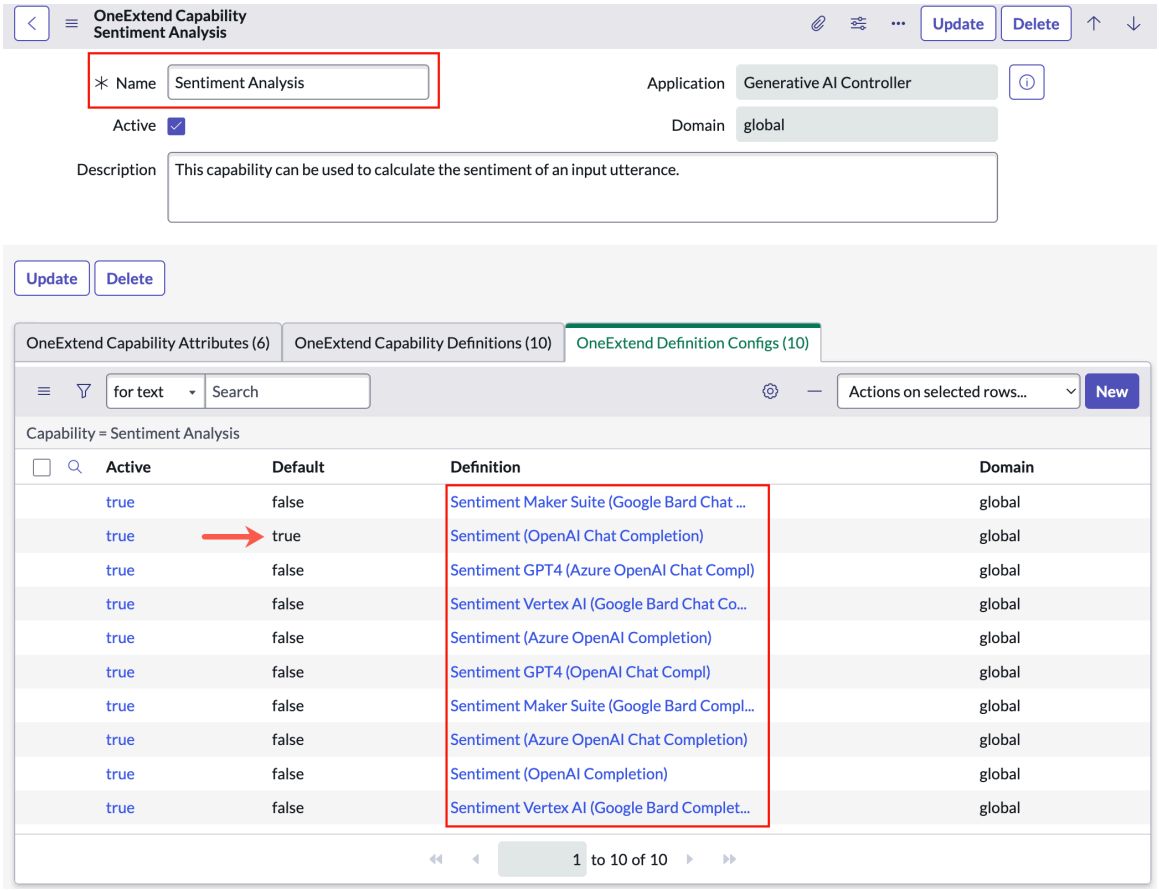
Generative AI Controller capabilities

Capability	Definition
Generate Content	Generate texts about a given topic with Workflow Studio and Virtual Agent Designer.
Generic Prompt	Create your own use case and prompt.
Sentiment Analysis	Analyze user sentiment with Workflow Studio and Virtual Agent Designer.
Summarize	Create summaries of topics with Workflow Studio and Virtual Agent Designer.

3. In the OneExtend Definition Configs related list, set **Default** to `true` for your preferred capability provider.

Note: By default, you can choose only one provider for a capability. For example, if **Default** is `true` for Sentiment (OpenAI Completion), you must set **Default** to `false` before changing **Default** to `true` for Sentiment (Azure OpenAI).

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



Model Options

You have the option to choose between several different models for each capability.

Capability Definition	Model
○ OpenAI Completion ○ Azure OpenAI Completion	GPT-3
○ OpenAI Chat Completion ○ Azure OpenAI Chat Completion	GPT-3.5
○ GPT4 (OpenAI Chat Compl) ○ GPT4 (Azure OpenAI Chat Compl)	GPT-4
○ AI Studio (Google Cloud Completion) ○ AI Studio (Google Cloud Chat Completion) ○ Vertex AI (Google Cloud Completion) ○ Vertex AI (Google Cloud Chat Completion)	Google Gemini
Aleph Alpha Completion	Luminous
IBM watsonx	Granite

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

Result
Flows, topics, and scripts that use the generative AI capability use the provider you've specified.

Related topics
- Bring your own key for third-party AI provider integration
- Configure a custom resource path for BYOK models

Select a model for Amazon Bedrock
Choose which large language model (LLM) to use with Amazon Bedrock for custom skills.

Before you begin
You must have the latest version of Generative AI Controller and the Amazon Bedrock spoke installed on your instance. You must also set up your API credentials. For more information, see Configure API credentials for Amazon Bedrock.

Role required: admin

About this task
You may have multiple models configured to use with Amazon Bedrock. By default, the settings use Amazon's Titan model, but you can configure the spoke to use a different model by changing the Generative AI Model Config record.

Procedure
1. In the filter navigator, navigate to the Generative AI Model Config table by entering `sys_generative_ai_model_config.list`.
2. Search in the Provider field column for `Amazon Bedrock`, then open the resulting record.
3. In the Model field, enter the ID of the model that you want to use.
   You can find the model ID on the **Foundation models** > **Base models** page on Amazon Bedrock.
4. Select **Save** to update the record.
5. Map the chosen model to a custom skill.

   a. In the filter navigator, go to the Generative AI Config table by entering `sys_generative_ai_config.list`.

   b. Search for the custom skill you've created that uses Amazon Bedrock as the service provider, and open the record.
   The Definition should be OneExtend Capability Definition: <Name of Skill>.

   c. Set the Model field to the model ID you entered in step 3.

Result
Your chosen model with the Amazon Bedrock provider will be used for custom skills.

What to do next
You can create custom skills with the Amazon Bedrock provider in Now Assist Skill Kit and perform step 6 to set the new model.

Configure small talk filters
Redirect users to different Virtual Agent topics if small talk, such as greetings, expressions of gratitude, complaints, or requests to close, are detected in conversations.

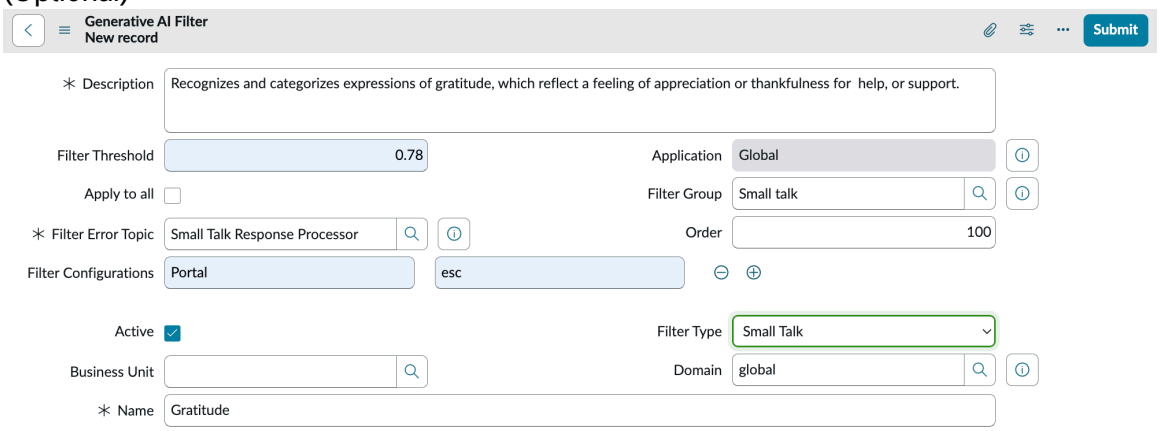
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

Before you begin
Role required: admin

About this task
When engaging with generative AI, many people send conversational messages that are ultimately unhelpful for the Virtual Agent when determining how to help the user achieve their goals. You can configure small talk filters that catch these conversational messages. When a user is sending a message such as a greeting or expressing complaints about the conversation, you can redirect them to serve their needs.

Procedure
1. In the navigator, go to the Generative AI Filter [sys_gen_ai_filter] table by entering `sys_gen_ai_filter.list`.
2. Open the record for the small talk filter that you want to configure or select **New** to create one. By default, the four available filters for small talk are Greetings, Gratitude, Complaint, and Closure. If you see other filters, those filters are sensitive topic filters that can be set up with the Now Assist Admin console. See Configure subject filters for generative AI for more information.
3. If creating a filter, add a **Name** and **Description** for the filter.
4. Choose the Filter Error Topic that you want to redirect users to when the small talk filter is detected.
   By default, the user is redirected to the Small Talk Response Processor topic.
5. For the **Order** field, choose the order to apply the filters.
   Filters are processed by lowest order to highest. A filter with order 100 is processed before a filter with order 200.
6. **Optional:** In the **Filter Configurations** field, use the name "portal" to determine where the filter should be run, such as setting "portal" to value "sp" to apply only the filter on the sp portal. Use a comma-separated list to select multiple values.

   (Optional)



7. Select **Save** to save any changes.
8. In the navigator, go to the Generative AI Filter Sample [sys_gen_ai_filter_sample] table by entering `sys_gen_ai_filter_sample.list`.
9. **Optional:** Review the sample phrases for each filter.
   Statements that match these sample phrases activate the selected filter. You can create sample phrases manually by creating records on the table or by using generative AI to generate samples for you.

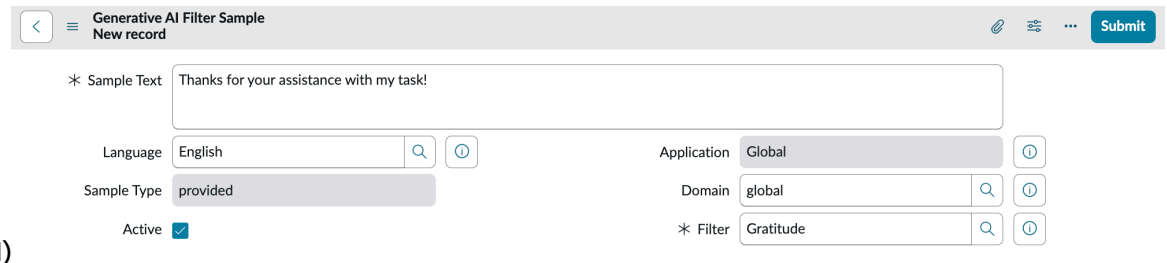
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

a. To create a sample manually, select **New**.

b. In the **Sample text** field, enter the sample text of what you want to be caught by the filter.

c. In the **Filter** field, select the filter you want the sample to be applied to.

d. If you want to generate more sample phrases with generative AI, open a record for a sample phrase for the desired filter.
An example might be the "Thanks for your assistance with my task" sample filter for the Gratitude filter.



(Optional)

e. Select **Generate Samples** to create additional example statements to trigger the filter.
The more sample statements that you provide, the more likely that the filter aligns with your expectations of its behavior.

Result

When engaging in conversations with generative AI, users are redirected to a new topic when filtered subjects or sentiments are detected.

Configure AI search answers capability for web search

AI web search is a OneExtend capability that enables end users to perform web searches and receive AI-powered answers. The capability supports multiple AI providers and integrates with virtual assistants and workflows.

Before you begin

When web search is turned on, Gemini is the default AI provider. To switch AI providers, admins must configure the AI Search answers OneExtend capability. The instance or skills default AI provider that is set in Now Assist Admin can differ from the web search AI provider. For more information on the default instance or skills AI provider, see Manage model providers.

Some of the AI providers require API keys whereas others don't require API keys because they use OEM. The following table highlights the four AI providers and whether they use API keys or OEM.

AI providers that require API keys	AI providers that use OEM
OpenAI	Gemini

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

AI providers that require API keys	AI providers that use OEM
	Note: Gemini is the default web search AI provider in the AI Search answers OneExtend capability. If you selected Gemini as the LLM provider in Now Assist Admin and you want to use Gemini as your web search AI provider, you don't need to configure the AI Search answers OneExtend capability.
Perplexity	Azure **Note:** Azure doesn't currently support web search.

Role required: admin

Procedure

1. In the navigator filter, go to the OneExtend Capability list by entering `sys_one_extend_capability.list`.
2. Search for the `AI Search answers` OneExtend capability.
3. Select the AI Search answers record.
4. In the OneExtend Definition Configs related list, view the Default column.
   When the AI provider's Default column is set to `true`, that is the AI search provider used for the web search. Gemini defaults to `true`. To change the AI search provider to an AI provider that requires the use of API keys, such as Perplexity or OpenAI, continue with the following steps.

   Note: Although Azure OpenAI answer is an option for an AI search provider, Azure doesn't support web search and should remain set to `false`.

5. In the OneExtend Capability Definitions related list, select the value in the Connection And Credential Alias column for the AI search provider you want to use.
   For example, if you want to switch the web search AI provider to OpenAI, select **sn_openai.OpenAI** from the Connection And Credential Alias column.
   The Connection & Credential Alias form opens. If no connections are available within the Connections related list, you must create a connection.
6. Complete one of the following actions based on the following scenarios.

Scenario	Action required
Connection is available	If a connection is already available within the Connections related list, continue to the next step.
Connection isn't available and must be created	**a.** Select **New** within the Connections related list to create a connection. **b.** In the **Name** field, enter the connection name. For example, `Open AI connection`.

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

Scenario	Action required
	Note: For more information about the fields on the HTTP(s) Connections New record form, see Create an **c.** Select the Lookup using list icon ( ) next to **Credential**. **d.** Select **New** to create a credential for your connection. **Note:** You may see an existing credential that you can select from the Credentials list, but some AI providers, like OpenAI, do not come with default credentials. **e.** Select **API Key Credentials**. **f.** Enter a name in the **Name** field. For example, `Open AI key`. **g.** Enter the API key in the **API key** field. **Note:** For more information about the fields on the API Key Credentials New record form, see API key creden **h.** Select **Submit**. You're redirected back to the HTTP(s) Connections new record form. **i.** Enter the connection URL in the **Connection URL** field. For example, enter `https://api.openai.com` for Open AI. **j.** Select **Submit**.

a. In the OneExtend Definition Configs related list, set the existing default AI provider from `true` to `false`.

b. In the OneExtend Definition Configs related list, set the AI provider that you want to use from `false` to `true`.

7. Select **Update**.
The AI provider for web search is now updated.

Result

End users can now perform an internet search to answer a query whenever in web search mode. For more information about web search mode for end users, see Web search .

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

What to do next
If you switched the AI provider to Perplexity, you must work with the *sn_ai_websearch.perplexity_model_name* system property. For more information on this system property, see Now Assist in Virtual Agent system properties.

Enable recursive summarization for large inputs

Use recursive summarization to break down the requests to the large language models (LLMs) into smaller pieces so that you can maintain the context for generative AI capabilities.

Before you begin
Role required: admin

About this task
LLMs have a maximum number of tokens that can be processed in a single request. Certain fields, such as activity fields, can have more information than can fit in within those restrictions. Recursive summarization breaks the information given to an LLM into chunks, summarizes each chunk individually, and then processes the original request with the summarized chunks. The chunks are organized with overlaps between the pieces so that the context is retained across every piece.

Note: Enabling recursive summarization may cause the capabilities to process large inputs more slowly because they must make multiple calls to the LLM instead of just one call.

Procedure
1. In the navigator filter, go to the OneExtend Capability list by entering `sys_one_extend_capability.list`.
2. Open the record for the OneExtend capability that you want to change.
3. In the OneExtend Definition Config related list, set **Enable Large Input Support** to true for the OneExtend Definition that you want to enable recursive summarization for.
4. In the OneExtend Capability Attributes related list, set **Contains Large Input** to true for the fields that you want to add recursive summarization in.
   The fields that are most likely to contain a large amount of data, such as an activities field, should have their values set to true. You can also select the **Contains Large Input** check box on the OneExtend Capability Attribute record and save the record to set the value to true.

Result
Recursive summarization is enabled for the OneExtend Capability for the fields specified in this procedure.

Enable Dynamic Translation for capabilities in Generative AI Controller

Use Dynamic Translation to add multiple language support for generative AI capabilities to support users who speak languages other than English.

Before you begin
You must have the Dynamic Translation and Conversational Dynamic Translation plugins installed and enabled on your instance. For more information on setting up Dynamic Translation, see configuring Dynamic Translation.

You must have a default Translation Configuration set in the Translation Configuration table (sn_dt_translator_configuration).

Role required: admin

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

About this task
By default, Dynamic Translation is inactive for capabilities with Generative AI Controller.

Dynamic Translation isn't available for capabilities in Virtual Agent or Now Assist panel.

Procedure
1. Navigate to the OneExtend Capability (sys_one_extend_capability) table by entering `sys_one_extend_capability.list`in the navigator.
2. Open the record for the capability that you want to enable Dynamic Translation for.
3. In the OneExtend Capability Attributes related list, open the record for the attribute you would like to be translated, such as response.
4. Select **Translate** to enable translation.
   If you don't see the field on the form, you must edit the form to display the **Translate**. One way of adding the field to the form is by selecting the menu icon by the table name of the form and choosing **Configure** > **Form Layout**. Add the **Translate** field to the Selected list. Then, select **Save** to save the form. You're then redirected back to the record with the field available to select.
5. Select **Submit** to save the record.

Result
Dynamic Translation is installed for your generative AI capability on the attribute that you selected.

Disable Dynamic Translation for LLM Virtual Agent conversations

Enable dynamic translation of chat messages into English before they are sent to the ServiceNow large language model (Now LLM) in generative AI topics to support users who speak other languages.

Before you begin
You must have Dynamic Translation for Virtual Agent installed and active for your Virtual Agent. For more information, check out .

Role required: admin

About this task
Dynamic Translation is enabled by default.

There are certain limitations of Dynamic Translation for Virtual Agent. Catalog item names are not automatically translated, and only languages configured for language detection will be recognized.

Procedure
1. Navigate to the System Properties table by entering `sys_properties.list` in the navigator.
2. Select the magnifying glass icon () to expand the column search row.
3. In the Name column, enter `sn_generative_ai.disable_dynamic_translation` and press **Enter** to search, and then open the matching record.
4. Set the property value to `true`.
5. Select **Save** to save the record.

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

Result
Chat requester utterances are translated into English before being sent to Now LLM and output messages are translated back into the requester's preferred language.

Configure rate limiting for providers

Configure rate limiting to control the traffic flow to the one_extend request by restricting the number of requests that can be made within a certain time frame.

Before you begin
Role required: admin

About this task
To rate limit the one_extend request, you must define the rate limit rules for the incoming requests. You can create your own rule in the One Extend Rate Limit Rules table.

Procedure
1. On your ServiceNow instance, navigate to the One Extend Rate Limit Rules [sys_one_extend_rate_limit_rules] table and select **New**.
2. On the form, fill in the fields.

One Extend Rate Limit Rule form

Field	Description
Name	Name of the rate limit rule. Identify the rate limit rule with a name.
Application	Application scope that must be set to Global.
Resource	Resources to select for your rate limit rule. Use one of the following resources for your rate limit rule: ○ **LLM Provider**: For the LLM Provider, select the generative AI provider mapping. ○ **Capability Definition**: For the Capability Definition, select the One Extend Capability Definition. ○ **All Capabilities**: Applies to all providers and all capability definitions.
Rule Type	Rule type for the rate limit rule. Define the rule type using the following options: ○ **Instance**: Accumulates data from all the users of that instance. ○ **User**: Keeps the count of the requests per user.
Request limit per hour	Maximum number of requests to be received per hour.
Active	Status of the rate limit rule.

3. Select **Submit**.

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

Bring your own key for third-party AI provider integration

The bring your own key (BYOK) feature enables you to use your own API credentials from supported cloud AI providers, such as Azure OpenAI, Amazon Bedrock, and Google Gemini, to run Now Assist skills and AI agents.

When BYOK is enabled, Now Assist sends model requests through your cloud AI provider account instead of ServiceNow AI Platform infrastructure. This approach lets you maintain ownership of your AI configuration, security controls, and data processing environment while continuing to use Now Assist capabilities.

BYOK supports the same AI models that the ServiceNow AI Platform supports. After you configure your API credentials, Now Assist uses your provider resources to process requests instead of using the default ServiceNow AI provider. BYOK keeps your keys under your control and AI workloads running within your cloud environment.

Note: Using your organization's cloud AI provider account might result in additional charges from that provider. Review the provider's pricing and terms of service before enabling BYOK.

ServiceNow integrated model providers

BYOK supports the following cloud AI providers:

- Azure OpenAI
- Anthropic Claude on AWS, which is accessible through Amazon Bedrock
- Google Gemini

Configuring a BYOK connection with a cloud AI provider

To use your cloud AI provider resources within your ServiceNow instance, configure a connection between ServiceNow and your provider. Configuration steps vary by provider, but follow a similar high-level configuration workflow as the following steps:

1. Create connection and credentials for the provider, such as endpoint, API key, and region.
2. Enable BYOK for the provider in Now Assist Admin.
3. Map ServiceNow model aliases to provider model versions. For example, map `gemini_large` → 2.5 Pro.
4. Optionally, override variants in each skill in Version Management.
5. Validate that requests are using your AI provider account by reviewing outbound HTTPS logs.

For step-by-step instructions for each supported cloud AI provider, see Configure Bring Your Own Key (BYOK) for Now Assist and AI Agents [KB2666298] article in the Now Support Knowledge Base.

Example: High-level BYOK configuration for Azure OpenAI

1. Set up Azure OpenAI connection with your own key to use Azure OpenAI within your ServiceNow instance.

   For more information, see Configure API credentials for Azure OpenAI.

2. Configure your Azure OpenAI model and deployment name in Generative AI Controller so your ServiceNow instance can recognize and use it.

   For more information, see .

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

3. Modify the prompt configuration to write your own prompts to send to the Azure OpenAI LLMs.

   For more information, see Configure a generic large language model (LLM) connector.

4. Switch providers to use BYOK for AI processing.

   For more information, see Manage Integration.

Related topics

Configure a custom resource path for BYOK models

Configure a custom resource path for BYOK models

Enter a custom resource path in your bring your own key (BYOK) model configuration so that Generative AI Controller can connect to AI service providers, such as Azure OpenAI, that use a different web address than the default.

Before you begin

Role required: admin

About this task

When you configure a bring your own key (BYOK) model provider in Generative AI Controller, it uses the connection and credential alias to send requests to your AI provider. The alias contains the base URL and API key for your provider. Some AI providers, such as Azure OpenAI, include an extra path segment in their web address that Generative AI Controller does not add by default. Without this segment, Generative AI Controller requests do not reach your provider correctly.

The Resource Path field in the model configuration record is where you provide that extra segment.

Note: The custom resource path configuration is currently supported for Azure OpenAI only.

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connections & Credential Aliases**.
2. Open the Generative AI provider record for Azure OpenAI.
3. In the Connections Attributes related list, select the **Chat Completions Resource Path** attribute.
4. Enter the custom resource path for your provider's endpoint in the **Default value** field.
   For example, for Azure OpenAI, if your full endpoint URL is:

   ```
   https://<your-resource-name>.openai.azure.com/openai/deployments/<your-deployment-name>
   ```

   Enter the following resource path in the **Default value** field:

   ```
   /openai/deployments/<your-deployment-name>
   ```

   Note: You can find your resource name and deployment name in the Azure portal.

5. Select **Update**.

Result

Generative AI Controller now connects to your AI provider using the resource path you entered. Your agentic AI skills are ready to use this provider.

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

Reference for Generative AI Controller

Reference topics provide information about Generative AI Controller flow actions, tables, and properties.

Tables installed

Name	Table	Description
OneExtend Capability	sys_one_extend_capability	Generative AI Controller capabilities that include Summarize, Record Summarization, Generate Content, and Generic Prompt.
OneExtend Capability Definition	sys_one_extend_capability_definition	Attribute configuration for input and output variables for Workflow Studio subflows.
OneExtend Capability Definition Attribute	sys_one_extend_definition_attribute	Input and output variables for Workflow Studio subflows. Variable names can't be changed if the capability is active and used on the instance. You can check whether a capability is used by going to the OneExtend Usages table.
OneExtend Builder Config	sys_one_extend_builder_config	Determines which capability and provider is related to each builder component for Workflow Studio and Virtual Agent Designer.
OneExtend Builder Capability	sys_one_extend_builder_capability	Definitions for a capacity and its provider for builder components.
OneExtend Usage	sys_one_extend_usage	Each usage of a capability in a Workflow Studio or Virtual Agent Designer topic, as well as any scripts such as business rules or UI actions.
Gen AI Log Metadata	sys_gen_ai_log_metadata	Log data about requests to the LLMs, including information about definition, errors, user, and feedback provided. AI-generated content can be tracked for a duration beyond six months with Now Assist configuration option. You can export historical data by writing a script to copy it into a different table without deleting the information.
Generative AI Metric	sys_generative_ai_metric	Logs various metrics to help evaluate the performance and accuracy of LLM responses, such as edit score, edit distance, guardrail activity, and details about the LLM model used.

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

Properties

Generative AI Controller properties

Property	Description
com.sn.generative.ai.provider	Default provider when capability definition has no default. Type: choice list No default value
com.sn.generative.ai.ais.message	Message that is displayed when AI Search fails to find an answer to a query. Type: string Default value: No answer found.
com.sn.generative.ai.log_prompt	Prompt that determines whether generative AI API calls are logged. Type: true | false Default value: true
com.sn.generative.ai.moderation.message	Message that is displayed if the OpenAI or Azure OpenAI moderation tools identify the content that goes against their terms of service. Type: string Default value: The response cannot be displayed as it’s deemed inappropriate by OpenAI.
com.glide.one.extend.token.buffer	Buffer that checks the request for the number of tokens before a OneExtend capability is executed. The maximum allowed request tokens are calculated based on the maximum tokens that are permitted by the AI provider's API minus the response token and buffer value that is specified in this system property. Type: integer Default value: 250
domain.llm.usage.entitled	Determines whether to use the LLM for a specific domain to process data or restrict its use. Type: true | false Default value: true

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

External links

Provider	Data policy	Usage policy
Amazon Bedrock	Data protection	AWS Service Terms
Aleph Alpha	Data privacy	Terms and Conditions
Google Cloud	Google Cloud Platform Terms of Service	Google Cloud Platform Terms of Service
IBM watsonx	Keeping your data secure and compliant	Foundation model terms of use in watsonx.ai
Microsoft Azure OpenAI	Data, privacy, and security for Azure OpenAI Service	Code of conduct for Azure OpenAI Service
OpenAI	API data usage policies	Usage policies

Generative AI Controller tables

Generative AI Controller use dedicated tables to log AI activities and track Now Assist usage across Now Assist capabilities.

Generative AI Log [sys_generative_ai_log] table

Logs generative AI prompts, responses, and edited responses for debugging LLM calls. HR-related log entries are restricted to HR administrator to protect sensitive information.

Generative AI Log [sys_generative_ai_log]

Column	Data type	Description
Additional Data	String	Contains concatenated values for actions performed on the generated LLM response, such as copied or edited. This field is populated through the feedback API.
Caller	String	Contains information about the calling application that generated the log entry.
Completed At	Date/Time	Timestamp when the generative AI execution completed.
Conversation	Reference	Reference to the conversation this log is part of.
Created	Date/Time	Date and time when the log entry was created.
Created By	String	Contains user information that initiated the execution or flow.
Definition	Reference	Reference to the capability definition responsible for generating the log entry.

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

Generative AI Log [sys_generative_ai_log] (continued)

Column	Data type	Description
Derived Scope	Reference	Associated scope of the application that generated the log entry.
Domain	Domain ID	Domain associated with the execution that generated the log entry.
Domain Path	Domain Path	Path associated with the domain from where the capability is executed. When domain separation is enabled, this column is automatically populated by the platform.
Edited Response	String	Captures the edited version of the response when a user modifies it.
Error	String	Error message, if the request encountered an error.
Error Code	String	Error code associated with the error message.
External	True/False	Error code associated with the error message.
Feedback	Choice	Feedback on the LLM response. Possible values: Accepted, Rejected, Ignored, SlightlyPositive, SlightlyNegative, LikelyPositive.
Feedback Timestamp	Date/Time	Timestamp when the feedback was submitted.
Gen AI Usage Log	Reference	Reference to the Gen AI Usage Log.
Metadata Document	Document ID	Document ID for metadata used for the requested execution, if only a single document is used.
Metadata Documents	Glide List	Document IDs for metadata used for the requested execution, if multiple documents are used.
Metadata Table	Table Name	Table name for the metadata.
Model Name	String	Name of the model used for processing the generative AI request, such as gpt_large, claude_small, or gemini_small.

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

Generative AI Log [sys_generative_ai_log] (continued)

Column	Data type	Description
Model Version	String	Version of the ServiceNow LLM model used.
Output Metadata	String	Metadata received from the model or LLM, including tokens, processing time, and response time.
Prompt	String	Final prompt passed to the model.
Prompt Config	Reference	Reference to the prompt configuration, which contains settings such as model and temperature used in the request to the LLM.
Prompt Token Count	Integer	Number of tokens used in the prompt.
Response	String	The response generated by the LLM.
Response Token Count	Integer	Number of tokens in the response.
Skill Config Id	Reference	Reference to the skill configuration.
Source	String	Source feature responsible for the execution, such as ChatSummary, Now Assist in AI Search, or WWNA: Refine.
Started At	Date/Time	Timestamp when the generative AI execution started.
Status	String	Status of the AI log. Possible values: Success, Error.
Target Language	String	Target language for the AI log.
Time Taken	Integer	Time taken for the AI log process.
Untranslated Prompt	String	The untranslated version of the prompt.
Updated	Date/Time	Date and time when the log entry was last updated.
Updated By	String	User who made updates to the log.
Updates	Integer	Number of updates made to the log record.

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

Gen AI Usage Log [sys_gen_ai_usage_log] table

Tracks each use of a Now Assist capability triggered from Workflow Studio, Virtual Agent Designer topics, or scripts such as business rules and UI actions.

Gen AI Usage Log [sys_gen_ai_usage_log]

Column	Data Type	Description
Assists	Integer	Number of assists used or consumed.
Created	Date/Time	Date and time when the log entry was created.
Created By	String	Contains user information that initiated the execution or flow.
Document	Document ID	Document ID for metadata used for the requested execution.
Document Table	String	Table name for the document table.
Domain	Domain ID	Domain associated with the execution that generated the log entry.
Domain Path	Domain Path	Path associated with the domain from where the capability is executed.
Execution Type	Choice	Type of execution. Possible values: Test, Evaluation, Auto Prompt Tuning.
Feature	Reference	Source feature for the assist usage.
Line Item Name	String	License line item name associated with the assist usage.
License Name	String	License name associated with the assist usage.
Plugin Name	String	Plugin name that the feature responsible for the assist usage is part of.
Skill Config Id	Reference	ID of the skill responsible for the assist usage.
Status	Choice	Status of the usage log. Possible values: Completed, Pending.
Strategy	Choice	Strategy used for pricing the assist usage.
Trial Assists	Integer	Number of assists counted as trial assists.

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

Gen AI Usage Log [sys_gen_ai_usage_log] (continued)

Column	Data Type	Description
User	Reference	User associated with the assist usage.

Now Assist AI agents

The ServiceNow® Now Assist AI agents are entities that mimic human-like intelligence by using large language models (LLMs). AI agents can perform tasks that range from simple automated responses to complex problem solving. By using AI agents, you can reduce the workloads of your live agents and help increase their productivity.

https://player.vimeo.com/video/1164069631?h=c5eed6a89e&badge=0&autopause=0&player_id=0&app_id=58479

Learn more about the concepts behind AI agents and how they can help fit your specific business needs.

Get started

Select a tile to get started with Now Assist AI agents.

Explore  Learn more about Now Assist AI agents	Configure  Configure the agentic AI experience in AI Agent Studio	Create an AI agent  Create your own custom AI agents with advanced multi-agent reasoning frameworks
Create an agentic workflow  Create your agentic workflows with planned instructions and triggers	Reference  Learn more about the user roles, tables, and properties in Now Assist AI agents	

There might be AI agents installed with the Now Assist application that are not used in agentic workflows. To learn how to see all agents that are available to you, see Find AI agents.

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

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.

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.

For more information, see the Now Assist documentation.

Explore Now Assist AI agents

Learn how the Now Assist AI agents enhance live agent productivity by mimicking human-like intelligence to manage tasks ranging from automated responses to complex problem solving.

Now Assist AI agents overview

The Now Assist AI agents are virtual agents within the ServiceNow AI Platform ecosystem. They can perform specific tasks and functions, often using natural language instead of traditional code.

AI agents can perform the following tasks:

- Analyze data.
- Retrieve information from knowledge bases and enterprise systems.
- Execute automated actions.
- Collaborate with other agents to resolve issues or accomplish requests.

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

Example:#Imagine you need to reset your password. Multiple AI agents might be involved:

1. One agent understands the request.
2. Another verifies your account details.
3. A third agent conducts a security check.
4. A fourth agent resets the password and notifies you.

Agentic AI framework

AI Agent Studio

The AI Agent Studio enables you to create, manage, and test AI agents and agentic workflows in a unified environment.

agentic system

The ServiceNow AI Platform's agentic system uses AI agents, orchestrated by an Orchestrator, to execute Agentic workflows.

AI agents

On the ServiceNow AI Platform agentic system, an AI agent contains a set of large language model (LLM) instructions with the tools to accomplish tasks.

agentic workflow

Agentic workflows are smart, context-aware, automated processes designed and executed by AI agents. They represent dynamic, human-like processes where specialized AI agents collaborate to accomplish complex goals. Here are a few key characteristics of Agentic workflows:

- **Autonomous**: Operate independently without constant human intervention.
- **Adaptive**: Adjust to changing situations and business needs.
- **Collaborative**: Work together to complete complex tasks.
- **Intelligent**: Use large language models and business data to make informed decisions.

Orchestrator

The AI agent Orchestrator is a central management system that coordinates AI agents to ensure they collaborate effectively to complete complex workflows.

For more information, see Understand the Now Assist AI agents.

Guided Setup

The guided setup process in the helps you configure your AI Agent triggers, data sources, and display locations based on your specific business needs.

Now Assist panel

The Now Assist panel is a user-facing interface, commonly accessed through the context menu or chat interfaces. It provides quick access to agents' capabilities and enhances productivity by summarizing records, creating content, and analyzing alerts.

Tools

Tools equip your AI agents with the necessary capabilities to complete their tasks. Adding a tool to an AI agent enables various functionalities and helps the AI agents achieve their objectives.

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

How to put Now Assist AI agents to work

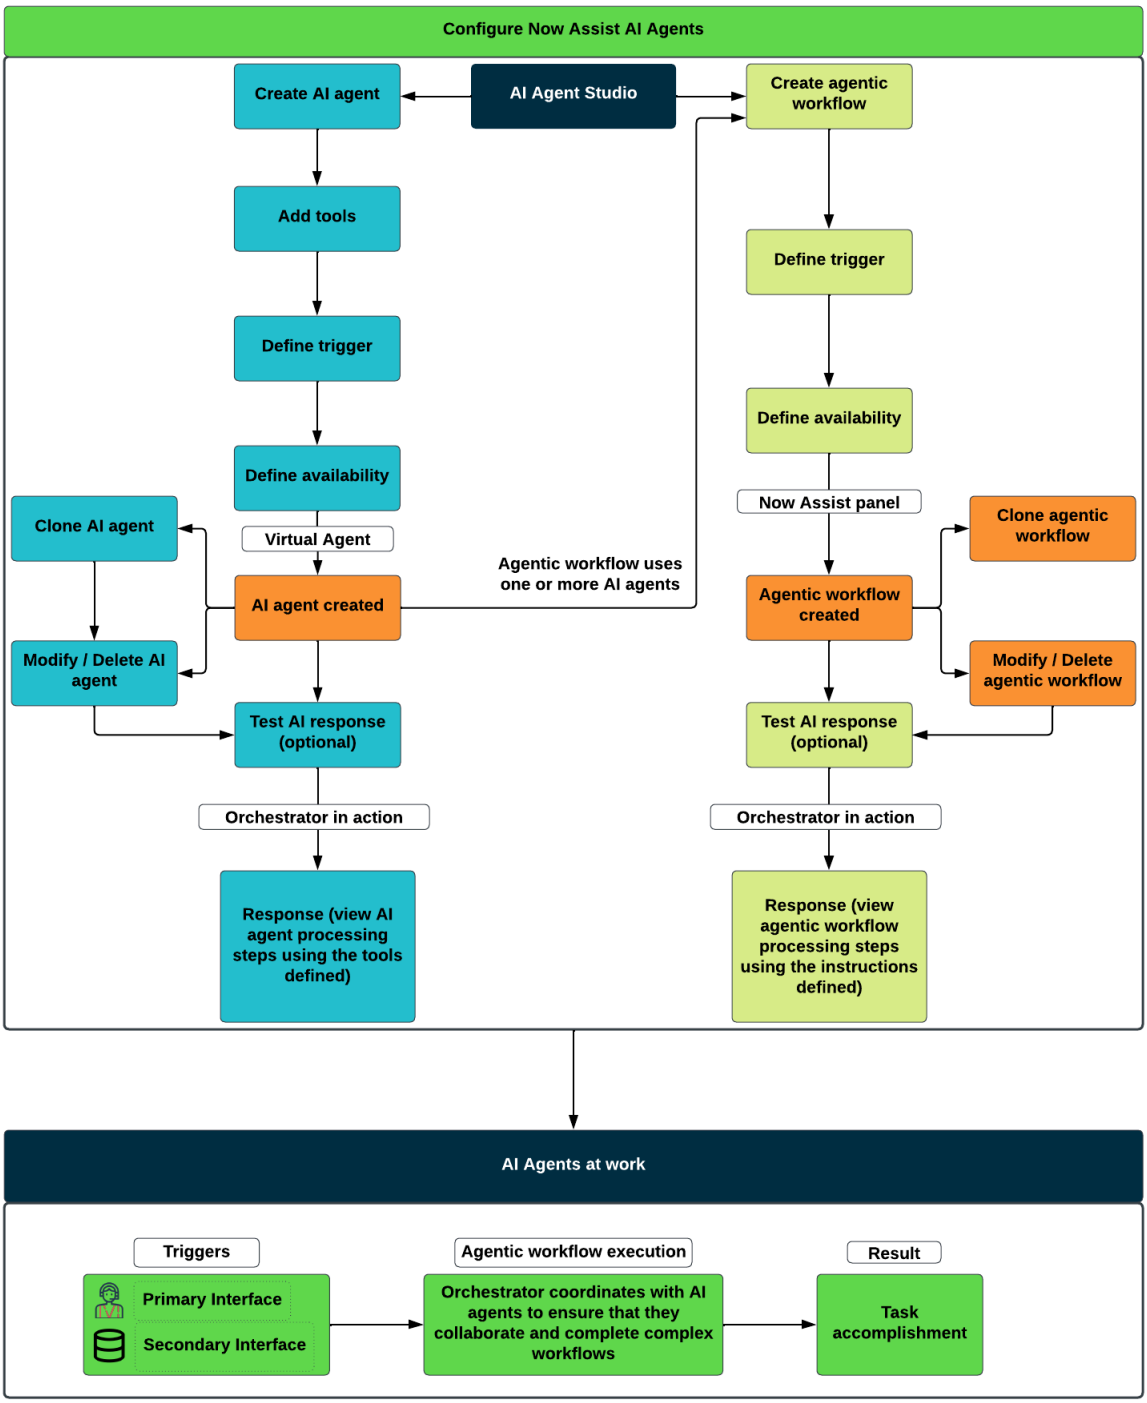
AI agents can use a variety of tools, such as web searches, record operations, and flows, to work on agentic workflows that you define.

Navigate to **All** > **AI Agent Studio** > **Overview** > **Get an overview of what to expect** and select **View steps** to see how you can successfully put your AI agents to work.

1. Create an agentic workflow: agentic workflows define the problem and situation that you want your AI agents to work on. Without an agentic workflow, your agents don't know when or what to work on.
2. Create AI agents: AI agents use tools to execute tasks. Consider the tasks needed to solve the agentic workflow you created. Create one or more agents to execute those tasks.
3. Test the agentic workflow: See how your AI agents perform. Make adjustments to either the agentic workflow or AI agents. Test until everything's working the way you want.
4. Deploy: Put your AI agents to work across your organization.

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

AI agents configuration and execution



Now Assist AI agents have two components: AI agents and agentic workflows. By using Guided Setup in AI Agent Studio, you create AI agents and workflows, add tools or define triggers, and define the availability or display location. After they're created, you can duplicate, modify, and test them in AI Agent Studio.

In run time, triggers in the primary and secondary interfaces cause agentic workflow execution, where one or more AI agents execute the plan with the help of the AI agent Orchestrator to accomplish the task. After agentic workflow execution, the output appears in the Now Assist

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

panel. To learn more about an AI agent Orchestrator is, see the Understand the Now Assist AI agents.

AI Agent Studio

Create, manage, or test AI agents and agentic workflows so that you can create self-executing workflows to help you achieve your business goals.

https://player.vimeo.com/video/1096896405?h=81ff7b3861&badge=0&autopause=0&player_id=0&app_id=58479

AI Agent Studio overview

With the AI Agent Studio application, you can create, manage, or test AI agents and agentic workflows all in one place. To enable the agentic AI experience, you must first install Now Assist AI agents. For more information, see Install Now Assist AI agents.

The Overview page has three sections where you can find the information that you must understand, begin, and continue developing AI agents and agentic workflows. When you first go to the AI Agent Studio, tour points are available to guide you through the experience.

In the Ready-made agentic workflows and AI agents section, you can find the templates and ready-made AI agents and agentic workflows that you can incorporate as-is or in custom workflows that you create. You can even explore more templates to find the ones that best fit your needs.

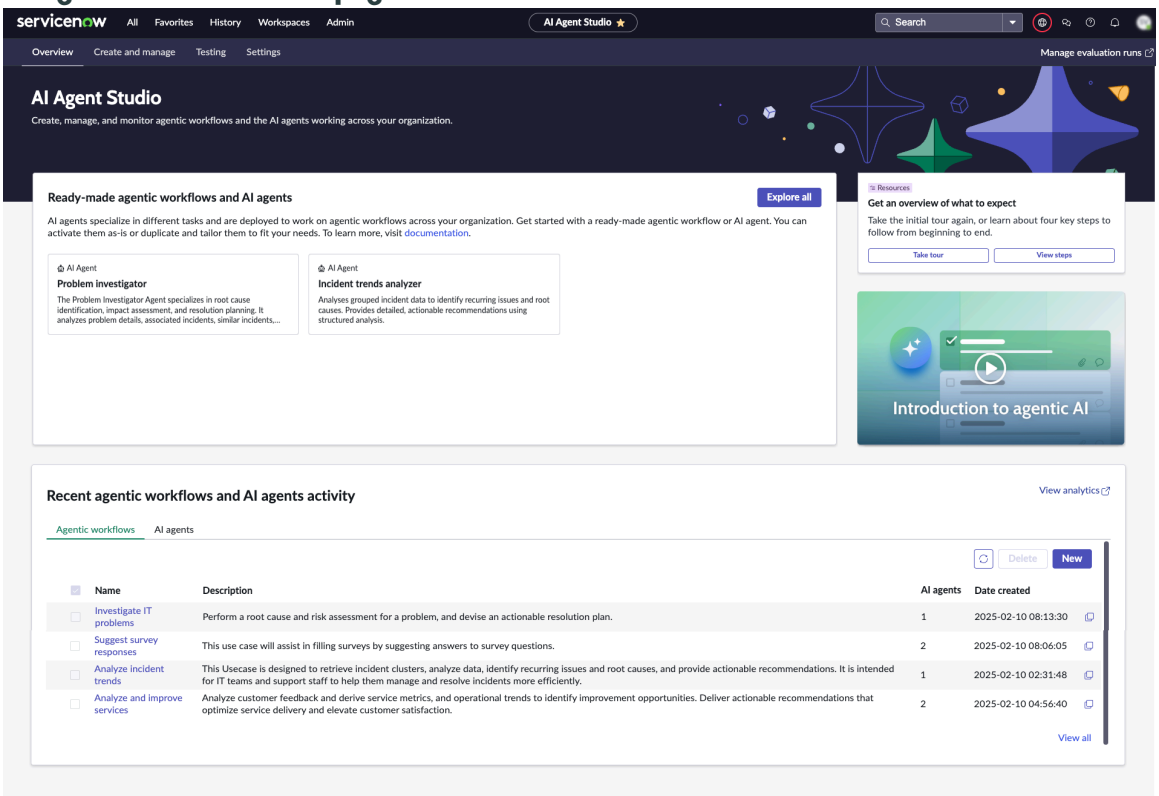
In the Recent agentic workflows and AI agents activity section, you can find the agentic workflows or AI agents that have been created or configured most recently. On both List views, you can create, duplicate, or delete agentic workflows or AI agents. The tab also includes a link to the AI Agents Dashboard where you can review the details about AI agent usage and performance.

The third section is a card to open a modal with a journey checklist that describes the steps to incorporate the AI agents into your workflows successfully and a video that provides an overview of the AI agents and workflows to get you started. Select **View steps** to open this section.

The following example shows the Overview page of AI Agent Studio.

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

AI Agent Studio overview page

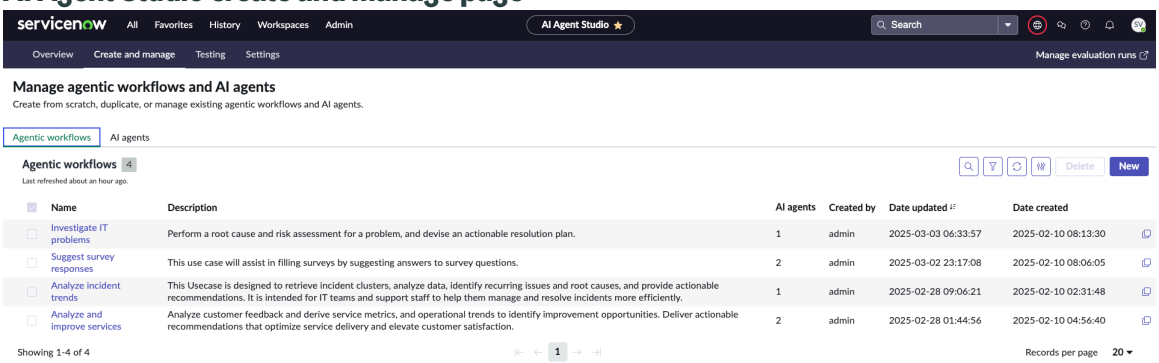


Managing agentic workflows and AI agents

From the AI Agent Studio Create and manage page, you can create, duplicate, or manage the existing agentic workflows and AI agents. This page has two tabs, one for agentic workflows and one for AI agents. You can edit the columns of the List view to change what information is displayed. You can also search or filter the Lists to find the agentic workflows and AI agents that you're looking for quickly. By selecting the name of the agentic workflow or AI agent, you can open Guided Setup to configure or reconfigure the agentic workflow or AI agent.

The following example shows the AI Agent Studio create and manage page after several Now Assist applications are installed.

AI Agent Studio create and manage page



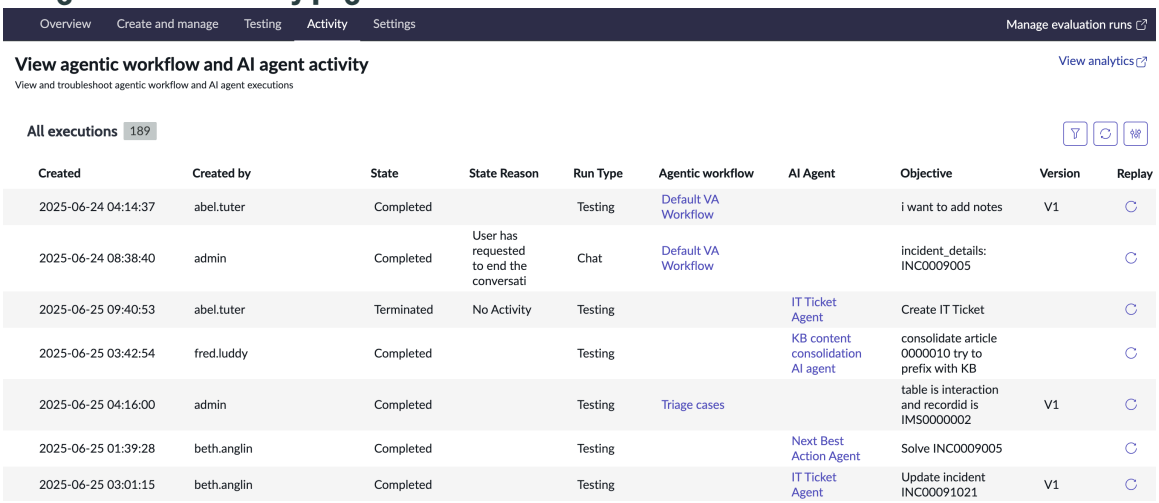
Agentic AI activity

The activity page contains execution logs for both agentic workflows and AI agents. The list allows you to filter based on various fields, including the version. For more information on creating multiple versions of the *`List of steps`* field, see Version control for AI agents and agentic workflows.

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

The following example shows several execution logs in AI Agent Studio.

AI Agent Studio activity page



Created	Created by	State	State Reason	Run Type	Agentic workflow	AI Agent	Objective	Version	Replay
2025-06-24 04:14:37	abel.tuter	Completed		Testing	Default VA Workflow		i want to add notes	V1	
2025-06-24 08:38:40	admin	Completed	User has requested to end the conversati	Chat	Default VA Workflow		incident_details: INC0009005		
2025-06-25 09:40:53	abel.tuter	Terminated	No Activity	Testing		IT Ticket Agent	Create IT Ticket		
2025-06-25 03:42:54	fred.luddy	Completed		Testing		KB content consolidation AI agent	consolidate article 0000010 try to prefix with KB		
2025-06-25 04:16:00	admin	Completed		Testing	Triage cases		table is interaction and recordid is IMS0000002	V1	
2025-06-25 01:39:28	beth.anglin	Completed		Testing		Next Best Action Agent	Solve INC0009005		
2025-06-25 03:01:15	beth.anglin	Completed		Testing		IT Ticket Agent	Update incident INC00091021	V1	

Testing agentic AI

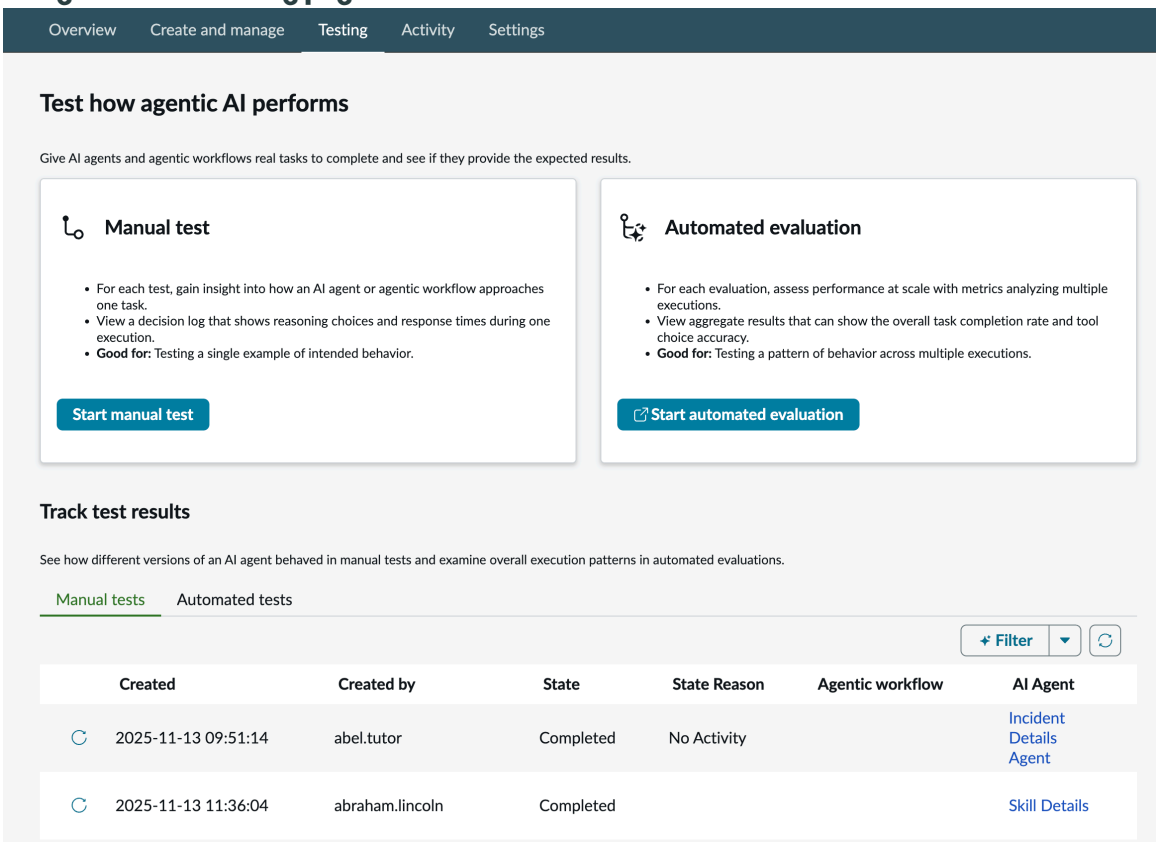
From the AI Agent Studio testing page, you can review the different tests your AI agents and agentic workflows, both manual and automated. You can test the performance of your agentic AI by simulating a single execution manually, or you can use automated agentic evaluations for testing multiple executions. Single tests are best for evaluating whether the AI agent or agentic workflow does what you expect it to. Agentic evaluations are better at finding underlying patterns and trends that may not be noticeable one execution at a time.

There are two types of manual tests you can do: *AI agent or agentic workflow* to test execution or *Test access* to test security controls. You can view your executed tests in the two tabs of the testing page. For manual execution tests, you can select the reply button to repeat the test. You can open the full details page of an automated test by selecting the test's name in the list in the *Automated* tab.

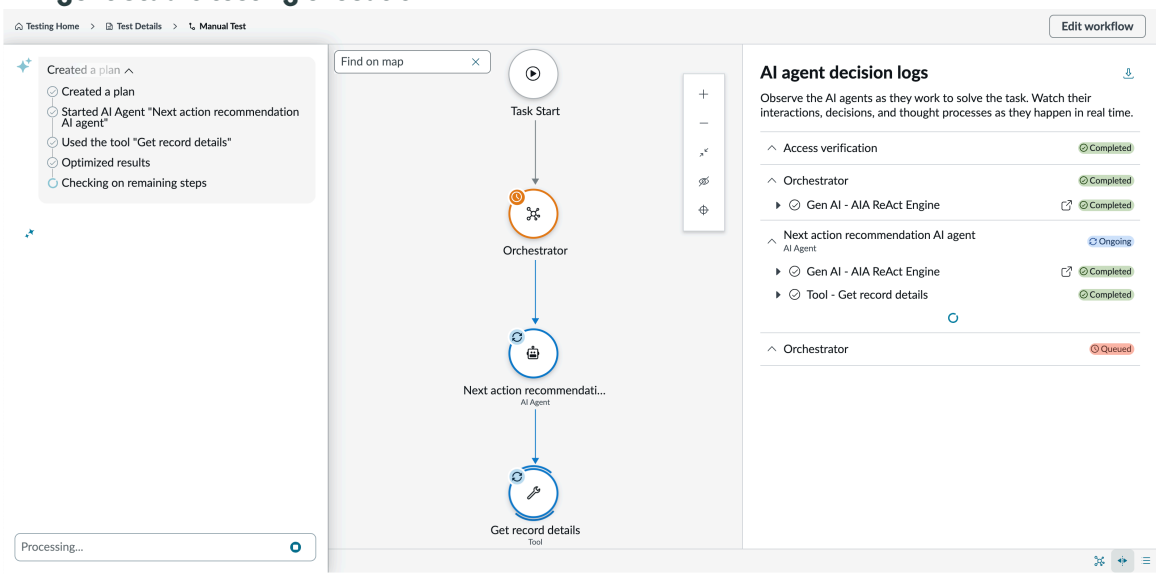
The following example shows the inputs for a Generate Resolution Plan agentic workflow testing setup of AI Agent Studio.

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

AI Agent Studio testing page



AI Agent Studio testing execution



AI Agent Studio settings

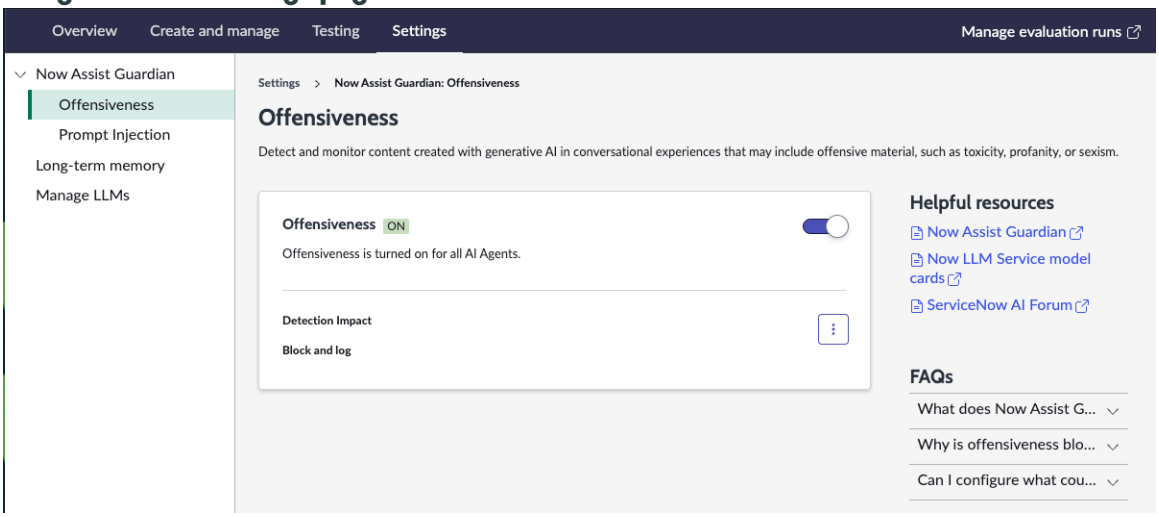
From the AI Agent Studio Settings page, you can enable Now Assist Guardian for your AI agents. By using Now Assist Guardian, you can configure:

- Offensiveness detection
- Prompt injection attempt decision
- Long-term memory for AI agents

The following image shows the AI Agent Studio settings.

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

AI Agent Studio settings page

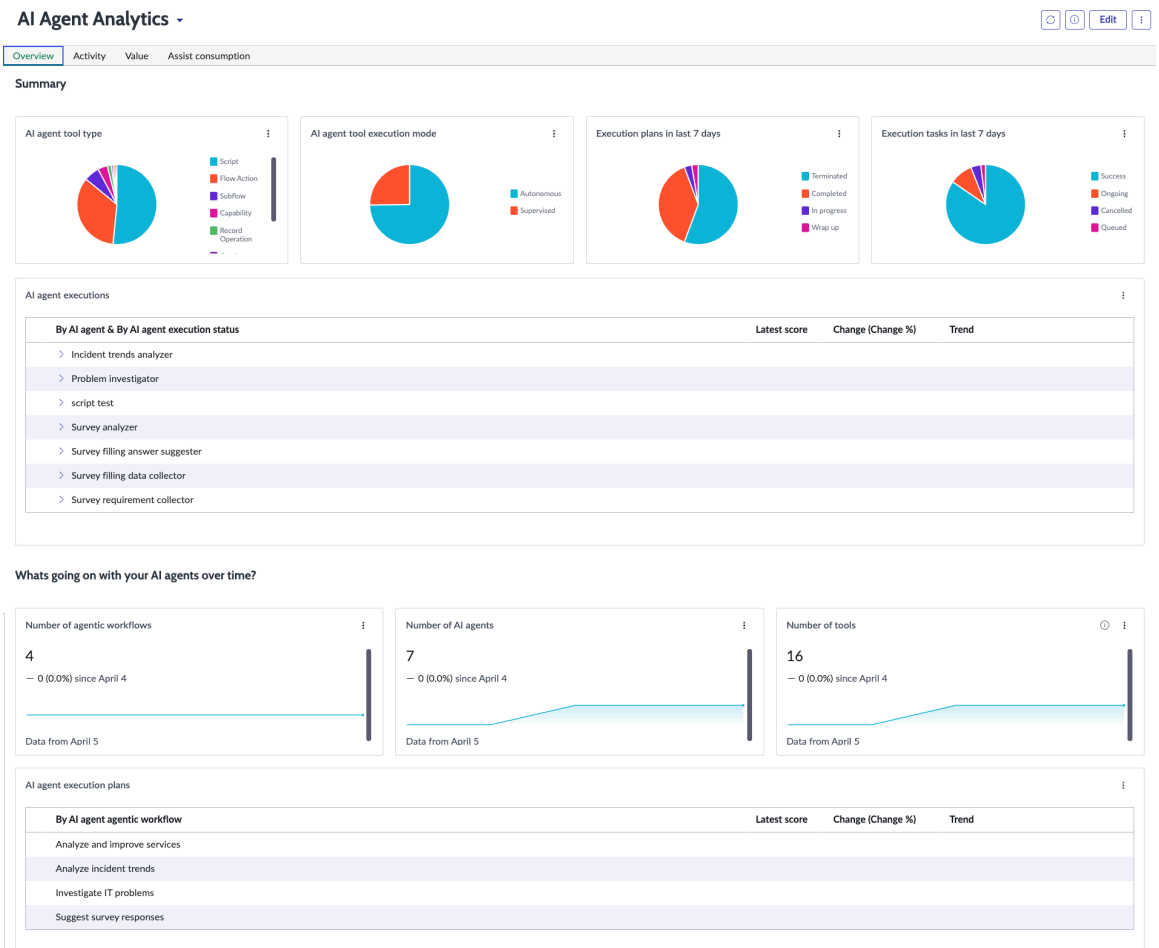


AI Agent Analytics dashboard

From the AI Agent Analytics dashboard, you can review the vital statistics about your AI agents and agentic workflows. You can see the data about the number of agents and agentic workflows used, time-to-resolution efficiency gain, and the number of executions.

The following example shows the overview page for the AI Agent Analytics dashboard.

AI Agent Analytics dashboard overview page



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

Understand the Now Assist AI agents

The Now Assist AI agents application is designed to securely leverage your data, workflows, and integrations directly within the ServiceNow AI Platform. AI agents can dynamically adjust their actions based on the progress and changing conditions of incidents or cases to help ensure that they stay focused on achieving their objectives.

Now Assist AI agents Orchestrator

The AI Agent Studio guides users in leveraging AI agents effectively, while AI Orchestrator ensures seamless collaboration among them. The AI agent Orchestrator is a central management system that coordinates AI agents to ensure that they collaborate effectively to complete complex workflows. It can also retrieve missing context from users when needed. AI agents operate iteratively, seeking assistance from the Orchestrator if they encounter difficulties. The Orchestrator enhances inter-agent communication and centralizes coordination. The Orchestrator allows AI agents to share information effectively and transition tasks seamlessly, regardless of where the process originates. As a result, the Orchestrator is essential for managing intricate workflows, facilitating real-time adaptability, and efficient task delegation among agents.

The role of Orchestrator in AI Agent Studio:

1. Task Coordination & Execution: manages the interaction between multiple **AI agents**, ensuring that tasks are executed in the correct sequence.
2. Dynamic Decision-Making: Can invoke AI models dynamically based on the scenario.
3. Seamless Integration with Enterprise Systems: Works with ServiceNow workflows, CMDB, and third-party APIs to enable end-to-end automation.
4. Handling Multi-Step Workflows: Enables chained execution of AI-driven tasks among multiple AI Agents.
5. Policy & Governance Enforcement: Enforces governance rules to ensure that AI agents adhere to compliance and security standards.

The Orchestrator works with ReAct prompt, where the previous reflection prompt has been merged with the planning prompt for a streamlined execution experience improving the efficiency and visibility of planner actions. With the ReAct prompt, a scheduling mode called Route has been introduced that enables asynchronous, parallel tool execution within agents to reduce latency and improve performance when an agent needs assistance from another agent during the execution process.

Dynamic Orchestrator for agentic workflows

The dynamic orchestrator helps map the right agents to the user's agentic workflow to provide better performance and accuracy. When there are more than the ideal 8 to 10 AI agents available for an agentic workflow, then the dynamic orchestrator helps identify the agents that need to be considered for executing that agentic workflow in the planning phase.

User impersonation in Now Assist AI agents

The agentic workflow executes tools as the logged-in user in the Now Assist panel. Any operations that are performed within the tool in this flow are also executed as the logged-in user. After impersonation is enabled, testing an AI agent uses the instance-level impersonation.

Administrators can see logs with individual AI agent names as a record of who approved the agentic action in an agentic workflow. The logs help when determining a point of contact if there’s an issue with the approved agentic action by the AI agent.

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

With impersonation, the fulfiller in the Now Assist panel and the requester in Virtual Agent can see the transactions recorded in the name of the AI agent that performed the agentic workflow execution.

Note: Virtual Agent can be turned on only for AI agents at the Define availability section in the AI agent guided setup. For more information, see Create an AI agent.

Best practices for AI agents configuration

AI agents can dynamically adjust the actions that are based on the progress and changing conditions of incidents or cases to help ensure that they stay focused on achieving their objectives.

AI agents can dynamically adjust the actions that are based on the progress and changing conditions of incidents or cases to help ensure that they stay focused on achieving their objectives.

- You can create, duplicate, and modify AI agents and agentic workflows. The default AI agents and agentic workflows are available in read only mode and to utilize them, you are required to duplicate these resources.
  - For more information about duplicating an AI agent, see Duplicate an AI agent.
  - For more information about duplicating an agentic workflow, see Duplicate an agentic workflow.
- You can also modify the duplicated AI agent and agentic workflow to meet your requirements.
  - For more information about modifying an AI agent, see Modify an AI agent.
  - For more information about modifying an agentic workflow, see Modify an agentic workflow.

Note: The agentic workflows available by default in the system are inactive and you must activate them to utilize. For more information about activating an agentic workflow, see agentic workflow template.

Now Assist AI agents capabilities

Lists the features supported by Now Assist AI agents, including configuration options and functional capabilities.

Invoke Conversations with AI Agent Background Channel

The AI Agent Background Channel helps you to invoke AI Agent or agentic workflow execution from the Workspace. Use the AI Agent Background Channel associated with the AI Agent Background Provider to invoke conversations. The AI Agent Background Provider is based on the Custom Adapter Framework from Virtual Agent. For more information, see Configure a provider for your custom chat integration.

Create a channel identifier in the Provider Channel Identities table [sys_cs_provider_application] to add any additional conversational capabilities to your own provider application and get a new inbound ID that allows for customization. For more information, see Create a channel identifier for your custom chat integration.

To start a conversation, trigger the flow using the `sn_aia.AiAgentRunttimeUtil().startAiAgentConversation(request)` API in the Script Include (sys_script_include) of the AIAgentBackgroundProvider and select **Run Script**. When the Script execution status indicates Success, the conversation begins in the order of utterances defined in the Script.

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

Conversations that are invoked for executing an AI agent are logged in the Execution Plans [sn_aia_execution_plan] table. Open the conversation record to confirm the device type as **AI Agent Background**. Open the execution record to see the **Execution Tasks**, **Messages**, and the **Tools Executions** used to execute the AI agent.

You can also see the entire execution steps on the AI Agent Studio Testing page by copying the execution plan record's [sys_Id] and testing it. On the **Chat responses** tab, in the AI agent decision logs, you can see the AI agent details and the tools it used to resolve the issue.

Interactive and Non-interactive AI agents

The Interactive AI agents reach out to users for information when there is a fallback in the execution process, and the AI agent re-triggers the flow.

The Non-interactive AI agents do not reach out to the user at any fallback stage in the execution process. When the AI agent needs user information, it takes the dynamic prompt approach using the ReAct layer, where the prompt of the ReAct will change based on the execution mode of the AI agent or agentic workflow. Therefore, in the Non-interactive execution, the reach fallback options do not have to collect input from a user as a fallback option. However, the output of the AI agent or agentic workflow will still need to be presented to the user, and in any execution failure scenario, a message in the Now Assist panel or Virtual Agent is shown.

To implement the Non-interactive execution, the **Execution Mode** field is added in the Execution Plans [sn_aia_execution_plan] table, where the execution mode can be **Interactive** or **Non Interactive** based on the given runtime parameter.

You can run the AI agents and agentic workflows concurrently in the AI Agent Background Channel and in Non-interactive mode where the background execution allows AI agents to operate with any chat panel like Now Assist panel or Virtual Agent.

Multilingual support

You can leverage multilingual support for AI agents across languages for better translation quality to:

- Tune system prompts for native translations.
- Implement dynamic translation strategies when native support is unavailable.
- Provide extensive testing via automated and manual evaluations.

General guidelines for creating AI agents and agentic workflows

By following some general guidelines for creating AI agents and agentic workflows, you can create clear and effective instructions that help maximize their efficiency and effectiveness.

https://player.vimeo.com/video/1098393879?h=d37199a090&badge=0&autopause=0&player_id=0&app_id=58479

Creating AI agents and agentic workflows overview

AI agents and agentic workflows rely on large language models (LLMs), so the language that you use for their instructions is important. There are many steps that you can take to help improve the quality and consistency of your agentic AI solutions.

Follow these guidelines when writing AI instructions for an AI agent or agentic workflow.

Clarity

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

- Use instructions that are concise, clear, and precise. Simple language can help remove ambiguous situations that could stall or disrupt an agent's progress. Example: "You're an assistant tasked with helping the caller or requester by suggesting answers to survey questions." instead of "You help callers with questions."
- Use specific action verbs. Example: "Analyze" instead of "look at."
- Add explicit conditions when a certain step must be completed. Example: "If priority = High, then escalate immediately" instead of "escalate urgent issues."
- Clarify boundaries and limitations. Example: "NEVER modify incident status without supervisor approval."
- Limit technical jargon. Technical jargon can limit applicability because it may not be accessible or universal.

Context

- Embed requirements within context. Specify when certain requirements should be used within the task. Example: "When generating answers to present to the user, apply standard quality controls."
- Only include the context that affects decision making. Avoid extraneous information to help prevent the AI agent or agentic workflow from incorporating unwanted details.
- Define what good outcomes look like. Include examples. Missing or vague descriptions of results could cause AI agents or agentic workflows to exit before the end state is reached. Example: "Present the user with a report that includes a minimum of 3 relevant graphs pertaining to the list of records."

Constraints

- Denote hard requirements. Use strong language to insist on the importance of the requirements. While output may not always follow these requirements, being specific about standards can help reduce deviations.
- Describe preferences that can be adjusted based on context.
- Add conditional restraints. Example: "If more than one record is returned, present the results to the user in order of creation date. If more than five records are returned, present the results to the user in order of relevance."

Coherence

- Create steps that flow logically. Each step should build on the results of the previous step. Example: "Step 1 gathers a list of incident records. Step 2 performs record operations on those list of records."
- Use consistent terminology throughout.
- Design instructions that serve the overall objective. All of your tools and AI agents work together to solve problems.
- Break complex goals into individual pieces. Example: "Step 1: Analyze customer data systematically. Step 2: Check priority and validation permissions. Step 3: Generate and format comprehensive report."

Structure

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

- Use consistent formatting. Similar instructions should use similar language patterns.
- Group related steps together.
- Create a smooth flow from instruction to instruction. Along with grouped instructions and logical step progression, this provides a systematic and structured process for the AI agent or agentic workflow to follow.

Guidelines for creating AI agents

The description, AI agent role, and list of steps give the LLM the context and instructions to perform a task. Together, they form the blueprint that's necessary for the LLM to complete its role in a complex workflow. Follow these guidelines to improve the accuracy, adaptability, and optimization of the AI agent:

AI agent description

- Specify the key areas or tasks that you want the agent to handle. Example: "Specializes in handling inquiries and resolving customer issues."
- Use clear, focused language and avoid vague terminology.
- Define the agent's inputs, outputs, and context.
- Differentiate the agent's unique role from other agents. Provide distinct and detailed descriptions of what that specific agent should do that is different than other agents.

AI agent role

- Clearly state the primary function of the AI agent in one or two sentences. Example: "The AI agent acts as a customer service assistant."
- State the specific business challenge that you want the AI agent to address. Example: "Reducing customer wait times by 50%."
- Provide a brief scenario of how the AI agent is to be used. Example: "Automating responses to common queries and escalating complex issues to human agents."

AI agent list of steps

- Create a logical progression from step to step.
- Use action-oriented language. Use verbs like the following:
  - Fetch
  - Retrieve
  - Filter
  - Analyze
  - Extract
  - Parse
  - Update
  - Send
  - Notify
  - Generate
  - Validate

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

- Include outputs for each step to measure the task completion. Example: "Return the total incident count to other agents or processes that need it."
- Add contingencies to account for unexpected scenarios. Example: "If you encounter an error while looking up records, try again. If you still get an error, report that an error occurred."
- Define success and completion. Provide end states so that the AI Agent Orchestrator can determine whether an AI agent has finished its objective.
- If you have created a tool for your AI agent, refer to the tool by name. However, verify that the instructions change if the tool is ever renamed. If you don't, then the AI agent might not complete its task.

Tools

- Write detailed tool descriptions to help the AI agent understand what the tool is and how to use it.
- Create tools that work together. Build your tools around achieving the list of steps for the AI agent itself.
- Use the output transformation strategy field to define what the output of a tool should look like. Specifying how a tool should present its output can help the AI agent use that output when employing other tools or when sharing information between other agents.

Guidelines for creating agentic workflows

Follow these agentic workflow guidelines to provide the detailed information and the steps needed to accomplish a task:

Agentic workflow list of steps

- Write each step sequentially and number each step so that there's a logical and actionable flow.
- Account for as many possibilities as you can to avoid gaps. Provide enough detail that the agentic workflow can adapt if it encounters an edge case.
- Define the starting conditions, actions, decision points, and end states.
- Refer to end users as "the user."
- Use verbs like "show," "display," or "inform" to describe steps where the user needs to be shown something but they do not need to provide input.

Additional tips

Associate a maximum of 10 agents per agentic workflow. Adding more agents might not provide better or faster results. Instead, use small, well-defined scopes.

> **Note:** You can assign up to 100 agents to an agentic workflow. However, not all 100 agents may be involved in resolving the agentic workflow. The AI Agent Orchestrator decides which agents should execute the plan. For example, an incident resolution gets impacted by the size of the agentic workflow and the number of agents that are assigned to it.

When creating agentic workflows with more than one assigned agent, make sure that you clearly define the agents with non-overlapping responsibilities. Include explicit limitations in the agents' roles. For example, one agent could handle user record details while another agent handles the incident record details.

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

Common issues

Problem: Agent consistently uses the wrong tool

Potential solutions:

- Create clear, specific instructions with action keywords.
- Write detailed tool descriptions that specify when and why to use each tool.
- Include explicit input and output guidance in tool descriptions.
- Evaluate tool selection logic by testing against various scenarios.

Problem: Inconsistent output quality

Potential solutions:

- Embed standards as actionable requirements.
- Use phrases such as "implementing professional standards."
- Add quality verification steps.

Problem: Quality checks or other steps skipped

Potential solutions:

- Convert checklists to actionable steps.
- Use phrases such as "analyze completion against criteria."
- Add conditional progression logic.

Additional resources

For more prompting recommendations, see the Community Now Assist AI Agents prompting guide.

Example AI agent

Use the example AI agent with clear name, description, AI agent role, and list of steps fields to use as a guide when creating your own AI agents.

Name

Categorize ITSM incident AI agent

Description

Categorize ITSM incident AI agent assigns the appropriate category, subcategory, and configuration item (CI) to an incident.

AI agent role

You're an expert in analyzing incidents and are tasked with assigning an appropriate category, subcategory, and configuration item (CI).

List of steps

Objective: Incidents represent disruptions, unavailability of a service, or issues faced by a user. The agent's task is to understand the issue and assign the appropriate:

- Category
- Subcategory
- Configuration Item (CI)

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

Step 1: Retrieve Incident Details.

- Fetch the details of the given incident.
- DO NOT PROCEED if incident details are missing.
- If no details are found, verify the correctness of the incident number.
- Only proceed when incident details are successfully retrieved.

Step 2: Assign Category, Subcategory, and CI.

Step 2.1: Determine the Category.

- Fetch the available categories.
- Recommend a category only if the incident details contain clear, specific, and meaningful information that directly corresponds to one of the available categories.
- If the details are too vague, generic, lack context, or do not clearly indicate a relevant category, then set the category and subcategory to "Not determined" (field_value = null). Do not edit the subcategory later.
- Do not infer or guess a category based on common or placeholder terms (for example, "test", "check", "issue") unless they are accompanied by enough context to confidently justify a match.
- You must be 100% certain that the incident clearly belongs to a category before assigning it.

Step 2.2: Determine the Subcategory.

- Fetch the list of available subcategories for the chosen category.
- Choose the most appropriate subcategory based on the incident details.
- If no suitable subcategory can be determined or if the category is not determined, set subcategory to: "Not determined" (field_value = null).

Step 2.3: Determine the Configuration Item (CI).

- Fetch the list of CIs assigned to the caller of the incident.
- Pick the most relevant CI based on
  - Name
  - Manufacturer
  - Asset Name
- If no relevant CI is found or the list is empty, set CI to: "Not determined" (field_value = null).
- If multiple options exist, select the most relevant one.

Step 3: Present the Recommendation.

- PRESENT the recommended values to the user in the following bullet-point format:

  Recommended incident ({incident number}) categorization details:

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

- Category: {recommended category}, (Reason: {reason for recommendation})
- Subcategory: {recommended subcategory}, (Reason: {reason for recommendation})
- Caller's CI: {recommended CI}, (Reason: {reason for recommendation})

- Ensure this information is explicitly displayed before proceeding to the next step.

Step 4: Update the Incident. Update the incident record with the selected Category, Subcategory, and CI values.

Key Considerations:

- NEVER ask the user for input.
- Never proceed without retrieving incident details.
- If NO appropriate value is found, assign "Not determined" (field_value = null).
- Do NOT ask the user to manually determine category, subcategory, or CI.
- Always present recommendations before updating the incident.
- Ensure that the recommendations are logical and relevant to the incident details.

Example agentic workflow

Use the example agentic workflow with clear name, description, and list of steps fields to use as a guide when creating your own agentic workflows.

Name

Investigate incident

Description

This workflow streamlines incident management by analyzing incidents, linking them to existing problems, and identifying relevant catalog items. It leverages three specialized AI agents: the ITSM Incident Resolution Investigation AI Agent retrieves incident details, identifies similar incidents, and fetches relevant knowledge articles; the Link Incident to Problem AI Agent searches for and links incidents to existing problems based on contextual similarity; and the Find Catalog Item AI Agent locates relevant catalog items and integrates their details into incident records or provides direct links. This structured orchestration ensures efficient incident resolution, reduces redundancy, and enhances decision-making by leveraging historical data, organizational knowledge, and catalog resources, ultimately improving enterprise incident management effectiveness.

List of steps

Step 1: Incident Analysis:

Analyze incident details and retrieve relevant historical data and knowledge articles.

Step 1.1: Retrieve Incident Context:

- Use the ITSM Incident Resolution Investigation AI Agent to fetch incident details (short description, description, caller, state).
- If retrieval fails, inform the user: "Unable to retrieve incident details. Please verify the incident number." and terminate workflow.

Step 1.2: Identify Similar Incidents and Knowledge Articles:

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

- Use the ITSM Incident Resolution Investigation AI Agent to find similar incidents and relevant knowledge articles.
- Confirm with the user before adding resolutions to incident comments.

Step 1.3: Prepare Incident Documentation:

Update incident additional comments with approved resolutions and knowledge article references, preparing clear documentation for next phase.

Step 2: Problem Linking:

Link the incident to an existing relevant problem, if applicable.

Step 2.1: Receive Incident Documentation:

Use Link Incident to Problem AI Agent to validate incident details and confirm no existing problem linkage.

Step 2.2: Identify Relevant Problems:

- Execute search for relevant problems using incident short description.
- If no relevant problems found, inform the user: "No relevant problems found to link." and proceed to Catalog Item Identification.

Step 2.3: Link Incident to Problem:

- Present most similar problem details to the user for approval.
- Upon user approval, link incident to problem using "Set problem to incident" tool; confirm successful linkage to the user.

Step 3: Catalog Item Identification:

Identify and integrate relevant catalog items into the incident record.

Step 3.1: Receive Incident and Problem Context:

Use Find Catalog Item AI Agent to validate incident context and prepare for catalog item search.

Step 3.2: Execute Catalog Item Search:

- Perform targeted search for relevant catalog items based on incident description.
- If no catalog items found, inform the user: "No relevant catalog items identified." and proceed to final delivery.

Step 3.3: Integrate Catalog Item Details:

- Present catalog item choices to the user; upon selection, either add catalog item details to incident additional comments or provide direct catalog item link.
- Confirm successful integration or delivery to the user and complete workflow.

Group Action Framework

Group Action Framework (GAF) is an intelligence feature on the ServiceNow AI Platform that groups related records and applies actions to them using LLMs.

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

GAF overview

GAF is composed of two processes. The grouping process identifies clusters of similar records (incident, cases, KB articles, etc.) and selects a set of representative records for the cluster. The actioning process maps new records to clusters and executes large language model (LLM) instructions for those clusters to achieve certain tasks, such as summarizing the contents of a cluster or generating resolution steps based on steps that worked for records in that cluster. GAF processes together benefit your AI agents and Now Assist generative features in multiple ways.

- Improves consistency and quality of agentic and generative AI features by using the best examples from groups of records.
- Reduces the cost of LLM calls by only executing on the representative records.
- Scales to accommodate large amounts of data because selected records can represent any size cluster.

Skills used in GAF

Multiple skills are involved in GAF setup and execution. They are modular, so not all executions will use all skills, but they can be used together in tandem. These skills are used exclusively for GAF and currently cannot be included on their own in custom agentic workflows.

Grouping skill

Clusters related records using machine learning techniques.

Topic-labeling skill

Adds human-readable names to the clusters using an LLM to make the clusters easier to identify.

Action strategy skill

Selects representative records from each cluster for the mapper and reducer skills to use.

Action mapper skill

Runs LLM inference calls for the selected representative records, producing a record summary for the selected records.

Action reducer skill

Uses the generated summaries created by the mapper skill to produce a single summary for the entire cluster.

GAF and Now Assist in AI Search

GAF uses AI Search to improve its effectiveness and can use it as a fallback option in case GAF does not return any results. GAF can work without Now Assist in AI Search, but if it is enabled then GAF has optimized prediction. The optimized prediction feature increases clustering capacity up to 500,000 records and improves recall speed.

See Install Now Assist in AI Search and Set up AI Search for Group Action Framework for more information on configuring AI Search for GAF.

Version control for AI agents and agentic workflows

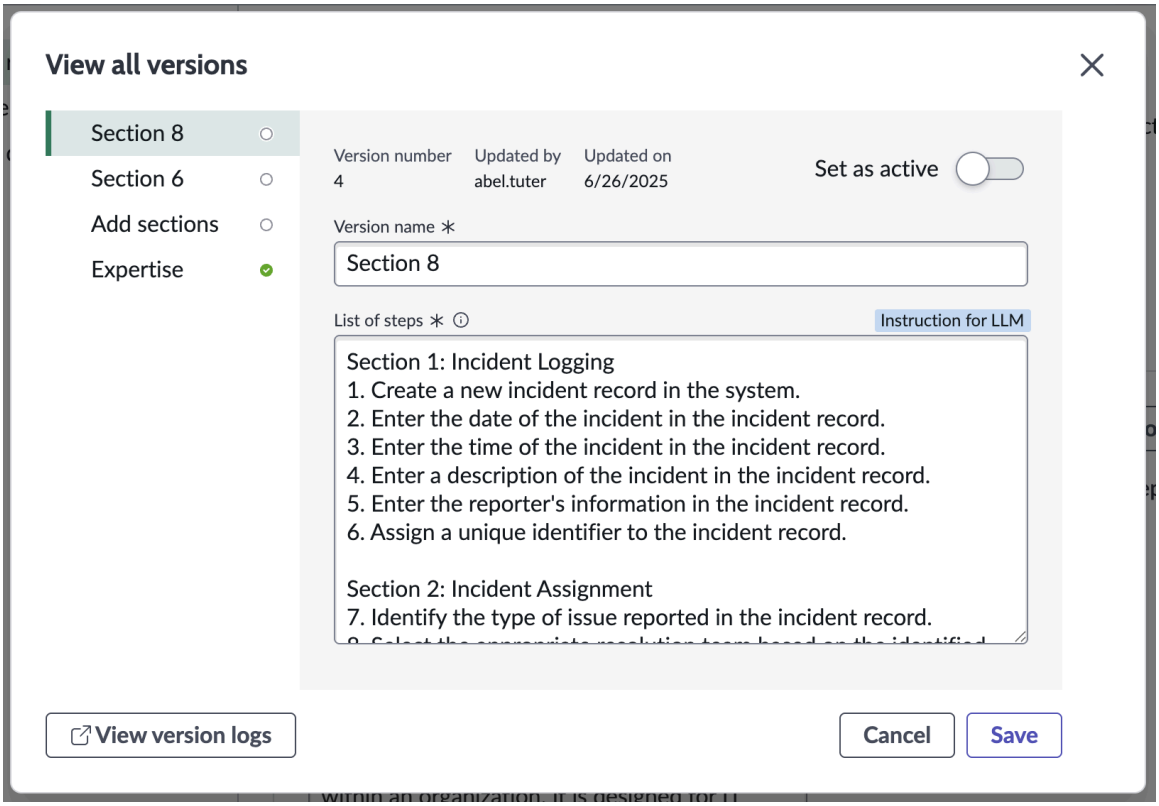
Version control enables you to track changes made to instructions sent to the LLM for AI agents and agentic workflows.

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

Version control overview

When creating an AI agent or agentic workflow, the *List of steps* field is important because it provides the context necessary for the large language model (LLM) to accomplish tasks. You can use version control to create multiple versions of the *List of steps* field in the guided setup. Tracking your changes with version control helps enable you to experiment and test new instructions. You can create versions, edit and refine versions, and revert to previous versions as needed.

Version control modal



Creating versions

You can create a version by selecting **View versions** > **New version** in the guided setup of an AI agent or agentic workflow. Doing so brings up the *Create a new version* modal where you can name the version and make changes. Naming your version enables you to identify what changes were made to make it easier to track.

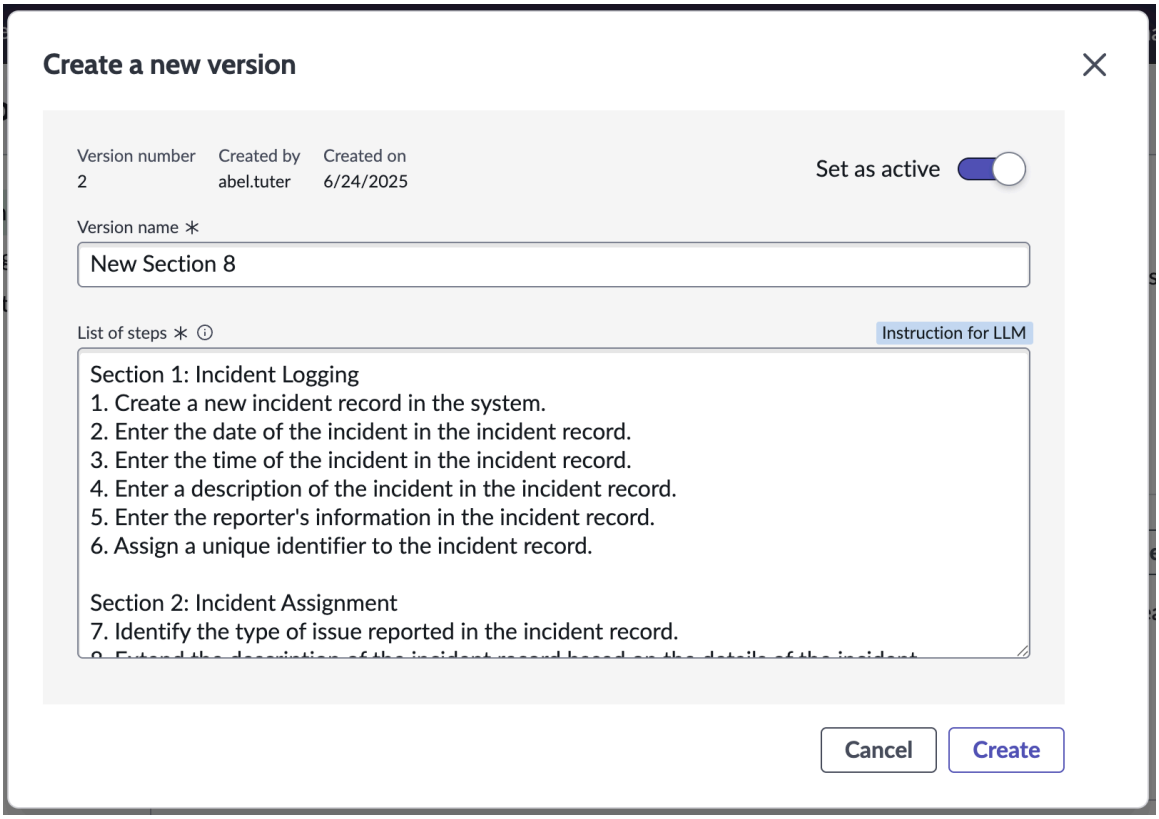
You can choose to make the new version active immediately by selecting the **Set as active** toggle.

For more information about writing effective descriptions and instructions, see General guidelines for creating AI agents and agentic workflows.

When creating a new version, you also have the option to use Now Assist to refine your current *List of steps*. You must already have an active version for this option to be available. Selecting **View versions** > **Refine version** brings up the modal where Now Assist generates and displays a new *List of steps* for you to review. You can change any of the generated text, or you can use it as it was generated. Select **Create** to use the new, refined version.

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

Create a new version modal



Changing between versions

When you have multiple versions of the same *`List of steps`* field, you can switch which one is active at any time. Choosing to make a new version active enables you to test it and evaluate its effectiveness.

If you're updating an agentic workflow, you can execute an agentic evaluation run to measure performance.

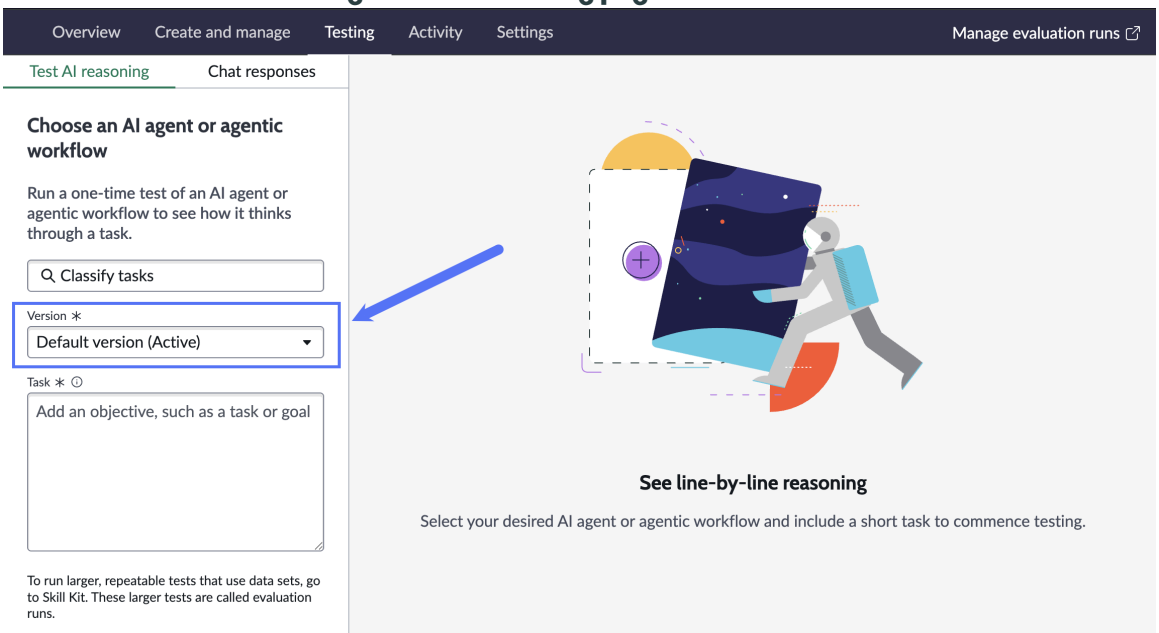
You can decide which version is active by selecting **View versions** in the guided setup. Toggle the **Set as active** field and save your changes to make the new version active.

Testing different versions

You can test different versions of an agentic workflow or AI agent from the AI Agent Studio **Testing** page. The **Version** selector dropdown displays both active and inactive versions with the currently active version marked as *`(Active)`*. You can also start evaluation runs for specific versions by selecting **Set up evaluation run**.

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

Version selector in the AI Agent Studio Testing page

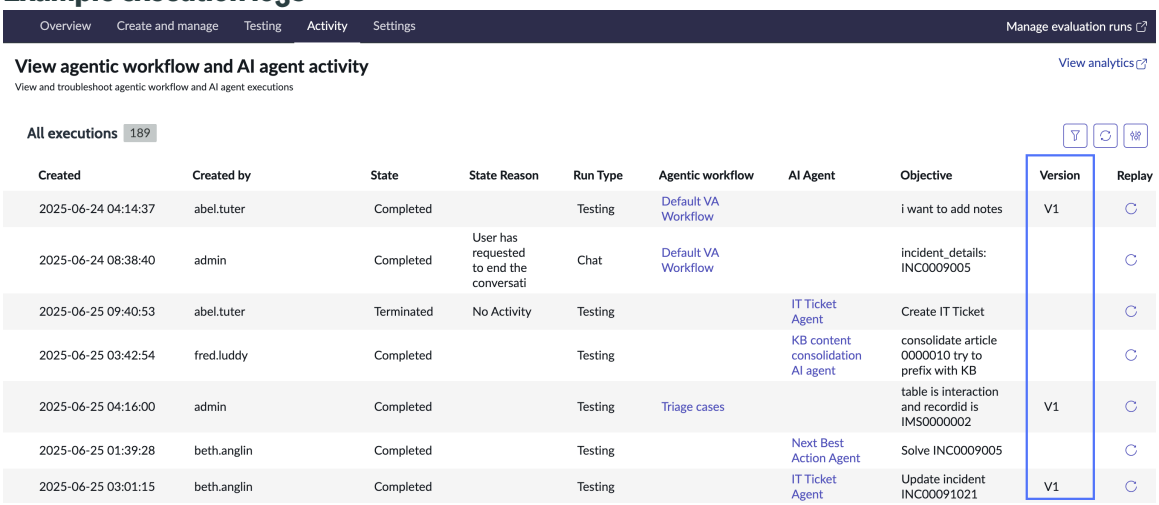


Version logs

The AI Agent Studio **Activity** page enables you to track AI agent and agentic workflow executions by version. When in the guided setup, you can select **View version logs** in the *`View all versions`* modal. Choosing to view the version logs opens the **Activity** tab in a new window with the list filtered for the version you selected.

When viewing the **Activity** tab at any time, you can choose to filter by version. The more specific your version names, the easier it is to identify those changes in the list.

Example execution logs



Created	Created by	State	State Reason	Run Type	Agentic workflow	AI Agent	Objective	Version	Replay
2025-06-24 04:14:37	abel.tuter	Completed		Testing	Default VA Workflow		i want to add notes	V1	
2025-06-24 08:38:40	admin	Completed	User has requested to end the conversati	Chat	Default VA Workflow		incident_details: INC0009005		
2025-06-25 09:40:53	abel.tuter	Terminated	No Activity	Testing		IT Ticket Agent	Create IT Ticket		
2025-06-25 03:42:54	fred.luddy	Completed		Testing		KB content consolidation AI agent	consolidate article 0000010 try to prefix with KB		
2025-06-25 04:16:00	admin	Completed		Testing	Triage cases		table is interaction and recordid is IMS0000002	V1	
2025-06-25 01:39:28	beth.anglin	Completed		Testing		Next Best Action Agent	Solve INC0009005		
2025-06-25 03:01:15	beth.anglin	Completed		Testing		IT Ticket Agent	Update incident INC00091021	V1	

Implement access control in Now Assist AI agents

Implement security controls for AI agents and agentic workflows through access control lists (ACLs) and user identities to increase alignment with the access control-based security measures in the agentic system.

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

Security for AI agents overview

Access controls for agentic AI on the ServiceNow AI Platform comprises two major components: Access control lists (ACLs) and user identities. The interaction between these two components at the agentic workflow, AI agent, and tool levels within the AI Agent Studio influences their overall security and functionality.

Access control lists

The access control lists (ACLs) in Now Assist AI agents determine users who can invoke an agentic workflow or an AI agent. ACLs must be configured individually for each agentic workflow, AI agent, and certain AI agent tools.

The ACLs added to an AI agent and agentic workflow are available in the respective related lists for reference.

Important: ACLs configured in AI Agent Studio only determine the roles required for users to invoke an agentic workflow or an AI agent. They don't determine the access that the agentic workflow or an AI agent has once it's invoked.

User identity

The user identity determines the roles that the AI agent or an agentic workflow operates with and the data it can access depending on the permissions assigned to the .

After configuring the access control lists (ACLs), you must configure the User identity (also called as *Run as*) using which the AI agent or agentic workflow gets executed. There are two possible user configurations to select from:

- **Dynamic user**: The logged-in user who invokes the execution of an AI agent or an agentic workflow. Dynamic user is the default user identity, and you can use the dynamic user unless there's a specific need that justifies an AI user.
- **AI user**: A dedicated user who executes the AI agent or an agentic workflow with assigned roles that remain consistent regardless of who or how the execution is invoked. For example, an AI agent or an agentic workflow may need to be run with elevated privileges that the dynamic user might not have.

If you do not have an AI user but want to use the *AI user* identity, you need to create a new record on the User table. See Create a user. Select *AI user* as the identity type.

Note:

- An AI user can be configured as part of user identity and user identities are configured at the agentic workflow and AI agent levels.
  - For more information about user identity in an AI agent, refer to Create an AI agent.
  - For more information about user identity in an agentic workflow, refer to Create an agentic workflow.
- The ACLs are checked with the actual conversational user, a user who has invoked the agentic workflow or the AI agent. Once the ACL check is complete, the user identities can be applied.

Configure ACLs in AI Agent Studio

ACLs configured in the AI Agent Studio for AI agents and agentic workflows are role-based and are of two types:

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

- **Allow-If**: Grants access to data or resources when all the specified conditions in the ACL are met and the ACL doesn't prevent other ACLs from granting access to the same resource even if it doesn't.
- **Deny-Unless**: Grants access only when the roles meet a specified condition and no other ACLs can override or grant access to that resource.

There are three possible options for ACLs created in AI Agent Studio:

- *`Any authenticated user`*: Grants access to any user who is authenticated on the instance, regardless of the role.
- *`Users with specified roles`*: The default ACL option that requires you to select the roles to invoke an AI agent or an agentic workflow.
- *`Public`*: Grants access to all users, including guests who aren't signed in.

Each AI agent and agentic workflow must have its own unique ACL.

- To configure an ACL in the AI Agent Studio for an AI agent, see the Create an AI agent guided setup.
- To configure an ACL for an agentic workflow, see the Create an agentic workflow guided setup.

Note: If there are conflicting security requirements between agentic workflows, AI agents, and AI agent tools, or if the invoking user meets the criteria for some ACLs but not others, your agentic AI fails to execute. When configuring these security settings, consider all aspects of the agentic system- including the agentic workflow, AI agents, and tools.

Supervised execution mode for AI agents

You can minimize potential negative impact of an AI agent not executing as expected by configuring AI agents' tools to run in supervised mode. This will ensure human oversight for the tool's actions. You can use the Supervised mode to enhance security for agents with the capability to perform sensitive or critical actions.

You can set the supervised execution mode when creating a tool in the AI agent guided setup. For example, choose Supervised as the Execution mode when adding a catalog item tool. For reference, see Add a catalog item to an AI agent.

AI agent learning

Enhance AI agent learning through episodic memory, enabling agents to improve by learning from past successful interactions.

Episodic memory overview

The represents historical experiences and their associated learnings in the form of feedback. It's the ability to reflect on and learn from them by applying memory of similar episodes to new experiences. This capability enables agents to extract meaningful patterns and insights from experiences and use them effectively in the future.

Understanding episodic memory in AI agents

When an agent learning is enabled, the AI agent:

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

- Focuses on learning from execution outcomes, building episodic memory of what went wrong and how it was fixed.
- Acts as a feedback refinement loop, converting agent failures into structured learning for future tasks.
- Memory is heavily based on experience correction rather than just context extension.
- Memory entries carry relevance weights (like attention scores) → aligns with adaptive retention seen in human-like learning.

To configure agent learning at the AI agent level, see Define the specialty of an AI agent.

To configure agent learning at the AI Agent Studio level, see Set up long-term memory.

Configure Now Assist AI agents

Configure the Now Assist AI agents to execute agentic workflows with AI agents and mapped tools.

AI agents follow your instructions and act toward a specific goal and outcome by using the tools that you configure for those agents. By using the context of your record and your searchable content, AI agents can plan and analyze the task with a business logic that is combined with the instructions that are sent to large language models (LLMs) that suggest the next best action to be taken.

Note: Make sure that your record data and knowledge base have the latest accurate information for the best results.

Preparing for AI agents configuration

Prerequisites

By making a plan, you can improve your AI agent performance and result quality. When you have a solid foundation of what you want to build, you can minimize creating redundant agents and maximizing the efficiency of your existing ones. Before you send instructions to your AI agents, make sure that you follow these prerequisites:

- Have a good idea of the different kinds of tasks that your agentic workflow should be able to handle.
- Understand the general flow for your agentic workflow and agents.
- Use agentic tools with well-written descriptions.

Configurable elements

Instruct the agentic workflows and AI agents through the following elements within the framework:

- **Base plan**: Instructions to the AI Agent Orchestrator for the initial planning procedure that is configured at the agentic workflow level.
- **Role**: Clear identity of the AI agent that includes these elements:

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

- **Agent reasoning**: When a role is added to each reasoning prompt, it provides a sense of identity to the content that is generated by the LLM.
- **Agent proficiency**: An LLM-generated description of what an agent is capable of, including the content from the role, instructions, and the descriptions from the tools that are assigned to the AI agent.

  Note: The agent proficiency is auto-generated.

- **Instructions**: Clear directives for the AI agent. Write instructions as a step-by-step algorithm that describes the operational flow for the AI agent.

Understanding the tools for agentic workflows configuration

Define the procedure to build functional tools for your agentic workflow with the following three elements:

Functionality

What an AI agent contributes to the agentic workflow. Configure the tools with a single purpose. Multipurpose tools can cause a problem for the agents for the following reasons:

- Multipurpose tools are harder for the AI agent to reason through and determine when to use the tool. If a tool can be used for more than one purpose, the AI Agent Orchestrator has to determine which purpose is most applicable, which can decrease your AI agent's performance by increasing the runtime.
- The tool description must be comprehensive enough to account for all the scenarios for the usage of the tool that is being defined.

Note: Don't use tools that can operate in different modes. Instead, configure your tools as the solution to a singular problem for a scenario.

Tool description

Natural language descriptions that describe the utility provided by the tool. Make sure that you define the scope and limits of the tools clearly to help ensure that the tools are picked for the appropriate scenarios in the following ways:

- Provide a description of what the tool is supposed to do.
- Describe the scenarios where the tool can be called. Include the specific agentic workflows and tasks where the tool and its functionality can be used.
- Explore the scenarios where the tool is explicitly not useful but an AI agent can confuse the tool as being useful.
- Explain the terms that are being used in the preceding cases. For example, if you have a tool for assigning a role to a user, you must explain what the role is in the agentic system of the given instance.

Error messages

An AI agent operates through trial and error. For example, an error message about an execution that accidentally ran incorrect tools can help the AI agent reach more valid conclusions in the future. Error messages offer an AI agent a chance to reflect and explore other options.

Understanding the scenarios where the tool can go wrong can help the AI agent with keeping the execution on track.

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

Invoke Conversations with AI Agent Background Channel

The AI Agent Background Channel helps you to invoke AI Agent or agentic workflow execution from the Workspace. Use the AI Agent Background Channel associated with the AI Agent Background Provider to invoke conversations. The AI Agent Background Provider is based on the Custom Adapter Framework from Virtual Agent. For more information, see Configure a provider for your custom chat integration.

Create a channel identifier in the Provider Channel Identities table [sys_cs_provider_application] to add any additional conversational capabilities to your own provider application and get a new inbound ID that allows for customization. For more information, see Create a channel identifier for your custom chat integration.

To start a conversation, trigger the flow using the `sn_aia.AiAgentRunttimeUtil().startAiAgentConversation(request)` API in the Script Include (sys_script_include) of the AIAgentBackgroundProvider and select **Run Script**. When the Script execution status indicates Success, the conversation begins in the order of utterances defined in the Script.

Conversations that are invoked for executing an AI agent are logged in the Execution Plans [sn_aia_execution_plan] table. Open the conversation record to confirm the device type as **AI Agent Background**. Open the execution record to see the **Execution Tasks**, **Messages**, and the **Tools Executions** used to execute the AI agent.

You can also see the entire execution steps on the AI Agent Studio Testing page by copying the execution plan record's [sys_Id] and testing it. On the **Chat responses** tab, in the AI agent decision logs, you can see the AI agent details and the tools it used to resolve the issue.

Interactive and Non-interactive AI agents

The Interactive AI agents reach out to users for information when there is a fallback in the execution process, and the AI agent re-triggers the flow.

The Non-interactive AI agents do not reach out to the user at any fallback stage in the execution process. When the AI agent needs user information, it takes the dynamic prompt approach using the ReAct layer, where the prompt of the ReAct will change based on the execution mode of the AI agent or agentic workflow. Therefore, in the Non-interactive execution, the reach fallback options do not have to collect input from a user as a fallback option. However, the output of the AI agent or agentic workflow will still need to be presented to the user, and in any execution failure scenario, a message in the Now Assist panel or Virtual Agent is shown.

To implement the Non-interactive execution, the **Execution Mode** field is added in the Execution Plans [sn_aia_execution_plan] table, where the execution mode can be **Interactive** or **Non Interactive** based on the given runtime parameter.

You can run the AI agents and agentic workflows concurrently in the AI Agent Background Channel and in Non-interactive mode where the background execution allows AI agents to operate with any chat panel like Now Assist panel or Virtual Agent.

Multilingual support

You can leverage multilingual support for AI agents across languages for better translation quality to:

- Tune system prompts for native translations.
- Implement dynamic translation strategies when native support is unavailable.
- Provide extensive testing via automated and manual evaluations.

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

Install Now Assist AI agents

Install Now Assist AI agents on your ServiceNow instance to enable the agentic AI experience.

Before you begin

To get started with AI agents, you must have:

- A ServiceNow Pro Plus or Enterprise Plus license.
- An instance on Xanadu Patch 7+ or Yokohama.
- AI Search enabled on your instance.

Role required: sn_aia_admin

About this task

AI agents aren't standalone applications that you can install directly. To enable AI agents on your instance, you must install and activate other Now Assist applications that include AI agents, such as Now Assist for IT Service Management (ITSM) or Now Assist for Customer Service Management (CSM).

Procedure

1. Navigate to **All** > **System Definition** > **Plugins**.
2. Search for and select a Now Assist application, such as Now Assist for IT Service Management (ITSM) or Now Assist for Platform.
3. Select **Install**.

Result

AI agents are associated with that application are installed on your instance.

What to do next

You can access AI agents in the Now Assist panel and in Virtual Agent. To enable the Now Assist panel, see Turn on the Now Assist panel.

Set up Now Assist AI agents

The default (out-of-the-box) AI agents provide preconfigured agentic workflows that address common business challenges across ServiceNow applications. Before activating the default AI agents, you must ensure that your instance meets the prerequisites and complete the required configuration steps.

Before you begin

Role required: sn_aia_admin

About this task

Prerequisites

Platform version:

- Minimum: Yokohama Patch 1+ or Xanadu Patch 7+
- Recommended: Yokohama Patch 8 or Zurich Patch 4

Required plugins and store Apps

- Now Assist AI agents store app (version 6.0.17 or later as of December 2025)
- At least one Now Assist workflow plugin:

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

- Now Assist for ITSM
- Now Assist for HRSD
- Now Assist for CSM
- Now Assist for Security Operations

- Generative AI Controller plugin (minimum version 11.0.0)

Note: Select **Load demo data** during installation to access preconfigured examples.

Licensing Requirements

Now Assist Pro Plus or Enterprise entitlement

Required User Roles

- sn_aia.admin - Provides access to AI Agent Studio
- sn_nowassist_admin.nsa_admin - Enables Now Assist administration

Procedure

1. Enable AI Search.
   a. Navigate to **AI Search** > **AI Search Status**.
   b. Verify that AI Search is enabled.
2. Enable Now Assist panel.
   a. Navigate to **Now Assist Admin** > **Experiences**.
   b. Turn on the **Panel** option.
3. Access the AI Agent Studio.
   a. Navigate to **All** > **AI Agent Studio** > **Overview**.
   b. Review the available default (out-of-the-box) agentic workflows.
4. Activate the default agentic workflows.
   a. Select the agentic workflow that you want to use from the available default workflows.
   b. Activate the workflow. The default workflows are set to read only and must be activated before use.

   The default (OOTB) agentic workflows include:

   - Generate resolution plans
   - Analyze incident trends
   - Classify tasks
   - Investigate IT problems
   - Process emails for tasks
5. Configure chat assistants.
   a. Navigate to **Conversational Interfaces** > **Assistants**.
   b. Determine which chat assistants can access your AI agents.
   c. Complete the configuration wizards for Now Assist panel (NAP) or Now Assist Virtual Agent (NAVA).

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

6. Configure LLM Connections (optional).

   If using external LLM providers instead of the default OpenAI GPT-4.1:

   a. Navigate to **Now Assist Admin Console** > **Settings** > **Manage Integration**.

   b. Configure API credentials for your chosen LLM provider.

   Note: The default Orchestration layer uses OpenAI GPT-4.1 on ServiceNow-managed Azure servers.

7. Test the default AI agents.

   a. Navigate to **AI Agent Studio** > **Testing** > **Test AI reasoning**.

   b. Verify that the orchestrator canvas displays agent execution properly.

   c. Review AI agent decision logs to confirm interactions and thought processes are recorded.

Enable Now Assist Guardian for AI agents

Identify and block offensive messages that are sent by human agents automatically by enabling Now Assist Guardian in AI agents. With this capability, you can help reduce your agentic workflow or test from being exposed to harmful content.

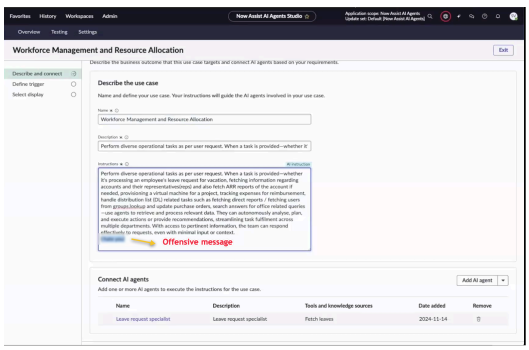
Before you begin

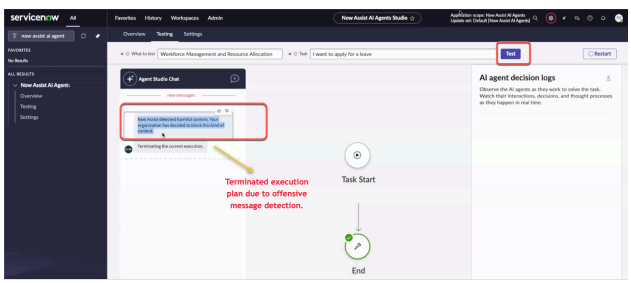
Role required: admin

About this task

The Now Assist Guardian, which is a ServiceNow AI Platform capability in the Now Assist panel, is a set of guardrails that are designed to intercept and mitigate offensive, sensitive, or security-related issues that may arise during interactions with the Now Assist application.

For example, let's say that Now Assist Guardian detects an offensive message in the execution plan of an agentic workflow. When you try to trigger the plan or test it, Now Assist Guardian can step in to terminate the plan or test because it detected harmful content at the first step of the execution plan.





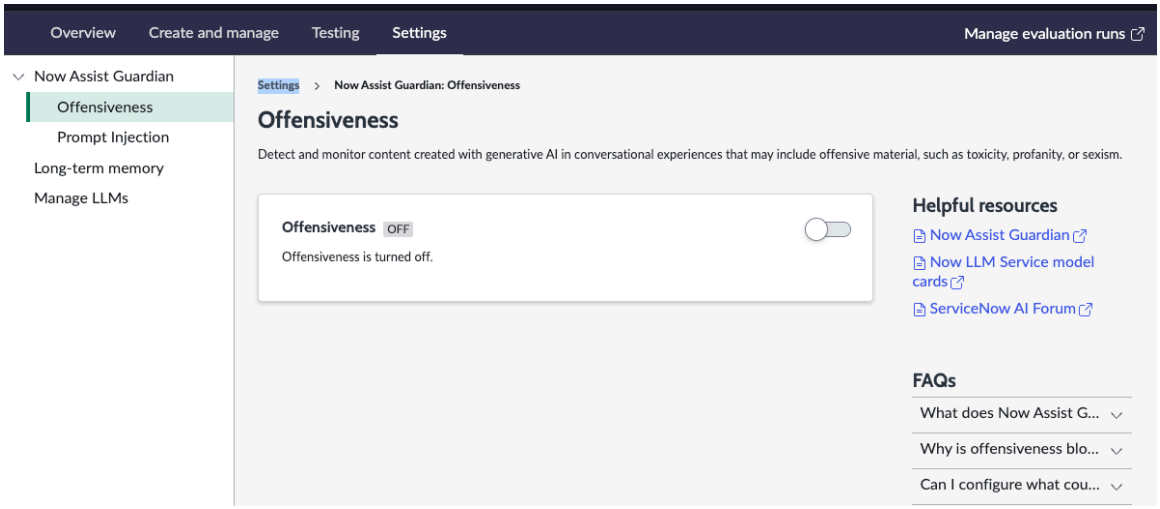
For more information about the different guardrails, see Now Assist Guardian.

Procedure

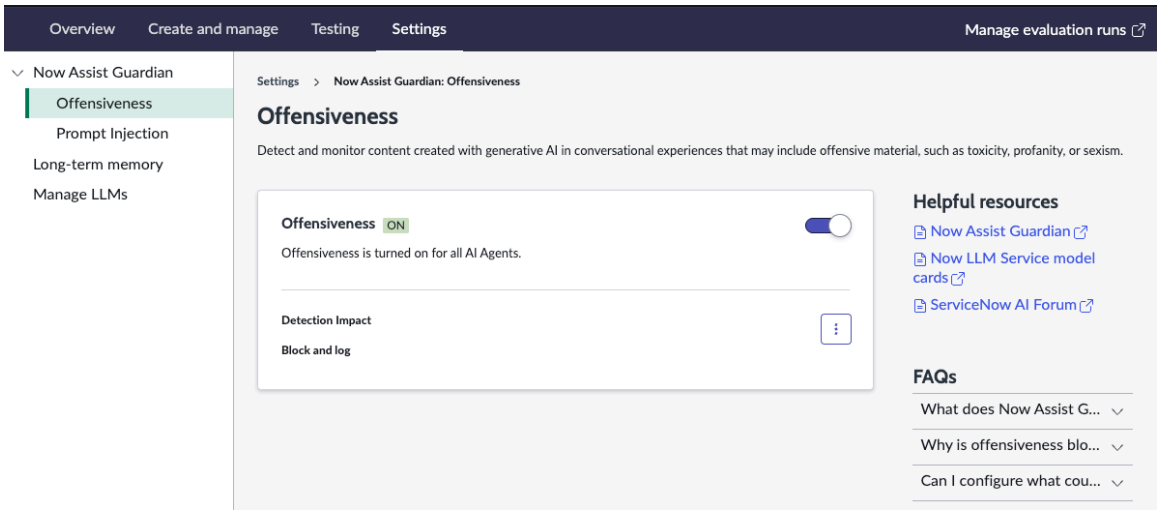
1. Configure Offensiveness for AI agents.

   a. In AI Agent Studio, navigate to the **Settings** tab.
   You're directed to the Offensiveness page.

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

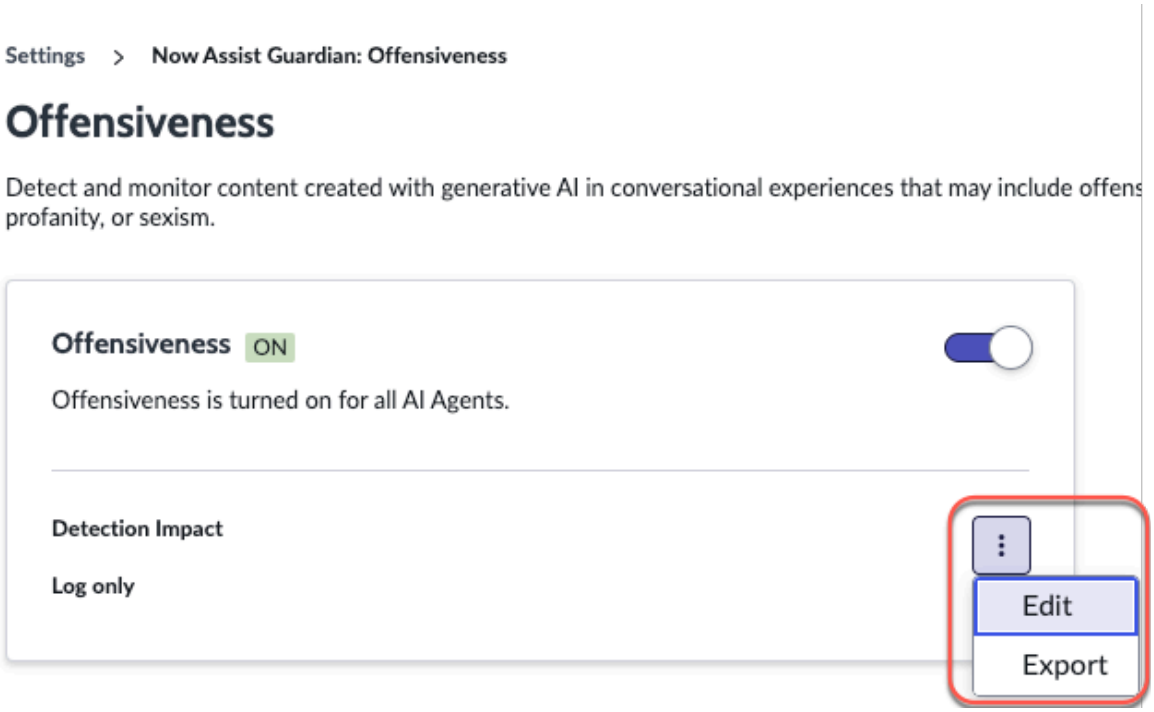


b. Turn on the Offensiveness setting for AI agents by using the toggle button.



c. Configure the detection impact by selecting the options icon (⋮) to enable the detection impact to use the following options:

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



- **Edit**: Choose the detection impact between logging or both blocking and logging.

  Enable the detection impact:

  i. Select the **Edit** option.

  ii. On the Detection impact page, select either the **Log** or **Block and log** button based on your requirements.

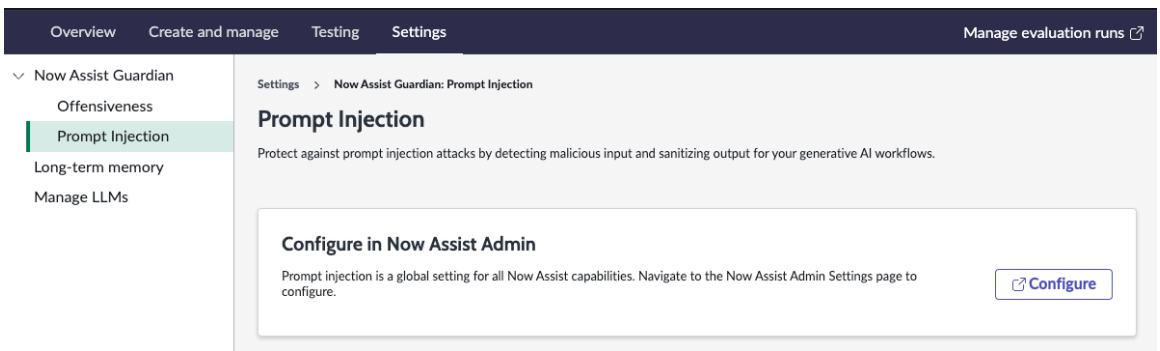
  > **Note:** You can switch the detection impact from **Log** to **Block and log** or from **Block and Log** to **Log** at any time.

  iii. Select **Save**.

- **Export**: Exports the offensiveness-based logs as a . CSV file. Logs can be analyzed for the types of content that are being identified so you can take other preventive measures, such as changing conversational questions.

2. Configure Prompt Injection for AI agents.

   a. In the AI Agent Studio, navigate to **Settings** > **Prompt Injection** and select **Configure**. You're directed to the Now Assist Admin page to configure the Prompt Injection.



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

Note: For more information about configuring the Prompt Injection, see Configure prompt injection attack protection.

When you configure the Prompt Injection for an agentic workflow by using the required instructions, the system is designed to detect harmful content and block the conversation.

Select the LLM for AI agents and agentic workflows

Choose the large language model (LLM) service provider for Now Assist AI agents in AI Agent Studio.

Before you begin

Role required: sn_aia.admin

About this task

LLMs are part of the foundation of agentic AI. Different LLMs can provide slightly different performance and responses. You can select which LLM to use at a global level for agentic AI from the AI Agent Studio to help adjust the response quality to best fit your agentic workflows.

Note: If you already have agents built and you change the global LLM, then you must test the agents after making the change.

Depending on your region, you may have to consent to using a different service provider.

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Settings**.
2. Navigate to the **Model provider** tab.
3. Select **Configure** to be redirected to Now Assist Admin.
4. In the *Choose a default model provider* field, select

   Your choices are the following:

   - Now LLM Service
   - AWS Claude
   - Azure OpenAI
   - Google Gemini
5. Review the impacted AI agents.

   Certain AI agents may require specific model providers. Select the AI agents tab in the **Impact overview** section to see which AI agents are affected by any changes to the model provider.
6. Select **Save**.

Result

Your chosen LLM is used globally for all AI agents and agentic workflows.

Set up long-term memory

Make AI agents remember your preference or facts from previous interactions and use memories for more focused conversations.

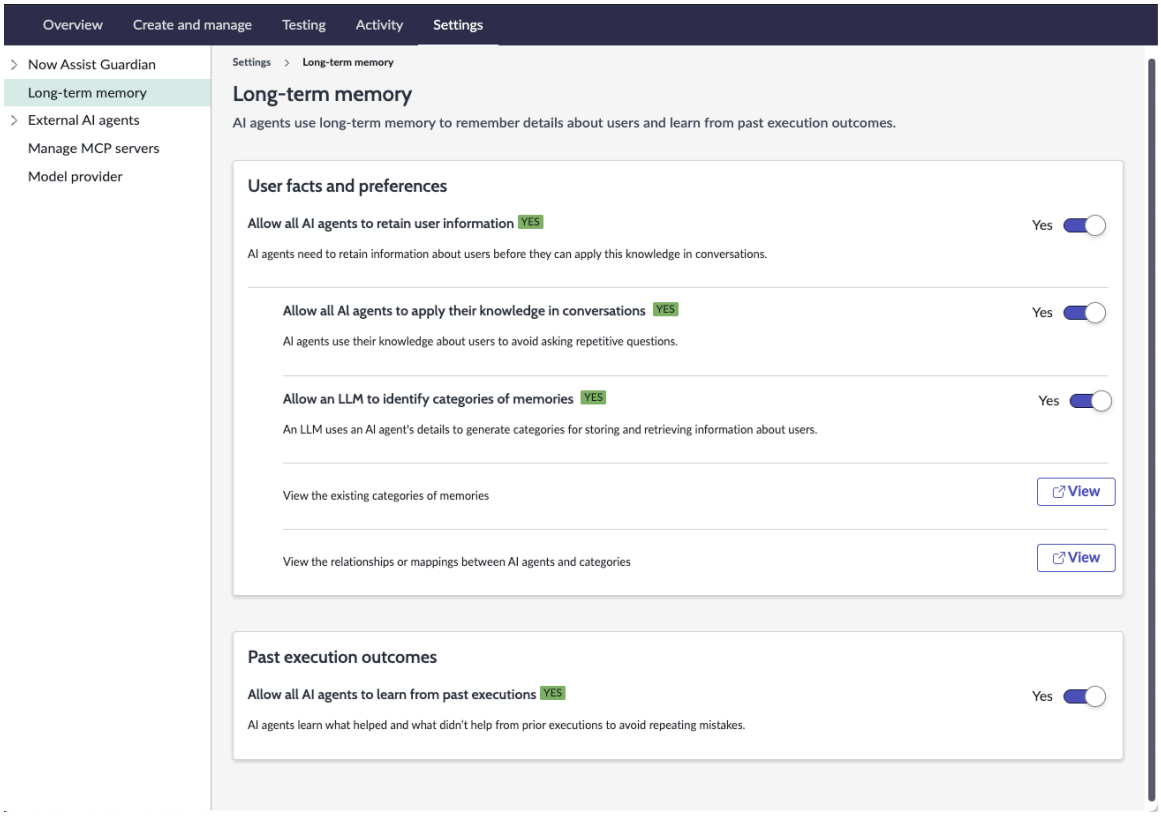
Before you begin

Role required: sn_aia.admin

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

Procedure

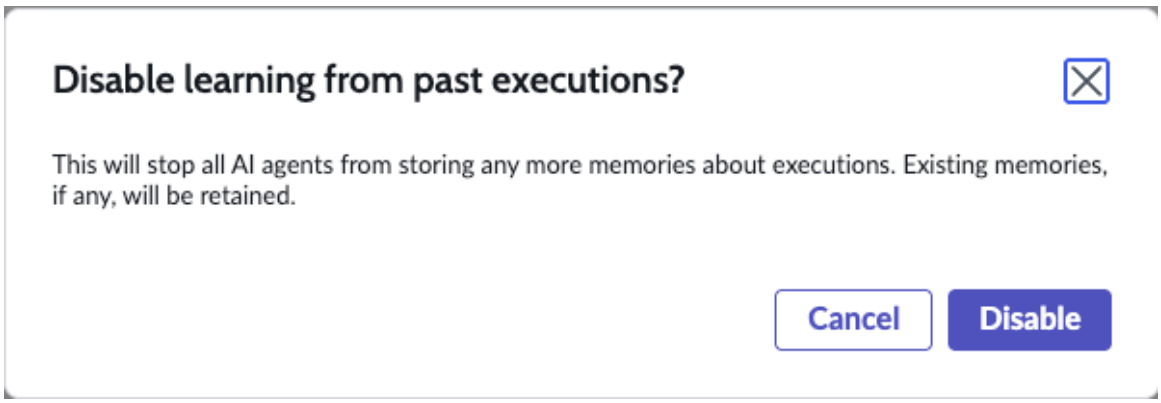
1. Navigate to **All** > **AI Agent Studio** > **Settings**.
2. Select **Long-term memory**.



3. Configure the User facts and preferences.
   a. Turn on the Allow all AI agents to retain user information using the toggle button.
   b. Turn on the Allow all AI agents to apply their knowledge in conversations using the toggle button.
   c. Turn on the Allow an LLM to identify categories of memories using the toggle button.
   d. To view the existing categories of memories, select **View**.
   e. To view the relationships or mappings between AI agents and categories, select **View**.
4. Configure Past executions outcomes by turning on the Allow all AI agents to learn from past executions using the toggle button.

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

When you want to turn off agent learning from past executions, you will see a confirmation pop up. Selecting **Disable** will disable agent learning.



Long-term memory categories

Long-term memory (LTM) categories define the types of semantic information that a Now Assist AI agent can learn and retain about users over time. You can add new categories and map them to specific agents to personalize agent responses based on accumulated user context.

Semantic memory in the AI agent Memories table (`sn_aia_memory_list`) is organized by LTM categories. Each category represents a distinct type of user-specific information, such as software preferences, workplace context, or communication style. By mapping categories to an agent, you control what the agent learns and retains across interactions.

How LTM categories work

When an agent executes, the platform evaluates the interaction for user-specific facts that match the agent's configured LTM categories. Matched facts are stored as semantic memory records in the AI Agent Memories [sn_aia_memory] table, scoped to the user and category. On subsequent interactions, the agent retrieves relevant semantic memories and uses them to personalize its responses without requiring the user to repeat context.

LTM categories are defined globally and can be mapped to one or more agents. An agent only learns and retrieves memories for the categories explicitly mapped to it.

Default LTM categories

The platform includes the following default LTM categories:

Software and tools

Captures information about applications, tools, and software configurations relevant to the user, such as operating system version or approved applications.

Work context

Captures facts about the user's role, department, location, and workplace preferences, such as remote work setup or team structure.

User preferences

Captures communication preferences and interaction style, such as preferred language, response format, or notification settings.

You can extend this list by creating custom categories suited to your organization's use cases.

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

How categories affect memory extraction

During memory extraction, the platform runs an LLM prompt that evaluates the agent interaction against the descriptions of each mapped LTM category. If the interaction contains information that matches a category, a semantic memory record is created or updated in the AI Agent Memories table with the following fields:

Category

The LTM category that the memory is associated with.

User

The user whose interaction generated the memory.

Memory

The extracted user-specific fact, stored as a JSON object.

Type

Set to **Semantic** for category-based memories.

Semantic memories are retrieved at runtime using retrieval-augmented generation (RAG) and injected into the agent prompt to personalize the response for the current user.

Considerations

- Category descriptions directly influence LLM extraction quality. Use specific, unambiguous language to reduce false positives or missed extractions.
- Mapping too many categories to a single agent can increase extraction processing time. Map only the categories relevant to the agent's use case.
- To verify that semantic memories are being extracted, open the AI Agent Memories table and filter by **Type** = **Semantic** and the relevant agent or user.

Create long-term memory category

Add a long-term memory category to add it an AI agent while setting up long term memory

Before you begin
Role required: sn_aia.admin

Procedure

1. Navigate to **AI Agent Studio** > **Settings** > **Long-term memory**.
2. Select **View All** against View the existing categories of memories.
3. On the AI Agent Memory Categories page, select **New**.
4. On the form, fill in the fields.

Field	Description
Name	A descriptive label for the category, such as `Hardware preferences` or `Accessibility needs`.
Description	A brief explanation of what type of user information this category captures. The description is used by the LLM extraction prompt to identify relevant facts during agent

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

Field	Description
	execution, so precise language improves extraction accuracy.

5. Select **Save**.
   After saving the record, the category is activated by default and is available to map to one or more AI agents.

Map Long-term memory category

Map the created long-term memory category to the agents that should use it. Categories must be mapped individually to each agent; they are not applied globally by default.

Before you begin
Role required: sn_aia.admin

About this task
The mapping of LTM categories is done directly while creating an AI agent at the Define the specialty page.

Procedure
1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **AI agents**.
2. Navigate to the Define the specialty page of an AI agent form.

   Note: You can add map the categories for a new or existing AI agent.

3. In the manage long-term memory section, select categories that you want to map to your AI agent.
   You will see the list of categories that were created in the AI Agent Memory Categories page. For more information, see Create long-term memory category.
4. Select **Save and continue**.

Result
The selected categories are mapped to your AI agent.

Configure Group Action Framework

Set up Group Action Framework (GAF) to improve the response quality, recall speed, and consistency of AI agents.

Before you begin
There must be a ml_platform read ACL for GAF to be configured. If there is no read ACL present, you must create it and grant the appropriate role access, such as admin, ml_admin, or sn_aia.admin. This is required for GAF to access Machine Learning tables and services. If your ACL is not configured correctly, you may see the error "Failed to initialize pipeline: Failed to load message_content dataset. No columns to parse from file."

To access GAF's optimized prediction feature, you can enable Now Assist in AI Search. For more information, see Setup AI Search for Group Action Framework.

Role required: sn_aia.admin

About this task
You can activate GAF to have AI agents use indexed clusters that perform LLM executions on representative records rather than all records. GAF is used by some AI agents and agentic

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

workflows to work optimally. For more information about GAF and how it works, see Group Action Framework.

You can have different GAF configurations for different agentic workflows and Now Assist applications. You must configure each agentic workflow or application separately.

Important: Setting up GAF can take some time, between 10 minutes up to an hour, depending on the number of records in the grouping. The script runs in the background.

Procedure

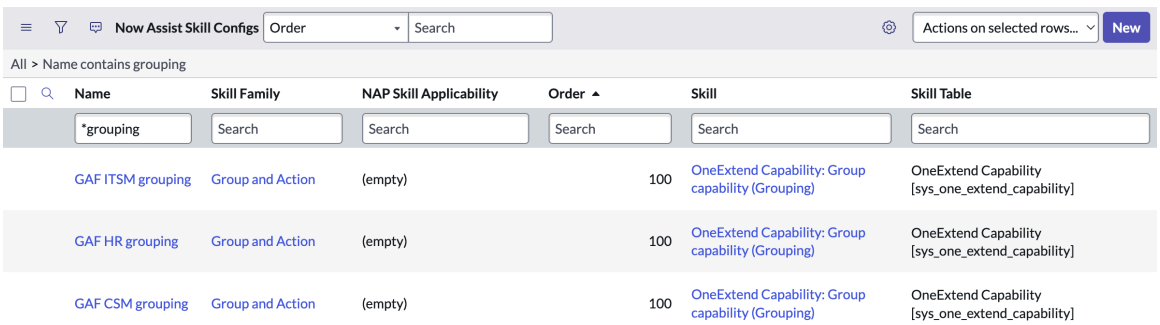
1. Navigate to **All** > **System Definition** > **Scripts - Background**.
2. Paste the following code into the text area.

   Do not run the script until the groupSkillID and the actionSkillID variables have values. The values are acquired in the following steps.

   ```
   var groupSkillId = "";
   var actionSkillId = "";
   var topicSkillId = "43bce9e477e012103f075cea5b5a998f";
   new sn_gaf.GAFUtils().activate(groupSkillId, topicSkillId,
    actionSkillId, "run_once");
   ```

   The topicSkillId is the same for all GAF setups.

3. In a new browser tab, navigate to the Now Assist Skill Config [sn_nowassist_skill_config] table by entering `sn_nowassist_skill_config.list` in the filter navigator.
4. In the **Name** field, enter `*grouping` and search to see the grouping records that are associated with your agentic workflows and Now Assist applications.

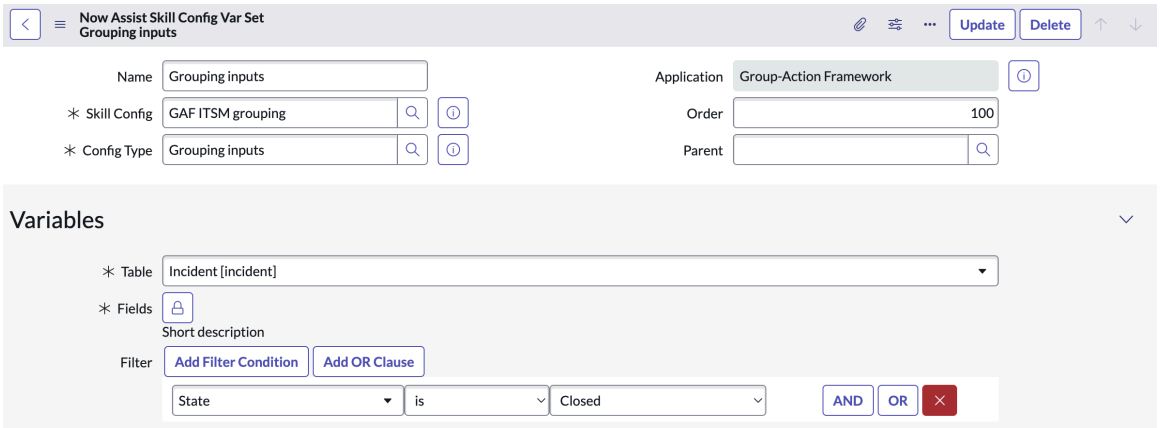


5. Open the grouping record for the agentic workflow or application that you want to configure.
6. In the Now Assist Config Var Set related list, open the Grouping inputs record.
7. Confirm that the filters for the grouping include all the records you want to index.

   If you want to add filters to change what records are included in the GAF setup, do so here. The more records that you include means that you have a longer setup time, but they can help increase the quality of your results. You should have at least 2000 records for successful clustering.

   Note: Your current scope should be Group-Action Framework to make changes.

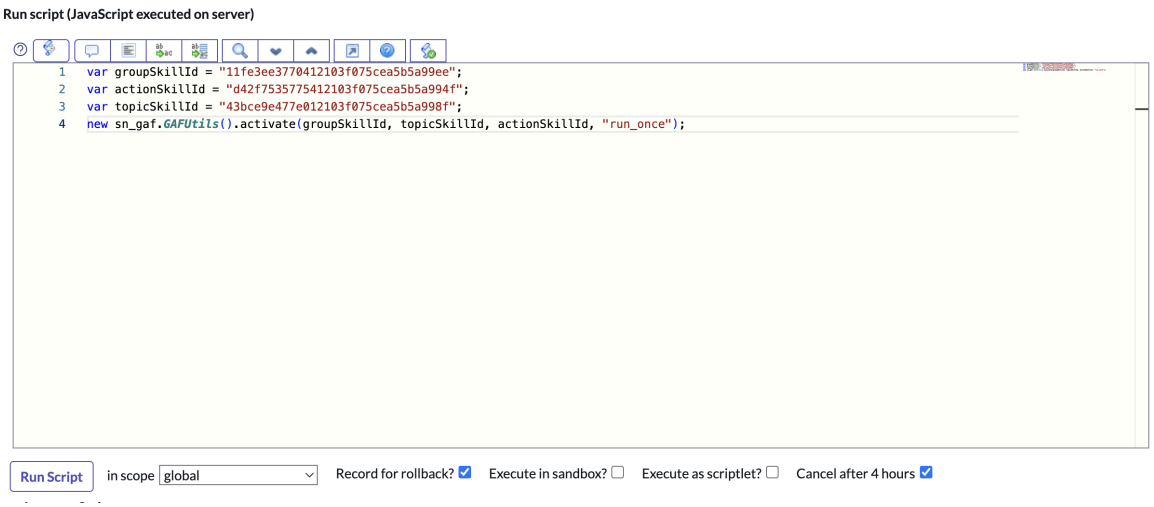
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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



8. Return to the Now Assist Skill Config Grouping record and copy the sys_id.
   You use this sys_id and two other sys_ids in the background script function call.
9. In the browser tab with the background script, paste the sys_id in the script between the quotation marks for the groupSkillId variable.

   Your `groupSkillId` variable should be in the same format as the `topicSkillId` variable.
10. In the tab with the Now Assist Skill Config record, return to the Now Assist Skill Config table and search for `*action strategy` to find the action strategy skill config for the application. Make sure that the record is for the same application. If you pasted the sys_id of the GAF ITSM grouping, you must open the GAF ITSM action strategy.
11. Open the action strategy record for the application that you're configuring.
12. Copy the sys_id of the action strategy record.
13. In the browser tab with the background script, paste the sys_id in the script between the quotation marks for the actionSkillId variable.

    Your `actionSkillId` variable should be in the same format as the `groupSkillId` and `topickillId` variables.



14. Run the background script by selecting **Run script**.

    Running this background script creates a scheduled job called GAF - Run Offline Flow. You can view the scheduled job on the Scheduled Script Executions [sysauto_script] table.

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

Result
GAF is configured on your instance for that Now Assist application and can be used by AI agents to find related records.

What to do next
To verify that grouping and action outputs have been generated, go to the ML Solution [ml_solution] table to check if the clustering solution is running or completed. If it is complete, check the following tables to see records for groups, clustered records, and down-sampled records per group are present.

- GAF record group [sn_gaf_record_group]
- GAF record group detail [sn_gaf_record_group_detail]
- GAF action strategy result [sn_gaf_action_strategy_result]

You can repeat this procedure for additional agentic workflows and Now Assist applications.

You may see the error "Failed to initialize pipeline: Failed to load message_content dataset. No columns to parse from file." when trying to configure GAF. This could be because your instance lacks data on the table or filters you configured excluded records from the table. You can reach out to Now Support for additional assistance if you cannot resolve this error.

Set up AI Search for GAF

Configure AI Search to enable Group Action Framework (GAF) to improve quality and consistency of agentic AI and Now Assist generative AI on the Now Platform.

Before you begin
Role required: admin

About this task
GAF is a feature on the Now Platform that clusters and indexes related records and executes actions on them in agentic AI and Now Assist generative AI. See Group Action Framework for more information about GAF's role in intelligent experiences and how it works.

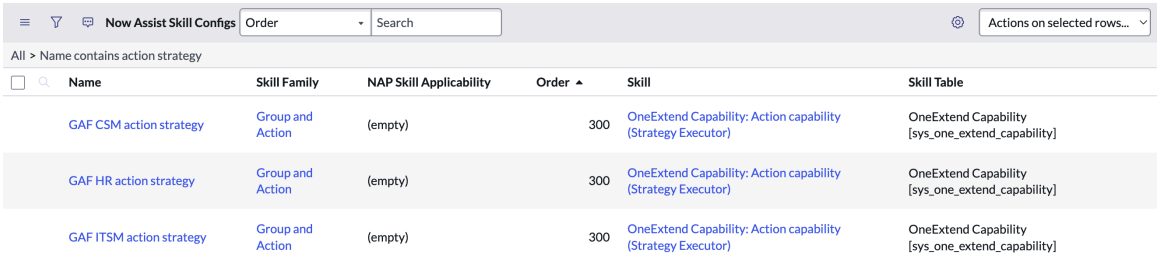
Now Assist in AI Search is the foundation for GAF's optimized prediction feature. AI Search is the backup search for certain workflows, and if it is not enabled and ready, GAF will not return any results.

You must index every table for each workflow or application you'd like to configure GAF for. For example, if you are configuring GAF for Now Assist for IT Service Management (ITSM), you must index the Incident and related tables.

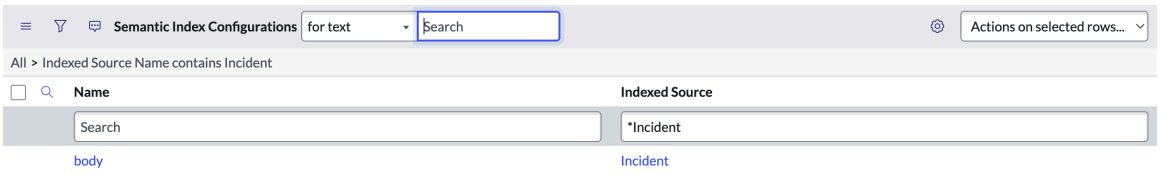
Procedure
1. Navigate to **All** > **AI Search** > **AI Search Status**.
   If you see “AI Search is ready,” proceed. If not, request AI Search installation, which can take 2–24 hours.
2. Install Now Assist for AI Search by installing a Now Assist application.
   For more information, see Install Now Assist in AI Search.
3. Navigate to the Now Assist Skill Config [sn_nowassist_skill_config] table by entering `sn_nowassist_skill_config.list` in the filter navigator.
4. Search the Now Assist Skill Config table for action strategy skills by searching *Name* is `*action strategy`, and select the action strategy skill for the application you want to configure.

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

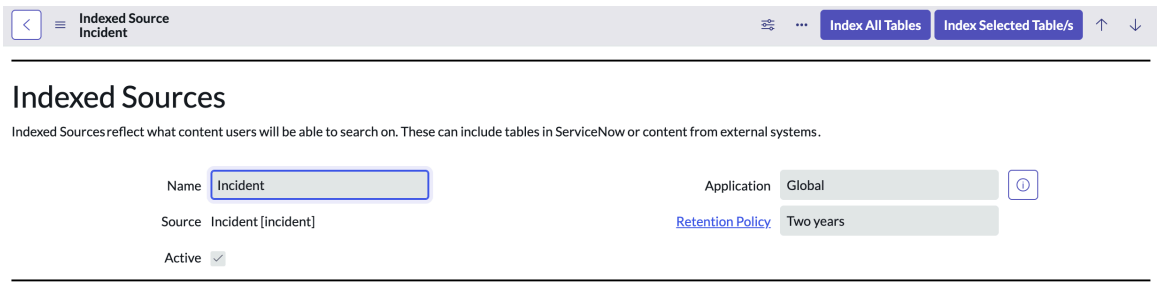
For example, if you are configuring GAF for Now Assist for ITSM, select the GAF ITSM action strategy record.



5. In the related list, select the Now Assist Skill Config Var Set record.
6. Set *`Semantic Index Name`* to the appropriate field for your table.
   To find the field for your table, go to the Semantic Index Configuration [ais_semantic_index_configuration] table by entering `ais_semantic_index_configuration.list` in the filter navigator, filter the *`Indexed Source`* for the table you want to index, and use the value in the *`Name`* field. For example, the value is *`body`* for the Incident table.
7. Navigate to the Semantic Index Configuration [ais_semantic_index_configuration] table by entering `ais_semantic_index_configuration.list` in the filter navigator, filter the *`Indexed Source`* for the table you want to index and the *`Semantic Index Name`* you configured in the previous step, and set *`Active`* to `true` to activate the embedding model.



8. Navigate to **All** > **AI Search** > **AI Search Index** > **Indexed Source**, filter the list so that the *`Name`* is the name of the table you are configuring, and select the record.
9. Select **Index All Tables**.



Result

The tables for Now Assist in AI Search for your Now Assist application is indexed based on the appropriate field.

What to do next

To confirm that indexing has occurred successfully, check the Indexed Source History related list and ensure that both the *`Keyword Ingestion State`* and *`Semantic Ingestion State`* are both set to *`Indexed`*.

Once your tables have been indexed, you can continue to Configure Group Action Framework to set up GAF for each application.

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

Multiple conversations in Now Assist AI agents

Multiple active conversations enable live agents to maintain separate conversations for different records. You can preserve the context of multiple conversations and enable multiple AI agents to interact at the same time through the Now Assist panel.

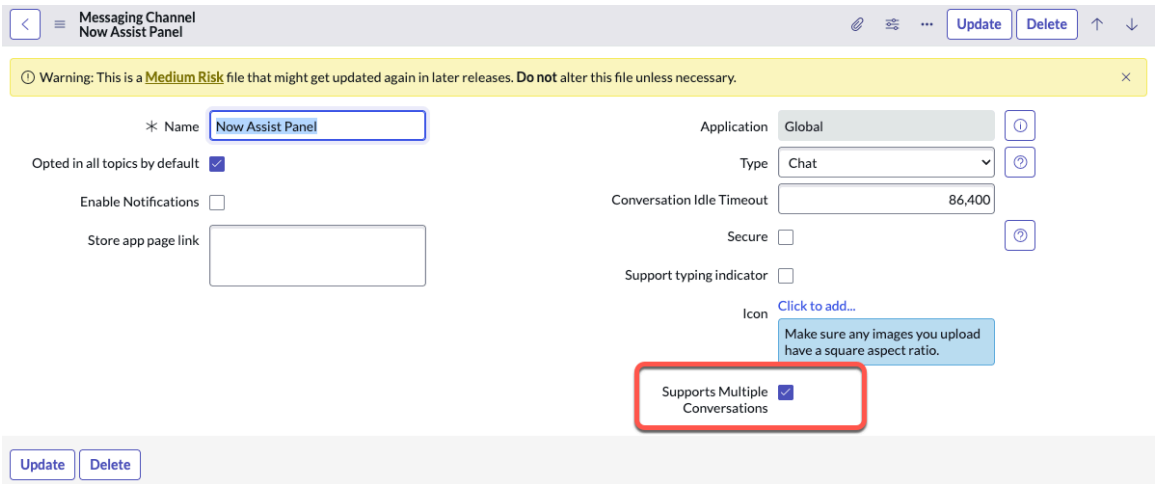
Multiple conversations between live agents and AI agents overview

Once configured, live agents can interact with multiple AI agent conversations through the Now Assist panel.

Note: By default, the Now Assist panel currently supports only a single conversation and it is required for you to enable the multiple conversation support if you intend to use it

Enabling multiple conversations

Enable multiple conversations on a ServiceNow instance in the Messaging Channels [sys_cs_channel.list] table by setting the value of the **Supports Multiple Conversations** field on the Now Assist panel record to **true**. If you don't see this field, make sure you are in the Global scope and are using the Default view. By saving the record and refreshing the instance, you can see the multiple active conversations features in the Now Assist panel.



Title for the conversation

With the multiple conversations feature enabled in the Now Assist panel, the first utterance from the live agent is set as the title for that conversation. For example, if the first utterance is `Explain change risk`, then Explain change risk is set as the title for that conversation in the Now Assist panel.

Starting a new conversation

Start a new conversation by selecting the plus icon ( ) and choosing a topic. See the conversation list by selecting the All chats icon ( ), to see the Active chats and Closed chats.

Unread chats on the conversation list

See the batch count of the unread chats on the conversation list icon ( ). For example, if number 6 appears on the All chats icon ( ), that means there's six unread conversations on the Now Assist panel.

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

Citations in Now Assist AI agents

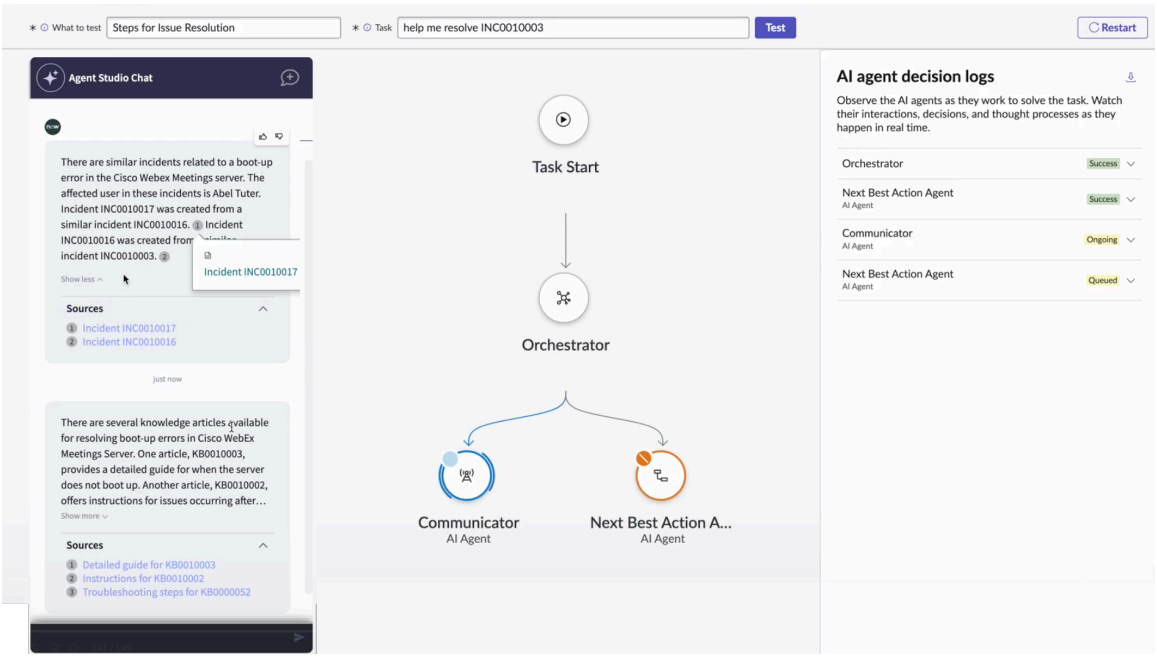
You can see citations on the Now Assist panel when you execute AI agents and agentic workflows in Now Assist AI agents that provide summaries to get similar incidents and relevant knowledge articles.

Generating citations:

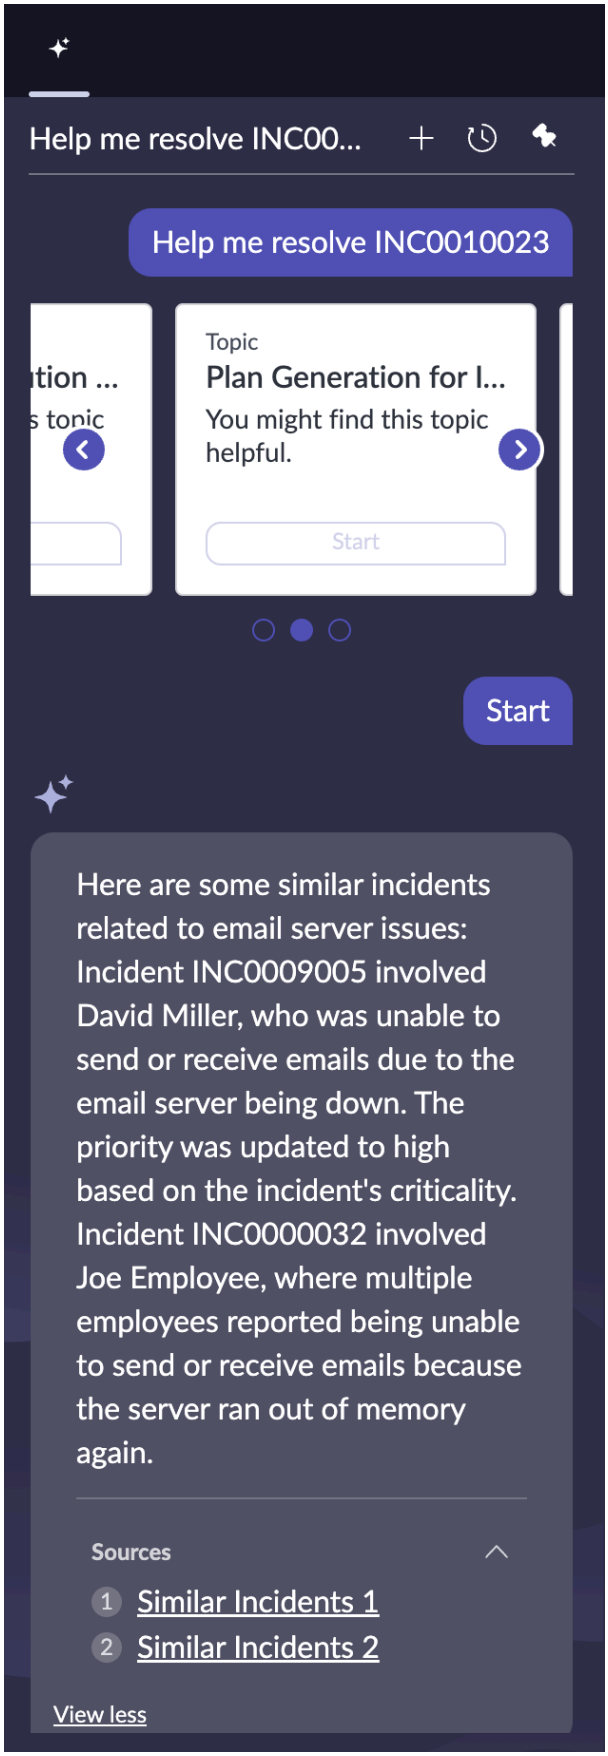
- Citations are generated by the LLM (Large Language Model), which use the information available to it.
- URLs can't be generated by LLM if relevant information isn't provided.
- LLMs need provide proper URLs for better results.
- Citation URLs return directly from the tool's outputs.
- The LLM is instructed not to generate or assume citations if they aren't present.

Note: Invalid or malformed URLs won't be displayed to the user.

Citations with similar incidents and relevant knowledge articles appear in the Sources section of the AI Agent Studio Chat and on the Now Assist panel and in with click-able links that direct you to the incident or knowledge article.



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



You can hide citations for specific agentic workflows or AI agents where it's necessary or confidential. For more information, see Disable citations in AI Agent Studio.

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

Disable citations in AI Agent Studio

Disable citations for specific agentic workflows or AI agents in AI Agent Studio where citations aren't required or involve confidential information.

Before you begin

Role required: sn_aia_admin

About this task

By default, the citations are enabled for all agentic workflows and AI agents. The Agent Properties [sn_aia_property] table enables disabling citations at the agentic workflow or AI agent level. Create a record only when you want to disable citations for specific agentic workflows or AI agents when the **Output transformation strategy** field in the Agent tool record is set either to **Summary for Search Results** or **Verbose**.

- All citations are generated by the LLM (Large Language Model), which uses the information available to it.
-

Procedure

1. Navigate to **All** > **System Definitions** > **Tables** and open the Agent Properties [sn_aia_property] table.
2. Select **New**.
3. On the form, fill in the fields.

Agent Property form

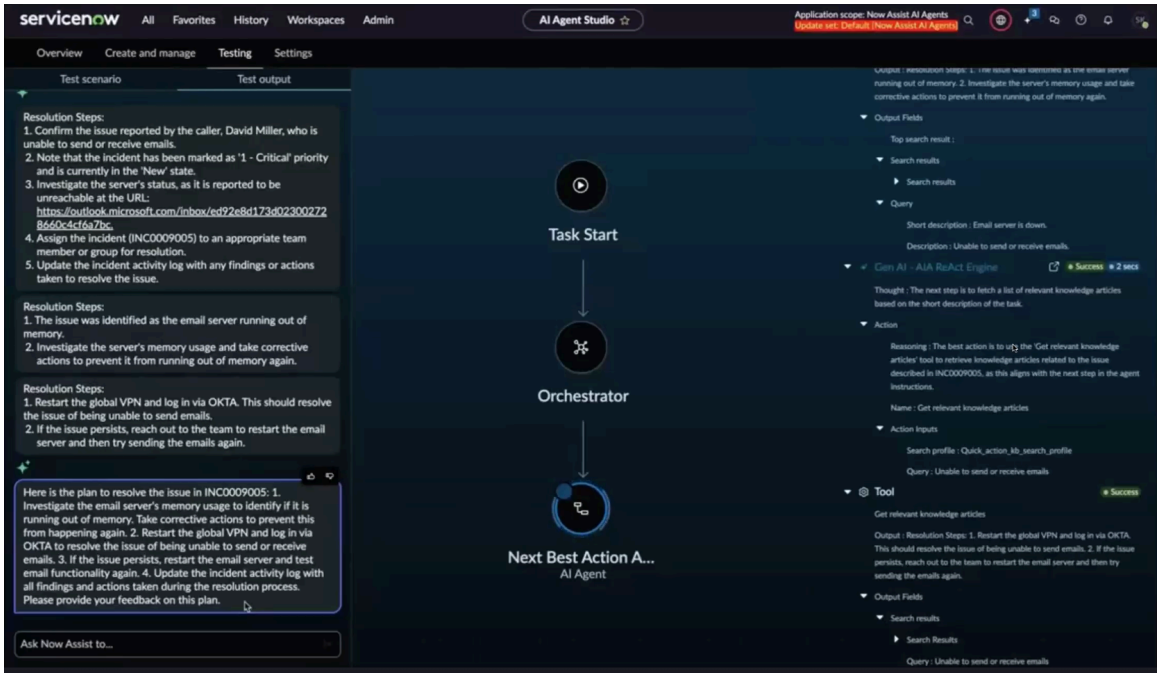
Field	Description
Name	Name for the agent property record used to hide citations for the agentic workflow or AI agent.
Value	Provide **false** as the value to hide the citations.
Target table	The table where the target agentic workflow or AI agent resides.
Target record	The record of the target agentic workflow or AI agent.
Application	The application scope for the agent property record: Now Assist AI agents.

4. Select **Submit**.

Result

Execute an AI agent or an agentic workflow for which you have disabled citations. You'll see the output without citations. See the following example for reference.

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



Follow-up conversations

AI agents continue with follow-up conversations after the AI agent execution is complete.

Before you begin

Role required: sn_aia_admin

About this task

When the [sn_aia.enable_follow_up] system property is set to **true** and if there is no record for the agentic workflow created in the Agent properties [sn_aia_property] table, then the AI agents continue with the `How else can I help you?` follow-up message after the agentic workflow execution.

The AI agent conversation doesn't get closed after the execution is complete by default. If you would like to override the follow-up behavior for your AI agents for a specific agentic workflow, you must create a record in the Agent Properties [sn_aia_property] table.

Procedure

1. Navigate to **All** > **System Definitions** > **Tables** and open the Agent Properties [sn_aia_property] table.
2. Select **New**.
3. On the form, fill in the fields.

Agent Property form

Field	Description
Name	Name for the agent property record for the agentic workflow.For example, `follow_up_behaviour`
Value	Provide the following values as per the requirement:

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

Field	Description
	○ **no_followup_close_conversation**: Closes the conversation after the use case is complete. ○ **no_followup_open_conversation**: Keeps the conversation open after the use case is complete.
Target table	The table where the target agentic workflow resides: Agentic workflows [sn_aia_use case] table.
Target record	The record of the target agentic workflow or AI agent.
Application	The application scope for the agent property record: Now Assist AI agents.

4. Select **Submit**.

Add recipients to usage spike, execution failure, and latency error email notifications

Add or change recipients to email notifications triggered by unexpected or undesired behavior in AI agent and agentic workflow executions.

Before you begin

Set your application scope to Now Assist AI Agents.

Role required: admin

About this task

There are different email notifications that can be configured: one for Assist usage spikes, one for failed AI agent executions, and two for latency errors. The properties for configuring the thresholds for sending these notifications can be found in the Now Assist AI Agents Reference.

Procedure

1. Navigate to **All** > **System Notification** > **Email** > **Notifications**.
2. In the *Name* field, search for `AI Agent`, and select the notification you want to edit.

   There are different Notifications records for the different sources of alert. If you want to change the recipients for all types of alerts, you must edit each of the Notification records individually.
3. Open the **Who will receive** tab.
4. Select the unlock icon for the type of recipient you want to add or change.

   To quickly add yourself to the list of recipients, select the Add me icon next to the *Users* field.
5. Select the users and groups you want to receive email alerts.

   If you select **Send to event creator**, the owner of the AI agent or agentic workflow that triggered the alert receives a notification.

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

You can select **Exclude delegates** if you don't want delegated users to receive the notification. By default, delegates will receive notifications.

6. Update or save the record.

Result
The specified users or groups can now receive notifications for a specific triggered condition of AI agent or agentic workflow executions.

What to do next
Repeat for any other AI agent alert notifications you want to configure.

Role masking in Now Assist AI agents

Role masking for AI agents and agentic workflows helps users enhance security by enabling them to limit their roles during tool execution and verify that AI agents run with least-access privileges.

Role masking overview

Role masking lets AI admins in the AI Agent Studio to limit permissions of agentic workflows or AI agents set to run as dynamic users by defining an allow-list of roles they can inherit from invoking users, enforcing least-privilege access.

Use role masking to:

- Empower users to follow least-access principles when an agentic workflow or AI agent executes on behalf of a dynamic user.
- Limit roles that agentic workflows, AI agents, and skills inherit from users and can be applied when a user invokes them.

  For more information about configuring skills, see Now Assist Skill Kit.
- Reduce the risk of AI solutions accessing the resources that they shouldn't, therefore helping prevent agentic overreach into sensitive data or capabilities beyond what the user is allowed.
- Expand the security configurations to increase functionality of agentic products while reducing security risks by enforcing governance for elevated or scoped roles.

Prerequisites

To configure role masking on your ServiceNow instance, you must have:

- Now Assist for Platform version 10.0.2-SS.
- The sn_aia.admin privileges.

Role masking behavior

Role masking behavior in agentic workflows, AI agents, and tools controls which of the invoking user's roles are available to AI agents and tools during workflow execution. The roles are applied sequentially across layers in the agentic workflow, AI agent, and Tool sequence to verify that tools execute with minimum required privileges.

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

Role masking rules

1. Role masking limits the roles with which an agentic workflow, AI agent or Skill can execute to the intersection between the roles assigned to the invoking user and the roles included in the role masking approved roles list.
2. AI user vs Role mask:

   The AI admin can choose for the component to run as either an AI user or a dynamic user. If set to run as a dynamic user, the AI admin can configure role masking for the component. Role masking cannot be configured for agentic workflows or AI agents set to run as AI users.

   - If an AI user is selected, all roles assigned to the AI user are available to the agentic workflow or AI agent. This can be used to provide elevated access to the agentic workflow AI agent.

     Note: Tools always run as dynamic users.

   - If Role masking is applied to an agentic workflow, AI agent, or tool running as a dynamic user, the component runs with roles with roles limited to the intersection of the current invoking user's roles and the roles included in the role masking approved roles list.
3. Multiple role masks can be configured and applied in the agentic workflow - AI agent - tool sequence, but each mask still follows the intersection rule.
4. Role masking only restricts roles that components can execute with, but never grants components roles which exceed those of the invoking user. Thus, if the invoking user has a role that is not included in the approved roles list, the component will not be allowed to execute with that role. And, if the approved roles list includes a role that is not assigned to the invoking user, the component will not be able to execute with that role again.
5. When a workflow running as an AI user invokes an AI agent running as a dynamic user with role masking configured, then the roles available to the AI agent will be the intersection of the role masking configuration of the AI agent and the roles assigned to the AI user of the agentic workflow. The same holds true for an AI agent running as an AI user that then invokes a skill with role masking configured.

ACL and role masking evaluation sequence

The following sequence defines how ACLs and role masks are evaluated across the agentic workflow, AI agent, and tool execution contexts:

Step 1: Agentic workflow ACL validation

ACLs configured for workflows are evaluated against the invoking user (automated or conversational) session roles.

Step 2: Agentic workflow role mask application

If the invoking user meets the agentic workflow's ACL criteria (and the agentic workflow is set to run as a dynamic user with role masking configured), the agentic workflow role masking is applied to the invoking user's roles (there by restricting roles from the user session based on the intersection with the configured role masking).

Step 3: AI agent ACL validation

When an agentic workflow invokes an AI agent, the AI agents' ACLs are validated against one AI agent ACLs are validated against one of the roles with which the agentic workflow was approved to execute. Thus:

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

- If the agentic workflow was set to run as an AI user, the AI agent ACL will validate against the AI user session configured at the workflow.
- If the agentic workflow was set to run as a dynamic user with role masking, the AI agent ACL will check whether the effective remaining roles after applying the workflow role masking meets the ACL criteria.

Step 4: AI agent role masking application

Similar to the agentic workflow above, either the AI user or the AI agent role mask is applied:

- If an AI user is selected, all roles of the AI user are enforced (no masking).
- If role mask is applied, then the roles are limited further based on intersection with the effective roles after applying the workflow role masking.

Step 5: Tool ACL validation

If a tool uses ACLs, these are checked against the roles that the AI agent—assigned to the tool—is permitted to use. This means that if role masking is set up, only the roles left after masking are considered during validation.

Step 6: Tool role masking application

If the tool is a skill and has role mask configured, then the approved roles will be applied to roles with which the AI agent was approved to run, thereby limiting roles for the tool's execution.

Summary of ACL and role masking evaluation order:

1. **Agentic workflow ACLs** → validated with invoking (conversational or automated) user session's roles.
2. **Agentic workflow role masking** → applied to the invoking user session.
3. **AI agent ACLs** → validated with agentic workflow's approved roles (agentic workflow's AI user OR roles after workflow role masking).
4. **AI agent role masking** → applied to agentic workflow's approved roles.
5. **Tool ACLs** → validated with AI agent's approved roles (AI agent's AI user or roles after agent role masking).
6. **Tool role mask (Skills only)** → applied to AI agent's approved roles.

Note: When evaluating ACLs and role masks, the admin can identify where and why execution failed due to either ACL or role mask restrictions.

Configuration

- To configure role masking for an AI agent, see Define security controls for an AI agent.
- To configure role masking for an AI agent, see Define security controls for an agentic workflow.

Create an AI agent

Create an AI agent in AI Agent Studio to solve problems for your users and coordinate with other AI agents while executing the agentic workflows.

Before you begin
Role required: sn_aia.admin

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

About this task

In the ServiceNow agentic ecosystem, an AI agent is a set of large language model (LLM) instructions and tools that can perform specific tasks.

An AI agent can collaborate with other agents to achieve better results by using fewer LLM calls. AI agents can also reach out to the user if they need any help or information.

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **AI agents**.
2. From the **Add** drop-down menu, select one of the following agent creation options:**Chat** or **External**.

   a. **Chat**: Select the Chat option to create AI agents in the AI Agent Studio to connect with internal AI agents within the ServiceNow AI Platform to execute agentic workflows.

      The following instructions are for Chat AI agents. See the following links for guides to create other types of AI agents.

   b. **Voice**: Select the Voice option to create an AI voice agent.

      See Deploy AI agents for voice for more information on setting up and creating AI voice agents.

   c. **External**: Select the External option to create external AI agents in the AI Agent Studio to connect the ServiceNow AI Platform with the third-party agentic AI providers.

      For more information about creating external AI agents in the AI Agent Studio, see Create an external AI agent.

3. Define the AI agent's specialty.
4. Define the AI agent security controls.
5. **Optional:** Add a trigger to automatically invoke your AI agent if a specified event occurs..

   (Optional) If you want your AI agent to be used only in chats, you don't need to add a trigger.

6. Determine the chat assistants to access your AI agent, set the processing messages, and activate your AI agent..
7. Select **Save and test** to complete the configuration steps or review a previous step by selecting **Back**.

   Selecting Save and test leads you to the AI agent testing page, where you can test the AI agent that you created. For more information, see Manually test the execution of an AI agent.

   To test the AI agent, you must have the sn_aia.admin role and any roles the ACLs configured for the AI agent and its tools require, if applicable.

Result

You can see the AI agent that you created in the **Manage agentic workflows and AI agents** page.

If you selected **Users with specific roles** for your ACL, you can change roles at any point. Return to the guided setup and double-click the pills to make changes. All other security changes must be done on the ACL [sys_security_acl] table by a user with elevated permissions.

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

Find AI agents

Find available AI agents in AI Agent Studio to explore your options when creating agentic workflows.

Before you begin

Role required: sn_aia_viewer.

About this task

AI agents are autonomous systems that interact with their environment to gather data, make decisions, and complete tasks that would otherwise need to be done by a human.

Procedure

1. View the available AI agents in AI Agent Studio.

Current location	Navigation option
AI Agent Studio Overview page	Select the **Explore all** button in the Ready-made agentic workflow and AI agents section, and then select the **AI agents** tab.
AI Agent Studio Overview page	View the Recent agentic workflows and AI agents activity section on the **AI agents** tab to see the most recently added or changed agents. **Note:** You see a list of the AI agents on ly when there's recent activity of the AI agents on your instance.
Anywhere else	Navigate to **All** > **AI Agent Studio** > **Create and manage**, and then select the **AI agents** tab.

2. View all AI agents, including those AI agents that were installed with Now Assist applications by navigating to the **AI agents** tab.

Result

Navigating through any of the preceding options directs you to the AI agents that are available on your AI Agent Studio instance.

Define the specialty of an AI agent

In the guided setup for an AI agent, write a clear description defining your agent and its role. You can also configure supported LLMs, enable third-party access, and manage long-term memory.

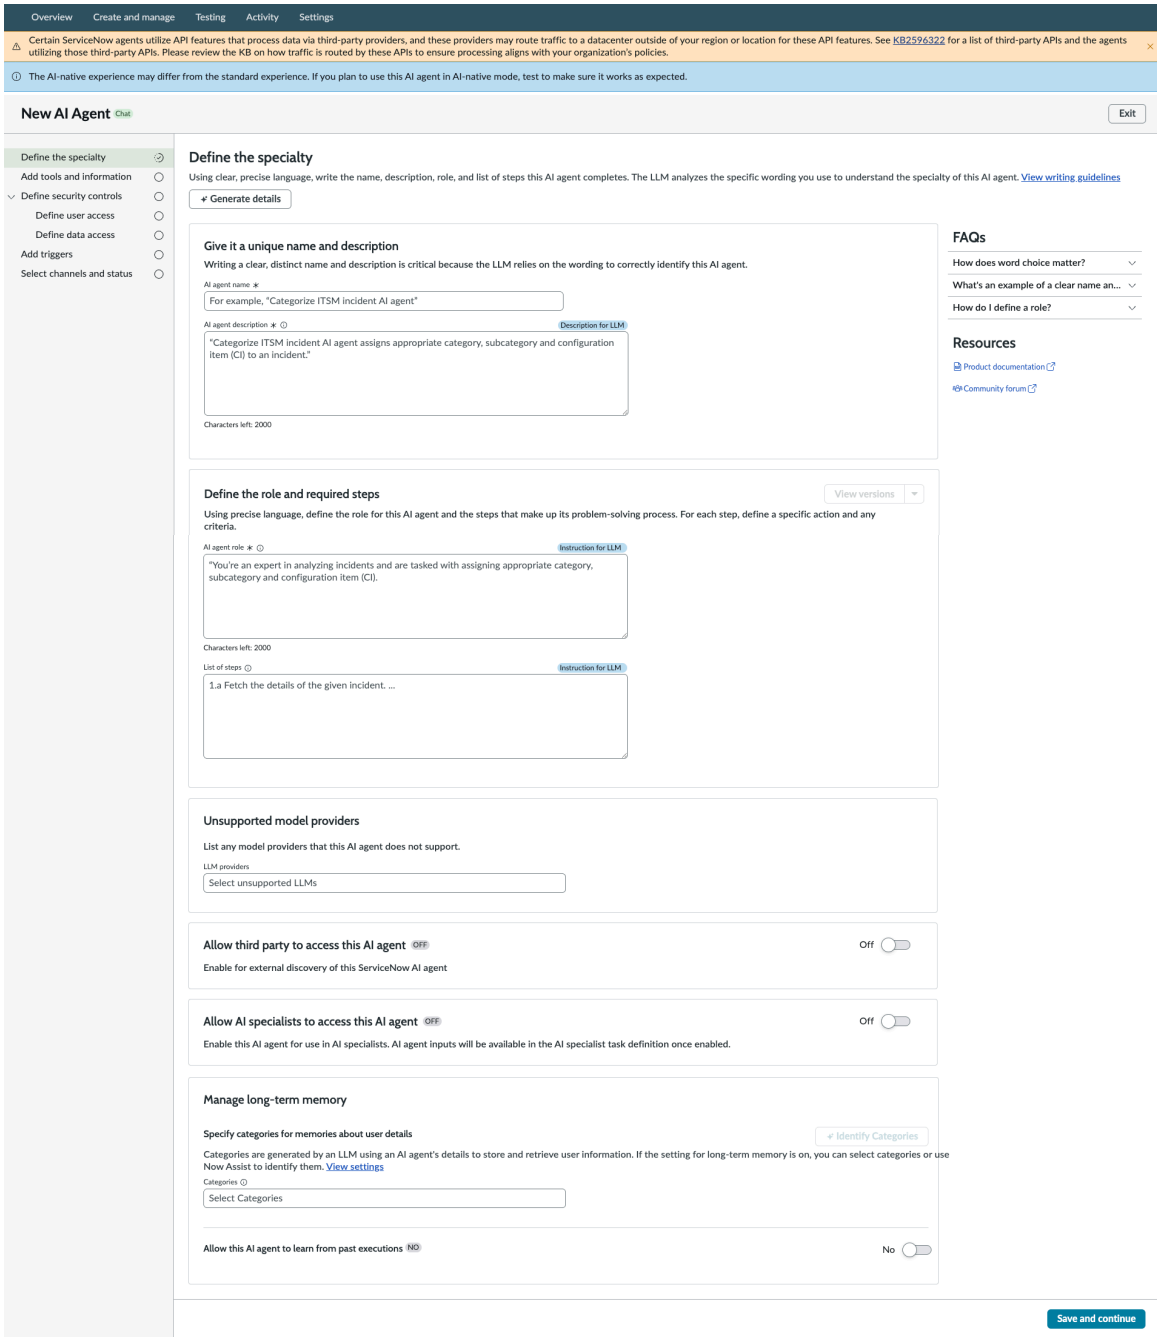
Before you begin

Role required: sn_aia.admin

About this task

The first step of the guided setup enables you to define the fundamentals of the AI agent. The description, AI agent role, and list of steps fields are used by the LLM to understand how to use the AI agent, by itself or as part of an agentic workflow. Descriptions, AI agent roles, and list of steps should be clear and well-defined. For guidelines for writing these fields, see Writing effectively for agentic AI. For an example AI agent, see Example AI agent.

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



Note: If you have the Off Glide Conversation Server plugin (com.glide.cs.offglide) installed on your ServiceNow instance, then the AI-native experience may differ from the standard experience. If you plan to use an AI agent in the AI-native mode, make sure to test so it works as expected.

Procedure

1. Enter a name for the AI agent.
2. Craft an AI agent description.

   The description should clarify the purpose of the AI agent, including its inputs, outputs, and context.

3. Craft the AI agent role.

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

The difference between the description and the agent role is that the agent role provides more detail about the function that the AI agent serves within a greater context. The greater context may be as part of an agentic workflow, or it could be the specific function it serves in another type of work process.

4. Craft a list of steps.
   You can also use an old version of a list of steps. Select **View versions** for a full list of previous versions of the field. See Version control for more information.
5. Determine unsupported model providers.

   You may find that some models use your AI agent better than others. If you want to avoid users from having a less optimal experience with a certain model, you can mark a model as unsupported. If an administrator changes the instance-level settings to an unsupported model, your AI agent won't be available to users.
6. Determine if the AI agent can be accessed by third parties.

   Make your AI agent discoverable by turning on the **Allow third-party to access this AI agent** toggle button that enables you to use your ServiceNow AI agent on other platforms. You can still use it on your ServiceNow instances as well.

   See ServiceNow agents as secondary agents for more information about configuring agents to use on other AI platforms.
7. Configure long-term memory options.

   Long-term memory for AI agents enables agents to access data outside of the specific conversation or trigger that executes it. For example, you can select the *`Device and Software`* category to give the AI agent access to user information about the kinds of hardware devices and software that are available to the invoking user. The AI agent can remember this information in later conversations.

   - For more information about creating LTM categories, see Create long-term memory category.
   - For more information about selecting the categories to map them to you AI agent, see Map Long-term memory category.
8. Configure AI agent learning.

   You can enable the AI agent to learn from past interactions with the user. For example, if the user encounters a similar issue, enabling this feature could help the AI agent understand what didn't work previously, so it can recommend other options.

   Note: When the agent learning capability is inactive at the AI Agent Studio level, the option to enable at the AI agent level is also inactive. To enable agent learning at the AI Agent Studio level, see Set up long-term memory.

Result
You have defined the basics of your AI agent.

What to do next
Select **Save and continue** to move to the next step, Add tools and information.

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

Add tools and information to an AI agent

Add a tool to an AI agent to enable different functionalities and help your AI agents achieve their objectives.

Tool overview

Tools provide your AI agents with the capabilities necessary to complete their tasks. When adding tools, consider how the AI agent uses them to achieve its objectives as well as how those tools interact with each other. Providing your AI agents with the appropriate tools help promote the robustness and quality of their performance.

Some additional guidelines for adding and creating tools are the following:

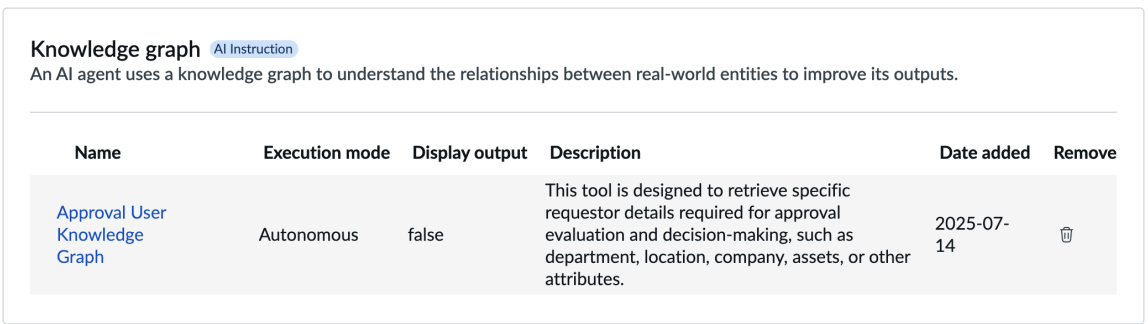
- Design tools to work together. Your AI agent should solve a specific, discrete task, and the tools should give the AI agent the capabilities necessary to achieve its goal.
- Write detailed tool descriptions. Tool descriptions are used by the AI agent to determine a tool's function. Thorough descriptions about what the tools are and how they work give the AI agent the best chance to succeed.
- Consider the outputs of tools when creating them. Use the tool description or output transformation strategy fields to describe how to process the tool outputs. For example, if you have a tool that gathers records from a large number of tables, provide the tool with plans for how to handle the large number of records.

Note: If you select Google as your web search tool provider, the web search tool leverages Grounding with Google Search, offered under a Global Standard deployment. Because grounding isn't data resident, Google's global infrastructure routes traffic to a global datacenter for each web search request. This processing may be different than your data processing location chosen for your ServiceNow instance. Consider your organization's data policies before enabling AI agents that use Google web search tools.

Once you have added the tools to your AI agent, you can select **Save and continue** to move to the next step in the guided setup.

Knowledge graphs

In this step, you can also add Knowledge graphs. Knowledge graphs give the AI agent information to understand the relationships between real-world entities to improve its outputs. For example, you could add a Knowledge Graph to an approval AI agent that maps users to their location, company, and department to help the AI agent understand the specific approval process to apply.



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

Supervised execution mode for AI agents

You can minimize the potential negative impact of an AI agent not executing as expected by configuring AI agents' tools to run in supervised mode. Running in supervised mode means that the tools use human oversight when executing actions. You can use the Supervised mode to enhance security for agents with the capability to perform sensitive or critical actions.

You can set the supervised execution mode when creating a tool in the AI agent guided setup. For example, choose Supervised as the Execution mode when adding a catalog item tool. For reference, see Add a catalog item to an AI agent.

Add a catalog item to an AI agent

Add a Service Catalog to an AI agent in AI Agent Studio so that your users can access conversational catalog items.

Before you begin

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > *AI Agent Studio* > **Create and manage** > **AI agents**.
2. Open the AI agent that you want to add a catalog item to and navigate to the Add tools and information section.
3. In the Add tool drop-down list, select **Catalog item**.
4. On the form, fill in the fields.

Add a catalog item form

Fields	Description
Name	Name that you want to specify for your catalog item.
Description	Description of the catalog item and what it's going to do to assist your AI agent. **Note:** This description is sent to the large language model (LLM).
Select catalog item	Catalog item that you want to add. If there is a tool used by an AI agent for the same catalog item, the rest of the form will populate the fields based on the other tool. You can make any changes specific to the current AI agent to suit your needs, and it will not affect the tool of the existing AI agent.
Execution mode	Mode of execution for your selected catalog item:

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

Fields	Description
	○ **Supervised**: Inputs from your human agent are required during the execution of this catalog item while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this catalog item while the AI agent runs.
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No** If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that can be used in assistants built with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down. ○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this transformation step, the tool output will be directly mapped to the widget. ▪ **Yes** ▪ **No** ○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No** **Note:** The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
zz	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats:

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

Fields	Description
	○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution. ○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

5. Select **Add**.
   A catalog item is added in the Catalog items list on the Add tools and information page.

Add a conversational topic to an AI agent

Add a Virtual Agent topic to an AI agent in AI Agent Studio so that you can use conversations to get additional information from the user. For example, a conversational topic could be used to let a user select a date range for surveys.

Before you begin

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > *AI Agent Studio* > **Create and manage** > **AI agents**.
2. Open the AI agent that you want to add a conversational topic to and navigate to the Add tools and information section.
3. In the Add tool drop-down list, select **Conversational topic**.
4. On the form, fill in the fields.

Add a conversational topic form

Fields	Description
Name	Name that you want to specify for your conversational topic.
Description	Description of the conversational topic and what it's going to do to assist your AI agent. **Note:** This description is sent to the large language model (LLM).
Select topic	The Virtual Agent topic that you want to add. After it's selected, the description of the topic displays underneath the drop-down list.

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

Fields	Description
	If there is a tool used by an AI agent for the same conversational topic, the rest of the form will populate the fields based on the other tool. You can make any changes specific to the current AI agent to suit your needs, and it will not affect the tool of the existing AI agent. **Note:** Only large language model (LLM) topics can be selected.
Execution mode	Mode of execution for your selected conversational topic: ○ **Supervised**: Inputs from your human agent are required during the execution of this conversational topic while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this conversational topic while the AI agent runs.
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No** If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that can be used in assistants built with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down. ○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this transformation step, the tool output will be directly mapped to the widget. ▪ **Yes** ▪ **No** ○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No**

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

Fields	Description
	Note: The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
Advanced settings	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats: ○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution. ○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

5. Select **Add**.
   A conversational topic is added in the Conversational topics list on the Add tools and information page.

Add a desktop action to an AI agent

Add a desktop action as a tool to an AI agent in AI Agent Studio so that AI agents can perform desktop actions for repetitive tasks.

Before you begin

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > *AI Agent Studio* > **Create and manage** > **AI agents**.
2. Open the AI agent that you want to add a desktop action to.
   For creating an AI agent, see Create an AI agent.
3. Navigate to the Add tools and information step.
4. In the Add tool drop-down list, select **Desktop action**.
5. On the form, fill in the fields.

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

Add a desktop action form

Fields	Description
Select a type of desktop action	○ **UI block**: These actions include UI interactions, such as clicking buttons, typing into text boxes, selecting from drop-down menus, and more. Example: Filling in fields on a payroll application, updating inventory in a point-of-sale system, or submitting a request through a desktop insurance claims app. ○ **Non UI blocks**: These actions are prebuilt components that don't interact with the UI, but run in the background. Non UI block actions enable your AI agents to interact with various applications and system components and offer prebuilt actions for common tasks, reducing the need for complex scripting. Example: Reading data from Microsoft Excel, extracting emails from Microsoft Outlook, generating a PDF, compressing files into a ZIP, fetching records from a database, or sending a system notification.
Select a desktop action	Desktop action that you want to add to the AI agent from the list. **Note:** This field appears when **UI block** is selected as the type of desktop action.
Select an application	Select an application from the list to use the associated actions. The following applications are supported for non-UI block actions: ○ Microsoft Excel ○ Microsoft Outlook ○ Microsoft Word ○ PDF ○ PowerShell Connector ○ SQL ○ SSH ○ SystemAction **Note:** This field appears when **Non UI block** is selected as the type of desktop action.

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

Fields	Description
Desktop action description	Read-only. Description of the desktop action or application that you selected.
Applications	Read-only. Applications that this desktop action interacts with. **Note:** This field appears when **UI block** is selected as the type of desktop action and a desktop action is selected from the list.
Created by	Read-only. User who created the desktop action. **Note:** This field appears when **UI block** is selected as the type of desktop action and a desktop action is selected from the list.
Last updated	Read-only. Date when this desktop action was last updated. **Note:** This field appears when **UI block** is selected as the type of desktop action and a desktop action is selected from the list.
Inputs	Read-only. Input parameters associated with this desktop action. **Note:** This field appears when a desktop action or application is selected from the list.
Outputs	Read-only. Output parameters associated with this desktop action. **Note:** This field appears when a desktop action or application is selected from the list.
Name	Name that you want to specify for your selected desktop action.
Tool description	Description of the desktop action tool and what it's going to do to assist your AI agent. **Note:** This description is sent to the large language model (LLM).
Execution mode	Mode of execution for your selected desktop action:

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

Fields	Description
	○ **Supervised**: Inputs from your live agent are required during the execution of this desktop action while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this desktop action while the AI agent runs.
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No** If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that can be used in assistants built with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down. ○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this transformation step, the tool output will be directly mapped to the widget. ▪ **Yes** ▪ **No** ○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No** **Note:** The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
Advanced settings	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats:

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

Fields	Description
	○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution. ○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

6. Select **Add**.
   The desktop action is added in the Desktop actions list on the Add tools and information page.

Add a file upload to an AI agent

Upload files for analysis by an AI agent in AI Agent Studio to grant your AI agent access to specialized knowledge.

Before you begin

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > *AI Agent Studio* > **Create and manage** > **AI agents**.
2. Open the AI agent that you want to add a file upload to and navigate to the Add tools and information section.
3. In the Add tool drop-down list, select **File upload**.
4. On the form, fill in the fields.

Add file upload form

Fields	Description
Name	Name that you want to specify for your file upload.
Description	Description of the file upload and what it's going to do to assist your AI agent. **Note:** This description is sent to the large language model (LLM).
Execution mode	Mode of execution for your selected file upload:

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

Fields	Description
	○ **Supervised**: Inputs from your human agent are required during the execution of this uploaded file while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this uploaded file while the AI agent runs.
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No** If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that can be used in assistants built with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down. ○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this transformation step, the tool output will be directly mapped to the widget. ▪ **Yes** ▪ **No** ○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No** **Note:** The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
Attachments	For analysis by an AI agent, attach up to 5 files with a maximum size of 5 MB each in these file formats: PDF, DOCX, or TXT formats.

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

Fields	Description
	Note: A user who interacts with the AI agent with the file upload tool can see the information contained within the files.
Advanced settings	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats: ○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution. ○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

5. Select **Add**.
   A file upload is added to the AI agent.

Add a flow action to an AI agent

Add a flow action to an AI agent in AI Agent Studio. Define the flow action to use it as a reusable operation in automating the ServiceNow AI Platform features without having to write code.

Before you begin

When an AI agent uses a flow action tool, the user the AI agent is running as must pass the ACL of the flow action. Ensure that the security configurations for the flow action are met by the AI agent and agentic workflow. For more information, see Security for AI agents.

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > *AI Agent Studio* > **Create and manage** > **AI agents**.
2. Open the AI agent that you want to add a flow action to and navigate to the Add tools and information section.
3. In the Add tool drop-down list, select **Flow action**.
4. On the form, fill in the fields.

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

Add a flow action form

Fields	Description
Name	Name that you want to specify for your selected flow action.
Description	Description of the flow action and what it's going to do to assist your AI agent. **Note:** This description is sent to the large language model (LLM).
Select flow action	Existing automated process to be selected from the list. If there is a tool used by an AI agent for the same flow action, the rest of the form will populate the fields based on the other tool. You can make any changes specific to the current AI agent to suit your needs, and it will not affect the tool of the existing AI agent.
Inputs	Flow action inputs. The values are filled in by the LLM at runtime unless you specify a value override. **Note:** If the agent uses multiple tools, you can choose to use another tool's output as an input value override. Select the data picker icon ( ) to review the available options.
Execution mode	Mode of execution for your selected flow action: ○ **Supervised**: Inputs from your human agent are required during the execution of this flow action while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this flow action while the AI agent runs.
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No** If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that can be used in assistants built

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

Fields	Description
	with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down. ○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this transformation step, the tool output will be directly mapped to the widget. ▪ **Yes** ▪ **No** ○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No** **Note:** The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
Advanced settings	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats: ○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution. ○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

5. Select **Add**.
   A flow action is added in the Flow actions list on the Add tools and information page.

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

Add a Knowledge Graph to an AI agent

Add a Knowledge Graph to an AI agent in AI Agent Studio that uses the structured and unstructured data from different ServiceNow records to enhance the performance of AI agents.

Before you begin

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > *AI Agent Studio* > **Create and manage** > **AI agents**.
2. Select **New** or open the AI agent that you want to add a Knowledge Graph to and navigate to the Add tools and information section.
3. In the Add tool drop-down list, select **Knowledge graph**.
   You may need to scroll down.
4. On the form, fill in the fields.

Add a Knowledge Graph form

Fields	Description
Name	Name that you want to specify for your selected Knowledge Graph.
Description	Description of the Knowledge Graph and what it's going to do to assist your AI agent. **Note:** This description is sent to the large language model (LLM).
Select Knowledge Graph	Existing Knowledge Graph to be added to the AI agent. If there is a tool used by an AI agent for the same Knowledge Graph, the rest of the form will populate the fields based on the other tool. You can make any changes specific to the current AI agent to suit your needs, and it will not affect the tool of the existing AI agent. You can utilize Enterprise Graph and Enterprise Graph (Small) as Knowledge Graph resources while creating the Knowledge Graph tool.
Query instruction	The search query. Translate your request into a search query, including a verb. The query instruction is passed on to the LLM to generate a structured query for the Graph from the inputs selected in the tool form.
Execution mode	Mode of execution for your selected Knowledge Graph:

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

Fields	Description
	○ **Supervised**: Inputs from your human agent are required during the execution of this Knowledge Graph while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this Knowledge Graph while the AI agent runs.
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No** If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that can be used in assistants built with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down. ○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this transformation step, the tool output will be directly mapped to the widget. ▪ **Yes** ▪ **No** ○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No** **Note:** The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
Advanced settings	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats:

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

Fields	Description
	○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution. ○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

5. Select **Add**.
   The Knowledge Graph tool is added to the AI agent.

Add a Now Assist skill to an AI agent

Add a generative AI skill to an AI agent in AI Agent Studio. You can customize the skills to meet the needs of your users in different workflows.

Before you begin

If you want to add a custom skill to an AI agent, the skill must be published and activated on the Now Assist Admin console. For more information on deploying custom skills, see Finalize and publish a custom skill and Activate a custom skill.

When an AI agent uses a skill as a tool, the user the AI agent is running as must pass the ACL of the skill. Ensure that the security configurations for the skill are met by the AI agent and agentic workflow. For more information on setting skill-level ACLs, see Configure access control lists for a skill.

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > *AI Agent Studio* > **Create and manage** > **AI agents**.
2. Open the AI agent that you want to add a skill to and navigate to the Add tools and information section.
3. In the Add tool drop-down list, select **Now Assist skill**.
4. On the form, fill in the fields.

Add a Now Assist skill form

Fields	Description
Name	Name that you want to specify for your skill.
Description	Description of the skill and what it's going to do to assist your AI agent.

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

Fields	Description
	Note: This description is sent to the large language model (LLM).
Select skill	Existing skill to be selected from the list. If there is a tool used by an AI agent for the same skill, the rest of the form will populate the fields based on the other tool. You can make any changes specific to the current AI agent to suit your needs, and it will not affect the tool of the existing AI agent.
Inputs	Skill inputs. The values are filled in by the LLM at runtime unless you specify a value override. **Note:** If the agent uses multiple tools, you can choose to use another tool's output as an input value override. Select the data picker icon () to review the available options.
Execution mode	Mode of execution for your skill: ○ **Supervised**: Inputs from your human agent are required during the execution of this tool while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this tool while the AI agent runs.
Display output	Permission to display the output of the tool execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No** If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that can be used in assistants built with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down. ○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this

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

Fields	Description
	transformation step, the tool output will be directly mapped to the widget. ▪ **Yes** ▪ **No** ○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No** **Note:** The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
Advanced settings	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats: ○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution. ○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

5. Select **Add**.
   A Now Assist skill tool is added in the Skills section on the Add tools and information page.

Add a record operation to an AI agent

Add a record operation to an AI agent in AI Agent Studio to create, update, look up, or delete records.

Before you begin

Role required: sn_aia.admin

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

Procedure

1. Navigate to **All** > *AI Agent Studio* > **Create and manage** > **AI agents**.
2. Open the AI agent that you want to add a record operation to and navigate to the Add tools and information section.
3. In the Add tool drop-down list, select **Record operation**.
4. On the form, fill in the fields.

Add a record operation form

Fields	Description
Select a type of record operation to add	If you select **An existing one**, you can select a record operation tool used by the current AI agent or any other AI agent. You can make changes to the settings of that tool to fit the needs of your specific AI agent.
Name	Name that you want to specify for your record operation.
Description	Description of the record operation and what it's going to do to assist your AI agent. **Note:** This description is sent to the large language model (LLM).
Inputs	○ **Input name**: Name of the input for the LLM to use. ○ **Description**: Description of the input to give the LLM context. ○ **Value override**: Value for the input. If you leave it blank, generative AI fills in the value for you. **Note:** If the agent uses multiple tools, you can choose to use another tool's output as an input value override. Select the data picker icon () to review the available options.
Table	Table that you want to perform the record operation on.
Select operation	Inputs should be referred to as {{inputname}}. Camelcase is not supported. ○ **Create record**: Requires values for fields in the new record. ○ **Delete records**: Requires the conditions to determine which records to delete. ○ **Look up records**: Requires the conditions to determine which records to look up,

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

Fields	Description
	number of results specified, which fields to return, the result order, and the result sort type. ○ **Update records**: Requires the values for fields in the updated record and conditions to determine which records to update.
Execution mode	Mode of execution for your selected record operation: ○ **Supervised**: Inputs from your human agent are required during the execution of this tool while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this tool while the AI agent runs.
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No** If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that can be used in assistants built with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down. ○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this transformation step, the tool output will be directly mapped to the widget. ▪ **Yes** ▪ **No** ○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No**

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

Fields	Description
	Note: The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
Advanced settings	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats: ○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution. ○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

5. Select **Add**.
   A record operation is added in the Record operations list on the Add tools and information page.

Add a script to an AI agent

Create a script to add it to an AI agent in AI Agent Studio. With scripts, you can use the scriptable APIs and back-end integration to support the AI agent.

Before you begin

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > *AI Agent Studio* > **Create and manage** > **AI agents**.
2. Open the AI agent that you want to add a script to and navigate to the Add tools and information section.
3. In the **Add tool** drop-down list, select **Script**.
4. On the form, fill in the fields.

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

Add a script form

Fields	Description
Select a type of record operation to add	If you select **An existing one**, you can select a record operation tool used by the current AI agent or any other AI agent. You can make changes to the settings of that tool to fit the needs of your specific AI agent.
Name	Name that you want to specify for your script.
Description	Description of the script and what it's going to do to assist your AI agent. **Note:** This description is sent to the large language model (LLM).
Script inputs	Input name and description to use in the script. The LLM uses the name and description to determine what value should be used at runtime. Example: Input name is *`task_record_sys_id`* and description is *`sys_id of the Task Record, this is mandatory`*.
Script	JavaScript code that includes an object class or a function for your script. **Note:** For improved security, use GlideRecordSecure instead of GlideRecord and addUserEncodedQuery() instead of addEncodedQuery().
Execution mode	Mode of execution for your script: ○ **Supervised**: Inputs from your human agent are required during the execution of this tool while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this tool while the AI agent runs.
Display output	Permission to display the output of the tool execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No** If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that

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

Fields	Description
	can be used in assistants built with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down. ○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this transformation step, the tool output will be directly mapped to the widget. ▪ **Yes** ▪ **No** ○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No** **Note:** The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
Advanced settings	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats: ○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution. ○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

5. Select **Add**.
   A script tool is added in the Scripts section on the Add tools and information page.

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

Add a search retrieval to an AI agent

Add a search retrieval to an AI agent in AI Agent Studio. Leveraging the Retrieval-Augmented Generation (RAG) enables an AI agent to retrieve and incorporate relevant information from an external source.

Before you begin

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > *AI Agent Studio* > **Create and manage** > **AI agents**.
2. Open the AI agent that you want to add a search retrieval to and navigate to the Add tools and information section.
3. In the Add tool drop-down list, select **Search retrieval**.
4. On the form, fill in the fields.

Add search retrieval form

Fields	Description
Select a type of record operation to add	If you select **An existing one**, you can select a record operation tool used by the current AI agent or any other AI agent. You can make changes to the settings of that tool to fit the needs of your specific AI agent.
Name	Name that you want to specify for your search retrieval tool.
Description	Description of the search retrieval tool and what it's going to do to assist your AI agent. **Note:** This description is sent to the large language model (LLM).
Search profile	Name of the search profile that you want to add to your RAG-based tool from the list. To learn more about a search profile, see Search profiles in AI Search.
Search sources	One or multiple sources that you want to add to your search profile. To learn more about the search sources, see Search profiles in AI Search.
Fields returned	One or multiple fields that you want your search profile to return from the search sources.
Results limit	Number of records to be retrieved from the defined search source. The default value is 10.
Search criteria	Type of search that needs to be used:

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

Fields	Description
	○ **Semantic**: This search is carried out with the logical meaning of the search query. ○ **Keyword**: This search is carried out with the defined keywords. ○ **Hybrid**: This search is a combination of both keyword and semantic searches.
Semantic indexes	Fields on the source table that are indexed for a semantic search. **Note:** Semantic indexed fields are required if the search criteria are semantic or hybrid.
Document matching threshold	Cosine similarity score between 0 and 1 (exclusive). Default value is 0. The higher the number, the more variation in search results.
Execution mode	Mode of execution for your search retrieval tool: ○ **Supervised**: Inputs from your human agent are required during the execution of this tool while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this tool while the AI agent runs.
Display output	Permission to display the output of the tool execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No** If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that can be used in assistants built with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down. ○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this transformation step, the tool output will be directly mapped to the widget. ▪ **Yes** ▪ **No**

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

Fields	Description
	○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No** **Note:** The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
Advanced settings	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats: ○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution. ○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

5. Select **Add**.
   A search retrieval tool is added in the Search retrievals section on the Add tools and information page.

Add a subflow to an AI agent

Add a subflow to an AI agent in AI Agent Studio. Subflows are reusable sequences of processing steps that can be called from within a flow.

Before you begin

When an AI agent uses a subflow tool, the user the AI agent is running as must pass the ACL of the subflow. Ensure that the security configurations for the subflow are met by the AI agent and agentic workflow. For more information, see Security for AI agents.

Role required: sn_aia.admin

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

Procedure

1. Navigate to **All** > *AI Agent Studio* > **Create and manage** > **AI agents**.
2. Open the AI agent that you want to add a subflow to and navigate to the Add tools and information section.
3. In the Add tool drop-down list, select **Subflow**.
4. On the form, fill in the fields.

Add a subflow form

Fields	Description
Name	Name that you want to specify for your selected subflow.
Description	Description of the subflow and what it's going to do to assist your AI agent. **Note:** This description is sent to the large language model (LLM).
Select subflow	Existing process automation capability you want to add. If there is a tool used by an AI agent for the same subflow, the rest of the form will populate the fields based on the other tool. You can make any changes specific to the current AI agent to suit your needs, and it will not affect the tool of the existing AI agent.
Inputs	Subflow inputs. The values are filled in by the LLM at runtime unless you specify a value override. **Note:** If the agent uses multiple tools, you can choose to use another tool's output as an input value override. Select the data picker icon () to review the available options.
Execution mode	Mode of execution for your selected subflow: ○ **Supervised**: Inputs from your human agent are required during the execution of this tool while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this tool while the AI agent runs.
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No**

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

Fields	Description
	If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that can be used in assistants built with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down. ○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this transformation step, the tool output will be directly mapped to the widget. ▪ **Yes** ▪ **No** ○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No** **Note:** The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
Advanced settings	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats: ○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution.

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

Fields	Description
	○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

5. Select **Add**.
   A subflow tool is added in the Subflows section on the Add tools and information page.

Add a web search to an AI agent

Add a web search to an AI agent in AI Agent Studio using a third-party search API such as Microsoft Bing or Google.

Before you begin

Note: If you select Google as your web search tool provider, the web search tool leverages Grounding with Google Search, offered under a Global Standard deployment. Because grounding is not data resident, Google's global infrastructure routes traffic to a global data center for each web search request. This processing may be different than your data processing location chosen for your ServiceNow instance. Please consider your organization's data policies before adding a web search tool with Google as the provider.

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **AI agents**.
2. Open the AI agent that you want to add a web search to and navigate to the Add tools and information section.
3. In the Add tool drop-down list, select **Web search**.
4. On the form, fill in the fields.

Add a web search

Fields	Description
Name	Name that you want to specify for your web search.
Resource	Resource for the web search.
Provider	Third-party search API provider. **Note:** You can configure the default provider for the Web Search API capability on the OneExtend Definition Config table [sys_one_extend_definition_config].
Inputs	Information for the search API to include in the web search. You can select value

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

Fields	Description
	overrides or leave them blank for generative AI to fill in the details for you. ○ **Country**: Country where results come from. ○ **Max tokens**: Maximum number of tokens to include in the response. ○ **Number of results**: Total number of results acquired. ○ **Search query**: Value to search for ○ **Sites or domains**: Websites where you want to search. **Note:** If the agent uses multiple tools, you can choose to use another tool's output as an input value override. Select the data picker icon ( ) to review the available options.
Execution mode	Mode of execution for your selected web search: ○ **Supervised**: Inputs from your human agent are required during the execution of this tool while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this tool while the AI agent runs.
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No** If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that can be used in assistants built with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down. ○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this transformation step, the tool output will be directly mapped to the widget.

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

Fields	Description
	▪ **Yes** ▪ **No** ○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No** **Note:** The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
Advanced settings	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats: ○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution. ○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

5. Select **Add**.
   A web search is added in the Web search list on the Add tools and information page.

Define security controls for an AI agent

In the guided setup for an AI agent, define security controls for who can access the AI agent and what data the AI agent has access to.

Before you begin

Role required: sn_aia.admin

About this task

The **Define security controls** step is divided into two parts: **Define user access** and **Define data access**. The former creates an ACL that determines who can discover or invoke the AI agent. The latter defines the data that the AI agent has access to once it's invoked.

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

See Security for AI agents for more information about creating ACLs and user identities for security for AI agent.

Procedure

1. Select which users can access the AI agent.

   The drop-down menu options are:

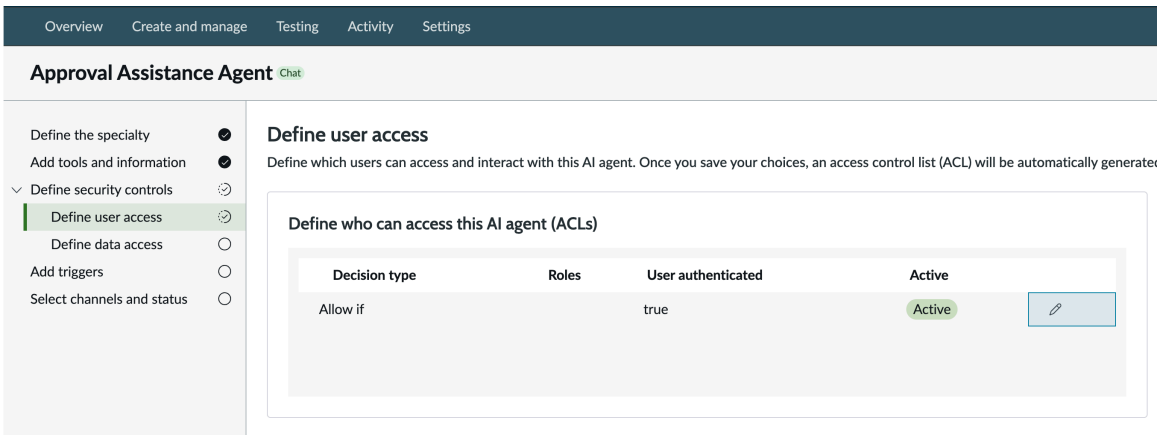
   - Users with specified roles
   - Authenticated users
   - Public

   If you select *Users with specified roles*, you can select exactly which roles can access the AI agent.

   AI agents installed with Now Assist applications each require specific roles. To learn which roles they need, consult the documentation for the AI agent or the agentic workflow that uses the AI agent.

2. Select **Save and continue** to move to the next step.

   Saving and moving onto the next step triggers the creation of an ACL for your AI agent. If you want to make changes later, you can return to the guided setup and change the options here. If you have the correct elevated role, you can also make edits directly on the ACL table.



3. Define the user identity of the AI agent to determine what data it has access to.

   The two options are *Dynamic user* and *AI user*. The dynamic user is the user invoking the AI agent or the dynamic user of the agentic workflow calling on the AI agent. An AI user is a dedicated user that has its own specified roles that allow access, which could be more than the dynamic user.

   If you do not have an AI user but want to use the *AI user* identity, you need to create a new record on the User table. See Create a user. Select *AI user* as the identity type.

   If you select *Dynamic user*, you can select the *Approved roles* that the AI agent runs with. By default, an AI agent runs as a dynamic user and has the roles of the invoking user. Select the approved roles to limit the data access that an AI agent could have. Role masking must be applied for all AI agents and agentic workflows to run as dynamic users.

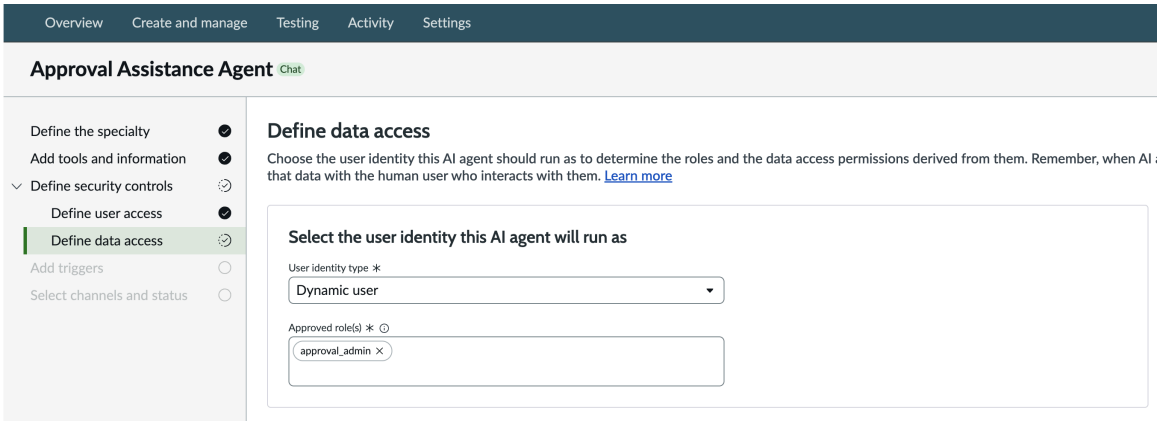
   To override the role masking requirement for a specific agentic workflow or AI agent, admins with the correct elevated access can create an approved list of roles for a given agentic workflow or AI agent. Then, they can access that role masking record in the Agent Access Role

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

Configurations table [sys_agent_access_role_configuration], and select the “allow all roles” check box. Taking these steps deactivates the requirement for a role masking approved roles list in AI Agent Studio, so the AI admin can return to AI Agent Studio and continue to configure the agentic workflow or AI agent without role masking applied.

Note: Role masking should be applied as security best practice and adherence to the principle of least privilege. Overriding the role masking requirement isn’t recommended.

If you select *AI user*, the list of roles that the AI user has is displayed.



Result

You have created an ACL that determines who can discover and access your AI agent, and you have assigned a user identity (and role masking, if relevant) to the AI agent to determine what data it can access.

What to do next

Select **Save and continue** to move to the next step, Adding a trigger. Adding a trigger is optional. You can also skip to the final step, Select channels and access.

Add a trigger to an AI agent

In the guided setup for an AI agent, add triggers to run the AI agent automatically when certain conditions are met.

Before you begin

Role required: sn_aia.admin

About this task

Adding a trigger is optional. If you want your AI agent to be used only in chats, you don't need to add a trigger. Only add a trigger if you want to invoke the AI agent automatically when some event occurs.

If you don't want to add a trigger, skip to the final step, Select channels and access.

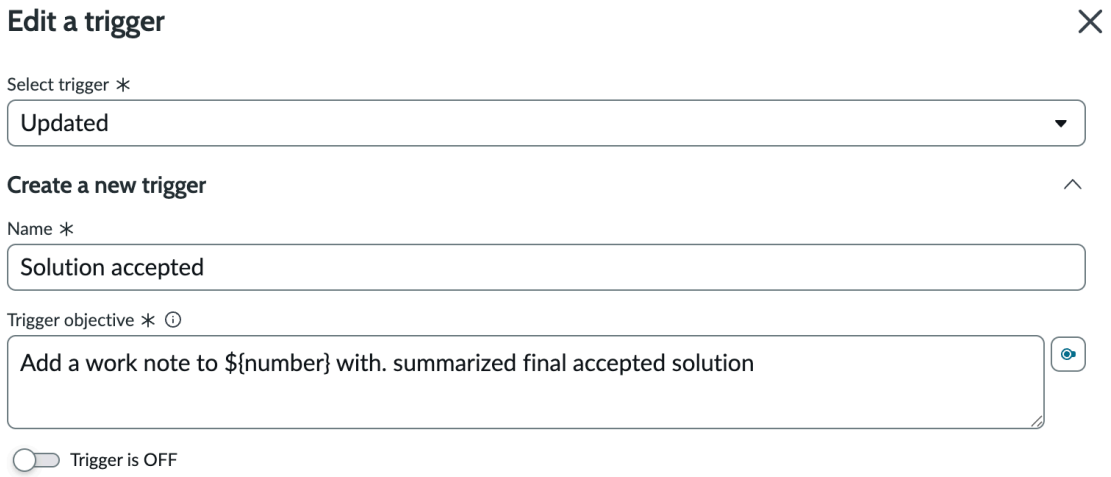
Procedure

1. Select **Add trigger**.
2. On the form, fill in the fields.

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

Create a new trigger

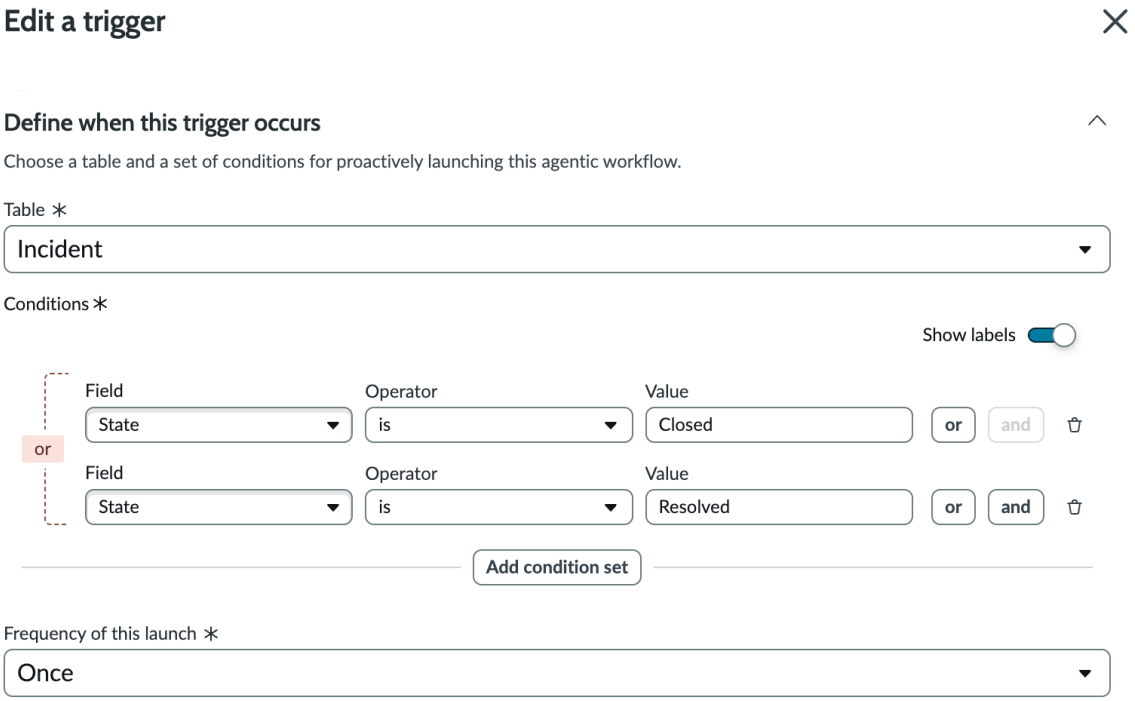
Fields	Description
Select trigger	List of triggers that are available in the instance.
Name	Name of the trigger.
Trigger objective	Additional user statements or sample utterances that help guide when to trigger this agentic workflow.



Define when this trigger occurs

Fields	Description
Table	Name of the applicable table for your agentic workflow.
Conditions	Conditions that you can add to control the trigger configuration. You must have at least one condition. Select **+ Add condition set** to add conditions to your agentic workflow trigger.
Active trigger toggle	Only enable the trigger once you're confident in the execution of your AI agent. Try testing the AI agent execution and user access. To review overall trends over many executions, try an automated evaluation.

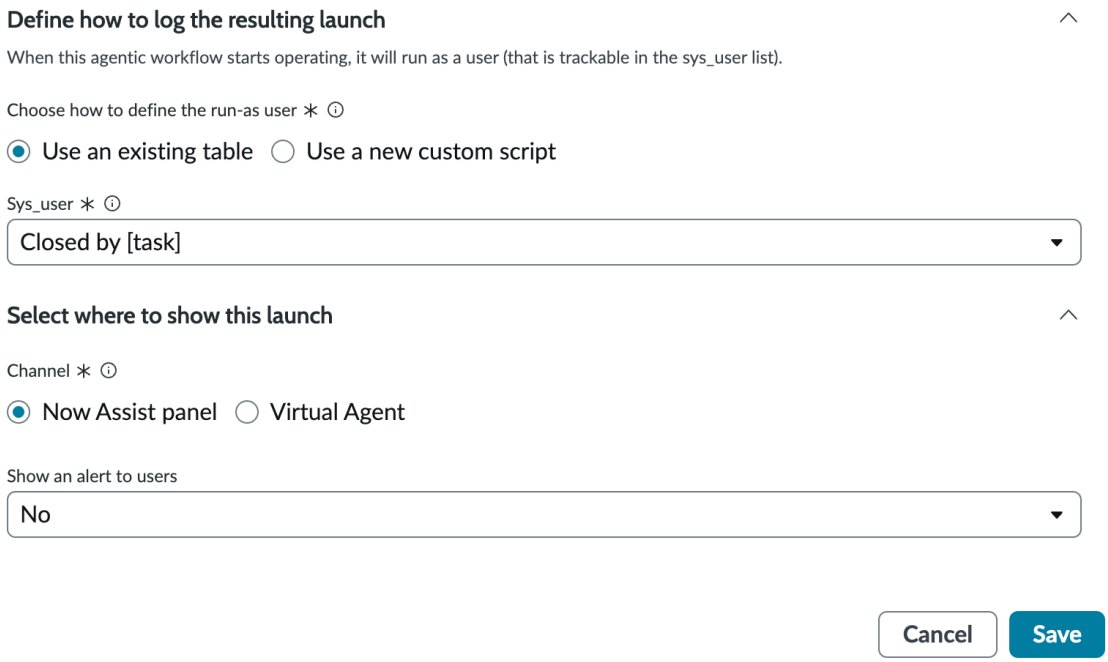
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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



Select where to show this launch

Fields	Description
Channel	Medium for the agentic workflow output: Now Assist panel or Virtual Agent. **Note:** To view the output from a triggered agentic workflow in the Now Assist panel, you need the now_assist_panel_user role.
Show an alert to users	Alerts appear in the selected channel.

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



If you choose a scheduled trigger, additional options are available, such as the day of the week and time when you want the trigger to run.

Note: When running a scheduled trigger, not every record is included in the execution. By default, the value is 10. If you want to change this number, you must set the *sn_aia.max_scheduled_trigger_query* system property to a different value.

If you choose an email trigger, the target emails must exist on the Reply [sys_reply] table. New emails aren't available as triggers.

3. Select **Add**.
4. Repeat the preceding steps for additional triggers.

Result
You have added triggers to your AI agent to run it automatically under the specified conditions.

What to do next
Select **Save and continue** to move to the final step, Select channels and access.

Select channels and status for an AI agent

In the guided setup for an AI agent, activate the AI agent to use in an assistant in Now Assist for Virtual Agent, and set the processing messages.

Before you begin
Role required: sn_aia.admin

About this task
The final step of the AI agent guided setup includes options for where you can invoke the agentic workflow as well as set the processing message. You must select which assistants for Now Assist in Virtual Agent you want to be able to invoke the AI agent. Your processing messages appear when the AI agent is running and when the AI agent's task is completed.

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

Procedure

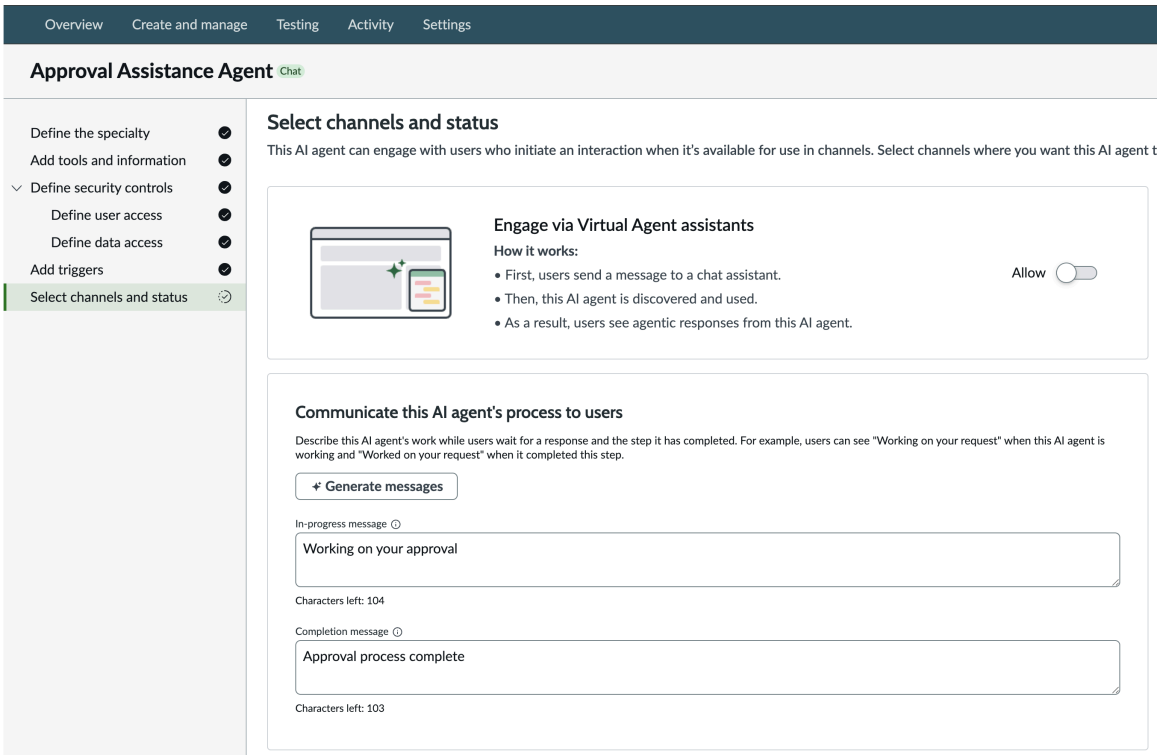
1. Select whether you want users to use Now Assist in Virtual Agent to invoke the AI agent.

   If enabled, you must select which chat assistants have access to the AI agent. You can edit assistants using Assistant Designer.

2. **Optional:** Set the processing and completion messages for your AI agent.

   (Optional) You can set a processing message, completion message, or both. If you don't want to use a specific type of message, unselect the toggle next to the message field.

   You can also use Now Assist to generate the messages for you by selecting **Generate messages**. You can change the messages after they're generated.

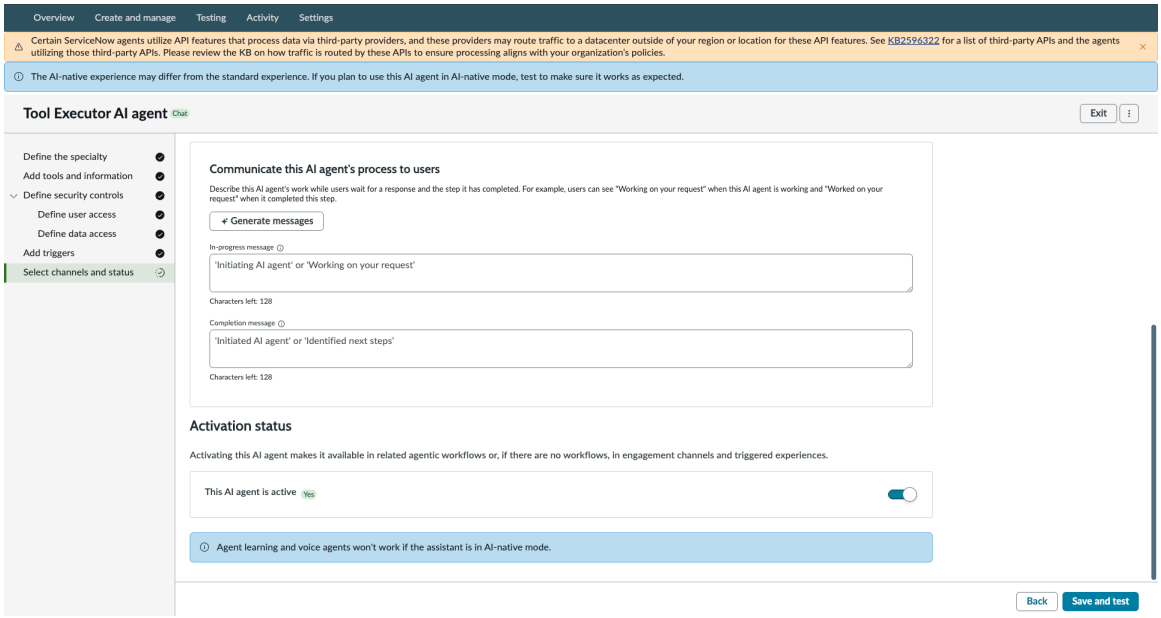


3. Activate the AI agent.

   Select the **This AI agent is active** toggle to activate the AI agent only when you want users to be able to use it. If you want to test your AI agent first, wait until after you have tested it before activating it. You can return to the guided setup later.

   If you don't see this option, you may need to scroll.

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



Note: If you have installed the Off Glide Conversation Server plugin (com.glide.cs.offglide) on your ServiceNow AI Agent Studio instance, then agent learning and voice agents won't work if the assistant is in AI-native mode.

4. Select **Save and test**.

Result

You have completed the guided setup for creating an AI agent. Your new AI agent can be edited at any time using AI Agent Studio.

What to do next

Move to the **Testing** playground to test an AI agent execution using example utterances or to test user access.

Duplicate an AI agent

Duplicate an existing AI agent in AI Agent Studio so that you can save time by not having to manually configure or create AI agents.

Before you begin

Role required: sn_aia.admin

About this task

Duplicate the AI agents to do the following tasks:

- Duplicate the record.
- Disallow any AI agents with existing names.

Custom columns, such as the Tools and Knowledge sources, Status, and a column with the Duplicate icon ( ) are available for the AI agents list.

Note: The duplicated AI Agents use the same tools as the original agent and modifying the agent tools in the duplicated AI agent affects the agent tools in the original AI agent. To use the tools in a duplicated AI agent, you can either use the duplicated agent tools without making changes to them or add a new tool."

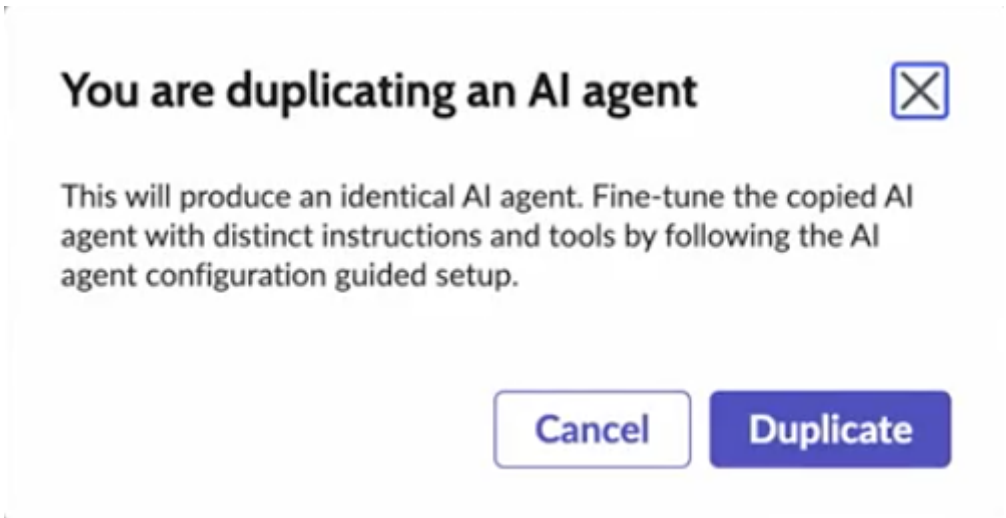
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **AI agents** and open the AI agents list in AI Agent Studio.
2. Duplicate the AI agent from either the Manage agentic workflows and AI agents page or the AI agent form.

Current location	Navigation option
Manage agentic workflows and AI agents page	On the AI agents list, select the duplicate icon ( ) for the AI agent that you would like to duplicate.
AI agent form	Open the AI agent that you want to duplicate, select the menu icon ( ) next to **Exit** on the Describe and instruct form, and select **Duplicate**.

You see a confirmation message in a pop-up window.



3. Create a copy of the AI agent with the same information from the original AI agent's record by selecting **Duplicate**.

   Note: The name of the duplicated AI agent is appended with the suffix `(Copy)` so that you can clearly identify the duplicated AI agent from the original AI agent. For example, `Knowledge Article Agent (Copy)`. You can rename the duplicated AI agent and also edit the other information that is duplicated from the original AI agent.

4. After editing the details in the Describe and instruct section, select **Save and continue**.
5. In the Add tools and information section, edit the details that are from the original AI agent or add new tools for your requirements and select **Save and continue**.
6. In the Toggle display section, turn on the Status for the duplicated AI agent and select **Save and continue**.

Result

You're redirected to the AI Agent Studio home page. In the Manage and configure AI Agents to solve agentic workflow section on the **AI Agents** tab, you see the duplicated AI agent.

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

Modify an AI agent

Modify an AI agent in AI Agent Studio to suit your changing business needs.

Before you begin

Role required: sns_aia.admin

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **AI agents**
2. Select the AI agent that you want to modify.
3. Modify the fields of the different sections of the Guided Setup, making changes to the fields for your requirements.

   The guided setup is the same as the one for creating an AI agent:

   - Define the AI agent's specialty.
   - Define the AI agent security controls.
   - (Optional) Add a trigger to automatically invoke your AI agent if a specified event occurs..
   - Determine whether you want to use the Now Assist in Virtual Agent chat assistants and/or as UI action, set the processing messages, and activate your AI agent..

   Note: Some fields aren't editable if the agent is associated with a Now Assist application. If you want to make more modifications, duplicate the agentic workflow and make changes to the duplicate.

   - For the *List of steps* field in the **Define specialty** step, you can create multiple versions of the same AI agent without losing previous versions. Creating versions enables you to test different instructions to evaluate performance. See Version control for AI agents and agentic workflows for more information.
   - For more guidance on creating effective instructions, see the general guidelines for creating AI agents and agentic workflows.
   - For the access control lists (ACLs), you can edit the security fields and define who can access the AI agent and edit the entity to run the AI agent as a dynamic user or an AI user. For more information, see Implement access control in Now Assist AI agents.

   You can navigate through the steps of the Guided Setup with the **Continue** and **Back** buttons.
4. Navigate to the last step and select **Test** to save your changes and begin testing your modified agent.

Result

Your agent is modified and ready to test.

What to do next

You can test an execution of your AI agent manually or test the user access. Once you've determined that the AI agent has the basic functionality you expect, you can evaluate the AI agent using automated tests.

Manually test the execution of an AI agent

Analyze the execution of an AI agent so you can see that it functions the way that you defined it.

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

Before you begin

Role required: sn_aia_admin and either admin or at least one role required by the ACL of the AI agent being testing and each of the ACLs of its downstream components

About this task

After you create an AI agent, test it to see that it functions the way that you defined it. You can choose to run a manual test with a single instruction to test its basic functioning, or you can evaluate the AI agent's performance across multiple executions.

If you want to test multiple executions using execution logs, you can run an automated test. If you select **Start automated evaluation**, the agentic evaluation guided setup opens in a new browser tab. See Evaluate an AI agent for more details about running a new automated evaluation.

When manually testing performance, you can see how the AI agent interacts with the AI Agent Orchestrator and the Communicator AI agent. The Orchestration is an agent that directs different AI agents, and the Communicator AI agent facilitates the communication between the user and other AI agents.

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Testing**.
2. Select **Start manual test**.

   If you want to start an automated test, see Evaluate an AI agent for more details on that process.

3. In the *Choose a test type* drop-down menu, select *AI agent or workflow*.

   If you want to test user access security controls, see Test AI agent user access.

4. Select an agentic workflow or AI agent that you want to test by searching the name of a workflow or choosing from the drop-down menu.
5. Under *Choose a testing mode*, select the testing mode either as **AI Native** or **Standard**.

   Note: The AI-native testing mode is exclusively accessible when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard testing playground.

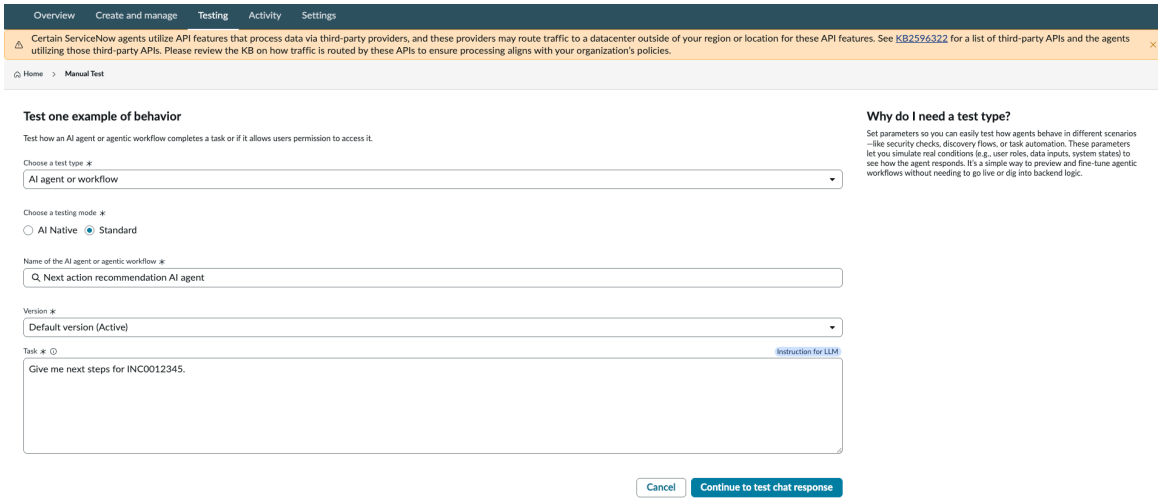
6. In the Version drop-down list, select the version of the AI agent you want to test.

   See Version control for AI agents and agentic workflows for more information about creating and changing versions.

7. In the **Task** field, provide a concise summary of the task to be achieved.

   Note: In the task summary, include a reference number or specific record for better results during your testing.

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



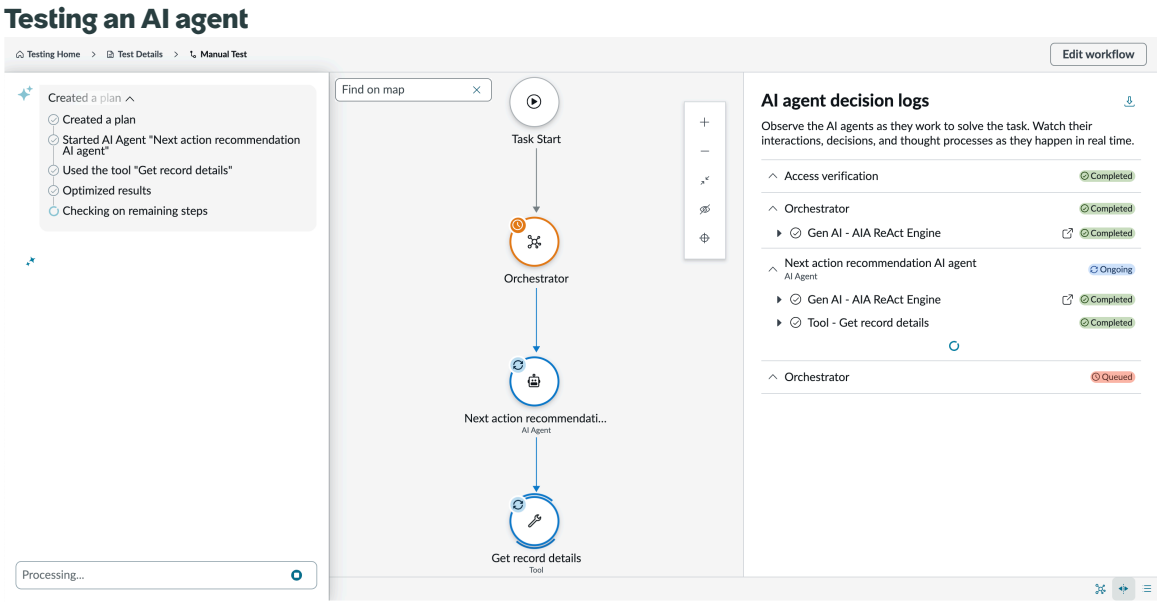
8. Select **Continue to Test Chat Response**.

Result

When testing an AI agent, you can see the AI agent Orchestrator and Communicator working together to organize and manage the AI agents like a team. The AI agent Orchestrator assigns the individual, specialized agents to complete the subtasks. The AI agent Communicator lets the AI agent Orchestrator know the status of each agent.

Note: If you don't have the roles necessary to pass the ACLs of the AI agent and all of its tools, you'll be notified that you don't have the necessary access and the test won't execute.

Your AI agent starts to execute the test autonomously to resolve the task.



- A simulated chat experience begins on the Now Assist panel between your invoking user and AI agent.
- At the top of the canvas, you can see information about the AI agent you're testing, including its name, version, and description.
- A diagram shows the tasks and communication of the AI agents that are working together to solve the case.

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

- A decision log records the thought process of each AI agent that is involved in solving the agentic workflow.

Note: You can view the entire decision log by selecting **Download logs**.

You can restart the entire testing process at any time by selecting **Restart**.

Test user access to an AI agent

Run a manual test to verify that only the users you want to access the AI agent can do so.

Before you begin

Role required: sn_aia_admin and either admin or at least one role required by the ACL of the AI agent being testing and each of the ACLs of its downstream components

About this task

You can establish the security settings for an AI agent in the guided setup to establish which users can access it. See Define security controls for instructions on how to change the user access settings. When you select **Save and continue** on that step of the guided setup, an ACL is created that establishes limitations on which users can access the AI agent.

Once you have created these ACLs, you can verify that they work as intended by using the *Test access* test type of a manual test of an AI agent.

To see instructions for performing manual tests to evaluate performance, see Test performance manually. For more information about automated tests, see Evaluate an AI agent.

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Testing**.
2. Select **Start manual test**.
3. In the *Choose a test type* drop-down menu, select *Test access*.

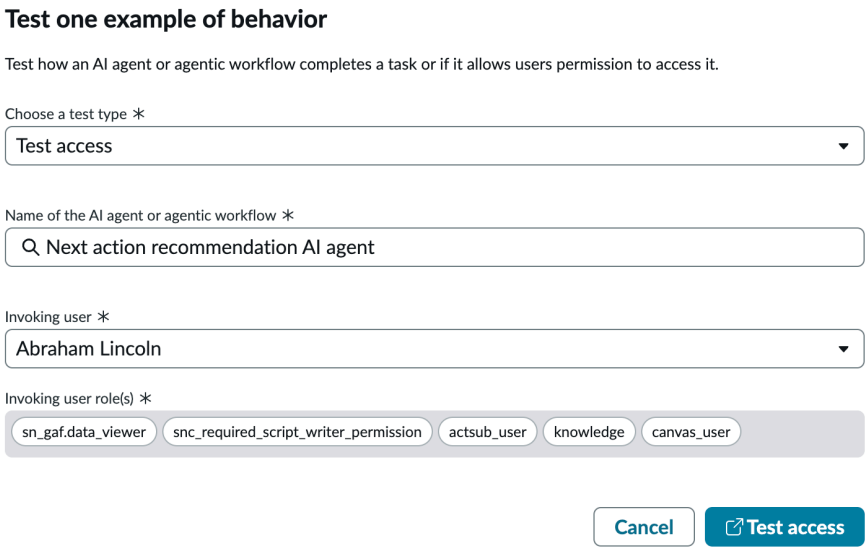
   If you want to test the basics of how an AI agent works, select *AI agent or workflow*. The full instructions for that test type can be found in Test performance manually.
4. Search for or select the name of the AI agent you want to test.
5. Select an invoking user.

   The invoking user can be the user that triggers the AI agent or it can be the invoking user of an upstream component, such as an agentic workflow. For more information about how the invoking user works, see Security for AI agents.

   When you select an invoking user, the user roles are populated in the *Invoking user roles* field. The field is read-only. If you want to change a user's roles, you must change the user's User record. See Assign a role to a user.

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

6. Select **Test access**.



Result
The results of the access test are opened in the Access Analyzer in a new browser tab.

Access Analyzer identifies all the ACL calls made in the execution of the AI agent, including its tools, if any are present. You can see the results of each ACL to verify that each one has the correct user access defined.

What to do next
If the results are different than what you expect or want, you can redefine the security controls of the AI agent.

You can also test an AI agent's performance with either a manual test or automated evaluations. See Test performance manually or Evaluate an AI agent.

Delete an AI agent

Delete an AI agent from AI Agent Studio if you no longer need it.

Before you begin
Role required: sn_aia_admin

About this task
You must assign appropriate permissions by using the access control lists (ACLs) to delete AI agents on AI Agent Studio.

Note: You can't delete the AI agents that come installed with Now Assist applications.

Procedure
1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **AI agents**.
2. Select the check box of the AI agent that you want to delete and select **Delete**.

   You can also delete multiple AI agents by selecting multiple AI agent records at a time.

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

Note: Some AI agents installed with Now Assist applications can't be deleted.

3. In the confirmation pop-up window, select **Delete**.

Result
The selected AI agents are deleted from the AI agents list in AI Agent Studio.

Create an external AI agent

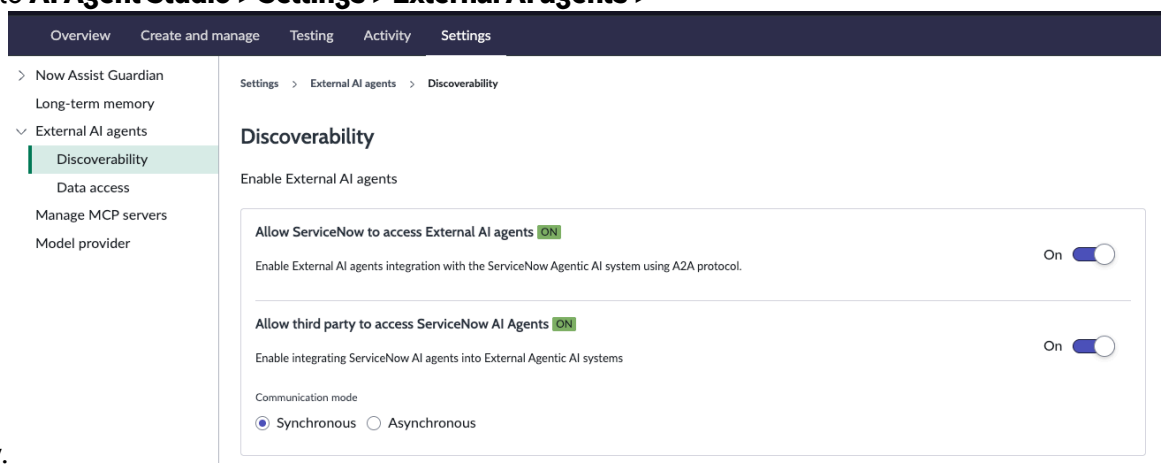
Create external AI agents in AI Agent Studio to connect the ServiceNow AI Platform with third-party agentic AI providers as primary agents.

Integrating external AI agents

Integrate and configure external agents with the ServiceNow agentic AI system using Agent2Agent (A2A) protocol integration to use in agentic workflows created in the AI Agent Studio.

External agent discovery on the ServiceNow AI Platform

You can enable external AI agents on the AI Agent Studio via the Settings page. Navigate to **AI Agent Studio** > **Settings** > **External AI agents** > **Discoverability**.



- **Allow ServiceNow to access External AI agents**: Turn on the toggle to enable external AI agents integration with the ServiceNow agentic AI system using the A2A protocol.
- **Allow third-party to access ServiceNow AI agents**: Turn on the toggle to enables integrating ServiceNow AI agents into external agentic AI systems.

When creating a new external AI agent in AI Agent Studio, you can connect your agent to the ServiceNow agentic AI system using the Agent2Agent (A2A) protocol integration. To do this, you must have a Connection & Credential alias record that connects to your agentic AI provider.

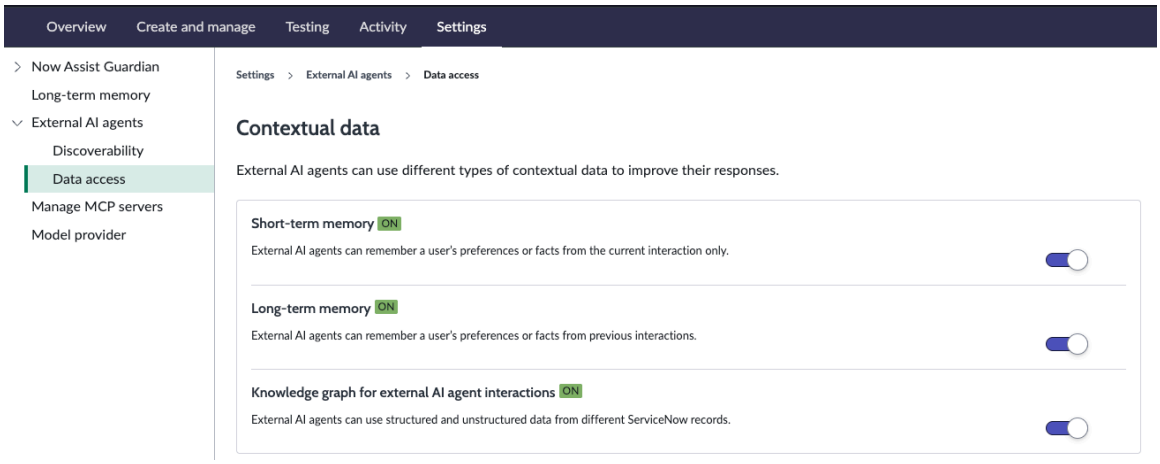
After connecting to the external AI agent, you can add details about its role and instructions to provide context for the AI Agent Orchestrator. Additional context helps differentiate your AI agent from other AI agents so that the AI Agent Orchestrator can decide about which agent to use.

Contextual data access for external agents on the ServiceNow AI Platform

Enable contextual data access for external agents to improve AI agent response during execution.

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

- **Short-term memory**: Turn on the toggle to enable external AI agents remember your preferences or facts from the current interaction only.
- **Long-term memory**: Turn on the toggle to enable external AI agents remember your preferences or facts from previous interaction.
- **Knowledge graph for external AI agent interactions**: Turn on the toggle to enable external AI agents to use structured and unstructured data from different records across the ServiceNow AI Platform.



Agent2Agent protocol overview

Agent2Agent (A2A) is an open standard that enables AI agents across different platforms to communicate with each other. The standard relies on every AI agent having an Agent Card associated with it that provides basic information for providers like the ServiceNow AI Platform, Azure, and Google to use them.

Version 0.3 of A2A is supported.

An AI agent's Agent Card uses standardized JSON to help different providers understand its capabilities. The Agent Card is accessed by a specific type of endpoint from a provider's server. Execution plans are communicated through an execution endpoint so that both the provider's server and the ServiceNow AI Platform can track what the external AI agent is doing.

See Create an external AI agent with the Agent2Agent protocol for instructions for using this protocol to create an AI agent.

Configuring A2A authentication

A2A connections require authorization from the source platform to execute on the ServiceNow AI Platform. Authentication is established by creating two Connection & Credential Alias and Connection records, one with an Agent Card endpoint and one with an execution endpoint.

Configuration Properties

The following properties configure different aspects of how your agents can interact with the ServiceNow AI Platform.

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

External AI agent properties

Name	Description
sn_aia.external_agents.enabled	Set to `true` to enable external agent features
sn_aia.external_agents.parallel_conversations.enabled	Enables or disables multiple simultaneous conversations per user

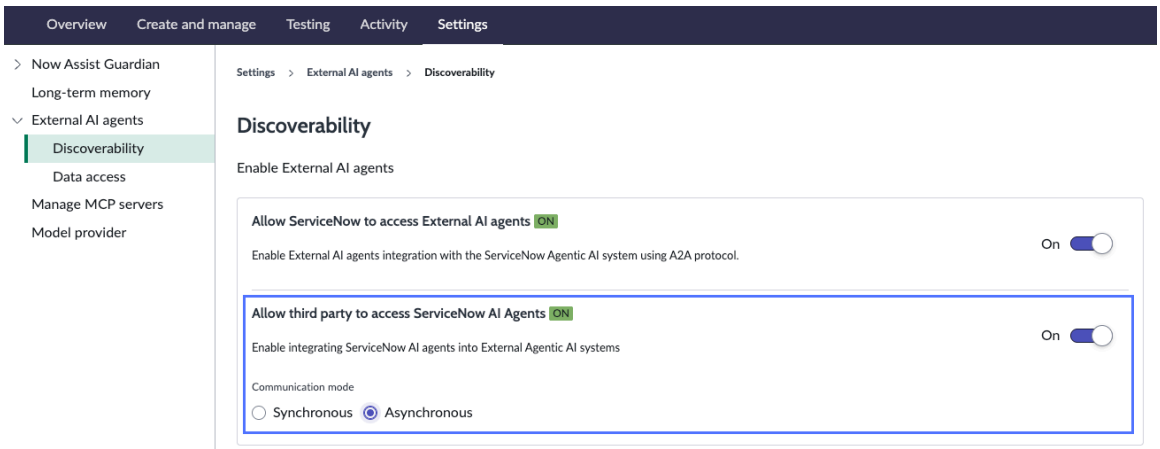
ServiceNow AI agents as secondary agents

Integrate ServiceNow AI agents into other agentic AI systems, such as Google Cloud or Azure OpenAI.

Enabling discovery of ServiceNow mobile agents

In AI Agent Studio, on the **Settings** page, under **External AI Agents** > **Discoverability**, you can enable the discovery of ServiceNow AI agents to use on other AI platforms. To do so, toggle **Allow third party to access ServiceNow AI Agents**.

You can also choose between **Synchronous** and **Asynchronous** communication between your external AI agent and the agentic AI provider.



ServiceNow AI agents as secondary agents overview

You can connect your ServiceNow agent to other agentic AI model providers using the Agent2Agent protocol.

After creating your AI agent in AI Agent Studio, you can point to the Agent Card URL that is displayed for secondary agents for easy access so the admins can view, copy, and consume the URL.

The endpoint to point to for the actual execution of the AI agent is in the format `{{instance}}.service-now.com/api/sn_aia/a2a/v2/agent/id/{{agent-id}}`.

You can use the same OAuth or API key for authenticating the agent discovery and the agent execution.

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

To verify that your AI agent is running from the ServiceNow side, during a conversation with the AI agent, you can go to the *Execution Plan [sn_aia_execution_plan]* table. From the Execution Plan table, you can identify the execution plan based on the *Objective* field that contains the prompt from the conversation on the other platform.

For more information about sample payloads for Google A2A with ServiceNow AI agent as Secondary agent, see Sample payloads for Google A2A.

Asynchronous connection

Asynchronous connection involves communication where the sender and receiver don't have to be active simultaneously unlike the synchronous communication mode.

To establish asynchronous connection, you must obtain a callback URL for push notifications to function.

Once you have obtained a callback URL, you must create a record on the External Agent Callback Registries [sn_aia_external_agent_callback_registry] table. Go to the table, select **New**, and enter the callback URL. Save the record.

Once you save the record, a Connection & Credential Alias [sys_alias] record is created for you. To add authentication, you can open the Connection record associated with the sys_alias record and add a credential to the *Credential* field.

When the record is created, you can go back to the External Agent Callback Registry record you created and select **Verify URL** to test the connection works as expected.

Create an external AI agent with the Agent2Agent protocol

Create external AI agents in AI Agent Studio to connect the ServiceNow AI Platform with third-party agentic AI providers.

Before you begin

If you don’t see the option to create an external agent, confirm that your administrator has selected **Allow ServiceNow to access External AI Agents**. This option is available on the **Settings** page in AI Agent Studio, under **External AI Agents** > **Discoverability**.

Your instance must be at least on Zurich patch 4.

Role required: sn_aia.admin

About this task

You can integrate third-party AI providers into the ServiceNow AI Platform by creating an external AI agent in AI Agent Studio.

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **AI agents**.
2. Select **New** > **External**.
3. In the **Discover and activate** step, select the AI agent provider for your external agent and discover your agent.
   Select **Save and continue** to navigate to the next step.

   a. Select an existing provider or add a new provider.

      If you select **Add new provider**, you must fill out the fields and then select **Save**.

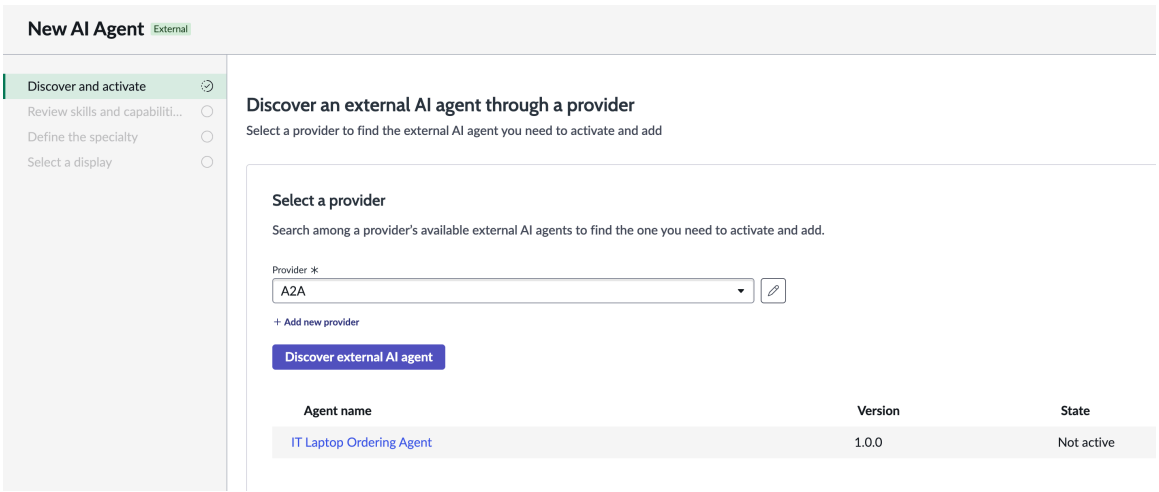
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

Add new provider form fields

Name	Description
Name	Name of the agentic AI provider.
Agent card URL	The URL that points to the external AI agent's Agent Card. The URL should include `well_known/agent_json`.
Under **Advanced settings**, Connection & Credential alias	Credentials to access your external AI agent's Agent Card. You can select an existing alias or create one.
Under **Advanced settings**, Select subflow	Subflow that establishes Agent2Agent protocol. The default subflow should handle the majority of cases, but you can also create your own.

b. Select **Discover external AI agent** to validate the connection to your external agent.

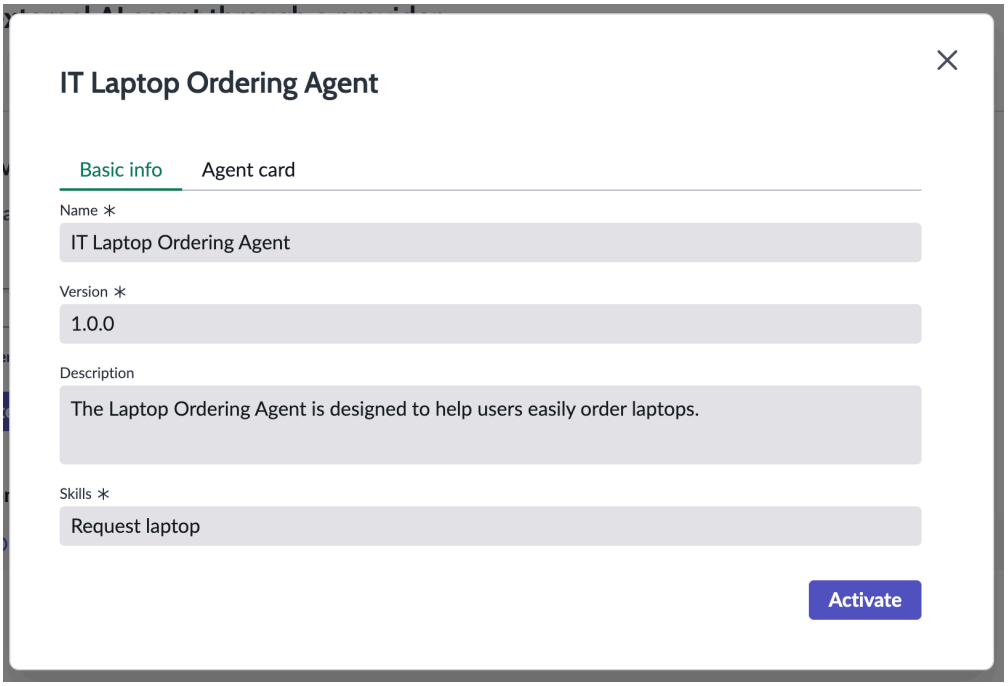
If discovery is successful, the name of your agent and the version number are added to the page below the Discover external AI agent button.



c. Select the name of the agent to verify that the *`Name`*, *`Version`*, *`Description`*, and *`Skills`* fields are populated correctly, then select **Activate**.

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

The **Agent card** tab shows the entire JSON for the agent.



4. In the **Review skills and capabilities** step, verify that the skill names are correct.
5. In the **Define the specialty** step, add details for how the external agent fits in the ServiceNow agentic system.

   a. Review your AI agent description.

      You can leave your AI agent description as it is, or you can add a longer description to help differentiate the agent from other AI agents. This helps enable the AI Agent Orchestrator to use your external AI agent more effectively. For suggestions for writing AI agent descriptions, see General guidelines for creating AI agents and agentic workflows.

   b. Set your communication mode to either **Synchronous** or **Asynchronous**.

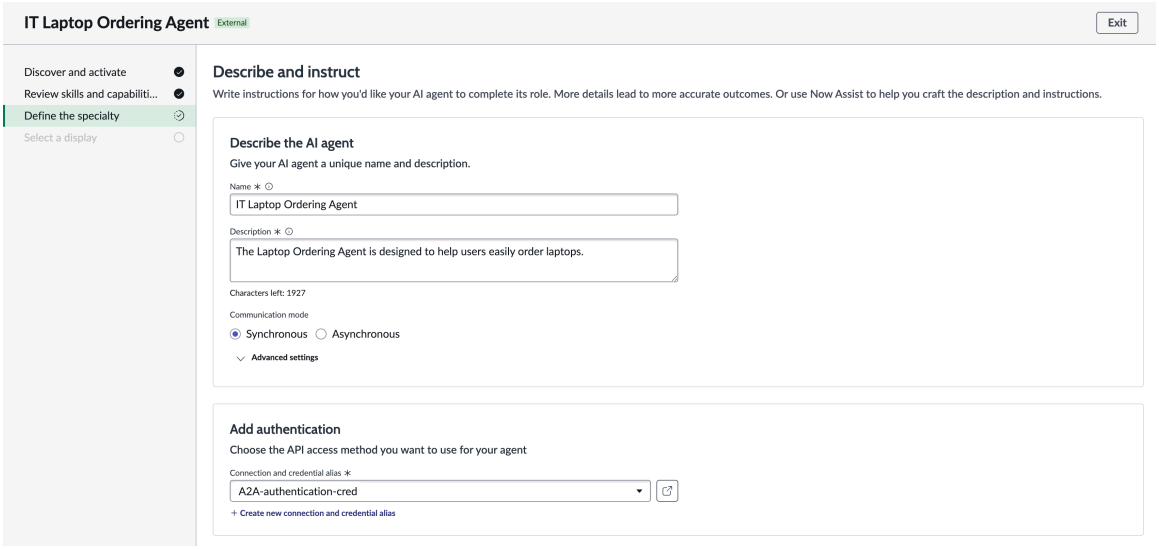
      Some agents don't support asynchronous communication.

      Under **Advanced settings**, you can use the default subflow, an existing subflow, or create a subflow. The default subflow should work for the majority of cases.

   c. Select your AI agent's execution Connection & Credential alias for authentication.

      Choose which credentials access your external AI agent's execution endpoint. You can select an existing alias or create one. If you create one, a modal displays with options for configuring an OAuth or API Key authentication.

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



d. Configure access control lists (ACLs) for the AI agent.

> **Note:** The ACLs determine who has access to discover and execute the AI agent. To learn more about the ACLs you can create in AI Agent Studio and how to add more advanced security configurations, see Implement access control in Now Assist AI agents.

This is a required step. If you have previously configured an AI agent without creating an ACL, you must generate an ACL before you can make other modifications.

Define who can access the AI agent

Field	Description
User access	The type of users whose access for the AI agent is defined by the following options: ▪ **Any authenticated user**: Any user who is logged in can access the AI agent. ▪ **Users with specific roles**: Users that have at least one of the roles assigned to them can access the AI agent. This option is the default. **Note:** If a user doesn't have access to an AI agent or if the user doesn't have access to at least one of the AI agents in the respective agentic workflow execution, then the whole execution aborts before the first AI agent is initiated. ▪ **Public**: Any user can access the AI agent even without logging in. Use this option only when you want guests to be able to access the AI agent.

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

Field	Description
Role	Assign one or more specific roles from the drop-down menu. **Note:** Selecting the role is possible only when you chose the **Users with specific roles** user access.

6. In the **Select a display** step, choose where you want the AI agent to appear and what message users see when the AI agent is running.

   a. Select your AI agent's availability.
   Set the *`Status`* to active if you want the AI agent to be available to users with the correct role.

   b. Select a display channel.

   You can choose to use a Virtual Agent and specify which assistant you want to access the AI agent. If you want to test the AI agent first, you don't need to select a display channel yet.

   c. Choose processing messages to display to the user when the AI agent is executing.

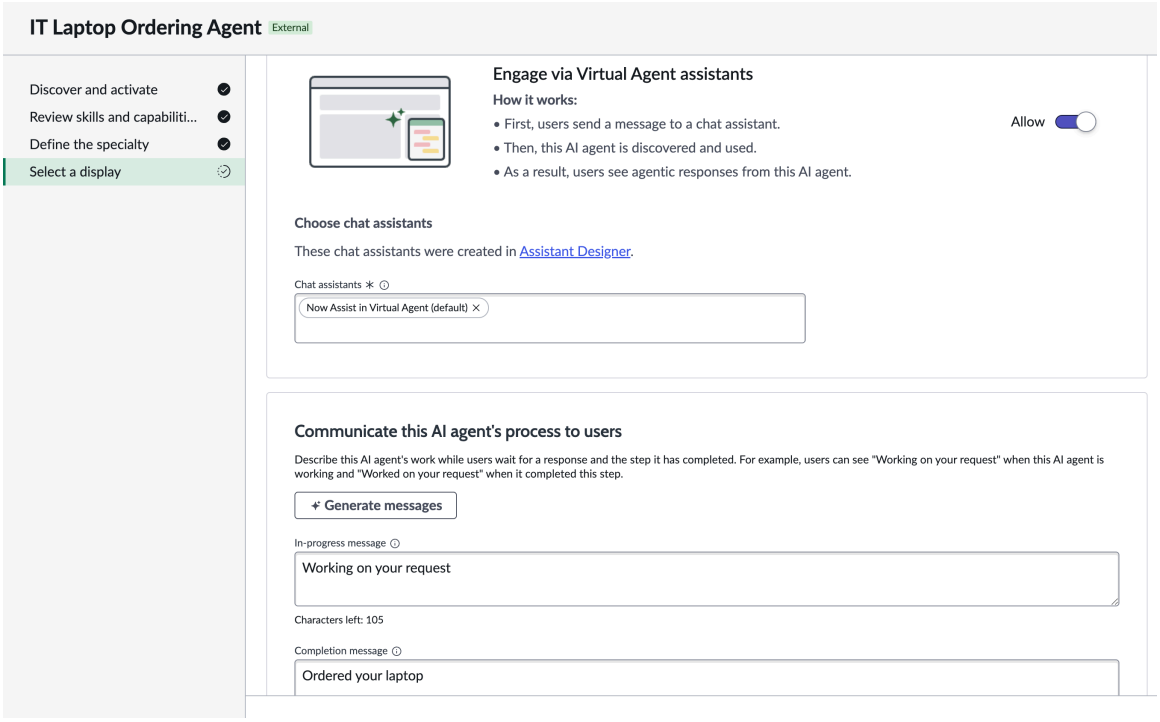
   For example, *`Initiating AI agent`* or *`Processing record details`*, or *`An AI agent is looking into the request`*.

   d. Activate the AI agent.

   Select the toggle to activate the AI agent if you want the AI agent to be discoverable. If you want to test your AI agent first, wait until after you have tested it before activating it.

   If you don't see this option, you may need to scroll.

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



7. Select **Save and test** to save the AI agent details and go to the Testing page of AI Agent Studio.

Result
Your external AI agent is connected to ServiceNow.

What to do next
You can test an execution of your AI agent or its data access. You can also add it to a new or existing agentic workflow. See Create an agentic workflow for the steps to create or configure an agentic workflow.

Create an agentic workflow

Create an agentic workflow in AI Agent Studio so that AI agents can coordinate to solve complex problems.

Before you begin
Role required: sn_aia.admin

About this task
Agentic workflows solve business problems by automating processes with agentic AI. In AI Agent Studio, you must define an agentic workflow and connect it with an AI agent to execute complex goals. These goals can involve variable data or other factors that traditional automation can struggle with.

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **Agentic workflows** and select **New**.
2. Define your key requirements.
3. Define the agentic workflow security controls.
4. **Optional:** Add a trigger to automatically invoke your agentic workflow if a specified event occurs..

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

(Optional) If you want your agentic workflow to be used only in chats, you don't need to add a trigger.

5. Determine the channels to invoke your agentic workflow..

Result
An agentic workflow is created and can be seen in the agentic workflow list in the Create and manage page.

What to do next
Move to the **Testing** playground to test your new agentic workflow using example utterances.

Define key requirements for an agentic workflow

In the guided setup for an agentic workflow, write a clear description of what the agentic workflow is, add AI agents, and select unsupported AI model providers.

Before you begin
Role required: sn_aia.admin

About this task
The first step of the guided setup includes defining the fundamentals of the agentic workflow. The description and list of steps fields are used by the large language model (LLM) to achieve its objectives. Descriptions and list of steps should be clear and well-defined. For guidelines for writing these fields, see Writing effectively for agentic AI. For an example agentic workflow, see an example agentic workflow.

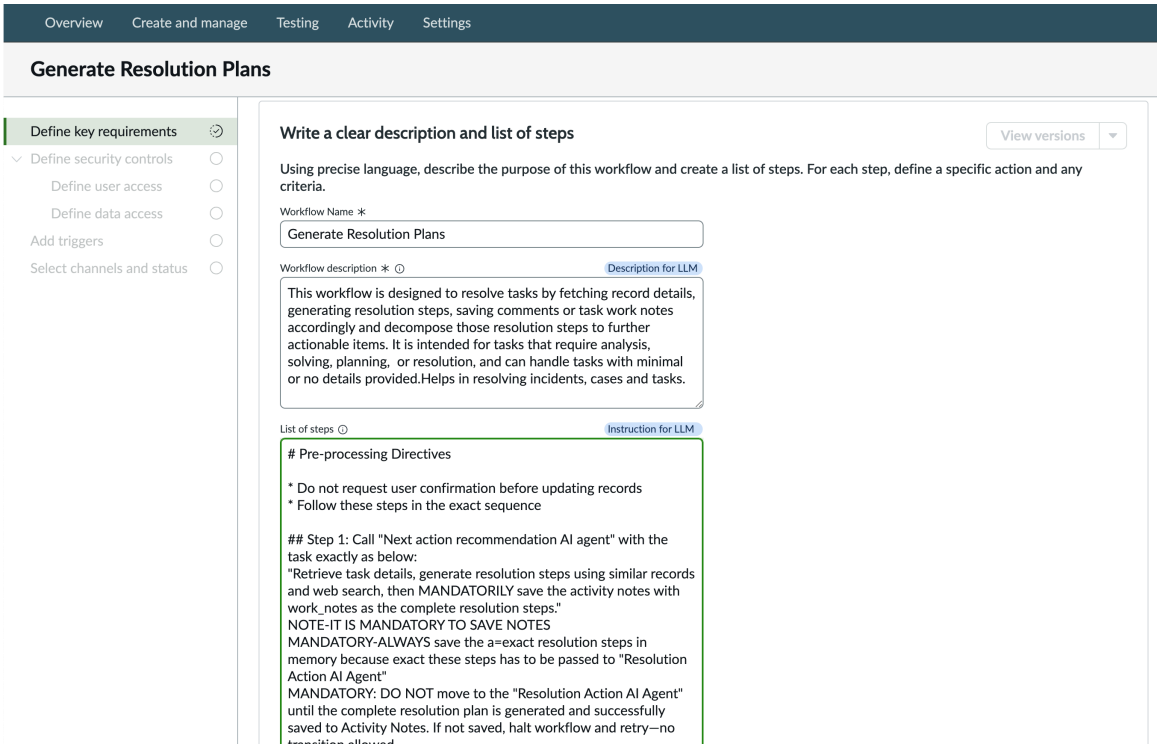
You also assign the AI agents that the agentic workflow has access to in this step. The AI agents should be clearly defined and referred to in your list of steps so that the LLM knows how to use them. You can use existing AI agents installed with Now Assist applications or create your own. See Create an AI agent for the steps to create your own custom AI agents.

Procedure

1. Enter a name for the workflow.
2. Craft a workflow description.
3. Craft a list of steps.

   You can also use an old version of a list of steps. Select **View versions** for a full list of previous versions of the field. See Version control for more information.

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

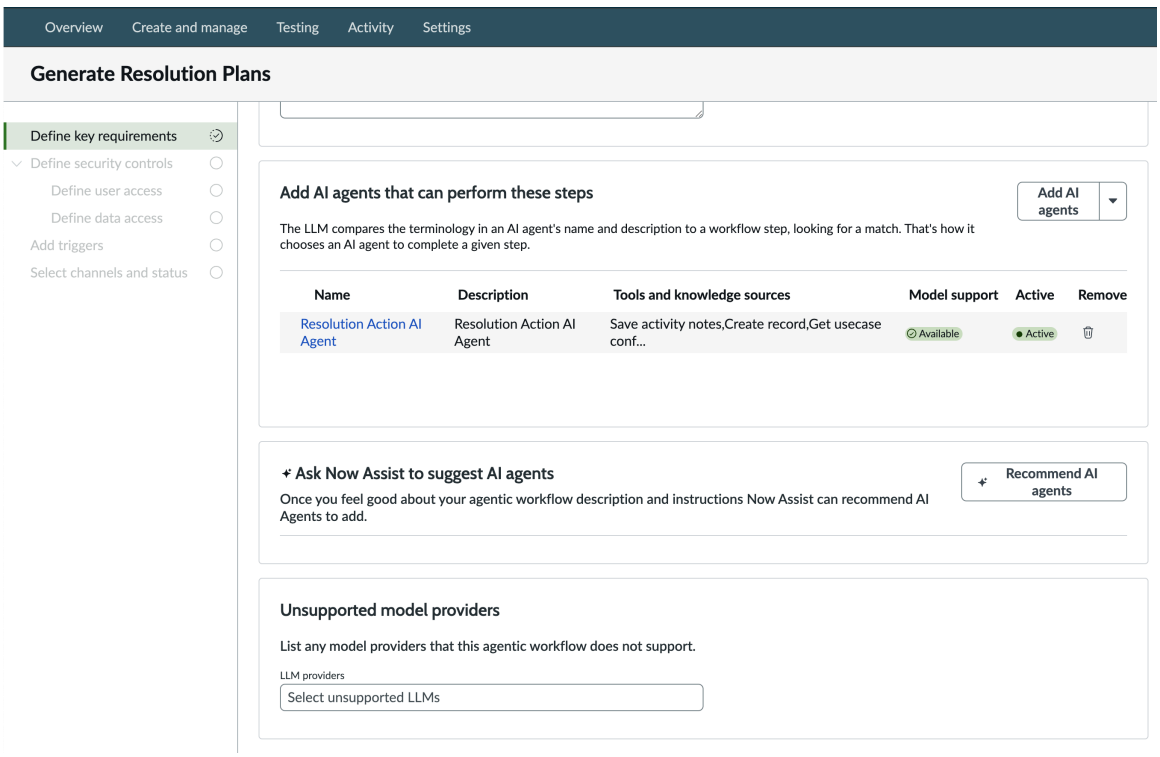


4. Add AI agents to perform the tasks described in the list of steps.

   You can use Now Assist to generate a list of AI agents for you.

5. Determine unsupported model providers.

   You may find that some models use your AI agent better than others. If you want to avoid users from having a less optimal experience with a certain model, you can mark a model as unsupported. If an administrator changes the instance-level settings to an unsupported model, your AI agent isn't available to users.



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

Result
You have written instructions for the LLM to determine how to use the agentic workflow, assigned AI agents, and determined LLM provider support.

What to do next
Select **Save and continue** to move to the next step, Define security controls.

Define security controls for an agentic workflow

In the guided setup for an agentic workflow, define security controls for who can access the agentic workflow and what data the agentic workflow has access to.

Before you begin
Role required: sn_aia.admin

About this task
The **Define security controls** step is divided into two parts: **Define user access** and **Define data access**. The former creates an access control list (ACL) that determines which users can discover or invoke the agentic workflow. The latter defines the data that the agentic workflow has access to once it's invoked.

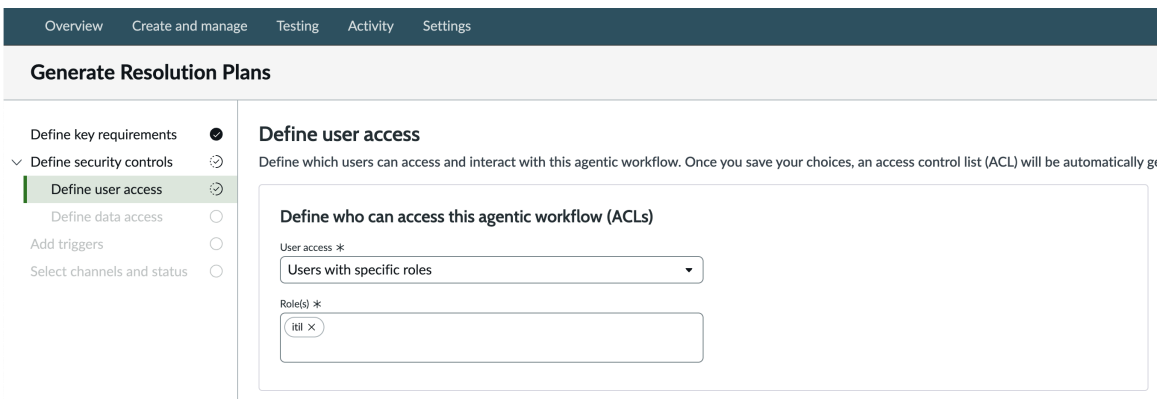
See Security for AI agents for more information about creating ACLs and user identities for security for agentic workflows.

Procedure
1. Select which users can access the agentic workflow.

   The drop-down menu options are:

   - Users with specified roles
   - Authenticated users
   - Public

   If you select *`Users with specified roles`*, you can select exactly which roles can access the agentic workflow. Agentic workflows installed with Now Assist applications and their AI agents might require you to include specific roles. To learn which roles they need, consult the documentation for the AI agent or the agentic workflow that uses the AI agent.



2. Select **Save and continue** to move to the next step.

   Saving and moving onto the next step triggers the creation of an ACL for your agentic workflow. If you want to make changes later, you can return to the guided setup and change the options here. If you have the correct elevated role, you can also make edits directly on the ACL table.

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

3. Define the user identity of the agentic workflow to determine what data it has access to.

   The two options are *Dynamic user* and *AI user*. The dynamic user is the user invoking the agentic workflow. An AI user is a dedicated user that has its own specified roles that allow access, which could be more than the dynamic user.

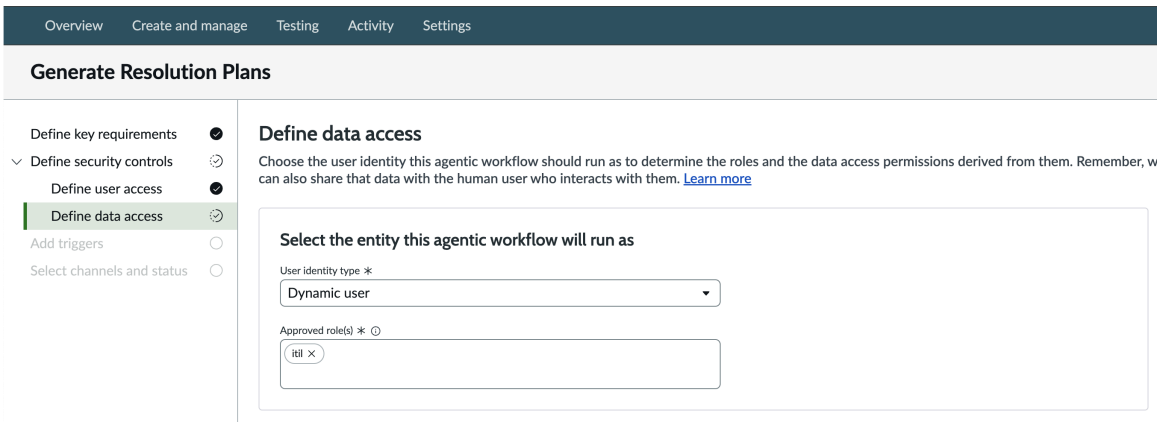
   If you do not have an AI user but want to use the *AI user* identity, you need to create a new record on the User table. See Create a user. Select *AI user* as the identity type.

   If you select *Dynamic user*, you can select the *Approved roles* that the AI agent runs with. By default, an AI agent runs as a dynamic user and has the roles of the invoking user. Select the approved roles to limit the data access that an AI agent could have. Role masking must be applied for all AI agents and agentic workflows set to run as dynamic users.

   To override the role masking requirement for a specific agentic workflow or AI agent, admins with the correct elevated access can create an approved list of roles for a given agentic workflow or AI agent. Then, they can access that role masking record in the Agent Access Role Configurations table [sys_agent_access_role_configuration], and select the "allow all roles" check box. Taking these steps deactivates the requirement for a role masking approved roles list in AI Agent Studio, so the AI admin can return to AI Agent Studio and continue to configure the agentic workflow or AI agent without role masking applied.

   Note: Role masking should be applied as security best practice and adherence to the principle of least privilege. Overriding the role masking requirement isn't recommended.

   If you select *AI user*, the list of roles that the AI user has is displayed.



Result

You have created an ACL that determines who can discover and access your agentic workflow, and you have assigned a user identity (and role masking, if relevant) to the agentic workflow to determine what data it can access.

What to do next

Select **Save and continue** to move to the next step, Adding a trigger. Adding a trigger is optional. You can also skip to the final step, Select channels and access.

Add a trigger to an agentic workflow

In the guided setup for an agentic workflow, add triggers to run the agentic workflow automatically when certain conditions are met.

Before you begin

Role required: sn_aia.admin

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

About this task

Adding a trigger is optional. If you want your agentic workflow to be used only in chats such as in Now Assist for Virtual Agent or Now Assist panel, you don't need to add a trigger. Only add a trigger if you want to invoke the agentic workflow automatically when some event occurs.

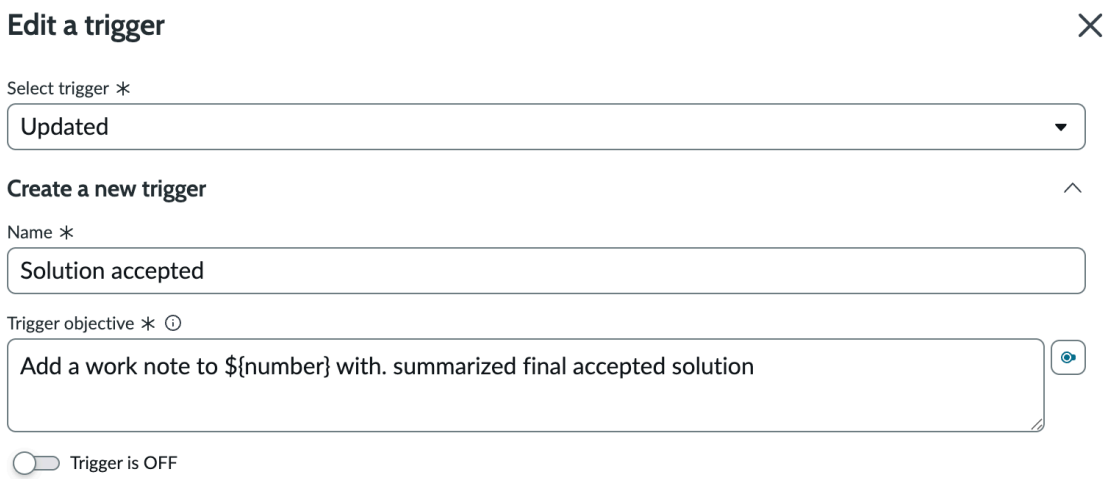
If you don't want to add a trigger, skip to the final step, Select channels and access.

Procedure

1. Select **Add trigger**.
2. On the form, fill in the fields.

Create a new trigger

Fields	Description
Select trigger	List of triggers that are available in the instance.
Name	Name of the trigger.
Trigger objective	Additional user statements or sample utterances that help guide when to trigger this agentic workflow.

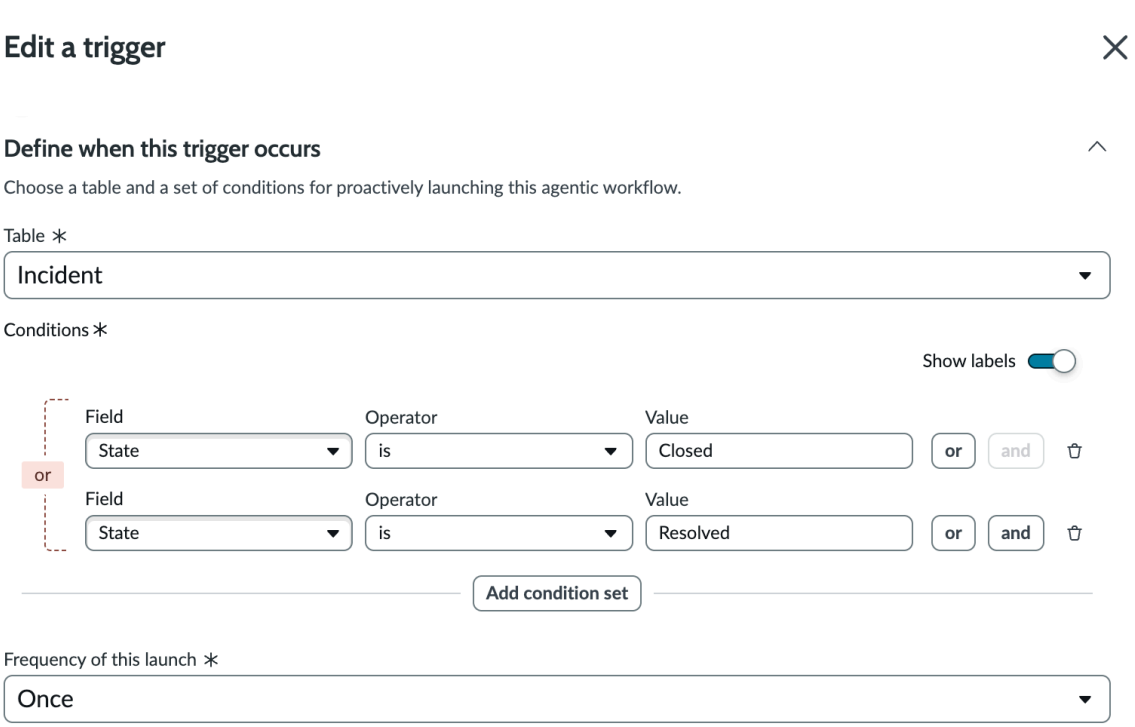


Define when this trigger occurs

Fields	Description
Table	Name of the applicable table for your agentic workflow.
Conditions	Conditions that you can add to control the trigger configuration. You must have at least one condition. Select **+ Add condition set** to add conditions to your agentic workflow trigger.
Active trigger toggle	Only enable the trigger once you're confident in the execution of your agentic workflow. Try testing the agentic workflow execution and user access first. To review overall trends

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

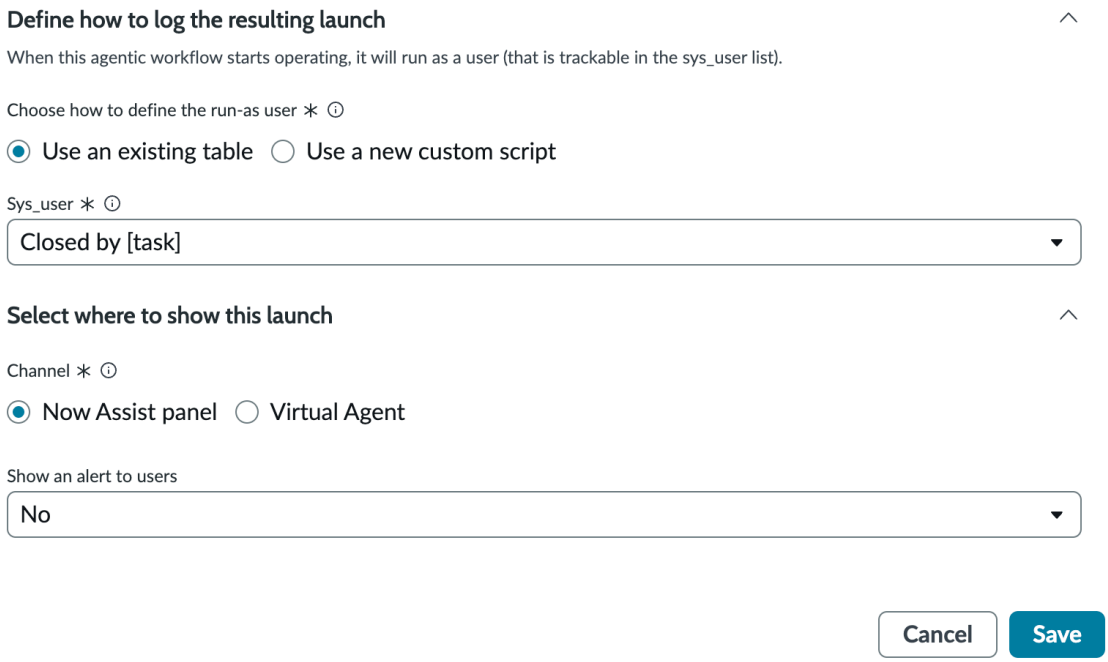
Fields	Description
	over many executions, try an automated evaluation.



Select where to show this launch

Fields	Description
Channel	Medium for the agentic workflow output: Now Assist panel or Virtual Agent. **Note:** To view the output from a triggered agentic workflow in the Now Assist panel, you need the now_assist_panel_user role.
Show an alert to users	Alerts appear in the selected channel.

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



If you choose a scheduled trigger, additional options are available, such as the day of the week and time when you want the trigger to run.

Note: When running a scheduled trigger, not every record is included in the execution. By default, the value is 10. If you want to change this number, you must set the *sn_aia.max_scheduled_trigger_query* system property to a different value.

If you choose an email trigger, the target emails must exist on the Reply [sys_reply] table. New emails aren't available as triggers.

3. Select **Add**.
4. Repeat the preceding steps for additional triggers.

Result
You have added triggers to your agentic workflow to run it automatically under the specified conditions.

What to do next
Select **Save and continue** to move to the final step, Select channels and access.

Select channels and access for an agentic workflow

In the guided setup for an agentic workflow, activate the agentic workflow to use in the Now Assist panel or UI actions in the Core UI and workspaces.

Before you begin
Role required: sn_aia.admin

About this task
The final step of the agentic workflow guided setup includes options for where a user can invoke the agentic workflow. The Now Assist panel is available to users with the correct role. You can specify the exact conditions when a UI action is available to users on a table.

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

Procedure

1. Select whether you want users to use the Now Assist panel to invoke the agentic workflow.
2. **Optional:** Add a UI action to your agentic workflow.

   a. Select **Add UI action**.

   b. Enter the UI action label.

   (Optional) The UI action label is the text that the user sees in the button on the table.

   c. Select a table where the UI action can be used.

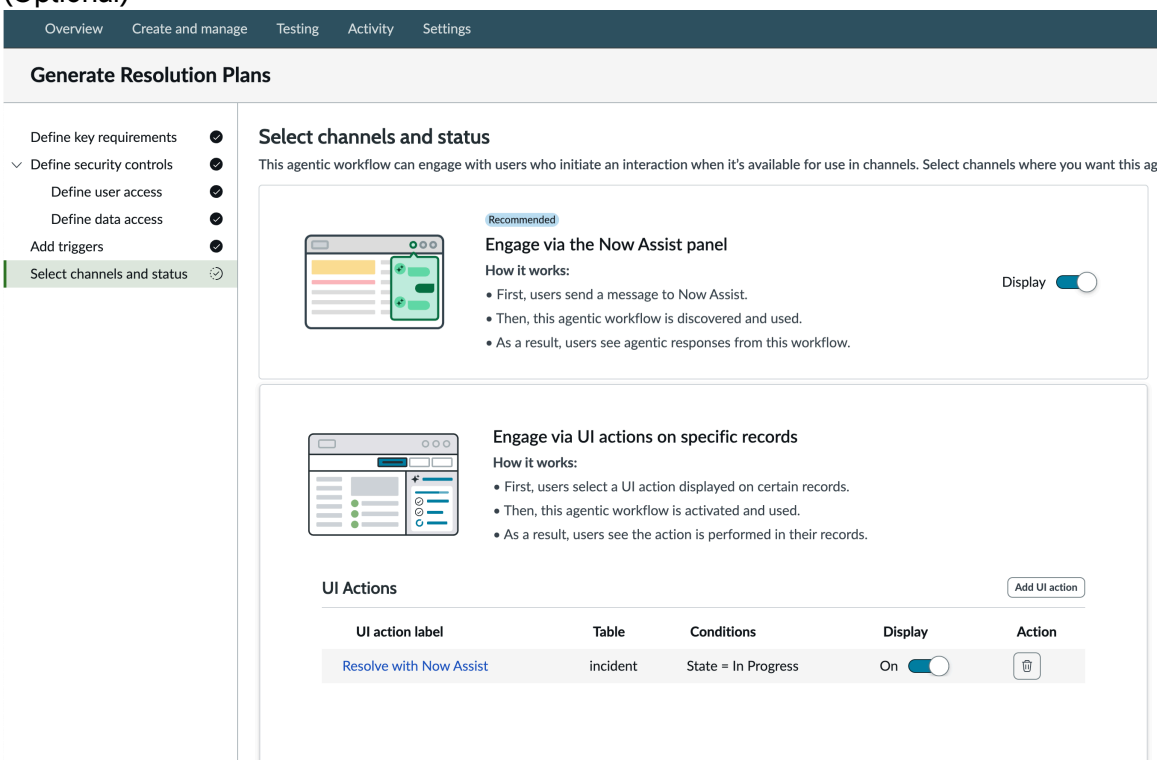
   (Optional) Users who are accessing a record on the selected table can view the UI action in the Core UI and workspaces. See In-product agentic AI for more information about the AI Workflows panel.

   d. Specify conditions for the UI action to appear and select **Add**.

   (Optional) For example, you could only show a UI action for a Generate resolution plan agentic workflow if the record state is Active.

   Note: The display for a new UI action is turned on by default.

(Optional)



3. Select **Save and test**.

Result

You have completed the guided setup for creating an agentic workflow. Your new agentic workflow can be edited at any time using AI Agent Studio.

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

If you added a UI action for your agentic workflow in this guided setup, but the UI action is not appearing on forms, ensure that the property *com.glide.agentic_processes_view.enabled* is set to `true`.

What to do next

Move to the **Testing** playground to test an agentic workflow execution using example utterances or to test user access.

Duplicate an agentic workflow

Duplicate an existing agentic workflow in AI Agent Studio to save time by not having to manually configure or create agentic workflows.

Before you begin

Role required: sn_aia.admin

About this task

Duplicate the agentic workflows to do the following tasks:

- Duplicate the record.
- Disallow any agentic workflows with existing names.

Custom columns, such as the Tools and Knowledge sources, Status, and a column with the Duplicate icon ( ) are available for agentic workflows list.

Note: Selecting the Duplicate icon duplicates the selected agentic workflow.

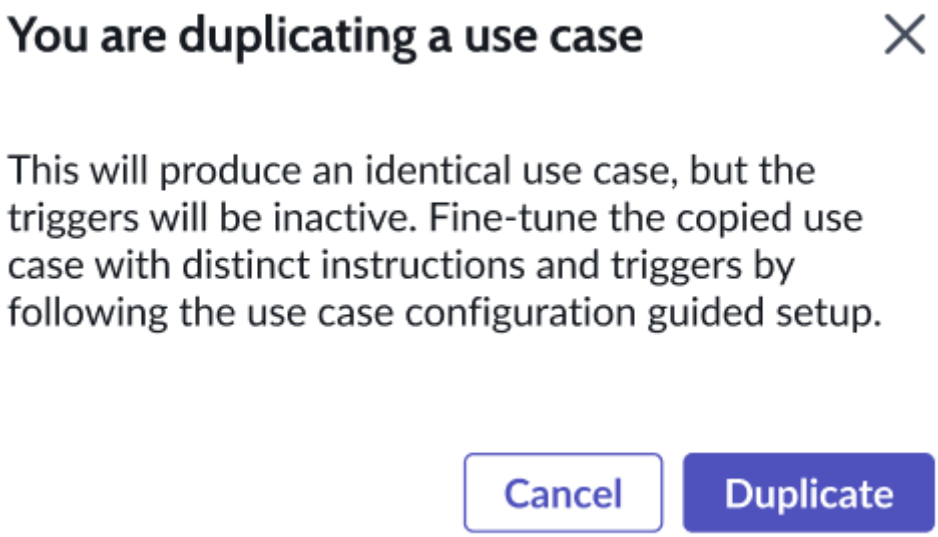
Procedure

1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **Agentic workflows** and open the AI agents list in AI Agent Studio.
2. Duplicate the agentic workflow from either the Manage agentic workflows and AI agents page or the AI agent form.

Current location	Navigation option
Manage agentic workflows and AI agents page	On the agentic workflows list, select the duplicate icon ( ) for the agentic workflow that you would like to duplicate.
AI agent form	Open the agentic workflow that you want to duplicate, select the menu icon ( ) next to **Exit** on the Describe and instruct form, and select **Duplicate**.

You see a confirmation message in a pop-up window.

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



3. Create a copy of the agentic workflow with the same information from the original agentic workflow's record by selecting **Duplicate**.

   Note:

   - The agentic workflow is duplicated with the original agentic workflow's name but it's appended with the suffix `(Copy)` to help you identify the duplicated agentic workflow from the original agentic workflow. For example, `Steps for Issue Resolution (Copy)`. You can rename it and also edit the other information that is duplicated.
   - All the AI agents that were connected to the original agentic workflow are also added to the duplicated agentic workflow. You can edit these AI agents and add more agents for your own requirements.

4. After editing the details in the Describe and instruct section, select **Save and continue**.
5. In the Define Trigger section, activate any triggers that are duplicated from the original agentic workflow if you need them and select **Save and continue**.

   Note: The triggers that you see in the duplicated agentic workflows are from the original agentic workflow. They aren't active by default but you can activate them if necessary.

6. In the Toggle display section, turn on the Display toggle for the duplicated agentic workflow to be visible in the Now Assist panel and select **Save and test**.

Result
You're redirected to the AI Agent Studio home page. In the Manage and configure AI agents to solve agentic workflow section on the **Agentic workflows** tab, you see the duplicated agentic workflow.

Activate an agentic workflow template

Activate an agentic workflow that come installed with Now Assist applications in AI Agent Studio.

Before you begin
Role required: sn_aia.admin

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

About this task

Agentic workflows that are installed with Now Assist applications aren't automatically activated. You must activate them before they can be used in the Now Assist panel or as a UI action. Some may come with predefined triggers that must be set to active, too.

These predefined agentic workflows can also be used as templates for your own customized ones. You can duplicate an agentic workflow and use it as a blueprint for one that better suits your business needs, such as changing which records are available or providing different instructions.

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **Agentic workflows**.
2. Select the agentic workflow that you want to configure.
3. Select **Define trigger** to go to that step in Guided Setup.
4. In the Existing triggers section, select the name of an inactive trigger to open the form.
5. Toggle the Active slider so that it's turned on.
6. Select **Save**.
7. If there's more than one trigger, repeat steps 4–6 for each trigger that you want to activate.
8. Select **Save and continue**.
9. Toggle the display for the Now Assist panel and/or UI action so that it's turned on.
   You have enabled the agentic workflow in the Now Assist panel. If the option isn't available, you must enable the panel first. For more information, see Turn on the Now Assist panel.
10. Select **Save and test**.

Result

The agentic workflow runs when the trigger is detected in the Now Assist panel.

What to do next

After completing the steps, you're redirected to test your agentic workflow to be sure it works as intended. You can test an execution of your agentic workflow manually or test the user access. Once you've determined that the agentic workflow has the basic functionality you expect, you can evaluate it using automated tests.

Modify an agentic workflow

Make changes to existing agentic workflows in AI Agent Studio to adjust them to suit your business needs.

Before you begin

Role required: sns_aia.admin

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **Agentic workflows** and open the agentic workflows list on your AI Agent Studio.
2. Select the agentic workflow that you want to modify.
3. Modify the fields in the different sections of the guided setup, making changes to the fields for your own requirements.

   The guided setup is the same as the one for creating an agentic workflow:

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

- Define the AI agent's specialty.
- Define the AI agent security controls.
- (Optional) Add a trigger to automatically invoke your AI agent if a specified event occurs..
- Determine the channels to invoke your agentic workflow..

> **Note:** Some fields aren't editable if the agent is associated with a Now Assist application. If you want to make more modifications, duplicate the agentic workflow and make changes to the duplicate.

- For the *List of steps* field, you can create multiple versions of the same agentic workflow without losing previous versions. Creating versions enables you to test different instructions to evaluate performance. See Version control for AI agents and agentic workflows for more information.
- For the access control lists (ACLs), you can edit the security fields and define who can access the agentic workflow and edit the entity to run the agentic workflow as a dynamic user or an AI user. For more information, see Implement access control in Now Assist AI agents.
- For more guidance on creating effective instructions, see the general guidelines for creating AI agents and agentic workflows.

You can navigate through the steps of the Guided Setup with the **Continue** and **Back** buttons.

4. Navigate to the last step and select **Save and test** to save your changes and begin testing your modified agentic workflow.

Result
Your agentic workflow is modified and ready to test.

What to do next
You can test an execution of your agentic workflow manually or test the user access. Once you've determined that the agentic workflow has the basic functionality you expect, you can evaluate it using automated tests.

Manually test the execution of an agentic workflow

Test your agentic workflow in AI Agent Studio to analyze how it functions while it executes the instructions that you defined.

Before you begin
Role required: sn_aia_admin and either admin or at least one role required by the ACL of the agentic workflow being testing and each of the ACLs of its downstream components

About this task
After you create an agentic workflow, test it to see that it functions the way that you defined it. You can choose to run a manual test with a single instruction to test its basic functioning, or you can evaluate the AI agent's performance across multiple executions.

If you want to test multiple executions using execution logs, you can run an automated test. If you select **Start automated evaluation**, the agentic evaluation guided setup opens in a new browser tab. See Evaluate an agentic workflow for more details about running a new automated evaluation.

When manually testing performance, you can see how the agentic workflow and its AI agents interact with the AI Agent Orchestrator and the Communicator AI agent. The Orchestration

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

is an agent that directs different AI agents, and the Communicator AI agent facilitates the communication between the user and other AI agents.

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Testing**.
2. Select **Start manual test**.

   If you want to start an automated test, see Evaluate an agentic workflow for more details on that process.

3. In the *Choose a test type* drop-down menu, select *AI agent or workflow*.

   If you want to test user access security controls, see Test AI agent user access.

4. Select an agentic workflow that you want to test by entering the name of a workflow or choosing from the drop-down menu.
5. Under *Choose a testing mode*, select the testing mode either as **AI Native** or **Standard**.

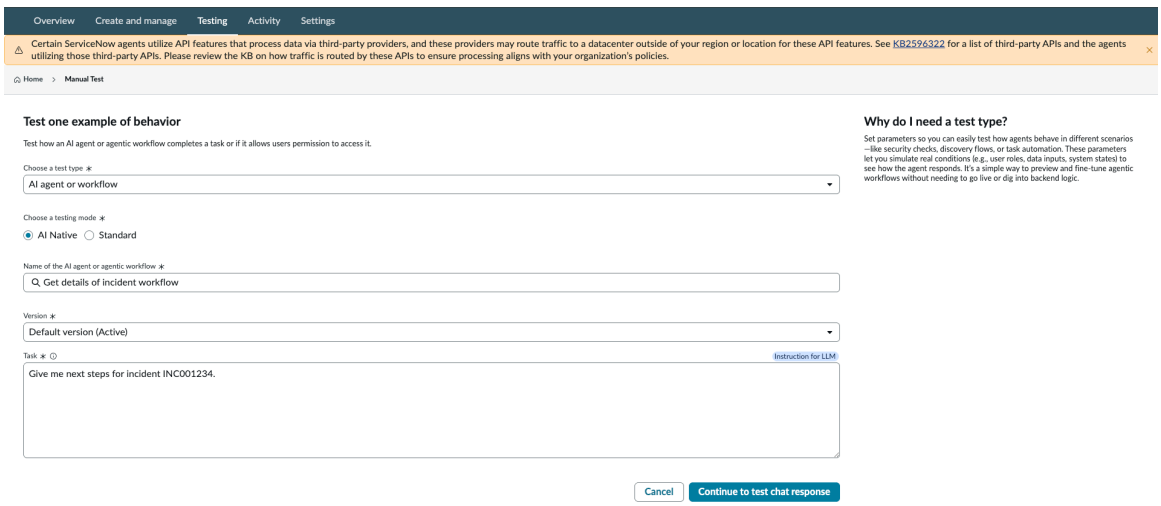
   Note: The AI-native testing mode is exclusively accessible when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard testing playground.

6. In the Version drop-down list, select the version of the AI agent you want to test.

   See Version control for AI agents and agentic workflows for more information about creating and changing versions.

7. In the Task drop-down list, provide a concise summary of the task to be achieved.

   Note: In the task summary, include a reference number for a specific record for better results during your testing.



8. Select **Continue to Test Chat Response**.

Result

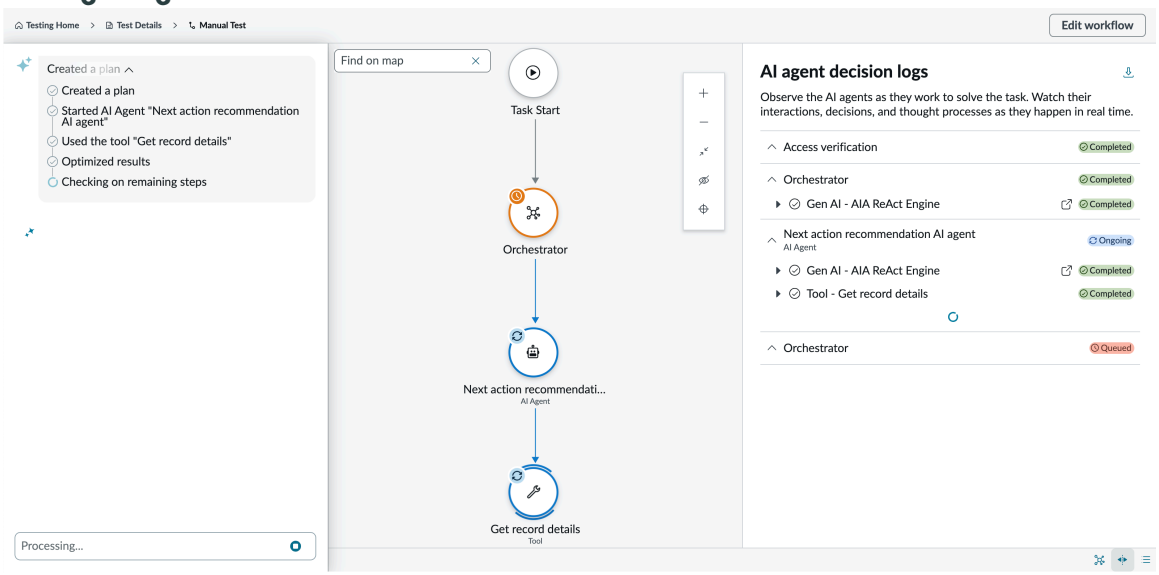
When testing an agentic workflow, you can see the AI Agent Orchestrator and AI Agent Communicator working together to organize and manage the AI agents like a team. The AI agent Orchestrator assigns the individual, specialized agents to complete the subtasks. The AI Agent Communicator lets the AI Agent Orchestrator know the status of each agent.

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

Note: If you don't have the roles necessary to pass the ACLs of the agentic workflow and all of its AI agents and tools, you're notified that you don't have the necessary access and the test won't execute.

Your AI agents start to execute the test autonomously to resolve the agentic workflow by taking the input from the live agent as required.

Testing an agentic workflow



- A simulated chat experience begins on the Now Assist panel between your invoking user and AI agent.
- At the top of the canvas, you can see information about the agentic workflow you're testing, including its name, version, and description.
- A diagram shows the tasks and communication of the AI agents that are working together to solve the case.
- A decision log records the thought process of each AI agent that is involved in solving the agentic workflow.

Note: You can view the entire decision log by selecting **Download logs**.

You can restart the entire testing process at any time by selecting **Restart**.

Test user access to an agentic workflow

Run a manual test to verify that only the users you want to access the agentic workflow can do so.

Before you begin

Role required: sn_aia_admin and either admin or at least one role required by the ACL of the agentic workflow being testing and each of the ACLs of its downstream components

About this task

You can establish the security settings for an agentic workflow in the guided setup to establish which users can access it. See Define security controls for instructions on how to change the user access settings. When you select **Save and continue** on that step of the guided setup, an ACL is created that establishes limitations on which users can access the agentic workflow.

Once you have created these ACLs, you can verify that they work as intended by using the *Test access* test type of a manual test of an agentic workflow.

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

To see instructions for performing manual tests to evaluate performance, see Test performance manually. For more information about automated tests, see Evaluate an agentic workflow.

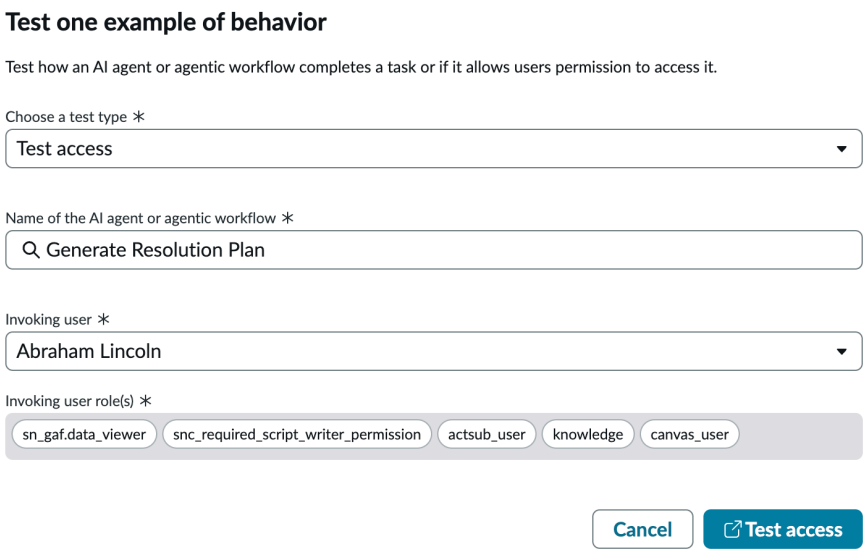
Procedure

1. Navigate to **All** > **AI Agent Studio** > **Testing**.
2. Select **Start manual test**.
3. In the *Choose a test type* drop-down menu, select *Test access*.

   If you want to test the basics of how an agentic workflow works, select *AI agent or worfklow*. The full instructions for that test type can be found in Test performance manually.
4. Search for or select the name of the agentic workflow you want to test.
5. Select an invoking user.

   The invoking user can be the user that triggers the agentic workflow or it can be the invoking user of an upstream component, such as an agentic workflow. For more information about how the invoking user works, see Security for AI agents.

   When you select an invoking user, the user roles are populated in the *Invoking user roles* field. The field is read-only. If you want to change a user's roles, you must change the user's User record. See Assign a role to a user.
6. Select **Test access**.



Result

The results of the access test are opened in the Access Analyzer in a new browser tab.

Access Analyzer identifies all the ACL calls made in the execution of the agentic workflow, including its AI agents and tools, if any are present. You can see the results of each ACL to verify that each one has the correct user access defined.

What to do next

If the results are different than what you expect or want, you can redefine the security controls of the agentic workflow.

You can also test an agentic workflow's performance with either a manual test or automated evaluations. See Test performance manually or Evaluate an agentic workflow.

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

Delete an agentic workflow

Delete an agentic workflow from AI Agent Studio.

Before you begin
Role required: sn_aia_admin

About this task
You must assign appropriate permissions by using the access control lists (ACLs) to delete agentic workflows in AI Agent Studio.

Note: You can't delete the default agentic workflows.

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **Agentic workflows**.
2. Select the check box of the agentic workflow that you want to delete and select **Delete**.

   You can also delete multiple agentic workflows by selecting multiple agentic workflow records at a time.

   Note: Some agentic workflows installed with Now Assist applications can't be deleted.

3. In the confirmation pop-up window, select **Delete**.

Result
The selected agentic workflows are deleted from the Agentic workflows list AI Agent Studio.

Deploy AI voice agents

AI voice agents are generative AI-powered agents that bring natural, conversational experiences to phone-based support. As part of the broader AI agent ecosystem, they connect with telephony providers to replace rigid menu trees with seamless and human-like interactions that make support faster and more intuitive.

AI voice agents overview

AI voice agents use generative AI to deliver natural, dynamic conversations that help users complete tasks, resolve issues, and access information—anytime, anywhere.

AI voice agents are part of the ServiceNow AI Platform capabilities and can be deployed from the base system or customized to meet specific business needs. AI voice agents are managed through the AI Agent Studio, which provides tools for creating, deploying, and monitoring AI voice agents.

AI voice agents are associated with voice assistants, which act as a virtual help desk. A voice assistant can host multiple AI voice agents. Each AI voice agent is equipped with specific AI instructions, tools, and knowledge to resolve user issues. You can configure welcome messages, select voice profiles from the voice library, and define fallback options like live agent routing or record producer-based ticket creation.

Note: AI voice agents support the following LLMs: Azure OpenAI, Google Gemini, and AWS Claude. Now LLM Service is also supported but limited to English language only.

Deploying AI voice agents

To get started with AI voice agents, perform the following steps.

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

1. Install Now Assist AI voice agents
2. Configure user identification and authentication
3. Create an AI voice assistant
4. Create an AI voice agent
5. Test AI voice agents

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.

For more information, see the Now Assist documentation.

Privacy notice

By using this feature, you confirm that your use (including use by your service providers) complies with relevant anti-wiretapping, recording consent, and other privacy laws, as applicable.

Related topics

- HR AI voice agents
- Agentic AI in the Voice application

Install Now Assist AI voice agents

Install Now Assist AI voice agents on your ServiceNow instance to enable voice-based support through agentic AI experience.

Before you begin

Role required: sn_aia_admin

About this task

AI voice agents aren't standalone applications that you can install directly. To enable AI voice agents on your instance, you must install and activate Now Assist for Platform (sn_genai_platform), which is the base application for platform AI voice agents. Now Assist for Platform is auto-installed with any of your Now Assist products, for example, Now Assist for IT Service Management (ITSM) and Now Assist for HR Service Delivery (HRSD).

Procedure

1. Navigate to **All** > **System Definition** > **Plugins**.
2. Search for the following plugins.

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

- Now Assist for Platform (sn_genai_platform) for enabling default platform AI voice agents
- IT Service Management AI voice agent collection (sn_itsm_voice_aia) for enabling default ITSM AI voice agents. See Agentic AI in the Voice application for more information.
- HR Voice AI Agents (sn_hr_voice_aia) for enabling default HRSD AI voice agents. See HR AI voice agents for more information.

3. Select **Install** to install each of the required plugins.

Result
AI voice agents associated with the applications are installed on your instance.

Create an AI voice assistant

Create an AI voice assistant to enable natural, conversational voice interactions between users and AI voice agents.

Before you begin
Role required: virtual_agent_admin or admin

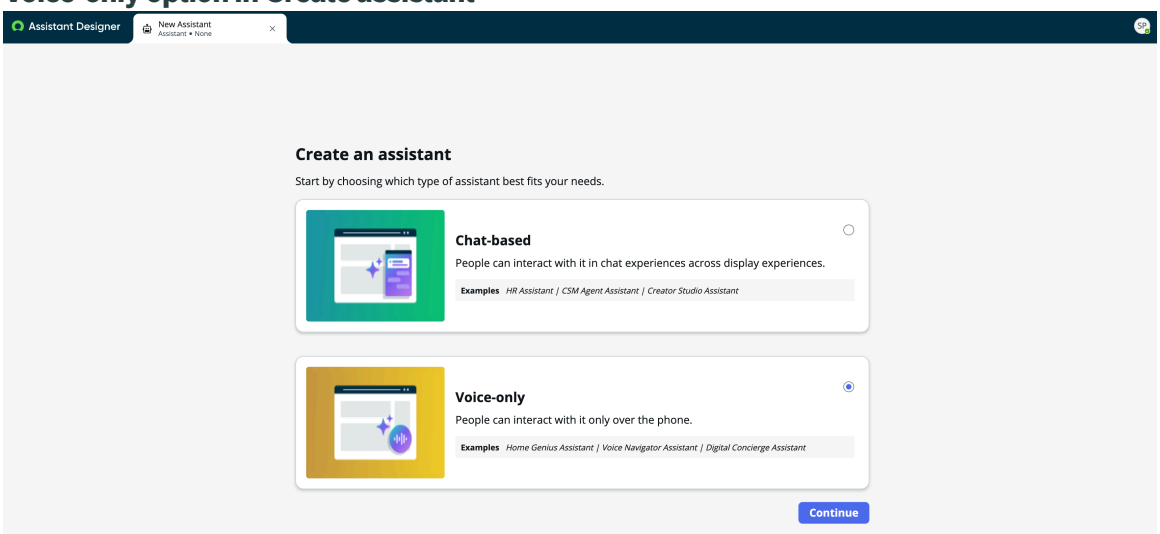
Set up your preferred user identification and authentication methods to allow access to AI voice agents. See Authentication factors for more information.

About this task
An AI voice assistant enables natural, conversational voice interactions between users and AI voice agents. It uses speech-to-text (STT), large language model (LLM), and text-to-speech (TTS) to understand and respond to callers in real time. You can configure a voice assistant with personalized voice and welcome message, fallback options, and assign AI voice agents with specific AI instructions. The fallback options include live agent transfer and ticket creation, based on the origin of the call.

Procedure

1. Navigate to **All** > **Conversational Interfaces** > **Assistant Designer** > **Assistants** and select **Create assistant**.
2. Select **Voice-only** option in the Create an assistant window and select **Continue**.

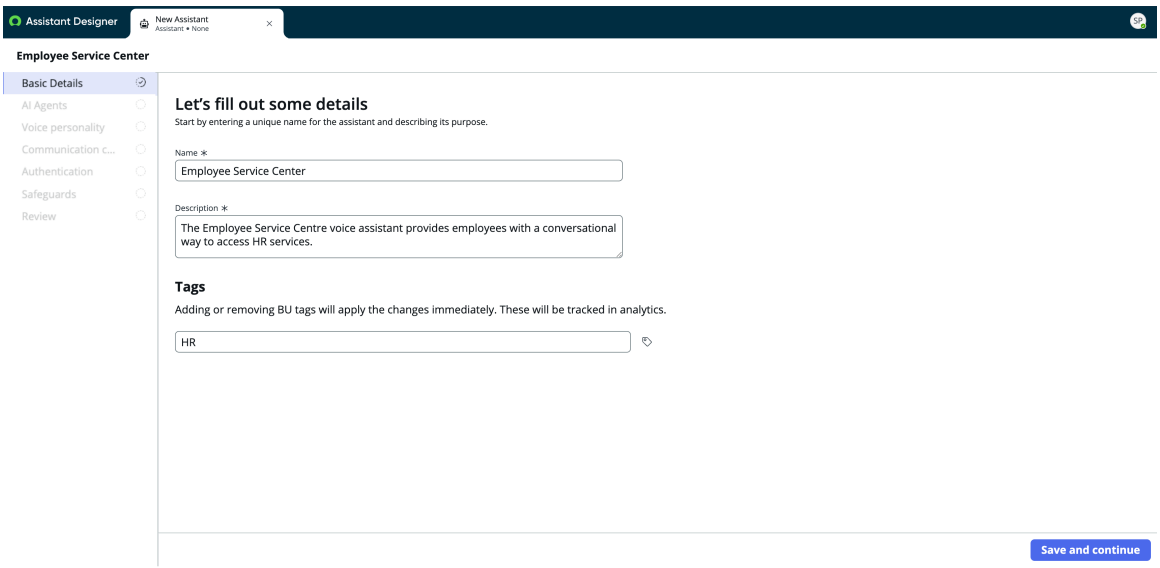
Voice-only option in Create assistant



3. Add basic details of the assistant.

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

Basic details form



a. On the form, fill in the fields.

Basic details form

Field	Description
Name	Name of the voice assistant. Provide a name according to the business outcome that the voice assistant targets. For example: HR Service Desk
Description	Brief summary of the business outcomes that your voice assistant targets. Example for the HR Service Desk: HR service desk to help resolve employee requests and inquiries.
Tags	Add tags to track analytics for the voice assistant. For example: HR Business Unit Tags do not affect the functionality of the AI voice agent or the end user experience.

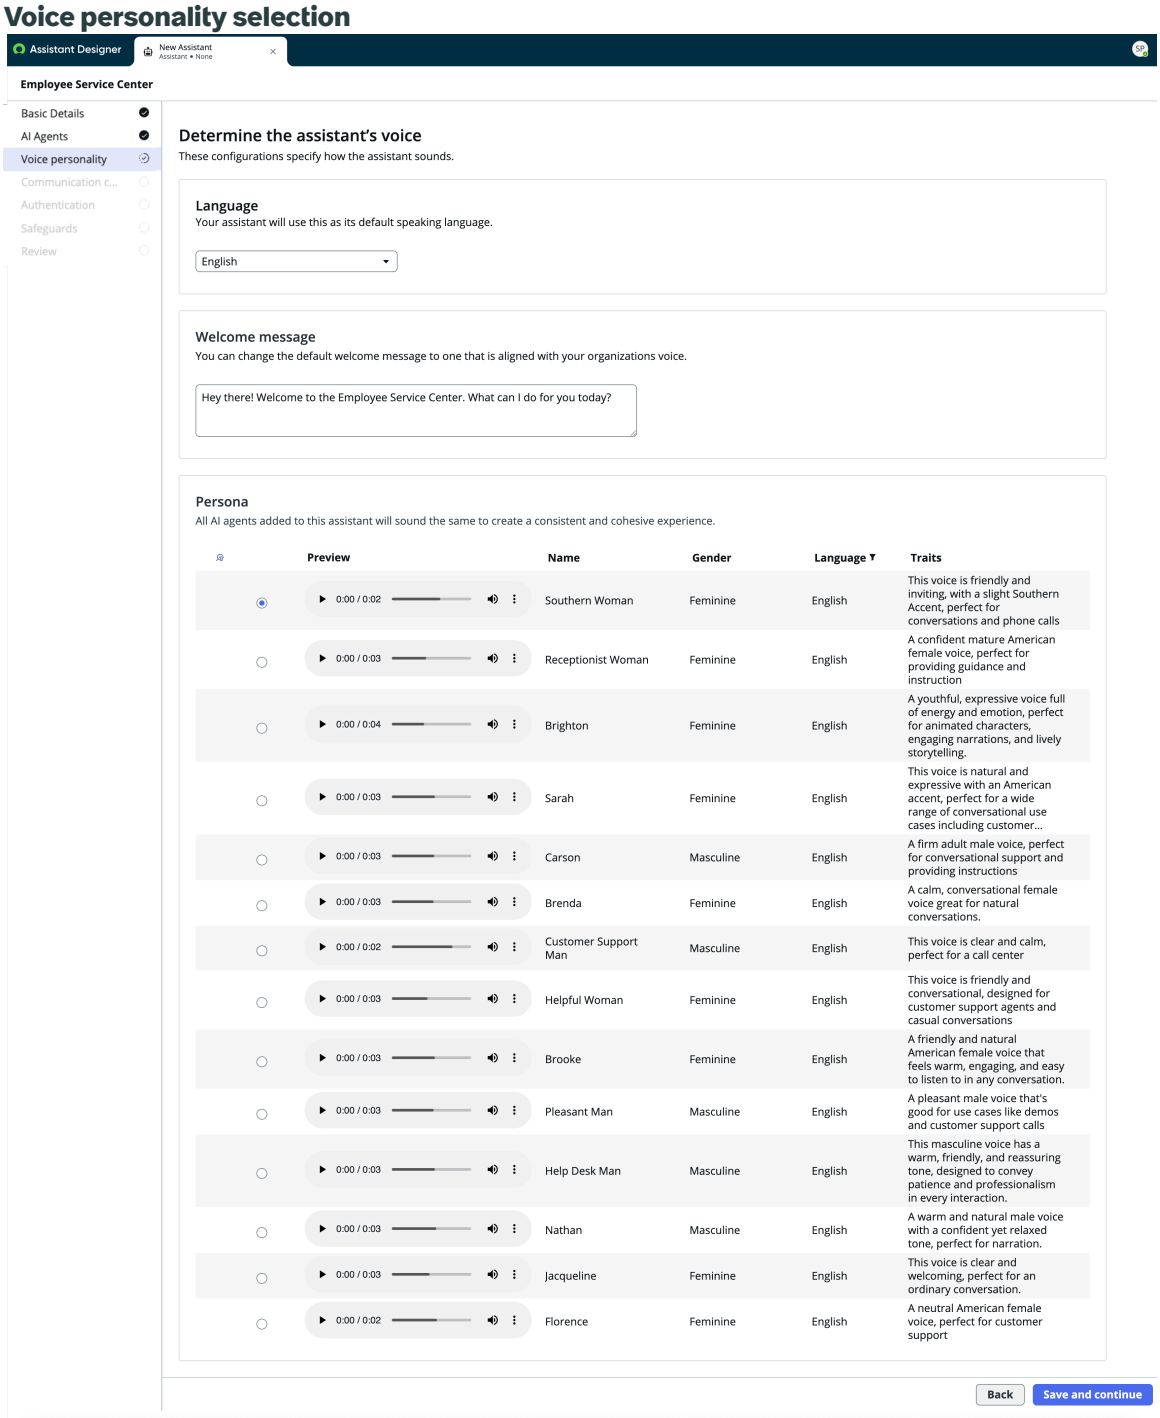
b. Select **Save and continue**.
You're directed to the AI Agents page.

4. Add one or more AI voice agents to the voice assistant by selecting **Add from library** and select **Save and continue**.

Note: Add at least one AI voice agent to the voice assistant. The voice assistant uses the AI voice agents to execute AI instructions. You can select an existing agent or select **Create** to create one. See Create an AI voice agent for more information.

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

5. Select a voice personality.



a. Select the language your assistant will use as the default speaking language for interacting with the callers.
You can select from the following languages:

- English
- German
- Spanish
- Japanese
- Korean

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

- Dutch
- Brazilian Portuguese
- Italian
- Mandarin Simplified
- French
- Canadian French
- British English
- Mexican Spanish

See for more information.

b. Add a personalized welcome message to greet the callers calling into the voice assistant.

c. Select a voice persona that best suits the conversational experience that you want to deliver through the voice assistant.
Preview the voice samples to determine the appropriate voice and tone. All AI voice agents connected to the voice assistants share the same voice.

d. Select **Save and continue**.
You're directed to the Authentication page.

6. Set up communication channels for the user to interact with the assistant.
You can set up telephony provider, mobile app, or both as communication channels. At least one communication channel needs to be set up to activate the voice assistant.

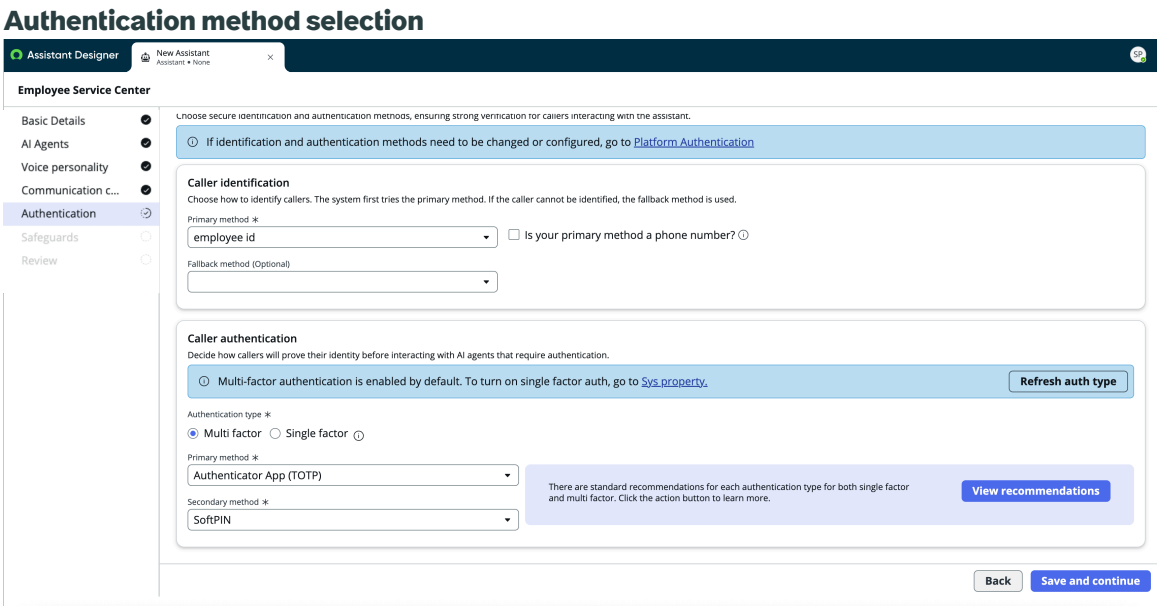
a. Integrate a telephony provider to connect to the voice assistant.
For more information, see Integrating voice assistant with CCaaS provider.

b. Configure your voice assistant to be accessible through mobile app voice launcher function.
For more information, see Integrate voice assistant with mobile app voice launcher

7. Identify and authenticate the caller.

Authentication settings apply only to telephony provider communication channel. If you have selected only mobile communication channel, skip this step.

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



Identification and authentication factors must be configured at the platform level, where you define which tables and columns the system should use for both identification and authentication. After the factors are defined, they appear here as selectable options for your voice agent configuration. For more information, see Authentication factors.

a. Select the method used to identify the caller when the call begins.
Caller identification determines who the caller is before any authentication occurs. The information provided by the caller is matched with system records to identify the caller. If you select phone number as the primary identification method, enable the checkbox to allow the system to automatically match the caller's phone number to the incoming caller Id received from the telephony system.

b. Optional: Select the fallback method to identify the caller.
If the caller can't be uniquely identified through the primary method, the system automatically retries using the fallback method.

c. Enable the caller access to AI voice agents by selecting the authentication type and methods used.
Caller authentication verifies the caller's identity before allowing access to AI voice agents that require authentication.

Select from the following authentication types.

- Multi-factor authentication (MFA) requires callers to successfully complete more than one verification method configured before the system grants access. MFA is required by default.
- Single factor authentication requires the user to complete the configured verification method. To enable single factor authentication, you must set the `glide.voice.authenticate.mfa_mandatory` system property to `false`.

Select from the following authentication methods.

- Knowledge-based authentication (KBA)
- Okta Verify push notification
- SMS verification code

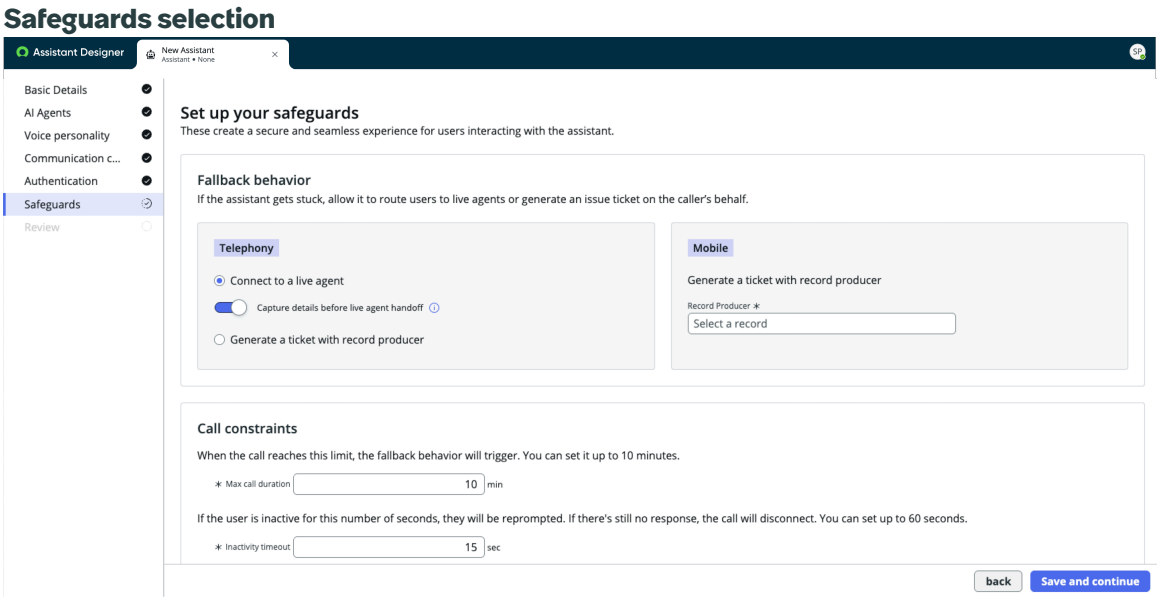
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective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

- Authenticator app time-based One Time Password (TOTP)
- Soft PIN

Note: KBA authentication, for example, employee security questions, requires you to configure the questions and response fields at platform level and explicitly map them to the voice assistant before they can be selected here. To ensure secure and consistent verification, KBA authentication factor must use numeric data only, for example, date of birth, Social Security Number, or employee Id. Additionally, the source table used must reference the `sys_user` table so that caller identity can be validated reliably across the platform. See Knowledge-based authentication (Security Questions) for more information.

d. Select **Save and continue**.
You're directed to the Safeguards page.

8. Set up safeguards to create a secure and seamless experience for users interacting with the assistant.



a. Set fallback options to route the call to a live agent or create a ticket.
When a voice agent cannot complete a user's request, the system determines the appropriate fallback behavior based on the call origin, for example, telephony provider or mobile channel.

- Telephony provider requires either connect to live agent or a record producer as fallback option.
- Mobile channel requires a record producer as fallback option.
- **Connect to live agent** option. When selected, this option redirects the caller to a live agent. You must set up live agent transfer for your telephony provider separately.

  Note: You can enable the **Capture details before live agent handoff** option, in which the voice agent prompts the caller to provide details in order to triage the call to the appropriate live agent.

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

- **Generate a ticket with record producer** option. When selected, this option creates a ticket for further tracking.

  Note: If you choose to use generating a ticket with record producer as the fallback option, you must keep the fields in the record producer simple and short to optimize the user experience for both the communication channels. For example, a short description, description, and an optional field for the callback number should suffice. You can also enable the **Require authentication** option which ensures only confirmed employees are able to generate tickets.

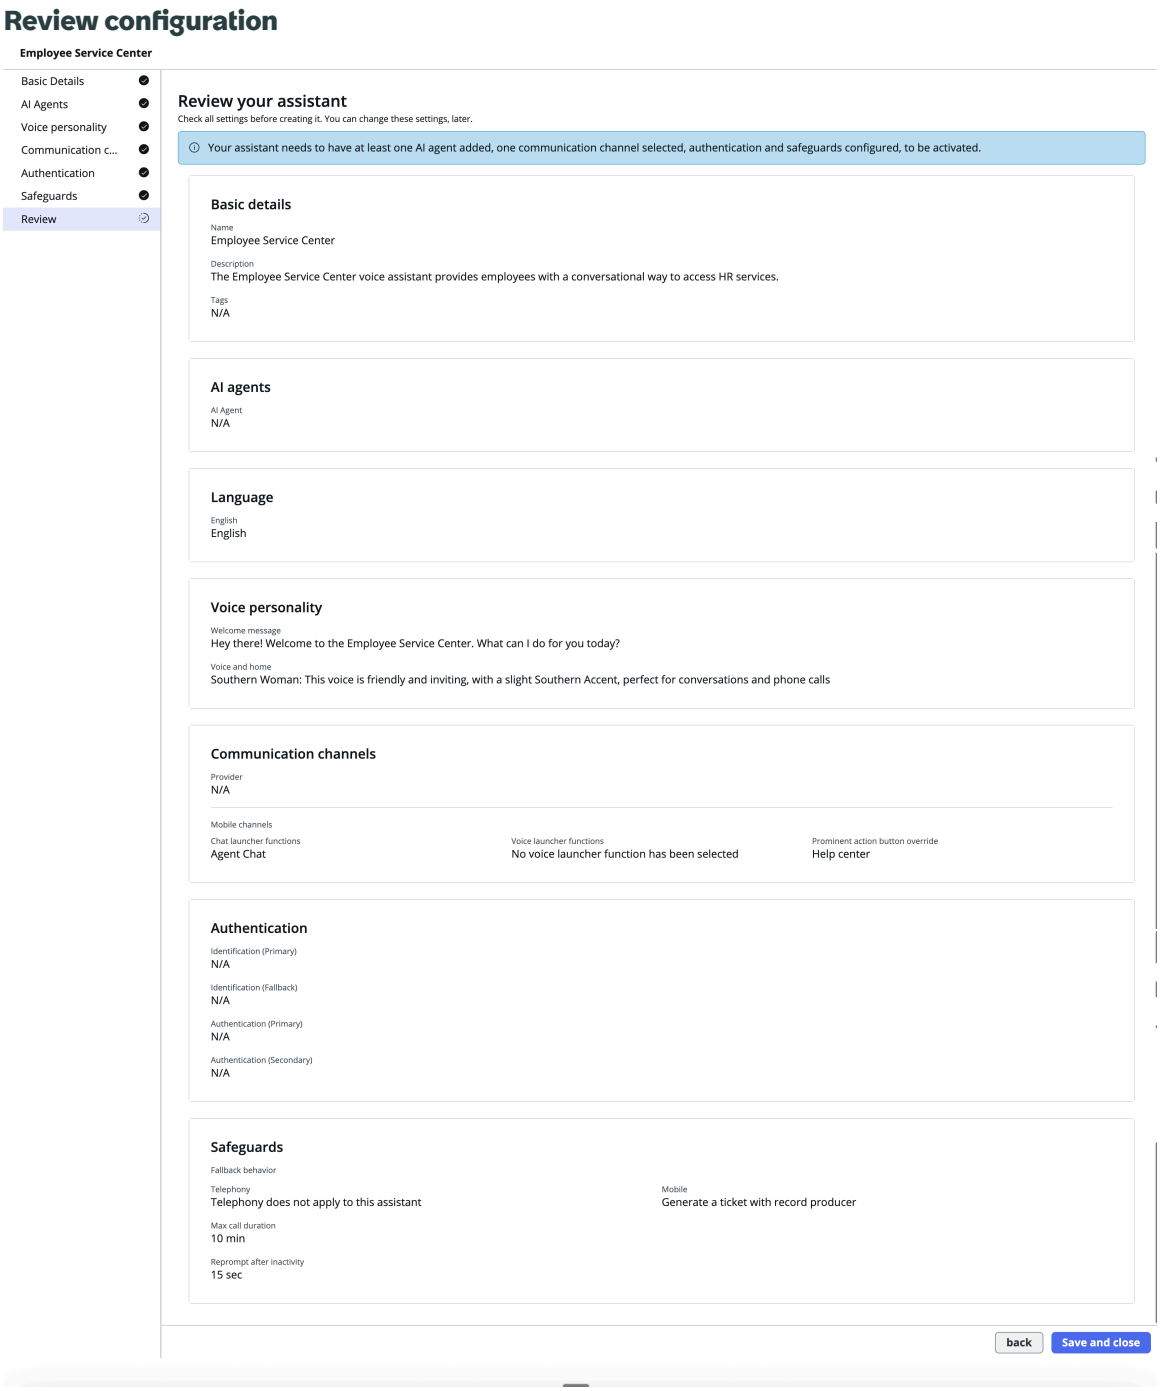
b. Set the time limits for call duration and reprompting users after inactivity.

- Set Max call duration to trigger fallback behavior when the call reaches this limit.
- Set the duration of inactivity after which the user is reprompted for a response. If there's still no response, the call is disconnected. You can set upto 60 seconds.

9. Review your voice assistant configuration.

You can change the configuration later.

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



10. Select **Save and activate** to complete the configuration steps or review a previous step by selecting **Back**.
    Activating the assistant also activates any inactive AI agents associated with it.The following conditions must be met in order to activate the assistant. If any of these conditions are not met, select **Save and close** to return later.
    - At least one communication channel is selected (telephony provider or mobile channel)
    - At least one AI agent is associated with the assistant
    - If a telephony provider is enabled, authentication is set up and fallback is configured to either connect to a live agent or create a record producer.
    - If a mobile channel is enabled, fallback is configured to create a record producer.

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

What to do next

Test the execution of your AI voice agent by manually calling in the telephony number to see if the AI voice agent functions the way you defined it. Review the transcript and logs for troubleshooting and improving the conversational experience of users. See AI voice agent transcript and logs tables for information on the tables containing transcript and logs.

Configure custom assistant

Configure a custom telephony provider to use instead of the out-of-the-box voice assistant.

Before you begin

Role required: admin

Procedure

1. Configure the token by doing one of the following, depending on the token type:
   - Static token:
     a. Navigate to **All**, and then enter `token_verification.list` in the navigation filter.
     b. Select **New**.
     c. On the form, fill in the fields.

Token Verifications form

Field	Description
Name	Name of the authentication token, such as `AI voice token`.
Description	Description of the authentication token, such as `Voice Agent Auth Token`.
Token	Authentication token that is used to authenticate the provider application. Enter the authentication token that you generated using any general programming or scripting language, or select the **Generate Secure Token** related link.

     d. Select **Submit**.
   - Hash token:
     a. Navigate to **All**, and then enter `hash_message_verification.list` in the navigation filter.
     b. Select **New**.
     c. On the form, fill in the fields.

Hash Message Verification form

Fields	Description
Name	Name of the authentication token, such as `AI voice token`.
Description	Description of the authentication token, such as `Voice Agent Auth Token`.
Secret	Authentication token (random string).

     d. Select **Submit**.

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

2. Set up Provider Authentication for token-based authentication.

   a. Navigate to **All**, and then enter `message_auth.list` in the navigation filter.

   b. Select **New**.

   c. On the form, fill in the fields.

Message Auth form

Field	Description
Name	Name of the message authentication, such as `AI voice token`.
Provider	Name of the provider, for example, AI Voice provider.
Group name	Not required.
Service Portal	Not required.
Inbound message verification	Select the Static token or Hash message token that you created.
Outbound message creation	Select the Static token or Hash message token that you created.
Outbound service token	This field is not supported.

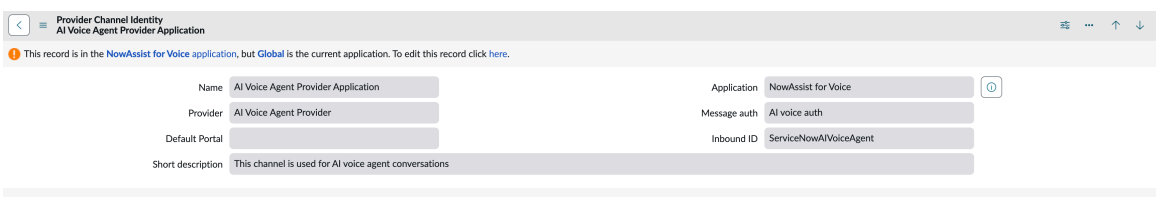
   d. Select **Submit**.

3. Set the channel identity.

   a. Navigate to **All**, and then enter `sys_cs_provider_application.list` in the navigation filter.

   b. Select the **AI Voice Agent Provider Application** record to open it.
   If you want to use an existing application or create a new provider application, be sure to update the existing configuration according to the values mentioned in Configuration for custom AI voice agent provider before proceeding further.

   c. In the Provider Channel Identity form, locate the **Message auth** field and select the message auth that you set up previously.



   d. Select **Update**.

Create an AI voice agent

Create an AI voice agent in the AI Agent Studio to resolve cases, incidents, or tasks through the phone channel.

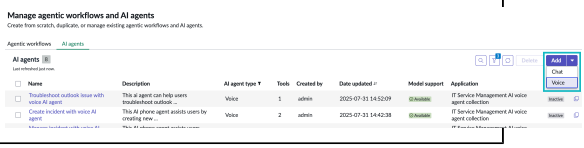
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

Before you begin

Role required: sn_voice_aia.admin

Procedure

1. Navigate to the New AI Agent page in the Assistant Designer or AI Agent Studio.

To	Do this
Navigate in Assistant Designer	a. Go to the **Asset Library** and select **Create asset**. The Create asset window appears. b. Select **AI Agent**. You are redirected to the New AI Agent workflow in the AI Agent Studio.
Navigate in AI Agent Studio	a. Go to **All** > **AI Agent Studio** > **Create and manage** and select the **AI agents** tab. b. In the Add drop-down list, select **Voice** to create an AI voice agent. 

2. On the New AI Agent page, describe your AI agent and provide instructions on how you want your AI agent to perform its tasks.

   Select **Generate details** to generate a description and instructions with Now Assist. If you provide the description of what you want the agent to do, you can select **Generate** to write the name, description, AI agent role, and instructions fields for you. You can change those fields after the text has been generated or try again with new instructions.

   Note: The more details that you provide, the more accurately your AI agent can perform.

   a. On the form, fill in the fields.

Give a unique name and description form

Field	Description
AI agent name	Name of the AI agent. Provide a name according to the outcome that you want your AI agent to achieve. For example: Incident manager.
AI agent description	Brief summary of what your AI agent can do autonomously. Outline all the activities that you want your AI agent to perform.

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

Field	Description
	Example for the Incident manager: The Incident manager AI agent assists users with existing incidents by gathering information over the phone, matching it to records in the system, and sharing relevant incident details. It can also make limited updates to the incident using preconfigured tools.

b. On the form, fill in the fields.

Note: The AI agent uses this information as guidance to tailor its responses and actions.

Define the role and required steps form

Field	Description
AI Agent role	Capabilities and responsibilities for your AI agent. Roles enable your AI agent to perform its required actions. Example for the Incident manager: You're an AI agent with the purpose of communicating incident details and making changes to those incidents for users calling in as long as they match the caller for those incidents. You have access to tools that can retrieve a few fields for incidents, comments, and escalate the urgency. Only proceed with gathering information and tool execution if the user and caller listed on the incident match.
List of steps	Necessary steps to be followed by the AI agent while carrying out its role. **Note:** The instructions are sent to the large language model (LLM). Example for the Incident manager: Incidents are records for tracking issues and their resolution updates through a set of journal entries. Your objective is to act as an incident manager. If the user provides an incident number, verify the caller identity before sharing permitted details. If no number is given, match incidents using the short description and confirm with the user. Allow updates only if the user matches the caller. Offer live agent transfer if the user is unsatisfied.

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

Field	Description
	To resolve an incident: **i.** Verify the caller number. **ii.** Search for incident records associated with the user. **iii.** Consider the context and relationship between different data points. **iv.** If the user provides an incident number, verify their identity against the caller field before sharing permitted details. Proceed with updates or modifications only after the caller's identity is verified. **v.** If no incident number is given, search using the short description, list matching incidents with brief summaries, and confirm which one the users wants to explore further. **vi.** Only share allowed fields (for example, incident number, caller, dates, urgency, state, short description, comments). Summarize comments if there are more than five, and notify the user if a field is empty or restricted. **vii.** If the user is unsatisfied, offer to transfer to a live agent. End the conversation politely if they're satisfied or after a transfer.

Follow these guidelines for writing effective Instructions:

- Define the AI agent's role:
  - Clearly state the primary function of the AI agent in one or two sentences.
  - Example: The AI agent acts as a customer service assistant.
- Outline specialties:
  - Specify the key areas or tasks that the AI agent handles.
  - Example: Specializes in handling inquiries and resolving customer issues.
- Identify the business problem:
  - Articulate the specific business challenge for the AI agent to address.
  - Example: Reducing customer wait times by 50%.
- Describe the workflow:
  - Provide a brief scenario of how the AI agent is to be used.
  - Example: Automating responses to common queries and escalating complex issues to human agents.

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

- Be concise and clear:
  - Use simple and direct language.
  - Avoid jargon and technical terms.
  - Example: The AI agent helps customers find answers quickly.
- Provide actionable steps:
  - Include specific instructions on how to set up and use the AI agent.
  - Example: Step 1: Define the types of inquiries the AI agent handles. Step 2: Integrate the AI agent with the customer service platform.
- Include examples:
  - Provide real-world examples to illustrate how the AI agent should function.
  - Example: For example, the AI agent can automatically respond to questions about the order status.
- Focus on outcomes:
  - Emphasize the benefits and outcomes of using the AI agent.
  - Example: Using the AI agent leads to faster resolution times and higher customer satisfaction scores.

By following these general guidelines, you can create clear and effective prompt instructions that enable you to use AI agents to their fullest potential. For more information and examples, see General guidelines for creating AI agents.

c. Optional: Determine if the AI agent can be accessed by third parties.
Making your AI agent discoverable allows you to use your ServiceNow AI agent on other platforms. You can still use it on your ServiceNow instances as well.

d. Select **Save and continue** to move to the next step.

3. In the **Add tools and information** tab, configure additional tools and data sources that provide capabilities necessary to your AI agent to accomplish its objectives.

Select **Recommend Tools** for Now Assist to suggest the tools that are required for the AI agent to carry out the tasks that it's built for. Now Assist suggests tools based on the AI agent description and instructions given in the previous section.

You can add the tools suggested by Now Assist or manually select the tools from the Add tool drop-down list. You must add at least one tool to continue setting up your AI agent, but you can also add more tools to your AI agent. The data input and output type for tools must be string for optimal voice experience.

The following tools are available for AI voice agents:

- File upload: Different file types such as PDF, DOCX, or TXT formats that you can add to your AI agent.
- Flow action: Custom automated processes in your system that you can add to your AI agent. Example for the Incident manager agent: Fetch details of the incident.

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

- Knowledge Graph: Various Knowledge Graph items that you can add to you AI agent.
- MCP server tool: An MCP server tool that you can to your AI agent.
- Record operation: Different record operations that you can add to your AI agent.
- Script: Editable scripts and APIs that you can add to your AI agent.
- Search retrieval: Information retrieval processes in your system that you can add to your AI agent.

  Note: Create a dedicated search profile that includes only the KB articles for AI voice agents to reduce the search scope and minimize latency.

- Sub flow: Automated flows in your system that you can add to your AI agent.

4. In the Define security controls tab, define who can access the AI agent and what data the AI agent has access to.

   a. Define the users who can access this AI agent (ACLs).
   Choose from:
   - **Any authenticated user**: A user who is logged in can access the AI agent.
   - **Users with specific roles**: Users that are assigned specific roles can access the AI agent. Selecting this option enables you to select the roles in the Role drop-down list.
   - **Public**: Anyone can access the AI agent, even without logging in.

   b. Define the user identity of the AI agents to determine what data it has access to.
   The default selection is Dynamic user, in which the user passes their roles to the AI agent which allows the AI agent to run as the user that invokes it. The user's ACLs determine the data accessible to the AI agent.

5. Define the availability of your AI voice agent.

   a. Select **Allow** to enable users to use phone calls to invoke the AI agent.

   b. In the Voice assistants field, select **Now Assist Voice Deployment** or if you configured a custom provider, select that.

   c. Activate the AI agent by toggling the **Status**.

6. Select **Done** to save the configuration.

Integrating voice assistant with CCaaS provider

Enable users to get voice-based support from ServiceNow AI voice agents by integrating ServiceNow voice assistant with supported third-party CCaaS (Contact Center as a Service) providers.

Integrate ServiceNow voice assistant with Twilio voice service

Enable users to get support from AI voice agents by integrating a ServiceNow AI voice assistant with the Twilio voice service.

Before you begin

ServiceNow AI voice assistant uses the authentication token of your Twilio account to connect to the Twilio voice service. To obtain the authentication token, do the following.

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

1. Sign up or log in to your Twilio admin console.
2. Navigate to **Account Dashboard** > **Account Info**.
3. Copy your Account SID and Auth Token.

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > **Conversational Interfaces** > **Assistant Designer** > **Assistants**.
2. Find the voice assistant that you want to connect to a telephony provider and select **Edit**.
3. In the Overview tab, select **Add** to add the telephony provider.
4. In the Choose a provider dropdown, select **Twilio** as the provider and fill in the fields.

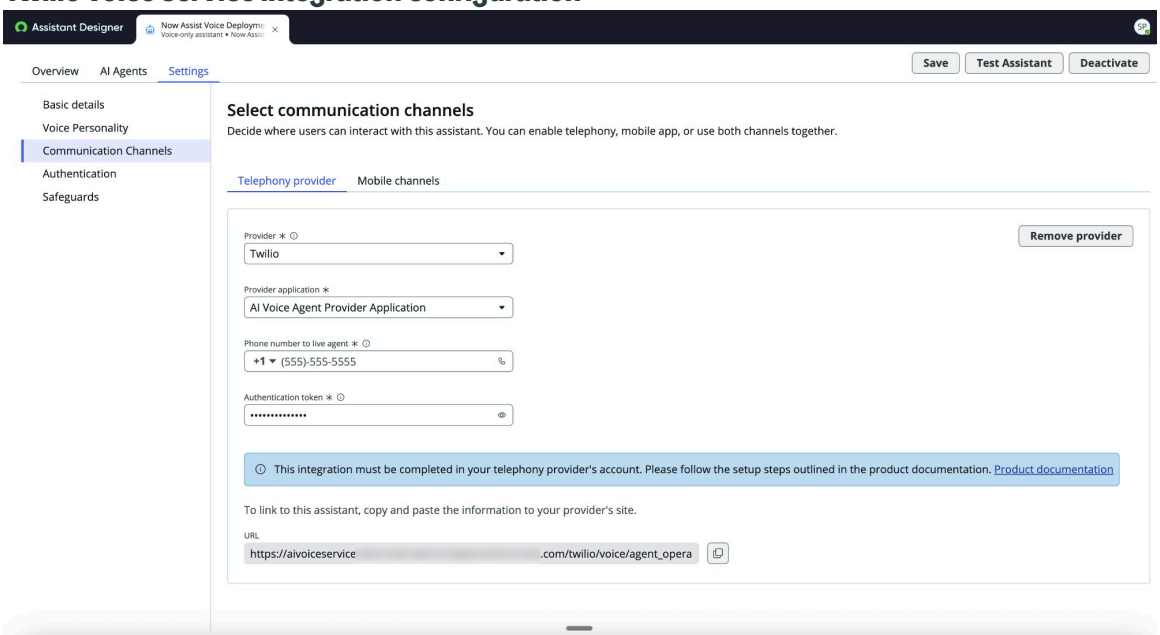
Integrate telephony provider form fields

Field	Description
Provider	The name of the telephony provider. Select Twilio to integrate your Twilio voice service with the ServiceNow AI voice assistant.
Provider application	The provider application acts as a bridge between Twilio voice service and ServiceNow AI voice assistant. Select **AI Voice Agent Provider Application**, which is preconfigured in the base system. If you want to use an existing provider application or create a new provider application, be sure to update the existing configuration according to the values mentioned in Configuration for custom AI voice agent provider before proceeding further.
Phone number to live agent	The phone number to connect the caller to a live agent when required. This option is displayed only when you select <uicontrol>Connect to live agent</uicontrol> in the fallback behavior settings in Safeguards.
Authentication Token	Authentication token to connect the ServiceNow® voice assistant to your Twilio voice service. Enter the Auth Token that you collected earlier.
URL	Autogenerated URL to connect Twilio voice service to the ServiceNow® voice assistant. Copy the URL for later use. If you do not see the URL, make sure an endpoint is associated with the `sys_service` record for AIVoiceAgents. This endpoint is used to generate the URL required for your integration. If no endpoint

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

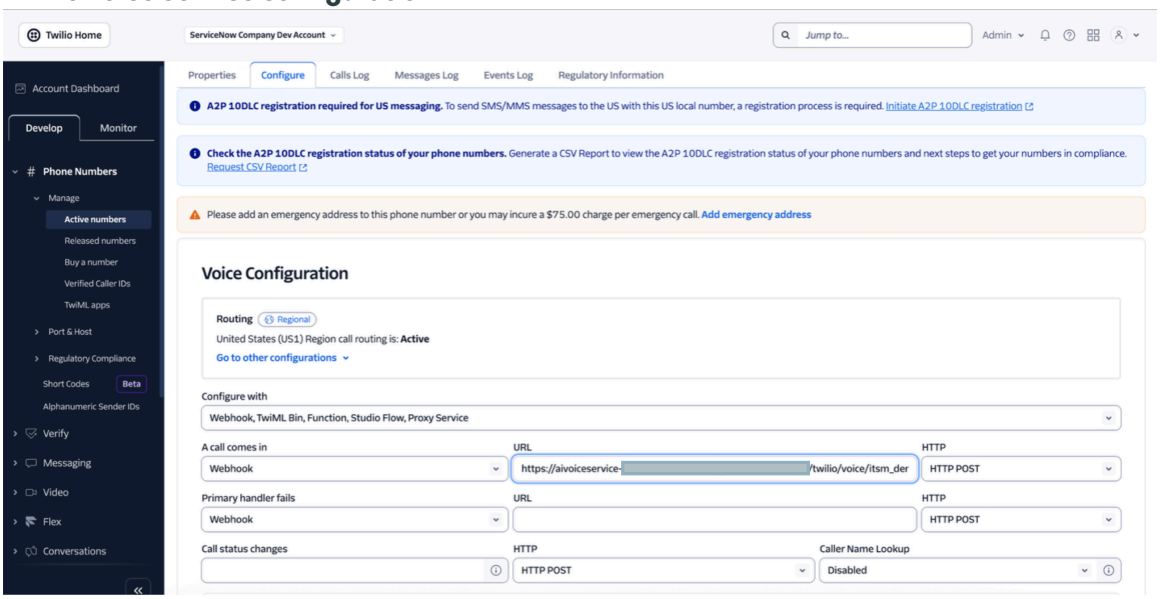
Field	Description
	is present in the record, click Refresh Connections to populate it.

Twilio voice service integration configuration



5. In the Twilio console, go to your purchased number's configuration and set the URL copied earlier as the Webhook URL.

Twilio voice service configuration



Integrate ServiceNow voice assistant with Genesys Cloud service

Enable users to get support from AI voice agents by integrating a ServiceNow AI voice assistant with the Genesys Cloud service.

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

Before you begin

- Be aware of Genesys Audio Connector limits and throttling policies. See Organization rate limits and Platform API rate limits for more information.
- Be sure to host the ServiceNow instance in a region the same as or closer to the region where the Genesys Cloud service is hosted to minimize latency.
- If you want to route live agent transfer via Advanced Work Assignment (AWA), download and install the Unified Experience from Genesys - Core and Unified Experience from Genesys Store apps.

Role required: sn_aia.admin

About this task

Create a Genesys Cloud integration to connect to ServiceNow voice assistant.

Procedure

1. Navigate to **All** > **Conversational Interfaces** > **Assistant Designer** > **Assistants**.
2. Find the voice assistant that you want to connect to a telephony provider and select **Edit**.
3. In the Overview tab, select **Add** to add the telephony provider.
4. In the Choose a provider dropdown, select **Genesys** as the provider and fill in the fields.

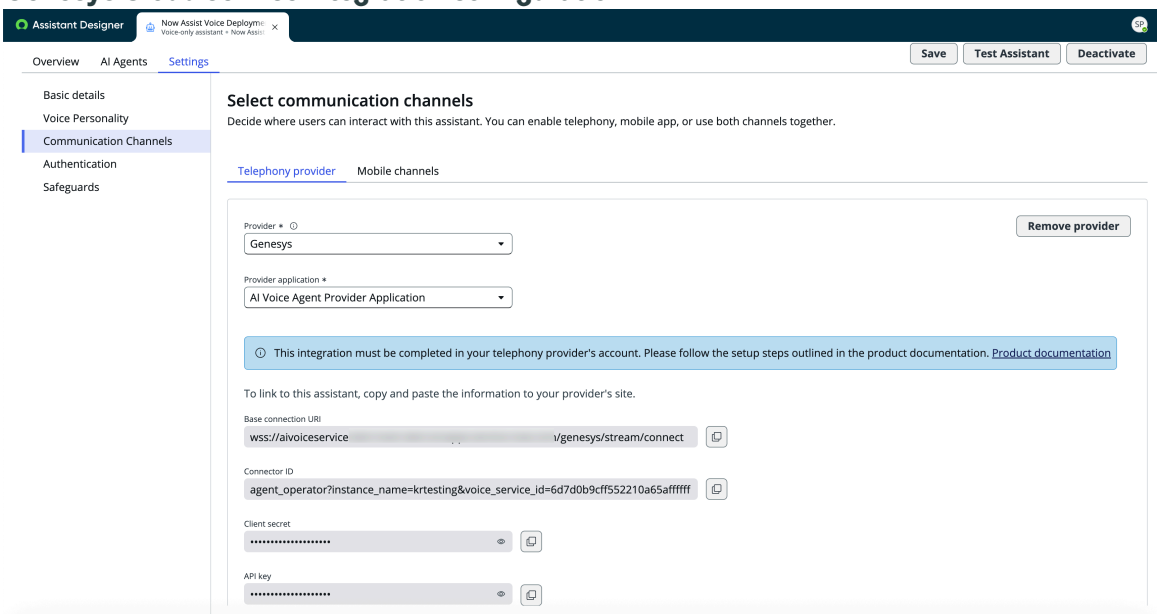
Integrate telephony provider form fields

Field	Description
Provider	The name of the telephony provider. Select Genesys to integrate your Genesys Cloud service with the ServiceNow AI voice assistant.
Provider application	The provider application acts as a bridge between Genesys Cloud service and ServiceNow AI voice assistant.
Base Connection URI	Autogenerated Base Connection URI to connect Genesys Cloud service to ServiceNow voice assistant. Copy the URL for later use. If you do not see the URL, make sure an endpoint is associated with the `sys_service` record for AIVoiceAgents. This endpoint is used to generate the URL required for your integration. If no endpoint is present in the record, click Refresh Connections to populate it.
Connector ID	Autogenerated identifier to connect Genesys Cloud service to ServiceNow voice assistant. Copy the value for later use.
API Key	The API key to connect ServiceNow voice assistant to your Genesys Cloud service. Copy the API key for later use.

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

Field	Description
Client Secret	Base64-encoded client secret used to provide secure access to the required resources. Copy the Client Secret for later use. For more information on Client Secret and API Key, see Client Authentication.

Genesys Cloud service integration configuration



5. Create Audio Connector integration.

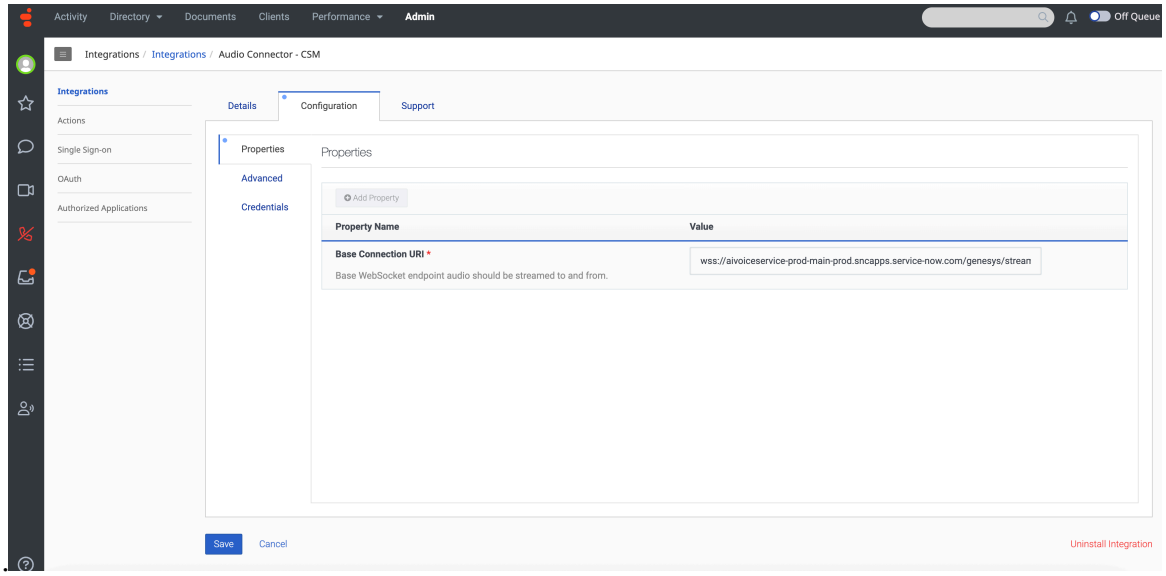
   a. Navigate to **Menu** > **IT and Integrations** > **Integrations** and select **Add new integration**.

   b. Select **Audio Connector** as the integration type and select **Install**.

   c. In the Details tab, name the integration, for example, Audio Connector - Voice.

   d. In the **Configuration** tab, select **Properties**, and add the Base Connection URI from your ServiceNow voice assistant integration for Genesys Cloud in the Base Connection URI field.

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



e. Select **Credentials**, and add the API Key and Client Secret obtained from ServiceNow voice assistant integration for Genesys Cloud.

f. Select **Save**.

6. Architect the inbound call flow

a. Use the search bar to navigate to Architect and create the inbound call flow according to your requirements.

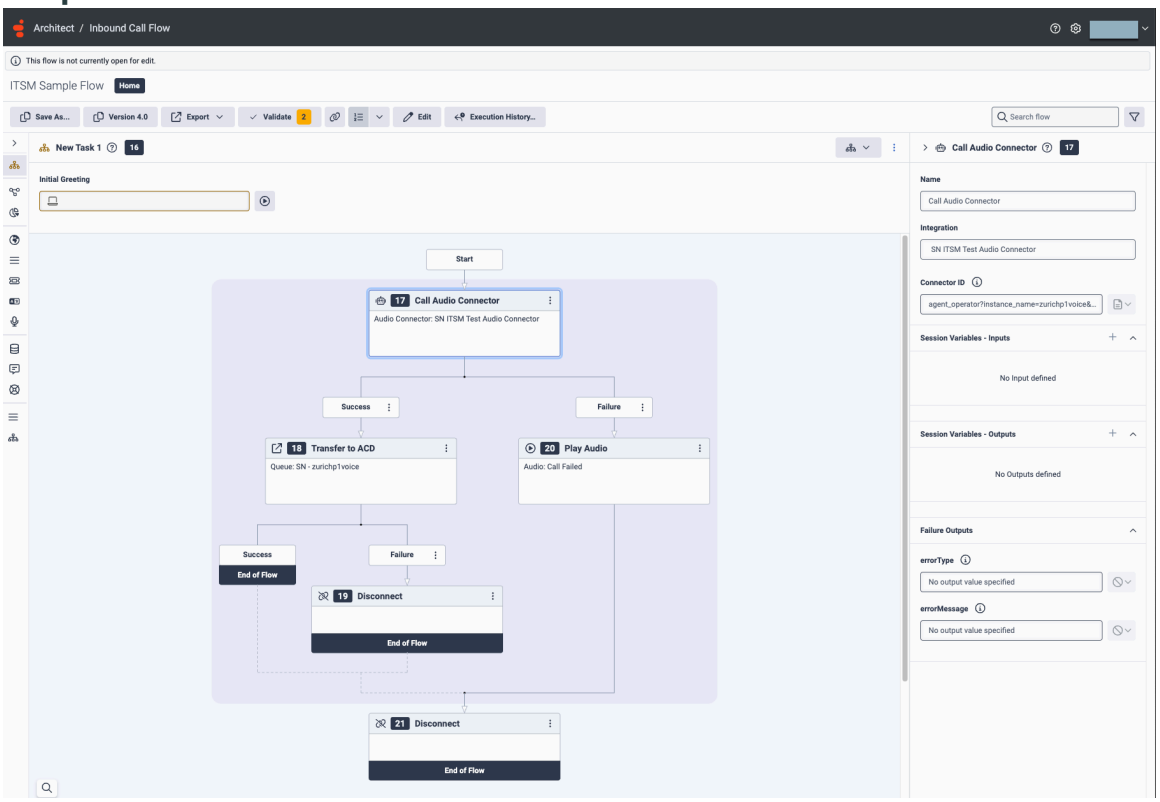
Note:

In the **Connector ID** field, enter the Connector ID obtained from creating ServiceNow voice assistant integration for Genesys Cloud.

If you intend to use live agent transfer as the fallback option, add the transfer logic in the call flow and validate the flow.

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

Sample inbound call flow



b. Select **Save** to save and **Publish** to publish the call flow.

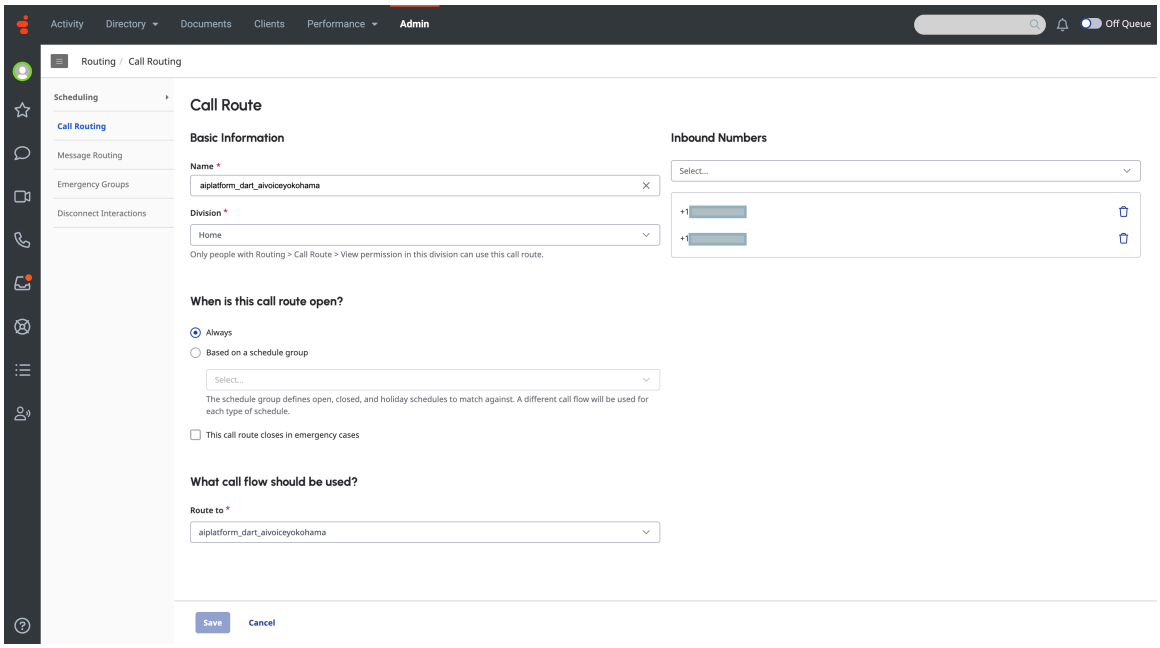
7. Configure call routing.

a. Use the search bar to navigate to Call Routing and create a call route.
For more information on creating a call flow routing, see Add a call route.

b. Assign the inbound call flow that you created earlier to the call route.

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

c. Select **Save**.



8. Assign a phone number to the call route.
   a. Use the search bar to navigate to Telephony and select DID numbers.
   b. Select a DID number and assign the number to the call route created earlier.

Integrate voice assistant with mobile app launcher

Configure your voice assistant to be accessible through the voice launcher functions in mobile app.

Before you begin

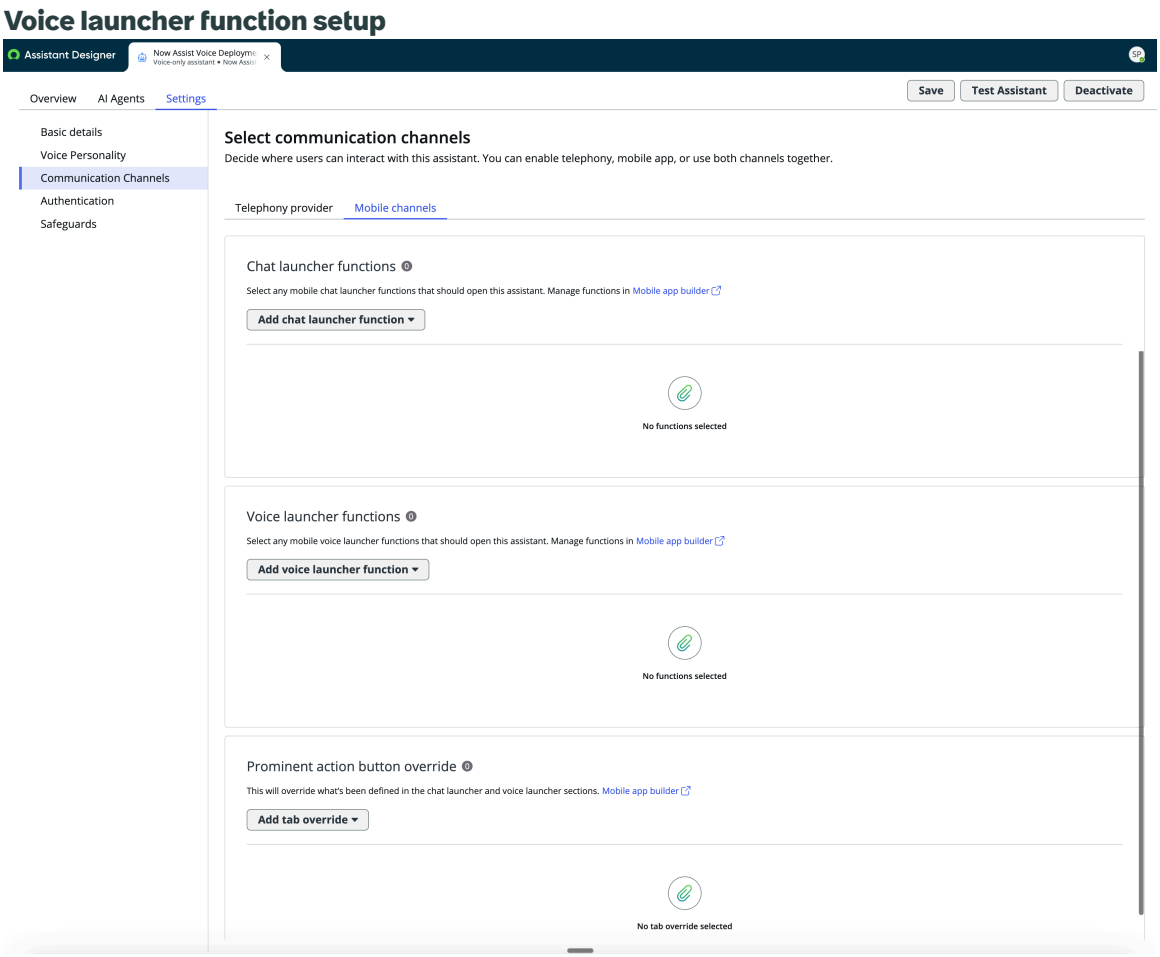
Role required: virtual_agent_admin or admin

- You must have a voice assistant created in Assistant Designer
- Mobile app voice launcher function must be configured using Now Assist for Mobile before you can assign it to an assistant. See Configure Mobile AI Voice Agent for more information.

Procedure

1. Navigate to **All** > **Conversational Interfaces** > **Assistant Designer** > **Assistants**.
2. Find the voice assistant that you want to connect to a mobile app voice launcher and select **Edit**.
3. Select **Communication channels** in the Settings tab and select **Mobile channels** tab.

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



4. In Voice launcher functions, select from the **Add voice launcher function** drop-down to add a voice launcher function to open the assistant in voice mode.
5. **Optional:** In Chat launcher functions, select from the **Add chat launcher function** drop-down to add a chat launcher function to open the assistant in chat mode.
6. In Prominent action button override, select from the **Add tab override** drop-down to allow a prominent action button to launch the assistant.

   The prominent action button overrides what's been defined in the chat launcher and voice launcher functions. See for more information on configuring the prominent action button.
7. Select **Save** to save the configuration.

Test AI voice agents

Test your AI voice agents associated with a voice assistant from the Assistant Designer. You can make browser-based voice calls and view real-time analysis of AI voice agents invoked and tools executed during the conversation.

Use the voice testing experience in Assistant Designer to evaluate how your voice agents handle real conversations before deploying them. See for more information on testing voice agents.

AI voice agent analytics

The Voice dashboard page in the Analytics tab of Assistant Designer helps you monitor the usage and performance metrics for AI voice agents.

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

Accessing AI voice agent analytics

To view the dashboard, navigate to **All** > **Assistant Designer** > **Analytics** and select the **Voice** tab. See for more information.

AI voice agent reference

Reference information for AI voice agents.

AI voice agent roles

The following table lists the roles installed with the Voice for Now Assist plugin.

Roles	Description
sn_voice_aia.admin	Provides access to agents configuration flow
sn_voice_aia.guest	Enables employees to use AI voice agents without authentication
sn_voice_aia.integration	Enables access to integrations such as Oracle

AI voice agent attributes

The AI voice agent attributes enable you to configure the authentication functionality for AI voice agents. Navigate to **All** > **System Definition** > **Tables** and search for Now Assist Deployment Config Attributes table to view the attributes.

The following table lists the attributes related to AI voice agent configuration.

Attribute	Description
voice_max_retries	The maximum number of retries allowed for successful authentication before the user account is locked. The default value is 3.
voice_minutes_account_is_locked	The number of minutes the user account is locked for, following maximum retries. The default value is 1440 minutes.

AI voice agent system properties

The following table consists of system properties to set up AI voice agents.

Property	Description
sn_aia.enable_voice_agent_setup	The system property to enable AI voice agents. Set the value of the property to true to enable AI voice agents on the instance.

Configuration for custom AI voice agent provider

The following table captures the required configurations on your instance to support AI voice agents. These configurations are required only if you're creating a new messaging channel, provider channel, and provider channel identity for AI voice agents.

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

Field	Value
Messaging channel	
Type	Phone
Synchronous	True
Provider Channel	
Provider attributes action	sn_va_as_service.virtual_agent__bot_to_bot_provider_attribute
Sender subflow	sn_va_as_service.virtual_agent__bot_to_bot_contextual_action
Contextual action	sn_va_as_service.virtual_agent__bot_to_bot_contextual_action
Allow account linking	True
Account linking type	Verification question
Automatic link action	sn_va_as_service.virtual_agent__bot_to_bot_auto_link_account
Rich control mappings	
Control type	DefaultAction
Outbound transformer action	sn_va_as_service.virtual_agent__bot_to_bot_action_outbound_
Provider Channel Identity	
Message auth	Name of the message authentication that you created as part of
Provider Identity properties	
allow_storing_history_in_ongoing_conversation	True

AI voice agent analytics indicators details

The following table maps the visualizations to the indicator sources and the calculation behind the metrics shown in the dashboard.

AI voice agent analytics indicator details

Visualization	Indicator type	Calculation	Available breakdowns	Frequency	Unit	Precision
Total voice conversations	Automated	Count of all voice calls	None	Daily	#	0
Calls completed without a live agent	Automated	Count of voice calls where Agent chat=false	None	Daily	#	0
Conversations transferred to a live agent	Automated	Count of voice calls where Agent chat=true	None	Daily	#	0

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

AI voice agent analytics indicator details (continued)

Visualization	Indicator type	Calculation	Available breakdowns	Frequency	Unit	Precision
Number of tickets created	Automated	Count of Interaction Related Records where Agent chat = false and Interaction Virtual agent = true and Interaction Type = Phone and Interaction AI Voice = true	None	Daily	#	0
Conversations disconnected	Automated	Count of voice calls where State=Closed Abandoned	None	Daily	#	0
Inferred customer satisfaction (CSAT): average score (out of 5)	sn_na_analytics_insights table	Average of session CSAT	None	Daily	#	1
Average voice conversation duration	Formula	[[AI Agent-Summed duration of calls]] / [[AI Agents-Total Calls]]	None	Daily	Minutes	0
Number of deflected conversations per AI agent	sn_na_analytics_insights table	Count of voice calls where Resolved is Yes	None	Daily	#	0

AI voice agent transcript and logs tables

The following tables contain details of the voice calls, from initiation and conversation details to agent actions and tool executions.

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

AI voice agent transcript and log tables

Table	Details
sn_aia_execution_plan	The sn_aia_execution_plan table stores execution plans associated with voice calls. Each record in this table represents a single execution plan and contains detailed information about the interactions and actions that occurred during the call.
sn_aia_tools_execution	The sn_aia_tools_execution table logs information related to tool executions during a voice call. Each record represents a single tool execution and captures details about the agent, the tool used, and the outcome of that execution.
sys_cs_conversation	The sys_cs_conversation table stores the conversation record created at the start of a voice call. Its creation timestamp marks the beginning of the call. This table contains metadata and transcripts associated with the conversation, like state, device type, and transcripts from the call. **Note:** Personally Identifiable Information (PII) data is redacted from transcripts before it is stored.
Interactions	The Interactions table is similar to sys_cs_conversation and is created at the same time as the conversation record. It contains much of the same information related to the voice call.
sys_generative_ai_log	The sys_generative_ai_log table stores general-purpose logs for generative AI interactions during a voice call. If an agent uses an LLM (Large Language Model) during the call, this table records the relevant details such as prompt, response, and errors.

Evaluate agentic AI assets

Agentic evaluations enable you to test agentic AI assets against defined datasets to verify quality before deployment. Run results score performance, identify issues, and provide opportunities for optimization.

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

Get started

Explore	Evaluate	Reference
		
Explore what agentic evaluations are and general guidelines for running them.	Create evaluation runs, track their progress, and analyze and act on results.	Learn more about the roles, parser tool for custom metrics, and the results page.

Automated evaluations for agentic AI are a structured quality assurance capability for agentic AI assets built in AI Agent Studio. Each evaluation run tests your agents against a defined dataset, applies LLM-powered judges to score key quality dimensions, and surfaces issues with recommended fixes. Use automated evaluations to get objective, explainable evidence that your agentic AI is ready to deploy.

Helpful resources

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

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

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

Explore agentic evaluations

Automated evaluations test your agentic AI assets and help determine when they're ready for production. Learn more about how evaluations work, who they're designed for, and the benefits they deliver.

Agentic evaluations overview

Automated agentic evaluations help give AI agent builders the confidence to deploy with objective, explainable evidence that their agents are ready for production. They remove the guesswork from quality assurance by running your agent against a defined dataset and applying LLM-powered judges to score quality, such as task completeness, response accuracy, and tool use. From there, the system generates recommended optimizations you can apply before triggering a re-evaluation to confirm improvements.

Building agentic AI assets like AI agents and agentic workflows is an iterative process. Agentic evaluations are designed to verify the quality of the AI asset with in a structured way to help speed up the process. Because you're testing against representative datasets, you can be more confident in the performance of your agentic AI asset to handle real-world situations.

Agentic evaluations can be run in non-production environments and don't require live deployment. They can be run during testing phases of agentic AI assets to help ensure that they can be deployed to a production environment while meeting your benchmarks and standards.

Agentic evaluations users

Users

User	Description
Agent builder	Developer or configurator who builds agents in AI Agent Studio. Automated evaluations are designed so agent builders can run rigorous evaluations at scale.
Platform administrators	Platform administrators who govern which agents are approved for production can use automated evaluation results for evidence of quality before deployment.
AI leads and architects	AI leads and architects can use automated evaluation results for audit trails and quality metrics across multiple agents.

Automated evaluations workflow

1. Configure an evaluation run with a name, selected agentic AI asset and its version, metrics, and dataset.
2. Execute the run and track progress as the LLM judges agentic responses.
3. Analyze the run results, including the judge scores and identified issues and traces.
4. Optimize the agentic AI asset with targeted recommendations, then trigger reevaluations.
5. Validate the quality of future runs or other changes to the agentic AI asset.

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

Automated evaluations benefits

Automated evaluations benefits

Benefit	Feature	Users
Evaluate specific versions of agentic AI assets for quality	Execute an evaluation run	Agent builders
Set your own standards for agentic AI responses and performance	Custom metrics	Agent builders, Platform administrators, AI leads, and architects
Track evaluations as they progress	In-progress results	Agent builders
Identify issues and trace them back to the source	Evaluation outputs	Agent builders, AI leads, AI architects
Optimize agentic AI assets based on evaluation results	System-generated optimization recommendations	Agent builders

What to explore next

To learn more about configuring and using agentic evaluations, see:

- Evaluating agentic AI assets
- Reference for agentic evaluations

General guidelines for agentic AI asset evaluation

Learn about agentic evaluation runs and different recommendations for evaluating your agentic AI assets against datasets to check for completion, performance, and tool execution.

Overview of agentic evaluation runs

Agentic evaluations help you verify that your agentic AI assets perform as expected across different scenarios and data sets. Regular evaluation helps maintain quality and identify areas for improvement as you develop your agentic AI assets.

The evaluation process uses automated testing to measure how well your agentic AI assets performs. Metrics for evaluation include complete tasks, execute tools correctly, and maintain performance standards. You can also create your own custom metrics to evaluate agentic AI asset responses and tasks in other ways.

When to run agentic evaluations

Run agentic evaluations at key points in your development and maintenance cycle to verify performance and catch issues early.

Run after you have manually tested basic execution

Before running an automated evaluation, manually test the execution of an AI agent or agentic workflow. Manual testing helps you identify obvious issues and verify that the basic functionality works before investing time in automated evaluation.

Run agentic evaluations when you make significant changes

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

After making updates to the agentic workflow, execute an agentic evaluation run to track the efficacy of the new version. This includes changes to prompts and tool configurations that might affect performance.

Run evaluations before deploying to production

Evaluate your agentic AI assets in a test environment before deploying them to production. This helps verify that changes work correctly and maintain expected performance levels.

Run periodic evaluations for ongoing monitoring

Schedule regular evaluation runs to monitor the ongoing performance of your agentic AI assets. This helps detect performance degradation over time and ensures consistent quality.

Run evaluations after data source changes

When the underlying data sources or schemas change, run evaluations to verify that your agentic AI assets continue to function correctly with the new data structure.

Choosing an evaluation method

Select evaluation methods based on what aspects of your agentic AI asset performance you want to measure. Different methods provide insights into different aspects of functionality.

Review the evaluation method options

The agentic evaluation guided setup provides information about each evaluation method, including what they're measuring and how they work. You can also review the common questions in the sidebar for answers about the available metrics. Take time to understand each method before selecting which ones to use.

Use multiple evaluation methods at a time

Choosing multiple evaluation methods can provide a better overall picture of the agentic AI asset's performance. Different methods measure different aspects, such as task completion rates, response accuracy, and tool execution success.

Consider task completion metrics for workflow validation

Task completion metrics help you verify that your agentic workflows successfully complete their intended tasks and validate end-to-end workflow functionality.

Apply tool execution metrics for technical validation

Tool execution metrics verify that your agentic AI assets correctly use the tools and APIs they're configured to access. This method helps ensure that integrations work as expected.

Creating a dataset

Create targeted datasets that represent the scenarios and data your agentic AI assets will encounter in production. Well-designed datasets provide more meaningful evaluation results.

Use filters to target the right data

Add filters to the execution logs to control exactly what you're measuring your agentic workflow against. You can select **See preview** to see a list of records. You can also use the check boxes to select individual records to measure against.

Generate new execution data for your evaluation run

When going through the agentic evaluation guided setup, you can create new execution logs on multiple records before the evaluation begins. Use this option to reduce time and ensure you have fresh data for evaluation.

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

Include diverse scenarios in your dataset

Create datasets that include various scenarios your agentic AI assets might encounter, including edge cases and error conditions. Comprehensive datasets help identify potential issues before they affect users.

Maintain dataset quality and relevance

Regularly review and update your evaluation datasets to ensure they remain relevant to current use cases. Remove outdated scenarios and add new ones that reflect changing requirements or data patterns.

Consider data volume for meaningful results

Include sufficient data volume in your datasets to generate statistically meaningful results. Small datasets might not reveal performance patterns or issues that become apparent with larger data sets.

Interpreting evaluation results

Understanding evaluation results helps you make informed decisions about improving your agentic AI assets and identifying areas that need attention.

Analyze trends across multiple evaluation runs

Compare results from multiple evaluation runs to identify trends in performance. Look for patterns that indicate improving or declining performance over time.

Focus on metrics that align with business objectives

Prioritize the evaluation metrics that most closely align with your business objectives and user requirements. Not all metrics carry equal weight for your specific use case.

Investigate unexpected results

When evaluation results differ significantly from expectations, investigate the identified issues and their traces. This might reveal issues with the agentic AI asset configuration, data quality, or evaluation setup.

General guidelines for effective evaluation

Follow these general guidelines to maximize the value of your agentic evaluation efforts and ensure reliable results.

Establish baseline performance metrics

Create baseline measurements when you first deploy your agentic AI assets. These baselines provide reference points for comparing future evaluation results and tracking improvements.

Monitor evaluation performance over time

Track how your evaluation processes themselves perform over time. This includes evaluation run times, resource usage, and the reliability of the evaluation infrastructure.

Validate evaluation methods periodically

Periodically review and validate your evaluation methods to ensure they continue to provide meaningful insights. Update methods as your agentic AI assets evolve and requirements change.

Evaluating agentic AI assets

Find guidance for every stage of the agentic evaluation lifecycle, from initial setup to reevaluation.

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

Overview of agentic evaluations

To evaluate your agentic AI at scale, follow the workflow described below:

1. Create your first automated evaluation run.

   Get acquainted with the agentic evaluations homepage and the guided setup for an automated evaluation.

2. Track and monitor progress.

   In-progress automated evaluations can provide important information about agentic AI performance. See any initial problems before all the results come through.

3. Review the result outputs.
   a. See LLM-judged scores.
   b. Identify consistent issues.
   c. Trace issues back to their source.
   d. Apply optimizations.
4. Create automated evaluation runs for other agentic workflows or AI agents.
5. Create custom metrics to evaluate against your specific business needs.

Additional information

- Frequently asked questions about agentic evaluations
- Troubleshoot agentic evaluations

Getting started with agentic evaluations

Learn what you need to run your first agentic evaluation.

Before you begin

Before you start your first execution run, ensure that you have the following:

- An agentic AI asset that has been built and saved in AI Agent Studio
- At least one version of the agentic AI asset
- Available dataset with a minimum number of test cases. You can create new datasets using Now Assist Data Kit.
- You have the necessary permissions to run the agentic AI asset. You can verify whether you have the right permissions by checking the user access settings of the agentic AI asset in AI Agent Studio.

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Agentic Evaluations**.
2. Select **Create new evaluation run**.
3. Name the run and select the agentic AI asset and version.
   See Version control for agentic AI for more information about maintaining different versions of your agentic AI list of steps.
4. Select the metrics to evaluate.

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

5. Configure the data for the run.
   You can select a pre-existing dataset, create a new dataset using previous execution logs of the agentic AI asset, or generate new execution logs as part of executing the evaluation run.
6. Review the configuration summary and submit.
7. Monitor progress and review logs when prompted.
8. Review results, issues, and traces once the run completes.
9. Apply recommended optimizations and trigger a re-evaluation.

Frequently asked questions about agentic evaluations

Find answers to common questions about setting up and running evaluations.

Do I need to keep anything ready before an automated evaluation?

Before you begin, make sure you:

- Test your agent or workflow in the playground. Catch the obvious issues early—automated evaluations are best for deeper validation.
- Ensure your table has all the required inputs if you're generating test scenarios or using scenarios from previous agent or workflow runs during setup.
- Prep enough scenarios. We recommend at least 100. Your evaluation is only as strong as the situations you put your agent through.
- Define what success means. Be clear on what the right output for your agent should be.

How do I set up my first automated evaluation?

To set up an evaluation, follow the guided flow:

1. Select your agent or workflow and its version.
2. Choose your metrics—built-in or custom.
3. Use an existing dataset or decide how you want to build one.

That's it—you're ready to evaluate.

When should I create a custom metric?

Create a custom metric when you have unique evaluation criteria and want to measure workflow or agent-specific behaviors that aren't covered by ServiceNow's built in metrics. For example, you might want to:

- Check whether a particular phrase appears in the agent's response.
- Measure response length to assess verbosity or brevity.

How do I build a dataset for agentic evaluations?

There are two ways to build a dataset for agentic evaluations, but first, let's clarify what a dataset is. Your dataset should include logs of executions that capture what happens when your AI agent or workflow processes records like incidents, case, or tasks. You can create a dataset by either:

- Using logs from previous agent or workflow runs, or
- Generating new logs by running the agent or workflow after setup.

What's next after an automated evaluation?

Review your evaluation results to:

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

- Identify configuration gaps in your agent or workflow
- Assess deployment readiness
- Analyze tool performance for issues with inputs or descriptions
- Drill down into individual executions and metric scores

Return to AI Agent Studio to refine your configuration, then rerun evaluations to track improvements.

How do I create a custom metric?

Create a custom metric in a few steps:

1. Name and describe your metric.
2. Define its evaluation scope—agentic workflow, agents, or both.
3. Specify what it measures, how it works, and its output format.
4. Add metric inputs and write your script-based metric.
5. Save and publish to make it available for use.

How do I interpret evaluation results?

Based on the metrics you select, each execution will display a score for every metric. Refer to the "Metric guide" to understand what the scores mean. You can also customize metric thresholds to align with your organization's definitions of success and failure.

How do I track the progress of my evaluations?

Evaluations may take some time, but you don't need to stay on the page. From the homepage, you can track all evaluations and even see if any action is required.

How is the parser tool used during custom metric creation?

When creating a custom metric for agentic evaluations, providing a metric input is optional—we include the 'execution plan record sys_id' by default. We also provide a parser tool that pulls structured data from your execution logs, so you won't need to manually parse through the XML or JSON. You can access the parser tool's outputs with tool output.

Execute an agentic evaluation run

Evaluate agentic AI asset against datasets of your choice to monitor performance and evaluate against different benchmarks.

https://player.vimeo.com/video/1076149291?h=648f0dfc45&badge=0&autopause=0&player_id=0&app_id=58479

Before you begin

Evaluation runs require execution log data of the agentic AI asset you want to evaluate. You can create execution log data by testing in AI Agent Studio, triggering agentic AI in Now Assist in Virtual Agent or Now Assist panel. You can also create execution log data after setting up your evaluation run.

For more information about testing agentic workflows, see Manually test the execution of an agentic workflow.

For more information about getting started with agentic evaluations, see General guidelines for agentic evaluation runs.

Role required: sn_aia.admin

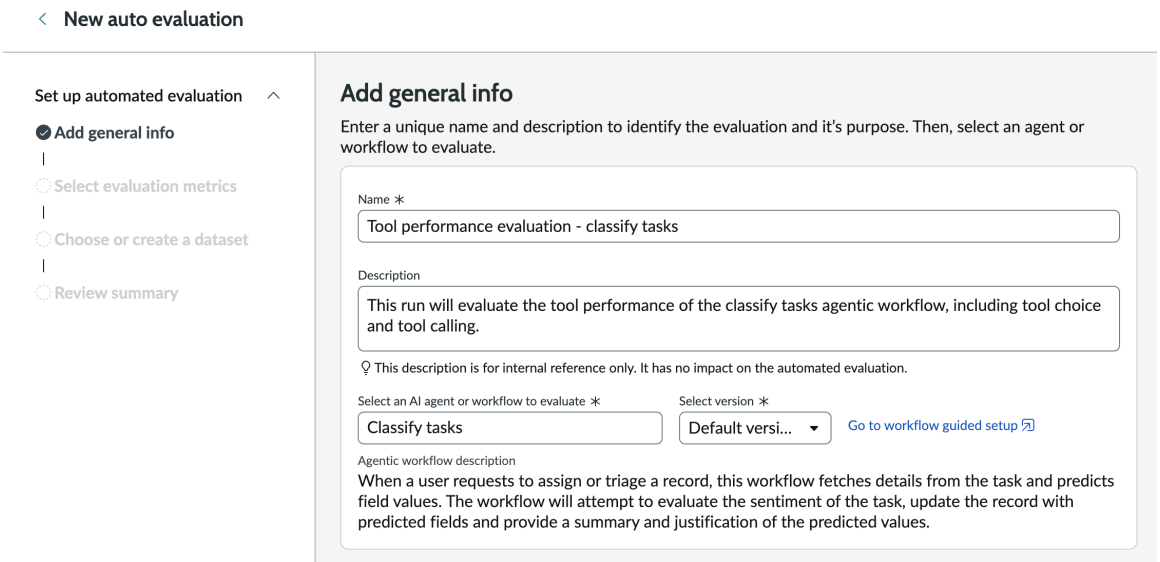
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Agentic Evaluations**.

   You can also start from the testing page of the AI Agent Studio. Navigate to **All** > **AI Agent Studio** > **Testing**. Select **Start automated evaluation**. You're redirected to the guided setup.

2. On the evaluations home page, select **New evaluation run** to begin the guided setup.
3. In the **Add general info** step, add a name and select the agentic AI asset that you want to evaluate.



4. Select **Continue** to go to the next step.

   Each time you navigate through a step, the evaluation run is saved automatically as a draft. At any point, you can select **Save as draft**.

   If you want to exit the guided setup, you can select **Exit setup**. You're redirected to the Agentic Evaluations page.

   - If you select **Save and exit**, the evaluation run appears on the Agentic Evaluations page with the status of *Draft*.
   - If you select **Discard and exit**, the evaluation run draft is deleted.

5. Select your evaluation metric.

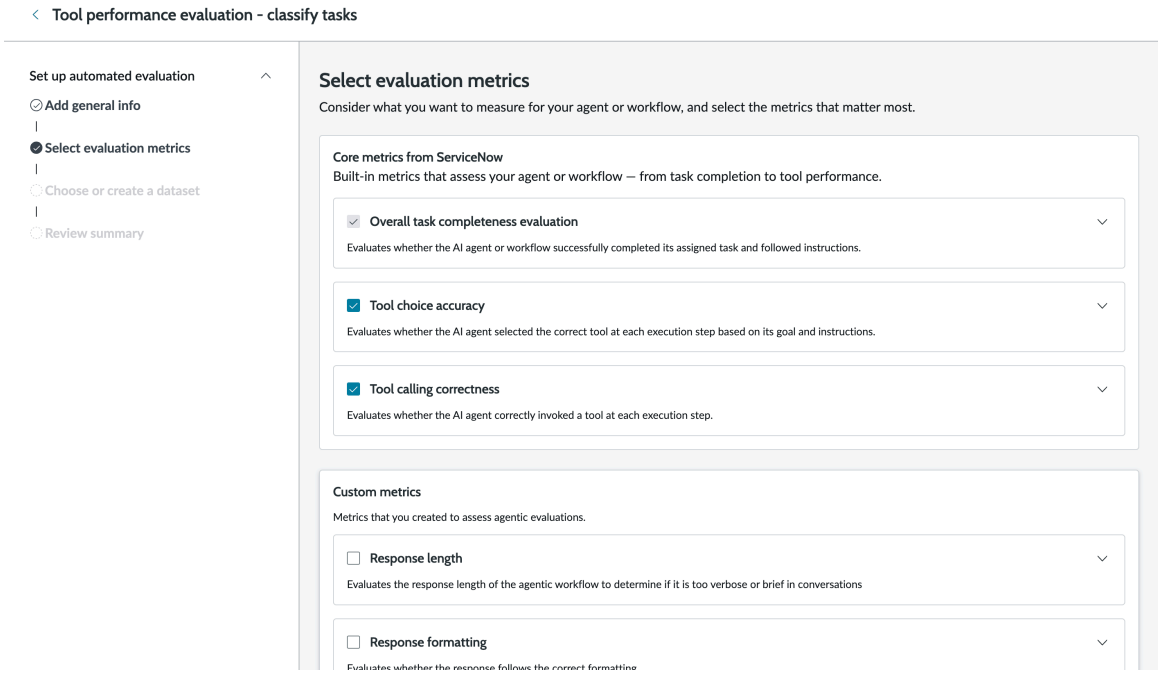
   Overall task completeness evaluation is selected by default. Running multiple evaluation metrics at a time can help provide a more comprehensive overview of the agentic AI asset's performance.

   To see more information about each plan, you can expand the card for each evaluation plan by selecting the chevron icon ⌄.

   Any custom metrics that you have published appear as options as well. If you don't see your custom metric, make sure that it's published. See Create a custom metric for more information.

   Note: The tool calling correctness metric is not available for AI voice agents.

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



6. Choose your dataset.

   a. Select an existing dataset or create your own.

   b. If you're creating a dataset, give it a name and description.

   c. To create a dataset, choose whether you want to run the agentic AI asset to generate new execution logs or use existing ones.

      Note: If you are evaluating AI voice agents, you must use existing execution logs.

Configure dataset form for new execution logs

Field name	Description
Table	The source table for records that the agentic AI asset uses to perform tasks and create executions.
No. of records	The maximum number of records within the dataset you want to run the evaluation on. If there are more records in the dataset than the maximum number of records, any records after the maximum number of records are ignored.
Filters	Conditions for narrowing down the list of records for the agentic AI asset to use to generate execution log data.
Starting phrase	Utterance given to the agentic AI asset to execute. Use the pill picker to select the specific dynamic inputs necessary to perform its task. For example, your starting instruction for evaluating an agentic workflow that generates resolution plans, you can set the starting instruction to `Help me resolve`

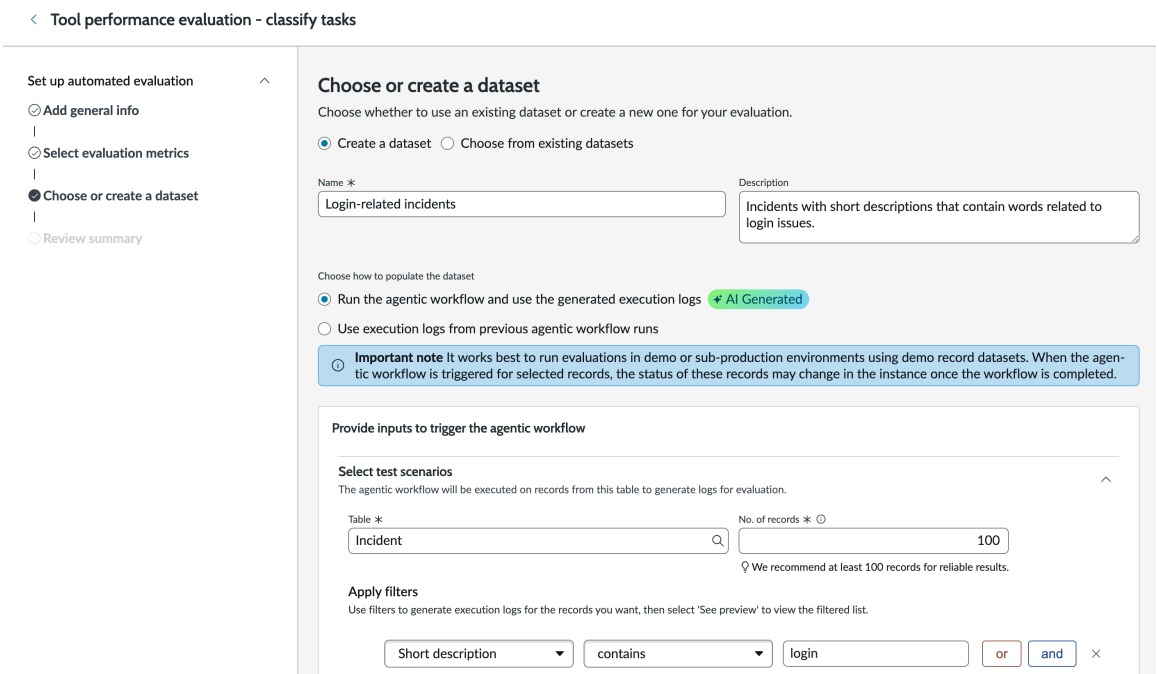
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

Field name	Description
	`{{incident.number}}`. Inputs from the record must be written between double curly braces.
Additional business context	Information given to the Large Language Model (LLM) acting as the invoking user that supplements information found within the table records. For example, a tuition reimbursement agentic workflow must know the normal reimbursement allowance, so the additional business context could be a knowledge article with that information.

Note: If you are creating new execution logs, the user that submits the evaluation must successfully pass the ACLs of the agentic AI asset and all of its components. If the correct role requirements aren't met, the execution logs will only report that the user does not have access, and the evaluation will fail. See Security for agentic AI for more information.

Configure dataset form for existing execution logs

Field name	Description
No. of records	The maximum number of records within the dataset you want to run the evaluation on. If there are more records in the dataset than the maximum number of records, any records after the maximum number of records are ignored.
Filters	Conditions for narrowing down the AI execution log records you want to include in the dataset. **Note:** Filter conditions are not supported for creating datasets of AI voice agent execution logs.

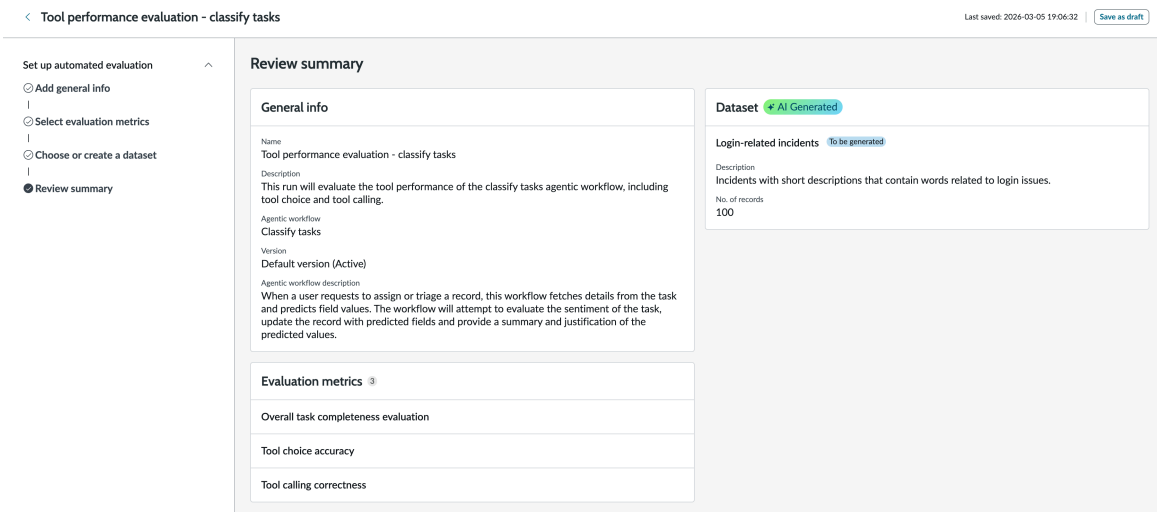


© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

d. Select **See preview** to see a list of records based on the conditions you specified. You can narrow down the records further by only selecting some of the records in the preview list. Unselected records won't be included in the dataset.

7. Review the agentic evaluation details in the last step of the guided setup.

   If you want to make changes, you can select **Back** to go to a previous step, or you can select the step in the sidebar.



8. Select **Start evaluation**.

Result

Your evaluation run executes. The time it takes for it to complete varies, but once completed you can select the evaluation from the Agentic Evaluations page to view the results.

For more information on the metrics on the results page, see Agentic evaluation run results.

Track and monitor agentic evaluation progress

Monitor the status of an active evaluation run to catch errors early and confirm when results are ready to review.

Before you begin

You must have an active evaluation run to monitor. For information about creating evaluation runs, see Execute an agentic evaluation run.

Role required: admin

About this task

Agentic evaluations can take time to complete, especially for large datasets. Monitoring progress helps you identify issues early and determine when results are ready for review.

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Agentic Evaluations**.
2. Select an evaluation with a trackable status.

   You can find evaluations to track in two locations:

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

- **Quick overview** section: Recent in-progress evaluations appear in the In-progress evaluations card
- **Automated evaluations** section: All evaluations, including older ones

Evaluations you can track have a *Run status* of **In progress** or **Action needed**.

3. Select the evaluation you want to monitor.
   The evaluation monitoring details page opens, showing the current status and progress information.
4. **Optional:** If the status is **Action needed**, review the generated execution logs.

   (Optional) The most common reason for the **Action needed** status is when execution logs have been generated but require approval before the evaluation phase can begin.

   a. Examine the dataset artifacts to understand how the agentic AI performed on specific records.
      You can open individual incidents or other records to see how the agentic AI asset interacted with them during the test.

   b. Select execution records to view detailed performance information.
      This opens the execution details in AI Agent Studio, where you can review the complete conversation between the simulated user and the agentic AI, including reasoning and processing messages from agents and tools.

   c. Review conversation records and timestamps to understand the interaction flow.
      The start phrase and conversation records provide detailed information about how the AI agent interacted with the simulated user, including timestamps for each message.

5. **Optional:** If you reviewed the execution logs and they meet your expectations, start the evaluation phase by selecting **Start evaluation**.

   (Optional) After you approve the logs, the LLM judging and scoring phase begins. This phase analyzes the execution logs and provides quantitative scores for the AI agent's performance.

   The evaluation status changes to **In progress** and the LLM evaluation begins.
6. **Optional:** Monitor the progress of the LLM evaluation phase.

   (Optional) During this phase, you can track:

   - Number of records evaluated
   - Estimated time remaining
   - Any errors or warnings that occur during evaluation
7. **Optional:** Check for completion notifications or status updates.
   When the evaluation completes, the status changes to **Completed** and results become available for review.

Result

You can monitor the evaluation progress and take action when required. When the evaluation completes successfully, you can review the detailed results to understand your agentic AI's performance.

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

What to do next

After the evaluation completes, review the results to identify areas for improvement in your agentic AI configuration. For information about analyzing evaluation results, see Review the results of an agentic evaluation.

Review agentic evaluation outputs

Assess your agent's overall performance after a run completes, including per-metric scores and issue counts. Use the results as your starting point for diagnosing quality issues and opportunities for improvement before deployment.

Before you begin

You must have a completed agentic evaluation.

Role required: sn_aia.admin or admin

About this task

Automated evaluations include scores and recommendations across the different metrics you chose. Each output provides information you can use to make decisions about development and deployment of the agentic AI asset. The evaluation results help you identify performance patterns, quality issues, and optimization opportunities before deploying your agent to production.

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Agentic Evaluations**.
2. Select the automated evaluation you want to review the results of.
   The evaluation details page opens, displaying the overall results and performance metrics.
3. Review the evaluation summary section to understand the overall performance.

   The summary provides a high-level overview of your agent's performance across all evaluated metrics. Key information includes:

   - Agentic AI asset information such as name and version
   - Total number of test cases evaluated
   - Average scores across all metrics
   - Number of issues identified by severity level
4. Review the overall LLM-judged scores for each metric.

   General LLM-judged scores for each metric demonstrate overall patterns and trends across the metrics you have evaluated against. These scores provide general recommendations for deployment based on the current version of the agentic AI asset. Detailed results include:

   - Numerical score
   - Performance rating (Excellent, Good, Moderate, or Poor)
   - Individual record evaluations
5. Investigate any issues and their associated traces.

   If problems with the agentic AI asset's performance are found, they are categorized by severity level, metric, and use case. Issues can be tracked down to their sources in specific interactions, called "traces." Review issues and their traces to diagnose underlying issues. Issues are classified by severity level:

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

- Critical: Issues that can prevent the agent from functioning correctly, resulting in a poor user experience
- High: Significant problems that impact user experience or accuracy
- Medium: Moderate issues that may affect performance in some scenarios
- Low: Minor issues that have minimal impact on overall functionality

6. Apply optimizations based on the findings.

   The automated evaluation can include recommended optimizations to address issues found in the evaluation. After you have applied the optimization, you can rerun the evaluation to see the changes in behavior and performance. Track improvements by comparing results across evaluation runs.

Result

You have a comprehensive review of your agent's performance across all evaluated metrics. Use these insights to make informed decisions about deployment readiness or identify areas requiring additional development work.

Review issues discovered in agentic evaluations

Identify and prioritize specific quality failures detected during an evaluation run, organized by severity. Use issue severity and metric relevance to decide which failures to address before deployment.

Before you begin

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Agentic Evaluations**.
2. From the list, select the automated evaluation you want to review.
   The evaluation details page opens, displaying the overall results and performance metrics.
3. If any issues are identified, select **Review and optimize** to view details.

   The optimization guided setup opens for that agentic AI asset.

   Each issue is associated with a specific tool or agentic AI asset.

4. From the issues list, select an issue to review.

   Each issue has a description of what the issue is, a root cause analysis, traces associated with it, and a space for adding notes about the analysis.

Section	Description
Issue score and short description	Brief overview of the identified issue and its impact on the agentic AI asset
Description	More detailed information about the identified issue than the short description

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

Section	Description
Root cause analysis	Insights into the cause of the symptoms seen in the agentic AI asset. Identifies specific components of the issue and which metric discovered it.
Traces with this issue	Specific execution logs where the issue is seen. You can analyze traces to understand how the issue manifests.
Analysis notes	Space for you to add details about the analysis and what you want to communicate to others about the issue. These notes aren't considered when applying optimizations. See Apply optimizations for more information.

5. **Optional:** Analyze the traces of the issue.
6. After reviewing the identified issues, select **Start optimization** to see the LLM-recommended fixes.

Result
Issues associated with the metrics for different components of the agentic AI asset are reviewed, analyzed, and documented.

What to do next
Proceed to the next step to apply optimizations to fix identified issues.

Analyze traces found in agentic evaluation run results

Investigate the full record of an agentic interaction to diagnose the root cause of a quality failure. Trace each step the agent took, including tool calls and outputs, to pinpoint where things went wrong.

Before you begin
You must have a completed agentic evaluation with issues identified.

Role required: sn_aia.admin

About this task
If you are already in the flow for identifying issues, skip to step 5.

Procedure
1. Navigate to **All** > **Now Assist Skill Kit** > **Agentic Evaluations**.
2. Select the automated evaluation you want to review the results of.
   The evaluation details page opens, displaying the overall results and performance metrics.
3. If any issues are identified, select **Review and optimize** to view details.

   The optimization guided setup opens for that agentic AI asset.

   Each issue is associated with a specific tool or agentic AI asset.
4. Select an issue to review.
5. Select **Show traces**.

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

Each execution plan where the issue occurred can be reviewed node-by-node with the inputs and outputs of each node displayed. You can navigate between nodes by selecting them in the center section.

6. Close the traces modal.

What to do next

Select **Apply optimization** to look at LLM-generated suggestions for optimizations to fix the issues.

Apply optimizations to agentic AI assets and reevaluate

Review and accept system-generated recommendations to improve agent quality based on detected issues. Apply optimizations before triggering a re-evaluation to confirm that changes resolved the failures.

Before you begin

You must have a completed agentic evaluation with issues identified.

Role required: admin

About this task

If issues are identified by the LLM judges, they can generate suggestions for possible fixes to address the issues. After reviewing the issue, you can select **Start optimization** to check the recommendations before applying them and starting a reevaluation.

If you are already in the flow for identifying issues, skip to step 6.

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Agentic Evaluations**.
2. Select the automated evaluation you want to review the results of.
   The evaluation details page opens, displaying the overall results and performance metrics.
3. If any issues are identified, select **Review and optimize** to view details.

   The optimization guided setup opens for that agentic AI asset.

   Each issue is associated with a specific tool or agentic AI asset.
4. Select an issue to review.
5. Go over the issue, including analyzing issue traces.
6. Select issues to optimize, then select **Start optimization**.

   You can choose to optimize for as many issues as you want. It is possible to generate specific changes to an agentic AI asset that address multiple issues at once.
7. Select **Yes** to generate the optimization suggestions.
8. Review the suggested optimizations.

   Optimizations occur at the prompt level for tool or agentic AI asset descriptions and lists of steps.

   If an optimization is to a description or list of steps, you can toggle to see the current instructions by itself next to the suggestion for the optimized version. You can also toggle to show the differences between the current and suggested version or to see just the optimized version by itself.

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

If you want to make edits to the suggested texts, you can toggle **Edit mode** to enter your own text or change the contents of the generated text.

To open the tool or agentic AI asset in AI Agent Studio, select the open link icon. The associated agentic AI asset opens in a new browser tab.

You can see the issues that the optimization is trying to address by selecting **Contributing issues**. You can select an issue to review all of the details, including the root cause analysis.

9. Select **Continue to re-evaluate** if you want to re-evaluate without applying the optimization or **Review and apply** if you want to make the suggested changes.
10. Review the details for reevaluation.

    The optimization summary includes the list of issues that are being trying to be addressed by the optimization.

    You can name your evaluation and select the version of the agentic AI asset. Because this is a reevaluation, the evaluation metrics and dataset are the same as the original run and can't be changed.

11. Select **Start re-evaluation** to begin the new agentic evaluation run.

    You can go back to a previous step in the issue-optimization flow by selecting **Back to evaluation setup**.

Result

A new evaluation run begins with any changes you applied to the agentic AI asset.

What to do next

You can track and monitor your agentic evaluation while the reevaluation is in-progress. After the evaluation is complete, you can review and compare the results against the original evaluation and make further optimizations if necessary.

Troubleshoot agentic evaluations

Find solutions to common evaluation errors, including run failures, data ingestion issues, and unexpected results.

When using agentic evaluations, you may see unexpected execution results or errors. The following discusses situations you could encounter and some of the reasons why those situations occur.

Evaluation run failed

There are a few different reasons why an evaluation run may fail to execute properly.

Agent version unavailable

Verify that the selected agent version still exists in AI Agent Studio. The version does not have to be the one currently active, but deleted or archived versions can't be evaluated.

User permissions

Confirm that your User record has the permissions required to execute evaluation runs in general and to use the specific AI asset. To check whether a certain user has access, you can perform an access test. See Test user access to an AI agent and Test user access to an agentic workflow.

Data format errors

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

Verify that the dataset conforms to the required format. Malformed records can cause the evaluation to fail. See Data requirements for agentic evaluations for the supported data types.

Metric and data mismatch

Confirm that all selected metrics have the required data inputs. Metrics that require ground truth will fail if the ground truth field is missing from the dataset.

Agentic AI asset underperformance despite no issues found

If the evaluation found no issues, but the specific agentic AI asset is still not performing to acceptable standards, consider the following:

Dataset coverage

The evaluation dataset may not include the types of inputs or scenarios that expose the agent's weaknesses. Review the dataset for any coverage gaps and add representative edge cases to more closely align what is being evaluated with real-world scenarios.

Metric selection

The selected metrics may not be measuring where the agentic AI asset is failing. Review whether additional or different metrics would better capture the performance gap. You can create custom metrics to evaluate other dimensions of the agentic AI asset responses or actions, such as length of response or whether a response meets certain formatting requirements.

Scoring thresholds

The pass threshold for a metric may be set at a level that does not reflect your requirements. Review threshold settings in the metrics configuration to redefine success and failure.

Optimization applied, but re-evaluation didn't improve

If the re-evaluation scores did not improve after applying optimizations, try the following:

- Review trace details for the issues that were targeted. The optimization may have only alleviated surface-level symptoms without resolving the underlying root cause.
- Check whether the optimization introduced a regression in a different metric. Score improvements in one area can sometimes degrade another, lowering the final scores.
- If the optimization was applied to the list of steps of an agentic AI asset, verify that the updated list of steps was applied to the version you are evaluating.

Data processing errors

If the data can't be processed because it doesn't meet data requirements, the evaluation can't execute properly. The following describes common causes of data processing errors:

Incorrect file format

The accepted file formats are CSV and structured JSON. Other file formats can't be processed.

Missing required fields

Datasets must include the fields required by the selected metrics. Check for missing or misnamed columns. If you're using a ground truth, you must include it in the dataset.

Encoding issues

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

Files must be UTF-8 encoded. Files with non-standard encoding may fail to be processed.

File size

Very large files or datasets may time out during processing. If this occurs, reduce the dataset size or contact your Platform Administrator.

Create a custom metric for evaluating agentic workflows

Create a custom metric for evaluating AI agents and agentic workflows to test the outputs against expected responses.

Before you begin

Role required: sn_aia.admin

About this task

Custom metrics enable you to choose you measure your AI agents and agentic workflow for effectiveness. You can create custom metrics that apply to one agentic AI asset or multiple ones. Once created, you can select custom metrics in the guided setup for agentic evaluations. You can select both or either custom metrics and metrics installed with Now Assist applications. Examples of custom metrics include:

- Testing that a certain phrase is in the response
- Measuring the length of a response to test verbosity

When writing the script, you have access to the output of a parser tool. See agentic evaluation parser tool for more details on how it works.

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Agentic Evaluations**.
2. Go to the Evaluation metrics tab and select **Create metric**.
3. In the **General information** step, enter a name and short description for your custom metric.

   You can navigate between steps in the guided setup with the **Continue** and **Back** buttons. After entering a name for your custom metric, you can skip directly to writing your script by selecting **Skip to Script Editor**. You can fill out the details later from the Metric settings tab in the Script Editor.

4. **Optional:** In the **Metric details** step, add information about the evaluation metric, how it works, and what the output format of the metric is.

   (Optional) If you select **Use this content**, then the details you provide are displayed along with the option in the guided setup for an agentic evaluation run.

   You do not have to fill in details, but you have the option to provide useful information and context to aid the custom metric development process. You can also add metric details from the **Metric settings** tab of the Script Editor.

5. After reviewing your metric information in the **Summary** step of the guided setup, select **Finish setup** to be redirected to the Script Editor.
6. **Optional:** Add metric inputs.

   (Optional) By default, the execution plan record sys_id is included as a metric input.

   The parser tool lets you access content from the execution plan log without having to parse through the XML or JSON content yourself. You can access the parser tool's outputs with `tool.output`.

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

a. In the Metric inputs section, select the plus icon to open the Add metric input modal.

b. Complete the form.

Field	Description
Input datatype	Drop-down selector to choose what data type and format the input is
Name	Short, user-friendly name for the input
Description	Description of the input. You can view the description of any input by selecting it in the Metric inputs section of the Script Editor.
Test values	Default input values used for testing your script in the Script Editor
Mandatory	Determines whether providing the input is necessary for the evaluation to run

c. Select **Add**.

7. Write your script-based metric.

   The default code block in the Script Editor contains commented code that can guide you through the script-writing process.

   At any time, you can select **Metric output template** to see the code necessary for the output to display correctly when looking at the evaluation run results in Skill Kit. This content is also included in the default code block when you first open the Script Editor. If necessary, you can copy the template and paste it into the script. Be sure to change the values to the results of your custom metric.

   You can review the details of the custom metric from the guided setup from the Metric settings tab.

   When you have made changes, you can select the **Save** button to save your changes.

   > **Important:** Even if you save, the metric won't be available to use in agentic evaluation runs until you publish it.

8. Select **Run test** to open the Run test modal, enter your metric test inputs, and run the test.

   After you have run at least one test, you can review the response and request of both the overall execution and the parser tool.

   After you have run at least one test, you can select the **Run test** drop-down button and select **Run test history** to open a new sidebar with past test results. You can filter test runs by the user who executed the test.

9. Select **Publish metric** to save your script and activate the metric.

   Before publishing, ensure that your output matches the metric output template. If not, your metric may not work as expected or display correctly on the evaluation run results page.

Result

You have a new custom metric that can be used for evaluating agentic AI assets.

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

What to do next

Execute an evaluation run and select your custom metric in the guided setup. See Evaluate agentic AI for more information on executing an evaluation run.

If you do not see your custom metric in the guided setup for an evaluation run, make sure that you have published it.

Reference for agentic evaluations

Find technical reference material for roles, metrics, and output formats of agentic evaluations.

Available metrics

Standard metrics available

Metric	What it measures	Ground truth required
Task completeness	Whether the agentic AI asset fully addresses the user need.	Optional
Response accuracy	Whether the agentic AI asset's response is factually accurate	Recommended
Groundedness	Whether the agentic AI asset's response is grounded in the specific context of the task	No
Coherence	Whether the agentic AI asset's response is logically structured and clear	No
Tool use accuracy	Whether the agentic AI asset selected and used the correct tool to execute its tasks	Optional
Goal adherence	Whether the agentic AI asset stayed within its defined scope and instructions	No

Issue types

Issues are broken down by behavior. Each metric has its own issues identified separately.

Issue categories

Category	Agentic AI asset behavior
Incomplete response	Response failed to address the user's full request

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

Issue categories (continued)

Category	Agentic AI asset behavior
Factual error	Response contained content that isn't factually correct
Hallucination	Response contained content not grounded in the specific context of the request
Incoherent output	Response was disorganized or difficult to understand
Incorrect tool use	Selected the wrong tool or passed incorrect parameters to a tool
Scope violation	Responded to a request outside its defined operating scope

Data requirements

Data requirements for datasets in agentic evaluations

Requirement	Description
Minimum test cases	A minimum number of test cases is required per run. The specific metrics you are using for the run may have their own minimum test cases. Ensure that your dataset meets the requirements for all metrics.
Supported formats	CSV and structured JSON are supported.
Ground truth field	If you're using a ground truth, it must be provided as a separate field in the dataset. The ground truth field must be aligned to each test case individually.
Data representativeness	Datasets should reflect all of the tasks that the AI agent or agentic workflow will handle. Include edge cases and failure-prone scenarios to help ensure that you're testing against common real-world scenarios.

Agentic evaluation parser tool

Use the outputs of the agentic evaluation parser tool in your scripts for custom metrics to customize the criteria for effective AI agents and agentic workflows.

Parser tool overview

The agentic evaluation parser tool extracts structured execution data from the execution logs of an agentic workflow or AI agent. You can use the information gathered by the tool to create custom metrics that use scripts to evaluate agentic workflows.

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

The parser tool returns structured AI agent or agentic workflow execution data within the *output.payload* object, which contains the following:

- *executionInputs*: A JSON object that contains initial workflow setup information, such as the names and instructions of agents and tools and the initial user utterance
- *executionOutputs*: A JSON object with AI agent actions and tool execution results
- *executionMessages*: A JSON object array of user-facing conversation flow and system responses
- *executionPlanDetails*: A JSON object of execution metadata, such as status, sys_ids, and configuration values

```
{
     "output": {
 "payload": {
  "executionInputs": { ... },
  "executionOutputs": { ... },
  "executionMessages": [ ... ],
  "executionPlanDetails": { ... }
 }
}
```

Accessing the parser tool's output

To view the complete parser tool output for testing and development, perform the following steps.

- Navigate to the Script Editor view in your custom metric guided setup.
- Select **Run Test**.
- Wait for the test to complete.
- View the complete JSON output in the Tools section of the test results.

Reviewing the parser tool's output before designing your custom metric enables you to inspect the data structure before implementing any specific logic.

executionInputs data structure

The executionInputs attribute contains a JSON object with the following structure:

```
"executionInputs": {
 "agenticWorkflow": "(name of agentic workflow)",
 "description": "(descriptions for agentic workflow)",
 "instructions": "(list of steps for agentic workflow)",
 "utterance": "(initial user utterance)",
 "agents": [
  {
   "name": "(AI agent name)",
   "instructions": "(list of steps for AI agent)",
   "tools": [
    {
     "name": "(tool name)",
     "description": "(tool description)",
     "executionMode": "(execution mode, either Autonomous or
 Supervised)",
     "inputs": { ... }
```

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

```
    },
    { ... }, ...
   ]
  },
  { ... }, ...
 ]
}
```

executionOutputs data structure

The executionOutputs attribute contains a JSON object with the following structure:

```
"executionOutputs": {
 "agents": [
  {
   "name": "(AI agent name)",
   "subTask": { ... },
   "tools": [
    {
     "name": "(tool name)",
     "inputs": { ... },
     "output": { ... }
    }
   ]
  },
  { ... }, ...
 ]
}
```

executionMessages data structure

The executionOutputs attribute contains an array of JSON objects with the following structure:

```
"executionMessages": [
 {
  "role": "(Message sender, either 'agent' or 'user')",
   "message": "(Content of message)",
   "order": "(Sequence number indicating order of message in the
 conversation)"
  },
  { ... }, ...
]
```

executionPlanDetails data structure

The executionPlanDetails attribute contains a JSON object with the following structure:

```
"executionPlanDetails": {
    "state": "(Current execution status)",
    "runType": "(Type of execution)",
    "conversationId": "(sys_id of conversation)",
    "relatedTask": "(sys_id of the associated task or record)",
    "relatedTaskTable": "(Table name where the related task is
 stored)",
    "context": { ... } (May be null)
    "builtInTools": [ { ... } ]
```

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

```
    }
```

This section provides execution metadata for tracking workflow performance, debugging issues, and correlating executions with specific tasks or conversations.

The possibilities for runType include the following:

- API
- Chat
- Evaluation
- Testing
- Trigger

Using parser tool output in metric scripts

The parser tool data is available in your metric script through the *context* parameter. Access the structured data using the following code:

```
// Access the parser tool output from context
var parserToolOutput =
 context['AgenticExecutionParserTool.output'];
if (typeof parserToolOutput == "string") {
 parserToolOutput = JSON.parse(parserToolOutput);
}
var parserToolPayload = parserToolOutput.payload;
var parserToolStatus = parserToolOutput.status;

// Extract individual sections from payload
var inputs = parserToolPayload.executionInputs;
var outputs = parserToolPayload.executionOutputs;
var messages = parserToolPayload.executionMessages;
var planDetails = parserToolPayload.executionPlanDetails;
```

Agentic evaluation run results

Learn about agentic evaluation runs and the meaning behind different evaluation scores from the agentic evaluation results page.

Agentic evaluations overview

Agentic evaluations measure how well AI agents and agentic workflows are accomplishing their objectives. A Now LLM Service model judges the AI agent or agentic workflow based on the execution logs. The results page of an evaluation run shows multiple metrics and scores measuring task completeness and tool use.

If you run an overall task completion evaluation, the results page shows recommended actions for the AI agent or agentic workflow. Recommended actions give you suggestions for deployment or improvement to help verify that the agentic workflows that you deploy are performing up to your standards.

After you've reviewed your evaluation results, you can archive your evaluation or copy it to run another evaluation with the same parameters and dataset.

You can export the evaluation results as a report. The report is formatted as a .csv file that includes the individual sys_ids of the execution records and the metric scores for each.

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

For more information on AI agent usage and other analytics, you can review the AI Agent Analytics dashboard in the AI Agent Studio.

Evaluation results overview

For each evaluation method that you execute, the results page displays an overall score for the agentic workflow with a percentage of successful record evaluations and a label of Excellent, Good, Moderate, or Poor. You can change the metric thresholds for each label by selecting **Customize metric thresholds**.

In addition to the overall task completeness results, you can review a summary of the results of the other metrics.

Overall task completeness evaluation run results

Label	Description	Recommended action	Default threshold
Excellent	Tasks were consistently performed at a high standard. The agentic workflow or AI agent is working well.	Proceed with confidence	90%–100%
Good	Most tasks were performed successfully, but some performance inconsistencies suggest areas for improvement.	Deploy with caution	70%–89%
Moderate	A significant number of tasks weren't fully completed. Performance is below the desired level.	Investigate the root causes of poor task completion	50%–69%
Poor	The agentic workflow is consistently failing to complete tasks adequately. Major issues are present.	Do not deploy	0%–49%

Individual record metric scores

Evaluations are run against the log tables of agentic workflow executions. Each record is individually scored for each evaluation plan that you run. Individual record evaluations are scored according to the following metrics.

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

Overall task completeness record metric scores

The overall task completeness metric assesses whether an AI agent successfully completes its assigned task. It evaluates the execution logs of the agent, ensuring that all required steps were taken and the task was logically and effectively completed.

Number	Score	Description
3	Successful	The main task was fully completed. All subtasks were resolved, and the steps followed a logical sequence with no critical errors.
2	Partially successful	The task was partially completed. Some subtasks remain unresolved or inefficiencies affected the process.
1	Unsuccessful	The task wasn't completed. Critical subtasks were abandoned or unresolved or the execution failed entirely.

Tool performance record metric scores

The tool performance evaluation metric assesses an AI agent's ability to select the most appropriate tool for each step while completing a task.

Number	Score	Description
1	True	The right tool was chosen for the action in the plan.
0	False	The right tool wasn't chosen.

Tool calling records metric scores

The tool calling evaluation metric assesses whether an AI agent correctly constructs tool calls by validating the accuracy, completeness, and formatting of the inputs it provides.

Number	Score	Description
1	True	Input key completeness, input value correctness, and input format correctness are all successful.

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

Tool calling records metric scores

The tool calling evaluation metric assesses whether an AI agent correctly constructs tool calls by validating the accuracy, completeness, and formatting of the inputs it provides.

(continued)

Number	Score	Description
		• *`Input key completeness`*: 1 - True – All required parameters are present with exact name matches, and no unexpected parameters are included. • *`Input value correctness`*: 1 - True – Tool input values are correctly mapped. • *`Input format correctness`*: 1 - True – Tool inputs are in the correct format.
0	False	One or more of input key completeness, input value completeness, or input format completeness wasn't successful. • *`Input key completeness`*: 0 - False – A mandatory parameter is either missing, its name doesn't match exactly, or an unexpected parameter was found. • *`Input value correctness`*: 0 - False – Tool input values are not correctly mapped. • *`Input format correctness`*: 0 - False – Tool inputs are not in the correct format.

Note: The values of the sub-metrics are aggregated using an AND operator. If any one value is 0, then the entire tool calling records metric score will be 0.

Examples of using AI agents

Review different ways that you can leverage the Now Assist AI agents application in agentic workflows across the platform.

AI agents Overview

AI agents in the Now Assist application provide automated solutions for a wide range of business processes and operational challenges. These agents can work independently or as part of hierarchical structures to handle complex workflows, reduce manual effort, and improve service delivery efficiency. You can leverage AI agents across multiple departments including IT Service Management, Customer Service Management, and Human Resources.

Incident Management and Resolution

AI agents can automate the entire incident lifecycle from initial triage to resolution. When an incident is created and assigned to a designated service desk user, the agent automatically retrieves incident details, searches for similar historical incidents, and analyzes relevant knowledge base articles. Based on this research, the agent determines an appropriate resolution path.

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

If the agent has high confidence in the solution, it posts the resolution in the work notes and closes the incident automatically. When confidence is lower, the agent documents its findings and reassigns the incident to a technical support specialist for review. This approach reduces mean time to resolution while ensuring quality control for complex issues.

Case Management Automation

Similar to incident management, AI agents can handle case workflows in customer service scenarios. Agents analyze case details, match them with historical cases, search knowledge articles, and provide recommended solutions or next steps. This automation helps customer service teams handle higher volumes of cases while maintaining consistent quality in responses.

Cross-Functional Process Automation

AI agents can orchestrate workflows that span multiple departments and systems. For instance, an onboarding agent might coordinate tasks across HR, IT, facilities, and finance departments, ensuring new employees receive proper access, equipment, workspace setup, and orientation materials. The agent can track progress, send reminders, escalate delays, and provide status updates to stakeholders throughout the process.

Knowledge Management and Content Discovery

AI agents can assist users in finding relevant information by analyzing queries, searching across multiple knowledge sources, and presenting contextualized results. These agents understand the user's intent, consider their role and permissions, and deliver targeted information that helps them complete tasks or resolve issues independently.

Virtual Agent Enhancement

AI agents can work behind the scenes to enhance Virtual Agent conversations by providing intelligent responses, retrieving relevant information, and executing actions on behalf of users. This integration creates more capable and helpful conversational experiences that can handle complex queries and multi-step processes.

Implementation Considerations

When implementing AI agents, organizations should consider the following factors:

- Start with well-defined, repetitive processes that have clear success criteria
- Ensure adequate training data and knowledge base content for agents to reference
- Define appropriate confidence thresholds for automated actions versus human review
- Establish monitoring and quality assurance processes to track agent performance
- Plan for iterative improvements based on agent interaction data and outcomes
- Configure memory sharing appropriately for different hierarchy structures
- Consider scalability requirements when designing complex agent hierarchies

Resolve an incident

Help your live agents resolve an incident faster with Now Assist AI agents by using the Now Assist panel.

https://player.vimeo.com/video/1098087911?h=d89ede16f9&badge=0&autopause=0&player_id=0&app_id=58479

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

Before you begin
Role required: none

Procedure

1. Navigate to **All** > **Incidents** > **Assigned to you** and open an incident that you would like to resolve.
   When you open an incident, the Now Assist application checks if a plan is available by using AI agents and displays the `Now Assist has a plan for solving INCXXXXXXX. Open Now Assist Panel to view the plan.` message in a banner.

   Note: You can select the banner and directly go to the conversation on the Now Assist panel to complete the task.

2. Open the Now Assist panel by using the Now Assist icon.
   Now Assist provides the resolution steps for the incident.
3. On the Now Assist panel, enter `Yes`.
   When the live agent enters `Yes`, they can proceed with the execution plan that is proposed by the Now Assist AI agents. The live agent checks the conversation and adds further inputs, if needed, to resolve the incident.

Review and update tickets with the Ticket Status AI agent

Review the status of your tickets and take standard ticket actions, such as adding comments, using the Ticket Status AI agent in Now Assist in Virtual Agent.

Before you begin
Role required: By default, any authenticated user with access to the portal the Ticket Status AI agent is enabled for can access the AI agent. There may be additional roles required depending on the configuration by your administrator.

Procedure

1. Navigate to your Service Portal and open Virtual Agent to start a new conversation.

   Depending on the configuration, you may also be able to access this AI agent in the Now Assist panel or Microsoft Teams with myBot.

2. Enter one of the following utterances:
   - Can you list all open tickets I've created?
   - What's the current status of my incident INC001?
   - Add comment to INC001 "I need your help in fixing this request ASAP. Please prioritize this."
   - Can you check the latest progress on my most recent request?
3. Answer any follow-up questions.

Result
The Ticket Status AI agent used various tools to check and update your tickets.

Now Assist AI agents reference

Find more information about user roles, tables, and the different properties that are installed in Now Assist AI agents.

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

Now Assist AI agents roles

The following roles are installed with Now Assist AI agents with a compatible Now Assist application.

Roles needed to configure or monitor AI agents

Role	Description
AI Agent Admin [sn_aia.admin]	Administrator of the application. A user with the sn_aia_admin role can create, read, update, and delete records.
AI Agent Viewer [sn_aia.viewer]	Read-only access to the application. A user with the sn_aia_viewer role has read and report access on all tables.
agent_role_config_admin	With this role, user can access and modify Agent role configurations with AI Agent Admin [sn_aia_admin] being the parent role.
agent_role_config_viewer	Can view the Agent role configurations with AI Agent Viewer [sn_aia_viewer] being the parent role.

Now Assist AI agents system properties

The following are system properties that define default values and behavior.

System properties to use for configuring AI agents

Property	Description
sn_aia.agent_llm_provider	Defines the (large language model) LLM service provider for AI agents. Default value: **azure_openai**.
sn_aia.agent_tool_supported_data_types	Defines a comma-separated list of supported data types for tools that are used by agents for IntegrationHub spoke. Each value corresponds to the name field of records in the Field classes table [sys_glide_object].
sn_aia.analytics_dashboard_sysid	Provides the sys_id for the AI Agents Analytics dashboard. read_roles: sn_aia.admin and sn_aia.viewer
sn_aia.continuous_communicator_output_limit	Defines the maximum number of continuous outputs that the AI Agent Orchestrator can trigger to show to users. Default value: **3**
sn_aia.continuous_tool_execution_limit	Defines the maximum consecutive number of uses for the same tool. Default value: **7**

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

System properties to use for configuring AI agents (continued)

Property	Description
sn_aia.enable_usecase_tool_execution_mode_override	Enables running agentic workflows fully autonomously, overriding any non-automated tools in the agentic workflow. Default value: **false**
sn_aia.ltm.category.auto_create	Enables AI-generated categories if no matching categories exist. Default value: **false**
sn_aia.ltm.enable_long_term_memory	Enables long-term memory for AI agents. All previous user interactions are used as context for the LLM. Default value: **false**
sn_aia.maximum_agent_tools	Defines the maximum number of tools that can be assigned to an AI agent. Default value: **20**
sn_aia.max_scheduled_trigger_query	Defines how many records are processed when a scheduled trigger is detected. Default value: **10** Write_role: admin
sn_aia.mid_skill_switch_enabled	Enables mid-skill switching. Default value: **false**
sn_aia.react_failure_retry_max_limit	Defines the maximum number of retries in case of an execution failure. Default value: **3**
sn_nowassist_va.router_redirect_va_agentic	Determines AI agent discovery in Virtual Agent. If set to *NEVER*, Virtual Agent continues Q&A without any agentic AI. Default: *ROUTER_DECISION*
com.glide.cs.dynamic.capability.timeout	Defines the Timeout for AIA Proficiency Descriptor. Default value: **180**. Write_role: admin.
sn_aia.enable_follow_up	Enables users to continue the conversation with follow-ups after the agentic workflow execution is complete.

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

System properties to use for configuring AI agents (continued)

Property	Description
	Default value: **true**.
sn_aia.follow_up_message	Defines a follow-up message sent after execution is completed. Default Value: **How else can I help you?**
sn_aia.allow_context_sharing	Enables the sharing of short-term memory, allowing context to persist across execution within the same conversation. Default Value: **true**
sn_aia.agent_strategy_choice_enabled	Enables to show the LLM reasoning strategy in the agent setup screen. Default Value: **false**
sn_aia.context_sharing_strategy	This property defines the strategy to use for storing short-term memory for an execution. Default Value: **summarise**
sn_aia.enable_agent_tool_input_value_overrides	Enables you to override the agent tool input value. Default Value: **true**.
sn_aia.follow_up_qna_failure_limit	Defines the limit to exit execution if the number of consecutive questions and answers aren't available in the follow-up. Default value: **1**.
sn_aia.ltm.use_memory_for_ai_agent	Enables long-term memory for AI agent interactions. When enabled, stored user memories are utilized in AI agent interactions. Default value: **true.**
sn_aia.quick_mode_failure_retry_max_limit	Defines maximum limit for retries in case of a failure in Quick Mode execution. Default value: **3**.
sn_aia.user_context_data	Defines a comma-separated list of user context data to be used with AI Agents. The list is used to pick the data available from knowledge graph API: getUserContext. List of available user information:

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

System properties to use for configuring AI agents (continued)

Property	Description
	• profile • manager • reportees • assets The user information can also be customized by overriding the method `getUserContext` via `UserContextUtil`. If customized, the property must define the comma separated list of keys generated by the customized `getUserContext` method. Default value: **profile**.
sn_aia.external_agents.enabled	Set to `true` to enable external agent features
sn_aia.external_agents.parallel_conversations.enabled	Enables or disables multiple simultaneous conversations per user
sn_mcp_client.cursor.max_iterations	Specifies the maximum number of cursor-based pagination iterations to perform when fetching MCP tools. Helps prevent excessive looping or API calls during data retrieval. **Note:** Set this value to 0 to disable pagination iteration limits.
mcp_guardian_check	Enables guardian check for MCP Client when the value is set to **true**. The default value is **false**. **Note:** To enable guardian check for MCP Client, ensure that you enable Now Assist guardian on **AI Agent Studio** > **Settings** page.
com.glide.agentic_processes_view.enabled	Enables the in-product experience for agentic workflows in the AI Workflows panel. Ensure that this is set to `true` if you plan on using UI actions to run agentic workflows.

Global properties used for episodic memory or agent learning

System property	Description
sn_aia.enable_episodic_memory	Enables episodic memory. When enabled, episodic memories from past interactions are collected and stored in sn_aia_memory and

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

Global properties used for episodic memory or agent learning (continued)

System property	Description
	utilized in future interactions to enhance the overall experience Default value: **true**. .
sn_aia.enable_memory_limit	Maximum number of episodic memories to inject into prompt when an agent is invoked. **Note:** The maximum number of values allowed is less than or equal to 5.
sn_aia.use_episodic_memory_for_ai_agent	Enables episodic memory injection for AI agent interactions. When enabled, stored user memories are utilized in AI agent interactions. Default value: **true**.

Properties on the Agent Properties table

The following properties are on the Agent Properties [sn_aia_property] table. These properties affect different AI agent behaviors.

General properties on the Agent Properties table

Name	Description	Default value
alert.assist_spike_hours_to_check	Time interval, in hours, between running the scheduled job that checks for spikes in assist usage	3
alert.assist_spike_usage_percentage_threshold	Percentage increase required from previous job results to trigger the spike in assist usage notification	0.5 (50%)
alert.assist_spike_usage_threshold	Minimum number of assists to	5000

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

General properties on the Agent Properties table (continued)

Name	Description	Default value
	trigger the spike in assist usage notification	
alert.consecutive_error_count	Consecutive number of latency errors to trigger the latency notification	3
alert.llm_latency_threshold	Time taken, in seconds, for the LLM to respond before sending a latency error	10
alert.tool_latency_threshold	Time taken, in seconds, for a tool to finish executing before sending a latency error	300
enable_agent_tool_input_value_overrides	Determines whether tool input values can be overriden by the AI agent or other tools for use by another tool	false
execution_task.latency_thresholds	Determines thresholds on LLM and tool execution times and response tokens/ character counts for tool calling	{ "llm": { "timeBoundaries": [5000, 10000], "tokenLimit": 500 }, "tool": {

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

General properties on the Agent Properties table (continued)

Name	Description	Default value
		<pre>"timeBoundaries": [200000, 300000], "outputCharLimit": 10000 } }</pre>
follow_up_behaviour	Actions to take by the agentic workflow after it has finished executing. Each record applies to only one agentic workflow.	no_followup_close_conversation
mcp_guardian_check	Determines whether Now Assist Guardian runs on MCP tool executions	false
show_citations	Determines whether agentic AI-generated responses in Now Assist panel or Now Assist in Virtual Agent add citations for their output	false

The following properties are used to detect and avoid infinite recursion or execution loops. If you reach the max executions that match the query within the time window, any new executions matching the query will abort. You can adjust these values if you want to lower the threshold for detecting recursion.

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

Recursive check properties on the Agent Properties table

recursive_check.create_max_executions	Maximum number of matching executions creating records	50
recursive_check.create_time_window	Time window, in minutes, for checking for matching executions creating records	15
recursive_check.query_for_create_record	Query for the fields of execution plans to create records to check if executions are repeating	objective={objective}^agent={agent}^ORusecase={usec
recursive_check.query_for_update_record	Query for the fields of execution plans to update records to check if executions are repeating	related_task_record={related_task_record}^objective={c
recursive_check.update_max_executions	Maximum number of matching executions updating a record	5
recursive_check.update_time_window	Time window, in minutes, for checking	15

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

Recursive check properties on the Agent Properties table

(continued)

	for matching execution updates

Now Assist AI agents tables installed

The following tables are installed so Now Assist AI agents works as expected:

Tables used for configuring AI agents

Table	Description
Agentic workflows [sn_aia_usecase]	List of configured agentic workflows.
AI Agents [sn_aia_agent]	List of configured AI agents.
Tools [sn_aia_tool]	List of tools used by an AI agent.
Strategies [sn_aia_strategy]	List of strategies used by an AI agent.
Teams [sn_aia_team]	Team that is a group of listed agents.
Team members [sn_aia_team_member]	List of teams mapped to an agent.
Agent Tools [sn_aia_agent_tool_m2m]	List of tools mapped to an AI agent.
AIA Trigger Configurations [sn_aia_trigger_configuration]	List of triggers created for an agentic workflow.
Execution Tasks [sn_aia_execution_task]	List of tasks by execution plan ID.
Messages [/sn_aia_message]	List of messages recorded in AI agent conversations to and from the human users.
Tools Executions [sn_aia_tools_execution]	List of tools executed by the plan ID. **Note:** The records in the Tools Executions table expire and become unavailable after a period of 13 months.
Execution Plans [sn_aia_execution_plan]	List of plan executions by conversation ID.
Agent Tools [sn_aia_agent_tool_m2m]	List of tools and maximum automatic executions.
AI Agent configs [sn_aia_agent_config]	List of active AI agents configured for the proficiency that they'll be used in.
Agent properties [sn_aia_property]	Various properties that affect AI agent behavior.
Gen AI Metadata M2M [sn_aia_gen_ai_m2m]	List of Gen AI metadata and the maximum automatic executions.Maintains the mapping between sn_aia_execution_task and Gen AI log metadata.

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

Tables used for configuring AI agents (continued)

Table	Description
	If two large language model (LLM) calls are made to the sn_aia_execution_task, then the sn_aia_gen_ai_m2m table has two records.
Report metrics [sn_aia_report_metric]	List of the report metrics.
Agent Access Role Configurations [sys_agent_access_role_configuration]	List of agent access roles. You can also create new agent access roles from this table.
Sys_generative_ai_configuration	

Tables used for Group Action Framework

See **Group Action Framework** *for more information.*

Table	Description
GAF record group [sn_gaf_record_group]	Stores the output of the grouping skill. Each record represents a cluster of related records. You can open each record group record to discover which records were included within the cluster.
GAF record group detail [sn_gaf_record_group_detail]	Contains the individual records that belong to each group GAF record group. You can also find if a record is marked to act as a representative of the cluster on these records.
GAF action strategy result [sn_gaf_action_strategy_result]	Holds the results of the Action Strategy skill, which selects representative records from each group for downstream processing.
GAF action mapper results [sn_gaf_action_mapper_result]	Stores the output of the Mapper skill, which maps new records to existing clusters.
GAF action reducer results [sn_gaf_action_reducer_result]	Stores the result of the Action Reducer skill. The results include insights for entire clusters. For example, how to resolve incidents similar to those gathered in a cluster.

Domain separation and Now Assist AI Agent Studio

Domain separation is supported for Now AssistAI Agent Studio. 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.

Domain Separation Overview

Now Assist AI agents use basic domain separation capabilities to help protect your users' data. Domain separation support for AI agents is applied at design time and run time.

Design-time support

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

Refers to creating or updating agentic workflows, agents, tools, trigger configurations, and so on. AI agent configurations can be made domain-specific for individual agents and the actual agentic workflows. Administrators can apply specific domains to those records. Similar to other basic domain separations, records in the AI agents tables are accessible if the user belongs to the same or a higher domain than those records.

Run-time support

Refers to the agentic conversation on the Now Assist panel, web client, or any conversational channel. In the agentic conversations, the user that the agent impersonates functions as an agent with any AI agents who initiate the conversation on demand. For example, if the conversation is happening via a trigger mentioned on the *Run as* field on the Trigger form of an agentic workflow. If the user that the agent impersonates belongs to the same or a higher domain, that agent can access and use configurations that are associated with that domain.

The domain visibility for an agentic workflow is resolved during run time based on the *Run as* attribute in the agentic workflow trigger condition. For more information, see defining a trigger for an agentic workflow.

When an agentic conversation is triggered on demand, the domain visibility is applied to the particular agent in action. When an agentic conversation is initiated through a trigger, the domain visibility is applied to the user who resolves the caller (in an incident record where the *Run as* attribute is set to *Caller*), when the conversation runs against the incident record.

Note: The **sys_domain** field is added to all AI agent tables to achieve domain separation in Now Assist AI agents. The sys_domain_path, which is available for domain separation, is enabled on your instance.

To understand more about the ServiceNow domain separation, see Exploring domain separation.

How domain separation works in Now Assist AI Agent Studio

Process separation is enabled through the use of the sys_overrides column in domain-aware tables. Any table that contains both the sys_domain and the sys_overrides fields can be configured to have different processes from the parent domain.

AI Agents support only configuration tables to be process separated. Below are the list of tables that are process separated:

- sn_aia_agent_config
- sn_aia_usecase_config_override

Domain separation in Now Assist AI agents supports:

- Agentic workflow discovery.
- AI agent and its tools can be active in the X domain and inactive in the Y domain.
- Memory category can be active in the X domain and inactive in the Y domain.
- sn_aia_property can be overridden in a different domain.
- Triggers can be overridden in different domain.

Note: AI agent and agentic workflow details can't be overridden in the different domains.

Related topics

Domain separation for service providers

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

AI Agent Analytics dashboard

Track the AI agent use and efficiency gain on your instance through the AI Agent Analytics dashboard. The dashboard can reveal trends in how AI agents are used to improve the time to resolution and the number of tasks closed.

Required ServiceNow AI Platform roles

To use the AI Agent Analytics dashboard, you must have either the sn_aia.viewer or the sn_aia.admin role.

If you want to change the dashboard, you must duplicate it and apply changes to the copy. You can redirect the Analytics page in AI Agent Studio to the new dashboard by replacing the dashboard sys_id in the system property *sn_aia.analytics_dashboard_sysid*.

Accessing the AI Agent Analytics dashboard

To open the dashboard, navigate to **All** > **AI Agent Studio** > **Analytics**.

You can also access the dashboard from the AI Agent Studio overview page. Go to the Recent agentic workflows and AI agents activity section and select the **View analytics** link.

Indicators

Most indicators are updated daily. The latency indicators are updated every 15 minutes.

Once you have installed Now Assist AI Agents, you can collect initial data by running the [Now Assist AI Agents] Historical Data Collection job. The other data collection jobs, [Now Assist AI Agents] Daily Data Collection and [Now Assist AI Agents] Periodic Data Collection, update the indicators.

Automated indicators for AI agents

Name	Description
Agentic workflow latency	Time taken for an agentic workflow to complete.
AI agent execution plan P90 latency	Time taken for 90% of AI agent execution plans in a system to be processed.
AI agent execution plan P95 latency	Time taken for 95% of AI agent executions plans in a system to be processed.
AI agent execution plan P99 latency	Time taken for 99% of AI agent execution plans in a system to be processed.
AI agent latency	Time taken for an AI agent in a system to be processed.
All agentic workflows	Total number of agentic workflows created before today.

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

Automated indicators for AI agents (continued)

Name	Description
All AI agents	Total number of AI agents created before today.
All tools	Total number of tools created before today.
LLM P90 latency	Time taken for 90% of LLM requests in a system to be processed.
LLM P95 latency	Time taken for 95% of LLM requests in a system to be processed.
LLM P99 latency	Time taken for 99% of LLM requests in a system to be processed.
Number of closed tasks	Total number of tasks closed.
Number of closed tasks with AI agent assist	Total number of tasks closed today that were closed with agentic workflows or AI agents.
Summed duration of closed tasks	Total time taken to close tasks.
Summed duration of closed tasks with AI agent assist	Total time taken to close tasks today that were closed with agentic workflows or AI agents.
Tool execution P90 latency	Time taken for 90% of tool executions in a system to be processed.
Tool execution P95 latency	Time taken for 95% of tool executions in a system to be processed.
Tool execution P99 latency	Time taken for 99% of tool executions in a system to be processed.
Tool latency	Time taken for a tool in a system to be processed.
Total number of CS conversations	Total number of conversations in Now Assist panel or Virtual Agent.

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

Automated indicators for AI agents (continued)

Name	Description
Total number of execution plans	Total number of execution plans created by agentic workflows and AI agents today.
Total number of execution tasks	Total number of tasks executed by agentic workflows and AI agents today.
Total number of tool executions	Total number of tool executions in agentic workflows and AI agents today.
Use case execution plan P90 latency	Time taken for 90% of agentic workflows in a system to be processed.
Use case execution plan P95 latency	Time taken for 95% of agentic workflows in a system to be processed.
Use case execution plan P99 latency	Time taken for 99% of agentic workflows in a system to be processed.

Formula indicators for AI agents

Name	Description
% of conversations with an AI agent or agentic workflow defined	Number of conversations in Now Assist panel or Virtual Agent with a defined AI agentic or agentic workflow divided by the total number of conversations.
Average time to close a task	Average time taken for a task to go from creation to closure.
Average time to close a task with AI agent assist	Average time taken for a task to go from creation to closure with an agentic workflow or AI agent involved.
Efficiency gain	Percentage efficiency gain comparing average time taken to close a task with AI agent assist against a task closed without AI agent assist.
Percentage of tasks closed using AI agents	Number of tasks closed using AI agents and agentic workflows divided by the total number of tasks.

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

Breakdowns

Different breakdowns enable you to divide the data differently to track specific aspects of the AI agent usage. Not all breakdowns are available for all indicators.

- Agentic workflow
- AI agent
- AI agent execution status
- Latency metric
- Tool
- Tool latency metric

Filters

You can filter data to review a subset of trends. The following table lists the available filters on the AI Agent Analytics dashboard.

Filters on the AI Agent Analytics dashboard pages

Page	Available filters
Status	Time created
Insights	• Time created • Agentic workflow • AI agent • AI Agent Type
Assist consumption	• Time created • Agentic workflow • AI agent • AI Agent Type
Performance	• Time created • Agentic workflow • AI agent • AI Agent Type • Tool
Troubleshooting	Time created

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

Data visualizations

Visualizations are graphic elements that can be used to see data trends, percentages, and scores. The AI Agent Analytics dashboard includes the following visualizations.

Data visualizations on the AI Agent Analytics dashboard overview page

Title	Type	Description
Average inferred CSAT for AI agents in last 7 days	Score card	AI-generated score estimating what a customer would rate the AI response. 1 being the user is extremely dissatisfied and 5 being the user is extremely satisfied with the interaction.
Execution tasks in last 7 days	Pie chart	Execution tasks divided by status (Success, Ongoing, Cancelled, Queued).
AI Agent errors in last 90 days	Line graph	Number of errors seen by AI agents within the last 90 days broken down by week.
AI agent tool type	Pie chart	Tools divided by type. Example: Script, Flow action, Subflow, and so on.
AI agent tool execution mode	Pie chart	Tools divided by either autonomous or supervised execution mode.
Execution plans in last 7 days	Pie chart	Execution plans divided by status (Completed, Terminated, Wrap Up, or In Progress)
AI agent executions	List	List of AI agent executions. The list is broken down by AI agent and AI agent execution status. Columns span for multiple weeks, and there are change percentages and trend lines comparing the current against the previous week.
Number of agentic workflows	Trend line	Total number of agentic workflows. The agentic workflows created in the last day and a trend line are shown.
Number of AI agents	Trend line	Total number of AI agents. The AI agents created in the last day and a trend line are shown.
Number of tools	Trend line	Total number of tools. The tools created in the last day and a trend line are shown.

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

Data visualizations on the AI Agent Analytics dashboard overview page (continued)

Title	Type	Description
AI agent execution plans	List	List of AI agent execution plans. The list is broken down by agentic workflow. Columns span for multiple weeks, and there are change percentages and trend lines comparing the current against the previous week.

Data visualizations on the AI Agent Analytics dashboard status page

Title	Type	Description
AI agent execution plans	Score cards	Four score cards, one for each execution plan status (Ready, In progress right now, Completed, and Terminated).
AI agent executions	Score cards	Four score cards, one for each execution plan status (In progress right now, Completed, Successful, and Errors/Cancelled).
AI agent/ tool executions	List	List of AI agent and tool executions. The list is broken down by AI agent and tool, and there are columns for each status (Cancelled, Ongoing, Queued, and Success).

Data visualizations on the AI Agent Analytics dashboard insights page

Title	Type	Description
Tasks closed using AI agents	Trend line	Percentage of tasks closed using AI agents. Trend line is shown.
Task closure efficiency gain using AI agents	Trend line	Percentage efficiency gained with agentic workflows and AI agents. Trend line is shown.
Average close time assisted by AI agents	Trend line	Average time taken for a task to go from creation to closure with the involvement of an AI agent or agentic workflow. Trend line is shown.
Inferred CSAT of AI agents	Score card	Estimated CSAT score calculated using AI based on the subjective quality of the AI agents' responses. 1 being the user is extremely dissatisfied and 5 being the user is extremely satisfied with the interaction.

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

Data visualizations on the AI Agent Analytics dashboard insights page (continued)

Title	Type	Description
Inferred Session CSAT	Score card	Estimated CSAT score calculated using AI based on the overall session interactions of a user. 1 being the user is extremely dissatisfied and 5 being the user is extremely satisfied with the interaction.
Top user intents when using agentic solutions	Bar chart	Most common user intents as interpreted by the LLM. Intents are grouped by Group Action Framework.
Average AI agent CSAT by intent	Bar chart	Average CSAT score for each user intent interpreted by the LLM. Intents are grouped by Group Action Framework.
Transfers and escalations using AI agents	Bar chart	An estimated score using AI that indicates whether the user requested an escalation or there was a transfer to another AI agent or human.
User effort required with AI agents	Bar chart	Counts of AI-generated scores that measure the time and energy users had to put in during their interaction with an AI agent. Values are on a 3-point scale of low, medium, and high.
AI agent recommended next steps	Bar chart	Counts of AI-generated scores that indicate whether the AI agent recapped the outcome and provided clear instructions or guidance on what the customer should do next, if applicable. Values are on a 3-point scale of low, medium, and high.
AI agent resolved the users request	Score card	An AI-generated score that indicates whether the AI agent resolved the user's request. Values can be yes, no or unknown.
User frustration with AI agents	Bar chart	An AI-generated score that indicates whether the user expressed frustration at any point in the conversation.
AI agent empathy	Score card	An AI-generated score that indicates whether the AI agent was empathetic and friendly. Values are on a 3-point scale of low, medium, and high.
User or AI agent confusion	Bar chart	Counts of an AI-generated score that indicates whether the user or AI agent is confused at any point of the conversation.
Top 10 agentic workflows	Bar chart	Top 10 agentic workflows by highest CSAT score.

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

Data visualizations on the AI Agent Analytics dashboard insights page (continued)

Title	Type	Description
with highest session CSAT		
Bottom 10 agentic workflows with highest session CSAT	Bar chart	Bottom 10 agentic workflows by CSAT score.
Top 10 AI agents with highest session CSAT	Bar chart	Top 10 AI agents by highest CSAT score.
Bottom 10 AI agents with highest session CSAT	Bar chart	Bottom 10 AI agents by CSAT score.
AI agent conversations daily	Trend line	Total number of AI agent conversations in Now Assist panel or Virtual Agent. Trend line is shown.
% of conversations using AI agents	Trend line	Percentage of AI agent conversations in Now Assist panel or Virtual Agent that used agentic workflows or AI agents. Trend line is shown.
Efficiency gain supporting indicators	List	List of indicators. The list is broken down by indicator. Columns span across a time span.

Data visualizations on the AI Agent Analytics dashboard assist consumption page

Title	Type	Description
Sum of AI agent assists	Score card	Total number of assists used on the instance for both AI agents and agentic workflows.
Sum of AI agent assists per tier	Pie chart	Totals and percentages of assists divided by tier. An assist is marked as small until the number of tool executions exceeds 20.

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

Data visualizations on the AI Agent Analytics dashboard assist consumption page (continued)

Title	Type	Description
Count of AI agent executions per assist tier.	Pie chart	Count of assist tiers for both AI agents and agentic workflows. An assist is marked as small until the number of tool executions exceeds 20.
Top 10 agentic workflows consuming assists	Pie chart	Percentages of assists consumed among the top 10 agentic workflows consuming assists.
Sum of AI agents assists	Bar chart	Total number of assists used by AI agents broken down by day.
Count of AI agent executions per assist tier	Bar chart	Count of assists used by AI agents broken down by assist tier and day.
Count of AI agent executions per agentic workflow	Bar chart	Count of assists used by AI agents broken down by agentic workflow.
Sum of AI agent assists per AI agent	Bar chart	Number of assists used by AI agents broken down by AI agent.
Average tool execution count per agentic workflow	List	Count of tools executed by an agentic workflow broken down by agentic workflow and day.
Average tool execution count per assist tier	List	Count of tools executed by an agentic workflow broken down by assist tier and day.

Data visualizations on the AI Agent Analytics dashboard performance page

Title	Type	Description
P90 agentic workflow latency	Score card	Time in seconds of 90% of agentic workflows in a system to be processed.
P95 agentic	Score card	Time in seconds of 95% of agentic workflows in a system to be processed.

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

Data visualizations on the AI Agent Analytics dashboard performance page (continued)

Title	Type	Description
workflow latency		
P99 agentic workflow latency	Score card	Time in seconds of 99% of agentic workflows in a system to be processed.
Total agentic workflows executed per hour	Bar chart	Visualization of what time of day agentic workflows have been executed by hour.
Total agentic workflows errors per hour	Bar chart	Visualization of what time of day agentic workflows have been executed and experienced an error by hour.
Agentic workflow latency	Line chart	Agentic workflow latency over time. Uses 3 different indicators: P90 tracks the fastest 90% of executions, P95 tracks the fastest 95% of executions, and P99 tracks the fastest 99% of executions.
Agentic workflow latency (weekly average)	Line chart	Agentic workflow latency over time. Uses 3 different indicators: P90 tracks the fastest 90% of executions, P95 tracks the fastest 95% of executions, and P99 tracks the fastest 99% of executions.
Agentic workflow latency (daily average)	Table	Average agentic workflow latency today. Each row is an agentic workflow and each column is a different latency metric. Uses 3 different indicators: P90 tracks the fastest 90% of executions, P95 tracks the fastest 95% of executions, and P99 tracks the fastest 99% of executions.
P90 AI agent latency	Score card	Time in seconds of 90% of AI agent executions in a system to be processed.
P95 AI agent latency	Score card	Time in seconds of 95% of AI agent executions in a system to be processed.

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

Data visualizations on the AI Agent Analytics dashboard performance page (continued)

Title	Type	Description
P99 AI agent latency	Score card	Time in seconds of 99% of AI agent executions in a system to be processed.
Total AI agents executed per hour	Bar chart	Visualization of what time of day AI agent executions have been executed by hour.
Total AI agent errors per hour	Bar chart	Visualization of what time of day AI agent executions have been executed and experienced an error by hour.
AI agent latency	Line chart	AI agent latency within the time filter. Uses 3 different indicators: P90 tracks the fastest 90% of executions, P95 tracks the fastest 95% of executions, and P99 tracks the fastest 99% of executions.
AI agent latency (weekly average)	Line chart	AI agent latency over time by week. Uses 3 different indicators: P90 tracks the fastest 90% of executions, P95 tracks the fastest 95% of executions, and P99 tracks the fastest 99% of executions.
AI agent latency (daily average)	Table	Average AI agent latency today. Each row is an AI agent and each column is a different latency metric. Uses 3 different indicators: P90 tracks the fastest 90% of executions, P95 tracks the fastest 95% of executions, and P99 tracks the fastest 99% of executions.
P90 tool latency	Score card	Time in seconds of 90% of tool executions in a system to be processed.
P95 tool latency	Score card	Time in seconds of 95% of tool executions in a system to be processed.
P99 tool latency	Score card	Time in seconds of 99% of tool executions in a system to be processed.
Total tools executed per hour	Bar chart	Visualization of what time of day tool executions have been executed by hour.

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

Data visualizations on the AI Agent Analytics dashboard performance page (continued)

Title	Type	Description
Total tool errors per hour	Bar chart	Visualization of what time of day tool executions have been executed and experienced an error by hour.
Tool latency	Line chart	Tool latency within the time filter. Uses 3 different indicators: P90 tracks the fastest 90% of executions, P95 tracks the fastest 95% of executions, and P99 tracks the fastest 99% of executions.
Tool latency (weekly average)	Line chart	Tool latency over time by week. Uses 3 different indicators: P90 tracks the fastest 90% of executions, P95 tracks the fastest 95% of executions, and P99 tracks the fastest 99% of executions.
Tool latency (daily average)	Table	Average tool latency today. Each row is a tool and each column is a different latency metric. Uses 3 different indicators: P90 tracks the fastest 90% of executions, P95 tracks the fastest 95% of executions, and P99 tracks the fastest 99% of executions.
P90 LLM latency	Score card	Time in seconds of 90% of LLM requests in a system to be processed.
P95 LLM latency	Score card	Time in seconds of 95% of LLM requests in a system to be processed.
P99 LLM latency	Score card	Time in seconds of 99% of LLM requests in a system to be processed.
LLM latency	Line chart	Tool latency within the time filter. Uses 3 different indicators: P90 tracks the fastest 90% of executions, P95 tracks the fastest 95% of executions, and P99 tracks the fastest 99% of executions.
LLM latency (weekly average)	Line chart	Tool latency over time by week. Uses 3 different indicators: P90 tracks the fastest 90% of executions, P95 tracks the fastest 95% of executions, and P99 tracks the fastest 99% of executions.
LLM calls per hour	Bar chart	Visualization of what time of day LLM calls have been made by hour.
LLM errors per hour	Bar chart	Visualization of what time of day LLM calls have been made and experienced an error by hour.

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

Data visualizations on the AI Agent Analytics dashboard troubleshooting page

Name	Type	Description
Agentic workflow errors	List	List of agentic workflow errors.

For more information on security controls for agentic AI, see Security for AI agents.

Data visualizations on the AI Agent Analytics dashboard security page

Name	Type	Description
Total blocked executions	Score card	Total number of agentic AI executions blocked because of security controls.
Total blocked executions	List	List of executions blocked.
Top agentic workflows with blocked executions	List	List of agentic workflows with the most blocked executions.
Top AI agents with blocked executions	List	List of AI agents with the most blocked executions.
Agentic Workflows w/o ACL defined	Score card	Total number of agentic workflows missing an ACL for user access.
Agentic Workflows ACLs	Pie chart	Pie chart displaying the percentage of agentic workflows with and without ACLs for user access.
Agentic Workflows w/o ACL defined	List	List of agentic workflows missing an ACL for user access.
AI agents with no ACLs defined	Score card	Total number of AI agents missing an ACL for user access.
AI agent ACLs	Pie chart	Pie chart displaying the percentage of AI agents with and without ACLs for user access.
AI agents w/o ACL defined	List	List of AI agents missing an ACL for user access.
Agentic workflows without role masking	Score card	Total number of agentic workflows without roles identified for role masking and data access.

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

Data visualizations on the AI Agent Analytics dashboard security page (continued)

Name	Type	Description
Agentic workflows without role masking	List	List of agentic workflows without roles identified for role masking and data access.
AI agents without role masking	Score card	Total number of AI agents without roles identified for role masking and data access.
AI agents without role masking	List	List of AI agents without roles identified for role masking and data access.
Agentic workflows running as dynamic user	Score card	Total number of agentic workflows running as a dynamic user.
Agentic workflows running as dynamic user	List	List of agentic workflows running as a dynamic user.
AI agents running as dynamic user	Score card	Total number of AI agents running as a dynamic user.
AI agents running as dynamic user	List	List of AI agents running as a dynamic user.
Agentic workflows running as AI user	Score card	Total number of agentic workflows running as an AI user.
Agentic workflows running as AI user	List	List of agentic workflows running as an AI user.
AI agents running as AI user	Score card	Total number of AI agents running as an AI user.
AI agents running as AI user	List	List of AI agents running as an AI user.

Agentic evaluation run results

Learn about agentic evaluation runs and the meaning behind different evaluation scores from the agentic evaluation results page.

Agentic evaluations overview

Agentic evaluations measure how well AI agents and agentic workflows are accomplishing their objectives. A Now LLM Service model judges the AI agent or agentic workflow based on the execution logs. The results page of an evaluation run shows multiple metrics and scores measuring task completeness and tool use.

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

If you run an overall task completion evaluation, the results page shows recommended actions for the AI agent or agentic workflow. Recommended actions give you suggestions for deployment or improvement to help verify that the agentic workflows that you deploy are performing up to your standards.

After you've reviewed your evaluation results, you can archive your evaluation or copy it to run another evaluation with the same parameters and dataset.

You can export the evaluation results as a report. The report is formatted as a .csv file that includes the individual sys_ids of the execution records and the metric scores for each.

For more information on AI agent usage and other analytics, you can review the AI Agent Analytics dashboard in the AI Agent Studio.

Evaluation results overview

For each evaluation method that you execute, the results page displays an overall score for the agentic workflow with a percentage of successful record evaluations and a label of Excellent, Good, Moderate, or Poor. You can change the metric thresholds for each label by selecting **Customize metric thresholds**.

In addition to the overall task completeness results, you can review a summary of the results of the other metrics.

Overall task completeness evaluation run results

Label	Description	Recommended action	Default threshold
Excellent	Tasks were consistently performed at a high standard. The agentic workflow or AI agent is working well.	Proceed with confidence	90%–100%
Good	Most tasks were performed successfully, but some performance inconsistencies suggest areas for improvement.	Deploy with caution	70%–89%
Moderate	A significant number of tasks weren't fully completed. Performance is below the desired level.	Investigate the root causes of poor task completion	50%–69%
Poor	The agentic workflow is consistently failing to complete tasks adequately. Major issues are present.	Do not deploy	0%–49%

Individual record metric scores

Evaluations are run against the log tables of agentic workflow executions. Each record is individually scored for each evaluation plan that you run. Individual record evaluations are scored according to the following metrics.

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

Overall task completeness record metric scores

The overall task completeness metric assesses whether an AI agent successfully completes its assigned task. It evaluates the execution logs of the agent, ensuring that all required steps were taken and the task was logically and effectively completed.

Number	Score	Description
3	Successful	The main task was fully completed. All subtasks were resolved, and the steps followed a logical sequence with no critical errors.
2	Partially successful	The task was partially completed. Some subtasks remain unresolved or inefficiencies affected the process.
1	Unsuccessful	The task wasn't completed. Critical subtasks were abandoned or unresolved or the execution failed entirely.

Tool performance record metric scores

The tool performance evaluation metric assesses an AI agent's ability to select the most appropriate tool for each step while completing a task.

Number	Score	Description
1	True	The right tool was chosen for the action in the plan.
0	False	The right tool wasn't chosen.

Tool calling records metric scores

The tool calling evaluation metric assesses whether an AI agent correctly constructs tool calls by validating the accuracy, completeness, and formatting of the inputs it provides.

Number	Score	Description
1	True	Input key completeness, input value correctness, and input format correctness are all successful.

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

Tool calling records metric scores

The tool calling evaluation metric assesses whether an AI agent correctly constructs tool calls by validating the accuracy, completeness, and formatting of the inputs it provides.

(continued)

Number	Score	Description
		• *`Input key completeness`*: 1 - True – All required parameters are present with exact name matches, and no unexpected parameters are included. • *`Input value correctness`*: 1 - True – Tool input values are correctly mapped. • *`Input format correctness`*: 1 - True – Tool inputs are in the correct format.
0	False	One or more of input key completeness, input value completeness, or input format completeness wasn't successful. • *`Input key completeness`*: 0 - False – A mandatory parameter is either missing, its name doesn't match exactly, or an unexpected parameter was found. • *`Input value correctness`*: 0 - False – Tool input values are not correctly mapped. • *`Input format correctness`*: 0 - False – Tool inputs are not in the correct format.

Note: The values of the sub-metrics are aggregated using an AND operator. If any one value is 0, then the entire tool calling records metric score will be 0.

Model Context Protocol Client

The ServiceNow® Model Context Protocol Client (MCP Client) enables you to access the Model Context Protocol tools that are hosted externally and published using an MCP Server in the ServiceNow® AI Agent Studio.

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

Get started

Explore



Learn more about MCP Client

Configure



Configure the Model Context Protocol Client in AI Agent Studio

Add MCP Server in AI Agent Studio



Add MCP Servers in AI Agent Studio

Reference



Learn more about the user roles, tables, and properties in MCP Client

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.

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.

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

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.

For more information, see the Now Assist documentation.

Exploring Model Context Protocol Client

The Model Context Protocol (MCP) is a standardized client-server protocol that enables AI applications to discover and interact seamlessly with external tools, data sources, and services. MCP facilitates communication between an AI host application (like AI Agent Studio), an MCP Client embedded in the host, and one or more MCP Servers that expose specific capabilities such as tools.

Model Context Protocol Client terminology

- **MCP Host**: The AI application environment where the tasks are initiated.
- **MCP Client**: Manages communication between the host and the MCP Servers, handling the requests and responses.

  The MCP Client architecture enables AI agents to extend their capabilities by invoking external systems dynamically, enabling tasks like accessing APIs, databases, or file systems in a standardized way.
- **MCP Server**: Provides external functionalities by exposing tools for AI agents to use.
- **MCP Tools**: The external tools exposed by the MCP Servers to be used by AI agents.

Why use MCP Tools

MCP tools enable seamless integration to:

- Connect your AI agent with a wide variety of external tools with minimal setup.
- Add multiple MCP tools to an AI agent to perform a broader range of tasks.
- Review and edit tool input parameters as needed before saving.

Configuring Model Context Protocol Client

Configure the MCP Client to execute agentic workflows with AI agents and mapped tools.

Configuring MCP Client enables you to:

- Incorporate external tools in AI agents built on ServiceNow.
- Access tools offered on external MCP Servers.
- Integrate MCP Client with the AI Agent Studio.

By setting up MCP Servers, AI agents in your environment can leverage external tools efficiently, enhancing their ability to perform complex tasks with up-to-date information.

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

Install Model Context Protocol Client

Install the MCP Client application on your ServiceNow instance to enable using the tools from the MCP Server in AI agents.

Before you begin

Role required: sn_mcp_client.admin

Verify that you:

- Install the AI Agent Studio plugin [com.snc.sn_aia] and the Generative AI Controller plugin [com.sn.generative.ai] to use the MCP Client.
- Enable the MCP Tool experience in your instance by setting the *sn_aia.enable_mcp_tool* system property to **true**.

Note: ServiceNow supports Protocol version 2025-03-26 of the MCP Servers for MCP Client.

Procedure

1. Navigate to **All** > **System Definition** > **Plugins**.
2. Search for Model Context Protocol Client plugin [sn_mcp_client].
3. Select **Install**.

Adding an MCP Server in AI Agent Studio

An MCP Server hosts the APIs and tools required by an AI application, enabling it to receive and process calls from MCP Clients to govern ingress traffic and to promote secure and efficient access to tools. Adding an MCP Server in the AI Agent Studio helps you to leverage the Model Context Protocol as a tool in an AI agent.

Connecting an MCP Server with the AI Agent Studio simplifies the integration process, making it easier to discover and invoke services. This connection enhances security and governance, providing a streamlined setup experience for administrators.

Adding an MCP Server requires you to add an MCP Server in the AI Agent Studio. You can add an MCP Server with one of the following authentication options:

1. **OAuth 2.1**: Helps add an MCP Server with an authentication code. For more information, see Add an MCP Server with OAuth 2.1.
2. **API Key**: Helps add an MCP Server with an API Key. For more information, see Add an MCP Server with API Key.
3. **Others**: Helps add an MCP Server in a manual way by selecting a Connection and Credential Alias record. For more information, see Add an MCP Server with Connection and Credential Alias.

Note: You must authenticate the users with the MCP Server to add the MCP tool to an AI agent and without prior authentication, you can't add the MCP Servers.

Add an MCP Server with OAuth 2.1

Add an MCP Server with OAuth 2.1 in the AI Agent Studio.

Before you begin

Role required: sn_mcp_client.admin

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

About this task

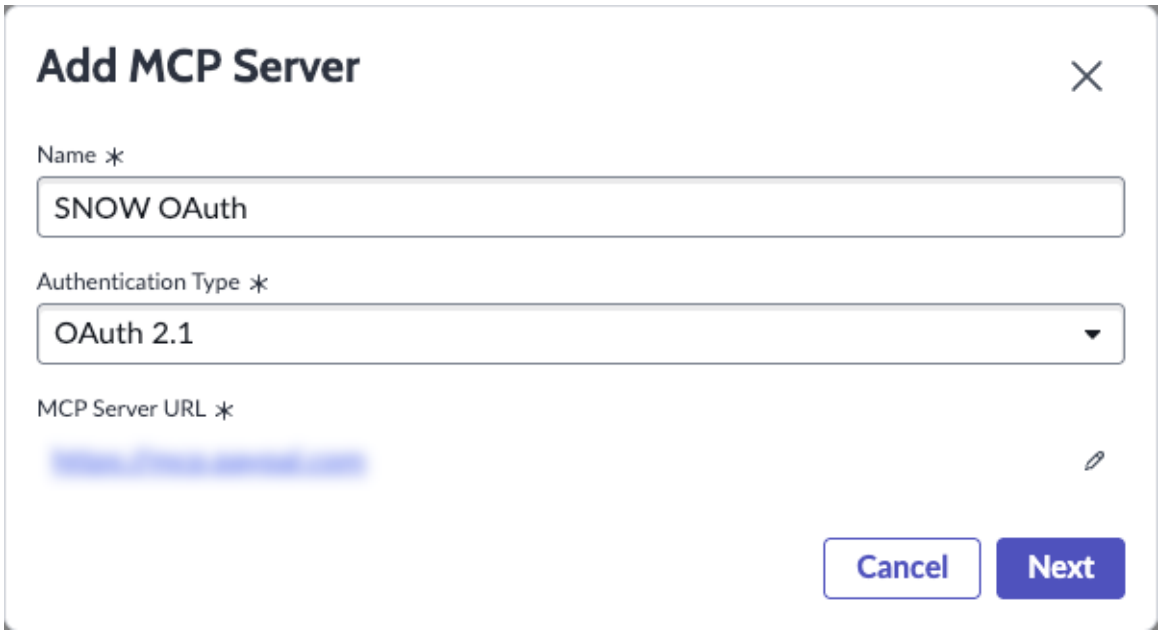
OAuth is the usual method to authenticate the MCP servers. If your MCP server supports it, choose OAuth 2.1 in the Authentication method.

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Settings**.
2. Navigate to **Manage MCP Servers** and select **New**.
3. On the form, fill in the fields.

Add MCP Server

Field	Description
Name	Name for your MCP Server.
Authentication Type	The authentication type with which you want to add your MCP Server. Select **OAuth 2.1**.
MCP Server URL	The web address of your MCP Server.



4. Select **Next**.
   The form extends with the registration details.
5. On the form, fill in the fields:

   The dynamic client registration auto-populates the data in the form fields.

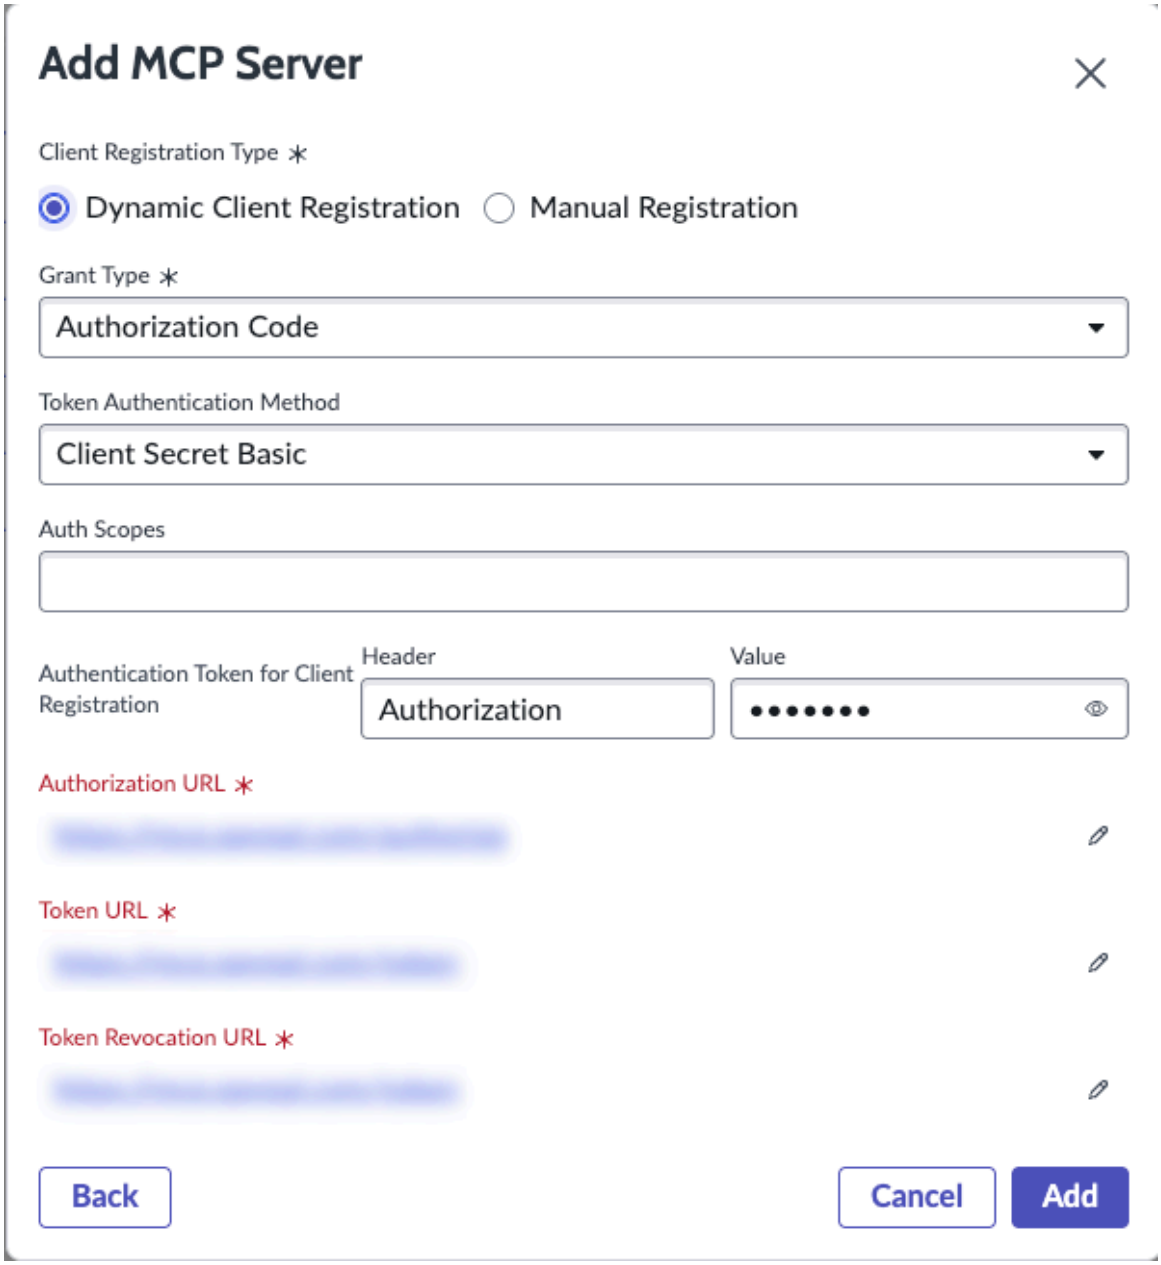
> **Note:** If dynamic registration is supported by the MCP Server, you'll be prompted to confirm the grant type and other inputs. If dynamic registration isn't supported, you can continue with the Manual Registration method in Step 6.

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

Dynamic Client Registration

Field	Description
Grant Type	The way you want to obtain an access token for accessing protected resources. One of the following two options is used here with auto-populated data. ○ **Authorization Code**: A system-generated code that is used for granting access to an application or resource. ○ **Client Credentials**: Access token provided by the administrator to access the application directly.
Token Authentication Method	Unique token granted to access resources. Select the value from the drop-down list.
Auth Scopes	Define the specific permissions for an application. **Note:** The Auth scopes must be comma-separated values.
Authentication Token for Client Registration	Authentication token used to verify client's identity and grant it permission to register with an authorization token. Provide the following values: ○ **Header**: Authorization header. ○ **Value**: Authentication password.
Authorization URL	Web address of your server to direct authorization.
Token URL	Web address containing the token.
Token Revocation URL	Web address to revoke the previously provided token.

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



The MCP Server with dynamic client registration is added as a simple connection and credential alias.

6. On the form, fill in the fields:

Note:

- If your MCP Server doesn't support Dynamic Client Registration, then you can register the OAuth client for AI agents in your MCP Server manually and update those client details in AI Agent Studio.
- The manual client registration doesn't auto-populate the data in the form fields and must be manually provided.

Manual Registration

Field	Description
Grant Type	The way you want to obtain an access token for accessing protected resources.

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

Field	Description
	You can select one of the two options: ○ **Authorization Code**: A system-generated code that is used for grating access to an application or resource. ○ **Client Credentials**: Access token provided by the administrator to access the application directly.
Token Authentication Method	Unique token granted to access resources. Select the value from the drop-down list.
Client ID	Unique ID for authentication and authorization purposes.
Client Secret	Unique credentials used for authentication.
Auth Scopes	Define the specific permissions for an application. **Note:** The Auth scopes must be comma-separated values.
Authentication Token for Client Registration	Authentication token used to verify client's identity and grant it permission to register with an authorization token. Provide the following values: ○ **Header**: Authorization header. ○ **Value**: Authentication password.
Authorization URL	Web address of your server to direct authorization.
Token URL	Web address containing the token.
Token Revocation URL	Web address to revoke the previously provided token.

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



7. Select **Add**.
   Selecting Add takes you to the MCP Server record where you must authenticate it for adding an MCP tool to an AI agent.
8. Select **Authenticate**.

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

Selecting **Authenticate** takes you to the third-party authorization page where you must select a workspace and select **Authorize**.

Note: When the OAuth access or refresh tokens aren't available, you must authenticate the OAuth configuration on the MCP Server record to add it as a tool to an AI agent.

Add an MCP Server with API Key

Add an MCP Server with an API Key in the AI Agent Studio.

Before you begin

Role required: sn_mcp_client.admin

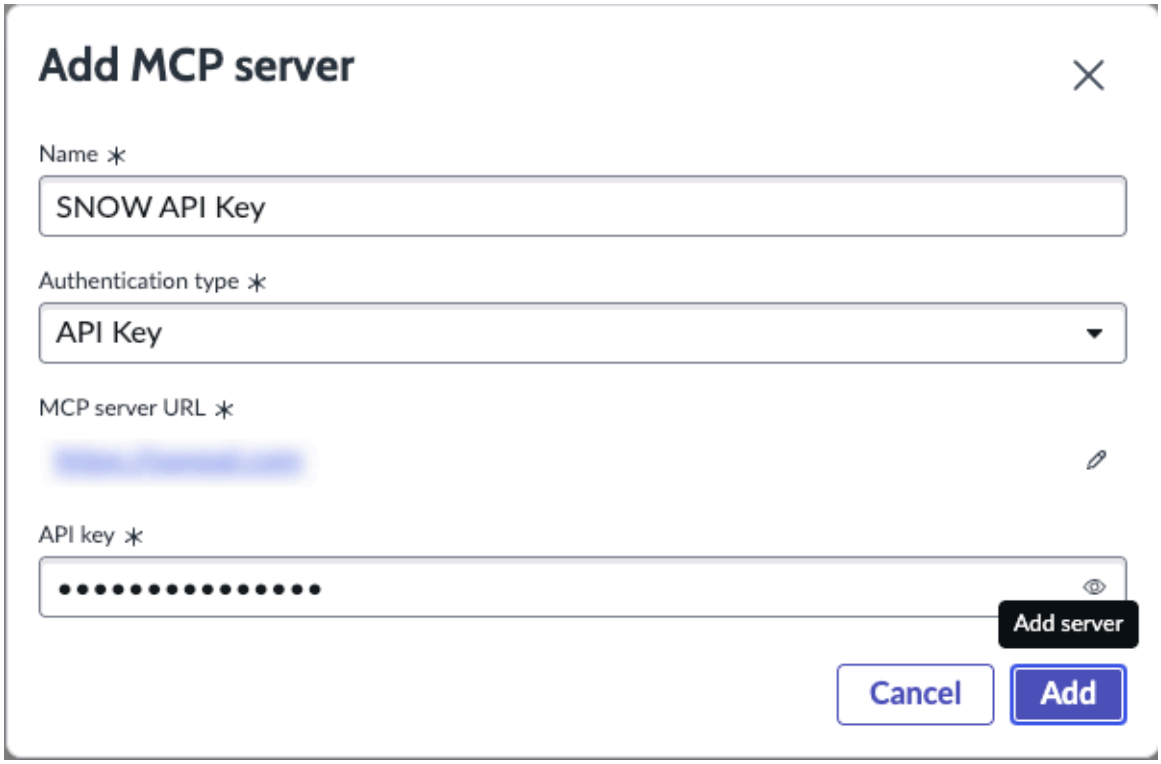
Procedure

1. Navigate to **All** > **AI Agent Studio** > **Settings**.
2. Navigate to **Manage MCP Servers** and select **New**.
3. On the form, fill in the fields.

Add MCP Server

Field	Description
Name	Name for your MCP Server.
Authentication Type	The authentication type with which you want to add your MCP Server. Select **API Key**.
MCP Server URL	The web address of your MCP Server.
API Key	A unique code or password to identify and authenticate the user or application when accessing the API. **Note:** The API Key adds a Connection alias dynamically at runtime and maps with the MCP Server.

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



4. Select **Add**.

Result
An MCP Server with the name you provided appears in the list of MCP Servers.

Add an MCP Server with Connection and Credential Alias

Add an MCP Server by selecting a Connection and Credential Alias record.

Before you begin

- Role required: sn_mcp_client.admin
- Verify that you have a Connection and Credential alias record created before adding an MCP Server with Others as the option.

  For more information, see Create a Connection & Credential alias.

About this task
For using the Authentication methods not supported in MCP Server, you can use the Connection and Credential Alias method to create the Alias manually and then update it in the MCP Server. For using OAuth based authentication method, first create the OAuth entity profile and then create the Connection and Credential Alias. For other authentication methods like API Key, Basic Auth, AWS Credentials, SSH Credentials, create the connection and credential Alias directly.

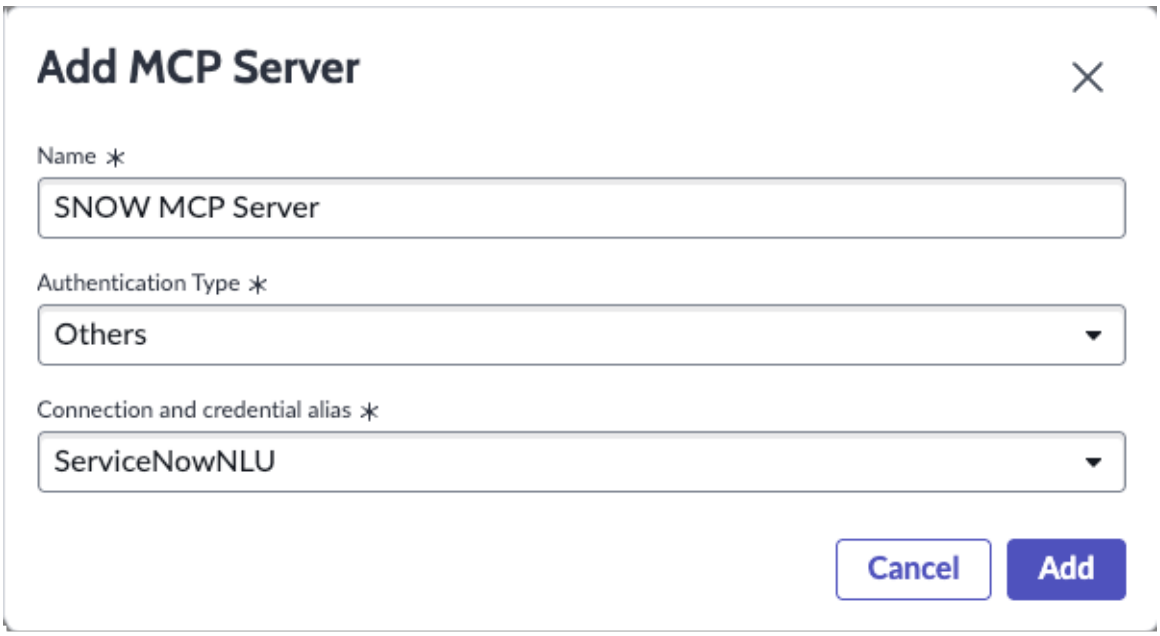
Procedure

1. Navigate to **All** > **AI Agent Studio** > **Settings**.
2. Navigate to **Manage MCP Servers** and select **New**.
3. On the form, fill in the fields.

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

Add MCP Server

Field	Description
Name	Name for your MCP Server.
Authentication Type	The authentication type with which you want to add your MCP Server. Select **Others**.
Connection and credential alias	Select a Connection and credential alias record to map with your MCP Server.



4. Select **Add**.

Note: Authenticate the MCP Server if you see the message requiring it.

Result
An MCP Server with the name you provided appears in the list of MCP Servers.

Add an MCP server to an AI agent

Add an MCP tool to an AI agent in the AI Agent Studio so that your users can access the MCP server.

Before you begin
Role required: sn_aia.admin

Procedure
1. Navigate to **All** > *AI Agent Studio* > **Create and manage** > **AI agents**.
2. Open the AI agent that you want to add a Model Context Protocol tool to and navigate to the Add tools and information section.
3. In the Add tool drop-down list, select **MCP server tool**.
4. On the form, fill in the fields.

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

Add a Model Context Protocol Tool

Field	Description
Select Model Context Protocol Server	The Model Context Protocol Server that you want to add to the tool. **Note:** Verify that the MCP Server is authenticated before adding the MCP tool to an AI agent. If the MCP Server isn't authenticated, then you receive an error message - "Authentication is needed to display tools. Open the MCP server record to authenticate."
Select Tool	The tool that you want to add from the Model Context Protocol Server.
Name	Name of the Model Context Protocol tool.
Tool and Description	Description of the Model Context Protocol and what it's going to do to assist your AI agent. **Note:** This description is sent to the LLM (LLM).
Execution Mode	Mode of execution for your selected Model Context Protocol tool: ○ **Supervised**: Inputs from your live agent are required during the execution of this Model Context Protocol tool while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this Model Context Protocol tool while the AI agent runs.
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No** If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that can be used in assistants built with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down.

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

Field	Description
	○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this transformation step, the tool output will be directly mapped to the widget. ▪ **Yes** ▪ **No** ○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No** **Note:** The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
Advanced settings	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats: ○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution. ○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

5. Select **Add**.
   A Model Context Protocol tool is added in the Model Context Protocol Tools list on the Add tools and information page.

Test an AI agent

Analyze the performance of an AI agent with a Model Context Protocol tool added to it to verify that it functions the way that you defined it.

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

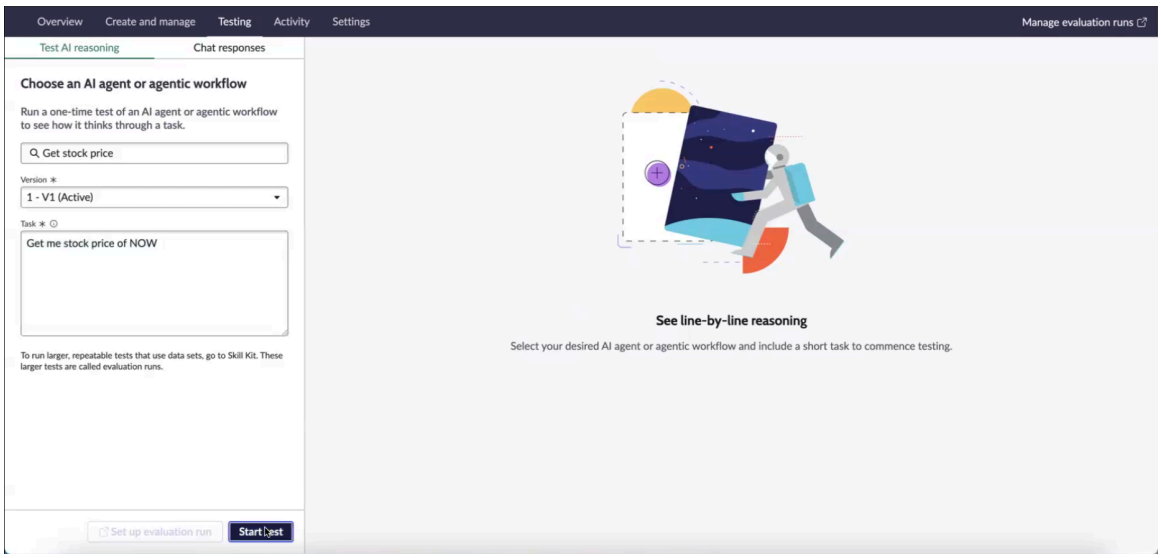
Before you begin
Role required: sn_aia.admin

About this task
After you create an AI agent and a Model Context Protocol to it as a tool, test it to see that it functions the way that you defined it. Search for the AI agent that you want to test and select the version from the Version drop-down and provide a task with a concise summary and a reference number in the **Task** field to begin testing the AI agent.

Procedure
1. Navigate to **All** > **AI Agent Studio** > **Testing**.
2. In the **Test AI reasoning** tab, use the Search bar to search and select the AI agent that you want to test.
3. In the Version drop-down list, select the version of the AI agent.
4. In the **Task** field, provide a concise summary of the task to be achieved.

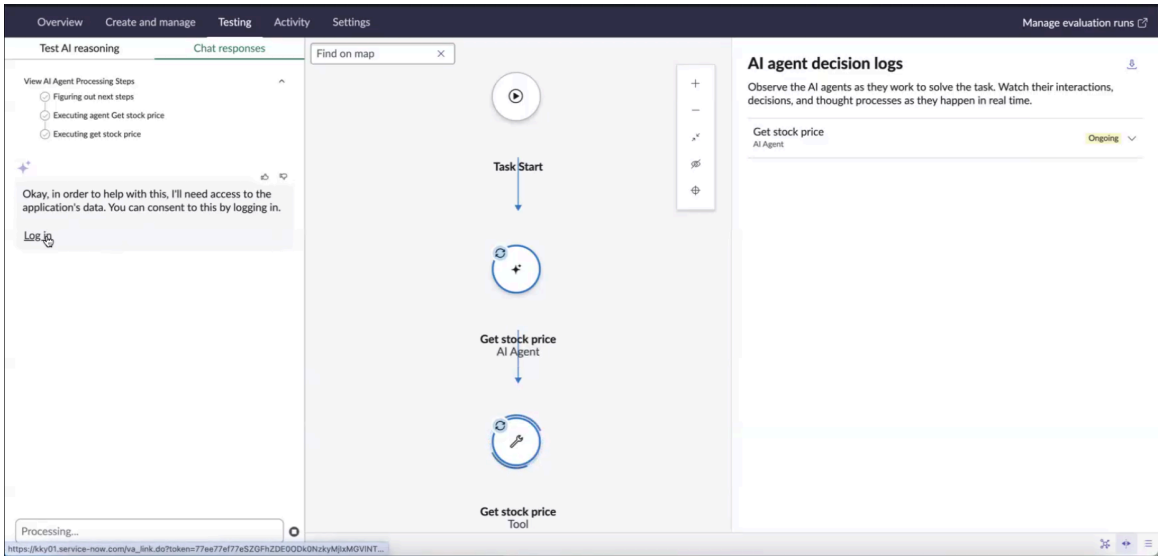
   Note: In the task summary, include a reference number or specific record for better results during your testing.

5. Select **Start test**.



You're directed to the **Chat responses** tab, where you can see Now Assist executing operations to test the AI agent.

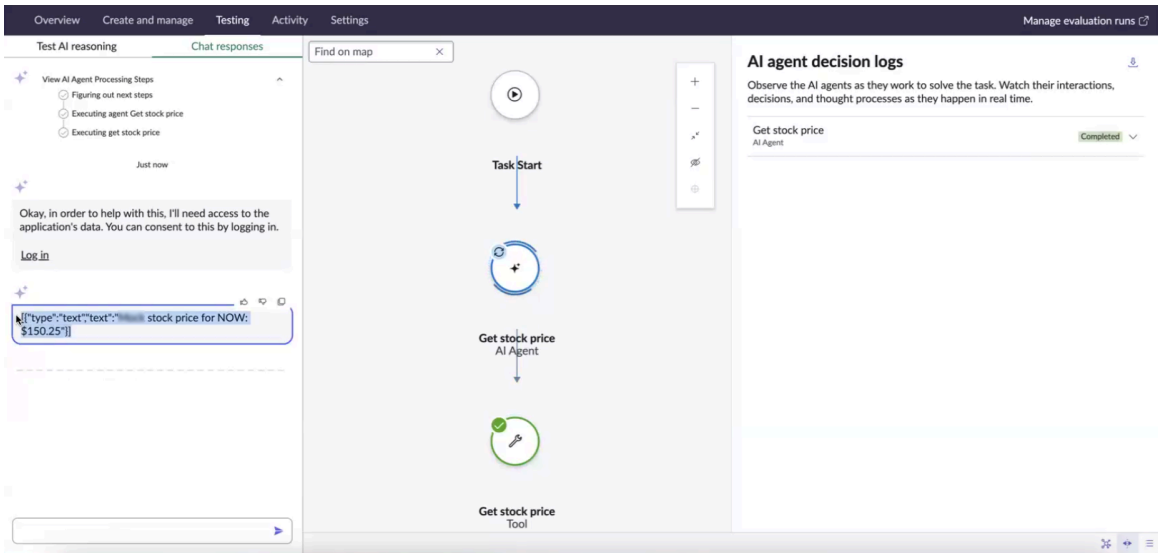
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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



6. Select **Log in** and authenticate the MCP server to complete the testing.

Note:

- Log in will be visible only if the server isn't authenticated.
- Enabling access to the MCP Server refreshes the whole Testing page and resumes testing in the Chat responses tab. The OAuth token used for the authentication is saved in the OAuth Credentials table [oauth_credential].



Model Context Protocol Client reference

Find more information about user roles, tables, and the different properties that are installed with the Model Context Protocol Client application.

MCP Client Roles

The following roles are installed with the Model Context Protocol Client plugin [sn_mcp_client].

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

Roles used in MCP Client

Role	Description
MCP Client Admin [sn_mcp_client.admin]	Administrator of the application. A user with the sn_mcp-client.admin role can create, read, update, and delete records. **Note:** The sn_aia.admin role contains the sn_mcp_client.admin role, a role inheritance on plugin activation.
MCP Client Viewer [sn_mcp_client.viewer]	Read-only access to the application. A user with the sn_mcp_client.viewer role has read and report access on all tables. **Note:** The sn_aia.viewer role contains the sn_mcp_client.viewer role, a role inheritance on plugin activation.

MCP Client system properties

The following system properties are installed with the Model Context Protocol Client plugin [sn_mcp_client].

System properties used for configuring MCP Client

System property	Description
sn_aia.enable_mcp_tool	Enables access to the MCP Client tool on your ServiceNow instance. Default value: **false**

MCP Client tables

The following tables are installed by default with the installation of the Model Context Protocol Client plugin [sn_mcp_client].

Tables used for configuring MCP Client

Table	Description
MCP Execution Logs [sn_mcp_execution_logs]	Lists the request and response, the method (tool call or tool list) that was carried out along with the mapped session. The execution logs table also helps to debug and shows if there are any errors.
MCP Client Server Session Mappings[sn_mcp_client_server_session_mapping]	Lists the sessions mapped to the server based on the capabilities support.
Model Context Protocol Severs [sn_mcp_server]	Lists the MCP Servers.

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

Now Assist Data Kit

Use ServiceNow®Now Assist Data Kit to add datasets to a data catalog. The curated data works with ServiceNow SDK to enable the AI skill development and evaluation.

https://player.vimeo.com/video/1088497625?h=40305c5af0&badge=0&autopause=0&player_id=0&app_id=58479

Get started

Now Assist Data KitNow Assist Data Kit



Explore

Learn how Now Assist Data Kit can help create custom datasets



Configure

Configure Now Assist Data Kit for prompts and skills



Use

Use Now Assist Data Kit to create datasets to add to data collections



Reference

Reference information for Now Assist Data Kit

Important: 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. Please check for availability updates in future releases.

Important: Some Now Assist products/features are currently available only for customers in some regions. Be sure to check for availability updates in future releases.

Troubleshoot and get help

- Ask questions and explore other resources for Now Assist Data Kit in the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

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-

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

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 Data Kit

The Now Assist Data Kit plugin for Now Assist enables you to add datasets to a data catalog and create collections for use in ServiceNow SDK.

Now Assist Data Kit overview

If the base system Now Assist skills don't fit your needs, you can use Now Assist Data Kit to create custom datasets and data collections that can be used in Now Assist Skill Kit for evaluation.

Now Assist Data Kit users

Users

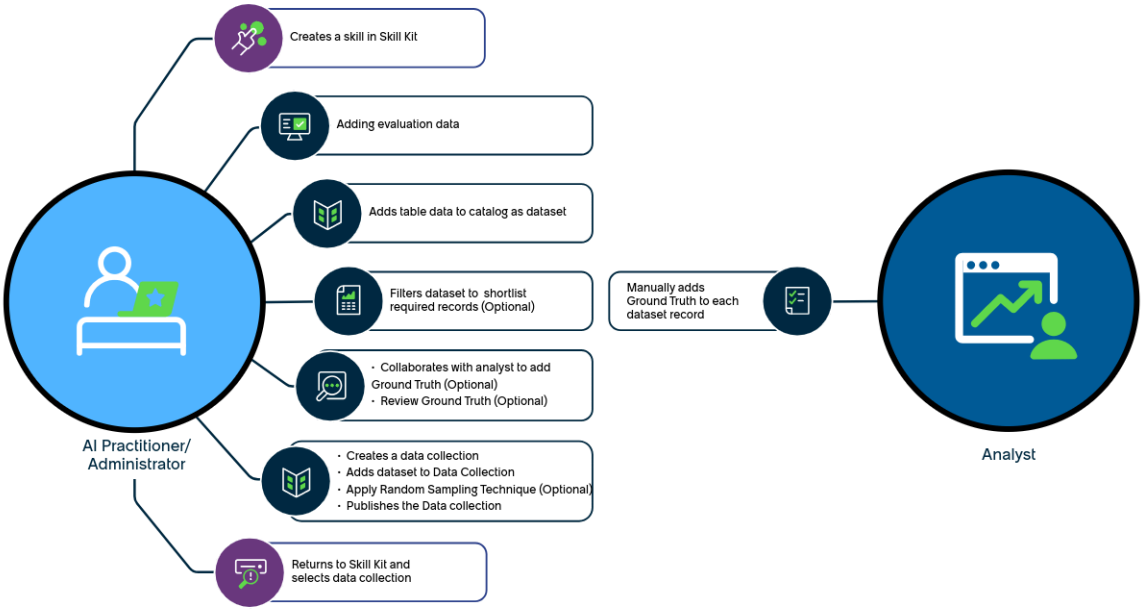
User	Description
AI practitioner	AI practitioners manage data set creation in Now Assist Data Kit. They develop and evaluate skills and other technical solutions for various use cases.
Analyst	Analysts confirm data quality for AI development and evaluation. They work with AI practitioners to follow data curation guidelines set for specific AI use cases.

Now Assist Data Kit workflow

The following diagram shows the user journey for Now Assist Data Kit.

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

User journey for Now Assist Data Kit



1. Create a skill in Now Assist Skill Kit
2. Create an evaluation dataset in Now Assist Data Kit
3. Add table data to the catalog as a dataset
4. Create a data collection
5. Add the dataset
6. Publish the data collection
7. Return to Now Assist Skill Kit and select the data collection to use in the evaluation tool
8. Run the evaluation and review the results to see if you must iterate upon the prompt
9. Publish the skill

Now Assist Data Kit benefits

Now Assist Data Kit benefits

Benefit	Feature	Users
Addition of datasets to a data catalog.	Add a dataset	AI developer
Create smaller datasets from existing datasets.	Create a derived dataset	AI developer
Merge multiple datasets and create a collection that can be used by Now Assist Skill Kit.	Activate a skill	AI developer

What to explore next

To learn more about configuring and using Now Assist Skill Kit, see:

- Configuring Now Assist Skill Kit
- Using Now Assist Skill Kit

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

Configuring Now Assist Data Kit

Configure prompts and skills for Now Assist Data Kit.

Configuration overview

1. Create a skill

   Create a skill in Now Assist Skill Kit.

2. Configure skill deployment settings

   Configure the deployment settings for the skill that you create. The deployment settings enable you to choose where the admin can find the skill in Now Assist Admin.

Install Now Assist Data Kit

Install Now Assist Data Kit to create datasets for a data collection.

Before you begin

Role required: admin or sn_data_kit.admin

Procedure

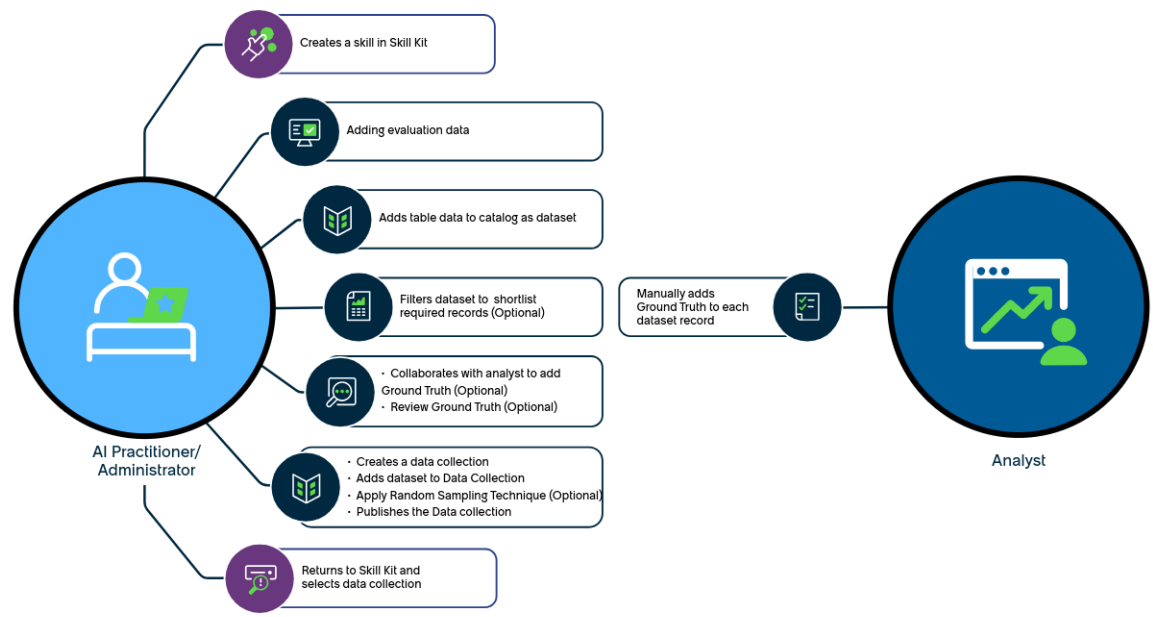
1. Navigate to **All** > **System Definition** > **Plugins**.
2. In the Store applications section, select Now Assist Data Kit.
3. Select **Install**.

Using Now Assist Data Kit

Use Now Assist Data Kit to add datasets to a data catalog to create collections for use in ServiceNow SDK.

Overview of Now Assist Data Kit

Now Assist Data Kit work flow



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

1. Create a skill

   Create a custom skill for Now Assist. Creating a custom skill enables you to have greater flexibility with the Now Assist generative AI capabilities.

2. Configure skill deployment settings

   Configure the deployment settings for the skill that you create. The deployment settings enable you to choose where the administrator can find the skill in Now Assist Admin.

3. Navigate to **Now Assist Data Kit** to Add a dataset.

Add a dataset

Add the data from a table to a data catalog as a dataset through generative AI by using the Now Assist Data Kit application. Adding a dataset is required to create and publish a data collection.

Before you begin

Role required: sn_data_kit.admin

Procedure

1. Navigate to **All** > *Now Assist Data Kit* > **Home**.
2. Navigate to Discover datasets and select **Get started**.
3. On the **Datasets** tab, select **New**.
4. Select where to curate data from.
   - I'll import data from Instance table
   - I'll import data from my computer
5. On the Choose data form, select the table and columns.
   If a column does not appear as a possible selection, then the field is not a supported data type. For example, Watch List fields are the glide_list datatype, which is not supported, so Watch List is not a selectable field.
6. **Optional:** Select **Add column via scripting** if you have columns, such as work notes or comments, that aren't stored in the table.
   This is an example script for adding worknotes.

```
(function generate(current) {
//Current is the Gliderecord object of current record.
// //Sample script for adding worknotes
// var gr = new GlideRecord("sys_journal_field");
// gr.addQuery('name', current.getTableName());
// gr.addQuery('element_id', current.getUniqueValue());
// gr.query();
// var records = [];
// while (gr.next()) {
//     var obj = {
//         "value": gr.getValue("value"),
//         "type": gr.getValue("element"),
//         "created": gr.getValue("sys_created_on"),
//         "created_by": gr.getValue("sys_created_by")
//     };
//     records.push(obj);
// }
```

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

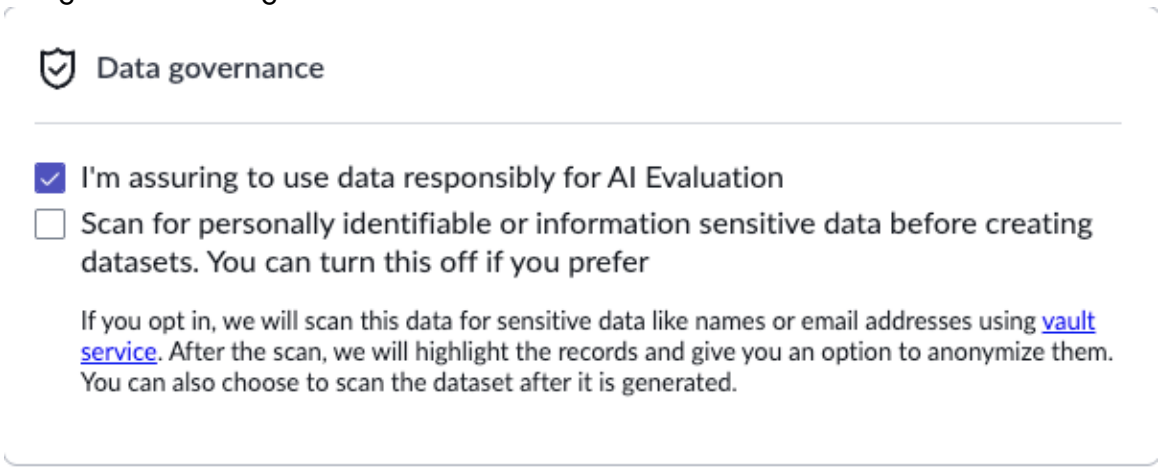
```
// return records;
})(current);
```

7. Select **Edit filter condition**.
8. Review the records and select **Continue**.
9. On the form, fill in the fields.

Add dataset info form

Field	Description
Dataset name	Name of the dataset.
Dataset description	Description of the dataset.

10. Add tags to identify the dataset.
11. Navigate to the Data governance section and select each check box.



   - I'm assuring to use data responsibly for AI Evaluation
   - Scan for personally identifiable or information sensitive data before creating datasets. You can turn this off if you prefer.

     Note: If you opt in, your data is scanned for sensitive data like names or email addresses using vault service. After the scan, records will be highlighted and give you an option to anonymize them. You can also choose to scan the dataset after it is generated.

12. Select **Generate dataset**.
    Fields of any data type are stored as strings when converted to a data asset.
    The dataset is added to the data assets.

What to do next

After your dataset is added to the data catalog, you can choose to create a smaller dataset by creating a derived dataset or adding a ground truth to your existing data set. For more information, see Create a derived dataset or Add a ground truth to each dataset record.

Create a derived dataset

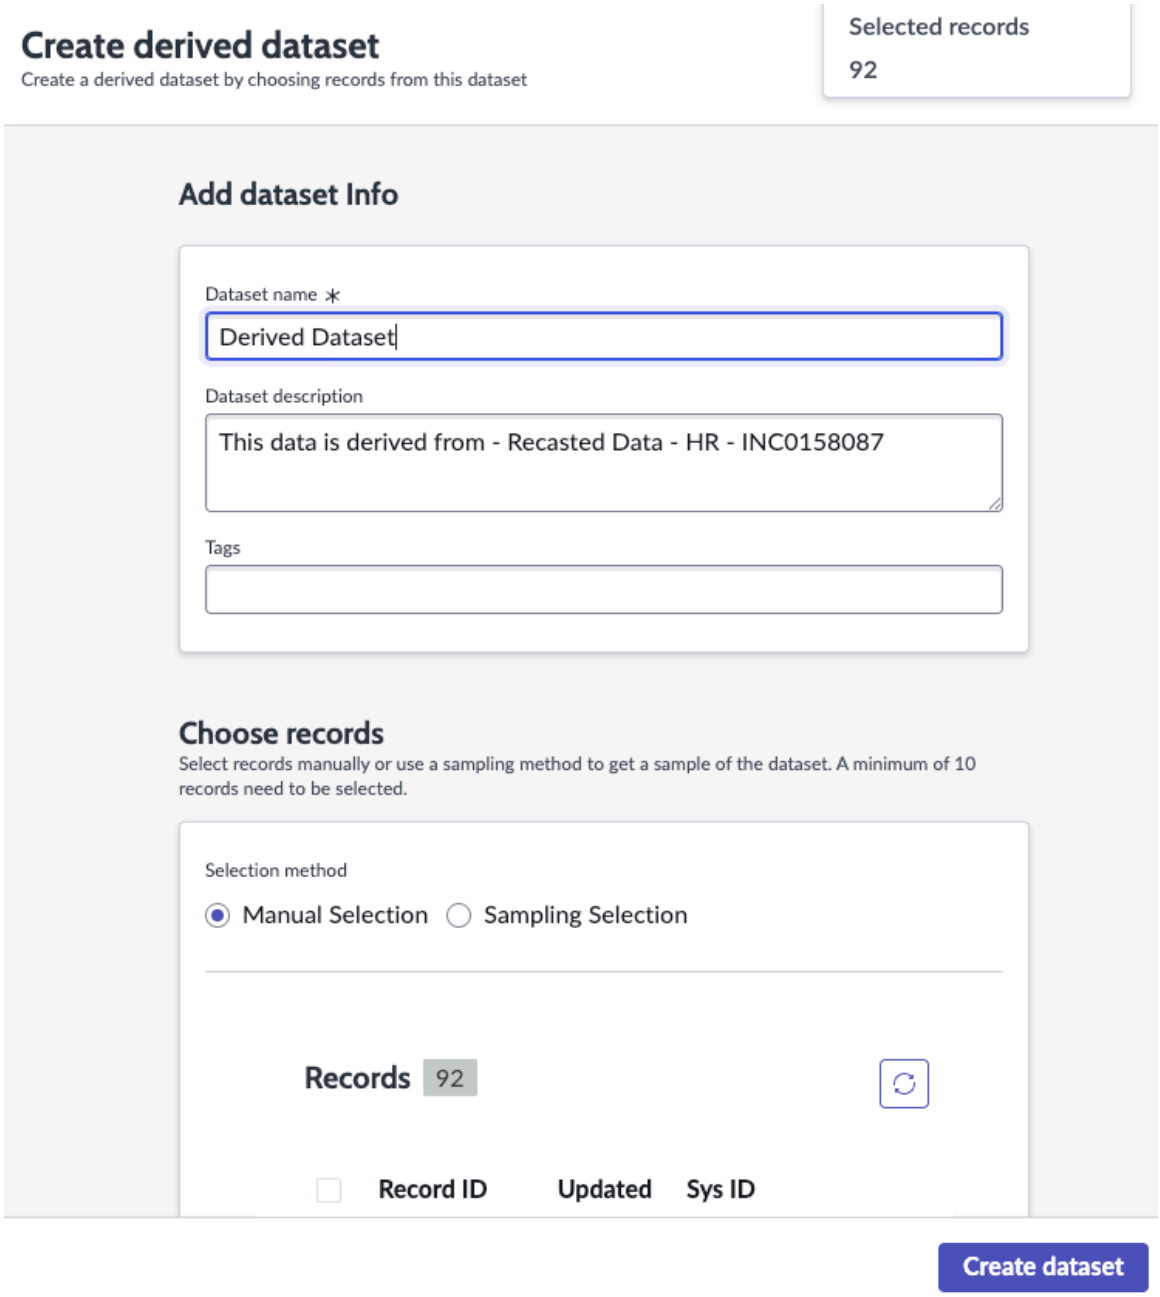
Create a smaller dataset from an existing dataset through generative AI by using the Now Assist Data Kit application. A derived dataset enables you to choose a smaller dataset.

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

Before you begin
Role required: sn_data_kit.admin

About this task
After you add a dataset to the data catalog, you can create a smaller dataset or derived dataset from the Create derived dataset page.

Procedure
1. Select the dataset where you want to create a derived dataset.
2. Select **Create derived dataset**.
3. Add information for the new derived dataset.



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

New derived dataset info form

Field	Description
Dataset name	Name of the dataset.
Dataset description	Description of the dataset.
Tags	Tags to identify the dataset.

4. Navigate to the Choose records section and select a data selection method for your user.
   You can select records manually or use a sampling method to choose the data. If you select a sampling selection, use the drop-down menu and then enter a number in the sample size box.
5. Select **Run**.
6. Preview the record results and select **Create dataset**.
   The derived dataset is added to the data catalog as a separate dataset.

What to do next
After the derived dataset is completed, you can add the ground truth.

Generate synthetic data

Create synthetic data with a sample dataset and a prompt through generative AI by using the Now Assist Data Kit application. You can use synthetic data to create training for a test model or an evaluation dataset.

Before you begin
Role required: sn_data_kit.admin

Procedure
1. Navigate to **All** > **Now Assist Data Kit** > **Home**.
2. Select the **Synthetic datasets** tab.
3. Select **Generate dataset**.
4. Choose how you want to generate data.

Way to generate synthetic data

Method	Description
Use a default generator	The Default Data Generator creates flat records based on a fixed schema with files like name, description, and short description. By providing a few examples, the model learns the structure and generates consistent, realistic data, making it suitable for typical use cases. These configurations can also be saved as templates for easy reuse.

5. On the Add job info page, fill in the fields.

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

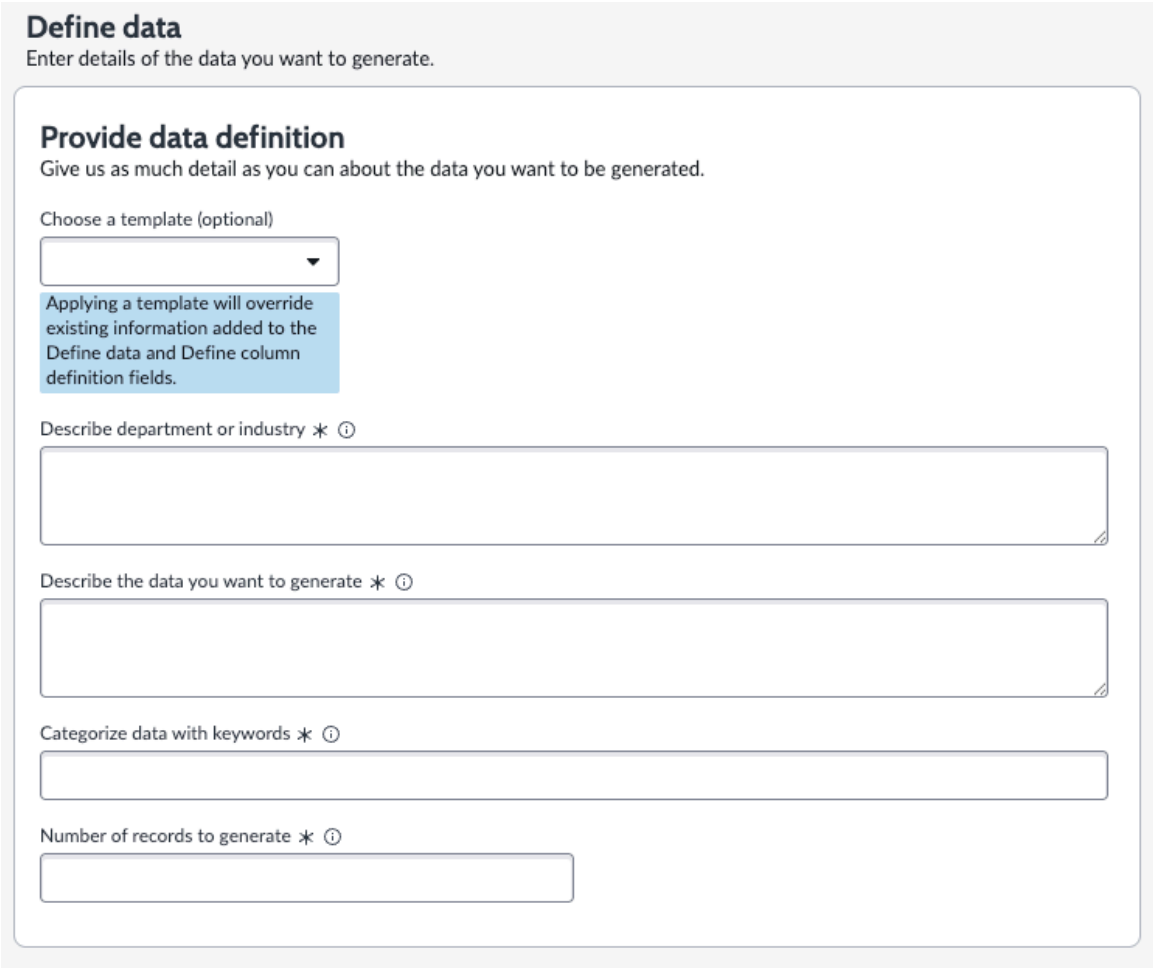
Add job info form

Field	Description
Job name	Name of the generated data.

6. Navigate to **Define data** .
7. **Optional:** Select a template.

Synthetic data templates

Template	Description
Catalog item	Service Catalog is a user-friendly interface that allows end-users to browse, request, and manage services and products offered by the organization, streamlining self-service and improving operational efficiency.
Incident data	Information Technology Service Management is a business function that involves managing IT services and processes to meet business needs effectively.



8. On the form, fill in the rest of the fields.

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

Define data form

Field	Description
Name	Department or industry where the data belongs.
Type	Data that you want to generate.
Category	Data that has been categorized with keywords.
Count	Number of records to generate.

Note: If you have sample data available, you can navigate to Select the sample data to enhance the accuracy of the generated data. If you don't have sample data, refer to the in-product help for guidance.

9. Select **Continue**.
10. Add your sample data by navigating to Define columns to generate data.
    The columns are populated when you add the sample data. If no sample data is available, you must manually populate the columns with data.
11. Select **Continue**.
12. **Optional:** Add additional rules to improve generated data.
13. Select the number of records to generate for testing.
14. Select **Start generation**.

Select the sample data

Add the available sample data to enhance the accuracy and relevancy of the generated data in the Now Assist Data Kit application. The sample data can be curated, added into the data collection, or published.

Before you begin

Role required: sn_data_kit.admin

Procedure

1. Navigate to **Select sample data**.
2. In the Select table and columns pop-up window, select an existing dataset in the Data collection field.
3. Navigate to **Continue** to select up to three records.
   The data is populated in the different columns.
4. Enter the descriptions for each column that you want to generate the data for.
   Describe the columns for the large language model (LLM) to generate values accordingly.
5. Set the key column toggle switch on one column.
   The key column is required and used as a baseline for creating the data in the other columns.
6. Navigate to **Test generation** to preview a small number of records before generating a full dataset.
   When previewing a small number of records, you can make adjustments before generating a full dataset.
7. Select the records to generate for testing.

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

8. Select **Start testing**.
   You see the preview generation data results page. If you want more refined results, select **Define data** or **Select sample data** to update the criteria.
9. If you're satisfied with the preview results, select **Start generation** for all the data.
   The generated data is completed.

Define columns to generate data

Provide detailed definitions to preview a test for each column that you want the data to generate the results for.

Before you begin
Role required: sn_data_kit.admin

About this task
For better quality data, use fewer than 20 columns.

Procedure
1. Enter a description for each column that you want to generate data for.
   Describe the columns for the large language model (LLM) to generate values accordingly. This description can include a range and any rules or limitations to be considered.
2. Enable one column as the key column.
   The key column is required and used as a baseline for creating data in other columns.
3. Select **Continue**.
   The test generation page opens.
4. Navigate to **Test generation** to preview a small number of records before generating a full dataset.
5. Select the records to generate for testing.
6. Select **Start testing**.
   You see the preview generation data results page. If you want more refined results, navigate to **Define data** or **Select sample data** to update the criteria.
7. Select **Start generation**.
   The generated data is completed.

Add a ground truth to each dataset record

Add a ground truth, which is the real-world data that is used to train and test AI models to each dataset record. You can do this task by using the Now Assist Data Kit application.

Before you begin
Role required: sn_data_kit.admin, sn_data_kit.analyst

Procedure
1. After you create a data set, select **Create ground truth guidelines**.
2. On the form, fill in the fields.

Add ground truth info form

Field	Description
Type	○ Manual

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

Field	Description
Ground truth type	○ Summarization ○ Content generation ○ Other
Instructions	Guidelines for labelers and linguists who manually add the ground truth.
Column type	If you chose AI generated ground truth, select the column type. ○ String ○ Number ○ JSON
Column label	The name of the column that is added for the ground truth in the dataset.

3. Select **Confirm**.
   A new Record detail page opens.
4. On the form, fill in the fields.

Record detail info form

Field	Description
Form label	Incident record number.
State	Time of the record creation.
Summarization	Manual entry of the ground truth.

You can manually enter the ground truth and rate the ground truth. These options appear in a separate column in the dataset record.

5. Select **Save and next**
   The ground truth column is added in a separate column in the dataset. Use the Add to data collection pop-up window to combine similar records from different datasets.
6. Add the dataset to an existing collection by selecting **Add to data collection**.
   a. Enter the data collection name, description, data collection category, and relevant tags.
   b. To create a new data collection, slide the toggle bar.
7. Select **Next**.
   A new data collection page populates.
8. Select the Choose columns form and then select the columns that you want to add to the data collection.
9. Select the Choose records form and select how you want to sample the dataset again (manually or sampling method).
   For a random sample, select **Run** to preview the record. Choose a sampling method, and select **Run** to preview the record.
   The selected records are added to the data collection.
10. Select the data collection to preview the records.
11. Select**Publish** to make the data available for validation.
    When you publish a collection, the data freezes curation and makes the dataset available for use through ServiceNow SDK.

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

12. Select **Confirm** to make your collection available.
    The data collection is published and available in Mobile SDK for evaluation.

View data insights

You can view data insights to see the completeness and distribution of your generated data.

Before you begin

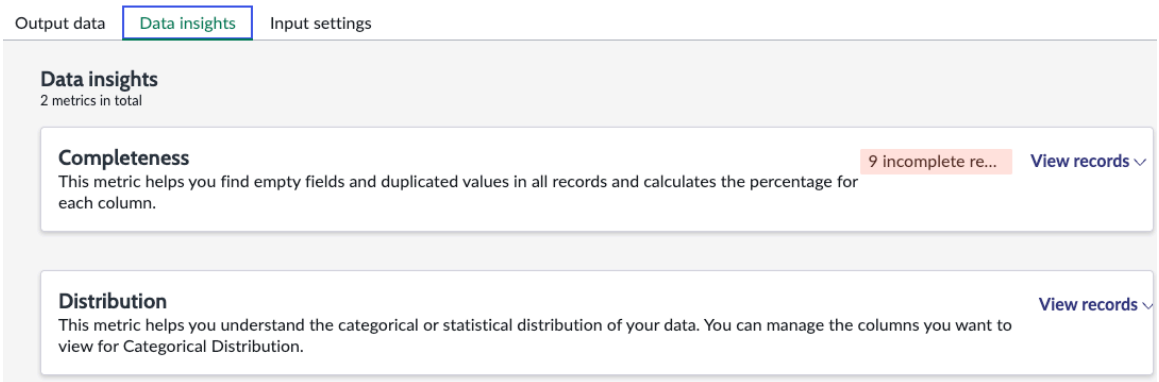
Role required: sn_data_kit.admin

Procedure

1. Navigate to **All** > **Now Assist Data Kit** > **Home**.
2. In the Generate Data section, select **Start now**.
3. Select the generated data job that you want to see data insights for.
4. Select the **Data insights** tab.
5. Select **View records** to see the records associated with each metric.

Data insight metrics

Metric	Description
Completeness	This metric helps you find empty fields and duplicated values in all records and calculates the percentage for each column.
Distribution	This metric helps you understand the categorical or statistical distribution of your data. You can manage the columns you want to view for Categorical Distribution.



Find and cleanse sensitive data

You can scan your datasets for sensitive data, like personally identifiable information (PII). If you find sensitive data, then you can cleanse it from your datasets.

Before you begin

Role required: admin

Plugins required: sn_data_discovery, sn_dp_store_app, com.glide.data_privacy.

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

Procedure

1. Navigate to **All** > **Now Assist Data Kit** > **Home**.
2. From the Discover datasets section, select **Get started**.
3. Select the dataset that you want to scan for sensitive data.
4. Select **Scan Data**.
5. If there is sensitive data in your dataset, select **Cleanse Data** to remove it.
   You can also see the records that contain sensitive data by viewing the Data Insights tab. To learn more, see View data insights.

Now Assist Data Kit reference

Reference topics for Now Assist Data Kit.

Now Assist Data Kit roles (sn_data_kit.analyst)

This user can edit and save ground truth in Now Assist Data Kit. This role has read only access within Now Assist Data Kit.

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.

Now Assist Data Kit roles (sn_data_kit.admin)

This user can create, update, and publish datasets in Now Assist Data Kit. This role is mandatory to use Now Assist Data Kit.

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.

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

Now Assist Skill Kit

Use ServiceNow® Now Assist Skill Kit to create and publish custom prompts and skills for Now Assist. Creating custom skills and prompts enables you to have greater flexibility with Now Assist's generative AI capabilities.

https://player.vimeo.com/video/985874233?h=2ccc18794b&badge=0&autopause=0&app_id=58479

Get started

Explore



Learn more about Now Assist Skill Kit and creating effective prompts.

Configure



Configure Now Assist Skill Kit prompt and deployment settings.

Use



Use Now Assist Skill Kit to create custom skills.

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.

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

Troubleshoot and get help

- Ask questions and explore other resources for Now Assist Skill Kit in the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

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 Skill Kit

The Now Assist Skill Kit plugin for Now Assist enables you to create and activate custom prompts and skills for Now Assist.

Now Assist Skill Kit overview

If the base system Now Assist skills don't fit your needs, you can use Now Assist Skill Kit to create custom skills. These custom skills enable you to have greater flexibility with Now Assist's generative AI capabilities.

Get Now Assist Skill Kit

To use Now Assist Skill Kit, you must update your Now Assist plugins in the Application Manager. For example, update your Now Assist for ITSM plugin to the Xanadu release.

You must also assign the sn_skill_builder.admin role to anyone who uses Now Assist Skill Kit.

Now Assist Skill Kit users

Now Assist Skill Kit Users

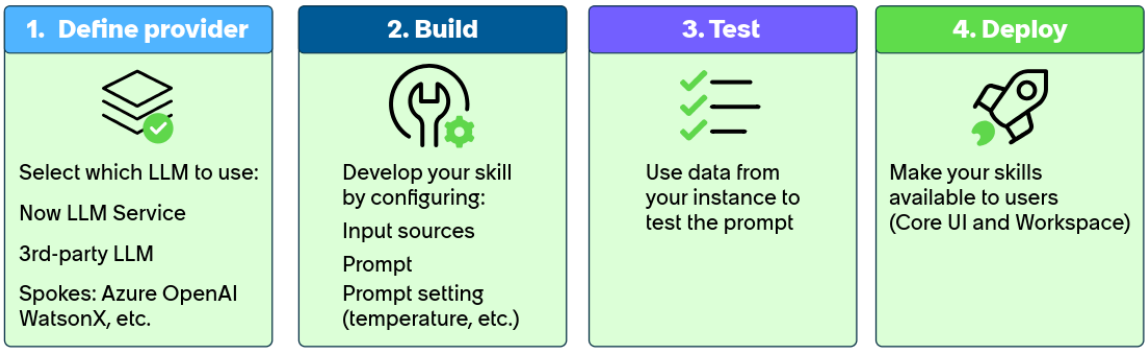
User	Description
AI developer	AI developers manage skill development in Now Assist Skill Kit. They create new skills, write the prompts, and configure the skill settings.
Now Assist admin	Now Assist admins activate published skills.

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

Now Assist Skill Kit workflow

The following diagram shows the user journey for Now Assist Skill Kit.

User journey for Now Assist Skill Kit



1. Define the provider

   You must understand the benefits and potential downsides of each large language model (LLM) that you are considering using.

2. Build the prompt

   You must have an understanding of the architecture of your Now Assist instance and be able to define where input data should come from. You should also have an understanding of LLM fundamentals to build an effective prompt.

3. Test the prompt

   Now Assist Skill Kit enables you to test your prompt from the editor.

4. Deploy the skill

   Now Assist Skill Kit enables you to deploy your skill to Now Assist panel, Now Assist Context Menu, Virtual Agent, Flow Action, or a UI Action.

Now Assist Skill Kit benefits

Now Assist Skill Kit enables you to design your own custom generative AI functionality that is then easily deployed into the ServiceNow platform. Custom skills can be used to augment workflows with generative AI, allowing for increased effectiveness and efficiency.

Now Assist Skill Kit benefits

Benefit	Feature	Users
Create custom solutions by building a custom skill or workflow.	Create a skill	AI developer
Create and edit prompts for skills and configure where you want to bring in data from to augment your prompt.	Create a prompt	AI developer
Test and iterate on your skill before activating it.	Test a prompt	AI developer

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

What to explore next

To learn more about configuring and using Now Assist Skill Kit, see:

- Configuring Now Assist Skill Kit
- Using Now Assist Skill Kit

General guidelines for Now Assist Skill Kit

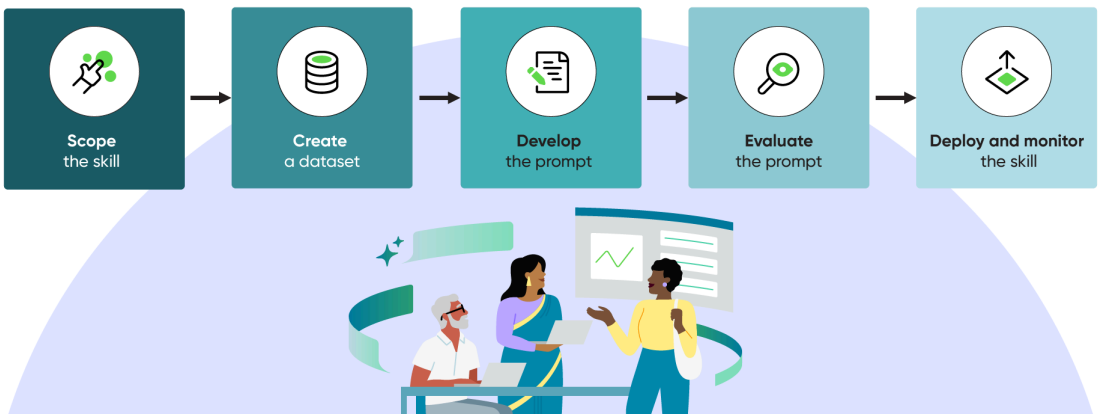
General guidelines are available to use Now Assist Skill Kit.

Now Assist Skill Kit guidelines overview

Think about the process of building a custom skill as having the following phases:

1. Scoping the skill
2. Collecting data and creating a dataset
3. Developing the prompt
4. Conducting a performance evaluation
5. Deploying and monitoring the skill

The stages of building a custom skill



You should adopt a data-driven mindset while you build a skill. A data-driven mindset means emphasizing evidence-based decisions and minimizing speculation about how a large language model (LLM) can interpret the specific wording of a prompt.

Now Assist Skill Kit needs intermediate skills in prompt engineering. If you're not a developer, you should equip yourself with the relevant knowledge before working with this tool.

There's information and helpful tooltips in many places throughout the product. Use this information to understand how to use the tool more effectively.

To ensure that there's a proper segregation of duties, a developer or generative AI practitioner creates the skill and then publishes it. After the skill is published, a Now Assist Admin will activate the skill, so that it shows up at the configured touch points and is available to use.

Scoping the skill

Scoping the skill before you build it helps determine the requirements needed for the skill and the expected outcomes.

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

Before you begin

Before you start to build a custom skill with Now Assist Skill Kit, you should first review the base system skills that are available in the Now Assist Admin console. Use these pre-existing skills whenever possible.

User prerequisites

Now Assist Skill Kit is meant to be used by a developer or someone with experience using generative AI. Because you must write the initial prompt template, you should:

1. Be knowledgeable about prompt engineering, including:
   - Having familiarity with the development and testing of machine learning systems.
   - Having informed expectations about the behavior and capabilities of a prompted large language model (LLM). You should understand:
     - The fundamentally probabilistic nature of LLMs.
     - That performance can vary greatly across different LLMs and different target tasks.
   - Having experience or training with writing and evaluating prompts for LLMs.
2. Have an in-depth understanding of the use case that they're trying to solve and the persona that they're building the skill for.

Design

Before you begin building a custom skill, you must think about the overall design and document the requirements. Consider the following:

1. What precisely do you want the skill to do?
   - What will be the inputs to the model?
   - What should be the outputs of the model? (Both content and format)
   - Who will be using the skill?
   - How can you characterize success?
2. What don't you want the model to do?
   - It's useful to think about the possible risks and downsides of using generative AI outputs so that you can try to guard against them during development and testing.
   - It's worth listing any specific model behaviors (LLM outputs) that would be detrimental or harmful in your specific use case.

Creating a dataset using Now Assist Skill Kit

Use these guidelines to create an effective dataset. Having an effective dataset provides better results for your prompt.

Now Assist Skill Kit dataset creation overview

A data-driven approach to skill development relies on the collection of a high-quality dataset to develop and test the skill. When you use Now Assist Skill Kit, you can also leverage the existing capabilities of the ServiceNow AI Platform to create a high-quality dataset.

When collecting data for this purpose, you should aim to create datasets that are:

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

1. Representative of the skill's intended deployment environment. The data should:
   - Seek to reflect the expected distribution of inputs in the deployment environment.
   - Capture variance along several identified axes, for example, input length, urgency.
   - Include any examples of inputs that are known to be important to the use case.
   - Consider edge cases (which may be rare) but that are suspected to cause problems, for example, long examples.
2. Sized appropriately for the team's risk appetite.
   - It's possible to develop and deploy a skill with little data. However, a lack of data creates more uncertainty about how the skill performs in deployment.
   - You should think like statisticians and produce confidence intervals for any associated performance scores and prompt comparisons.
3. Isolated from the data used for developing and writing the prompts.
   - You should split the data collected into development and testing sets. By splitting the data, you are protecting some data solely for evaluation purposes.
   - If you use all the data during the process of developing the prompt, your final evaluation of the skill is biased, meaning that it over-reports performance. This bias is because of a phenomenon known as prompt overfitting.

Developing the prompt

Use the guidelines to help create a prompt for your skill. A specific, clear, contextual prompt provides better results.

Prompt development overview

As a prompt engineer, you should make development decisions by looking at the model outputs that are generated in response to a prompt applied to many different inputs. However, there are still certain guidelines that may help users get started with prompt design.

1. Be specific

   Define your desired outcome clearly. Be specific about the task that you want the model to fulfill. Clearly identify the inputs that you're providing to the model, and specify the output you're expecting from the model (including formatting).

2. Include the right context

   Provide background information and context relevant for fulfilling the task. This information can generate a more focused response.

3. Use clear language

   Use precise and unambiguous language while writing the prompt.

4. Include demonstrations

   If possible, experiment with providing completed examples, or demonstrations, in the prompt after the instructions to illustrate what you want the model to produce. Demonstrations are a powerful way to increase the likelihood of generating a desirable output. However, the performance changes depending on the demonstrations selected.

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

5. Start simple and test variations

   Break down complex tasks into smaller and clearer instructions. Have a controlled and iterative approach. Experiment with different structures.

Other considerations

- Subtle differences in wording can lead to substantial differences in performance. Trying to reason about how a large language model (LLM) may “interpret” the instructions in a prompt only gets you so far. Which specific choice of prompt-wording works best depends on the underlying model and should ideally be chosen based on evidence (that is, looking at lots of outputs).
- In data-constrained settings, you should iteratively develop several candidate prompts using the development data, then measure the performance of each candidate prompt on the test set, choosing the best one.

Evaluating the prompt

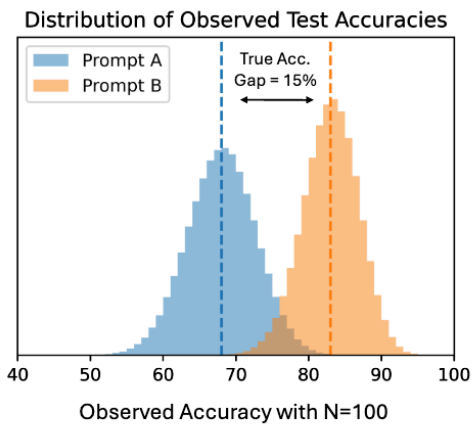
Evaluating the prompt is an ongoing process that occurs during and after prompt development and completion.

Prompt evaluation overview

To determine the effectiveness of your prompt, you should evaluate batches of test data. You should copy the model-generated responses and perform evaluations outside of Now Assist Skill Kit.

During prompt development

Ongoing, improvised evaluation should take place alongside the development of the prompt. This ongoing evaluation enables you to adapt the prompt based on observed model outputs. It may be tempting to test a change to a prompt against just one or two examples, however, to avoid reacting to noise, you should look at larger batches, and consider the statistical significance of the performance differences that you observed.



Chance of Seeing Better Performance from the Worse of Two Prompts

Data used in evaluation	"True" Underlying Accuracy Gap: 5%	10%	15%	20%
N=10	30%	21%	14%	9%
N=25	28%	16%	8%	3%
N=50	24%	10%	3%	0.6%
N=100	18%	4%	0.5%	0.02%
N=250	9%	0.4%	<0.01%	<0.001%
N=500	3%	0.01%	<0.001%	<0.0001%

The approximate probability of observing better performance from the less accurate prompt among two (2) candidates, tabulated as a function of the evaluation dataset size (N) and the "true" underlying accuracy gap between the prompts (as measured on "infinite" data).

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

Final performance evaluation

Before you deploy a skill, you should test the prompt on a representative batch of data that was isolated from the development process, that is, "test" data. You want to use isolated test data because of a phenomenon known as prompt overfitting. Iteratively editing a prompt based on the model outputs generated on the same data that is used for testing can lead to significant over-estimates of performance. This result is because the prompt can become overspecialized to the specific examples used in development. Even though the effect is typically less dramatic than what occurs when fitting machine learn model parameters to a test dataset, it's rooted in the same underlying principles, and should be avoided.

Evaluation metrics

Selecting the right metrics for evaluation is an important consideration. The following list provides a few approaches, each of which may be more or less appropriate depending on the use case.

- Classification-based assessment of short generations

  This approach requires labeled records, and it works best when the labels are short, well-defined "right answers," for example, true or false, multiple-choice, or category selection. In these cases, the model outputs can usually be parsed and formatted, then metrics like precision, recall, F1 scores, and so on can be directly calculated.

- Assessment of longer generations

  Many of the most interesting generative AI use cases require longer model generations, and there are many possible "right answers." In these cases, the output can be scored (by human evaluators) along several different axes, for example:

  - Faithfulness

    Is the generated text faithful to the context provided in the skill prompt? (The opposite of faithfulness is hallucination, which is to say that the model injects out-of-context information.)

  - Correctness

    Is the generated text correct relative to the skill instruction?

  - Helpfulness

    Is the generated text helpful relative to the task that the skill wants to accomplish? (Helpfulness is subjective but it's important to try to measure. Doing so properly requires a solid understanding of the needs of the people who will ultimately be using the skill.)

  - Fluency

    Is the generated text grammatically correct? Does it have any typos, issues with coherency, and so on?

Note: It's useful to score these properties on a scale, like 1-5, rather than with yes or no.

Deploying and monitoring the skill

After you evaluate the prompt, you can deploy the custom skill and monitor its effectiveness.

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

Now Assist Skill Kit deployment and monitoring overview

You should build and evaluate the skill in a subproduction instance. After you're satisfied with the final performance evaluation on a batch of test data in subproduction, you can begin to move a skill toward deployment.

1. Identify the touch points where the skills can be surfaced to be available for the end user. Configure the deployment settings accordingly.
2. Perform end-to-end user testing in subproduction for a while and with different users and inputs.
3. After the results are acceptable, deploy the skill to the production instance.
4. Monitor the skill.

Example use case for Now Assist Skill Kit

As an AI developer, you can create custom skills with Now Assist Skill Kit. For this example, create a custom skill for child incident summarization.

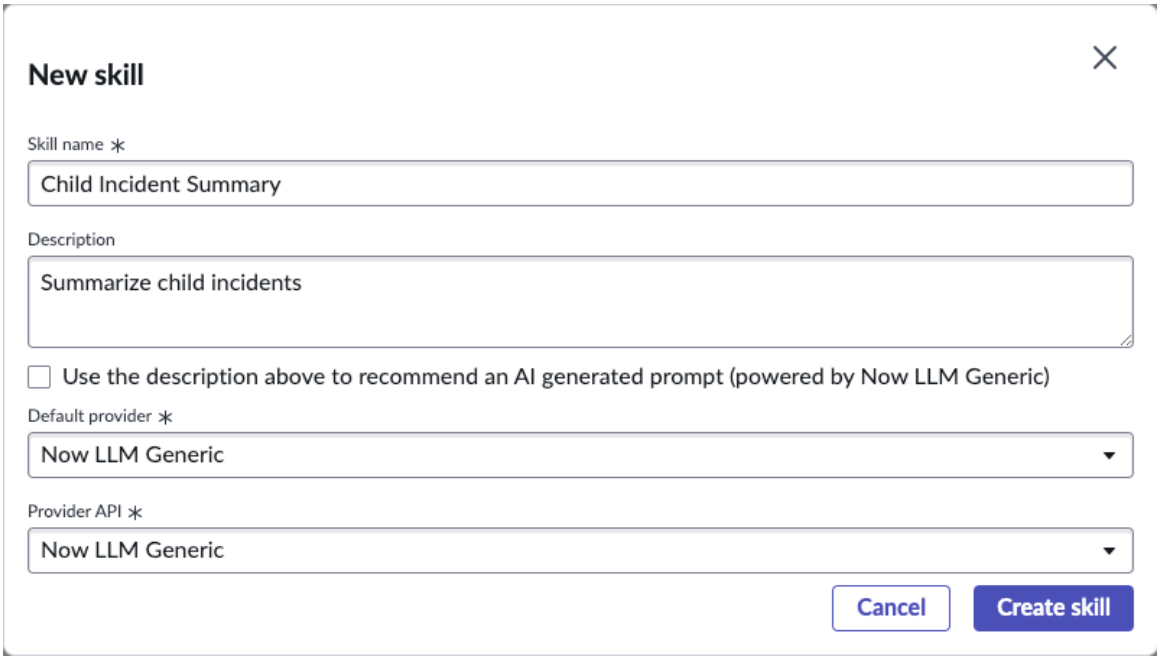
As an AI developer, you must create a skill for child incident summarization. Creating this skill can help you organize and understand multiple incidents that are related to the same parent.

Create and configure the skill

To create the skill:

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Select **Create new skill**.
3. On the form, fill in the fields. The following information is used to fill in the form for this example:
   - Skill name: `Child incident summarization`
   - Description: `Summarization of child incidents`
   - Default provider: `Now LLM Generic`
   - Provider API: `Now LLM Generic`

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



4. Select **Create skill**.

After you create the skill, you must configure the skill settings.

1. Select **Configurations**.
2. Select the model that you want to use. For this example, you can select `llm_generic`.
3. Select a temperature between 0-1 to determine the randomness and creativity of the output, such as `0.2`.

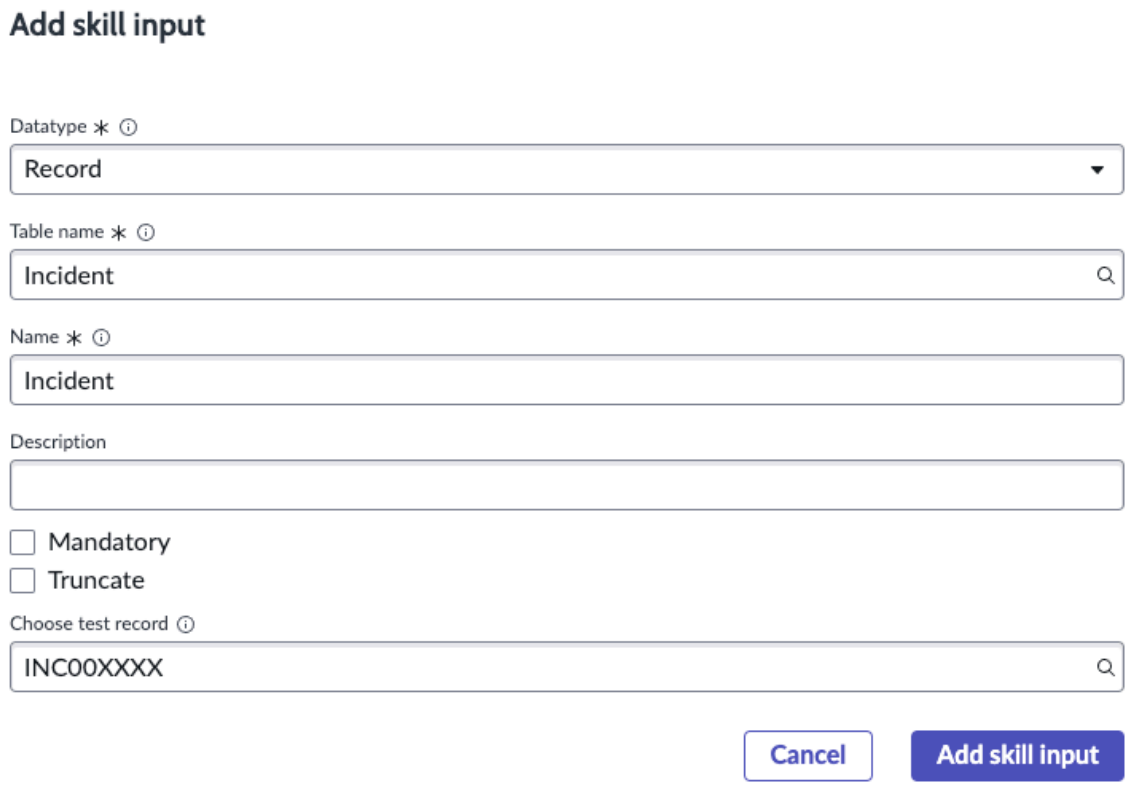
Develop the prompt

After you create the skill and configure the settings, you must develop the prompt. To follow this example, you can use the following prompt:

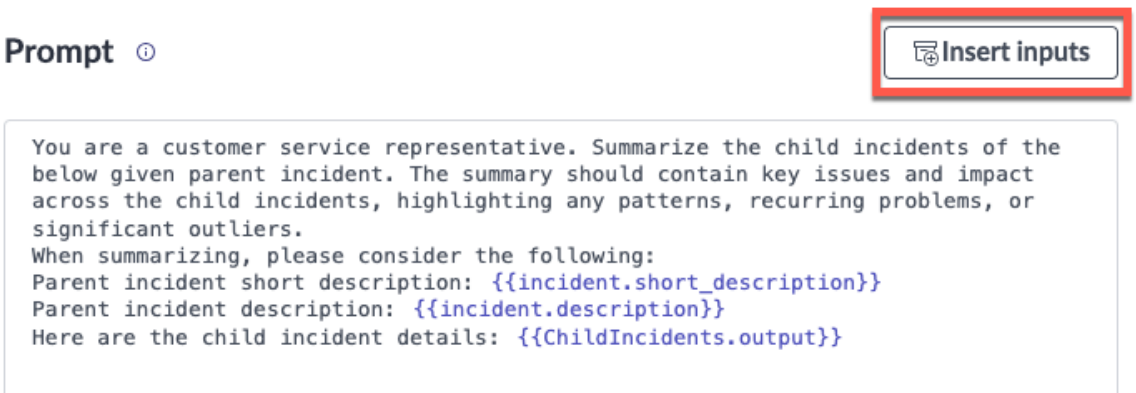
```
You are a customer service representative. Summarize the child
incidents of the below given parent incident. The summary
should contain key issues and impact across the child incidents,
highlighting any patterns, recurring problems, or significant
outliers. When summarizing, please consider the following: Parent
incident short description: {{incident.short_description}} Parent
incident description: {{incident.description}} Here are the child
incident details: {{ChildIncidents.output}}
```

Select the **Inputs** for the skill. For this example, the record is selected as the input.

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



Select insert inputs for the prompt.

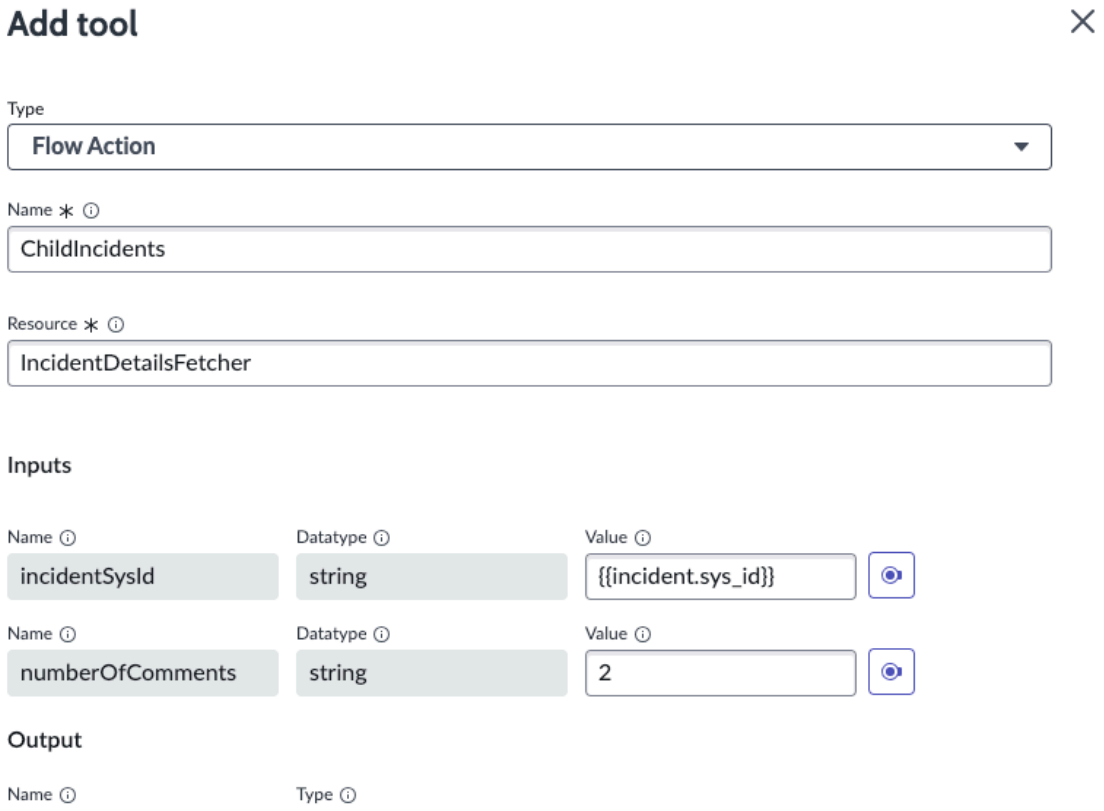


For this example, the following is used:

```
Parent incident short description: {{incident.short_description}}
Parent incident description: {{incident.description}} Here are
the child incident details: {{ChildIncidents.output}}
```

Select the tools for the skill. For this example, select Flow Action and the IncidentDetailsFetcher flow resource.

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



Configure the skill deployment options

The next step is to configure the skill deployment options. These options enable you to choose where to find the skill in Now Assist Admin.

1. Select the **Skill settings** tab.
2. Select **Deployment Settings**.
3. Enable the admin to enable the skill from the Core UI by selecting the **UI Action** check box.
4. Select **Save**.

Test and publish the skill

It is important to test your skill prompt to ensure that the correct type of data is being pulled in.

1. Select **Run tests**.
2. Choose a record or incident.
3. Select **Run test**.

Look at the response.

To see the data that was brought into the prompt from your skill inputs and tools, you can look at the grounded prompt tab.

If everything looks good, select **Finalize prompt**. After you finalize the prompt and you are ready to implement it, select **Publish**.

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

Activate the skill

After you test, finalize, and publish your skill, an admin must activate it in Now Assist Admin. To learn more about activating skills, see Activate a skill.

Configuring Now Assist Skill Kit

Configure prompts and skills for Now Assist Skill Kit.

Configuration overview

To use Now Assist Skill Kit, you must update your Now Assist plugins in the Application Manager. For example, update your Now Assist for ITSM plugin to the Xanadu release.

After you install the plugin, there are two parts to configuring a skill in Now Assist Skill Kit. First, you must configure how to deploy the skill. Next, you must configure the prompt.

Configure a skill prompt

Configure your skill prompt to set the model that is used and the randomness and creativity of the response.

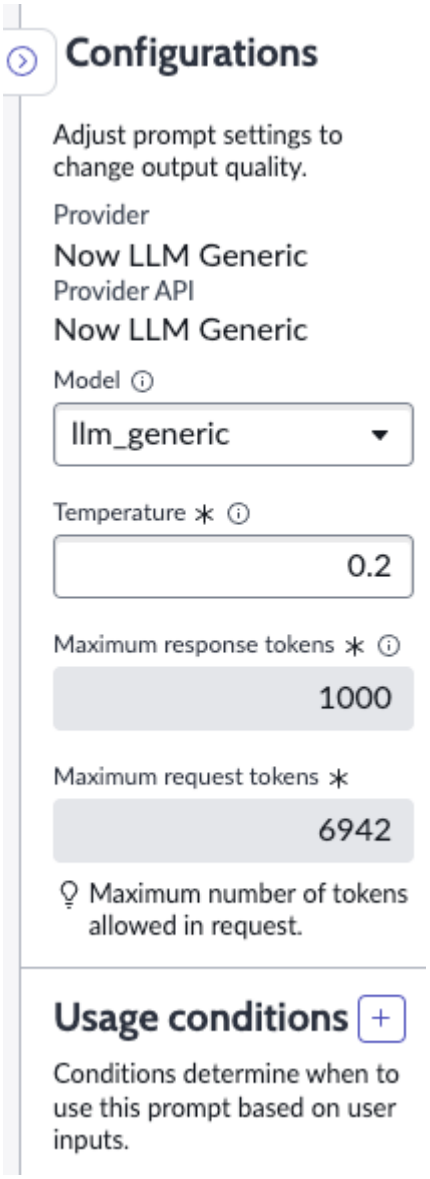
Before you begin

Role required: sn_skill_builder.admin

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Select the skill that you want to configure.
3. Select **Configurations**.

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



4. On the form, fill in the fields.

Configurations form

Field	Description
Model	The model is the large language model (LLM) that you want to use for the prompt.
Temperature	The temperature determines the randomness and creativity of the output. A higher value increases the randomness. The value must be between 0-1.
Thinking mode	Thinking mode controls how much reasoning effort a large language model applies when it generates a response. Higher levels improve output quality but increase latency. You can select thinking mode when you use *llm_generic_small_v2* or *llm_generic_large_v2*.

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

Field	Description
Maximum response tokens	The maximum number of tokens the model can return. If you're using Now LLM Service, the maximum is 1000.
Maximum request tokens	The maximum number of tokens allowed in a request.
Structured output	Structured outputs return prompt responses in a consistent JSON format. **Note:** Only Google Gemini and Azure OpenAI support structured output.

5. Add **Usage conditions** to determine when to use the prompt.

What to do next

After you configure your skill settings, you can test your skill. To learn more about testing skills, see Test a prompt.

To learn more about configuring the skill, including models and tokens, see Now Assist Skill Kit FAQs on the ServiceNow Community.

Configure skill deployment settings

Configure the deployment settings for the skill that you create. The deployment settings enable you to choose where the admin can find the skill in Now Assist Admin.

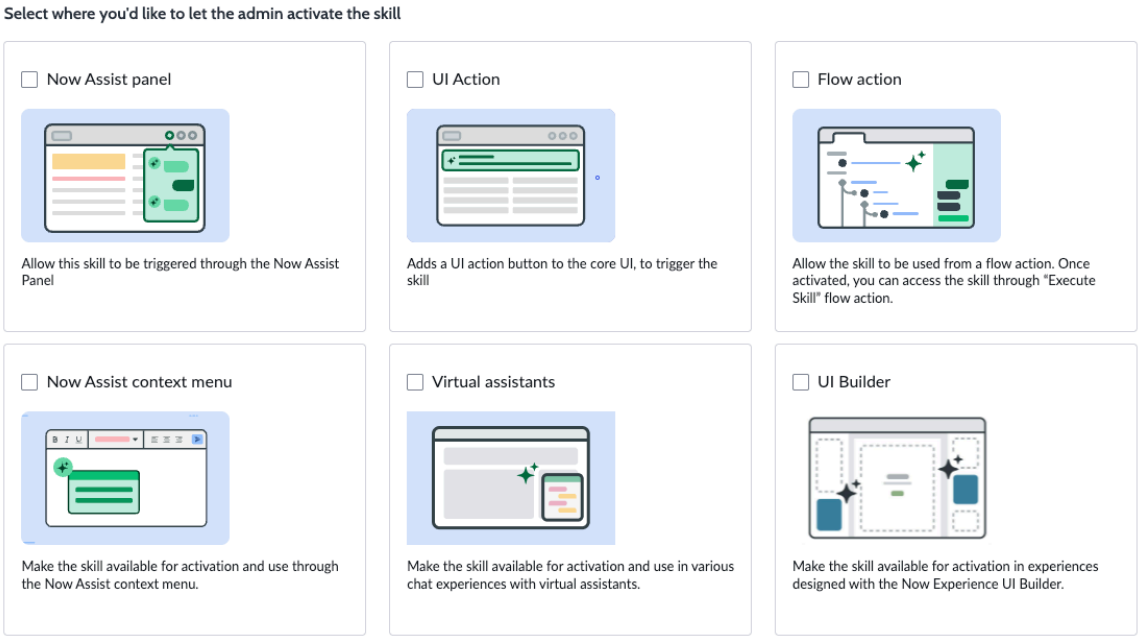
Before you begin

Role required: sn_skill_builder.admin

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Select the skill that you want to configure.
3. Select the **Skill settings** tab.
4. Select **General information**.
   This section shows the information that you added when you created the skill. You can edit the skill name and description here.
5. Select **Deployment Settings**.

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



6. On the form, fill in the fields.

Deployment settings form

Field	Description
Workflow	The high-level category that this skill pertains to, such as **Technology**, **Employee**, **Creator**, or **Platform**. You can also select **Other** if none of the categories fit. The workflow that you choose is where the skill appears in the Now Assist Admin console.
Product	The specific product that this skill operates within, such as ITSM, ITOM, HR Service Delivery, Now Assist Admin.
Feature	The feature that the skill is used on, such as Agent Chat, Knowledge, Virtual Agent. You can also define a custom feature if necessary.
Name	The name of the feature.
Description	A description of the feature.

7. Select where you want the admin to enable the skill from.
   - Now Assist Panel
   - UI Action
   - Flow action
   - Now Assist context menu
   - Virtual assistant

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

For more information about Now Assist in Virtual Agent, see Configuring assistants overview.

- UI Builder

8. Select **Save**.

What to do next

After you configure the skill settings, you can publish your skill. To learn more about publishing skills, see Finalize and publish a skill

Configure security controls for a skill

You must define an access control list (ACL) and role restrictions for all skills. An ACL enables you to restrict who is able to access and execute a skill to only users with the correct role. Role restrictions enable users to limit roles during skill execution.

About this task

If you have an existing skill that does not have an ACL, the execution of the skill is not disrupted. However, if you edit the skill and republish it, you must add an ACL.

If you try to execute a skill when you don't have permission, you see an error that you aren't authorized.

Before you begin

Role required: sn_skill_builder.admin

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
   A modal appears to explain ACLs. You can select **Got it** or **View skills without ACLs**.
2. Add an ACL and role restrictions to an existing skill.
   a. Select the skill that you want to add an ACL to.
   b. Select the **Deployment and skill settings** tab and then select **Security controls**.
   c. Select **Add ACL**.
   d. Select an option.

Options for ACL roles

Option	Description
Any authenticated user	As long as a user is logged in, they can access and execute the skill.
Select roles	Select the roles that a user must have to execute the skill. **Note:** If you select multiple roles, a user must only have one of the roles to execute the skill.

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

e. Add role restrictions.

f. Select **Apply**.

3. Add an ACL to a new skill.
   a. Create a skill.
   b. In the **Configure security controls** section, select an option for the access control list.
   c. Apply role restrictions to the skill.
   d. Continue creating the skill.

Using Now Assist Skill Kit

Use Now Assist Skill Kit to create and publish prompts and custom skills for Now Assist.

Create a skill

Create a custom skill for Now Assist. Creating a custom skill enables you to have greater flexibility with Now Assist's generative AI capabilities.

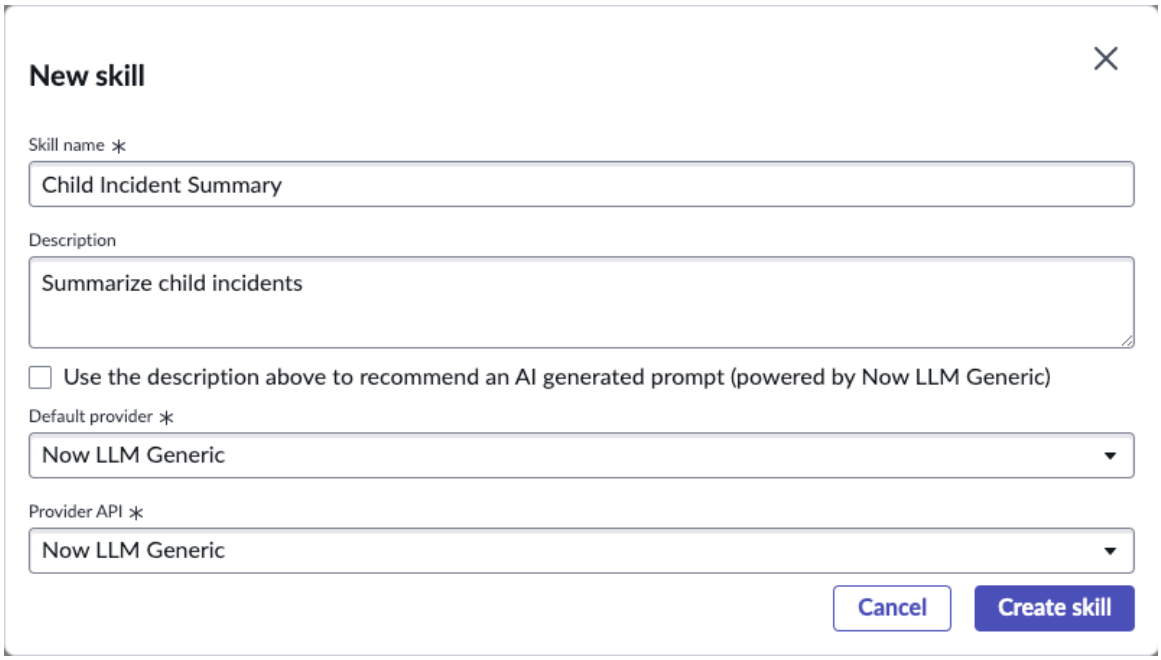
Before you begin

Role required: sn_skill_builder.admin

Note: When you update Now Assist Skill Kit you should also update Generative AI Controller to ensure that your skills continue to perform correctly.

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Select **Create new skill**.
3. On the form, fill in the fields.



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

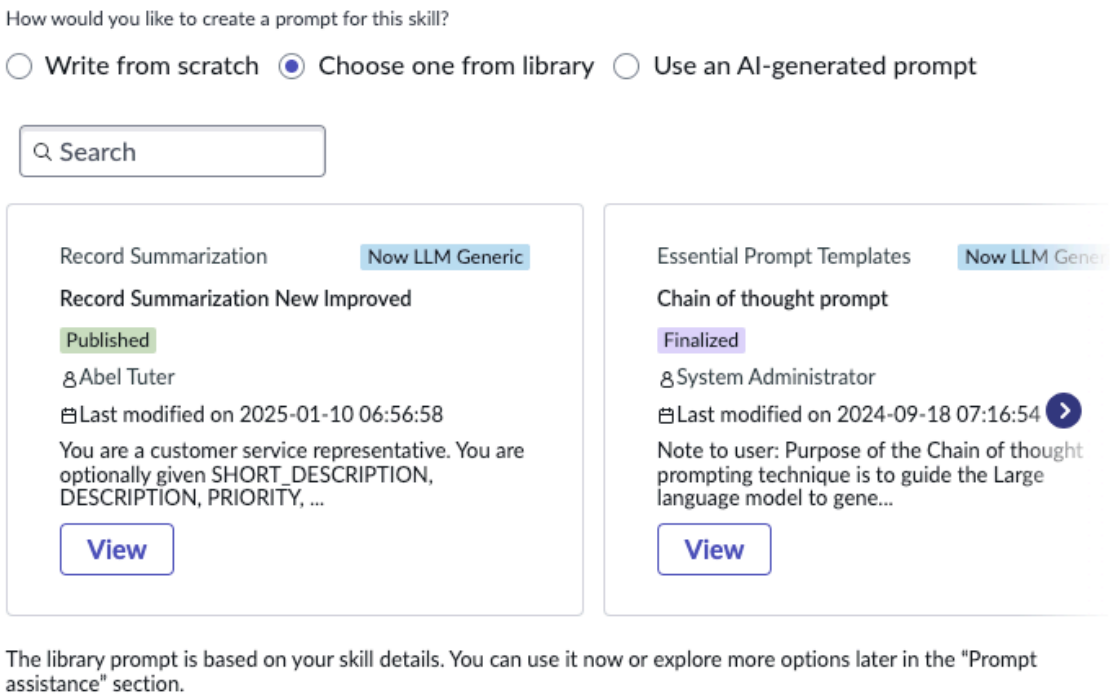
New skill form

Field	Description
Skill name	A name for the skill.
Description	A description of the skill.
Default provider	Available providers: ○ Now LLM Service ○ External LLM You can use Now LLM Service, Now LLM Long Term Stable models (LTS), Azure OpenAI, Google Gemini or Anthropic Claude on AWS as the AI model provider for all Now Assist skills and AI agents. Use the Configuration Controls in AI Control Tower to define which options are available, then set the skill-level preferences in the Now Assist Admin console. For more information, see Large language models on the ServiceNow AI Platform. ▪ Spokes ▪ Custom LLM For more information on setting up a custom large language model (LLM), see Configure a generic large language model (LLM) connector Available prebuilt spokes that enable you to connect with an external LLM: ○ Microsoft Azure OpenAI Generative AI Spoke ○ OpenAI Generative AI Spoke ○ Aleph Alpha ○ WatsonX ○ Google Gemini (MakerSuite and Vertex AI) **Note:** The spokes don't consume Integration Hub transactions. The spokes consume assists.
Provider API	The provider of the API for your chosen LLM.
User access	○ Any authenticated user As long as a user is logged in, they can access and execute the skill. ○ Select roles

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

Field	Description
	Select the roles that a user must have to execute the skill. **Note:** If you select multiple roles, a user must only have one of the roles to execute the skill.
Role restrictions	Role restrictions define the specific roles under which a skill in ServiceNow executes. While ACLs determine which user roles are permitted to trigger the skill, role restrictions determine the roles the skill will operate with during execution.

4. Select how you want to create the prompt for the skill.
   - Write from scratch
   - Choose one from library
     a. Find the prompt that you want to use.
     b. Select View.
     c. Select Use prompt.
   - Use an AI-generated prompt



5. Select **Next**.
6. **Optional:** Add skill inputs and outputs.
7. Select **Go to summary**.
8. Select **Finish**.

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

What to do next

After you create the skill, you must configure it. To learn more about configuring a skill, see Configure a skill prompt.

If you don't need to set any configurations for your skill, you can create your skill prompt and tools. To learn more, see Create a prompt and Add a tool.

Clone and edit a ServiceNow skill

Eligible skills provided in ServiceNow Now Assist applications can be cloned in Now Assist Skill Kit so that you can edit the prompt or change the AI service provider. Editing the prompt enables you to arrange the formatting and content of the LLM response.

Before you begin

Role required: sn_skill_builder.admin

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Select ***ServiceNow*** **skills**
3. Select the ***ServiceNow*** skill that you want to clone.
4. Select **Clone Skill**.
5. Fill in the fields on the form.

Clone skill form

Field	Description
Skill name	A name for the skill.
Description	A description of the skill.
Default provider	Available providers: ○ Now LLM Service ○ External LLM ▪ Spokes ▪ Custom LLM For more information on setting up a custom large language model (LLM), see Configure a generic large language model (LLM) connector Available prebuilt spokes that enable you to connect with an external LLM: ○ Microsoft Azure OpenAI Generative AI Spoke ○ OpenAI Generative AI Spoke ○ Aleph Alpha ○ WatsonX ○ Google Gemini (MakerSuite and Vertex AI)

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

Field	Description
	Note: The spokes don't consume Integration Hub transactions. The spokes consume assists.
Provider API	The provider of the API for your chosen LLM.

6. Select **Clone**.

   Note: When you clone a skill, you can't change the skill outputs, tools, or deployment settings.

7. Select the edit icon to add skill inputs.

Edit inputs form

Section	Description
Base input table fields	Each skill relies on a base input table and input fields with descriptions to provide context for the LLM to generate a response.
Rule conditions	Rule conditions determine when the input template is used. By default, record state determines which input template the LLM uses.
Additional input data sources	You can add input data sources like related tables, activity streams and relationships to provide more context to the LLM. You can also add rule conditions to these additional data sources.

8. **Optional:** Select **Clone prompt** to edit the prompt.
   You will see all of the prompts that use the same supporting skill for each provider.
9. Add **Prompt usage conditions**.
   These are conditions that determine when a prompt will be executed.
10. To test the prompt, select **Run test** and add test values.
11. Select the **Skill settings** tab.
12. **Optional:** Select **General information** to change the default provider by selecting the toggle.

What to do next

Create your skill prompt. To learn more about creating a prompt, see Create a prompt.

After you clone and edit the skill and prompt, you can evaluate your prompt.To learn more about evaluating a prompt, see Evaluate a prompt.

Create a prompt

After you create a custom skill, create a prompt. Creating a prompt enables you to choose what skill inputs to use, as well as the type of tool.

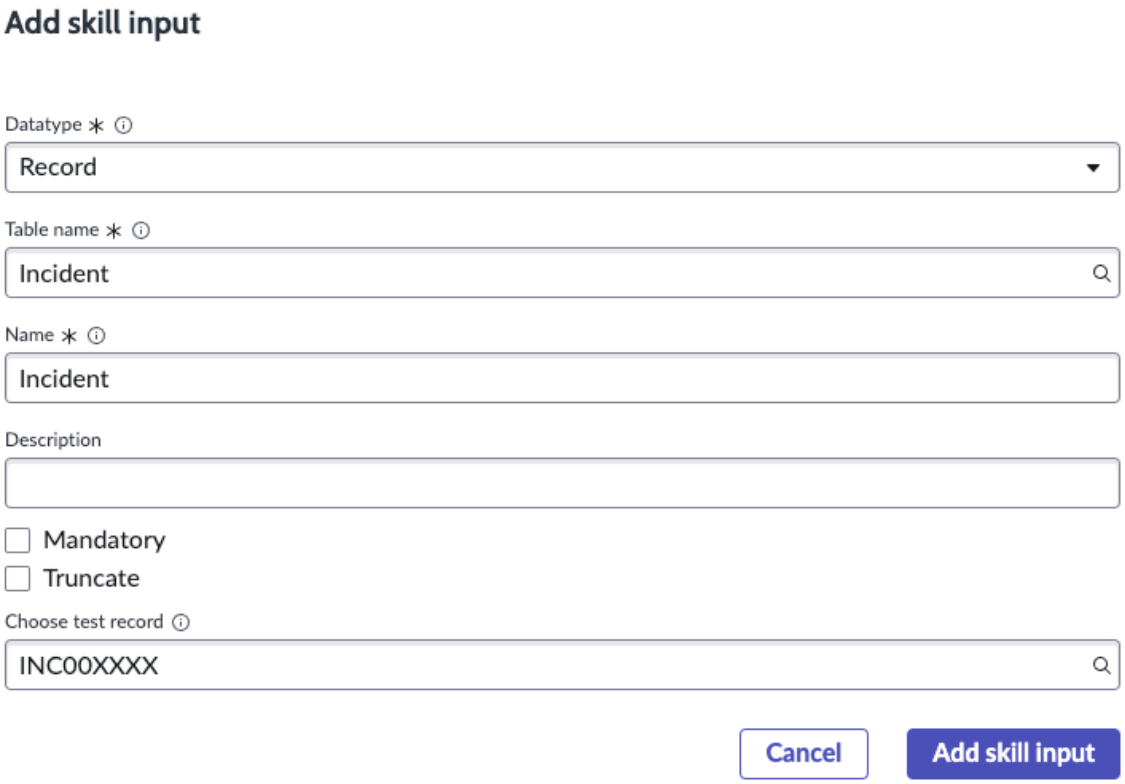
Before you begin

Role required: sn_skill_builder.admin

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

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Select the skill that you want to create a prompt for.
3. Select the edit icon ( ) and name the prompt.
4. Write the prompt.
5. Select **Skill Inputs**.



Skill input options

Input type	Description
Datatype	○ Record ○ String ○ Numeric ○ Boolean ○ Simple Array ○ JSON Object ○ JSON Array
Name	A name for the input.
Description	A description for the input.
Mandatory	The Mandatory option means that you must supply a value for the input when you run the skill.

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

Input type	Description
Truncate	The Truncate option means that, if your prompt is too large, the prompt context is shortened to fit the model context length.
For Records	
Table name	A name for the table.
Choose test record	The record that is used to test the prompt.
For Strings, Numeric. Boolean, Simple Array, JSON Object, JSON Array	
Test values	The values that are used to test the prompt.

6. Select **Add skill input**.
7. Select **Insert inputs**.
   The input options change depending on what kind of data type you choose.
8. Search for the inputs that you want to use for the prompt.
   For example, you can search for the incident short description or priority.
9. If you're not ready to finalize the prompt and publish the skill, select **Save** or **Save as**.

> **Note:** Skills can have multiple prompts. Usage conditions determine which prompt is executed. If no conditions are met, the default prompt is executed. To configure prompt usage conditions, see Configure a skill prompt.

What to do next
After you have created a prompt, you must test it. To learn more about testing your prompt, see Test a prompt.

Add a tool

Add and manage tools visually in the Tools editor, including decision branching, to execute different tools for your skill. Adding decision branches between tools enables you to define the conditions that need to be met for a tool to run. If no conditions are met, the default branch's step is executed.

Before you begin
Role required: sn_skill_builder.admin

About this task
A tool is a utility that is configured to convert skill inputs into skill outputs. The tool outputs can be referenced in a prompt template to introduce context to a prompt. The input for a tool can be a skill input or the output of another tool.

You can use the Tool editor to configure tools and link them to each other.

Decision nodes enable you to execute different tools, based on the logic of the branch. A decision node can contain multiple branches but will always need one default branch.

Procedure
1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Select the skill you want to add a tool to.
3. Select the **Tool editor** tab.
4. Select (+) icon to add a node.

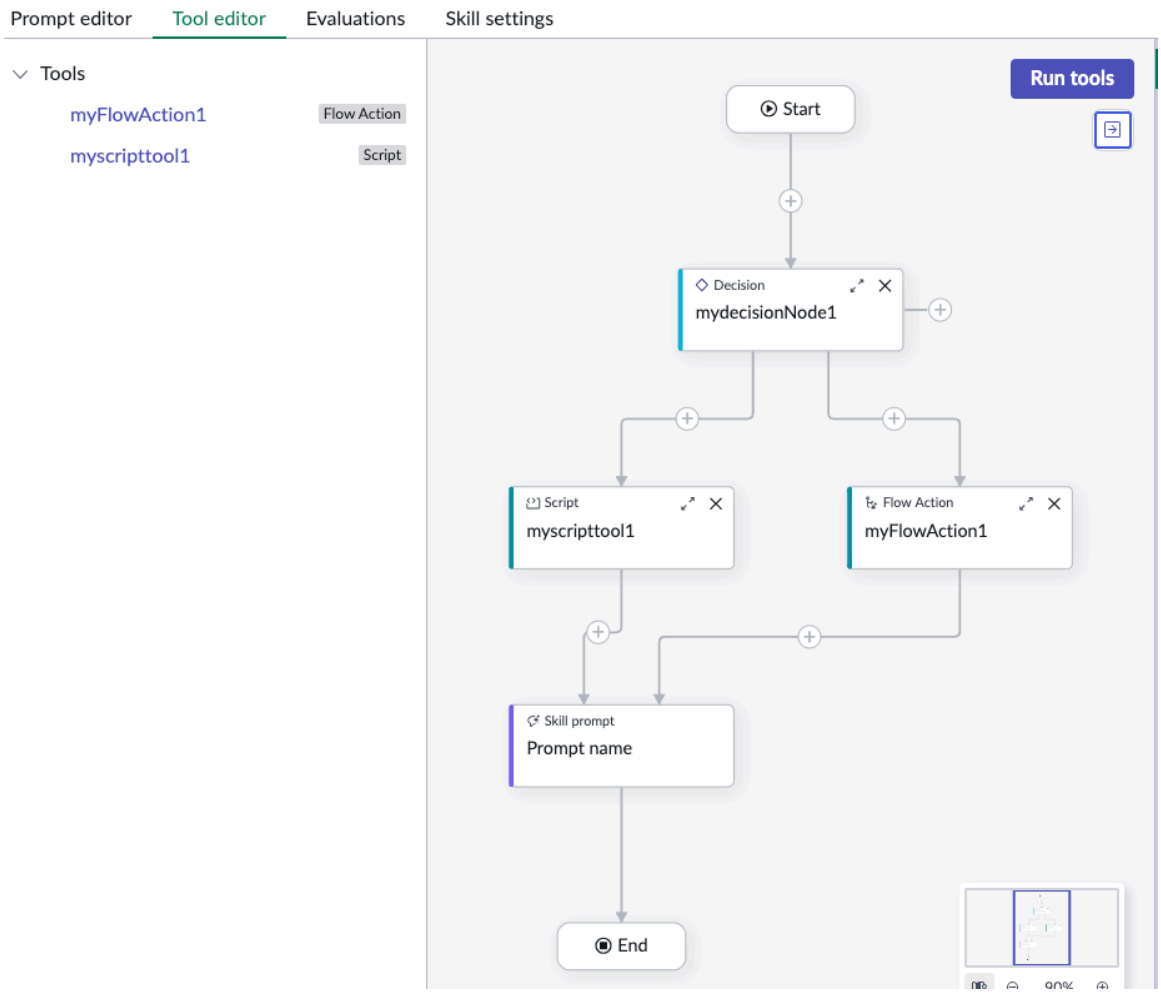
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

5. Select the type of node that you want to add.

Add a node

Node type	Steps
Tool node Types of tool: ○ Script ○ SubFlow ○ Flow Action ○ Retriever ○ Skill ○ Web search **Note:** If you select Google as your web search tool provider, the web search tool leverages Grounding with Google Search, offered under a Global Standard deployment. Because grounding is not data resident, Google's global infrastructure routes traffic to a global data center for each web search request. This processing may be different than your data processing location chosen for your ServiceNow instance. Please consider your organization's data policies before enabling skills that have Google web search tools. ○ Predictive Intelligence	**a.** Select a tool type. **b.** Name the tool. **c.** Select a condition.
Decision node	**a.** Name the node. **b.** Select **Add branch** **i.** Name the branch. **ii.** Select a Destination node. **iii.** Set the conditions. **c.** Select a Destination node.

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



6. Select **Add**.
7. Select **Run tools** to test the tools.

Add a retriever

Add a retriever to your prompt to augment and add context to your prompts with AI search results.

Before you begin

Role required: sn_skill_builder.admin

About this task

Using a retriever in a skill enhances the relevance and coherence of a response by pulling relevant information from a source and then feeding it to a large language model (LLM) to create the final output. This process leads to smoother conversational flow and better scalability as data grows, without requiring model fine-tuning.

A retriever enables the chatbot to access external knowledge by fetching relevant background information, resulting in more factual, in-depth, and informed responses.

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Create a new skill or select the skill that you want to add a retriever to.
3. Select the **Tool editor** tab.

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

4. Select (+) icon to add a node.
5. Select **Tool node**.
6. Select **Retriever** from the drop-down menu.
7. Select **Configure retriever**.
8. On the form, fill in the fields.

Add tool form

Field	Description
Name	The name for the retriever.
Search query	The information that you want to search for. It can be static text or a skill input.
Search space type	○ Table-based ○ Search-profile-based
Search profile	A group of search sources.
Search sources	Tables in ServiceNow that have been indexed and can be used for search.
Fields returned	The fields that you want returned from the search sources and sent to the large language model (LLM).
Limit	The maximum number of results that are returned.
Search criteria	○ Hybrid ○ Semantic ○ Keyword **Note:** If you choose Hybrid or Semantic, you can make selections for chunking and reranking. To learn more about chunking and reranking, see Retriever chunking and reranking.

9. Select **Next**.
10. Select an embedding model.
    To learn about embedding models, see .
11. If you selected **Hybrid** or **Semantic** search criteria, select a semantic index.
    A semantic index enables you to search for data based on contextual meaning.
12. **Optional:** If you selected **Hybrid** or **Semantic** search criteria, select **Advanced** to change the document matching threshold.
    The document matching threshold is the threshold for semantic matching. The default value is 0. You can input any value from 0 to 1. You can get more precise results with a higher value.
13. Select **Next**.

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

Note: If you selected **Hybrid** or **Semantic** search criteria, see Retriever chunking and reranking to complete setting up your retriever.

14. Review the retriever tool information.
15. Select **Add tool**.

Retriever chunking and reranking

When you're building a skill prompt that uses a retriever you can use chunking and reranking to enhance the accuracy and relevance of your responses.

Before you begin

Role required: sn_skill_builder.admin

About this task

To set up the chunking and reranking options for a retriever, you must have a retriever tool added to your skill with the **Hybrid** or **Semantic** search criteria. The following steps come after the semantic configuration in Add a retriever.

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Select the skill that you're adding a retriever to.
3. Add a retriever.
4. On the form, fill in the fields.

Chunking and reranking form

Field	Description
Max number of chunks per document	The maximum number of chunks that you want returned per document. The default value is 10.
Chunking strategy	○ Fixed size Fixed size breaks down the full text into smaller passages. You can choose the passage size limit in number of words or number of sentences. ○ Small to big Small to big enables you to choose the top K best matched indexed chunks and expand them to include surrounding chunks. Then the expanded chunks are linked together into a large text and broken into smaller passages. Don't use the small to big chunking strategy if you're using full text or truncate index chunking configuration.

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

Field	Description
Chunking unit	○ Words ○ Sentences
Chunk size	The size of the chunks that are returned. The default value for word size is 750. The default value for sentence size is 40.
Expanded snippet size	The size of the expanded chunks that are included when you select the small to big chunking strategy.
Top K results	The number of chunks the reranker returns. If you leave this field empty, the default number of chunks that are returned will be the limit you entered.

5. Select **Next**.
6. Select the type of condition that you want to be evaluated when the tool is executed.
7. Select **Next**.
8. Review the selections that you made for the tool and select **Add tool**.

Add a web search

Add a web search as a tool in Now Assist Skill Kit. Adding a web search as a tool enables you to add search results to your prompt.

Before you begin

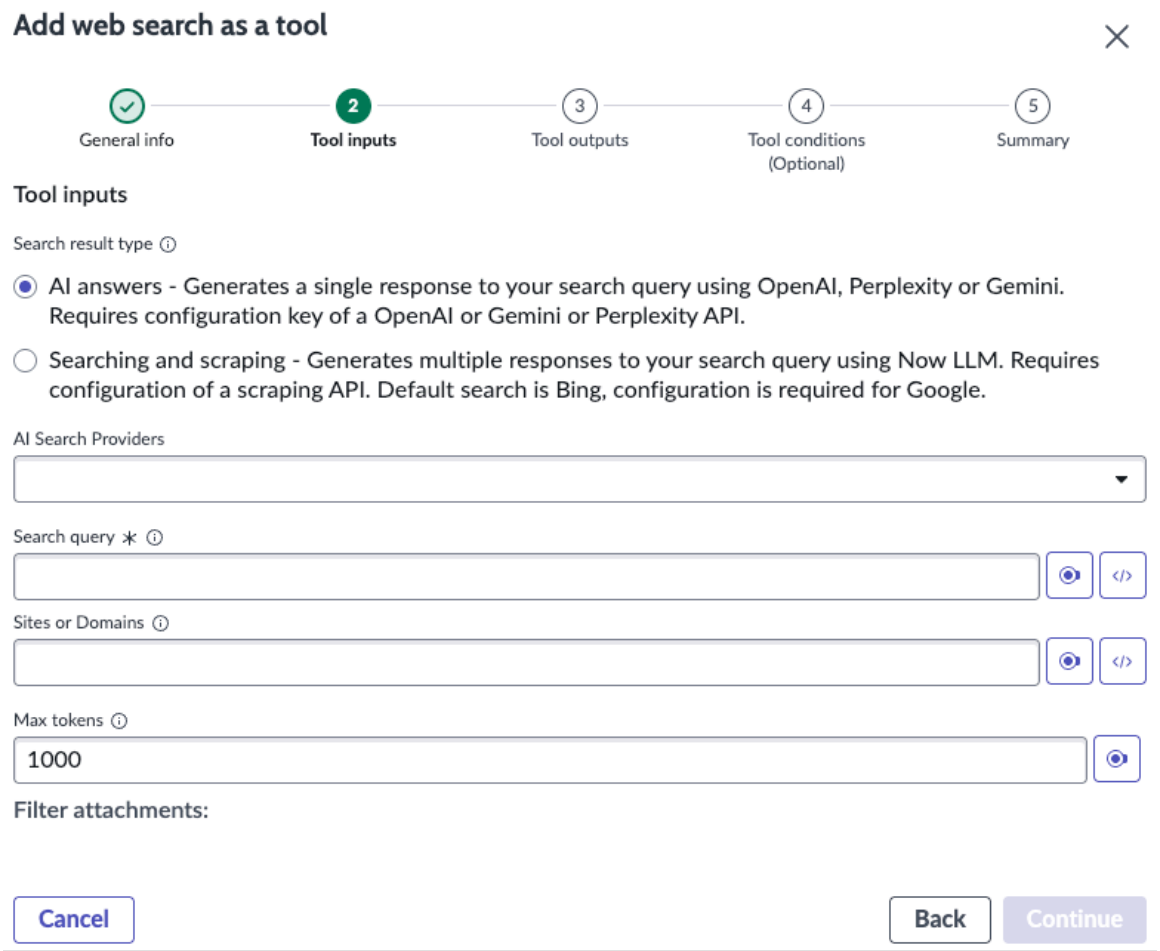
Role required: sn_skill_builder.admin or admin.

To use web search as a tool, you must bring your own search engine API key and configure it on the ServiceNow AI Platform.

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Create a skill or select the skill that you want to add web search to.
3. Select the **Tool editor** tab.
4. Select the (+) icon to add a node.
5. Select **Tool node**, and then **Add**.
6. Select **Web Search** in the **Tool type** field.
7. Select **Configure tool**.
8. On the General info form, enter a **Name** and select **Web Search** for the **Resource** field.
9. Select **Continue**.
10. On the Tool inputs form, fill in the fields.

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



Web search configuration form

Field	Description
Search result type	**AI answers** is the type of result that you want from the search. **AI answers** generates a single response to the search using either Perplexity, OpenAI, or Gemini.
AI search providers	The AI search provider used to perform the search.

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

Field	Description
	Note: If you select Perplexity or OpenAI, you must complete the API key setup in the AI Search answers OneExtend capability table. For more information on that process, see Configure AI search answers capability for web search. Although Azure Open AI is an option to select, Azure Open AI doesn't support web search. If you select Google as your web search tool provider, the web search tool leverages Grounding with Google Search, offered under a Global Standard deployment. Because grounding is not data resident, Google's global infrastructure routes traffic to a global data center for each web search request. This processing may be different than your data processing location chosen for your ServiceNow instance. Please consider your organization's data policies before adding a web search tool with Google as the provider.
Search query	The word, words, or phrase you're searching for.
Sites or domains	This field appears but it's not supported.
Max tokens	This field appears but it's not supported.

11. Select **Continue**.
12. On the Tool outputs form, select **Continue**.
13. On the Tool conditions (optional) form, select **Continue**.
14. On the Summary screen, select **Add tool**.

Use prompt assistance

Use prompt assistance to get a jump start with your prompt development by selecting an example from the prompt library or using Now Assist to generate one.

Before you begin

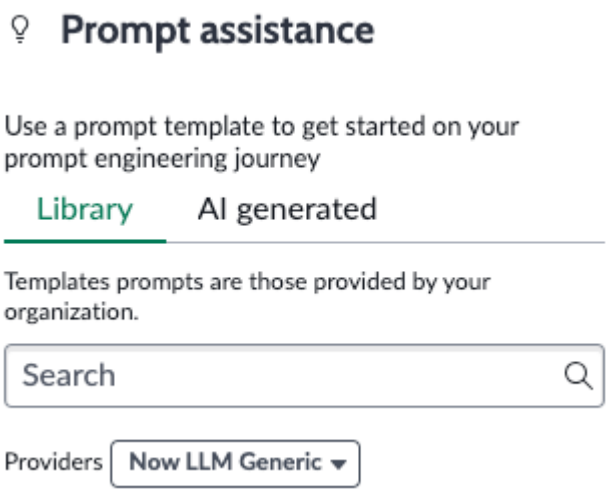
Role required: sn_skill_builder.admin

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Create a new skill or select the skill that you want to use prompt assistance for.
3. Select the prompt assistance icon.

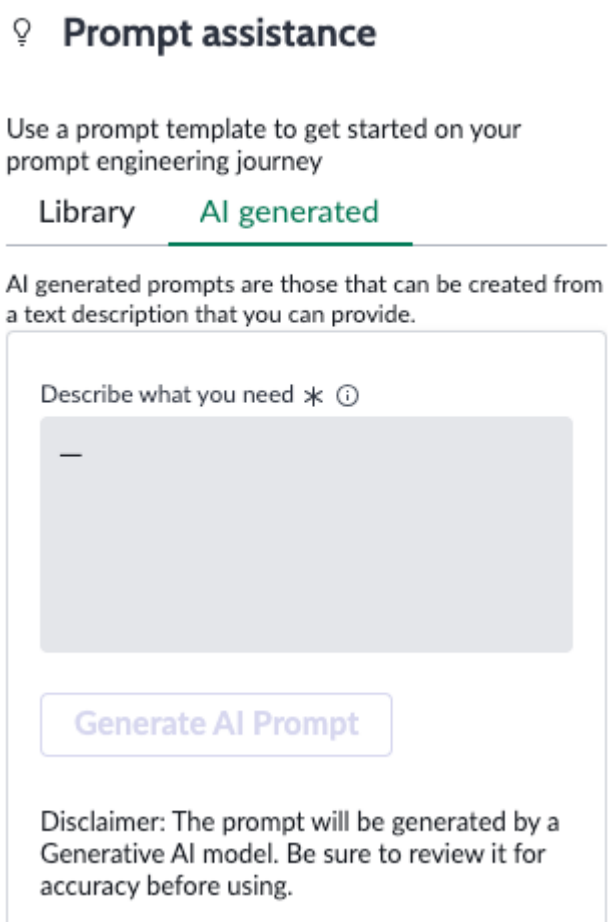
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

4. Select the **Library** tab to search for a prompt from the prompt library.



5.

6. Select the **AI generated** tab to have AI help you create a prompt based on your needs.



Test a prompt

After you create a prompt for your custom skill, test the prompt template before you finalize it. Testing the prompt verifies that you're seeing the expected prompt results before it's activated.

Before you begin

Role required: sn_skill_builder.admin

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

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Select the skill that you created the prompt for.
3. In the Test prompt section, select **Run tests**.
4. Choose a test incident or record.
5. Select **Run test**.

> **Note:** Testing your skill consumes an assist.

Prompt results

Tab	Description
Response	The response is the result that the large language model (LLM) sends back from the prompt.
Grounded prompt	The grounded prompt enables you to see the data that was brought into the prompt from your skill inputs and tools. With this view, you can see if your skill inputs and tools are returning the correct data.

6. **Optional:** If the skill is deployed as a flow, disable the system property com.glide.oneapi.fdih.async.quick.mode, and then enable flow reporting.
   This property allows the generation of flow execution details when running flows, subflows, and actions from a custom skill. You can use flow execution details to test and troubleshoot your flow, subflow, or action. For more information about the system property, see Workflow Studio flow system properties.
7. **Optional:** Refine the prompt if you want, and repeat testing as necessary.
8. **Optional:** Select the run test history icon to see the results from your previous run tests.

What to do next

After you test your prompt, you must finalize and publish it. To learn more about publishing a skill, see Finalize and publish a skill.

If you have not configured the deployment settings for your skill, see Configure skill deployment settings.

Evaluate a prompt

Use the Now Assist Skill Kit evaluation tools to evaluate the effectiveness of your skill prompts.

Before you begin

Role required: sn_skill_builder.admin

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Select the skill that you want to evaluate.
3. Select the **Prompt performance** tab.
4. Select the **Evaluation runs** tab.

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

5. Create a dataset from a table or data collection.

Create a data set

Method	Steps
Create a dataset from a table	**a.** Give the dataset a name and description. **b.** Select **Table**. **c.** Find the table that you want to use. **d.** Select the maximum number of records that you want to use. **e.** Add conditions. **f.** Select **Generate Preview**. **g.** Select the mappings. **h.** Select **Create**.
Create a dataset from a data collection	**a.** Give the dataset a name and description. **b.** Select Data Collection. **c.** Select a data collection that you created in Now Assist Data Kit. **d.** Select **Generate Preview**. **e.** Select the mappings. **f.** Select **Create**.

6. Select the add icon + for **Evaluation Runs**.
7. Give the evaluation run a name and description.
8. Select one or more prompts that you want to evaluate.
9. Select **Save & Next**.
10. Select a dataset.
11. Select **Save & Next**.
12. Expand the **Quality** tab.
13. Select the metrics that you want to evaluate.

Evaluation metrics

Evaluation method	Metric	Description
Human	Human Feedback	Human evaluation is the default option available for all prompt executions that generate a response. You can rate the response with a thumbs up or thumbs down, based on your satisfaction. You also have the option to provide more detailed feedback to explain your evaluation choice.
Automated	Correctness	The correctness metric assesses the generated response's accuracy, completeness, pertinence,

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

Evaluation method	Metric	Description
		and writing quality relative to the given instruction. This metric helps to check that the text accurately reflects the instruction, covers all important points, remains relevant, and is well written.
Automated	Correctness with Golden Response	The correctness with golden response metric uses a predefined reference to assess the generated response's accuracy, completeness, pertinence, and writing quality relative to the given instruction. This metric helps to check that the text accurately reflects the instruction, covers all important points, remains relevant, and is well written. You should use this metric whenever possible.
Automated	Faithfulness	The faithfulness metric assesses whether a generated response accurately reflects the information and context provided in the given instruction. This metric helps to check that the text contains no hallucinations, fabricated facts, or unsupported conclusions, maintaining alignment with the source material.

14. Select **Save & Next**.
15. Review the evaluation choices that you made.
16. Select **Save & Evaluate**.
17. **Optional:** Give a human evaluation.
    a. Select **Human evaluation**.
    b. Select a record to use in the evaluation.
    c. Expand the prompt and read the result.
    d. Select the thumbs up or thumbs down icon to give your evaluation.
    e. Add more information and select **Submit**.

Finalize and publish a skill

When you're satisfied with your prompt, you can publish your custom skill. Publishing the skill enables a Now Assist admin to activate it.

Before you begin

Role required: sn_skill_builder.admin

Procedure

1. Navigate to **All** > **Now Assist Skill Kit** > **Home**.
2. Select the skill that you want to publish.
3. Make any necessary changes to the prompt.
4. If you have changed the prompt, test it.
5. When you're satisfied with your prompt, select **Finalize prompt**.
6. Select **Publish**.

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

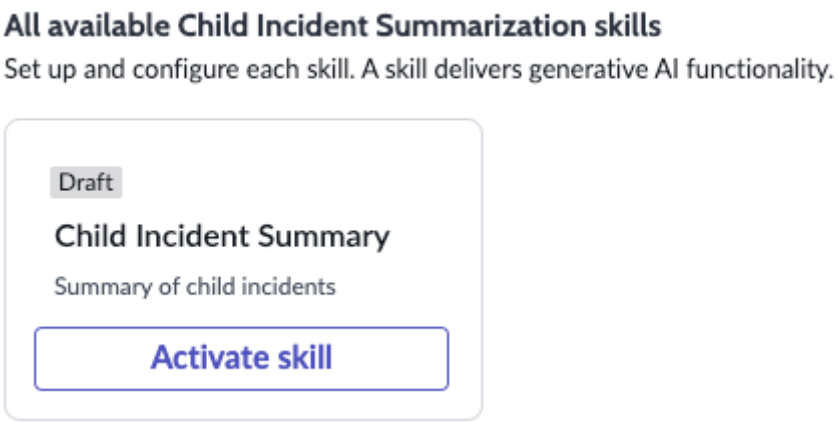
What to do next
A Now Assist admin must activate the skill. To learn more about how to activate a skill, see Activate a skill.

Activate a skill

After you create and publish a custom skill, you must activate it in Now Assist Admin. Activating the skill enables you to trigger the skill within the UI.

Before you begin
Role required: admin

Procedure
1. Navigate to **All** > **Now Assist Admin** > **Features**.
2. Find the custom skill that you created.
3. Select **Activate skill**.
   Make sure you are in the correct application scope when you activate the skill. To learn more about application scopes, see Application scope.



Note: Some skills might need to be reviewed and approved by a data steward before you can activate it.

What to do next
You can enable Now Assist Guardian for your skills. To learn more about Now Assist Guardian, see Now Assist Guardian

Call a custom skill from a script

You can use a script to call a custom skill.

Before you begin
Role required: admin

Procedure
1. Navigate to **All** > **System UI** > **UI Actions**.
2. Create a UI action.
   For more information on creating UI actions, see Create a UI action.
3. Add your script.

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

The following script is an example. You can replace the variables with your data.

```
var inputsPayload = {};

// create the payload to deliver input data to the skill

inputsPayload[‘input name’] = {

  tableName: 'table name',

  sysId: 'sys_id',

  queryString: ''

};


//create the request by combining the capability sys ID and the
 skill config sys ID

var request = {

    executionRequests: [{

        payload: inputsPayload,

        capabilityId: ‘capability sys id’,

        meta: {

            skillConfigId: ‘skill config sys id’

        }

    }],

    mode: 'sync'

};


//run the custom skill and get the output in a string format
try {
var output =
 sn_one_extend.OneExtendUtil.execute(request)['capabilities'][r
equest.executionRequests[0].capabilityId]['response'];
var LLMOutput = JSON.parse(output).model_output;
} catch(e) {
 gs.error(e);
 gs.addErrorMessage('Something went wrong while executing the
 skill.');
}
action.setRedirectURL(current);
```

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

Now Assist Skill Kit reference

Reference topics for Now Assist Skill Kit

Now Assist Skill Kit roles

Certain roles are required to use Now Assist Skill Kit functionality.

Skill kit admin (sn_skill_builder.admin)

This user can create, update, and publish skills in Now Assist Skill Kit. This role is mandatory to use Now Assist Skill Kit.

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.

Field of use for Now Assist Skill Kit

Now Assist Skill Kit is included in various Now Assist packages that cover a given customers' ability to build Now Assist skills.

The following table is provided to help clarify questions regarding how such custom skills can be created and deployed in each customer environment.

Expected field of use

Workflow Extensibility	Now Assist Skill Kit enables customers to supplement their use of licensed, pre-built Now Assist skills, like example Incident Summarization. Now Assist Skill Kit also enables customers to create new workflow related skills within the scope of their Now Assist purchases. For example, within ITx for ITSM Pro/Enterprise Plus, within ITx and HR for customers with ITSM, and HRSD Pro/ Enterprise Plus.
Non-duplicative	Refer to the Now Assist Overview to see the pre-built skills that are currently available exist for purchase in ServiceNow Now Assist packages. Customers must build from the purchased skills instead of attempting to create duplicative versions of licensable ServiceNow skills without purchase.

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

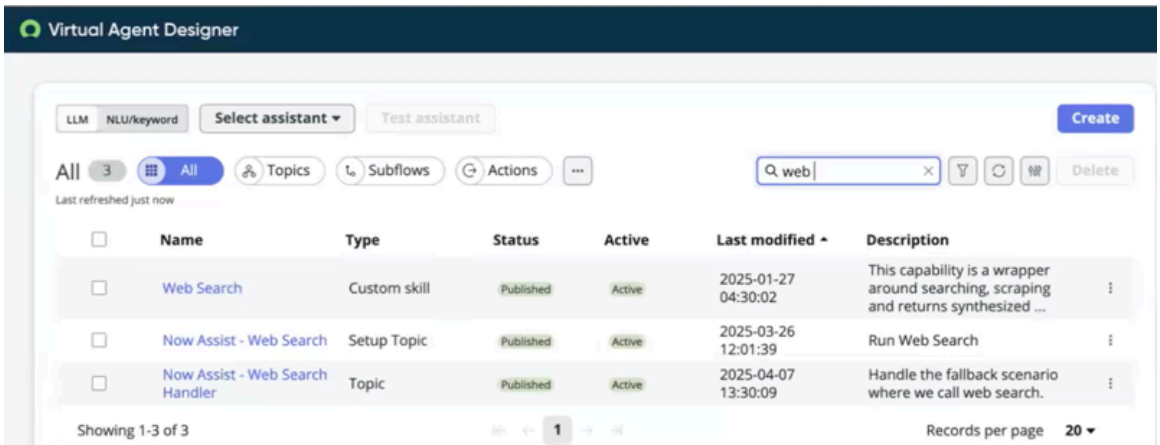
Expected field of use (continued)

Leveraging custom tables	Custom skills can be built using Now Assist Skill Kit to enable generative AI functionality on the custom tables that are provided in a customer's licenses.

Web search custom skill

The web search custom skill performs an internet search to answer a query. Web search is used whenever the LLM and AI Search are unable to provide results or whenever web search mode is activated.

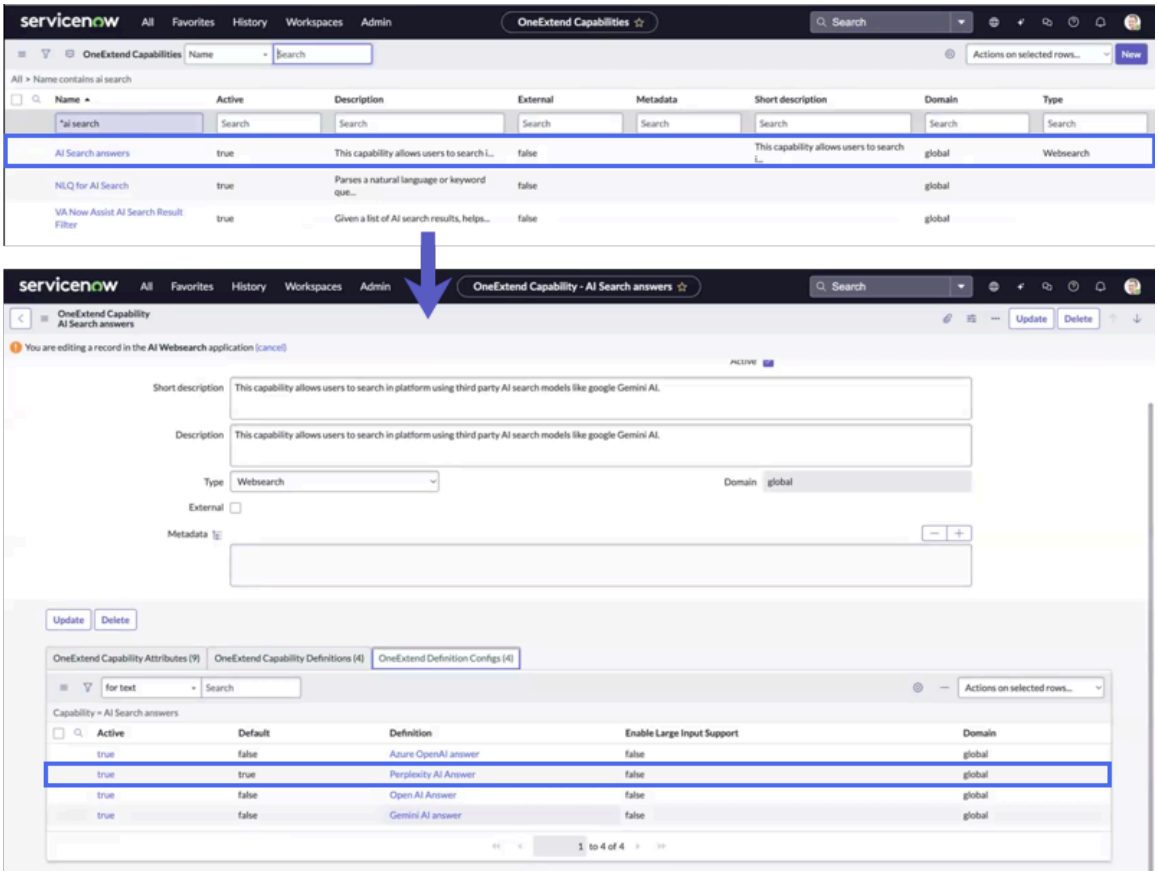
The web search custom skill is based on a prebuilt topic block in Virtual Agent Designer, and is available when you turn on Now Assist in Virtual Agent. This feature is automatically triggered when both AI Search and the LLM can't find information to answer an end user's query. This feature can also be manually triggered whenever end users select the Web search results icon ( ) to enter into web search mode. Once web search activates, a third-party AI performs the search and returns with an answer.

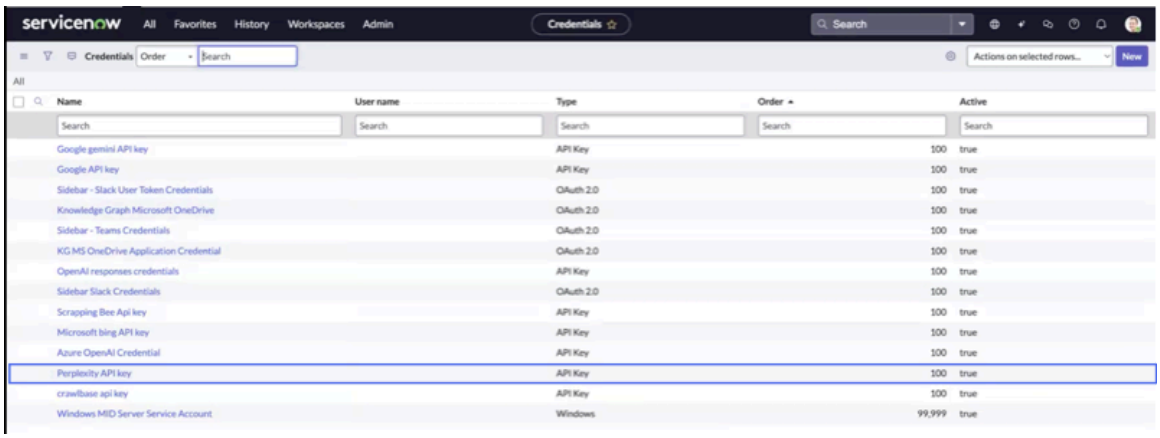


For more information about the web search mode end-user display and functionality, see Web search .

Web Search AI choices are available in the OneExtend Definition Configs related list within the *`AI Search answers`* entry in the OneExtend Capabilities [sys_one_extend_capability_list] table. Select one of the AI definitions there, such as Perplexity, and set its value to `true`. Activate a matching AI credential by navigating to **ALL** > **Connections & credentials** > **Credentials**, and setting the matching API key (such as *`Perplexity API key`* active value to `true`).

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





For more information on how to work with the *`AI Search answers`* OneExtend capability, see Configure AI search answers capability for web search.

For more information on how to add web search as a tool in the Now Assist Skill Kit, see Add a web search.

AI Control Tower

The AI Control Tower is a platform that connects different parts of an organization to speed up the AI adoption.

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

Get started

Explore



Learn how you use the AI Control Tower to manage and track all the AI systems, and other AI assets.

Configure



Configure all the AI Control Tower workflows

Use



Learn how to use AI Control Tower.

Reference



Get details about components such as roles and tables.

Risk and Compliance



AI Risk and compliance application to manage your artificial intelligence (AI) capabilities ethically, mitigate AI risks, and ensure compliance.

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

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.

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.

Note: We have controls in place to enable/disable the data collection and data processing.

Exploring AI Control Tower

Explore and learn about the AI Control Tower and its role in managing the AI asset inventory, overseeing configurations, integrating with other applications, and handling the asset approval workflow.

AI Control Tower overview

AI Control Tower provides enterprises with visibility into their AI footprint, manages the lifecycle of their AI investments, and identifies risks. This AI Control Tower's functionality ensures that their AI models, workflows, and agents are governed, optimized, and aligned with business strategy. It focuses on:

- **AI Strategy & Value:** Tracks how AI improves efficiency, revenue, and cost savings.
- **AI Governance & Compliance:** Ensures that AI systems follow company policies and global regulations, focusing on privacy, data governance, and ethical AI.

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

- **Security, Access & Privacy:** Manages AI agent identities, access controls, and sensitive data to ensure security.
- **Lifecycle Management:** Handles the entire lifecycle of AI assets, including onboarding, offboarding, skill management, and performance monitoring.
- **Health & Performance:** Evaluates how reliably and efficiently AI systems operate.
- **Enterprise AI Integrations:** Provides a platform to discover and use AI projects, models, and data from within the company and external partners.

AI asset Data Model

The AI asset data model catalogs different AI assets such as AI systems, AI models, datasets, and prompts. It also manages the relationship between different assets that helps manage the AI inventory effectively.

AI Control Tower benefits

- Visibility and management of enterprise AI in the organization.
- Manage the lifecycle of AI assets such as AI systems, AI models, prompts, and datasets.
- Identify risks and set controls.
- Proactively manage AI risks and compliance while protecting your business and speeding up AI innovation.

What to explore next

To learn more about AI Control Tower and other applications integrating with its workspace, see the following topics:

- AI Control Tower dashboard
- Configure AI Control Tower
- AI Risk and Compliance
- Explore Now Assist AI asset discovery
- Enterprise AI discovery: Unlock Visibility, Governance & Value
- AI Gateway
- Using AI Control Tower
- AI Control Tower reference

AI Control Tower dashboard

The AI Control Tower dashboard displays a comprehensive overview of the AI status, AI inventory, and other AI related metrics.

The AI Control Tower dashboard's left navigation view has **Home**, **AI assets**, and **Configuration** icons.

The table below illustrates how the navigation view in AI Control Tower appears for different persona.

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

Navigation view for users

Persona	Navigation view
AI steward	• Home page • AI assets • Configurations
Product owner/ asset owner	• Home • AI assets
AI Control Tower workspace user	• Home page • AI assets
Risk and compliance user	• Home page • AI assets

Note: The information and actions available to a user will vary based on their role. For instance, AI stewards have access to the state of AI across all inventory, including value and adoption insights. Product owners, on the other hand, can view only the assets they manage, along with details on their value and adoption.

AI Control Tower Home

The AI Control Tower home view on the dashboard features recommendations for actions at the top displaying task statuses, pending asset reviews, and newly added AI systems. The home page displays the Overview tab in AI Control Tower, AI strategy tab in AI Control Tower, AI asset inventory tab in AI Control Tower, Value tab in AI Control Tower, Adoption tab in AI Control Tower, Risk and compliance tab in AI Control Tower, AI cases tab in AI Control Tower and Security & Privacy tab in AI Control Tower tabs.

Each of these tabs includes widgets, and you can explore the data within each widget in more detail. For instance, the AI system trend widget displays trends based on historical data about AI system deployments. Similarly, the AI risk classification widget shows data according to the default risk classification settings for your skills.

AI Control Tower Home

Explore the AI Control Tower home page.

The AI Control Tower home page view on the dashboard features recommendations for actions at the top displaying task statuses, pending asset reviews, and newly added AI systems.

The home page displays the Overview tab in AI Control Tower, AI strategy tab in AI Control Tower, AI asset inventory tab in AI Control Tower, Value tab in AI Control Tower, Health tab in AI Control Tower Risk and compliance tab in AI Control Tower, AI cases tab in AI Control Tower, Security & Privacy tab in AI Control Tower.

Important: The AI strategy tab in AI Control Tower is available only with Strategic Portfolio Management (SPM) Professional license.

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

Each of these tabs includes widgets, and you can explore the data within each widget in more detail. For instance, the AI system trend widget displays trends based on historical data about AI system deployments. Similarly, the AI risk classification widget shows data according to the default risk classification settings for your skills.

The data appearing on these widgets consists entirely of transactional data. Data is populated once assessments are completed. When AI Control Tower is first installed, there's no data available.

Note: AI system trend's data is dependent on two types of scheduled job.

- AI Control Tower Core Monthly Data Collection
- AI Control Tower Core Historical Data Collection job

The AI Control Tower Core Monthly Data Collection is a job that runs each month to collect data. In AI Control Tower, this data is displayed on a quarterly basis through aggregation. So, if you configure AI Control Tower for the first time you don't see any data.

The AI Control Tower Core Historical Data Collection job, which is used to gather historical data. If there was any relevant data from previous months, it can be collected using this job. You can run this job manually, as the monthly job only collects data going forward, not from the past.

Overview tab in AI Control Tower

Explore the overview tab in AI Control Tower and the widgets available on the tab.

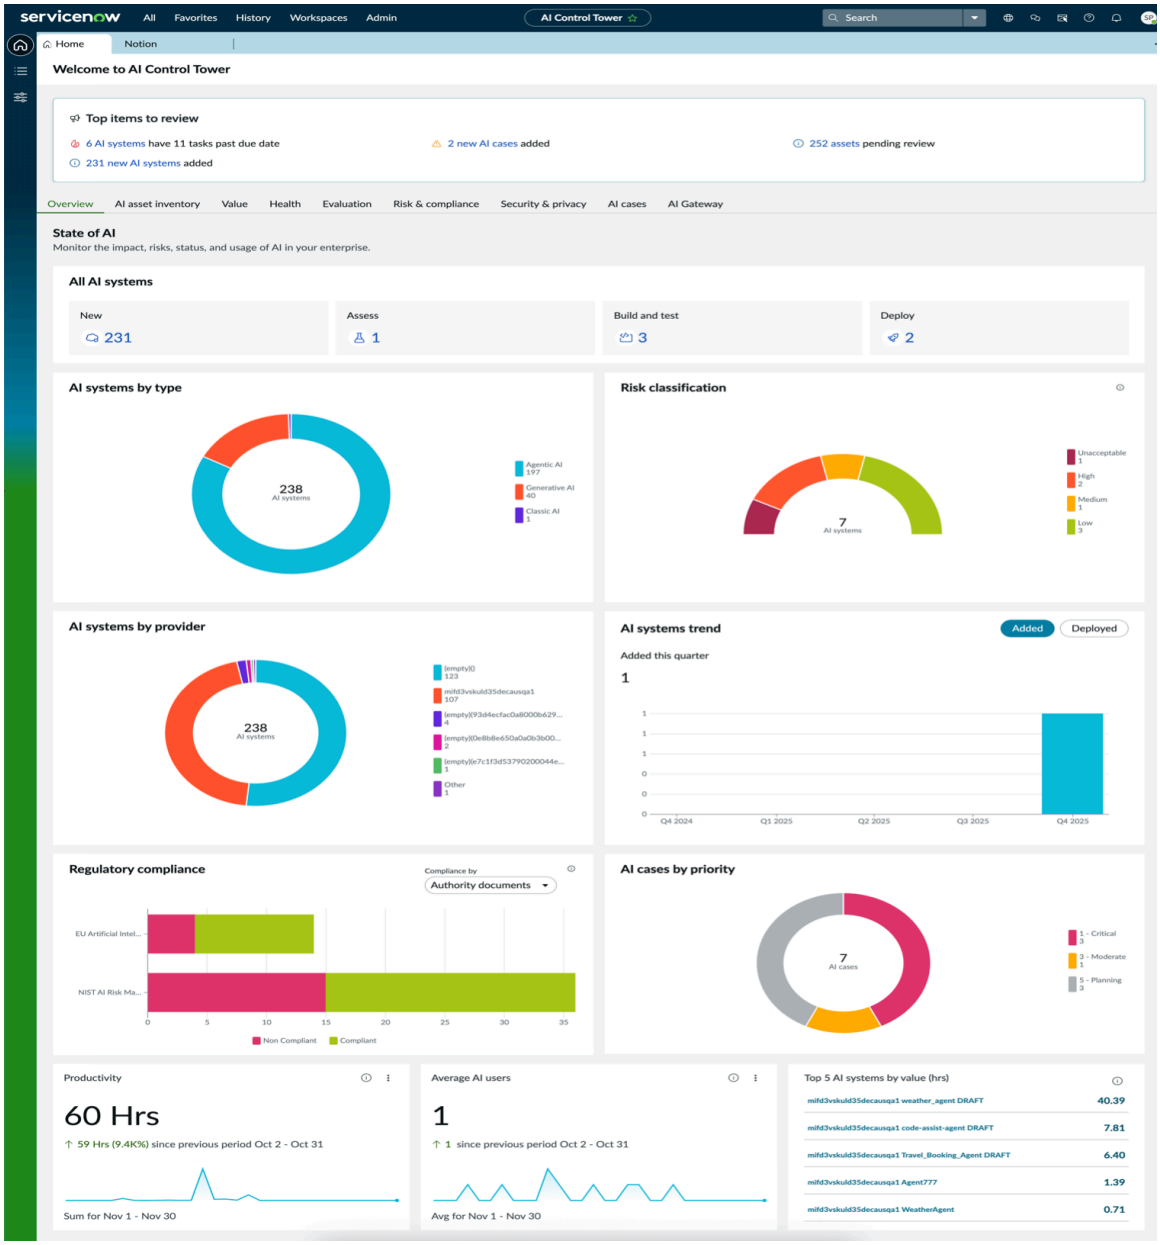
Overview dashboard

Prerequisites

Role required: AI steward [sn_ai_governance.ai_steward]

The **Overview** tab of the AI Control Tower dashboard provides a comprehensive view of all metrics of AI systems.

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



The following widgets are available, showing various metrics. You can view the data in detail on each widget.

All AI systems

This section shows all the AI systems classified into four stages:

- New
- Assess
- Build and test
- Deploy

For more information about the AI asset lifecycle stages, see AI asset lifecycle

AI systems by type

This section shows all the AI systems categorized into

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

- Agentic AI
- Generative AI
- Classic AI

Risk classification

This section shows the risk score of AI systems. The risks are qualitatively classified as

- High
- Low
- Medium
- Unacceptable

AI systems by provider

This section shows the number of AI systems, categorized by the providers.

> **Note:** The AI system by provider donut chart doesn't display the **Other** category when there are more than five providers, which can cause a mismatch between the number of metric labels and the items shown in the list when selecting a metric label.

Compliance

This section shows the compliance effectiveness percentage of different statuses of active controls applied on **Authority documents** and **Policies**. You can drill down respective sectors of the bar graph to access and view the list of active controls.

AI cases by priority

This section shows the total number of AI cases with breakdown analysis by priority rating. You can drill down to respective sectors of the donut chart to access and view a list of recorded AI cases.

AI strategy tab in AI Control Tower

The **AI strategy** tab includes information about all the AI-related priorities, goals, planning items, and execution items—projects and demands in the AI Control Tower workspace.

This tab helps you track and monitor AI-related work across various sections—Strategy, Costs, Prioritized AI work, and AI Risks, Issues, Decisions, Actions, and Request Changes (RIDAC)—enabling you to make informed decisions. The tab displays the AI-related work defined in ServiceNow® Goal Framework, the ServiceNow® Strategic Planning, and ServiceNow® PPM Standard.

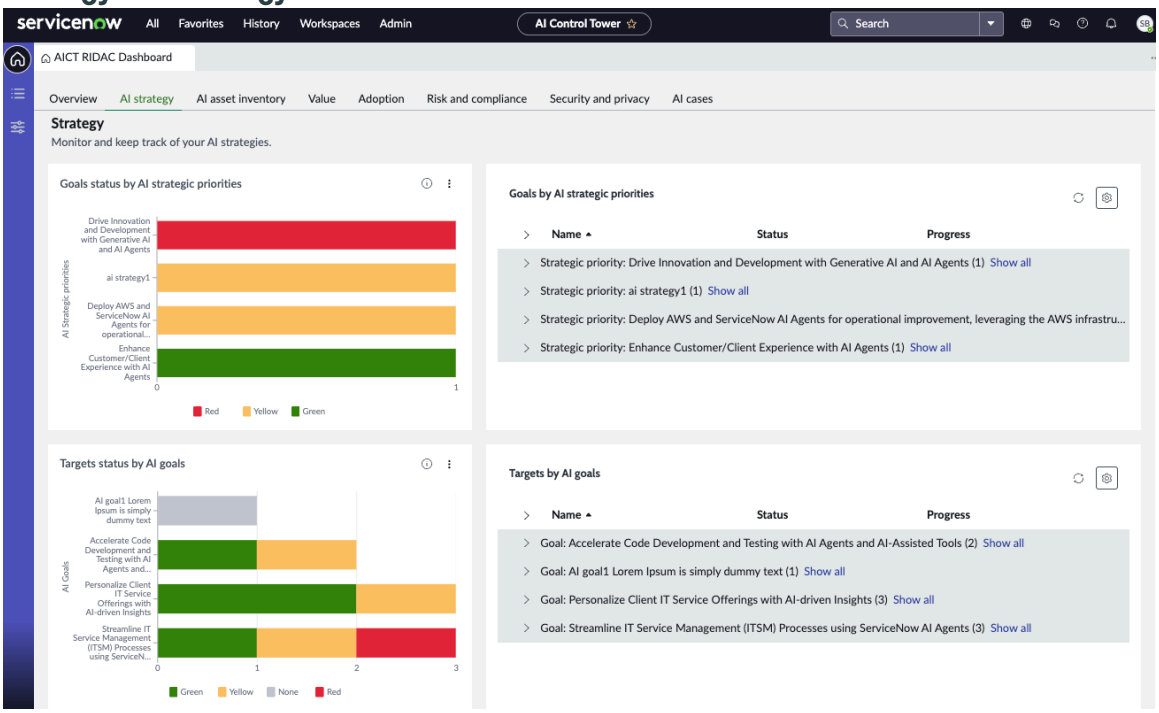
Strategy

The Strategy section helps you monitor and track your AI strategic priorities, goals, and targets associated with these goals.

Strategic priorities with the **Type** field set to **Artificial Intelligence** and goals with the **Category** field set to **Artificial Intelligence** in the Goal Framework application appear in the Strategy section in the AI Control Tower workspace, enabling you to monitor their progress effectively. For details, see Create a strategic priority for a strategic plan and Create a goal in Strategic Planning.

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

Strategy in AI strategy

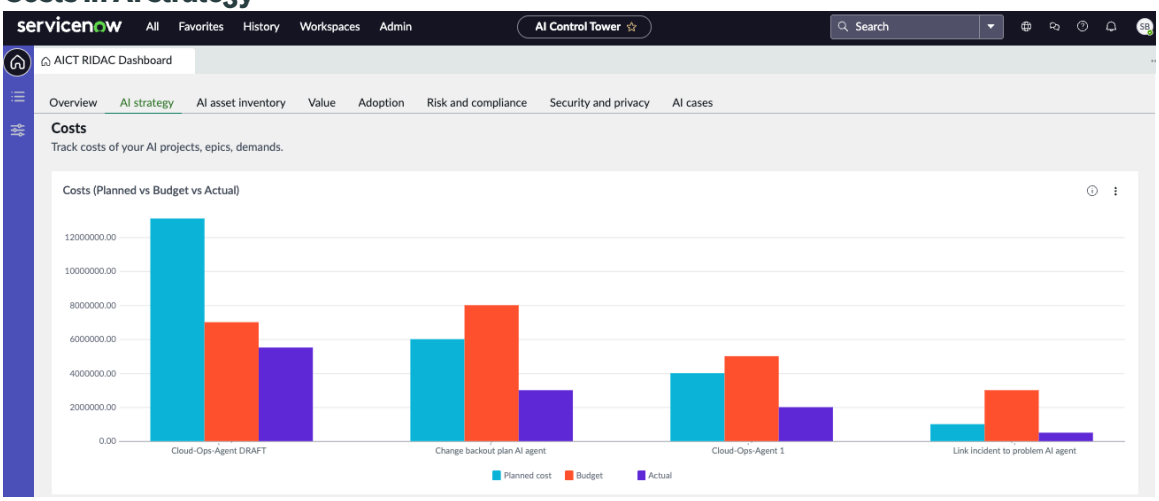


Costs

The Costs section helps you track the costs of your AI planning items—such as projects, epics, and demands—by comparing planned cost, budgeted cost, and actual cost.

Planning items with the **Investment type** field set to **Artificial Intelligence** in the Strategic Planning Workspace appear in the Costs section in the AI Control Tower workspace, enabling you to monitor their progress effectively. For details, see Add planning items to Prioritization.

Costs in AI strategy



Prioritized AI work

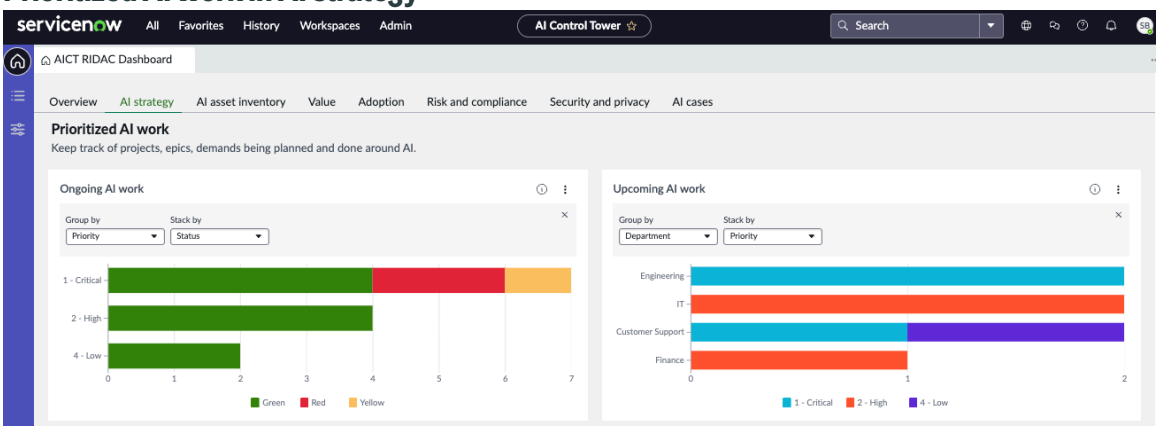
The Prioritized AI work section helps you track your projects, epics, demands being planned and done around AI.

Planning items with the **Investment type** field set to **Artificial Intelligence** and the **State** field set to **Prioritized** in the Strategic Planning Workspace appear in the Prioritized AI work section in the

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

AI Control Tower workspace, enabling you to monitor their progress effectively. For details, see Add planning items to Prioritization and Prioritize work.

Prioritized AI work in AI strategy



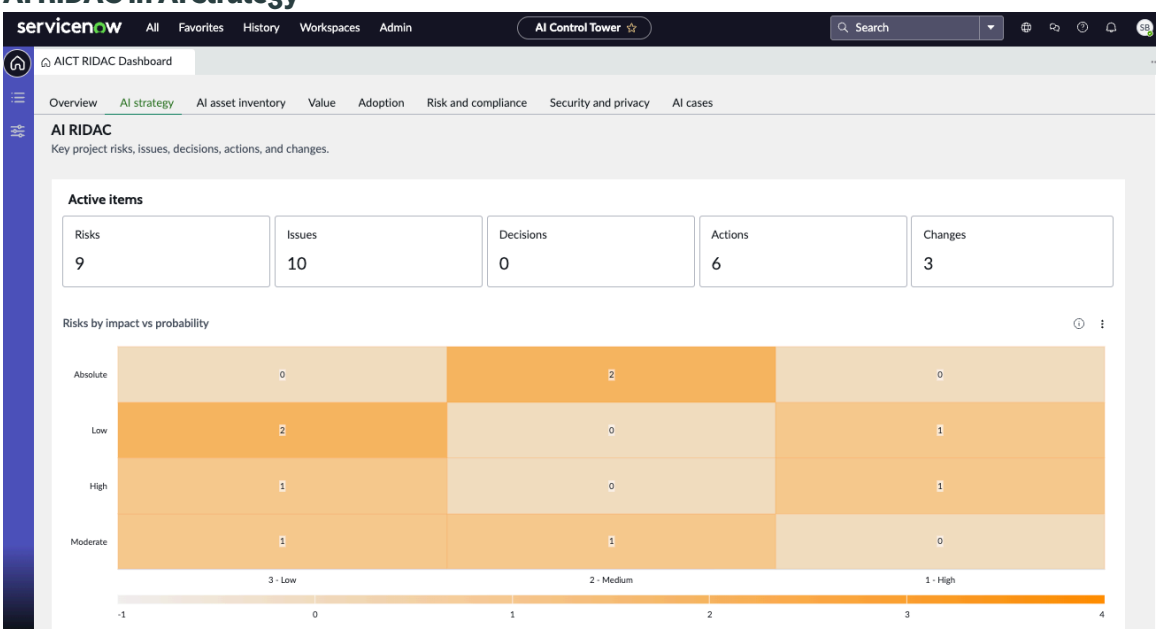
AI RIDAC

The AI RIDAC section helps you monitor key project risks, issues, decisions, actions, and changes, enabling you to make informed decisions.

Projects that meet any of the following criteria appear in the AI RIDAC section in the AI Control Tower workspace, enabling you to monitor their progress effectively:

- Projects with the **Investment type** field set to **Artificial Intelligence** in PPM Standard (Project Management or Project Workspace). For details, see Create a project from Project Workspace and Add a RIDAC record to a project in Project Workspace.
- The **Category** field on the associated primary goal record is set to **Artificial Intelligence** in Goal Framework or Strategic Planning Workspace. For details, see Create a goal in Goal Framework and Create a goal in Strategic Planning.
- The **Type** field on the strategic priority of the associated primary goal is set to **Artificial Intelligence** in Goal Framework or Strategic Planning Workspace. For details, see Create a strategic priority for a strategic plan.

AI RIDAC in AI strategy



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

AI asset inventory tab in AI Control Tower

The AI asset inventory includes all the AI-related assets used by an organization, including AI models, prompts, systems, and databases.

Explore AI asset inventory

The **AI asset inventory** tab displays the following widgets. You can view the data in detail on each widget.

AI systems by provider

This section displays the number of AI systems, categorized by the providers using a donut chart.

> **Note:** The AI system by provider donut chart doesn't display the **Other** category when there are more than five providers, which can cause a mismatch between the number of metric labels and the items shown in the list when selecting a metric label.

AI systems by type

This section displays all the AI systems categorized into **Agentic AI**, **Generative AI**, and **Classic AI** categories using a donut chart.

All AI systems

This section displays all the AI systems classified into four stages: **Steward review**, **Approved for development**, **Ready for deployment**, and **Deploy**.

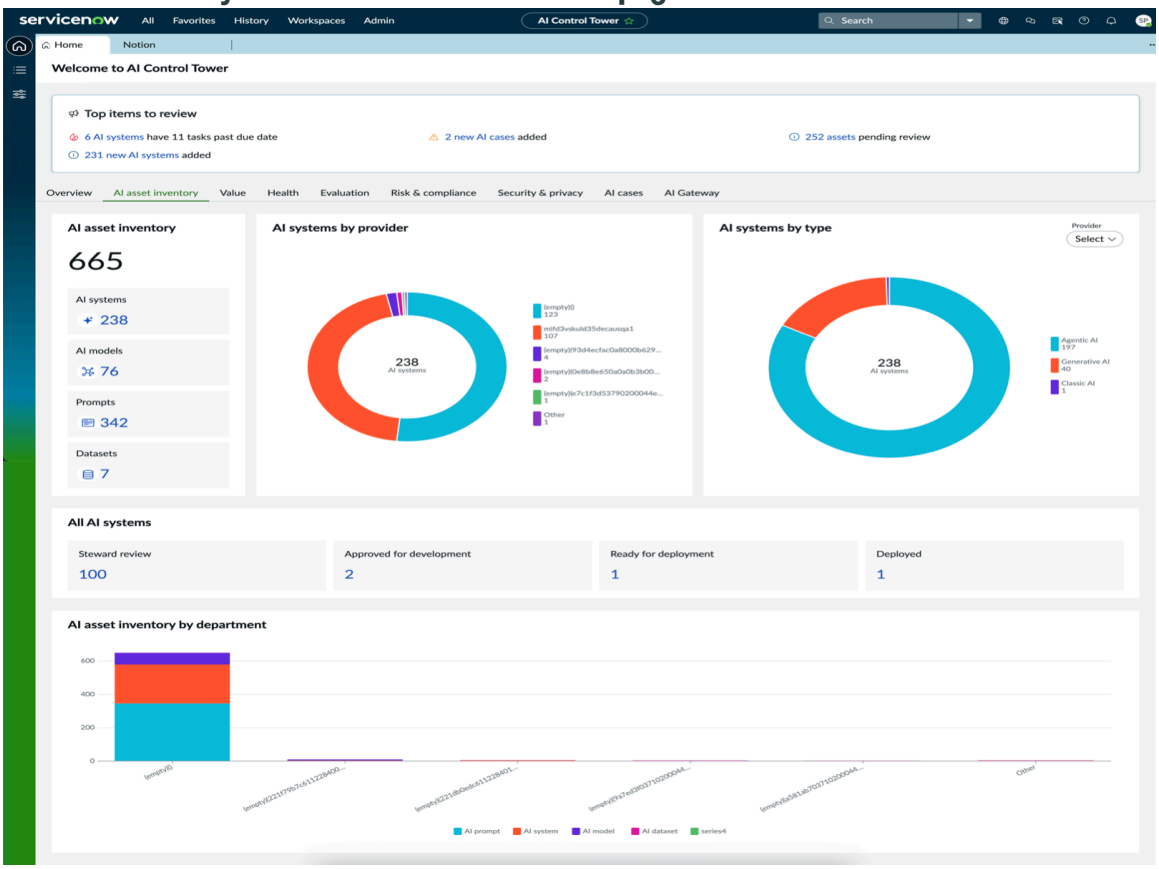
AI asset inventory by department

This section displays the number of AI systems classified in the **Providers** and **Components** using bar graph.

The following image shows the **AI asset inventory** tab on the AI Control Tower.

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

AI asset inventory tab in AI Control Tower home page



Value tab in AI Control Tower

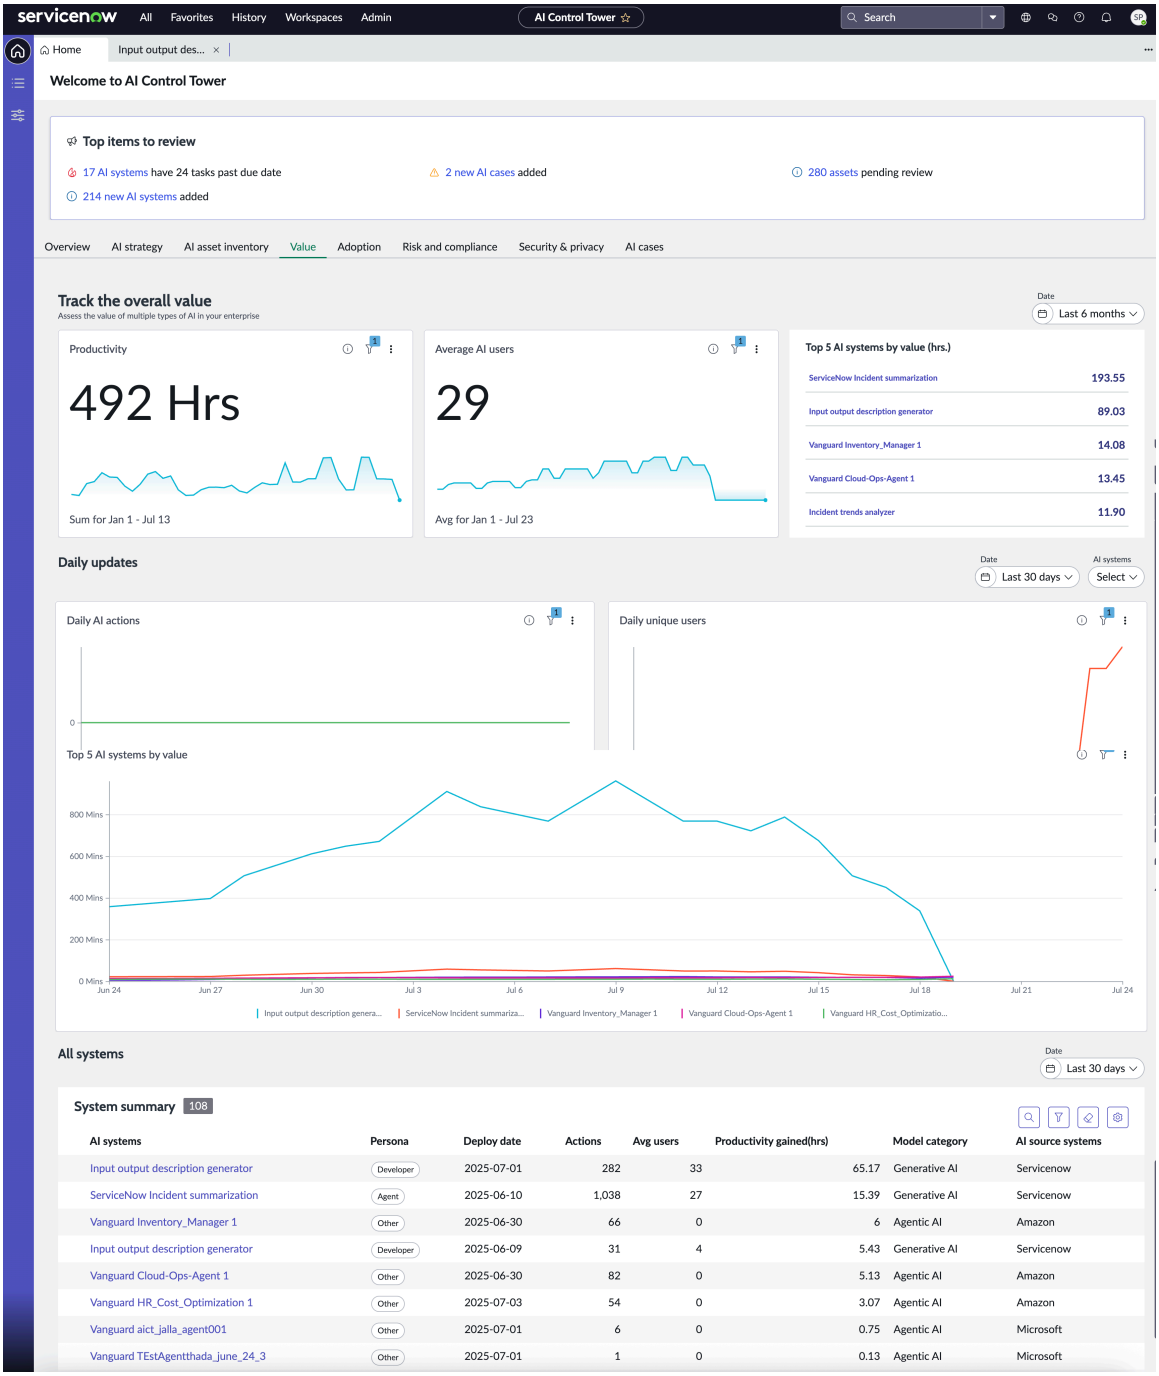
Use the Value dashboard page to gain insights into the value realized from multiple types of AI in your organization.

The Value dashboard page contains visualizations that help you assess the value of your AI assets. Creator skills are included in the Value dashboard page. Value metrics for creator skills are aggregated across non-production and production instances.

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

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

Value tab in AI Control Tower



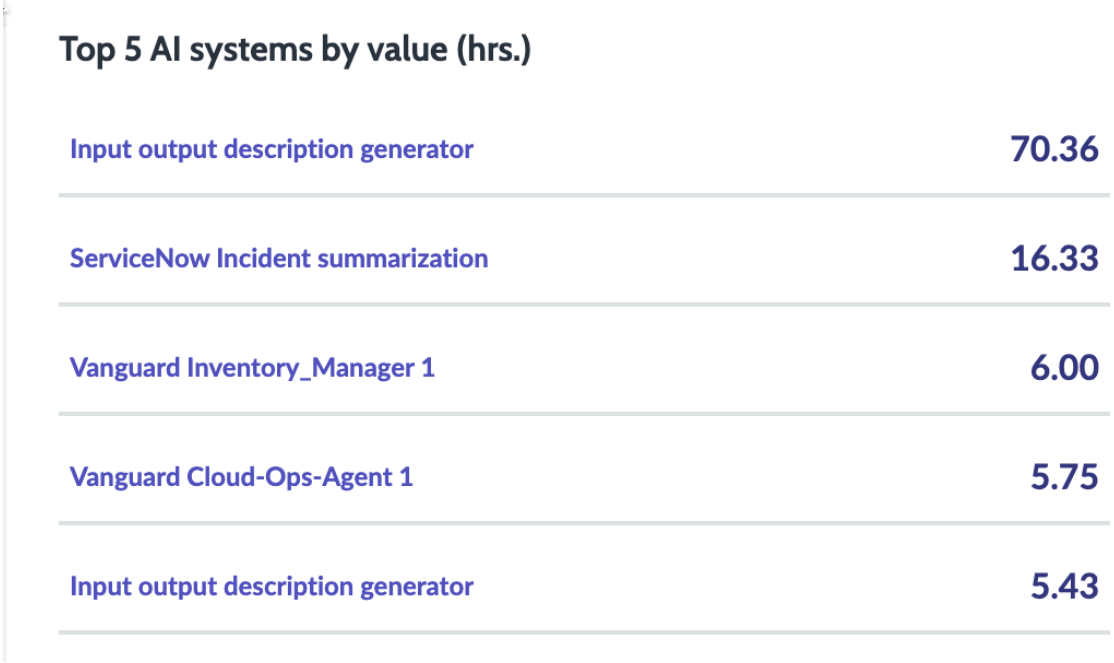
The indicators on the dashboard are based on the value templates associated with each AI asset. See Using value templates for more information.

Top 5 AI systems by value

This area of the dashboard shows the five AI systems that generated most value (in hrs).

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

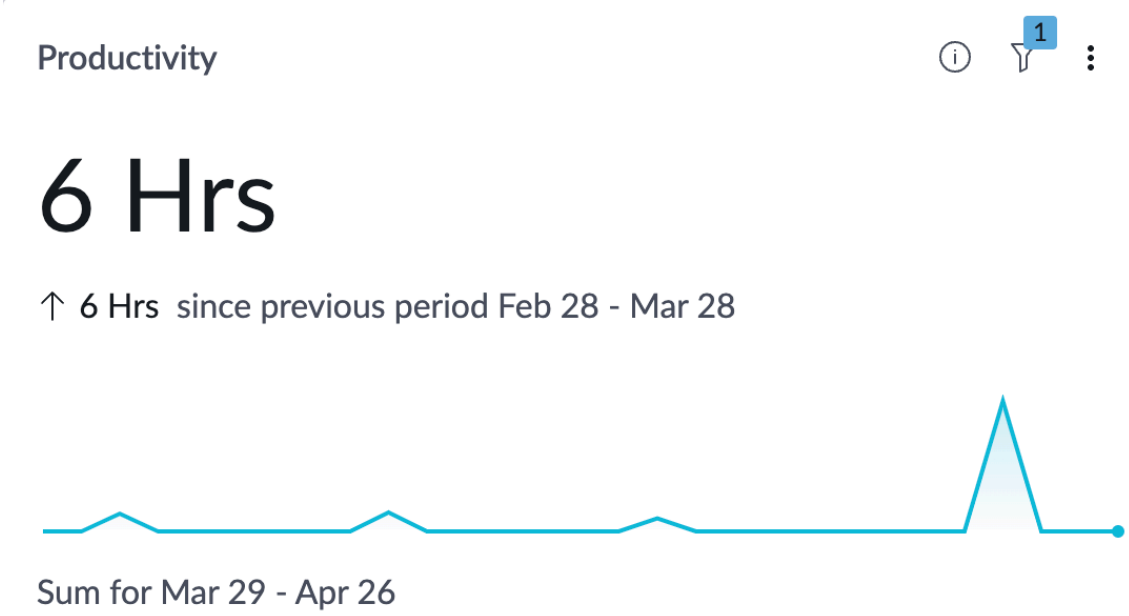
Top 5 AI systems by value



Productivity

This area of the dashboard shows the estimated productivity gains (in hours) in the selected period as a result of using various AI assets.

Productivity

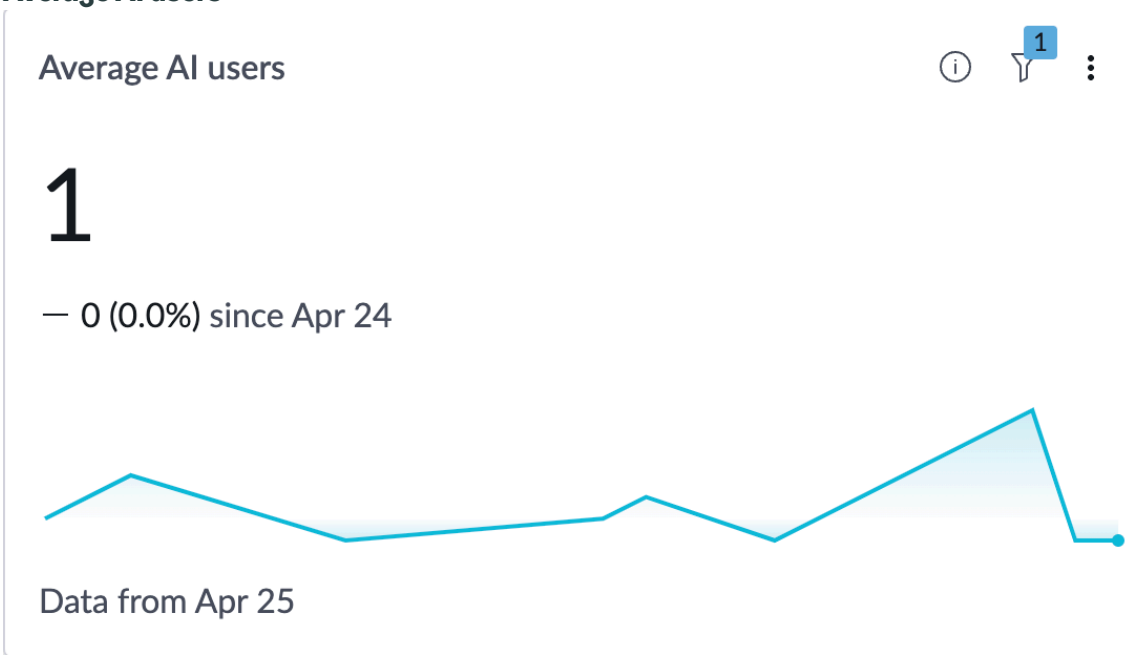


Average AI users

This area of the dashboard shows the average unique users who used an AI asset in the selected date range.

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

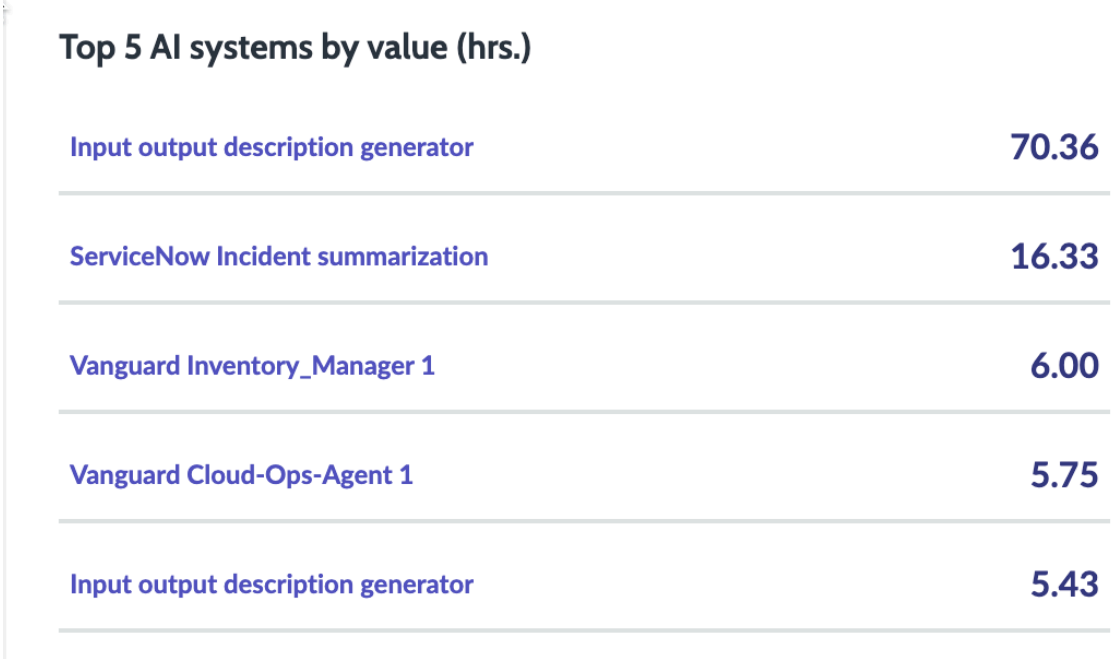
Average AI users



Top 5 AI assets by value (in hrs)

This area of the dashboard shows the top 5 AI assets that delivered the most value to the users in the selected date range.

Top 5 AI assets by value (in hrs)



Daily updates

Daily AI actions

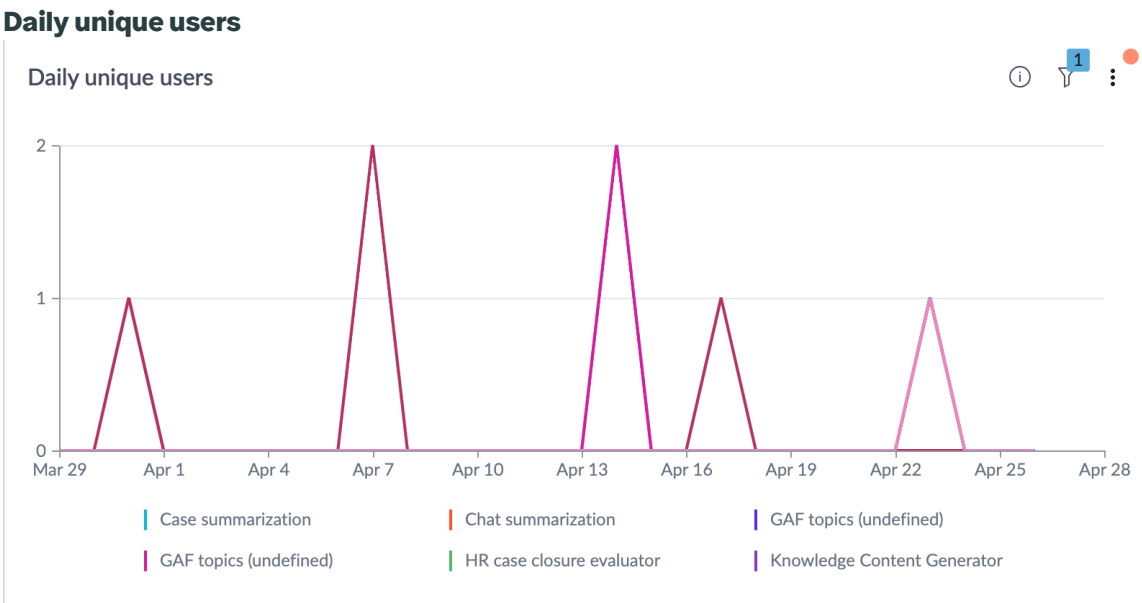
This area of the dashboard shows the number of AI actions per each AI asset.

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



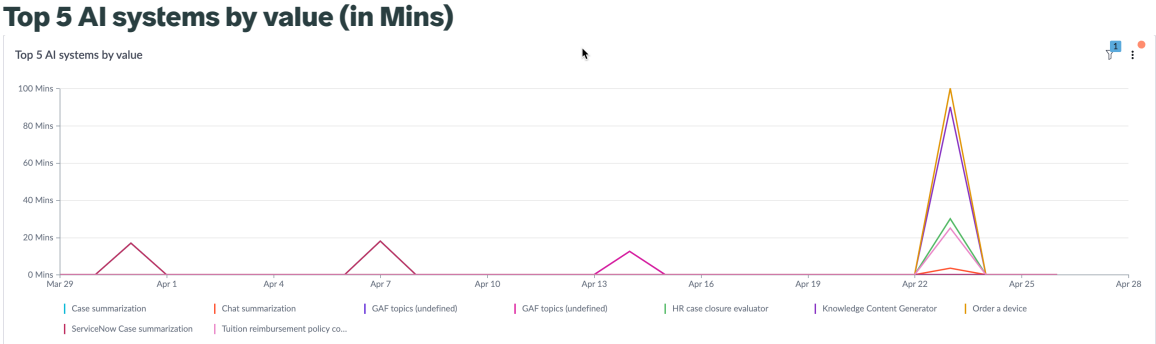
Daily unique users

This area of the dashboard shows the number of daily unique users per each AI asset.



Top 5 AI systems by value

This area of the dashboard shows the daily value delivered (in Mins) by the top 5 AI systems.



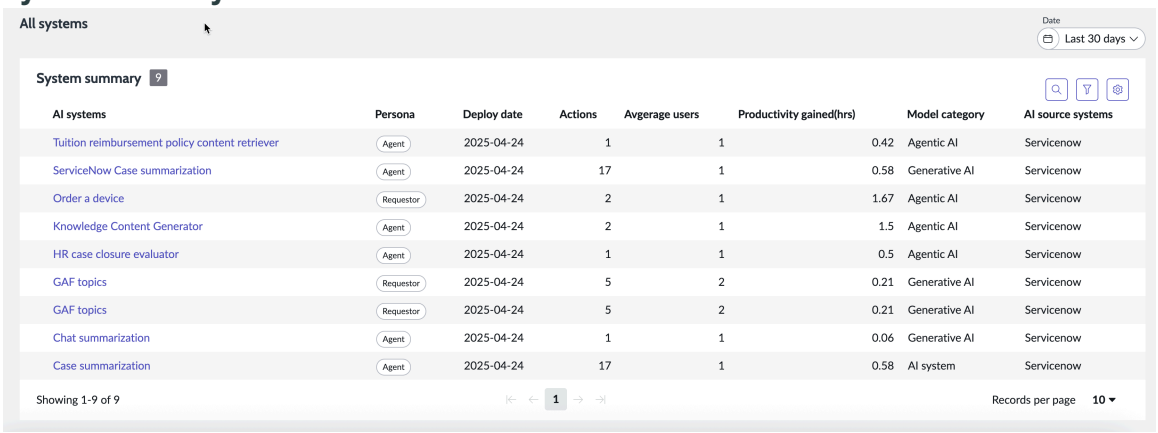
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

All systems

System summary

This area of the dashboard lists all the AI systems along with details of users and productivity gained.

System summary



Adoption tab in AI Control Tower

Use the Adoption dashboard page to view data about user engagement and feedback, and adoption of every type of AI.

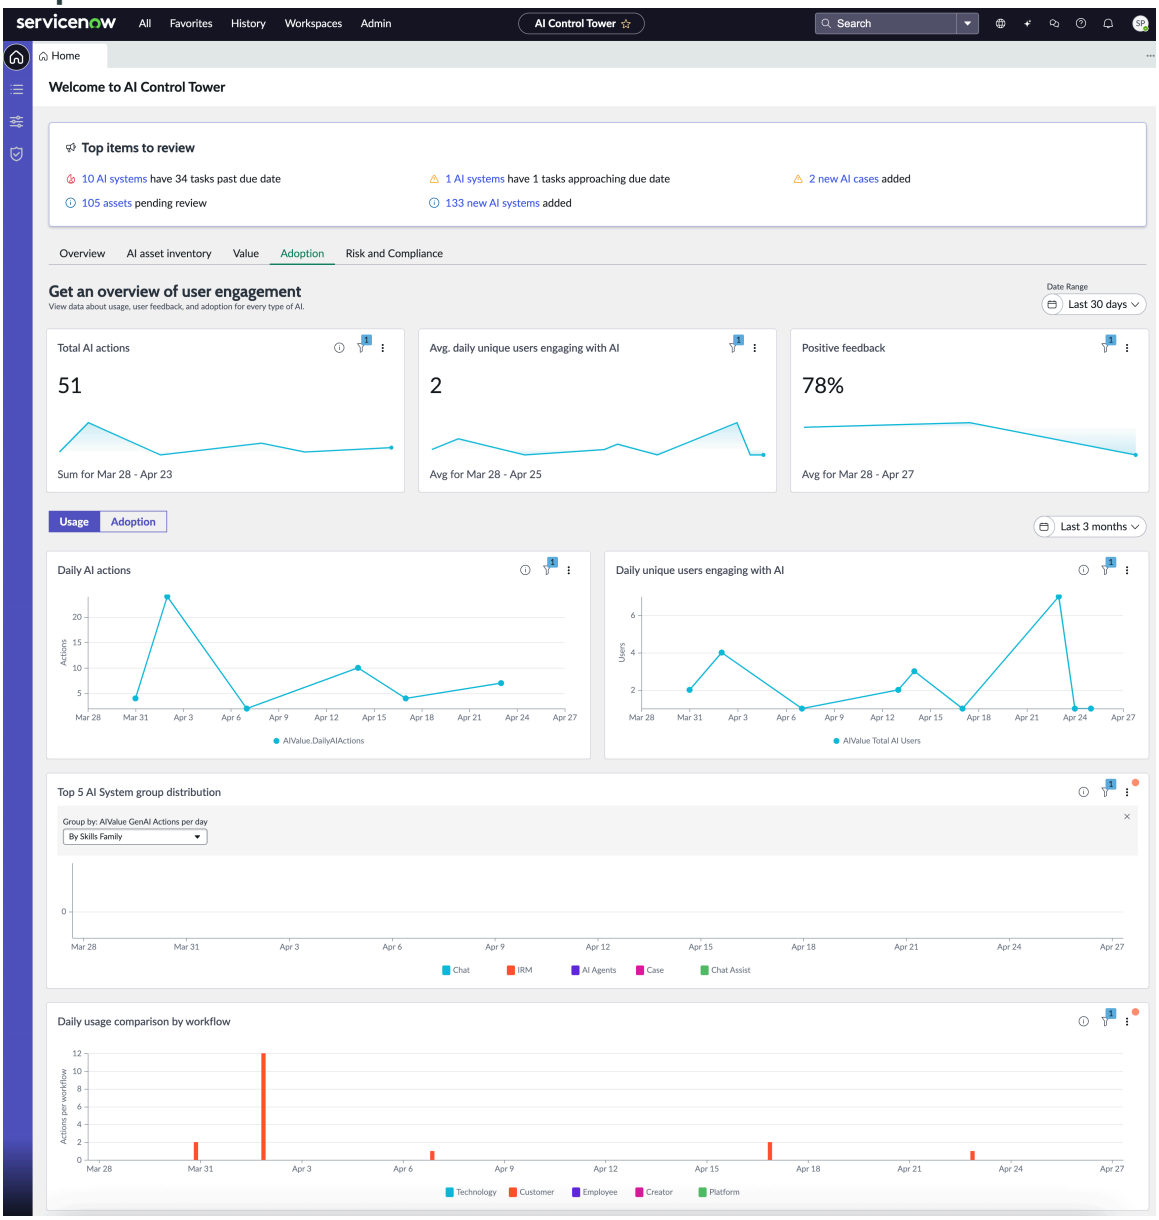
The Adoption dashboard page contains two tabs that provide insights on usage and adoption.

The **Usage** tab includes indicators on daily AI actions, daily unique users who have engaged with AI, AI system group distribution, and daily usage comparison by workflow. The **Adoption** tab tracks the departments with the highest AI usage, comparison of AI actions by user department, and feedback details. Use the date range filter to view usage and adoption indicators for the selected period. The default date range is 30 days.

The high-level indicators on the Adoption dashboard page provide an overview of user engagement.

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

Adoption tab in AI Control Tower

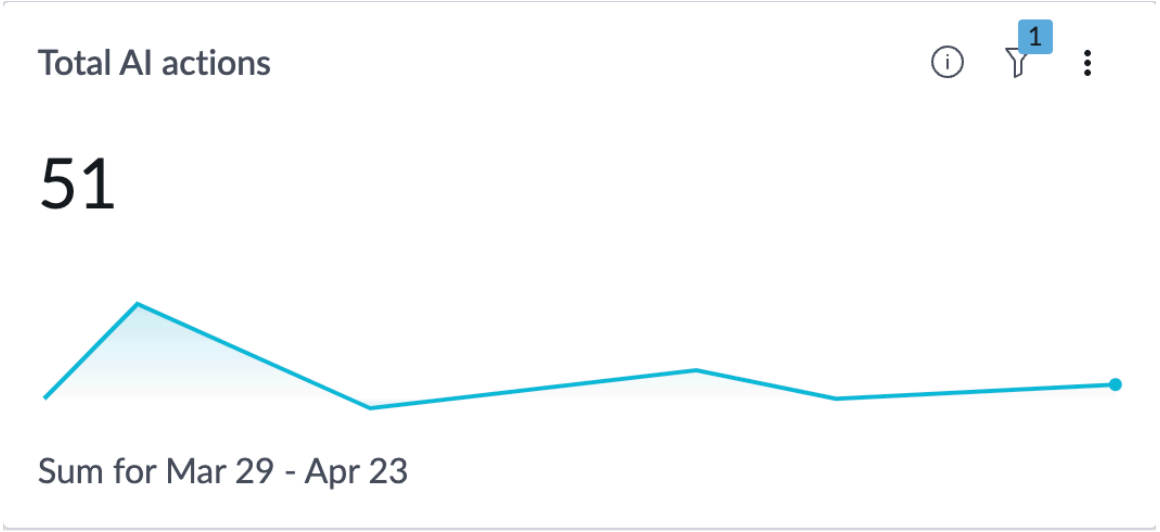


Total AI actions

This area of the dashboard shows the total number of AI actions in the selected date range. Total AI actions for a selected period can reveal the scale of AI actions executed. The trend line in the visualization shows increased or declining engagement.

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

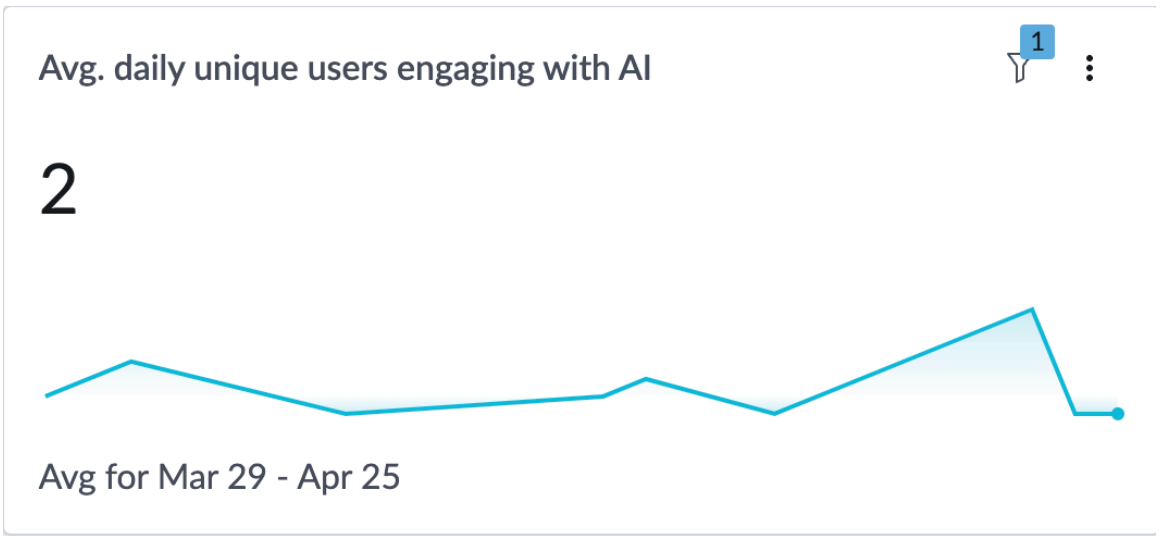
Total AI actions indicator



Average daily unique users engaging with AI

This area of the dashboard shows the average number of daily unique users who engaged with AI in the selected date range. The Average number of daily unique users is an indicator of the level of AI engagement.

Average number of unique users per day who triggered AI actions indicator

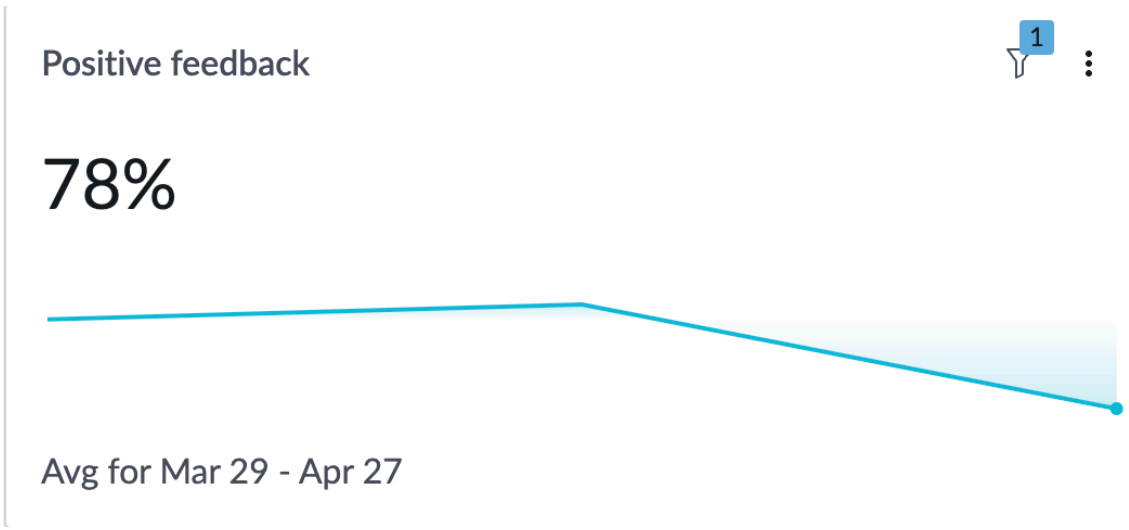


Positive feedback

This area of the dashboard shows the percentage of positive feedback received against executed AI actions. The Positive feedback indicator reveals the level of user satisfaction for executed AI actions.

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

Positive feedback indicator

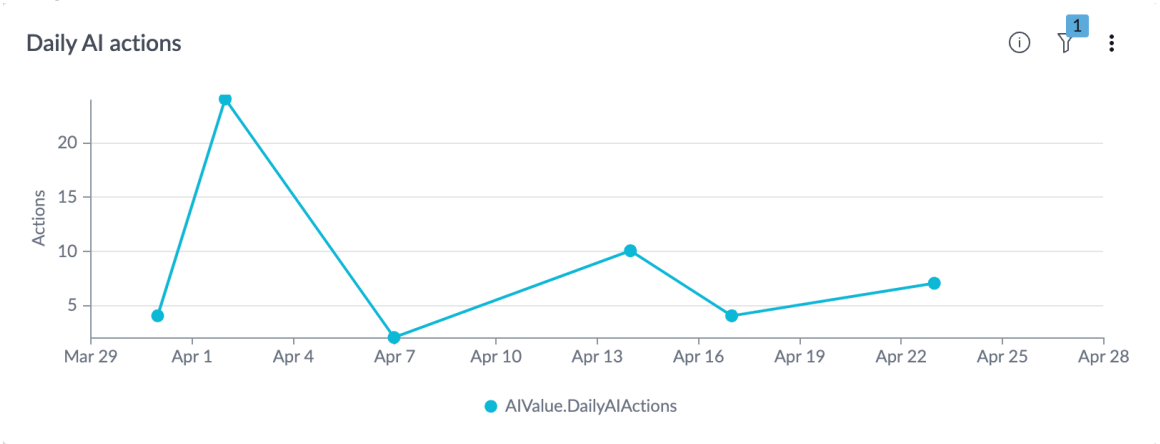


Usage indicators

Daily AI actions

This area of the dashboard shows the number of AI actions executed per day in the selected date range.

Daily AI actions indicator

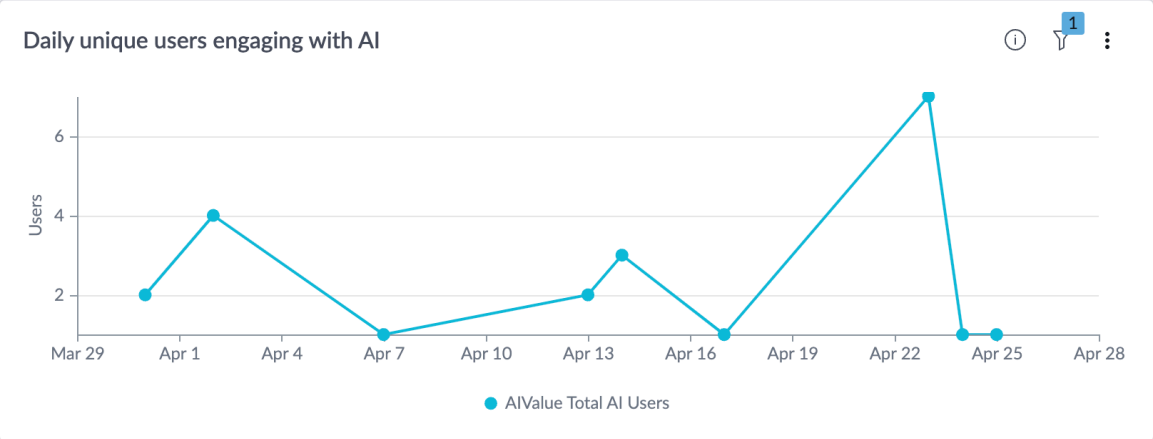


Daily unique users engaging with AI

This area of the dashboard shows the daily unique users who engaged with AI in the selected date range.

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

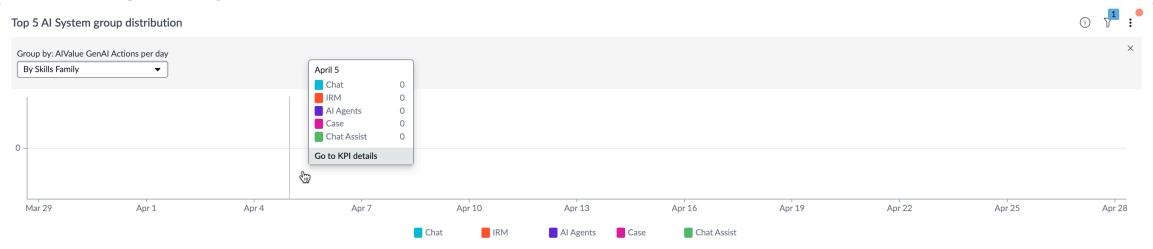
Daily unique users engaging with AI indicator



Top 5 AI System group distribution

This area of the dashboard shows the usage of the top 5 AI actions in a trend chart. The visualization is interactive. Use the Group by filter to apply predefined breakdowns. Hover over the trend lines to see the number of times that each AI action was used. Selecting one or more items in the legend removes the trend line for those items. You can compare and contrast AI actions usage distribution in a date range to understand in-demand actions.

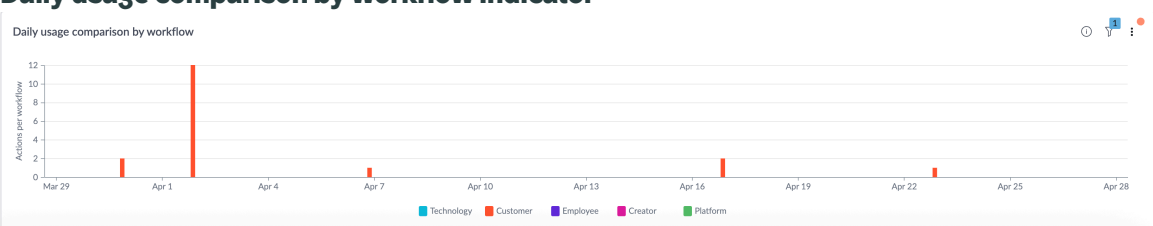
Top 5 AI System group distribution indicator



Daily usage comparison by workflow

This area of the dashboard shows the number of AI actions executed per day against each workflow in the selected date range. Hover over a bar to view the count of AI actions executed for a workflow on that day.

Daily usage comparison by workflow indicator



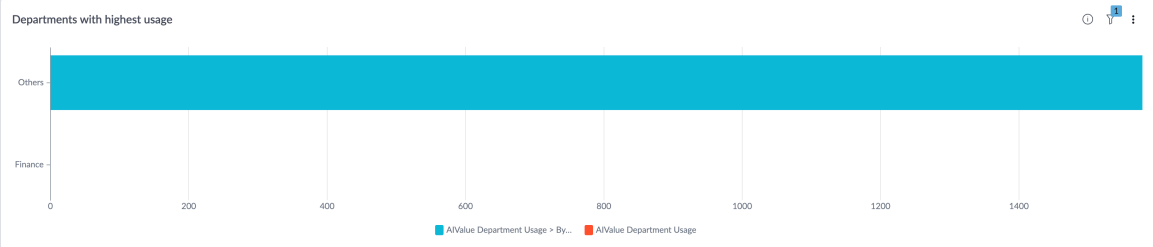
Adoption indicators

Departments with highest usage

This area of the dashboard shows the departments with the highest usage of AI actions in the selected date range. Hover over the horizontal bars to view the count of actions for a particular department.

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

Departments with highest usage indicator



AI actions comparison by user department

This area of the dashboard shows the number and type of AI actions executed against each department in the selected date range.

AI actions comparison by user department indicator

AI actions comparison by user department

By Department / By Skills Config	Others	Finance
Code Assist API Definition	1536	0
Case summarization	16	1
mikes new skill	6	0
GAF topics	5	0
Resolution notes generation	5	0

Feedback details

This area of the dashboard shows the numbers and percentage of feedback provided by users of AI actions in the selected date range.

Feedback details indicator

Feedback details

By Skills Family	Number of AI actions (30d running SUM)	Number of AI actions with feedback (30d running SUM)	Percentage of feedback submitted (30d running AVG)	Number of AI actions with positive fee
Others	122	9	117	
Platform	19	0	0	
CSM	10	0	0	
Creator	2	0	0	
Total	**153**	**9**	**78**	

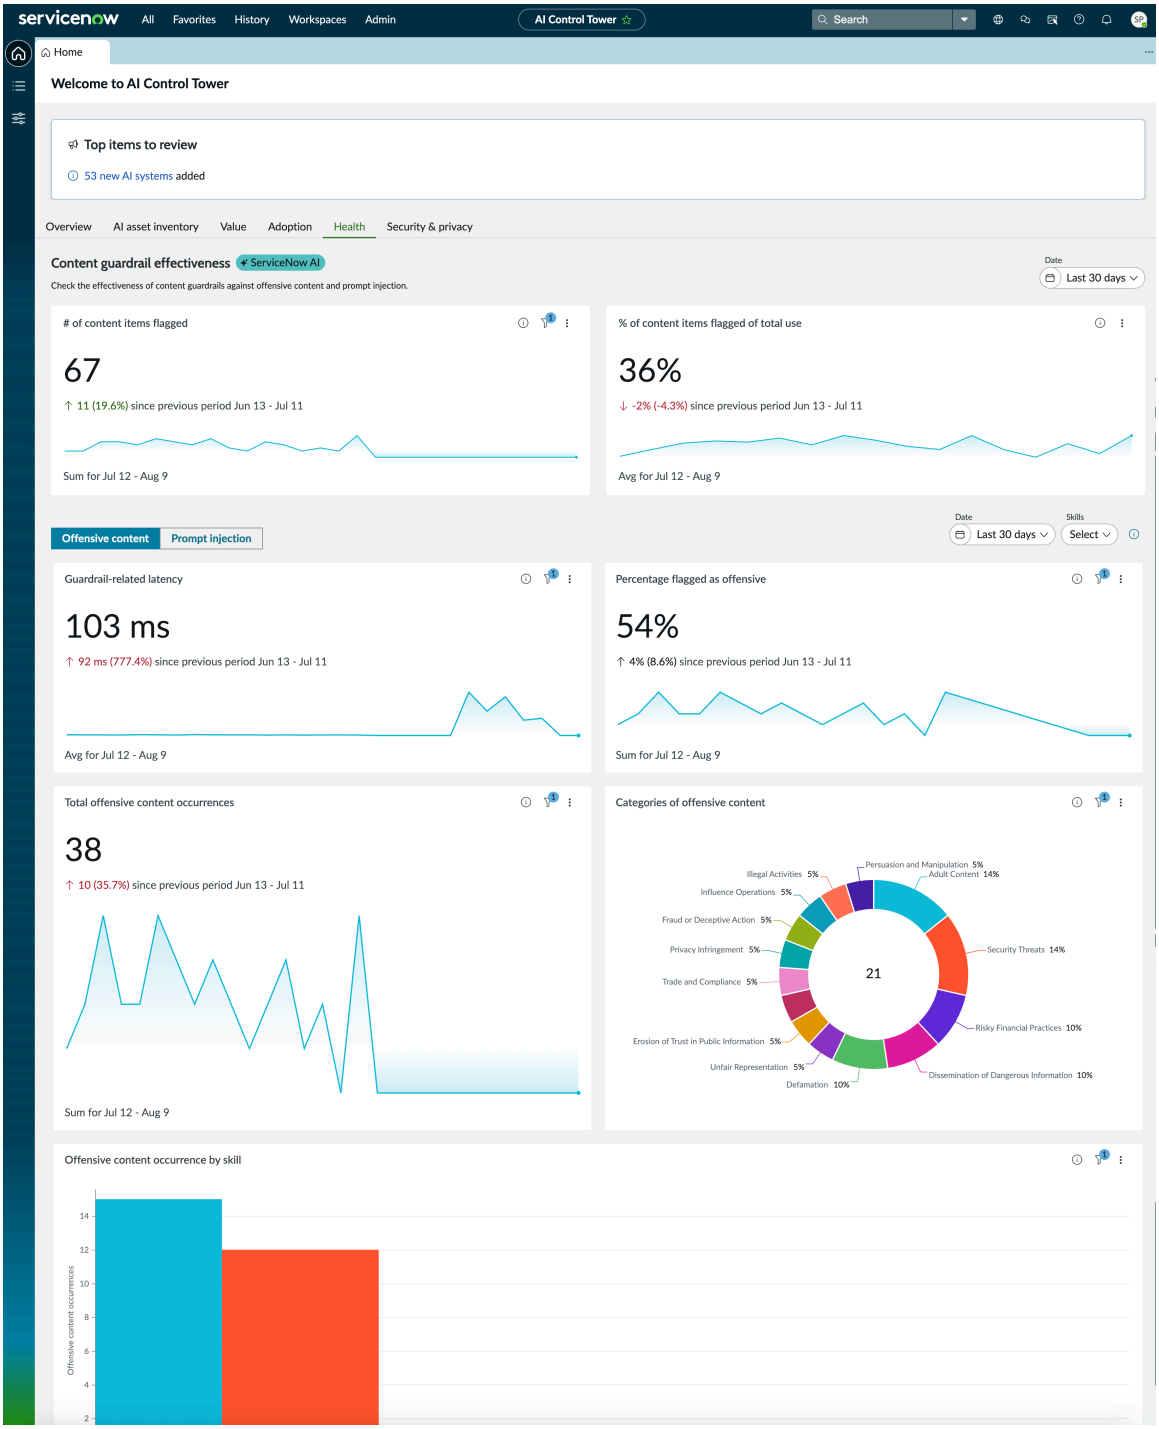
Health tab in AI Control Tower

Monitor the performance of guardrails enabled through Now Assist Guardian.

The Health tab in the AI Control Tower dashboard helps you monitor and evaluate the effectiveness of offensive content and prompt injection guardrails active on your ServiceNow AI assets.

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

Health tab in AI Control Tower



The visualizations on the Health tab provide the following insights.

- Average latency as a result of active offensive content and prompt injection guardrails. High latency could mean increased guardrail activity in the period.
- Count and percentage of offensive content and prompt injection occurrences.
- Skills where offensive content and prompt injection occurrences were detected.

The dashboard does not consider historical data for Health metrics.

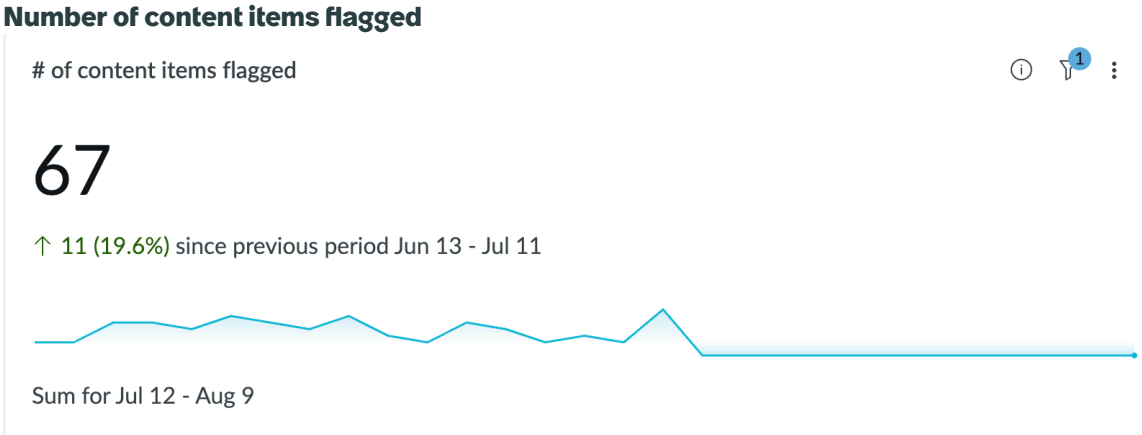
Apply the filters on the dashboard to view guardrail activity for skills in a date range.

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

Content guardrail effectiveness

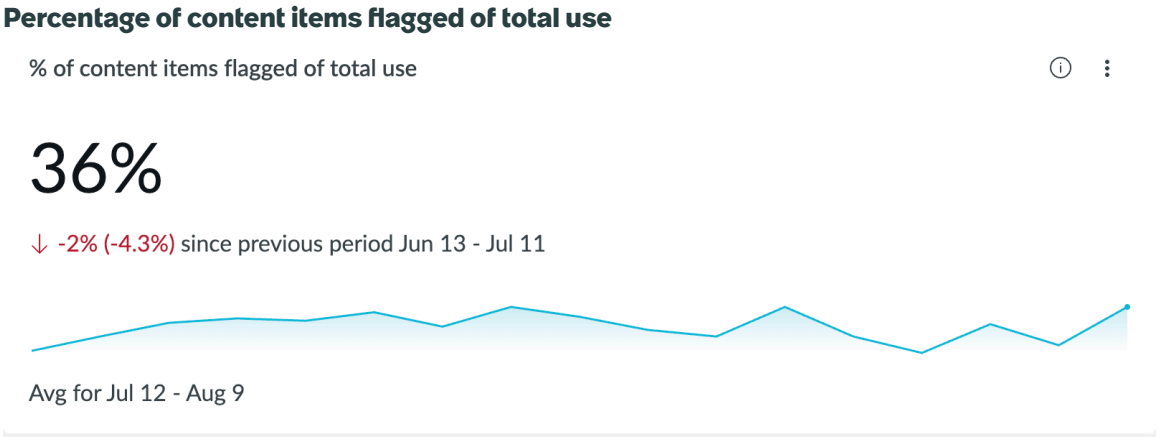
Number of content items flagged

This area of the dashboard shows the number of offensive content and prompt injection occurrences in the selected date range.



Percentage of content items flagged of total use

This area of the dashboard shows the percentage of requests and responses to and from the large language model (LLM) service that are flagged for offensiveness and prompt injection.



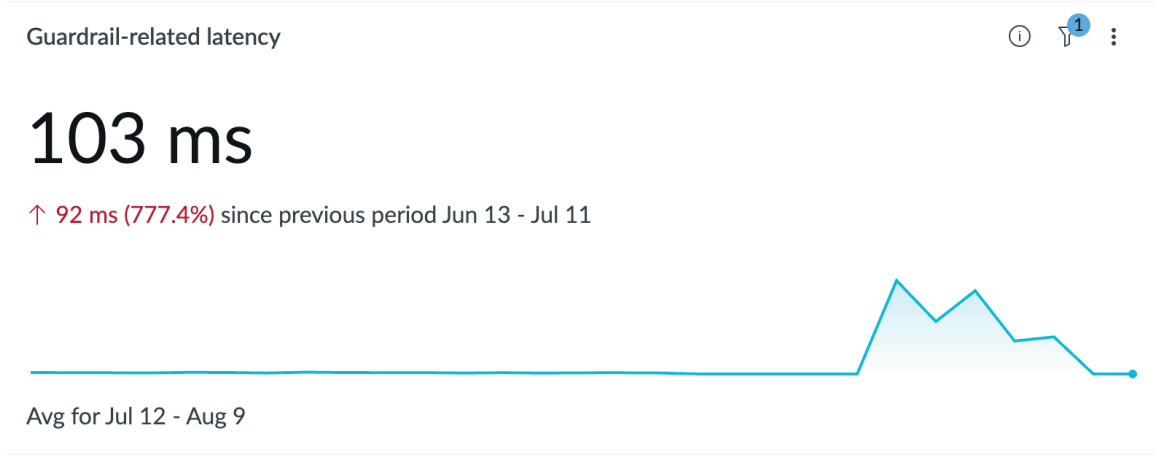
Offensive content visualizations

Guardrail-added latency

This area of the dashboard shows the average latency as a result of the active offensive content guardrail for the selected skills and date range.

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

Guardrail-added latency for offensiveness



Percentage flagged as offensive

This area of the dashboard shows the percentage of requests and responses to and from the large language model (LLM) service that are flagged for offensive content.

Percentage flagged as offensive

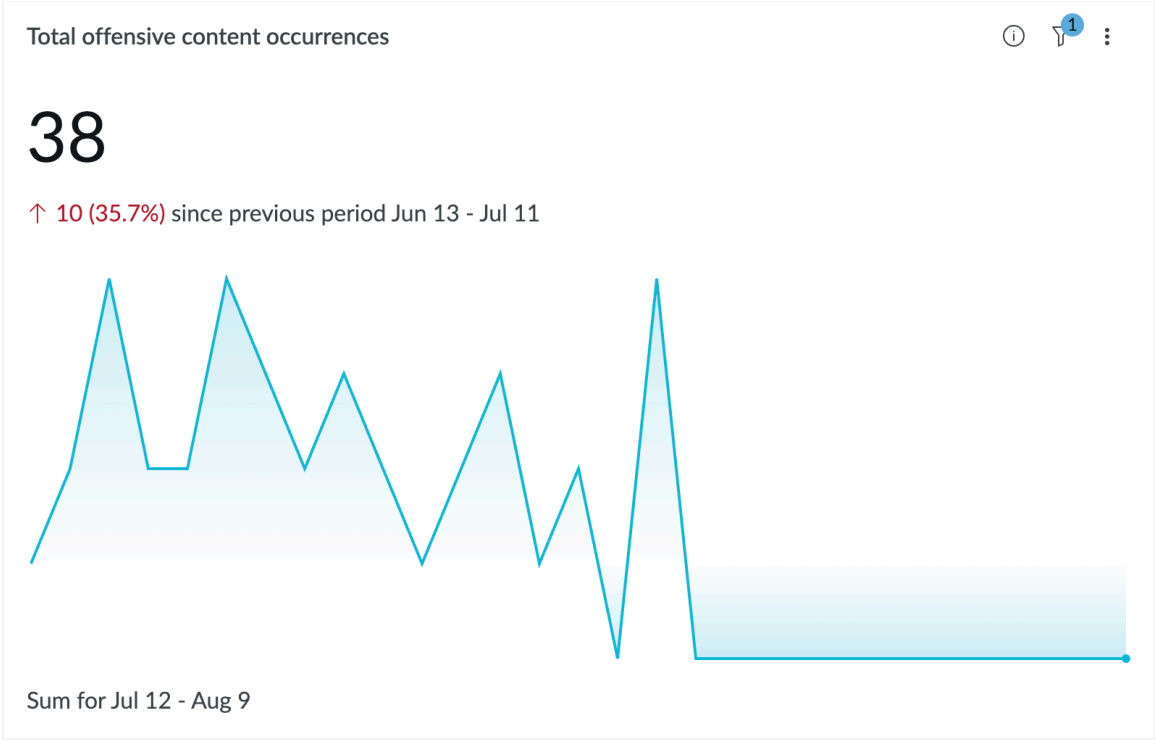


Total offensive content occurrences

This area of the dashboard shows the total number of offensive content occurrences for the selected skills and date range.

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

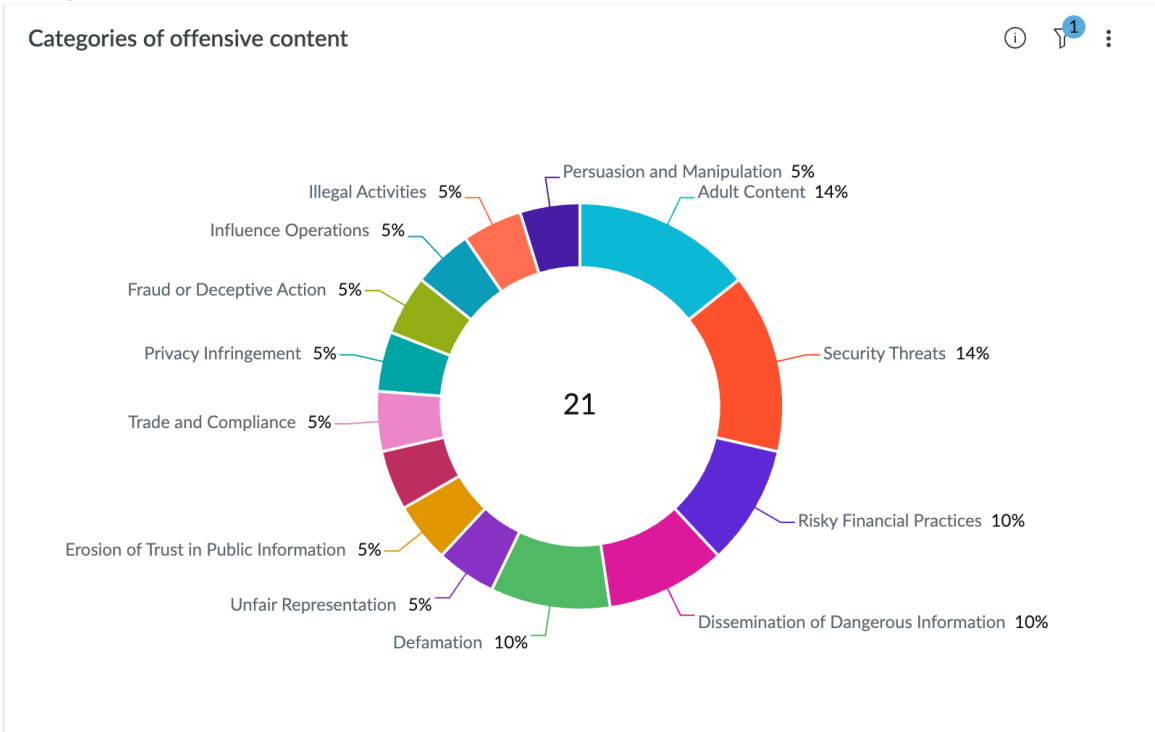
Total offensive content occurrences



Categories of offensive content

This area of the dashboard shows a breakdown of offensive content occurrences by the categories. If content is deemed to be offensive under more than one category, for example, toxic and defamatory, the occurrence is counted individually toward both the categories. For more information on offensive content categories, see Now Assist Guardian.

Categories of offensive content

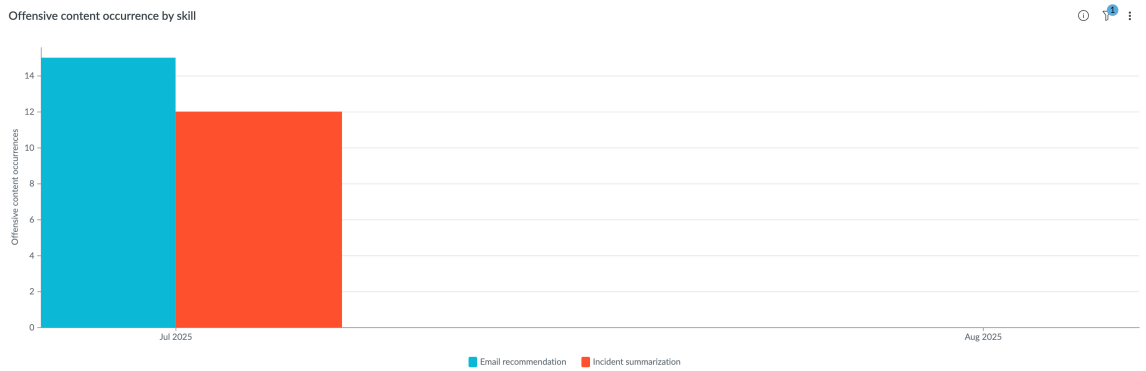


Offensive content occurrences by skill

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

This area of the dashboard shows the number of offensive content occurrences over time by the skills in which the content is detected.

Offensive content occurrences by skill

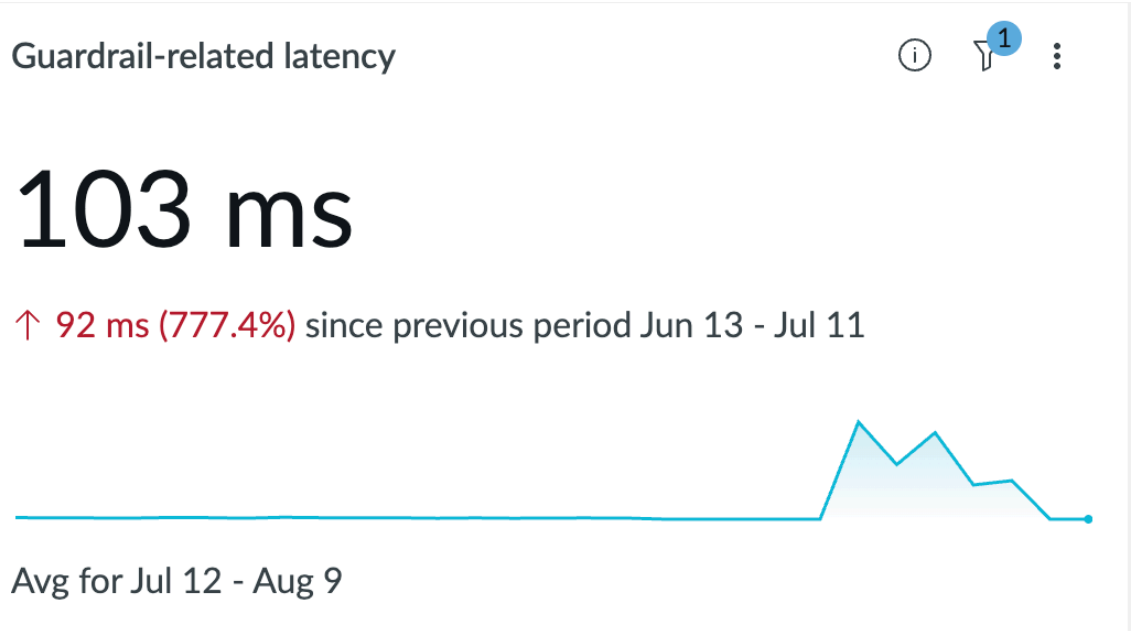


Prompt injection visualizations

Guardrail-added latency

This area of the dashboard shows the average latency as a result of the active prompt injection guardrail for the selected skills and date range.

Guardrail-added latency for prompt injection

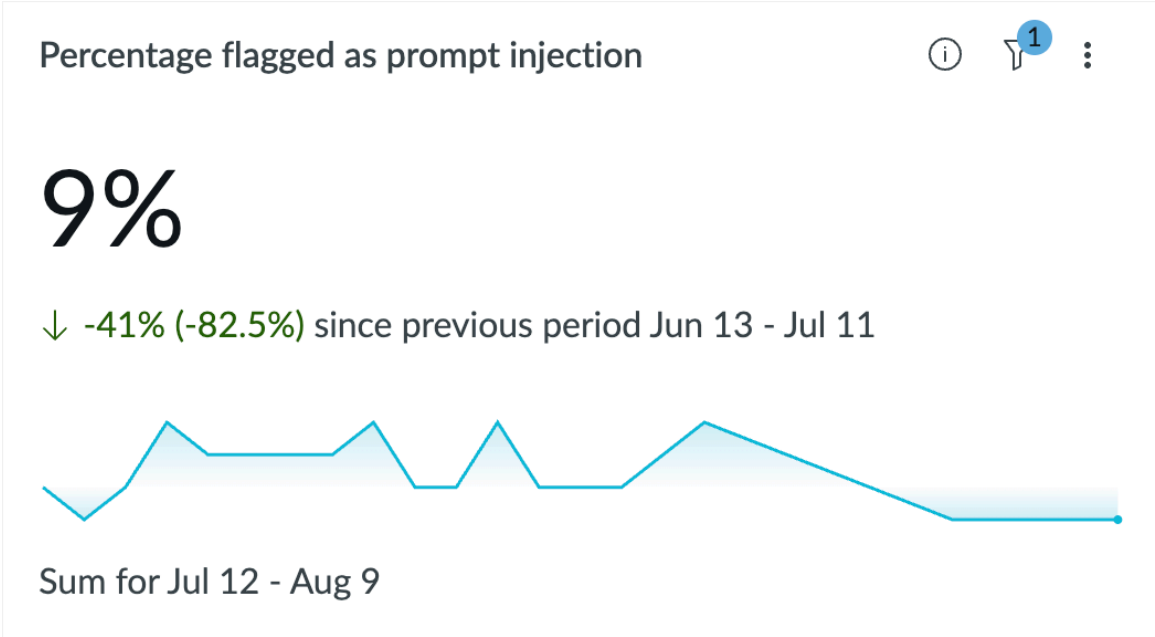


Percentage flagged as prompt injection

This area of the dashboard shows the percentage of requests and responses to and from the LLM service that are flagged for offensive content.

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

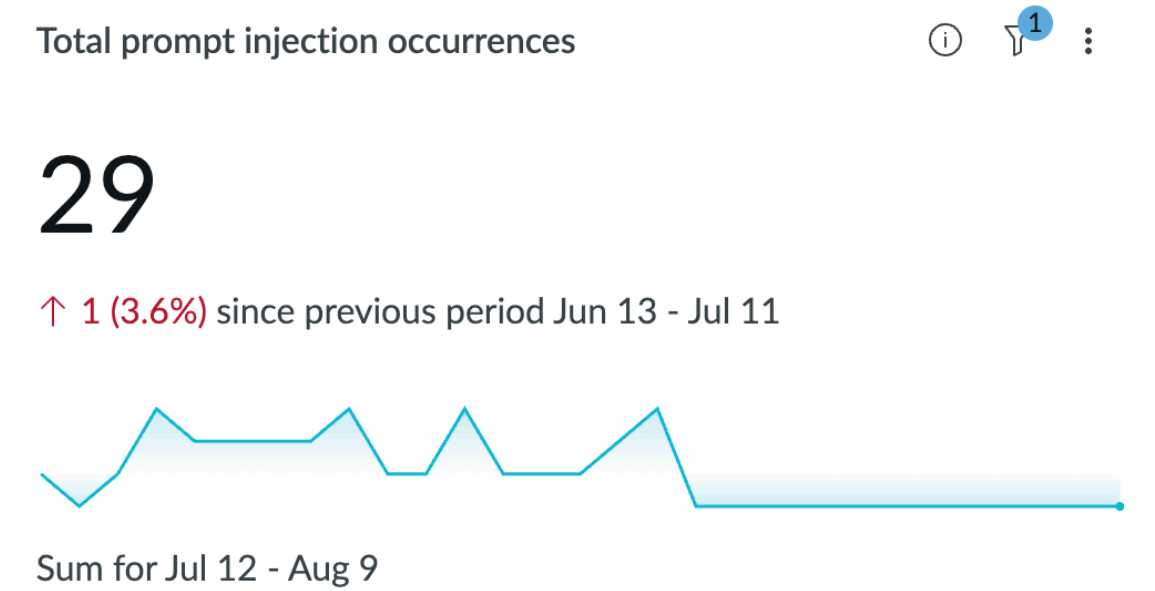
Percentage flagged as prompt injection



Total prompt injection occurrences

This area of the dashboard shows the total number of offensive content occurrences for the selected skills and date range.

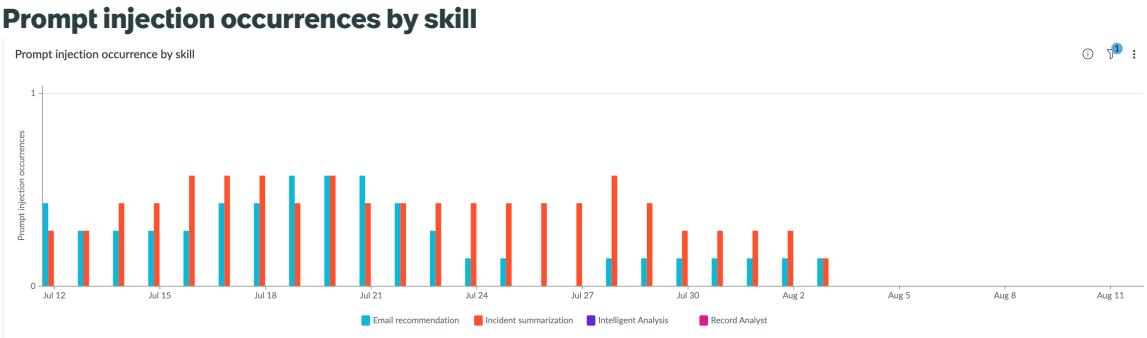
Total prompt injection occurrences



Prompt injection occurrences by skill

This area of the dashboard shows the number of prompt injection occurrences over time by the skills where prompt injection attempts were detected.

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



Risk and compliance tab in AI Control Tower

The Risk and compliance tab on the AI Control Tower displays the risk classification of AI assets and the compliance posture for selected authority documents and policies.

AI assets refer to the various components and resources that are essential for the development, deployment, and operation of artificial intelligence systems. These assets can include:

1. AI systems: The complete software or hardware infrastructure that runs AI algorithms and processes. This can include machine learning platforms, natural language processing systems, and other AI-driven applications.
2. AI models: The mathematical and computational models that are trained on data to perform specific tasks. These models can range from simple linear regression models to complex deep learning neural networks.
3. AI datasets: The collections of data used to train, validate, and test AI models.

For more information, see AI systems, AI models, and Datasets.

Understanding and managing these AI assets is crucial for ensuring that AI systems are effective, reliable, and compliant with regulatory and ethical standards as defined by your organization.

The Risk and compliance dashboard has the following sections. You can drill down into the data on each widget in any section.

Compliance overview

Regulatory risk classification

This section displays the risk classifications of AI systems, AI models, and Datasets using donut charts. The risks are qualitatively classified as **High**, **Low**, **Medium**, and **Unacceptable**. These classifications are based on the risk assessments of the AI assets.

Compliance by authority documents and policies

The section shows compliance based on controls implemented. By default, the compliance scores are displayed for the following frameworks that are provided in the library:

- NIST AI Risk Management Framework: This framework displays the four key associated citations, namely map, measure, manage, and govern. Each citation's compliance score is displayed based on its control attestations.
- EU Artificial Intelligence Act: This framework has multiple chapters that are displayed as citations and child citations. Each citation is mapped to a control objective to provide you with a compliance percentage score.

You can choose to view compliance data by selecting one of two options: **Authority Documents** or **Policies**. Additionally, you can view the overall compliance score

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

percentage, along with the number of compliant and non-compliant authority documents and policies, by using the drop-down filter to select specific authority documents or policies. You can also see all the issues that require immediate attention and AI cases related to each authority document or policy.

The authority documents are provided solely for informational and guidance purposes to assist with the initial setup of AI Risk and Compliance frameworks. It does not constitute legal advice or assurance of regulatory compliance. You are solely responsible for ensuring that all use of the content complies with applicable laws, regulations, directives, and industry standards in their jurisdictions.

Note: You can configure which authority documents and policies you want to display on the home page. For more information, see Set up properties for compliance posture.

Risk overview

This section monitors and tracks the risk posture of the AI assets in your organization. Using the **AI asset** filter, you can filter risk posture insights by the type of AI asset inventory.

AI systems by aggregated risk score

This section displays the classifications of AI systems by aggregated risk score using a donut chart. The risk scores are qualitatively classified as **High** and **Low**.

Risk heatmap

The Risk heatmap widget displays the visualization of all identified risks within the AI assets. By default, the Residual risk filter is applied, but you can filter the heatmap based on Inherent risk. The heatmap is segmented, and the segmentation changes based on the filter. The activities fall under the respective combination of risk and control effectiveness, or impact and likelihood. The combination is based on the selected risk classification filter. You can filter the risk heatmap by Risk Assessment Methodology (RAM)when more than one methodology is available.

Regulatory landscape overview

You need to install GRC: Regulatory Change Management application to see this section. For more information, see Installing Regulatory Change Management.

Overview

- Alerts

  This section displays the risk classifications of AI systems, AI models,and AI datasets using donut charts.

- Change tasks

  This section displays the distribution of regulatory change tasks by workflow state or lifecycle phase using a donut chart.

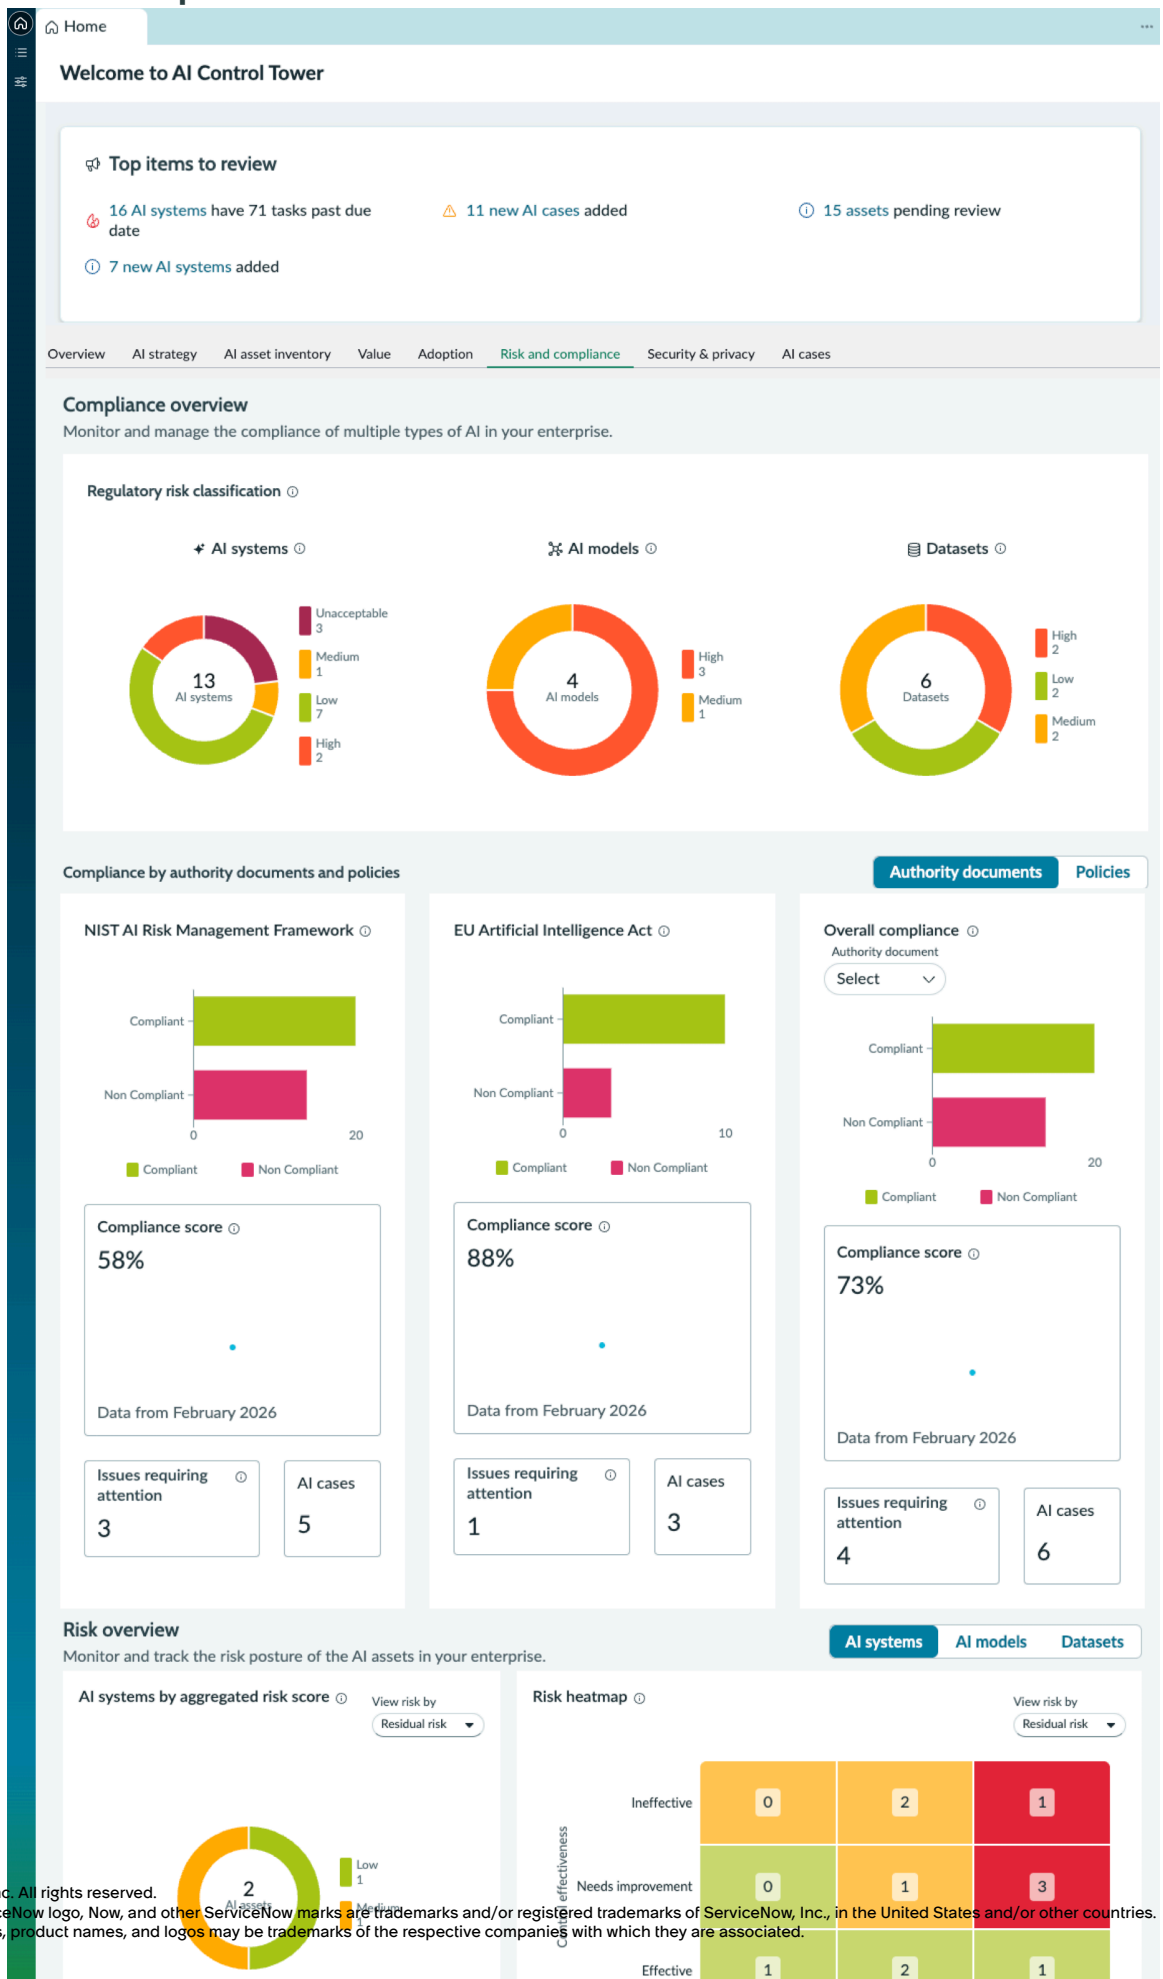
Assessments

This section displays the regulatory impact assessments and risk assessments linked to regulatory changes using a donut chart. By default, regulatory assessments filter is applied, but you can change it to risk assessments.

The following image shows the Risk and compliance dashboard.

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

Risk and compliance dashboard



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

For more information, see AI Risk and Compliance documentation.

AI cases tab in AI Control Tower

Track, monitor, and analyze your AI case workflows, identify your workflow bottlenecks, and check your accountability of your AI-related risks by using the AI risk and compliance dashboard. As an AI steward, you can also use the dashboard to track the status and trends of your AI-related inquiries.

AI cases

The **AI cases** tab on the AI risk and compliance dashboard provides the key metrics, trends, and actionable tasks that you can use to track your AI risks, case activities, compliance tracking, and governance enforcement. You can get case-level insights, task-level tracking, and accountability for your responsibilities and progress. As an AI steward, you can report a case when a violation is found. An AI analyst also follows the case and comes up with a resolution.

You can drill down into the data on each of the following widgets to learn more about an AI case.

AI cases status by state

Shows distribution of reported AI Cases by lifecycle state or workflow state.

AI cases status by priority

AI cases based on their priority: Critical, High, Moderate, Low, and Planning.

AI cases

Summarized list of AI cases reported:

- All: All the AI cases
- Overdue: AI cases that are past the due date.
- Due in 7 days: AI cases due in the next 7 days.
- Unassigned: AI cases that are not assigned to anyone.

Attributes of each case, such as Name, Sub-type, Case analyst, Priority, and State are available in a list view. Select each card to view that category of AI cases in a list view.

Trends

Visual patterns and historical data that are associated with the AI cases:

- AI cases by subtype widget: Trends by different categories of the AI case subtypes. You can group data by business unit, location, or primary cause to get contextual insight into the issues or AI cases patterns across the departments or systems. You can filter by AI assets.

- Open and closed cases widget: Historical trends of the case resolutions over a specified time frame. You can see the historical trends in the last 12 months. The historical trends help stakeholders to analyze the volume fluctuations and how efficiently the cases were closed.

  For example, in the last 12 months, consider that 4,200 cases were created and 3,960 cases were closed. The average time of case closure was 2.8 days, the peak case volume was in January 2025 (510 cases were created), and the fastest case resolution month was March 2025. You can see that over the past 12 months, the case closure efficiency has steadily improved, with a 12% decrease in the average resolution time compared to the previous year.

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

Tracking

AI case progress at the task and issue level:

- Action tasks
  - Number of tasks in the Overdue, Due in 7 days, or Unassigned states.
  - Task distribution by the status, such as Draft or Work in Progress.
- Issues
  - Counts of issues that are linked to AI cases by their states like New, Review, or Responds.
  - Counts of issues that are linked to AI cases by additional details like Overdue, Due in 7 days, or Unassigned.

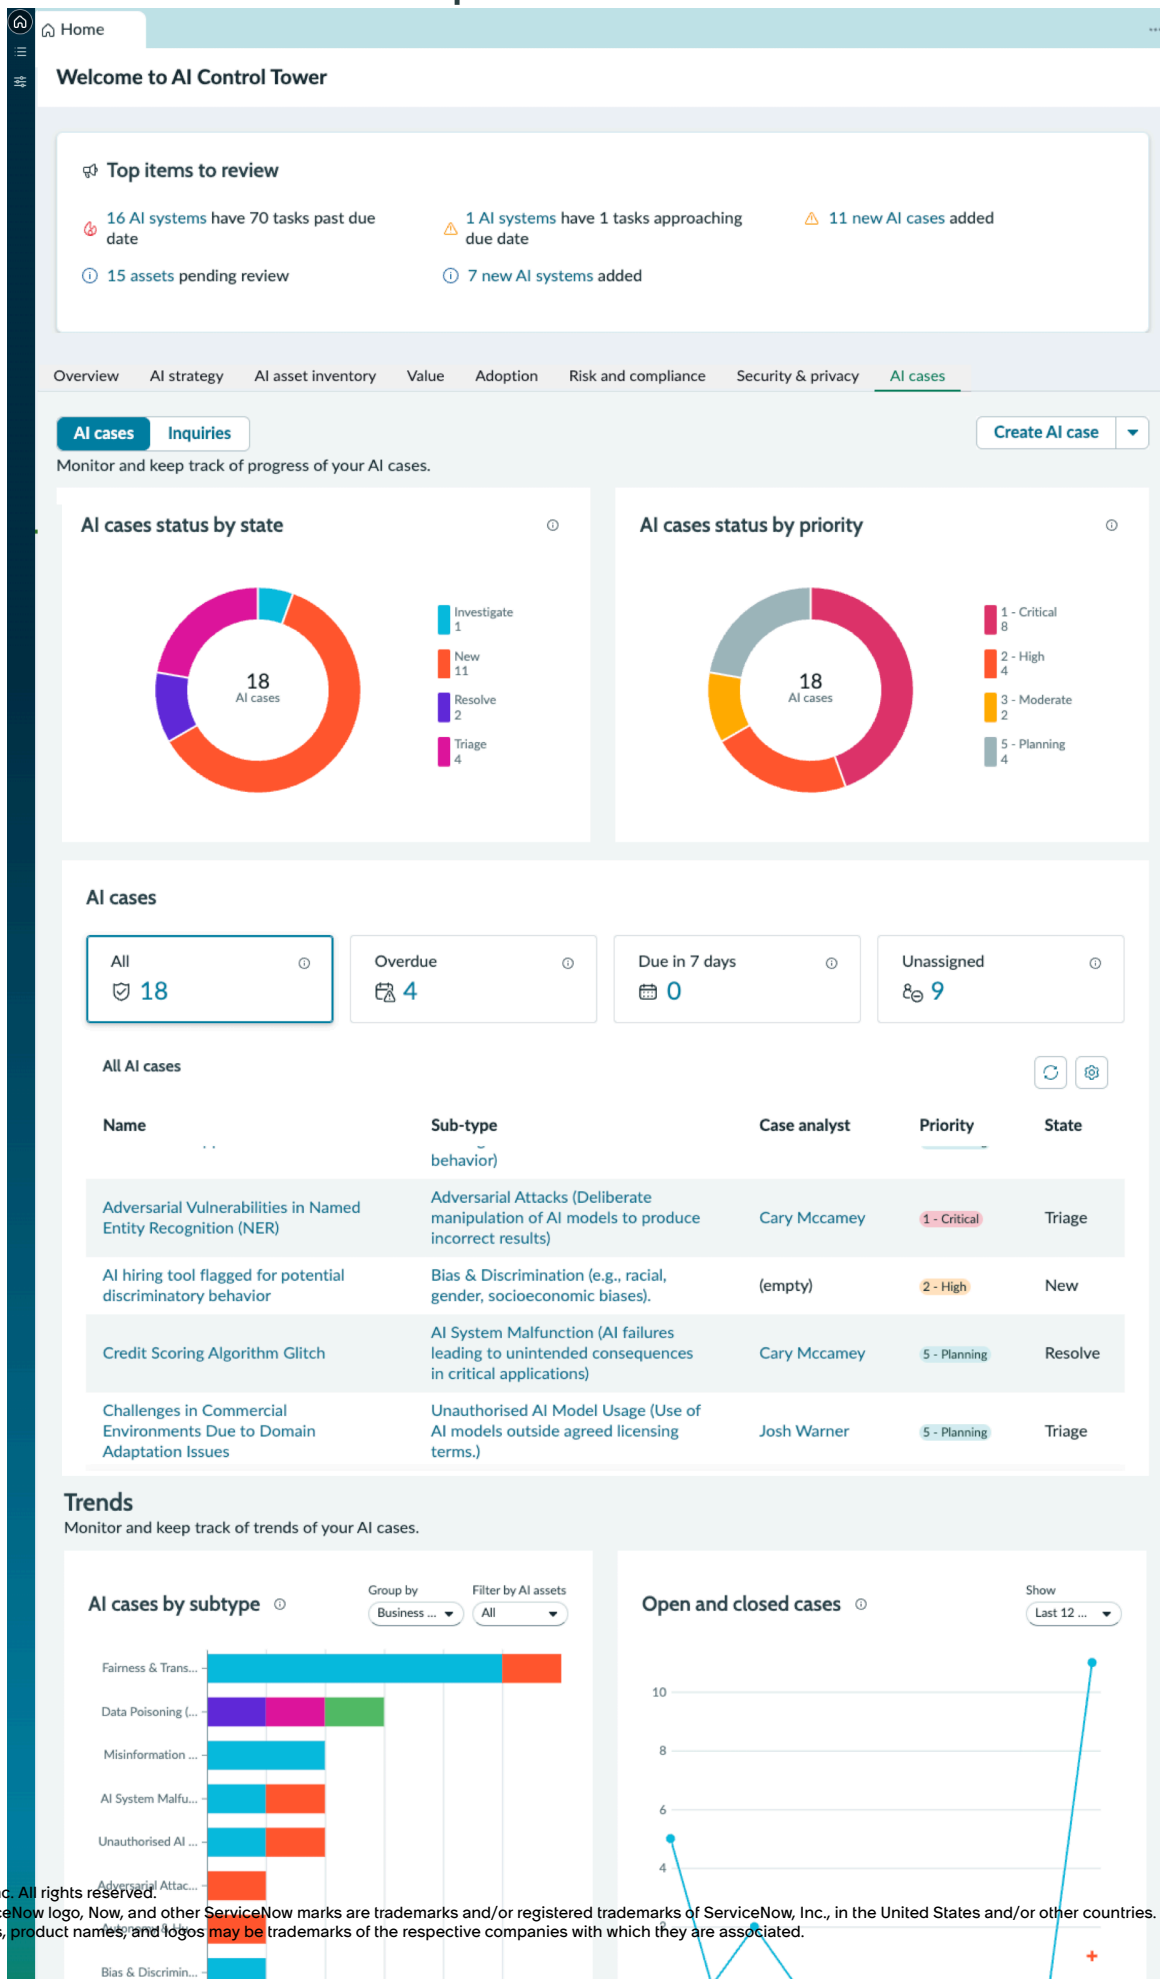
You can create AI cases and issues from the AI cases tab in both the AI Control Tower Workspace and the AI Risk and Compliance Workspace. The Create AI case button lets you initiate a new AI case record directly from the monitoring dashboard. Use this option to document, investigate, or track issues, risks, or compliance concerns related to AI models or datasets. Select the dropdown arrow next to Create AI case, and then select Create issue to create an issue related to an existing AI case.

For more information, see Create an AI case in the AI Risk and Compliance workspace and Create an AI issue in the AI Risk and Compliance workspace.

The following example shows the **AI cases** tab on the AI risk and compliance dashboard.

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

AI cases tab on the AI risk and compliance dashboard



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

Inquiries

The **Inquiries** tab on the AI risk and compliance dashboard provides you with a structured view of the AI inquiry workflow. Use the following widgets on the dashboard to track the compliance and then to monitor the status and progress of an inquiry:

You can drill down into the data on each of the following widgets to learn more about an AI inquiry.

Inquiries status by state

Counts of inquiries by state. New, Triage, In progress, and Approved inquiries are represented in the donut chart.

Inquiries status by priority

Counts of inquiries by priority. An inquiry can have Critical, High, Moderate, Low, or Planning priorities.

Inquiries

Summarized list of all AI inquiries. It's filterable based on parameters such as status, priority, and assignment.

The essential attributes of inquiries are based on the operational relevance or urgency. Each inquiry entry includes metadata such as:

- Name
- Requested by
- Assigned to
- State
- Priority

Trends

Trend-based insights into your inquiry activity across the organization. These visualizations support longitudinal analysis to improve planning and compliance responsiveness.

Note: Longitudinal analysis refers to the process of analyzing data points across departments, teams, or systems over a period of time.

Filters, such as the AI assets and the time range, help you to refine the analytics.

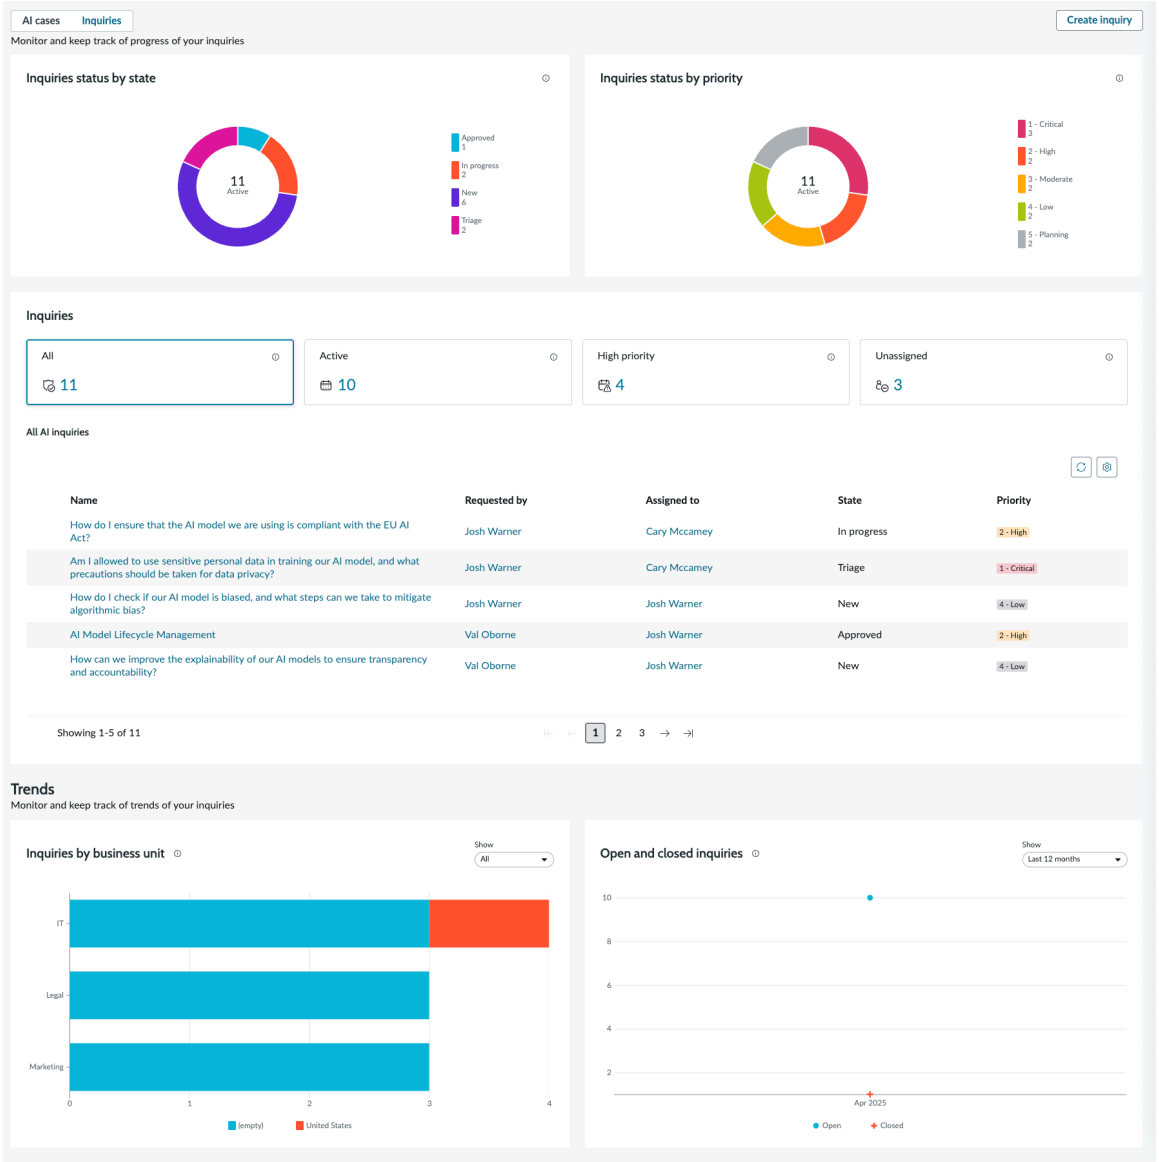
- Inquiries by business unit: You can see the distribution of inquiries across the various departments or functional areas of your organization. This distribution helps you to identify where most AI-related concerns originate so that you can plan on more targeted risk mitigation or awareness training.
- Open and closed inquiries: You can track the volume of inquiries that are initiated and resolved within a selected time range. This tracking enables you to track how your workflow is performing and to plan for future capacity.

Use the **Create inquiry** button to initiate a new inquiry. You can make timely reports of AI-related concerns or clarification requests.

The following example shows the **Inquiries** tab on the AI risk and compliance dashboard.

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

Inquiries tab on the AI risk and compliance dashboard



For more information, see AI Risk and Compliance documentation.

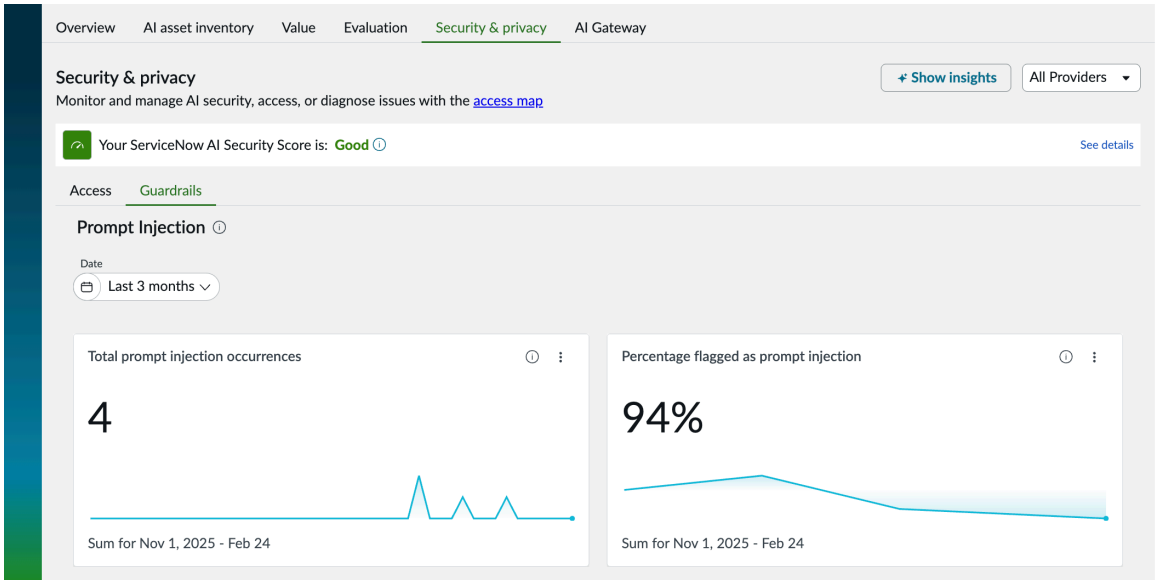
Security & Privacy tab in AI Control Tower

Review AI security metrics such as access issues, dormant and privileged AI agents, and map the relationships of your ServiceNow agents, agentic workflows, and tools.

The **Security & Privacy** tab of AI Control Tower offers a dashboard-based overview of your AI security metrics. The dashboard contains several visualizations detailing AI security metrics. In addition to tracking metrics the **Security & Privacy** tab contains the access map, a tool that gives an node-graph visualization of the relationships between your ServiceNow agents, agentic workflows, and tools. You can use the map to investigate the relationships between your AI agents and workflows further.

Note: You can drill down into the data on each widget by selecting the chart.

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



ServiceNow AI Insights

ServiceNow AI Insights require that the Now Assist AICT Security Posture Summarizer skill is enabled. For more details, see Activate a Now Assist skill.

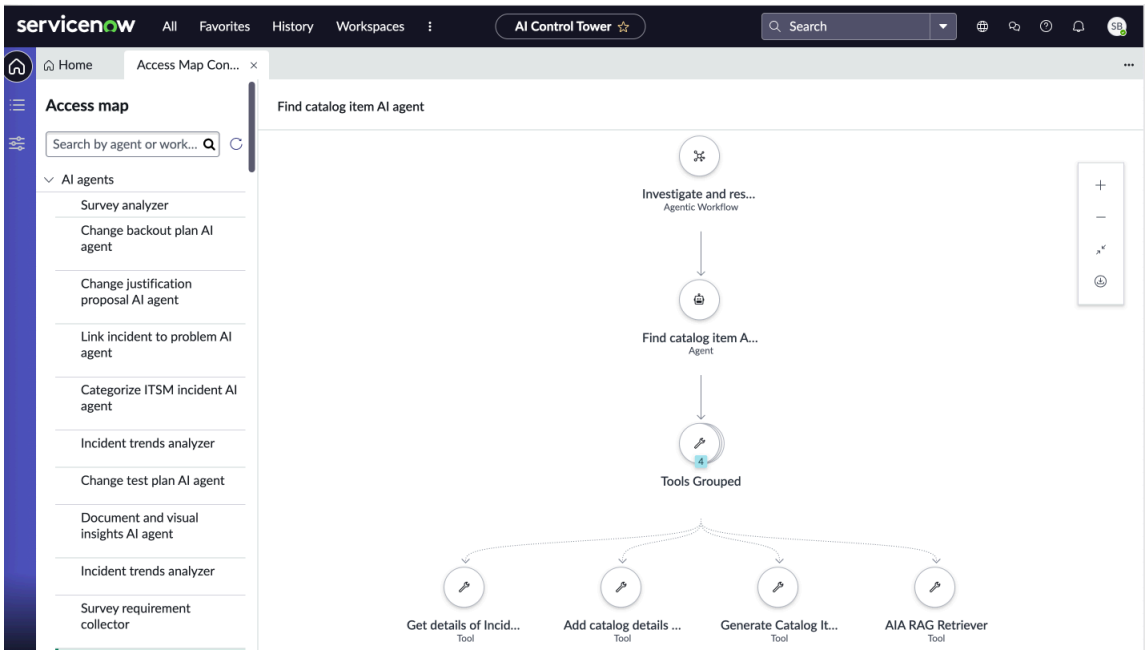
AI Control Tower AI insights summarize positives and potential issues to remediate to improve the overall security posture of your instance.

- Positives: Enabled settings and features that improve your security posture.
- Areas for Attention: Low- to medium-risk items to resolve.
- High Impact Observations: High-risk items to resolve.
- Actions: Suggested action items to address Areas for Attention and High-Impact Observations.

Access map

The **Access map** displays a node map detailing the relationships of your ServiceNow® agents, agentic workflows, and tools. You can use the map to review these relationships, configure agent details, and resolve access issues. The map includes filters for both agents and agentic workflows. You can open the access map by either navigating to **All** > **AI Security and Privacy** > **Access Map**, or selecting the link in the dashboard. See Using the access map to learn how to use access map.

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



If a warning icon appears on any agent, that agent has access issues. Select the warning icon to see details such as the workflow, agent, and tool associated with the access issue.

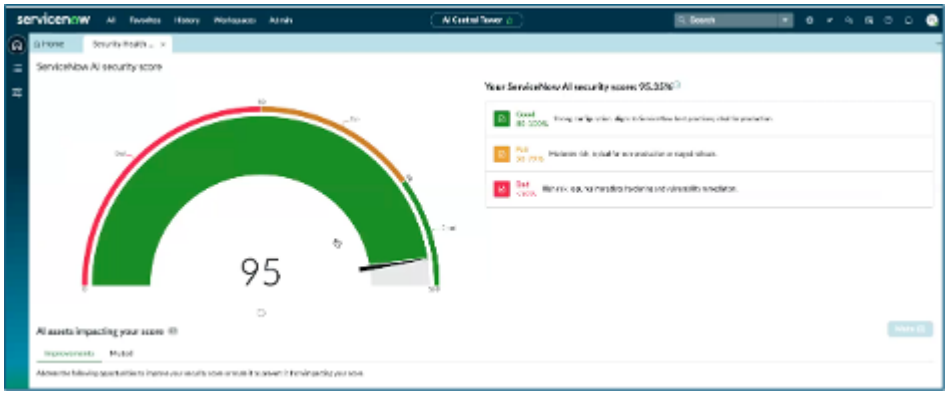
In Access issues, the User ID is the ID of the user who ran the agent.



ServiceNow AI Asset Security Score

The ServiceNow AI asset security score is a measure of the health of your ServiceNow AI assets in terms of access issues, privileged AI agents, and dormant AI systems.

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



AI assets impacting your score

To see more information about your score, select **See details** in the **Security & Privacy** tab. A list view shows the ServiceNow AI assets that are included in your AI asset security score calculation. Your score is the average of all managed ServiceNow AI assets listed.

You can exclude an asset from your score by selecting a row and selecting **Mute**. For example, you can mute an AI asset if you determine that remediating the asset's issue would be a risky change. You can also configure the score to remove security categories from the score or change the weights of security categories. For more information, see Data sharing, Data overflow processing, and Security & privacy in AI Control Tower.

AI asset details

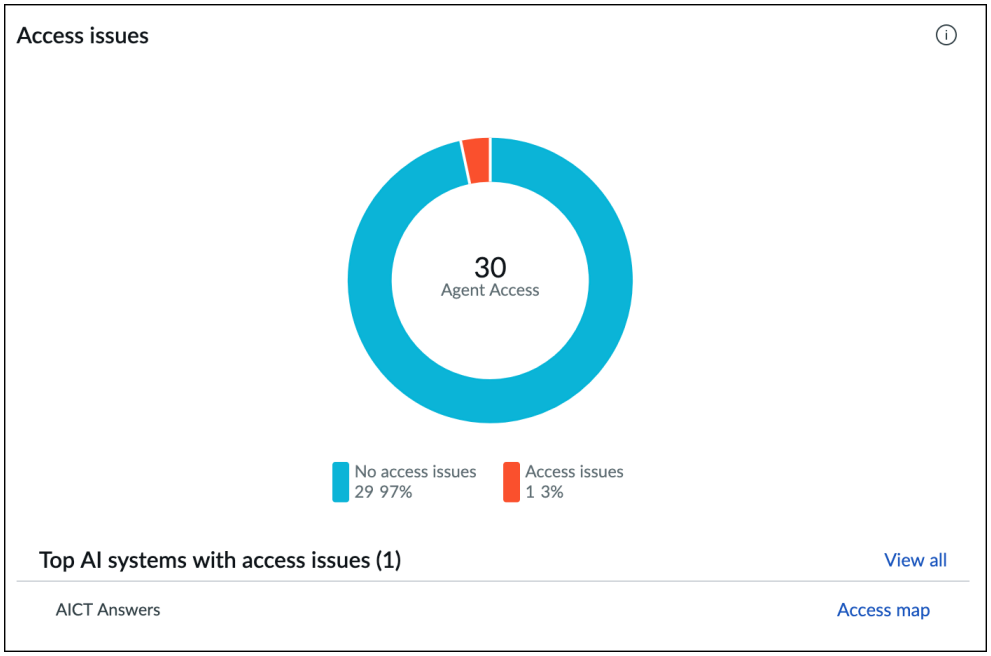
Column	Description
AI system	Name of the ServiceNow asset.
Category	Type of issue, such as dormant AI system, privileged AI agent, or access issue.
Provider	This value is always ServiceNow.
Score impact	The percentage impact to your AI asset security score.
Date	Date the issue occurred.
AI Task	The AI security task to remediate the issue, if applicable.

Access

Access issues

The Access issues chart displays the proportion of AI agents experiencing access-related issues and lists the top AI systems with access issues. AI agents with access issues may be unable to complete their workflows due to the access issue. Hover over a portion of the chart to see the exact proportion and count of agents.

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



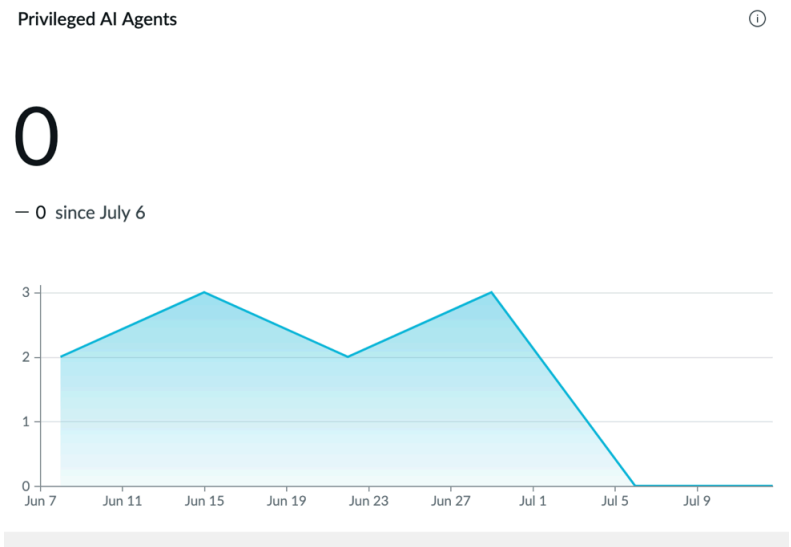
You can create AI security tasks directly from the list view by selecting **Create AI task**. See all active AI security tasks in **AI assets** in the AI Task section. Access to this section requires the sn_vsc.task_manager role.

Resolved AI security tasks that are over 180 days old are archived. Archival days can be configured in system properties.

Privileged AI Agents

The area chart shows AI agents with elevated permissions such as an agent with admin or security admin permissions that can perform critical actions. Some workflows require AI agents have elevated permissions to complete. Hover over a portion of the chart to see the exact number of privileged agents on that day.

To show AWS agent metrics, filter the metrics by selecting **AWS Bedrock** in the provider drop-down list. You must have an AWS account configured for your instance and the Now Assist AiSP AWS IAM Privileged Policy Checker skill enabled. For more details, see AI connections and Activate a Now Assist skill.



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

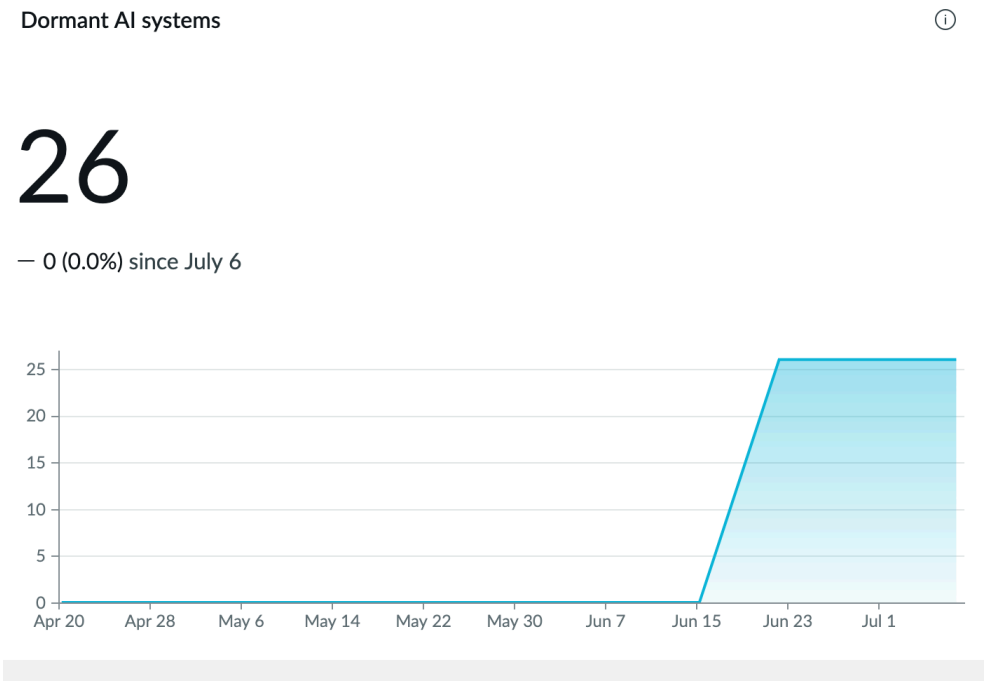
You can create AI security tasks directly from the list view by selecting **Create AI task**. See all active AI security tasks in **AI assets** in the AI Task section. (Role required: sn_vsc.task_manager.)

Resolved AI security tasks that are over 180 days old are archived. Archival days can be configured in system properties.

Dormant AI systems

The area chart shows AI agents that have not been active for over 90 days. Review dormant AI agent permissions to reduce security risk. Hover over a portion of the chart to see the exact number of dormant AI systems for that day.

To show AWS agent metrics, filter the metrics by selecting **AWS Bedrock** in the provider drop-down list. You must have an AWS account configured for your instance. For more details, see AI connections.



When an AI agent becomes dormant, an AI security task is created automatically to streamline your workflow, and quickly resolve issues. The AI security task is assigned to the agent's owner. See all active AI security tasks in **AI assets** in the AI Task section.

Resolved AI security tasks that are over 180 days old are archived. Archival days can be configured in system properties.

ServiceNow instance access to MCP servers

MCP server access metrics include all MCP client-server interactions routed through this instance's AI Gateway. Interactions that bypass the AI Gateway or are routed through another instance's AI Gateway aren't included.

Clients connecting to MCP servers

The Clients connecting to MCP servers chart shows the top 10 clients (ServiceNow AI agents or registered third-party MCP clients) connecting to MCP servers through this instance's AI Gateway. MCP server access metrics are captured for all client-server interactions routed through the AI Gateway. To see more clients, select the chart to drill down into the data.

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

Authorized access attempts to MCP servers

The Authorized access attempts to MCP servers chart shows successful access attempts from MCP clients to MCP servers through this instance's AI Gateway. Clients include ServiceNow AI agents and registered third-party MCP clients. To see more attempts, select the chart to drill down into the data.

Failed access attempts to MCP servers

The Failed access attempts to MCP servers chart shows unsuccessful access attempts from MCP clients to MCP servers through this instance's AI Gateway. Clients include ServiceNow AI agents and registered third-party MCP clients. To see more data, select the chart to drill down into the data.

Guardrails

Prompt injection

These charts show prompt injection data provided by Now Assist Guardian. To see data, enable Now Assist Guardian for your instance. For more details, see Now Assist Guardian analytics.

Offensive content

These charts show offensive content data provided by Now Assist Guardian. To see data, enable Now Assist Guardian for your instance. For more details, see Now Assist Guardian analytics.

Sensitive data

The Sensitive data detected chart shows sensitive data that was identified in user responses to Now Assist prompts. Exposure of sensitive data is limited to the LLM in your instance.

The Sensitive data anonymized chart shows prompt data that met configured data patterns. This data was anonymized based on the configuration for the pattern in Configuration Data Patterns in Data Privacy.

Security policy violations

The Security policy violations chart shows violations of ServiceNow security policies detected in LLM output. The data is collected by scanning generative AI logs to detect policy violations in LLM responses. ServiceNow security policies are based on industry best practices and can't be modified. For example, one policy prohibits linking to executable files or malware, directing users to run unverified downloads or macros, or asking users to disable their antivirus protection. For more information about policies and how to configure data for this chart, see Data sharing, Data overflow processing, and Security & privacy in AI Control Tower.

Agent goal hijack

The Agent goal hijack chart shows when AI agents deviate from their intended role or objective. For example, unauthorized actions or prompt injection attempts. The data is collected by scanning and analyzing agent execution history. For information on how to configure data for this chart, see Data sharing, Data overflow processing, and Security & privacy in AI Control Tower.

AI agent output with PII detected

The AI agent output with PII detected chart shows when agents' LLM output contains personally identifiable information (PII). The data is collected by scanning LLM output for default PII sensitive data patterns specified in Data Privacy and additional PII patterns. For example, U.S. phone number, credit card number, or U.S. passport ID. For information on how to configure data for this chart, see Data sharing, Data overflow processing, and Security & privacy in AI Control Tower.

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

High-risk AI agent output

The High-risk AI agent output chart shows when agents' LLM output contains security-vulnerable patterns. The data is collected by scanning LLM output for known vulnerable patterns and potential corresponding attack vectors. For example, HTML tags shouldn't have scripts associated with them for cross-site script attacks (XSS), or no stacked SQL queries that could result in SQL injection attacks. For information on how to configure data for this chart, see Data sharing, Data overflow processing, and Security & privacy in AI Control Tower.

AI Gateway tab in AI Control Tower

Explore the AI Gateway tab in the AI Control Tower.

AI Gateway overview

Prerequisites

Role required: AI steward [sn_ai_governance.ai_steward]

The AI Gateway tab shows the metrics at the MCP (Model Context Protocol) server level, listing all connected MCP servers along with the total number of transactions for each server and its success rate.

The AI Gateway displays a comprehensive view of all MCP servers categorized into:

- Name
- Total transactions
- Success rate

Total transactions are all tool calls, calculated by MCP method equal, which is ("tools/call").

The success rate is the number of such transactions that have an HTTP status from 200 through 299 divided by the total.

AI assets

An AI asset refers to a digital tool or resource that uses artificial intelligence to carry out particular tasks or address problems. Examples include machine learning models, chatbots, natural language processing systems, image recognition software, and similar technologies.

Details

For information on AI inventory, see AI asset inventory

For information on AI asset Lifecycle, see AI asset lifecycle

For information on approval playbook workflow, see Create an AI Control Tower Playbook workflow.

For information on creating Now Assist approvals task, see Create a Now Assist approval task.

For information on AI security tasks, see AI Task.

AI asset inventory

AI asset inventory includes all the AI assets such as AI models, AI systems, prompts, datasets, and MCP servers that an organization can use.

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

AI asset inventory view

AI systems

An AI system can make predictions, recommendations, or decisions that impact physical or virtual environments, with outputs or behaviors that aren't always predetermined by the developer or user. These systems are often trained using large amounts of structured or unstructured data and may function with different levels of autonomy, or none at all, to fulfill objectives set by humans.

AI models

AI models consist of functions and algorithms that are trained and evaluated on datasets to generate predictions or take decisions.

Prompts

Any data, including text or images, that is used to guide, instruct, or influence an AI model's output. Prompts can also refer to an element in a process that requires user input in a field before moving to the output screen.

Datasets

A dataset usually consists of records and fields organized into rows and columns. These datasets are stored within a database table.

MCP Server

Provides external functionality by exposing tools for AI agents to use.

AI assets are presented in a list format and organized by display name, provider, vendor, managed by, lifecycle phase, state, status, and risk classification.

For information on AI asset record details, see View AI asset details

For information on creating AI assets, see Creating AI assets

For information on process flow of the MCP servers, see Process flow of MCP servers Via AI Gateway

Assets list- managed and unmanaged assets

Learn about the AI assets list, managed and unmanaged assets.

Assets list

AI assets are presented in a list format and organized by display name, provider, vendor, managed by, lifecycle phase, state, status, and risk classification.

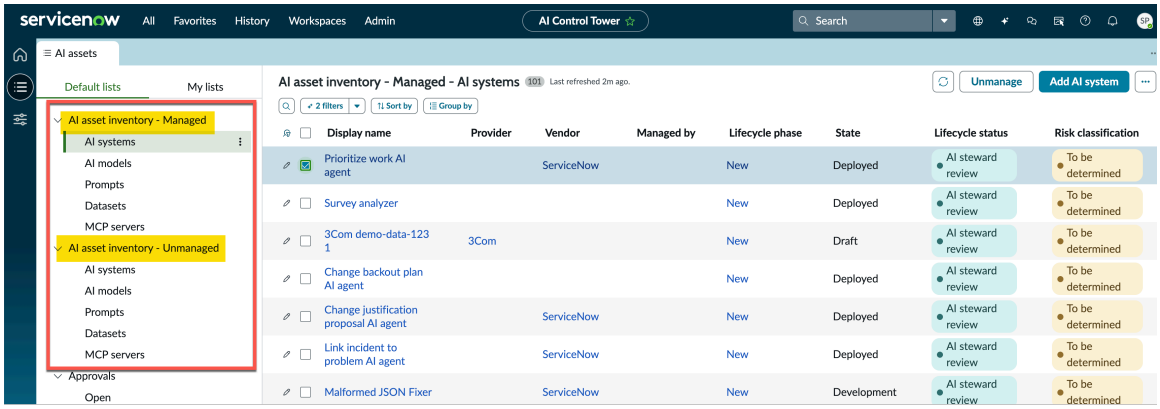
Managed & Unmanaged assets

Starting with the March (2026) release, you can designate assets as managed or unmanaged directly from the list. By default, all the assets in AI Control Tower are under unmanaged assets.

When an unmanaged asset is marked as managed, it gains access to AI Control Tower capabilities such as governance, lifecycle management, value assessment, risk classification, security, and privacy. If a managed asset is switched back to unmanaged, it loses all the before mentioned AI Control Tower capabilities.

Using Automation rules, AI assets are set as managed. For information on Automation rules page, see Automation rules

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



Note: Initiating a steward review for an unmanaged asset triggers the lifecycle process and automatically transitions the asset to a managed state.

View AI asset details

View the AI asset governance details to track the lifecycle's status, phase, and the install status.

Before you begin

Role required: AI steward

Procedure

1. Navigate to **Workspaces** > **AI Control Tower**.
2. From the AI Control Tower, open the AI assets view.
3. From the navigation menu of the AI assets view, under AI asset inventory, select an AI asset. AI asset details page is displayed.
4. On the page, review the AI asset details.

Field	Description
Version	Version number of the asset
Asset type	Type of assets: ○ Generative AI ○ Agentic AI ○ Classical AI ○ AI model ○ AI system ○ AI prompts ○ AI datasets
Provider	who has built the asset
Vendor	Who has sold the Asset
Department	Department where the asset is allocated
Managed by	Managed by the user who owns the asset
License details	License details of the asset

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

Field	Description
Supported locations	Shows the locations supported for the AI as set.
Lifecycle phase	Shows the asset's lifecycle phase. ○ New ○ Assess ○ Build and test ○ Deploy ○ Offboarding
State	The state of the asset. ○ Deployed ○ Retired ○ Development ○ Unknown ○ N/A
Lifecycle status	The lifecycle status of the asset. ○ In review ○ Approved ○ Rejected ○ Deployed ○ AI steward review ○ Approved for development ○ Ready for deployment ○ Approved for deployment ○ Canceled
Risk classification	Risk classification of the asset
Created	Creation date
Updated	Updated date

AI asset lifecycle

The AI asset life cycle defines the series of stages that you must follow to manage an AI system, AI model, prompt, or dataset throughout its useful life.

AI asset lifecycle stages

The AI asset lifecycle consists of the following stages:

Onboard

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

The onboard stage is the introduction of an AI asset into your organization. During this stage, you can define important details about the AI asset, including the asset version and documentation.

Assess

The Assess stage is the evaluation of an AI asset to determine its effectiveness, its efficiency, its reliability, and its alignment with your organizational goals. This evaluation includes assessments for the performance, business and risk impact, regulatory compliance, and overall value of each AI asset.

Build and test

The Build and test stage is the development and testing of an AI asset to prepare it for deployment. When you're developing an AI asset, you must create the asset, code any applicable algorithms, and integrate relevant data sources. After you develop the AI asset, you can run tests to verify that it functions correctly, meets your performance standards, and produces accurate results. You can also identify and resolve bugs.

Deploy

The Deploy stage is the integration of an AI asset into your existing workflows. During this stage, you can also set up monitoring to track the performance of the AI asset. You can choose to deploy each AI asset through either a gradual roll-out, in which the asset can be used only by a specific subset of users within your organization, or a full roll-out, in which the asset can be used by any user within your organization.

For more information on Completing AI lifecycle stages, see Complete AI asset lifecycle

For more information on view AI assets by lifecycle stage, see View AI assets by life-cycle stage

Approvals

The approvals are requests created to either approve or reject an asset. The Approvals menu displays a list of Now Assist approvals.

Approval record

The Now Assist approval record displays **Details** and **Playbook** tabs.

The **Details** section displays all the details of an approval and the asset in review. The Assigned to field is editable and the rest of the fields are auto-populated.

The Playbook section displays the workflow for the Now Assist approvals to approve or reject an asset.

The Playbook templates list has a record of playbooks to configure an approval record in three steps:

- Review asset
- Evaluate asset
- Approve or reject an asset

The AI steward can save an approval record by clicking **Save** option and select the more actions icon to **Cancel approval request** or **Delete approval request**.

You can post your work notes in the **Comment** section and track the approval task requests from the **Activity** section.

For enabling Now Assist governance, default Approval Playbook is available.

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

The Approval Playbook feature is available for all Now Assist customers.

Note: If the onboarding playbook is present, then this feature isn't available.

For information about creating AI Control Tower Approval Playbook to enable Now Assist governance, see Create an AI Control Tower Playbook workflow.

For information about creating Now Assist approvals task, see .

Cases

You can view all the assigned and unassigned AI cases in a list view in the AI Control Tower.

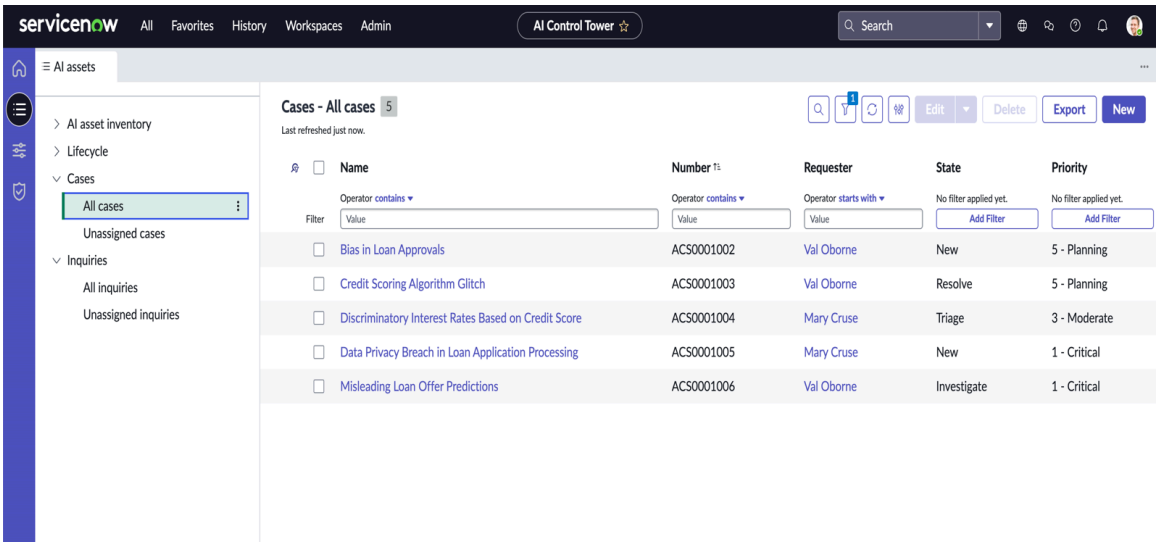
Cases list view

You can view all the cases in a list view. You can perform actions such as search, filter, refresh, and personalize list features on the list view.

Select any case from the list, select **Edit-More Actions** to assign the case to yourself, or **Delete** or **Export** the selected case. You can also select **New** to create a new AI case.

The following image displays all the AI cases in a list view.

Cases



For information about creating AI cases from the AI cases tab in the AI Control Tower workspace, see Create an AI case in the AI Control Tower

Inquiries

You can view all the assigned and unassigned inquiries in a list view in the AI Control Tower.

Inquiries list view

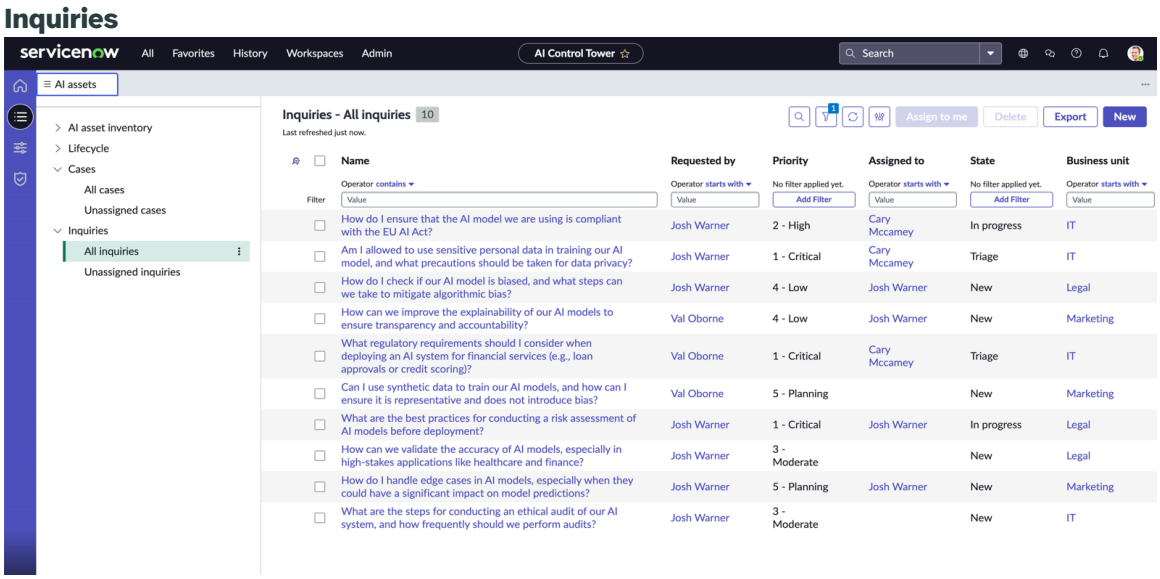
The inquiries section displays all the AI inquiries in a list view.

You can perform actions such as search, filter, refresh, and personalize list the inquiries on the list view.

Select any inquiry from the list, click **Assign to me** to assign the inquiry to yourself, or **Delete** or **Export** the selected inquiry. You can also click **New** to create a new AI inquiry.

The following image displays all the AI inquiries in a list view.

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



For information about creating AI inquiries from the AI cases tab in the AI Control Tower workspace, see Create an AI issue in the AI Control Tower

AI strategies and goals

You can track and monitor the progress of your strategic priorities, goals, planning items, and execution items—projects and demands classified as Artificial Intelligence in other ServiceNow applications.

Strategies

Actions of a strategic plan and its strategic priorities that help to orient your organizational strategic plan toward your goal. Strategic priorities are the objectives of your organizational plans. Strategic priorities help you to achieve your goals over a designated time period. The strategic priorities are also known as focus areas or long-term strategies.

Goals

Goals refer to objectives that an organization sets for itself to accomplish their organizational plans. Goals are typically qualitative in nature. Goals should be ambitious and are expected to motivate and challenge your teams.

Monitoring AI-related items

With Strategic Planning, you can track and monitor the progress of your strategic priorities, goals, planning items, and execution items—projects and demands classified as Artificial Intelligence in the Strategic Planning Workspace. For details, see Managing portfolio plan goals in Strategic Planning Workspace.

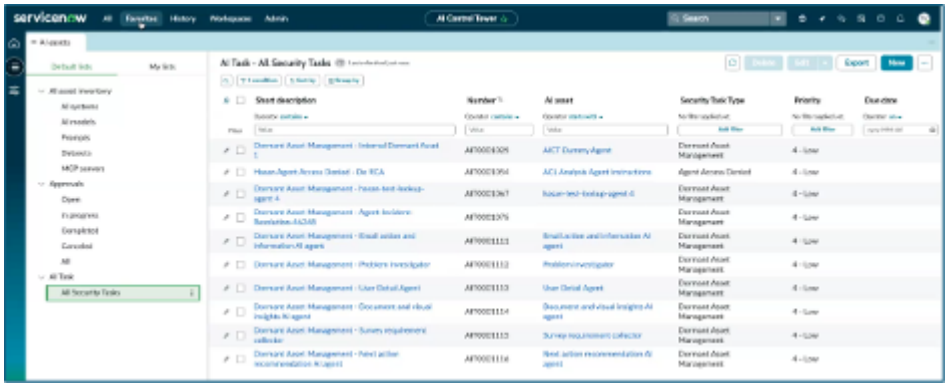
With Goal Framework, you can track and monitor the progress of your strategic priorities and goals classified as Artificial Intelligence in the Goal Framework application. For details, see Managing goals using Goal Framework or Goal Framework for SPM.

With PPM Standard, you can track and monitor the risks, issues, decisions, actions, and changes of your projects classified as Artificial Intelligence in Project Management or Project Workspace, or the primary goal or strategic priority associated with projects and categorized as Artificial Intelligence in Goal Framework or Strategic Planning Workspace. For details, see Create a project from Project Workspace and Add a RIDAC record to a project in Project Workspace.

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

AI Task

Use the **All Security Tasks** tab to view all AI security tasks for your instance. You can also create an AI task on this page.



Configurations

Explore the AI Control Tower Configurations page to manage and govern ServiceNow® AI assets.

The **Configurations** page contains the following sections to manage the workflow of AI Control Tower

- Data sharing, Data overflow processing, and Security & privacy in AI Control Tower
- Controls
- Multi-instance Setup
- AI connections
- Playbooks

For information on configuring AI Control Tower, see Configure AI Control Tower.

For information on completing the AI asset lifecycle, see Complete AI asset lifecycle.

Data sharing, Data overflow processing, and Security & privacy in AI Control Tower

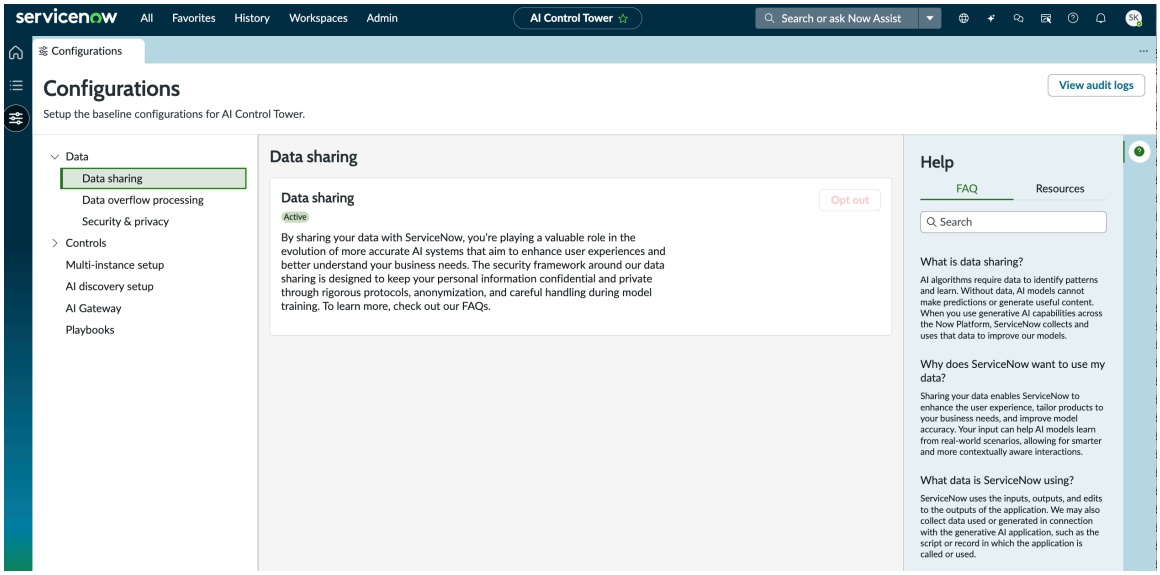
Explore the Data sharing, Data processing, and Security & privacy sections.

The section focuses on improving AI models, managing datacenter traffic, and enabling metrics to measure the integrity of your data model and monitor potential threats in large language model (LLM) input and output.

Data sharing

By default, Data sharing is active. You can opt out to deactivate AI Control Tower and share your data with ServiceNow to improve AI accuracy, enhance user experiences, and gain a better understanding of business needs.

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



Data sharing helps enhance ServiceNow products, but if you choose to opt out of the ServiceNow data sharing program, you'll no longer be able to contribute data to improve ServiceNow AI products.

For information on data sharing opt-out, see Opt out of data sharing.

Data overflow processing

By default, all Now Assist traffic is managed within ServiceNow datacenters. If there are traffic spikes, the system automatically redirects to Microsoft Azure datacenters to maintain performance. You can opt out of this feature to keep all Now Assist traffic exclusively within ServiceNow datacenters. By default, data overflow processing is inactive.

Note: The Data sharing and Data overflow processing features are available for a sub-prod instance in read-only mode, when Multi-instance setup is configured and active.

Security & privacy

Secure policies violation

These configuration settings control the Security policy violations chart, which shows violations of policies detected in LLM responses. To show data for this chart on the dashboard, select **Configure**, and then select **Active**. If you want to discontinue collecting data for the chart, deselect **Active**.

Note: If you inactivate the chart, past data shows on the chart for 90 days.

You can configure these settings:

- **Security policies** – Predefined ServiceNow policies to track for violations. These policies are based on industry best practices and can't be modified.
  - **Physical Chemical Biological Harm** – Detect recipes, parts lists, or methods that could cause real-world harm. For example, explosives, toxins, or bio agents.
  - **Policy Circumvention** – Detect jailbreak techniques. For example, disabling safety instructions or acting as an unfiltered model.
  - **Untrusted Links Or Downloads** – Detect linking to executable files or malware, directing users to run unverified downloads or macros, or asking users to disable their antivirus protection. Heuristic red flags include EXE, BAT, PS1, and APK files, along with macro-enabled documents.

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

- **Intellectual Property Proprietary** – Detect verbatim reproduction of paywalled or proprietary text or code beyond fair use.
- **Malware Exploits Hacking** – Detect exploited payloads, brute force attacks, evasion, persistence, keyloggers, or “ready-to-run” steps.
- **Refusal Appropriateness** – When a request conflicts with policy, refusal should be safe and brief. Detect under- and over-refusal. For example, refusing benign content.
- **Code And Command Safety** – Detect destructive and insecure commands. For example, don't allow “try this risky one-liner,” disabling security, or running unvetted code.
- **Social Engineering Deception** – Detect impersonation and phishing scripts, credential harvesting, and manipulative pretexts.
- **Fraud Misinformation Illegal Facilitation** – Detect how-to guides for illegal or deceptive practices. For example, for market manipulation, fraud templates, or evading audit. Benign corrections of misinformation are allowed, but not fabrication.
- **System Integrity Prompt Containment** – Detect output of system prompts, hidden policies, and internal configurations or tokens.
- **Legal Regulatory Compliance** – Detect guidance to evade regulations or practice without a license. For example, prescriptions, export control, or money laundering evasion.

- **Sampling rate** – The percentage of transactions that are evaluated. Selecting a rate lower than 100% results in fewer AI calls, but potentially less accurate data.
- **Max skill calls per execution** – The amount of AI usage per call. The minimum is 10 calls; the default is 1,000 calls. Entering a lower number results in fewer AI calls, but potentially less accurate data.
- **Single or multiple analysis** – Single analysis uses the default LLM to determine whether the model's output or behavior violates predefined security policies. Multiple analysis uses the results from three or more LLMs that ServiceNow supports to make a determination, using the majority result from the LLMs. Multiple analysis requires an odd number of LLMs.

Agent goal hijack

These configuration settings control the Agent output deviation chart, which shows when AI agents deviate from their intended role or objective. For example, unauthorized actions or prompt injection attempts. To show data for this chart on the dashboard, select **Configure**, and then select **Active**. If you want to discontinue collecting data for the chart, deselect **Active**.

Note: If you inactivate the chart, past data shows on the chart for 90 days.

You can configure these settings:

- **Sampling rate** – The percentage of transactions that are evaluated. Selecting a rate lower than 100% results in fewer AI calls, but potentially less accurate data.
- **Max skill calls per execution** – The amount of AI usage per call. The minimum is 10 calls; the default is 1,000 calls. Entering a lower number results in fewer AI calls, but potentially less accurate data.
- **Single or multiple analysis** – Single analysis uses the default LLM to determine whether the AI agent's or skill's response diverges from the expected output.

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

Multiple analysis uses the results from 3 or more LLMs to make a determination, using the majority result from the LLMs. Multiple analysis requires an odd number of LLMs.

Output screening

These configuration settings control the AI agent output with PII detected and High-risk AI agent output charts, which show when agents' LLM output contains PII or security-vulnerable patterns. To show data for these charts on the dashboard, select **Configure**, select **Active**, and then select a setting for the data to collect. If you want to discontinue collecting data for the charts, deselect **Active**.

Note: If you inactivate the charts, past data collected shows on the charts for 90 days.

You can configure these settings:

- **Output Security Vulnerability** – Collect and show data in the High-risk AI agent output chart. The data is collected by scanning LLM output for known vulnerable patterns and potential corresponding attack vectors. For example, HTML tags can have scripts associated with them to generate cross-site script attacks (XSS), or stacked SQL queries that can result in SQL injection attacks.
- **Output Extended PII** – Collect more PII data occurrences and show in the AI agent output with PII detected chart. The data is collected by scanning LLM output for additional PII data patterns beyond those specified in Data Privacy. These PII data patterns are U.S. CA drivers license, U.S. taxpayer ID, U.S. passport number, and vehicle ID number.
- **Output PII Violation** – Collect and show data in the AI agent output with PII detected chart. The data is collected by scanning LLM output for default PII sensitive data patterns specified in Data Privacy. For example, U.S. phone number or credit card number.

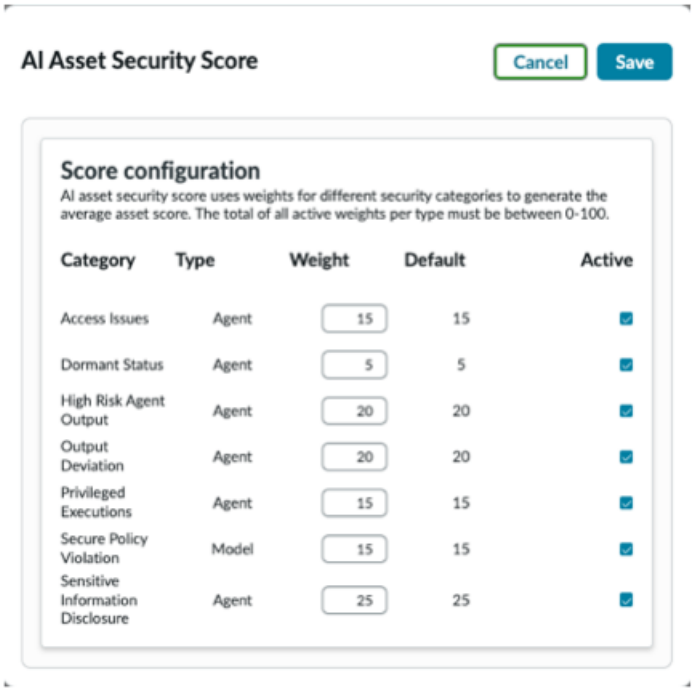
Sensitive data input and anonymization

This section shows the data patterns enabled in Data Privacy to detect and anonymize information in LLM prompts. Use this view as a quick reference when troubleshooting Sensitive data detected and Sensitive data anonymized charts. This feature requires the Data privacy plugin to be installed. For more information on how the data is sent and stored, see User data usage policy for Now Assist.

Score weight

This setting controls how the security categories that comprise the score are weighted. You can change the default weights or remove security categories from the score by deactivating them. The score formula is an average across all managed AI assets.

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



Controls

Explore the **Controls** section on the configurations page of the AI Control Tower.

Expand the **Controls** section to explore the sub sections:

- Explore Approvals
- Explore AI model providers
- Automation rules

Explore Approvals

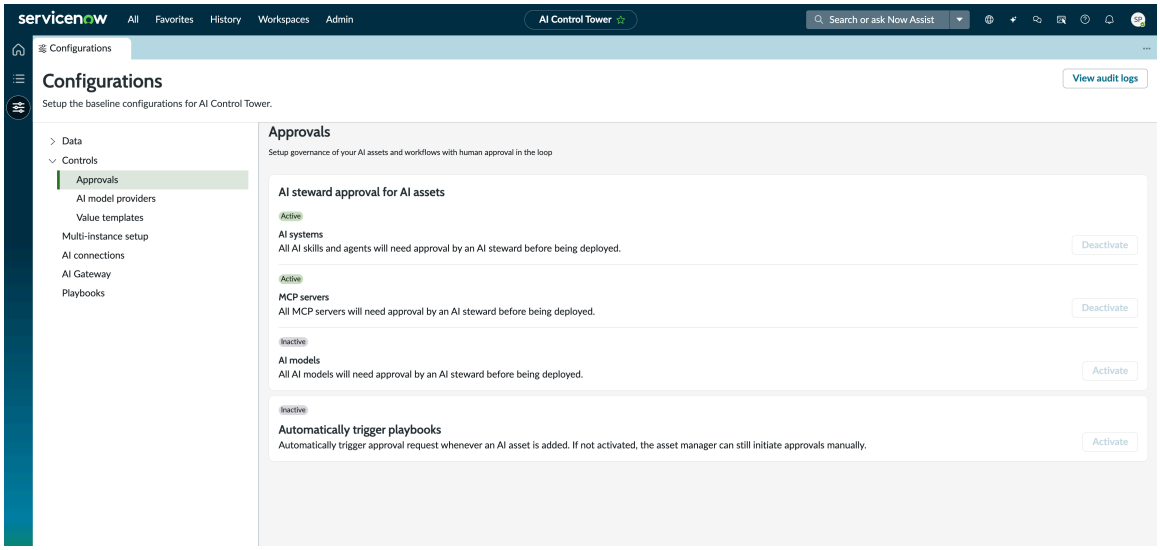
Explore the **Approvals** section on the Configurations page of AI Control Tower.

The approvals page has two features to set up governance of your approvals for AI assets and workflows.

- AI steward approval for AI assets- when the AI steward approval-required control is active, asset deployment is blocked until it receives an approval from the AI steward approves it. This control is inactive by default but can be activated for required approval before deploying an AI asset.
  1. AI systems- Activating AI systems, all the AI skills and agents need approval by an AI steward before being deployed.
  2. MCP servers- Activating MCP servers setting, AI Stewards can block unapproved MCP servers from being used in AI Agent Studio. The setting is deactivated by default and must be manually activated by the AI steward.

     When you attempt to add an unapproved server in Agent Studio, you'll see an error message: 'Approval is needed to display tools', helping prevent any interaction with that server.
  3. AI models- Activating this all AI models need approval by an AI steward before being deployed.
- Automatically trigger playbooks- enabling automatically trigger playbooks control triggers approval requests for AI assets. When inactive, only the AI steward can initiate requests manually. The recommendation is to have this control activated in the production environment.

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



Explore AI model providers

Explore the AI model providers section in AI Control Tower.

AI model providers enable you to select data routing, manage third-party LLMs (Large Language Models) and SLMs (Small Language Models), and configure the third-party LLMs by selecting the allowed model providers.

Data routing and model providers

Data routing is a technology by which LLM and SLM requests are routed to the most suitable datacenter. This technology helps to optimize data traffic, which reduces latency and speeds up response time.

The Data routing and model providers section enables you to route AI model requests and select the Allowed model providers.

There are two types of data routing:

Regional data routing

Regional data routing routes LLM and SLM requests to datacenters within your region. For example, if you're in the APJC (Asia Pacific, Japan, and China) region, these requests could be routed to the most suitable datacenter in the APJC region. Regional data routing can sometimes be mandated by governments of specific regions.

Global data routing

When you opt for Global data routing, LLM and SLM requests are routed to the most suitable datacenter globally.

You can configure the third-party LLM providers using the edit option by choosing either Regional or Global data routing and select all the Allowed model providers.

There are two sections of AI model providers:

- AI model providers supported by ServiceNow
- AI model providers configured by your organization

The AI model providers supported by ServiceNow contain providers such as Now LLM Service, AWS Claude, Now LLM-LTS model and so on.

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

The AI model providers configured by your organization such as Perplexity, IBM Watson and so on.

AI asset lifecycle

Note: You can select Now LLM Service- LTS (Long term stable) Model, which supports regulated industries, such as financial institutions, with stronger AI lifecycle management, governance, transparency, and compliance tools.

For more information on Now LLM Service- LTS model, see Long term stable models

For information on exploring the scenarios configuring third-party LLMs for all the regions, see Explore Third-party LLMs and regions

For information about configuring third-party LLMs through Data routing configuration for APJC region, see Configure third-party LLMs using AI Control Tower

Fallback and Spillover

Fallback

If you have active AI systems in ServiceNow® that aren't supported by your enabled model providers, the 'fallback mechanism' enables these systems to continue operating with their default providers. However, if you choose not to enable fallback, these AI systems must be deactivated. AI systems deployed on fallback providers conform with the list of approved providers.

Note: The Fallback is activated by default and can be modified.

Spillover

Regional deployments of AI models can experience limited capacity, which could lead to request rate limiting and impact performance. Enabling spillover can help prevent these performance issues. In ServiceNow®, only Azure OpenAI currently makes this switch.

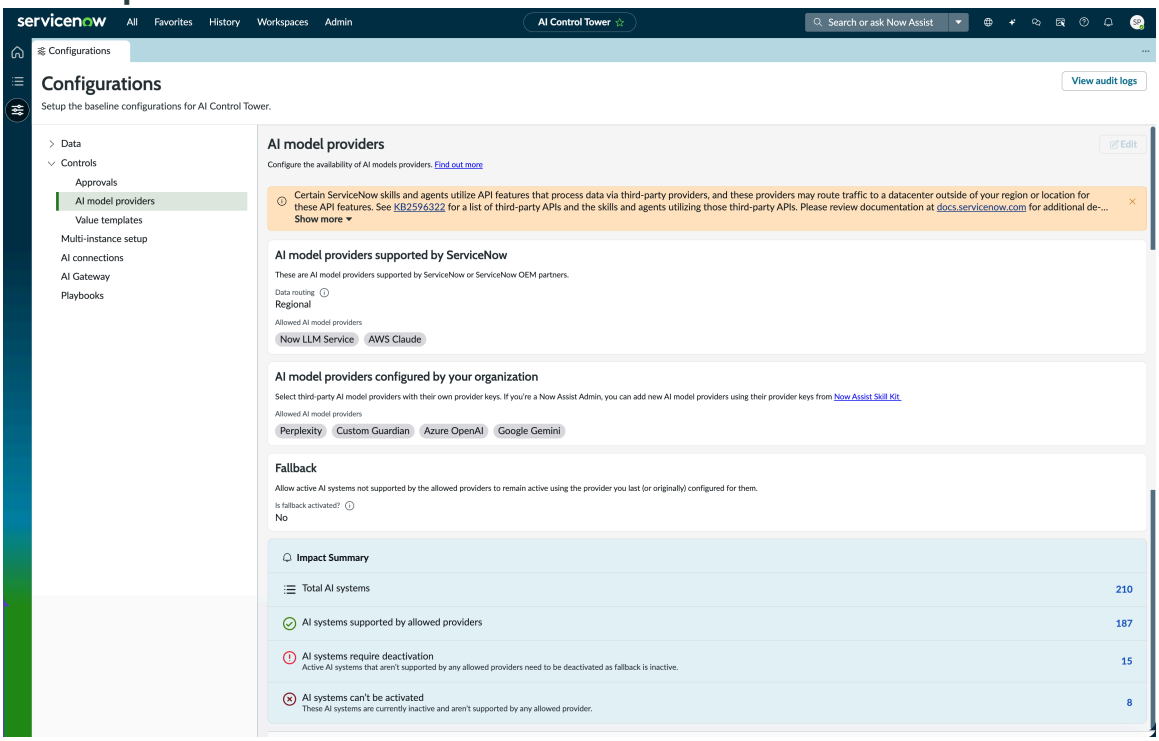
Note: The Spillover feature gets active or enabled when Azure OpenAI gets selected.

Impact Summary

The Impact Summary is determined by the chosen Allowed model providers and the status of the fallback, which is either active or inactive. The Fallback significantly affects how the Impact Summary data appears in the Impact Summary table.

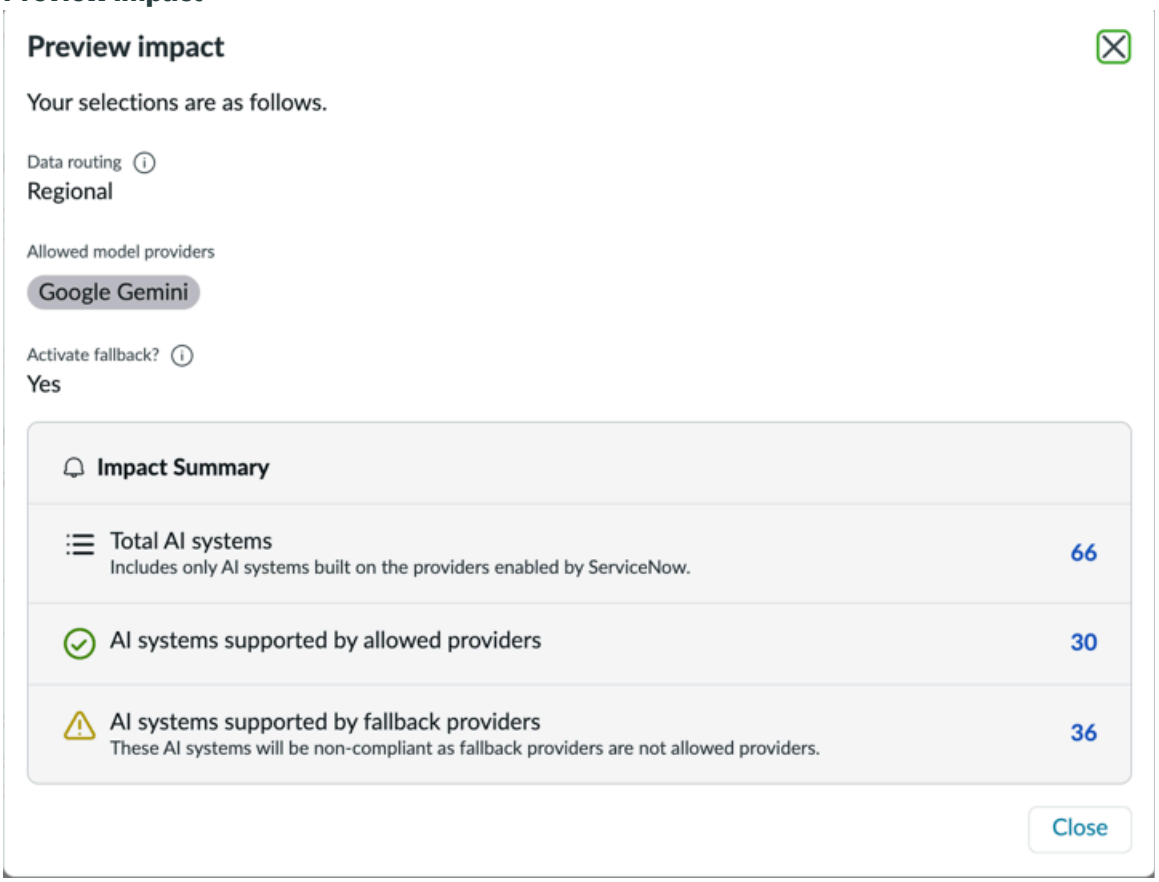
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective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

AI model providers



You can use the edit option to select Yes or No for activating the Fallback. Before saving, you can select Preview impact to review and confirm all your selections.

Preview impact



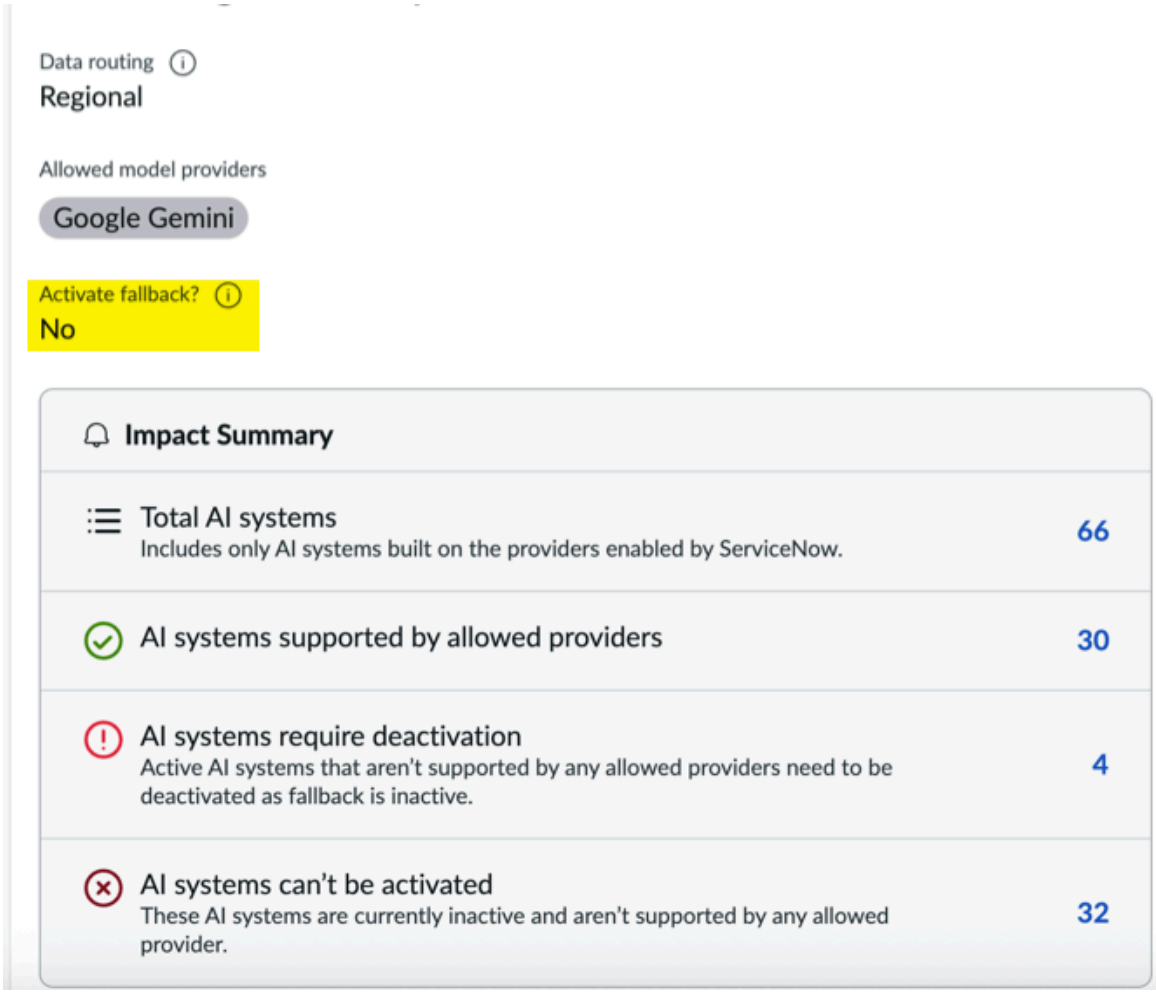
Let's review the Impact Summary table data for the following two scenarios.

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

Activate fallback- No

- Total AI systems- Shows all AI systems that are supported by the allowed model providers.
- AI systems supported by allowed providers- Shows AI systems with skill sets that are supported by the providers.
- AI systems require deactivation- Lists all active AI systems that lack provider support and must be deactivated because the fallback option isn't enabled.
- AI systems can't be activated- Shows all those systems, which are currently inactive and aren't supported by any provider.

Activate fallback- No

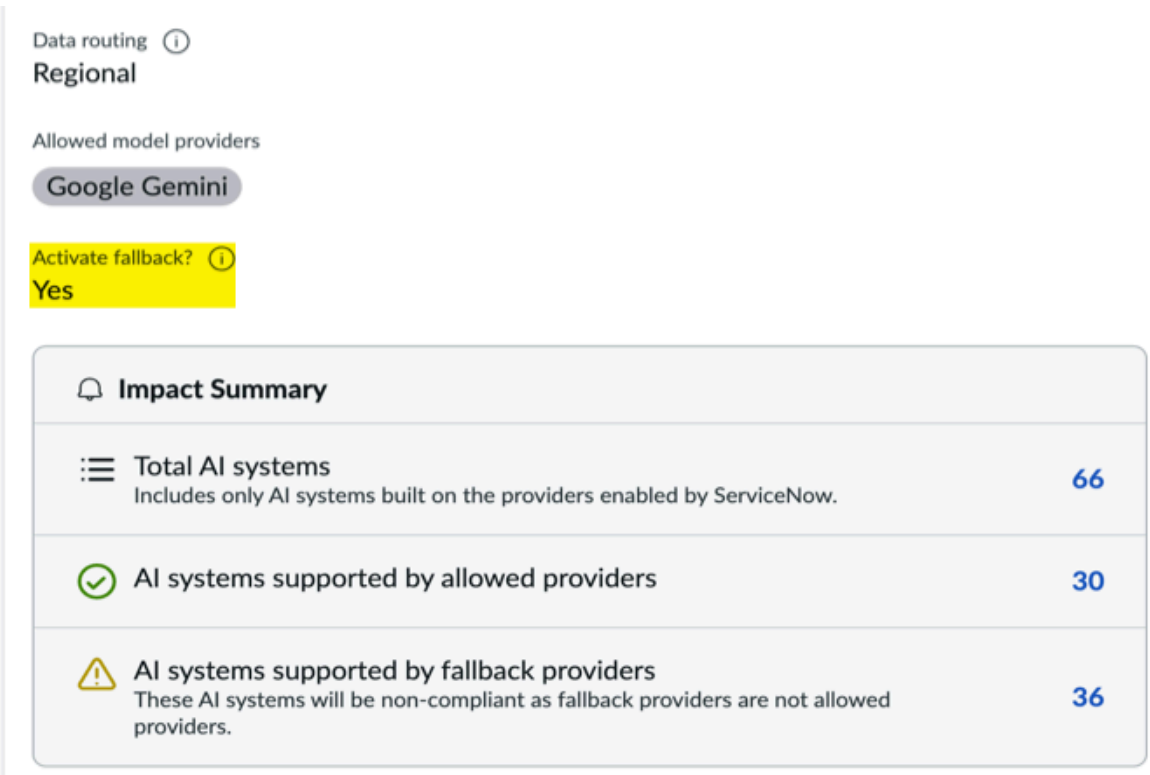


Activate fallback- Yes

- Total AI systems- Shows all AI systems that are supported by the allowed model providers.
- AI systems supported by allowed providers- Shows all AI systems with skill sets that are supported by the providers.
- AI systems supported by fallback providers- Shows AI systems that are non-compliant as fallback providers aren't permitted providers.

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

Activate fallback- Yes



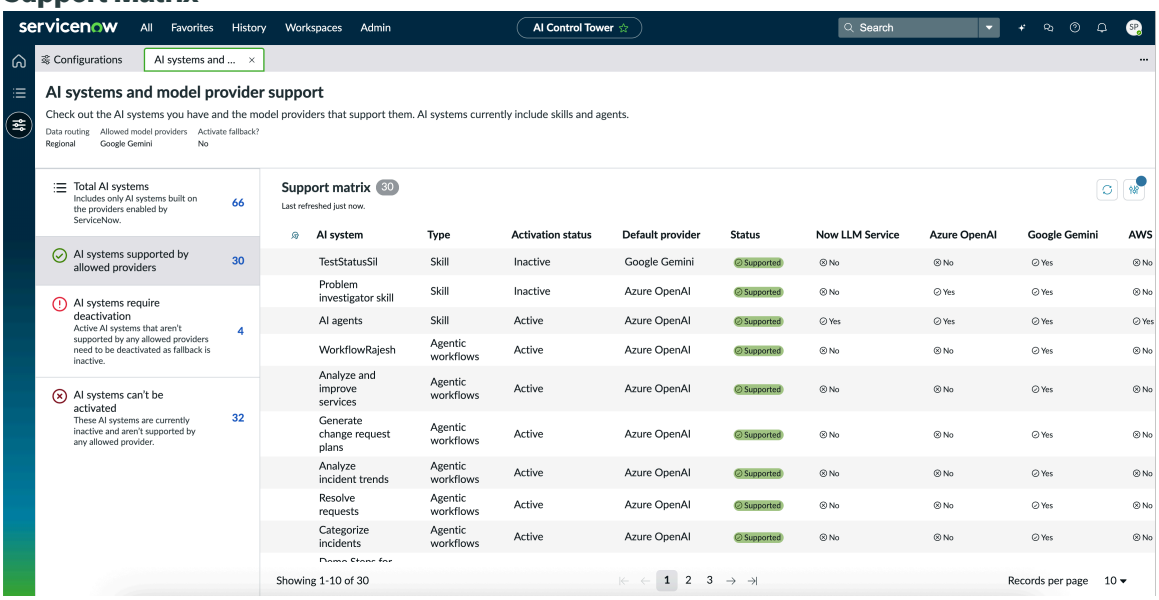
Note: The entries in the Impact summary table change based on the fallback status.

When you select an entry from the table, the support matrix page appears with those selected entries, allowing you to update your personalized list.

The support matrix presents all AI systems in a table format, along with their respective AI model providers. You are able to view the support matrix table and categories such as AI system, type, activation status and more, as well as the selected AI model provider.

If you have selected an AI provider, which is supported by your organization or a third party provider, the selected provider will show up in the AI systems and model provider support table.

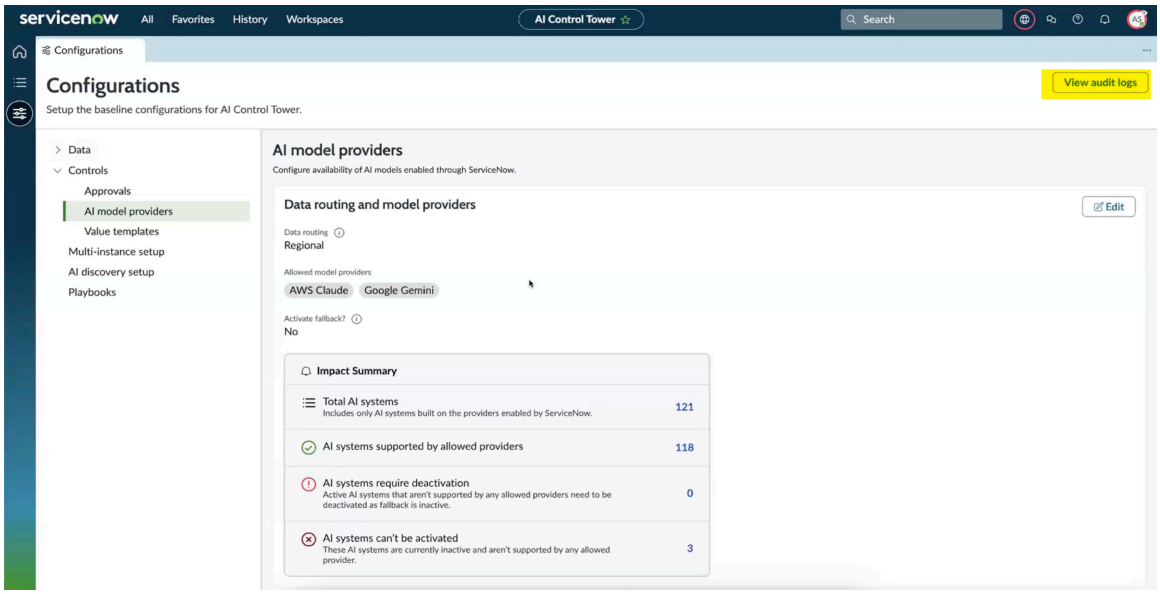
Support Matrix



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

Audit logs

Audit logs show configuration changes made on Data, Approvals, and AI model providers categories in AI Control Tower. You can select the **View audit logs** option to view the Audit logs.



The Audit logs page displays all the configuration changes details organized in the following categories:

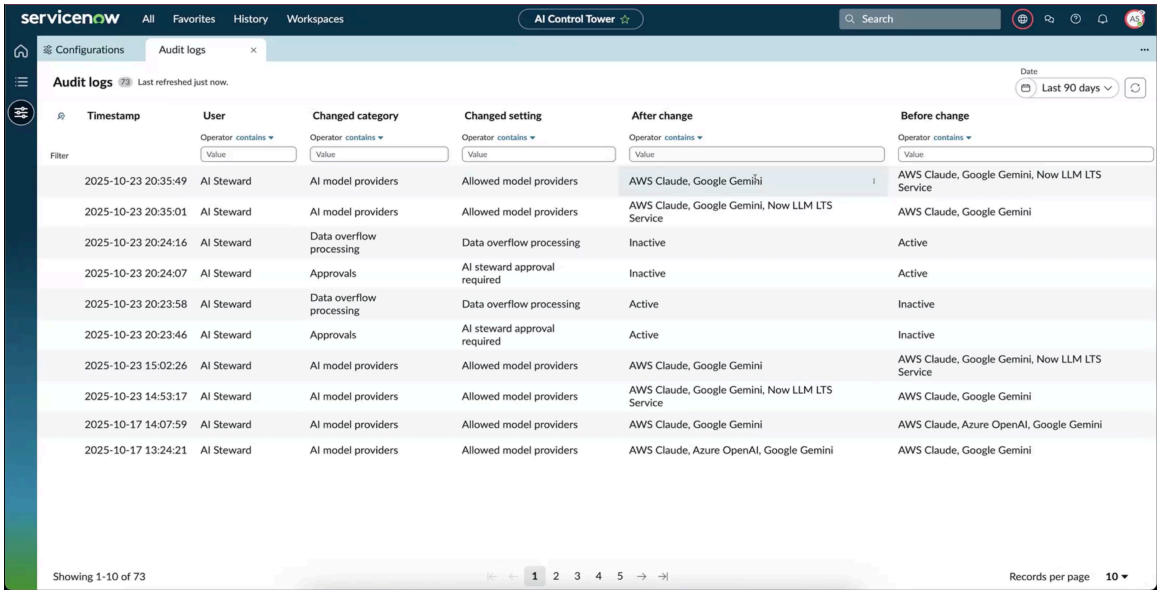
- Timestamp
- User
- Changed category
- Changed setting
- After change
- Before change

You can also filter the changes by selecting a date range, starting with the last 90 days.

Note:

In the Multi-instance setup, when a managed (sub-prod) instance is added or removed from the syncing instances in the AI inventory information to synchronize with a specific manager (prod) instance, the audit logs first display a record of all instances being removed, followed by a separate record indicating the instance being added or removed.

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



Automation rules

The Automation rules define how AI assets are set to be under managed assets.

Automation rules Overview

The Automation rules page displays all default rules available in AI Control Tower. The AI Stewards can modify the existing rules or create rules based on their organizational requirements.

Rule limits

The Administrators can configure system-enforced limits on asset management rules through system properties. The recommended defaults are 5 active rules and 10 total rules, controlled by the following system properties:

System properties

System properties	description
sn_ai_governance.management_rule.max_active_rules	Max active rules
sn_ai_governance.management_rule.max_total_rules	Max total rules

Rules execution

Only active rules are evaluated during scheduled runs; inactive rules are skipped. The scheduled job 'Run AI Asset Management Rules' is executed for every 6 hours, but AI stewards can trigger immediate execution on any active rule using the **Run Now** action. Rules that are no longer needed can be deleted. If an asset is qualified for a rule, it will be automatically moved under managed assets.

References

For information on managed and unmanaged assets, see Assets list- managed and unmanaged assets

For information on creating rules, see Create an Automation rule

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

Multi-instance Setup

The multi-instance setup enables a prod (manager) instance to control, manage, and communicate with multiple sub-prod (managed) instances for AI Control Tower.

AI asset Synchronization

Multi-instance setup uses the multi-instance framework, which helps the user to synchronize assets from sub-prod instances to prod instances for a faster review process.

Multi-instance setup synchronizes rules for the sub-prod instances from the prod instance.

Note: Starting with the February 2026 release, ensure that both the prod and sub-prod instances are running the same AI Control Tower core version (5.1.1), which is the minimum supported version.

If there's any upgrade to version 5.1.1 in a sub-prod, then it's advisable to upgrade the prod instance to 5.1.1 to ensure Multi-instance framework functions correctly.

AI inventory information

You can include the sub-prod instances that you want to synchronize with the prod instance. This synchronizes AI inventory information between the instances.

When configured, the scheduled job starts syncing AI systems, AI models, prompts, and datasets. From the September (2025) release, the job has been enhanced to include syncing AI agents as well.

Note: State of the assets while configuring Multi-instance management.

The AI inventory in production reflects the true state of your assets like models, datasets, or skills from a production standpoint. Even if a model or dataset is active in a sub prod (lower) environment, it's still considered as under development from a prod perspective, since it's being tested and not yet live.

For this reason, you don't synchronize asset states across environments. An asset's state changes to **deployed** only when the asset and its related records are activated in the production system.

In summary, the state represents the overall lifecycle of the asset, not its local status in a specific environment.

Data sharing preference

You have the option to enable the data sharing preference, when it is enabled the preferences of the data sharing from the production will be applied to all sub-prod instances. By default, the data sharing preference is turned off.

Data overflow processing and bursting preference

You have the option to enable the data overflow processing and bursting preferences, when it is enabled the preferences of the data overflow and bursting from the production will be applied to all sub-prod instances. By default, data overflow processing and bursting is turned off.

Note: All the preferences mentioned earlier for a sub-prod instance are available in read-only mode, when Multi-instance is configured and enabled.

For information about configuring Multi-instance management for AI Control Tower, see Configure multi-instance management for AI Control Tower

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

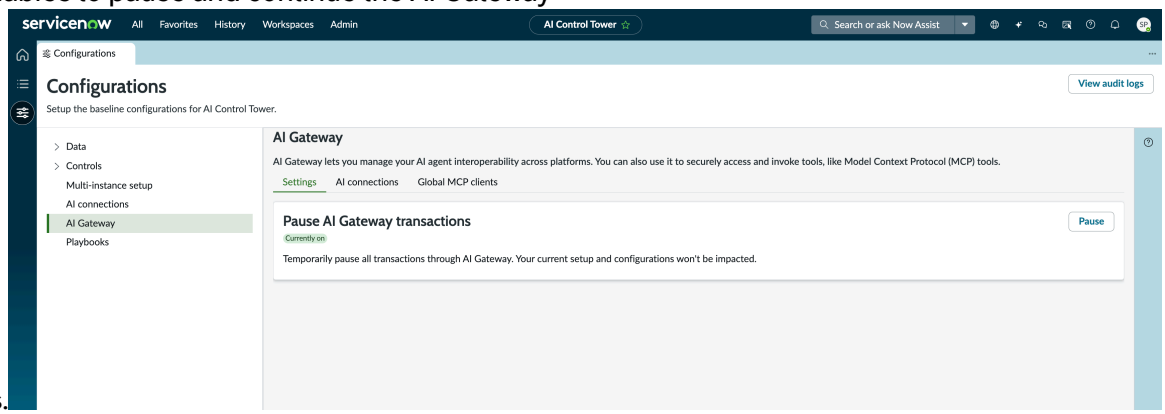
For information about Data section, see Data sharing, Data overflow processing, and Security & privacy in AI Control Tower

AI Gateway

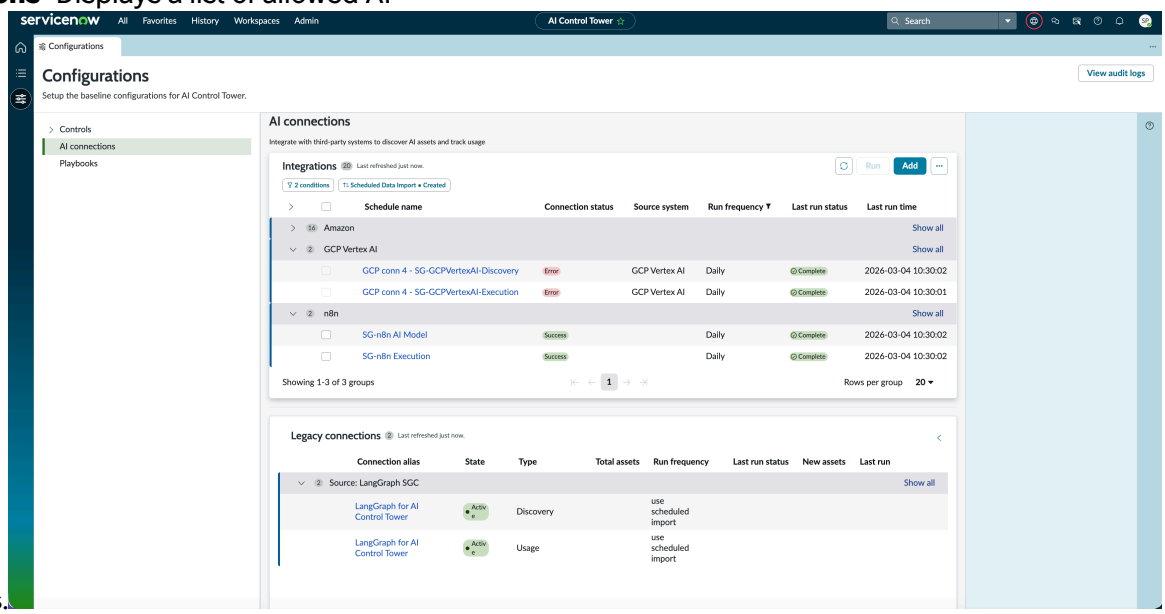
Explore the AI Gateway page and its tab in the AI Control Tower Configurations.

The AI Gateway page enables you to pause and resume AI Gateway transactions, as well as monitor AI connections. The AI Gateway page has the following tabs:

Settings- enables to pause and continue the AI Gateway transactions.



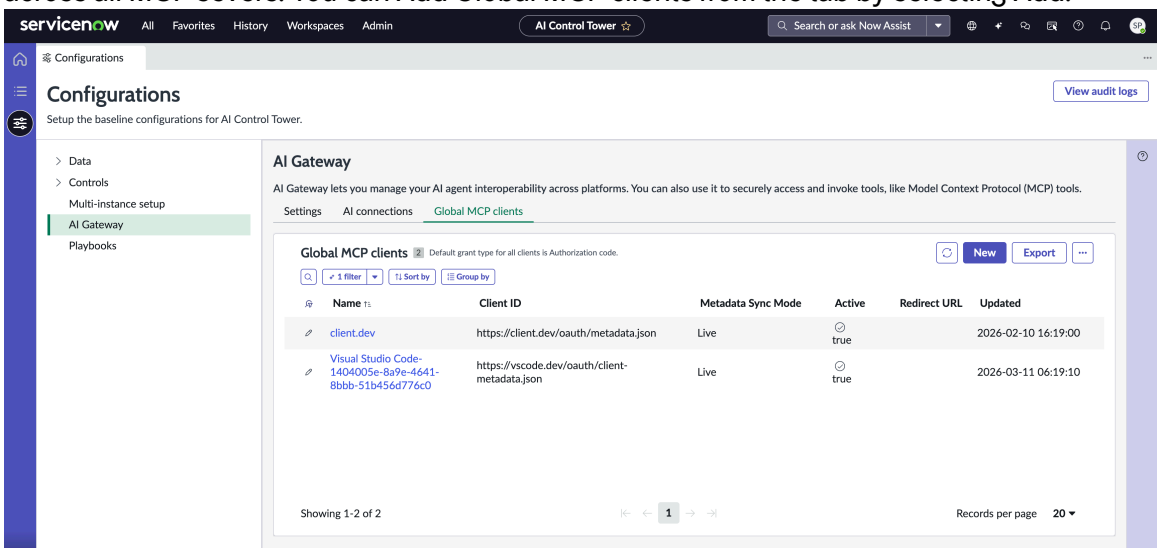
AI connections- Displays a list of allowed AI connections.



Global MCP clients- View the list of Global MCP clients (CIMD support), Where CIMD refers to Client ID Metadata document. These Global MCP clients once created can be used

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

across all MCP severs. You can Add Global MCP clients from the tab by selecting Add.



For more information on adding a Global MCP client, see Add a Global MCP client.

You can filter and export the list into PDF, JSON, Excel, and CSV format. The list of allowed AI connections is categorized into.

- Name
- Allowed on AI Gateway
- Last updated by
- Updated

For information on AI Gateway tab on the AI Control Tower Home, see AI Gateway tab in AI Control Tower

Playbooks

The Playbook templates list has a record of Approval Playbook for Now Assist approvals and Asset lifecycle playbook templates.

The Now Assist approval is a workflow that facilitates the AI asset approval process.

The AI asset lifecyle is completed by the following steps:

- Onboarding
- Assess
- Build and test
- Deploy

You can create your own playbook workflow by customizing the number of steps or rearranging them, as well as applying different security policies.

Note: To retire an asset, offboarding process is initiated.

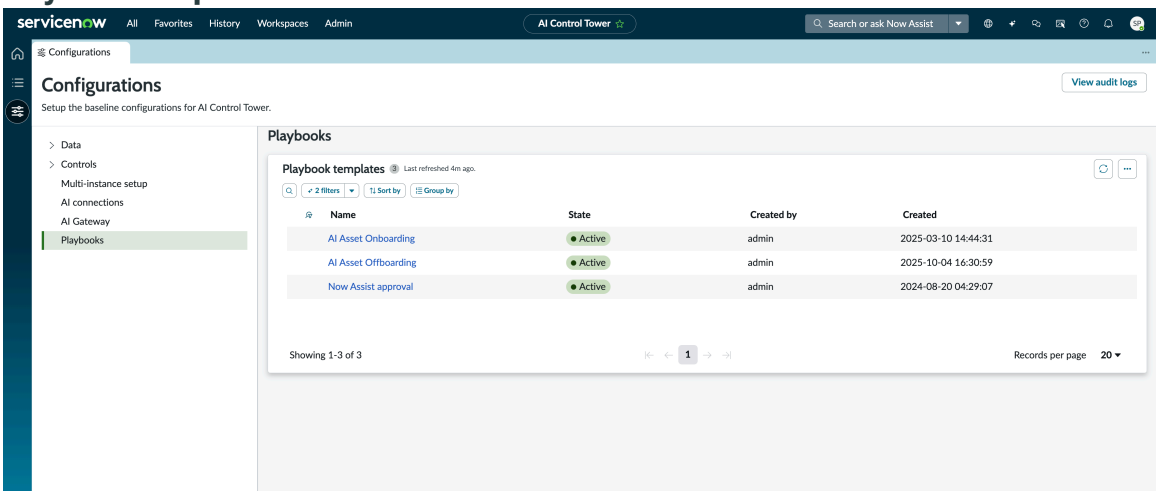
For information about AI asset lifecycle process, see AI asset lifecycle

To view AI assets by lifecycle stages, see View AI assets by life-cycle stage

For information about Off-boarding and creating offboard requests for AI assets, see Create offboarding requests for AI assets

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

Playbooks templates list



For information on creating a AI Control Tower playbook workflow for Now Assist approvals, see Create an AI Control Tower Playbook workflow.

To view the Asset lifecycle stage, see View AI assets by life-cycle stage and to view the completing of the Asset lifecycle stage, see Complete AI asset lifecycle

Product Owner portal

Explore the product owner view and navigate through the accessible tabs on AI Control Tower to ensure a comprehensive understanding of the features.

Product owner view

The AI Control Tower product owner view displays tabs for AI asset owner and workspace users. The AI asset owner can access **My overview**, **Value**, and **Adoption** tabs and the workspace user can access the **AI portfolio** tab.

For information about My overview tab, see My overview tab in AI Control Tower

For information about My portfolio tab, see AI portfolio tab in AI Control Tower

For information about Value tab, see Value tab in AI Control Tower

My overview tab in AI Control Tower

The **My overview** tab in the AI Control Tower home page displays all active assets to the asset owner and the AI stewards for assets that they manage or own.

Explore My overview tab

The **My overview** tab displays the following widgets. You can drill down into the data on each widget.

Assets

This area displays all the assets owned by the asset owner, organized into the following categories:

- Generative AI
- Agentic AI
- AI prompt

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

- AI dataset
- AI model

Assets by lifecycle status

This area displays the lifecycle status of the assets, organized into the following categories:

- Deployed
- In review
- AI steward review
- Ready for deployment
- Rejected
- Cancelled

AI cases by state

This area displays the active cases categorized into **New** and **Triage**.

The **Create an AI asset** button is added to the home page to track and manage the life cycles of your AI systems, AI models, Prompts, and Datasets.

Tasks and cases list

The **My overview** tab displays a list of tasks and cases. You can view all the active tasks that are assigned by default under tasks.

The tasks list includes the following columns:

- Number
- Priority
- Short description
- Assigned to
- State
- Due date
- Asset in review

The cases list includes the following columns:

- Name
- Number
- Case analyst
- Assignment group
- State
- Priority

AI portfolio tab in AI Control Tower

The **AI portfolio** tab in the AI Control Tower home page displays all AI assets, which includes (active, inactive, and completed) to the workspace users.

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

Explore AI portfolio tab

The AI portfolio tab AI shows all the active AI assets organized into the following categories:

- AI systems
- AI model
- Datasets
- Prompts

The lifecycle status of AI assets is divided into the following categories:

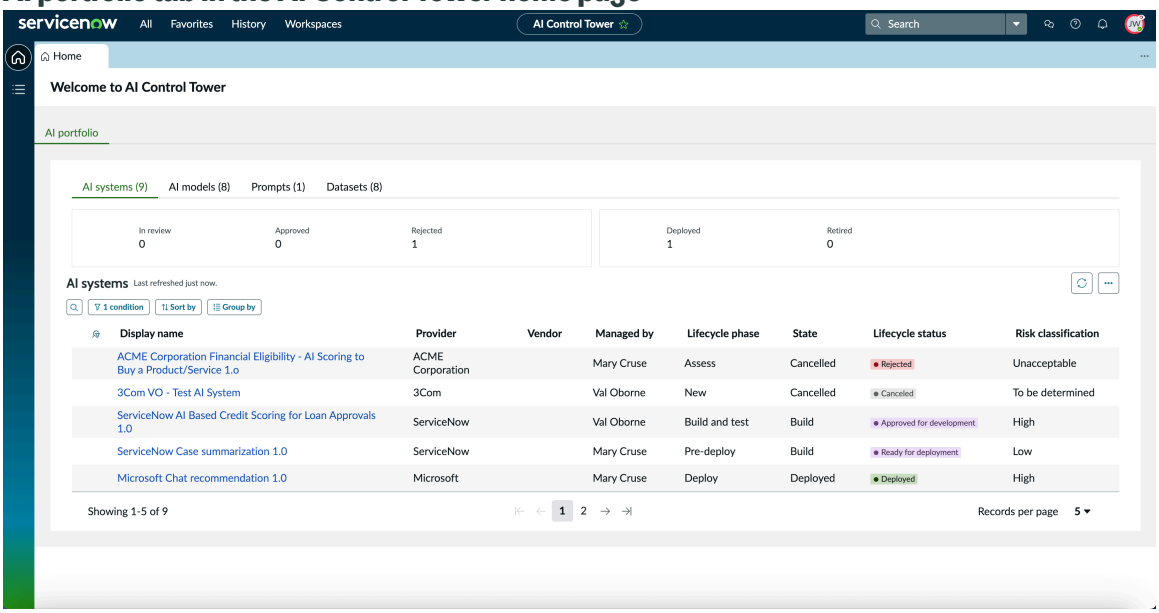
- In review
- Approved
- Rejected
- Deployed
- Retired

The AI assets list includes the following columns:

- Display name
- Provider
- Vendor
- Managed by
- Lifecycle phase
- Install status
- Lifcycle status
- Risk classification

The following image displays the AI portfolio tab as it appears to a workspace user.

AI portfolio tab in the AI Control Tower home page



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

Definitions of AI asset inventory

Explore AI asset inventory, and AI assets.

AI asset inventory overview

The AI asset Inventory provides a record of all components and their connections that contribute to an AI experience. In ServiceNow AI Control Tower, AI asset inventories can be organized in multiple ways. Organizations may opt to manage one or more of these AI assets as part of their overall AI governance strategy.

- AI systems
- AI models
- Datasets
- Prompts

For more information on, AI asset inventory tab, see AI asset inventory tab in AI Control Tower

For information on AI assets, see AI assets

AI systems

Explore AI systems in detail with examples.

AI system

An AI system is a type of computer program designed for performing tasks that typically require human intelligence, such as learning the language, deciding, or recognizing images. Unlike traditional software that operates based on set rules, AI systems learn from data, adapt their actions, and can decide on their own.

Hierarchical AI system

AI systems can be interconnected, with one system potentially depending on the outputs of another, referred to as a sub AI system." Recognizing these connections is important for monitoring AI usage and clarifying responsibilities.

Tools are considered attributes of an AI system or sub AI system. They aren't tracked as separate AI assets in the inventory and don't have a lifecycle managed by the AI Control Tower for governance or value assessment. The concept of Tools is relevant for agentic AI use cases, where they're linked to AI systems during discovery, often through cloud platforms like Azure Foundry, AWS, or internal discovery processes.

AI Control Tower helps identify and track all AI systems in an organization, including how they're connected. This supports AI risk assessments and audits.

Examples:

- Spam filter in Gmail- Determines whether emails are spam or not
- Netflix recommendation engine- Suggests shows that you might like
- Self-driving car software- Makes driving decisions
- Chatbot like ChatGPT-Responds to questions
- AI Issue Resolution Agent- uses other smaller AI systems like Issue Summarization and Owner Mapping

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

Example of an AI system

The AI Control Tower application and the AI Risk and Compliance application play a critical role in managing and governing the responsible use of AI systems, especially in high-stakes domains such as banking and finance.

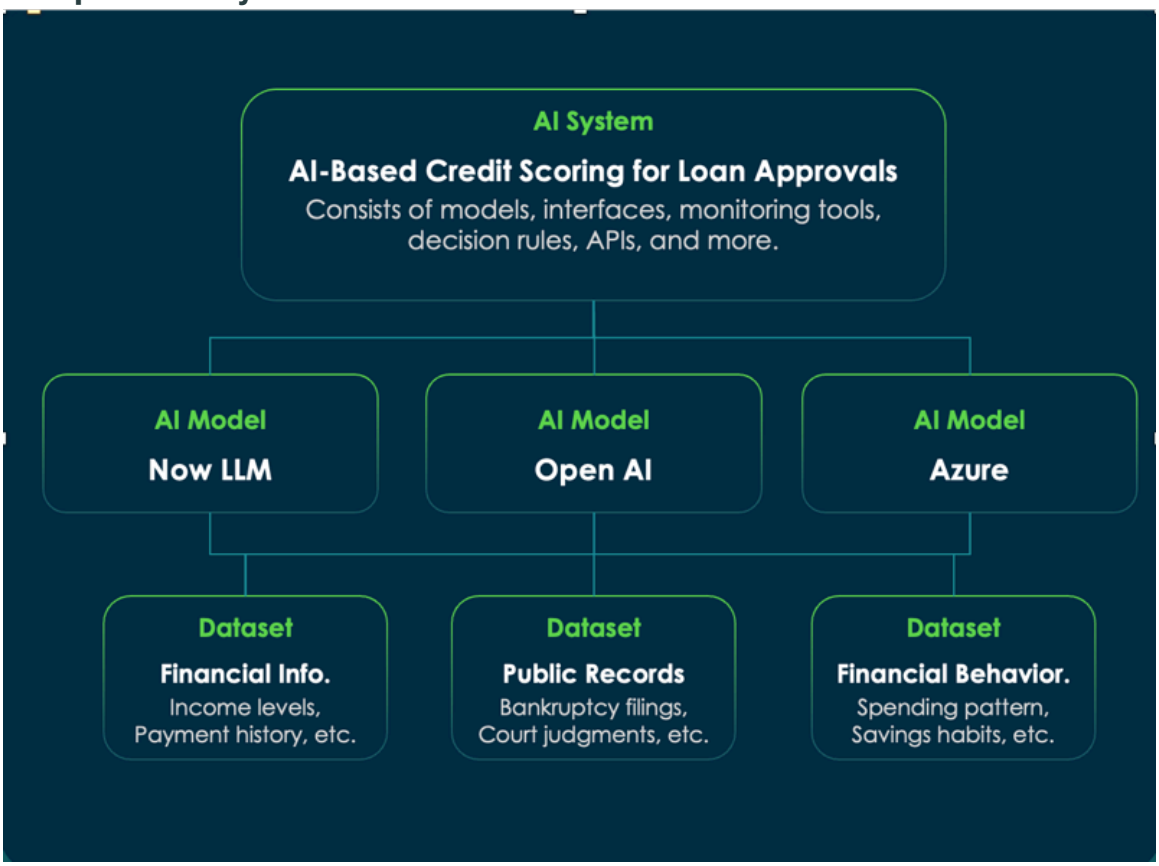
While AI risks extend beyond banking and finance, consider an AI system for credit scoring in loan approval as an example. This system interacts with other models, interfaces, monitoring tools, and decision rules. The AI Control Tower solution along with the AI Risk and Compliance application can help to ensure compliance and responsible usage when the AI system is created, built, and deployed.

For instance, an organization might use AI to generate credit scores for loan approvals and using those scores to make approval decisions. The AI system can be fed a variety of data from the following large language models (LLMs):

- Now LLM: Provides financial information, payment history, and so on.
- Open AI: Provides details such as bankruptcy filings, court judgments, and so on.
- Azure OpenAI: Provides spending and saving patterns.

The figure in the following example illustrates how an AI system is interrelated with datasets, models, and the broader ecosystem.

Example of an AI system



AI models

Explore AI models in detail with examples.

An AI model acts as the "brain" of an AI system. It learns from previous data and applies that knowledge to make predictions, respond to questions, or identify patterns. Models are developed using various training techniques, such as:

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

- Supervised learning
- Unsupervised learning
- Reinforcement learning

Learning which models power your AI systems is important for monitoring their accuracy, fairness, and potential risks.

Examples:

- Spam detection model- Flags emails as spam
- Language model (like ChatGPT)- Generates text
- Image model- Recognizes objects in photos
- Credit model- Predicts who might default on a loan

Prompts

Explore prompts in detail with examples.

A prompt is a question or instruction provided to an AI system, indicating what you want it to accomplish. Prompts are commonly used in tools like ChatGPT or image generators.

Certain AI systems, such as GenAI systems, require user prompts to function. Monitoring prompts can help analyze user interactions with the AI and the types of outputs produced.

Examples:

- Write a poem about the moon- AI writes a poem
- Translate hello into Spanish- Hola
- Show red dresses under $50- Filters e-commerce results
- Summarize this email- Creates a short summary

Datasets

Explore datasets in detail with examples.

A dataset is a collection of data used to train or evaluate an AI model. You can think of it as a large spreadsheet with rows representing examples and columns for features or labels. Datasets can also come in other formats, such as images, audio, or video. They help models learn patterns and decide.

Mapping datasets within the AI Control Tower is important for tracking data sources, ownership, and other key details, making it a vital part of the organization's AI governance program.

Examples:

- A set of labeled images (dogs, cats) for training a model
- Articles and books for training a language model
- Customer purchase history for recommendations
- Medical records for diagnosis tools

Inter-dependencies

Explore the inter-dependencies of AI systems, AI models, and datasets.

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

An AI system operates as part of a larger framework, closely connected to the AI model powering its capabilities and the datasets that shape its behavior. While these elements are interrelated, each plays a unique role and should be clearly identified to confirm effective governance.

Overlooking any aspect, such as the training dataset or the model supporting an AI feature can result in incomplete risk assessments, audit challenges, or compliance issues.

Comprehensive responsible AI governance requires attention to all these components:

- Inventory for each AI system
- Identify the associated AI model
- Document the dataset

For a visual example of how these elements interact, see Example of an AI system

For information on Model asset classes, see Model asset classes

For information on AI assets API, see AI Assets API

Explore Third-party LLMs and regions

Explore all the scenarios for configuring third party LLMs and SLMs based on the regions.

Configuration overview

These are the few scenarios to consider when configuring LLMs:

APJC, AMS & EMEA regions

LLMs follow regional data routing configurations, as there's no global routing option and you can view all the allowed providers.

Regulated markets

The Data routing configuration isn't available, and the list of allowed providers is limited. You can view the allowed provider, but only in a read-only format.

Activating Fallback enables AI systems supported by your approved providers to stay active by using their default providers. Manually deactivating Fallback ensures that these AI systems remain inactive.

For information about Data routing, Fallback, Spillover and Impact summary, see Explore AI model providers

For information on configuring of third party LLMs for APJC region, see Configure third-party LLMs using AI Control Tower

Explore Now Assist AI asset discovery

Explore the syncing process of AI assets including models, datasets, prompts, skills, and agentic AI components into the AI Asset Inventory of AI Control Tower.

The AI systems are synchronized from source Now Assist skill config record (sn_nowassist_skill_config).

1. The out of the box skill meets the following conditions for an AI system:
   a. If active, insert.
   b. If not active, insert only if it belongs to the custom skill category and is published.

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

2. If the skill exists in AI Control Tower and the source Now Assist skill config record is updated, those changes are applied to the AI system record in AI Control Tower. The state for the AI system is mapped as follows based on source Now Assist skill config record state:
   a. Active: Mark as Deployed
   b. Published or Draft: Mark as In Development
   c. Retired: Mark as Retired
   d. Deleted: No update is made

AI Asset Sync Mapping Overview

Type	Source Table	Target Tables	Field Tracked	Records included
Models	sys_generative_ai_model	alm_ai_model_config, cmdb_ai_model_product_model	active, model, supported_languages, domain_id	Active = true
Datasets	sys_one_extend_test_dataset	alm_ai_dataset, cmdb_ai_dataset_product_model	active, product_model, description	Published = true eval_run datasets
Skills	sn_nowassist_skill_config	alm_ai_system_digital_asset, cmdb_ai_system_product_model	active, datasets, prompts, model, sys_created_by, domain_id, name, description	Active = true or state = published for custom skills; skill family not platform or any of its children(non platform skills)
Prompts	sys_generative_ai_prompt_config	alm_ai_prompt_digital_asset, cmdb_ai_prompt_product_model	active, prompt, model, name, version	Active prompts with active model and linked skill
Tools	sn_aia_tool	sn_ent_ai_tool	name, description, type, active, sys_created_by, domain_id	
AI Agents	sn_aia_agent	alm_ai_system_digital_asset, cmdb_ai_system_product_model	name, description, role, instructions, sys_created_by, domain_id	Active = true
Agent Tool Mapping	sn_aia_agent_tool_m2m	sn_ent_ai_system_agent_tool_component_m2m	All agent-tool mapping	All agent-tool mappings
Agent Use cases	sn_aia_usecase	alm_ai_system_digital_asset, cmdb_ai_system_product_model	name, description, sys_created_by, domain_id	Active = true
Usecase to	sn_aia_team_member_m2m	sn_ent_ai_system_usecase_agent_m2m	usecase, agent mapping	All usecase-agent mappings

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

AI Asset Sync Mapping Overview (continued)

Type	Source Table	Target Tables	Field Tracked	Records included
Agent mapping				

For more information on Model, dataset, prompt, and skill sync process, see KB2674041 article in the Now Support Knowledge Base.

Enterprise AI discovery: Unlock Visibility, Governance & Value

Explore Enterprise AI discovery, its key benefits, and capabilities.

AI discovery overview

Enterprise AI discovery is a fundamental feature of the AI Control Tower offering a unified and comprehensive view of all AI assets—including AI systems, agents, models, prompts, and tools—across various hyperscalars, AI apps, and Agentic AI frameworks via Service Graph Connectors (SGC).

Key Benefits

- **Eliminate Blind Spots:** Aggregate fragmented AI assets for holistic visibility and control.
- **Accelerate Governance:** Centralize compliance, security, and operational efficiency.
- **Drive Innovation:** Enable responsible, scalable, and auditable AI ecosystems where governance and innovation coexist.

Why AI discovery

- Fragmented AI ecosystems create governance and compliance risks.
- Enterprises need real-time visibility into AI usage including shadow AI to ensure security, compliance, and operational efficiency.
- AI discovery empowers organizations to accelerate innovation while reducing risk.

Key Capabilities

- **Unified Registry:** Aggregate AI assets from hyperscalars, AI apps, and Agentic AI frameworks such as AWS Bedrock, Azure AI Foundry, Copilot Studio, Google Cloud Platform (GCP) Vertex AI, ServiceNow agents, and more.

  Note: We identify Assets and CIs (Configuration Item) using their unique source IDs, such as ARN (Amazon Resource Names) for AWS and Product Models are identified by the associated asset.

- **Automatic Discovery:** Seamlessly discover and incorporate agents, models, prompts, and tools into the AI Control Tower's AI asset inventory.
- **Comprehensive Meta data:** Capture detailed information for each asset, including relationships, versions, and operational status.
- **Visibility & Control:** Eliminates silos by aggregating AI assets from all major platforms, enabling holistic management and oversight, across cloud, on-premises, and hybrid environments. The AI steward has visibility into all the connections and integrations that the admin has set to be able to maintain oversight and get access to active connections. Enabling AI stewards to selectively synchronize assets and environments.

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

AI discovery challenges & solution

Enterprise AI adoption is growing rapidly, yet many organizations face challenges with fragmented visibility across internal developments, vendor offerings, and cloud deployments. The AI discovery feature in AI Control Tower automatically detects AI assets on the Now Platform, AWS Bedrock, Azure Machine Learning Services, Azure Cognitive Services, Copilot Studio, and GCP Vertex AI, then adds them to the ServiceNow's Configuration Management Database for unified governance.

This capability helps organizations achieve comprehensive visibility of their AI assets while supporting the expansion of initiatives across hybrid cloud environments.

Business Impact

AI discovery helps organizations reduce AI deployment risks by automatically finding AI assets and simplifying compliance. This feature improves governance by showing AI deployments across different environments, helping executives make informed decisions with full transparency. It also tracks AI usage and adoption, allowing organizations to measure productivity gains.

For customers using both ServiceNow and cloud platforms, AI asset discovery is unified. This simplifies operations and ensures strong visibility.

AI discovery in the AI Control Tower enables enterprises to automatically find, catalog, and manage AI assets across cloud platforms through AI connections.

This setup can be accessed through the AI connections page in AI Control Tower using ServiceNow Graph Connectors to discover AI assets, and usage.

AI connections

Explore the AI connections page and its features.

AI connections overview

The AI connections page enables you to integrate the third-party systems to discover AI assets and track data usage from hyperscalars, AI apps, and agentic AI frameworks via Service Graph Connectors (SGC) to set AI connections. AI stewards can access the AI connections page and on this page, you can find a list of data sources displayed in columns.

Note: The AI connections page appears only when the plugins com.sn_ai_disc and sn_sgc_central are installed.

Agents are discovered according to the set run frequency, but you can also manually discover agents or collect usage data by selecting a connection and choosing the run option. AI stewards are able to perform this action, including activating or deactivating the AI connection by selecting the State column in the list. Admins are able to adjust the run frequency by accessing the connection alias.

Note: Uninstall the AWS AI Discovery plugin before installing the AI Discovery plugin (sn_ai_disc) to use the AI connections.

There are two types of scheduled jobs on AI connections.

Discovery

To discover AI assets from hyperscalars, AI apps, and agentic AI frameworks.

Execution

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

To process usage data and the data is visible on the Performance analytics page for every individual discovered agent.

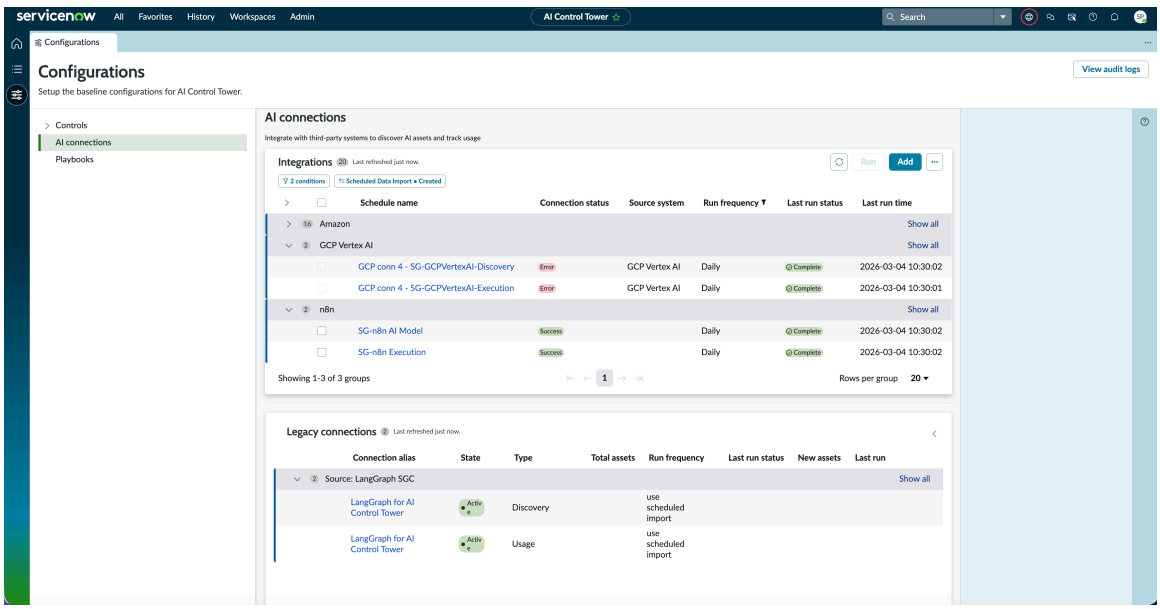
> **Note:** Confirm that the AI discovery daily data collection job is active, which is the key element in collecting the data.

AI connections are a combination of hyperscalars, AI apps, and agentic AI frameworks created using Service Graph Connectors. AI connections can discover AI agents and usage data from the hyperscalars, AI apps, and agentic AI frameworks.

Navigate to the AI connections page in the AI Control Tower, you're able to create AI connections as well as manage the existing ones. The existing connections appear under the section Legacy connections.

Starting March 2026, these AI Service Graph Connectors #are available.

- AWS
- GCP Vertex AI
- LangGraph
- Microsoft
- n8n
- Salesforce



AI connection record

The AI connection record has the following tabs:

Details- displays the connection details of the AI connection.

Data sources- represent a running unit, which defines what data is fetched from a third-party system.

Import schedules- Run the parent data import and execute the job to discover AI Agents and track usage data.

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

Service Graph Connectors for AI Control Tower

You can use Service Graph Connectors (SGC) to import and integrate third-party data into CMDB and non-CMDB tables to create AI connections.

How SGC works

When you access third-party APIs, they return data. This data must be stored in ServiceNow tables. Before saving it, we should transform the data from its JSON format into our preferred format and then store it in ServiceNow tables.

Benefits of SGC

Lets explore some of the major benefits:

- Gain trust in your CMDB data.
- Minimize operational risk.
- Minimize custom integrations.
- Accelerate time to value.

To understand benefits and features of SGC in detail, see Service Graph Connectors

AI Service Graph Connector for Amazon

Use the AI Service Graph Connector for Amazon to discover AI assets such as AI systems, agents, models, tools, and prompts and well as usage data for these AI agents within your ServiceNow AI Control Tower. This usage information is consumed by the AI Control Tower value dashboard.

Download apps from the Store

Visit the ServiceNow Store website to download the AI Service Graph Connector for Amazon app.

Supported ServiceNow versions

- Australia
- Zurich patch 7
- Yokohama patch 11

User roles

Roles required in the ServiceNow environment:

- sn_ai_disc.discovery_admin
- sn_cmdb_int_util.sgc_admin

Prerequisites from AWS

Role required: IAM user

AWS Prerequisites

Before proceeding, ensure you have:

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

- AWS Account- Active AWS account with access to the services you want to connect.
- IAM Credentials: AWS Access Key ID and Secret Access Key with read permissions for the services you plan to migrate.
- Service Access- API access enabled for Amazon Bedrock, Amazon SageMaker, Amazon CloudWatch, and/or Amazon Bedrock AgentCore.

Required IAM Permissions

Your IAM user role or role needs these permissions.

- Amazon Bedrock: bedrock:List*, bedrock:Get*
- Amazon SageMaker: sagemaker:List*, sagemaker:Describe*
- Amazon CloudWatch: logs:DescribeLogGroups, logs:DescribeLogStreams, cloudwatch:GetMetricData
- Amazon Bedrock AgentCore: bedrock:ListAgents, bedrock:GetAgent

AWS Setup documentation

Use these AWS resources to set up credentials and enable services:

- AWS IAM- Managing Access Keys
- Amazon Bedrock- Setting Up
- Amazon Bedrock Agents- Setup Guide
- Amazon CloudWatch- Setup Guide
- Amazon SageMaker- Get Started

Data mapping

The following table lists the data sources, the staging tables, and the target tables #CMDB#CI classes and non-CMDB#classes where data is stored for a#AWS #project.

Data Source	Staging Table	Target Table
SGawsBedrockAIAssetDSUtilSNC	sn_ai_disc_aws_sgc_bedrock_ai_asset	sn_ai_disc_aws_sgc (routes to other stag
SGawsBedrockAISystemDSUtilSNC	sn_ai_disc_aws_sgc_bedrock_ai_system	alm_ai_system_digit
SGawsBedrockAIModelDSUtilSNC	sn_ai_disc_aws_sgc_bedrock_ai_model	alm_ai_model_digita
SGawsBedrockAIToolDSUtilSNC	sn_ai_disc_aws_sgc_bedrock_ai_tool	sn_ent_ai_tool
SGawsBedrockAIPromptDSUtilSNC	sn_ai_disc_aws_sgc_bedrock_ai_prompt	alm_ai_prompt_digi
SGawsBedrockAISbcompM2mDSUtilSNC	sn_ai_disc_aws_sgc_bedrock_sbcomp_m2m	sn_ent_ai_system_s
SGawsBedrockAIUsageDSUtilSNC	sn_ai_disc_aws_sgc_bedrock_ai_usage	sn_ai_disc_ai_usage
SGAgentCoreDataSourceUtil (importAgentRuntimesByID)	sn_ai_disc_aws_sgc_agentcore_ai_system	alm_ai_system_digit
SGAgentCoreDataSourceUtil (importCodeInterpretersByID, importBrowsersByID, importTargetsByID)	sn_ai_disc_aws_sgc_agentcore_ai_tool	sn_ent_ai_tool

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

SGAgentCoreDataSourceUtil (getAWSAgentCoreUsage)	sn_ai_disc_aws_sgc_agentcore_ai_usage	sn_ai_disc_ai_usage
SGSageMakerAIModelDSUtilSNC	sn_ai_disc_aws_sgc_sg_awssagemaker_model	alm_ai_model_digita
SGSageMakerModelCardDSUtilSNC	sn_ai_disc_aws_sgc_sg_awssagemaker_model	alm_ai_model_digita

Create AI connection for Amazon

Create an AI connection for Amazon in AI Control Tower using the AI Service Graph Connector for Amazon.

Before you begin

Role required: sn_ai_disc.discovery_admin and n_cmdb_int_util.sgc_admin

Procedure

1. Navigate to **AI Control Tower** > **Configurations** > **AI connections**.
2. Select **Add**.
3. Select **AWS** from all the available connectors.
4. Click **Create connection**.

   Note: The Review the setup instructions page appears and verifies to follow all the prerequisites.

5. Click **Continue**.
   Setup page appears.
6. Select Source systems.
7. Choose the AWS services that you want to integrate with ServiceNow.
   - Amazon Bedrock
   - Amazon Bedrock AgentCore
   - Amazon SageMaker
8. Click **Submit**.
9. **Configure Bedrock**

   a. Enter the **Connection Name**.

   b. Enter the **Access Key ID**.

   c. Enter the **Secret Access Key**.
   Access keys are long-term credentials for an IAM user or the AWS account root user. Access keys consist of two parts: an access key ID and a secret access key. For detailed information, see how to get access and secret key

   d. Enter the **AWS Region**.

   Note: The region information is available in the navigation bar of the AWS management console.

   e. Click **Update and test connection**.

   f. Click **Continue**.

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

10. **Configure Bedrock import schedule**

    a. Open a parent schedule import.

    b. Select the Active check box.

    c. Select Run according to your preference.

    Note: This is an optional step as the schedule imports run according to the schedule.

11. **Configure CloudWatch logs forBedrock**

    a. Enter the **Connection Name**.

    b. Enter the **Access Key**.

    c. Enter the **Secret Key**.

    d. Enter the **AWS Region**.

    e. Enter the **Log Group Names**.

    f. Click **Create and test connection**.

    g. Click **Continue**.

12. **Configure CloudWatch logs import schedule for Bedrock**

    a. Open a parent schedule import.

    b. Select the Active check box.

    c. Select **Execute Now**.

    Note: This is an optional step as the schedule imports run according to the schedule.

    d. Click **Continue**.

13. **Configure SageMaker**

    a. Enter the **Connection Name**.

    b. Enter the **Access Key ID**.

    c. Enter the **Secret Access Key**.

    d. Enter the **AWS Region**.

    e. Click **Create and test connection**.

    f. Click **Continue**.

14. **Configure SageMaker import schedule**

    a. Open a parent schedule import.

    b. Select the Active check box.

    c. Select Run according to your preference.

    d. Select Execute Now.

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

Note: This is an optional step as the schedule imports run according to the schedule.

e. Click **Continue**.

15. **Configure CloudWatch monitoring for SageMaker**
    a. Enter the **Connection Name**.
    b. Enter the **Access Key**.
    c. Enter the **Secret Key**.
    d. Enter the **AWS Region**.
    e. Click **Create and test connection**.
    f. Click **Continue**.
16. **Configure CloudWatch monitoring import schedules for SageMake**
    a. Open a parent schedule import.
    b. Select Active check box.
    c. Select Run according to your preference.
    d. Select Execute Now.

Note: This is an optional step as the schedule imports run according to the schedule.

e. Click **Continue**.

Result
AI connection is created for AWS.

AI Service Graph Connector for Google Cloud Platform (GCP) Vertex AI

Use the#AI Service Graph Connector for GCP Vertex AI plugin to create AI connections to discover AI assets such as AI systems, models, prompts, and tools as well as usage data for these AI agents within your ServiceNow AI Control Tower. This usage information is consumed by the AI Control Tower value dashboard.

Download apps from the Store

Visit the#ServiceNow Store website to download the AI Service Graph Connector for GCP Vertex AI app.

Supported ServiceNow versions

- Australia
- Zurich patch 7
- Yokohama patch 11

User roles

Roles required in the ServiceNow environment:

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

- sn_ai_disc.discovery_admin
- sn_cmdb_int_util.sgc_admin

Prerequisites from GCP Vertex

Prerequisites:

Review the setup instructions to create a service account, assign roles, bind roles to the service account, enable APIs, create a JKS file, and register it in your ServiceNow instance. For setup instructions and used APIs, see the Service Graph connector for GCP Vertex AI- Setup Instructions [KB2731256]

Data Mapping

The following table lists the data sources, the staging tables, and the target tables#CMDB#CI classes and non-CMDB#classes where data is stored for a#GCP Vertex AI#project.

Data source	Staging table	Target tables
SG-GCPVertexAI-Execution	sn_ai_disc_gcp_sgc_sg_gcp_execution	sn_ai_disc_ai_usage
SG-GCPVertexAI-System	sn_ai_disc_gcp_sgc_sg_gcp_ai_system	cmdb_ai_system_component_ alm_ai_system_digital_asset cmdb_ci_function_ai cmdb_rel_asset_ci
SG-GCPVertexAI-Mode	sn_ai_disc_gcp_sgc_sg_gcp_ai_model	cmdb_ai_model_product_mod alm_ai_model_digital_asset
SG-GCPVertexAI-Tool	sn_ai_disc_gcp_sgc_sg_gcp_ai_tool	sn_ent_ai_tool
SG-GCPVertexAI-Prompt	sn_ai_disc_gcp_sgc_sg_gcp_ai_prompt	cmdb_ai_prompt_product_mod alm_ai_prompt_digital_asset
SG-GCPVertexAI-System Subcomponent M2M	sn_ai_disc_gcp_sgc_sg_gcp_ai_system_subcomponent_m2m	sn_ent_ai_system_subcompon

Create AI connections for GCP Vertex AI

Create an AI connection for GCP Vertex AI in AI Control Tower using the#AI Service Graph Connector for GCP Vertex AI.

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

Before you begin

Role required: sn_ai_disc.discovery_admin and sn_cmdb_int_util.sgc_admin

Procedure

1. Navigate to **AI Control Tower workspace** > **Configuration** > **AI connections**.
2. Select **Add**.
3. Select **GCP Vertex AI** from all the available connectors.
4. Click **Create connection**.
   Review setup instructions page displays
5. Verify to follow all the prerequisite steps.
6. Check the option **I have read the setup instructions**.
7. Click **Continue**.
8. Create **X.509** certificate.
9. Select **New**.
10. Enter the **Name**.
11. Enter the **Key store password**.
12. Click **+Add file**.
13. Select the JKS file.
14. Select **Upload** to upload the JKS file.
15. Click **Save**.
    The JKS file is added.
16. Select **Continue**.
    Setup page appears
17. Create and test connections

    a. Enter the **Connection Name**.

    b. Enter the **Cloud region**.

    c. Enter the **Service Account Email**.

    d. Enter the **Keystore**.

    e. Enter the **Keystore Password**.

    f. Enter the **Organization Id**.

    > **Note:** This field is needed only if you're looking to discover agents related to one organization. This is valid only if the service account email has organization level access.

    g. Click **Create and test connection**.

    h. Click **Continue**.

18. Configure import schedule

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

a. Ensure that both the parent scheduled jobs, Discovery and Execution are active as they are shipped out inactive.

Note: Ensure to execute the Discovery scheduled job first.

b. Set the run frequency.

c. To run frequency by demand, select **Execute now**.

Note: This is an optional step as the schedule imports run according to the schedule.

d. Click **Continue**.

e. Click **View all connections** to view the newly created connection.

Result
AI connection is created for GCP Vertex AI.

AI Service Graph connector for LangGraph

Use the#AI Service Graph Connector for LangGraph to discover AI assets such as AI systems, agents, models, tools, and prompts and well as usage data for these AI agents. This usage information is consumed by the AI Control Tower value dashboard.

Download apps from the Store

Visit the ServiceNow Store #website to download the AI Service Graph Connector for LangGraph app.

Supported ServiceNow versions

- Australia
- Zurich patch 7
- Yokohama patch 11

User roles

Roles required in the ServiceNow environment:

- sn_ai_disc.discovery_admin
- sn_cmdb_int_util.sgc_admin

Prerequisites from LangGraph

Enter the connection details, API Key credentials and update the credential record.

Create an API key for the AI Connector to authenticate with LangGraph. For more information on creating API Key in LangGraph, see Create an API Key

Note: To access the latest and most comprehensive guidelines for managing API keys, see About LangGraph

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

Data Mapping

The following table lists the data sources, the staging tables, and the target tables #CMDB#CI classes and non-CMDB#classes where data is stored for LangGraph project.

Data Source	Import Set Table	Target Table(s)
SG-LangGraph Agents	sn_langgraph_integ_agents	alm_ai_system_digital_asset alm_ai_prompt_digital_asset cmdb_ci_function_ai
SG-LangGraph Usage	sn_langgraph_integ_usage	alm_ai_model_digital_asset sn_ai_disc_ai_usage

Create an API Key for LangGraph

Create an API Key for the AI Connector to authenticate with LangGraph.

Before you begin

Role required: Workspace user (Organization user or Organization editor).

Procedure

1. Navigate to **Settings** > **API Keys in LangSmith**.
2. Click **+API Key**.
   API Key is created.
3. Save the API Key.

> **Note:** Ensure that the API Key grants access to the workspace where you intend to explore AI assets and usage details. Store the API key in a secure and private location.

Result

The API Key is created.

Create AI connection for LangGraph

Use the AI Service Graph Connector for LangGraph plugin to create AI connections to discover AI assets such as AI systems, models, prompts, and tools as well as usage data for these agents. This usage information is consumed by the AI Control Tower value dashboard.

Before you begin

Role required: sn_ai_disc.discovery_admin and n_cmdb_int_util.sgc_admin

Procedure

1. Navigate to **AI Control Tower** > **Configuration** > **AI connection**.
2. Select **Add**.
3. Select **LangGraph** from all the available connectors.
4. Click **Create connection**.
   Setup instructions page displays.
5. Verify to follow all the prerequisite steps.

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

Setup page appears

6. Create and test connection

   a. Enter the **Connection Name**.

   b. Enter the **Connection URL for LangSmith API**.

   c. Enter the **Connection URL for LangSmith Rest API**.

   d. Enter the **API Key**.

   > **Note:** Log in to LangGrapH instance and navigate to settings and select API Keys and create an API Key.

   e. To use a MID Server, select the MID Server check box.

   f. Click **Create and test connection**.

   g. Click **Continue**.

7. Configure import schedule

   a. Ensure that both the parent-scheduled jobs, Discovery and Execution are active as they're shipped out inactive.

   > **Note:** Ensure to execute the Discovery-scheduled job first.

   b. Set the run frequency.

   c. To run frequency by demand, select **Execute now**.

   > **Note:** This is an optional step as the schedule imports run according to the schedule.

   d. Click **Continue**.

   e. Click **View all connections** to view the newly created connection.

Result
AI connection is created for LangGraph.

APIs used for LangGraph connector

Explore the AWS APIs used in AI service Graph Connector for LangGraph.The table lists all the LangSmith API endpoints used by the connector.

The following table lists all the LangSmith API endpoints used by the connector.

API	Endpoint (example)	What it does
List Workspaces	https://api.smith.langchain.com / api/v1/workspaces/	Lists all workspaces at the given endpoint for the current API key or user

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

List Deployments	https://api.host.langchain.com/v2/deployments	Lists all agent deployments for a particular workspace
List Tracer Sessions	https://api.smith.langchain.com/api/v1/sessions	Lists all tracer sessions for a given workspace
Get Run Stats	https://api.smith.langchain.com / api/v1/runs/stats	Gets stats on a particular run, including LLM invocations

This table lists all the LangGraph Agent Deployment APIs used by the connector. The “endpoint_url” below will be specific to each deployment and will be discovered at runtime by the app.

API	Endpoint (example)	What it does
Search Assistants	{{endpoint_url}}/assistants/search	Lists all assistants in each deployment
Search Threads	{{endpoint_url}}/threads/search	Lists all threads for each deployment

Target tables for LangGraph

Target tables for storing Service Graph Connector for LangGraph data.

The LangGraph connector uses following main target tables:

Digital asset tables

- alm_ai_system_digital_asset - AI System digital assets
- alm_ai_prompt_digital_asset - AI Prompt digital assets
- alm_ai_model_digital_asset - AI Model digital assets

Usage table

sn_ai_disc_ai_usage - AI Usage/Execution data.

CMDB table

cmdb_ci_function_ai – AI Function (Agent) CI table.

AI Service Graph Connector for Microsoft

Use the #AI Service Graph Connector for Microsoft#plugin to create AI connections for Azure Foundry and Copilot to discover AI assets such as AI systems, models, prompts, and tools as well as usage data for these AI agents within your ServiceNow AI Control Tower. This usage information is consumed by the AI Control Tower value dashboard.

Download apps from the store

Visit the#ServiceNow Store #website to download the AI Service Graph Connector for Microsoft app.

Supported ServiceNow versions

- Australia
- Zurich patch 7
- Yokohama patch 11

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

User roles

Roles required in the ServiceNow environment:

- sn_ai_disc.discovery_admin
- sn_cmdb_int_util.sgc_admin

Prerequisites from Azure Foundry environment

Setting up Azure Foundry:

- To get the OAuth credentials, refer to the Azure documentation.
- The Azure client application needs at least 'User. Contributor' role on AI Hub (ML services) API. Refer documentation for Azure Hub roles. It also requires a workspace to be created.
- The Azure client application needs at least an 'Azure AI User' role on the Azure Foundry (Cognitive Services) API. Refer documentation for Azure Foundry roles. It also requires an account (resource name) to be created.
- Enable both Discovery and Execution import schedules to discover agents and retrieve execution data.

Prerequisites from Copilot environment

Setting up Copilot

- Register an application in Microsoft Entra ID and Register to obtain the Client ID and Client secret.
- Open the Power Platform Admin Center, select your Copilot environment, then go to Settings > Users+permissions > Application users, add a new app user using the Client ID and assign Basic user and System administrator roles.
- Find the Organization ID, Tenant ID under Settings > Session details.
- Usage is tracked only in non-developer environments.
- Enable both Discovery and Executions to discover agents and retrieve execution data.

Note: For Microsoft, OAuth authentication is supported.

Create AI connection for Azure Foundry

Create an AI connection for Azure Foundry in AI Control Tower using the#AI Service Graph Connector for Microsoft.

Before you begin

Role required: sn_ai_disc.discovery_admin and n_cmdb_int_util.sgc_admin

Procedure

1. Navigate to **AI Control Tower workspace** > **Configurations** > **AI connections**.
2. Select **Add**.
3. Select **AI connector for Microsoft** from all the available connectors.
4. Click **Create connection**.
5. Select Azure Foundry check box.
6. Review setup instructions page displays.

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

Note: Verify to follow all the prerequisite steps.

7. Click **Continue**.
8. Setup page appears
9. **Configure and test ML services connection**
   a. Enter the **Connection Name**.
   b. Enter the **Connection URL** (https://<region>api.azureml.ms).
      Note: The region is available under project details.
   c. Enter the **OAuth client ID**OAuth client ID.
   d. Enter the **OAuth client secret**.
   e. Enter the **OAuth token URL** (https://login.microsoftonline.com/<tenantid>/oauth2/v2.0/token).
      The tenant id can be found in the URL of every page. It's abbreviated as tid.
   f. Click **Update and test connection**
   g. Click **Continue**.
10. **Configure and test AI services connection**
    a. Enter the **Connection Name**.
    b. Enter the **Connection URL** (https://<resource-name>services.ai.azure.com)
       Note: To obtain the resource name, make sure that you're on New Foundry (Enable the New Foundry toggle) and select the project. Once you're on the home page, look for the Project endpoint to view the resource name.
       Starting from March 2026 onwards, ServiceNow provides support to the New Foundry along with the old Foundry.
    c. Enter the **OAuth client ID**.
    d. Enter the **OAuth client secret**.
    e. Enter the **OAuth token URL**.
11. **Configure import schedule**
    a. Ensure that both the parent-scheduled jobs, Discovery and Execution are active as they're shipped out inactive.
       Ensure to execute the Discovery-scheduled job first.

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

b. Select **Execute now** to run.

Note: This is an optional step as the schedule imports run according to the schedule.

c. Click **Continue**.

d. Click **View all connections** to view the newly created connection.

Result
AI connection is created for Azure Foundry.

Properties of Azure Foundry

System properties for AI Service Graph Connector for Microsoft (Azure Foundry specific).

System properties

Property	Description
sn_ai_msft_integ.usage_data_lookback	Number of days to look back for fetching threads before the run collection window start time. Type: Integer Default value: 3 Location: System Property [sys_properties] table
sn_ai_msft_integ.usage_first_run_lookback_days	Number of days to look back for usage data on first run (when no last_success_import_time exists). Type: integer Default value:30 Location: System Property [sys_properties] table

Create AI connection for Copilot

Create an AI connection for Copilot in AI Control Tower using the#AI Service Graph Connector for Microsoft.

Before you begin
Role required: sn_ai_disc.discovery_admin and n_cmdb_int_util.sgc_admin

Procedure

1. Navigate to **AI Control Tower workspace** > **Configurations** > **AI connections**.
2. Select **Add**.
3. Select **AI connector for Microsoft** from all the available connectors.
4. Click **Create connection**.

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

5. Select the Microsoft Copilot check box.
6. Review setup instructions page displays.

   Note: Verify to follow all the prerequisite steps.

7. Expand Create Copilot connection section.
8. Click **Continue**.
9. Setup page appears
10. **Configure and test Copilot connection**
    a. Enter the **Connection name**.
    b. Enter the **Connection URL** (https://<provider-domain-name>.com).
    c. Enter the **OAuth client ID**.
    d. Enter the OAuth client secret.
    e. Enter the OAuth token URL (Tenant id).
    f. Click **Create and test connection**
    g. Click **Continue**.
11. **Configure import schedule**

    a. Ensure that both the parent-scheduled jobs, Discovery and Execution are active as they're shipped out of the box inactive.

       Ensure to execute the Discovery-scheduled job first.

    b. Set the run frequency.

    c. To run frequency by demand, select Execute now to run.

       Note: This is an optional step as the schedule imports run according to the schedule.

    d. Click **Continue**.

    e. Click **View all connections** to view the newly created connection.

Result
AI connection is created for Copilot.

APIs used by Microsoft- Azure & Copilot

Explore the AWS APIs used in AI service Graph Connector for Amazon.

The table lists all the Amazon API endpoints used by the connector.

New Azure Foundry — Discovery & Usage

API	Endpoint	What it does
List Projects	https://management.azure.com/subscriptions/{subscription_id}/resourceGroups/{resource_group}/	Lists all Azure AI Foundry projects in a resource group

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

	providers/Microsoft.CognitiveServices/accounts?api-version=2025-06-01&kind=Project	
List Deployments	https://resource_name.services.ai.azure.com/api/projects/project_name/deployments?api-version=v1	Lists all model deployments within a Foundry project
List Agents	https://resource_name.services.ai.azure.com/api/projects/project_name/agents?api-version=2025-11-15-preview	Lists all agents in a Foundry project
List Agent Versions	https://resource_name.services.ai.azure.com/api/projects/project_name/agents/agent_name/versions?api-version=2025-11-15-preview	Lists all versions of a specific agent
List conversations	https://resource_name.services.ai.azure.com/api/projects/project_name/openai/conversations?api-version=2025-11-15-preview	Lists all conversation threads in a Foundry project
List Conv Items	https://resource_name.services.ai.azure.com/api/projects/project_name/openai/conversations/conv_id/items?api-version=2025-11-15-preview	Lists all messages/items within a specific conversation

AI Services (Classic)

API	Endpoint	What it does
List Projects	https://management.azure.com/subscriptions/{subscription_id}/resourceGroups/{resource_group}/providers/Microsoft.CognitiveServices/accounts?api-version=2025-06-01&kind=Project	Lists all AI Services projects in a resource group
List Deployments	https://resource_name.services.ai.azure.com/api/projects/project_name/deployments?api-version=2025-11-15-preview	Lists all model deployments within a project
List Assistants	https://resource_name.services.ai.azure.com/api/projects/project_name/assistants?api-version=v1	Lists all assistants (agents) in a project List
List Threads	https://resource_name.services.ai.azure.com/api/projects/project_name/threads?limit=limit&api-version=v1	Lists all execution threads in a project

ML Services (AI Hub)

API	Endpoint	**What it does**

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

List Subscriptions	https://management.azure.com/subscriptions?api-version=management_api_version	Lists all Azure subscriptions available to the authenticated principal
List Resource Group	https://management.azure.com/subscriptions/subscription_id/resourceGroups?api-version=management_api_version	Lists all resource groups within a subscription
List MLWorkspaces	https://management.azure.com/subscriptions/%7Bsubscription_id%7D/resourceGroups/%7Bresource_group%7D/providers/Microsoft.MachineLearningServices/workspaces?api-version=2023-04-01	Lists all Azure ML workspaces in a resource group
List Deployments	https://resource_name.services.ai.azure.com/api/projects/project_name/deployments?api-version=2025-11-15-preview	Lists all model deployments within an ML project
List Assistants from ML Services	https://region.api.azureml.ms/agents/v1.0/subscriptions/subscription_id/resourceGroups/resource_group/providers/Microsoft.MachineLearningServices/workspaces/workspace_name/assistants?api-version=v1	Lists all agents/assistants from an AzureML workspace
List Threads from ML Service	https://region.api.azureml.ms/agents/v1.0/subscriptions/subscription_id/resourceGroups/resource_group/providers/Microsoft.MachineLearningServices/workspaces/workspace_name/threads?api-version=v1	Lists all threads from an AzureML workspace
List Thread Runs from ML Service	https://region.api.azureml.ms/agents/v1.0/subscriptions/subscription_id/resourceGroups/resource_group/providers/Microsoft.MachineLearningServices/workspaces/workspace_name/threads/thread_id/runs?api-version=v1	Lists all thread runs within a specific thread in an AzureML workspace

AWS APIs used for Copilot

API	Endpoint	Description
List of all the agents in the Copilot environment	<org-url>/api/data/v9.2/bots	Returns metadata about all Copilot Studio bots (agents) registered in a Dataverse environment, including their name, status, authentication config, and ownership details. This is the primary endpoint to discover and inventory deployed copilots.
List of all the components per agent	<org-url>/api/data/v9.2/botcomponents	Exposes the internal authoring components of a Copilot, such as topics, entities, variables, and trigger phrases, stored in the botcomponent table. Essential for inspecting or programmatically analyzing a bot's conversational logic and structure.
List Conversations	<org url>/api/data/v9.2/conversationtranscripts	Retrieves the full chat/conversation transcripts between users and Copilot Studio bots, stored in the Dataverse conversationtranscript table.

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

Target tables for Microsoft

Target tables for storing Service Graph Connector for Azure and copilot data.

Both Azure Foundry and Copilot data sources use the following main target tables:

Digital Asset Tables

- alm_ai_system_digital_asset- AI System digital assets
- alm_ai_prompt_digital_asset- AI Prompt digital assets

Entity Tables

- sn_ent_ai_tool- AI Tools
- sn_ent_ai_system_subcomponent_m2m- AI System Subcomponent many-to-many relationships

Usage Table

sn_ai_disc_ai_usage- AI Usage/Execution data.

Additional information

- Staging tables are import set tables where data is initially loaded from external sources.
- Target tables are the destination where transformed data is stored.
- Data flows from Data Source → Staging Table → Transform Map → Target Table
- For the tools attached to AI Agents which are discovered, the tool names are suffixed to ensure uniqueness. This is the format for tool name- tool-name:tenantid:resource-name:project-name.

AI Service Graph Connector for n8n

Use the#AI Service Graph Connector for n8n to create AI connections to discover AI assets such as AI systems, models, prompts, and tools as well as usage data for these AI agents.

Download apps from the Store

Visit the#ServiceNow Store website to download the AI Service Graph Connector for n8n app.

Supported ServiceNow versions

- Australia
- Zurich patch 7
- Yokohama patch 11

User roles

Roles required in the ServiceNow environment:

- sn_ai_disc.discovery_admin
- sn_cmdb_int_util.sgc_admin

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

Prerequisites from n8n

Create a n8n API key using your n8n instance.

For more information on creating a n8n API Key, see Create an API Key in n8n environment

Create an API Key in n8n environment

Creating a API key in n8n environment using your n8n instance.

Before you begin
Role required: admin (authenticated user with the appropriate administrative permissions#in the n8n instance)

Procedure

1. Log in to your n8n instance.
2. Navigate to **Settings** > **n8n API**.
3. Select **Create an API Key**.
4. Choose a label and set an Expiration time for the key.
5. Copy the API Key and use this key to authenticate your calls.

Result
The n8n API Key is created.

What to do next
Create an AI connection for n8n

Create an AI connection for n8n

Use the#AI Service Graph Connector for n8n#to discover AI assets such as AI systems, models, prompts, and tools well as usage data for these AI agents. This usage information is consumed by the AI Control Tower value dashboard.

Before you begin
Role required: sn_ai_disc.discovery_admin and sn_cmdb_int_util.sgc_admin

Procedure

1. Navigate to **AI Control Tower** > **Configuration** > **AI connection**.
2. Select **Add**.
3. Select **n8n** from all the available connectors.
4. Click **Create connection**.
   Setup page appears
5. Configure and test connection.

   a. Enter the **Connection Name**.

   b. Enter the **Connection URL**(https://<n8n-instance>)

   c. Enter the **API Key**.

   d. Click **Create and test connection**.

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

e. Click **Continue**.
Setup page appears

6. Configure import schedule

a. Select a parent import schedule job.

b. Select the Active check box.

c. Select run and time to schedule the job.

d. Select all the additional settings if necessary.

e. Ensure that both the parent-scheduled jobs, Discovery and Execution are active as they're shipped out inactive.

Note: Ensure to execute the Discovery-scheduled job first.

f. To run frequency by demand, select **Execute now**

Note: This is an optional step as the schedule imports run according to the schedule.

g. Select **Save**.

h. Click **Continue**.

7. Select **View all connections** to view the newly created connection.

Result
AI connection is created for n8n.

Properties for n8n

Service Graph Connector for n8n#properties control the behavior of the connector.

System properties

Note: To open the System Properties [sys_properties] table, enter#sys_properties.list#n the navigation filter.

System properties

Property	Description
sn_n8n_integ.list_execution_page_size	The maximum number of items to return in each page of sn_n8n_integ.sgn8n_list_execution datastream. This is passed as a query parameter to n8n REST API via the datastream. **Note:** The maximum permitted size is 250 by n8n.

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

System properties (continued)

Property	Description
sn_n8n_integ.list_workflow_page_size	The maximum number of items to return in each page of sn_n8n_integ.sgn8n_list_workflow datastream. This is passed as query parameter to n8n REST API via the datastream. **Note:** The maximum permitted size is 250 by n8n.

Target tables for the n8n

When you complete setting up the connection, you can configure the integration to periodically pull data from a n8n project. The data is saved in tables that extend from the CMDB CI classes and other non-CMDB classes.

Target tables for AI systems

cmdb_ai_system_component_product_model

The following attributes in the n8n Datacenter [cmdb_ai_system_component_product_model] table are populated by collected data.

Attribute label	Attribute name	n8n attribute
Name	name	name (workflow) or name (agent)
Version	version	Type version

alm_ai_system_digital_asset

The following attributes in the n8n Datacenter [alm_ai_system_digital_asset] table are populated by collected data.

Attribute label	Attribute name	n8n attribute
External Ref id	external_ref_id	Connection Url + data or connection Url + data+ nodes
Display Name	display_name	Name
Install Status	Install_status	Data is Archived
AI Models	ai_models	AI Language model

cmdb_ci_function_ai

The following attributes in the n8n Datacenter [cmdb_ci_function_ai] table are populated by collected data.

Attribute label	Attribute name	n8n attribute
Name	name	display name
Object ID	object_id	external reference id

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

Install status	install_status	Install status
Model ID	model_id	model
Vendor	vendor	vendor
Sys id	asset	asset

Target tables for AI model

cmdb_ai_model_product_model

The following attributes in the n8n Datacenter [cmdb_ai_model_product_model] table are populated by collected data.

Attribute label	Attribute name	n8n attribute
Name	name	Model name
Version	version	Type version

alm_ai_model_digital_asset

The following attributes in the n8n Datacenter [alm_ai_model_digital_asset] table are populated by collected data.

Attribute label	Attribute name	n8n attribute
External Ref ID	external_ref_id	Connection URL, model name, type version

Target tables for AI prompts

cdmb_ai_prompt_product_model

The following attributes in the n8n Datacenter [cdmb_ai_prompt_product_model] table are populated by collected data.

Attribute label	Attribute name	n8n attribute
Name	name	Name
Prompt Info	prompt_info	Parameters system message

alm_ai_prompt_digital_asset

The following attributes in the n8n Datacenter [alm_ai_prompt_digital_asset] table are populated by collected data.

Attribute label	Attribute name	n8n attribute
External Ref ID	external_ref_id	Asset type, asset id

Target tables for AI tools

sn_ent_ai_tool

The following attributes in the n8n Datacenter [sn_ent_ai_tool] table are populated by collected data.

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

Attribute label	Attribute name	n8n attribute
External ref id	external_ref_id	workflow +node
Name	name	Nodes +toolType
Description	description	nodes
Active	active	nodes

Target tables for AI Subcomponents

sn_ent_ai_system_subcomponent_m2m

The following attributes in the n8n Datacenter [sn_ent_ai_system_subcomponent_m2m] table are populated by collected data.

Attribute label	Attribute name	n8n attribute
AI System	ai_system	connectionUrl + workflowId + +nodeId or n8n
AI Subcomponent	ai_subcomponent	connectionUrl + workflowId + nodeId (for sub sysems) or workflowId+ nodeId (for tools)
AI subcomponent reference table	ai_subcomponent_reference_table	agents or tools

Target tables for AI Usage

sn_ai_disc_ai_usage

The following attributes in the n8n Datacenter [sn_ai_disc_ai_usage] table are populated by collected data.

Attribute label	Attribute name	n8n attribute
AI System	ai_system	connectionUrl+ workflowId + nodeId
Session ID	session_id	data.id
Time	time	runData.executions
Count	count	runData.executions.length

AI Service Graph connector for Salesforce

Use the#AI Service Graph Connector for Salesforce to discover AI assets such as AI systems, agents, models, tools, and prompts and well as usage data for these AI agents. This usage information is consumed by the AI Control Tower value dashboard.

Download apps from the Store

Visit the ServiceNow Store website to download the AI Service Graph Connector for Salesforce app.

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

Supported ServiceNow versions

- Australia
- Zurich patch 7
- Yokohama patch 11

User role

- sn_ai_disc.discovery_admin
- sn_cmdb_int_util.sgc_admin

Salesforce AI Metadata mapping in ServiceNow tables

AI Agents

alm_ai_system_digital_asset -> Model table (cmdb_ai_system_component_product_model)

Required Fields	ServiceNow (Target)	Salesforce (Staging)
Model ID	Model	model_id
Agent ID	Product instance identifier	agent_id
Version	External record reference	version_id
Status	State (install_status)	u_status
Quantity	Quantity	quantity
Vendor	Vendor	u_vendor
Company	Company	cleansed_manufacturer_ref
Source System	Source System	vendor
Asset Type	Model Category	asset_type

AI Models

alm_ai_model_digital_asset -> Model table (cmdb_ai_model_product_model)

Required Fields	ServiceNow (Target)	Salesforce (Staging)
Asset Type	Model category	asset_type
Foundation Model ID	Base model	foundation_model_identifier
Model	Model	base_model
Status	State (install_status)	u_status
External Reference ID	External record reference	id
Vendor	Vendor	cleansed_vendor_ref
Quantity	Quantity	quantity

AI Tools

AI subcomponent reference table (sn_ent_ai_tool) -> sn_ent_ai_system_subcomponent_m2m -> AI system/Model table (alm_ai_system_digital_asset)

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

Required Fields	ServiceNow (Target)	Salesforce (Staging)
Description	Description	description
Type	Type	type
Source Information	Source information	source
Name	Name	developername
Table Name	Target document table	tablename

AI Prompts

alm_ai_prompt_digital_asset -> Model table (cmdb_ai_prompt_product_model)

Required Fields	ServiceNow (Target)	Salesforce (Staging)
Model Name	Model	model
Description	Prompt information	description
Table Name	Servicenow table	table_name
ID	External record reference	id

AI Agents Usage

sn_ai_disc_ai_usage -> AI system/Model table (alm_ai_system_digital_asset)

Required Fields	ServiceNow (Target)	Salesforce (Staging)
User	User	user
Time	Time	timestamp
Parent ID	Session ID	parentid
Total Invocations	Count	totalinvocations
Conversion Definition ID	AI system	conversationdefinitionid

Create AI connection for Salesforce

Use the AI Service Graph Connector for Salesforce plugin to create AI connections to discover AI assets such as AI systems, models, prompts, and tools as well as usage data for these AI agents within your ServiceNow AI Control Tower. This usage information is consumed by the AI Control Tower value dashboard.

Before you begin

Role required: sn_ai_disc.discovery_admin and n_cmdb_int_util.sgc_admin

Procedure

1. Navigate to **AI Control Tower** > **Configuration** > **AI connection**.
2. Select **Add**.
3. Select **Salesforce** from all the available connectors.
4. Click **Create connection**.
   Review setup instructions page displays

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

5. Verify to follow all the prerequisite steps.
   Setup page appears.
6. Create and test connection
   a. Enter the **Connection Name**.
   b. Enter the **Connection URL**.
   c. Enter the **OAuth Client ID**.
   d. Enter the **OAuth Token URL**.
   e. Click **Create and test connection**.
   f. Click **Continue**.
7. Configure import schedule
   a. Ensure that both the parent-scheduled jobs, Discovery and Execution are active as they're shipped inactive.

      Note: Ensure to execute the Discovery-scheduled job first.

   b. Select Run according to your preference.
   c. To run frequency by demand, select **Execute now**.

      Note: This is an optional step as the schedule imports run according to the schedule.

   d. Click **Continue**.
   e. Click **View all connections** to view the newly created connection.

Result
AI Connection is created for Salesforce.

APIs used by Salesforce

Explore the Salesforce APIs used in AI service Graph Connector for Salesforce. The table lists all the Salesforce API endpoints used by the connector.

The following table lists all the Salesforce API endpoints used by the connector.

API	Endpoint
Bot Definition, BotVersion	https://<domain_name>.salesforce.com/services/data/v65.0/query?q=SELEC BotVersions) FROM BotDefinition WHERE Type != 'Bot' LIMIT 100 OFFSET 100
Configured Einstein Models	https://<domain_name>.salesforce.com/services/data/v65.0/ssot/machine-le

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

GenAiFunctionDefinition	https://<domain_name>.salesforce.com/services/data/v65.0/query?q=SELEC Plugin.Planner.Id, Plugin.Planner.DeveloperName, Plugin.DeveloperName, Plu GenAiFunctionDefinition WHERE PluginId != null LIMIT 100 OFFSET 100
GenAiPluginDefinition, GenAiPluginInstructionDef	https://<domain_name>.salesforce.com//services/data/v65.0/query?q=SELE Description,id,GenAiPluginDefinition.DeveloperName,GenAiPluginDefinition.Id GenAiPluginDefinition.MasterLabel, GenAiPluginDefinition.Planner.Developer
ConversationDefinitionId, ConversationDefinitionEventLog	https://<domain_name>.salesforce.com/services/data/v65.0/query?q=SELEC ConversationDefinitionId, COUNT(Id) TotalInvocations FROM ConversationDe ParentId, CreatedById, ConversationDefinitionId, User ORDER BY MIN(EventD

AI Gateway

Explore the AI Gateway, its value and learn how to use and configure it for your MCP (Model context protocol) transactions.

Model Context Protocol

The Model Context Protocol (MCP) is a standardized client-server protocol that enables AI applications to discover and interact seamlessly with external tools, data sources, and systems. MCP facilitates communication between an AI host application (AI Agent Studio), an MCP Client embedded in the host, and one or more MCP servers that expose specific capabilities such as tools.

For more information on MCP, see Model Context Protocol Client

AI Gateway overview

AI Gateway enables governance, observability, and security for your MCP transactions. AI Gateway is designed to be platform agnostic, delivering the above-mentioned benefits both within ServiceNow (AI Agent studio) and in external agentic studios and hosts.

Why Gateway

Enterprises are building unified AI across vendor boundaries. Their agentic workflows operate across platforms requiring governance.

AI Gateway offers:

- Lifecycle management
- Asset registry
- Identity and access management
- Secure authentication method (OAuth 2.1)
- Observability

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

Value of AI Gateway

With the increasing need of interoperability across different AI solutions across enterprises. AI Gateway provides a way for interoperability governance, observability & security.

Features and benefits

- Enable tight governance and control of cross-platform resources such as MCP servers.
- Enable enterprises to securely connect their agents to remote MCP servers, which work for any agent and MCP server.
- Provide visibility into production metrics such as usage and latency.

Installation and prerequisites

Applications

- Now Assist subscription (required)
- AI Control Tower for Now Assist – Pro Plus license (auto-installs sn_awh_config)
- AI Agent Studio (required for MCP server intake in AI Gateway for AI Control Tower (1.0.5))

Plugins

- sn_ai_governance (version 5.0.6 or higher)
- sn_telemetry_data (version 1.1.10 or higher)

Roles

- AI steward [sn_ai_governance.ai_steward] - for approval workflows and policy management
- AI Agent admin [sn_aia.admin] – for configuring MCP servers as tools in AI Agent Studio

The AI Gateway application is automatically installed for customers who use AI Control Tower for Now Assist.

If you're using any generative AI features, you already have access to AI Gateway. The AI Gateway is included with all types of pro plus licenses. To use the features in AI Gateway, ensure that both AI Agent Studio and AI Control Tower are available in your instance.

What to explore next

- For information on the process flow of MCP servers, see Process flow of MCP servers Via AI Gateway
- For information on AI Gateway setup tab, see MCP server record
- For information on AI Gateway in the configuration page, see AI Gateway
- For information on the roles in AI Control Tower and their role with AI Gateway, see AI Control Tower roles

Process flow of MCP servers Via AI Gateway

Explore the process flow of configuring and introducing MCP servers via AI Gateway in AI Control Tower.

The following steps are performed in the following order to complete the process flow of MCP servers.

Role required: AI steward [sn_ai_governance.ai_steward]

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

1. Adding an MCP server (The MCP server can be added in three different ways)
   - ServiceNow AI Agent studio
   - AI Control Tower
   - MCP Catalog
2. MCP server approval workflow
3. AI Gateway setup and Client Registration

Add an MCP server from ServiceNow AI Agent studio

Add an external MCP server from ServiceNow AI Agent studio.

Before you begin

Role required: AI steward [sn_ai_governance.ai_steward]

Note: The AI steward (sn_ai_governance.ai_steward) role gets inherited with the following roles:

- (AI admin) sn_aia.admin
- (AIG admin) aig_admin
- (MCP Client admin) sn_mcp_client.admin

For more information on the AI steward role and its responsibilities, see AI Control Tower roles

Procedure

1. Navigate to **Workspaces** > **AI Agent Studio** > **Settings** > **Manage MCP servers**.
2. Select **New**.

   Note: In the Create and manage tab of the AI Agent studio, you can add only those MCP servers that have been approved in the AI Control Tower. To view these approvals, go to the Approvals tab on the MCP server record. For more information on Approvals tab and MCP server record, see MCP server record

3. Enter the Name.
4. Select the Authentication type: OAuth 2.1

   Note: AI Gateway supports OAuth 2.1.

5. Select **Next**.
6. Select the Client registration type: Dynamic Client Registration.

   AI Gateway registration in AI Agent Studio supports Dynamic Client Registration.

   If the server allows dynamic client registration, it will automatically retrieve the necessary details. Otherwise, you will need to enter the details manually.

7. Select the **Grant type**.
8. Select **Add**.

   The server gets added and displays the details of the server.

9. Select **Authenticate** to authenticate the server with your credentials.

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

Note: To view the approval request record for every MCP server in the AI Agent studio, click **View Approval record**. This action will direct you to the Related Variants tab of the respective MCP server

To know more about Related Variants tab in the MCP server record, see MCP server record

10. Select **Save** to create an MCP server record.

Note: A scheduled job named **AI Agent Studio to AICT-MCP server sync** runs every 15 minutes to synchronize the MCP servers from AI Agent Studio to AI Control Tower.

Result

After synchronizing, the MCP server shows up in the AI asset inventory with its Status listed as In review, and the lifecycle phase listed as New.

Note: After the MCP server is added to the AI Control Tower, it can't be added again. Additionally, AI Agent Studio will not discover a server that has already been added.

What to do next

The MCP server approval request must be approved. See MCP server approval workflow

Add an MCP server from MCP Catalog

Add an MCP server from the MCP Catalog in the AI Control Tower

Before you begin

Role required: sn_ai_governance.ai_steward

Note: The AI steward (sn_ai_governance.ai_steward) role gets inherited with the following roles:

- (AI admin) sn_aia.admin
- (AIG admin) aig_admin
- (MCP Client admin) sn_mcp_client.admin

Procedure

1. Navigate to **AI assets** > **AI asset inventory** > **MCP servers**.
2. Select **Add**.
3. Select **Choose from MCP Catalog**
4. Search for a server using the search box, or browse available MCP servers displayed as cards with icons, descriptions, and versions.
5. Select a server to view details including:
   - MCP server URL
   - Transport type
   - Available tools and capabilities
6. Click **Select** to import the server.

Note: The system auto-fills Name, MCP server URL, Authentication type, and Transport type.

7. Click **OK** to confirm your server and endpoint selections.

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

Note: Add any additional details.

8. Click **Submit**.

The MCP server is submitted to AI steward review.

Note: You can only add one MCP server at a time.

Result
The MCP server is added to AI Control Tower with status In review and synced to AI Agent Studio within 15 minutes.

Add an MCP server from AI Control Tower

Add an MCP server from the AI Control Tower.

Before you begin
Role required: sn_ai_governance.ai_steward

Note: The AI steward (sn_ai_governance.ai_steward) role gets inherited with the following roles:

- (AI admin) sn_aia.admin
- (AIG admin) aig_admin
- (MCP Client admin) sn_mcp_client.admin

For more information on the AI steward role and its responsibilities, see AI Control Tower roles

Procedure

1. Navigate to **AI assets** > **AI asset inventory** > **MCP servers**.
2. Select **Add**.
3. Enter the Name (Server name).
4. Enter the MCP server URL.
5. Select **Next**.
6. Select the Client registration type.

Note: Selecting Dynamic client Registration, AI Control Tower automatically receives the Client ID and Client Secret. Selecting Manual registration, the client must be created externally on the MCP provider side while the client ID and secret must be entered manually.

7. Review all the other details such as Grant type, and Token authentication method.

Note: There are three sections for the URL:

- Authorization URL: Where the user authenticates and grants access token\
- Token URL: Where the system exchanges and authorization code for an access token
- Token revocation URL: Used to invalidate issued tokens if supported.

8. Select **Submit**.

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

Result
The newly added MCP server is registered in the AI asset inventory MCP servers list and available for approval workflow.

MCP server approval workflow

Explore the MCP server approval workflow.

Before you begin
Role required: AI steward [sn_ai_governance.ai_steward]

Procedure

1. Navigate to **AI Control Tower** > **AI assets** > **Approvals**.
2. Select the MCP server from the asset type filter or navigate to Lifecycle and filter by the current state: **New**.
3. Select the MCP Server record with its Status: **New**.
4. Click **Start Review**.

   Note: Within AI Agent Studio, the MCP Server shows a status of AI Steward review during this phase. Product Owners can select View approval record in AI Agent Studio to go straight to the MCP Server record in AI Control Tower.

   A playbook is triggered to guide the approval process through three phases: Assess, Build and Test, and Deploy.
5. Click **Mark as Complete** for the Assess phase.

   Note: The Asses phase includes Architecture review, Risk Assessment, and Stakeholder review.

6. Complete the Build and Test phase activities.

   Note: The Build and Test phase involves out-of-application testing and validation activities such as running a test AI agent using the MCP Server as a tool and validating output received.

7. Click **Mark as Complete** for the Build and test phase.
8. Click **Mark as Complete** for the Deploy phase.

   Note: The Deploy phase is the final step to deploy and allow the MCP Server for usage.

9. Select **Save**.

   Note:

   The AI Gateway setup tab appears only after an MCP server request is approved. The AI steward can pause transactions whenever needed through this tab. For information on the tabs appearing in the MCP server record, see MCP server record

AI Gateway setup and Client Registration

Setting up AI Gateway tab and creating MCP Client Registration.

Before you begin
Role required: sn_ai_governance.ai_steward

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

Procedure

1. Navigate to **AI assets** > **AI asset inventory** > **MCP servers**.
2. Select and open an approved MCP server.
3. Navigate to the **AI Gateway setup** tab.

   If an AI steward wants to enable ServiceNow Agent Studio for the approved MCP server, they can create a new MCP Client integration.
4. Select **New** under **MCP Client integration**.
5. Enter the **Name**.

   Note: A redirect URL will be available after an agent is created and configured. The redirect URL is a callback point endpoint used during OAuth.

   The redirect URL for the MCP clients is unique to each:

   - AI Agent Studio: https://<instance_name>.service-now.com/oauth_redirect.do
   - Copilot studio: Unique to the tool or connection that customers create inside their Copilot agent.
   - Postman: https://oauth.pstmn.io/v1/browser-callback
   - Other platforms: Consult the platform's OAuth documentation
6. Select **Save**.

Result

A connection to the MCP server is set up via AI Gateway.

Note: The Client ID and Client secret are automatically generated and available in the MCP server record. These credentials are used to authenticate the client when connecting to the Gateway.

To connect with Copilot Studio via AI Gateway, see Connecting with Microsoft Copilot Studio Via AI Gateway

MCP server record

Explore the MCP server record and its associated tabs.

The MCP server record has the following tabs:

1. **Details**: Displays all the asset details along with the MCP server URL at the bottom of the page.
2. **Approvals**: Displays all the approval requests for the MCP server.

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

Note:

When MCP servers are synced to the AI Gateway with the AI Control Tower Core (with Now Assist), an option appears in AI Control Tower to start the approval process. However, with AI Control Tower Pro license version, the button doesn't appear after the synchronization job runs, and the Lifecycle Status is shown as 'None'.

This is a known behavior and as a workaround, confirm to set the lifecycle phase to "New" and the lifecycle status to "AI Steward review" for the relevant sn_ai_governance_asset_governance_details record.

This can be achieved by running a script manually:

```
var assetGovernanceRecord = new
 GlideRecord('sn_ai_governance_asset_governance_details');

if (assetGovernanceRecord.get('<sys_id of asset governance
 record>')) {

    assetGovernanceRecord.setValue('status', 5);


 assetGovernanceRecord.setDisplayValue('lifecycle_phase','N
ew');

    assetGovernanceRecord.update();

}
```

For more information on workflow of MCP server approval request, see MCP server approval workflow

3. **KPIs & Metrics**: Displays the MCP tools observability metrics for each server, which is approved and running. In the KPIs & Metrics tab, the AI steward can view different tools, which have been invoked across gateway usage from different hosts. The tab shows the total number of requests sent, the success rate, and any latency observed.

   Note: Similarly, at the AI Gateway tab in AI Control Tower tab on the landing page of AI Control Tower, you can view the various servers that have been used, along with the total number of requests or transactions and the corresponding success rate.

4. **AI Gateway setup**: The AI Gateway setup tab appears only after an MCP server request is approved. From the AI Gateway setup tab, you can pause the AI Gateway transaction of the MCP server and add an MCP Client integration. The AI Gateway tab is accessible for all AI Control Tower workspace users.

   The AI Gateway setup tab has three sections:

   - **AI Gateway MCP server details** - Displays the following URLs:
     - **AI Gateway MCP server URL**- The Gateway endpoint that clients connect to
     - **Authorization endpoint URL**- OAuth authorization endpoint.
     - **Token endpoint URL**- OAuth token exchange endpoint.

     Note: All the three URLs are generated automatically and can't be modified.

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

If an MCP server is identified as malicious for a certain period, the AI stewards can pause all transactions directed to that server through the AI Gateway. Once the risk is resolved, transactions can be resumed. Similarly, at the AI Gateway level, it's possible to pause all transactions across all servers from the AI Gateway page on the Configurations section of AI Control Tower.

- **MCP Client Integration**- Displays Registered clients for the server. Each agent builder requires its own client registration and MCP client Integration generates a client ID and client secret. If you want to enable ServiceNow Agent Studio for a linear MCP server, then you can create an MCP client integration by selecting **Add**.

  Note: For more information Gateway setup and MCP Client Integration, see AI Gateway setup and Client Registration

- **Global MCP Clients**- Global MCP clients are set up separately from server records. You can configure them on the Configurations page of AI Control Tower in the AI Gateway section. Once they're created, they're automatically available across all MCP servers, so there's no need to set up individual client integrations for each server.

5. **Related Variants**: The Related Variants tab provides a one-to-one mapping of all the MCP server entries from AI Agent Studio, so each entry is visible and trackable in AI Control Tower. In the Related Variants tab, you can have multiple MCP servers of the same variant.

Note: The AI steward can activate sensitive data protection for any transaction across servers by selecting **Activate PII check**. By activating, the AI gateway blocks those transactions, which involve sharing sensitive information to the MCP server.

Add a Global MCP client

Adding a Global MCP client from the AI gateway configuration page.

Before you begin

Role required: AI steward

Procedure

1. Navigate to **AI Control Tower** > **Configurations** > **AI Gateway** > **Global MCP clients**.
2. Select **Add**.
3. Enter then Client ID.
4. Select **Next**.
5. Select Metadata sync mode.

   It is preferable to select Live.

6. Change the Security controls according to the requirements of the agent builder.
7. Click **Submit**.
   A Global MCP client is created.

Result

The MCP client is created and appears in the Global MCP client tab.

Connect to MCP servers Via AI Gateway

Explore an example of connecting to an MCP servers via AI Gateway.

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

Connecting with Microsoft Copilot Studio Via AI Gateway

Connecting with a Microsoft Copilot Studio via AI Gateway.

Before you begin

Role required: Workspace user

Procedure

1. Log in to MCP client (https://copilotstudio.microsoft.com).
2. Creating an agent.
   a. Navigate to **Agents**.
   b. Select **+New agent**.
   c. While creating an agent, navigate to the **Configure** tab.
   d. Enter the Name.
   e. Enter the description.
   f. Click **Create**.
   An agent is created.
3. Adding tools to the agent in Microsoft Copilot Studio
   a. Open the Agent and navigate to **Tools** tab.
   b. Click **+Add a tool**.
   c. Select **Model Context Protocol** icon.
   d. Select **+New tool**.
   e. Select **Model Context Protocol**.
   Displays Add a model Context Protocol server (preview) window.
   f. Enter the Server name, Server description, and Server URL.

      The MCP server URL is the Server URL and the details are available in the AI Gateway setup tab.
   g. Select Authentication: OAuth 2.0.
   h. Select type: Manual.
   i. Enter the Client ID and Client secret.
   The Client ID and Client secret details are available in the MCP server record created for Microsoft Copilot Studio.
   j. Enter the Authorization URL and Token URL template.

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

The Authorization endpoint URL and Token URL template are available on the AI Gateway setup tab.

k. Enter the Refresh URL.
The Refresh URL is the same as the Token URL.

l. Click **Create**.
Your tool gets created successfully and generates a Redirect URL.

4. Copy the Redirect URL and paste in the Redirect URL field in the MCP server record.
5. Click **Save** to save the MCP server record.
6. Navigate to **ALL** > **Application registry** and look for the Copilot MCP server record.
7. Ensure to change the values of **Always use PKCE** and **Public Client** from **True** to **False** of the server record.
This step is only applicable for Microsoft Copilot Studio.
8. Save the Copilot MCP server record.
9. Navigate back to the Copilot Studio.
10. Select **Next**.
11. Select **Create new connection**.
12. Select **Approve**.
13. After getting connected, select **Add to agent**.
14. Open the connection manager and provide authentication to Copilot.
This step is only applicable for Microsoft Copilot Studio.
15. Click **Connect**.
16. Click **Submit**.

Connecting with ServiceNow AI Agent Studio Via AI Gateway

Connecting with ServiceNow AI Agent Studio via AI Gateway.

Before you begin

Role required: an_aia.admin

Note: Product Owners can be restricted to use only approved MCP Servers on AI Agent Studio

Procedure

1. Navigate to **All** > **AI Agent Studio**.
2. Open an existing agent or create agent.
3. Select **Edit**.
4. Select **Tools**.
5. Select **+Add a tool**.
6. Select **+ New tool**.
7. Select **Model Context Protocol**.
8. Select the MCP server from the drop-down.

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

Note: In AWH for AI Control Tower (2.0.0), unapproved MCP Servers appear in the list but tools cannot be fetched until the server is approved in AI Control Tower.

9. Select the specific tools from the MCP server that the agent should use.
10. Select **Save**.

Result
The agent is configured to use the MCP Server through AI Gateway. All requests from the agent to the MCP Server are routed through AI Gateway for governance, security, and observability.

Configuring AI Control Tower workflow

Configure the AI Control Tower workflow.

Configuration overview

1. Activation and installation of AI Control Tower
2. Configure AI Control Tower
3. Configure multi-instance management for AI Control Tower

Note:

Some AI capabilities are available only when the required plugins are installed.

- AI Risk and Asset Management capabilities in AI Control Tower with Now Assist require the AI Risk and Asset Management for Now Assist plugin (sn_aict_irm_aiam), which depends on:
  - AI Risk and Compliance Integration with Control Tower (sn_grc_ai_irm_intg)
  - AI Asset Management (sn_ai_asset_mgmt)
- AI Control Tower supports governance of both enterprise AI assets and ServiceNow AI assets, while AI Control Tower with Now Assist supports governance of ServiceNow AI assets only.
- When AI Control Tower Core (sn_ai_governance) is used with AI Risk and Compliance in a new IRM deployment, the IRM Standard (sn_irm_std) plugin is required to make AI intake request forms available. These intake forms are used to submit requests through the Employee Portal for registering AI systems, AI models, and datasets for governance and risk evaluation.

  This requirement applies only to AI intake request forms and does not apply to AI cases, inquiries, or the Anonymous Reporting Center.

Activation and installation of AI Control Tower

The following section provides an order of installing the applications and their dependent plug-ins.

Installation overview

Installing **AI Control Tower (sn_aict)**, the dependent plug-ins are installed in the following order:

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

1. AI Control Tower Core- This combines AI assets and controls in a central hub, ensuring comprehensive governance and management.
2. AI Asset management- The ability to collect information, track changes and manage lifecycle of AI artifacts like AI systems, models, datasets, prompts etc.
3. AI Risk and Compliance management- It is a comprehensive framework that facilitates end-to-end lifecycle management of AI risks. It supports activities such as risk classification of AI assets, mapping to regulatory authority documents, continuous monitoring, and policy compliance to ensure responsible and accountable AI usage.
4. AI Case management- This module enables tracking, triaging, and resolution of incidents or inquiries related to AI systems. It provides a structured case handling mechanism for AI stewards, compliance officers, and stakeholders to manage AI-related exceptions, investigations, and operational events efficiently.
5. AI Risk and Compliance integration with Control Tower- It establishes a seamless connection between AI Risk and Compliance and the AI Control Tower workspace. This integration ensures unified visibility into AI governance activities, allowing users to monitor risk postures, compliance statuses, and case workflows directly within the Control Tower interface for a centralized oversight experience.
6. AI Risk and Compliance content (optional)- Provides customers to ensure compliance with the applicable laws, regulations, directives and/ or standards and that any content provided in the product is accurate and up to date. Customer acknowledges that the content provided with the product is for ease of use and it is Customer's responsibility to replace the content with the applicable laws, regulations, directives and/ or standards at its own discretion.
7. AI Control Tower- Provides a single, consolidated view of all enterprise AI assets and their relationships, empowering organizations with enhanced integration, governance, and operational efficiency.

Application list

Application	Dependent plugin
AI Control Tower Core	Data Foundation Model (sn_cmdb_foundation:1.1.0)
AI Risk and Compliance management	• AI Control Tower Core (sn_ai_governance:4.0.2) • GRC Feature roles (sn_grc_ftr_role:21.0.1) • GRC: Common workspace elements (sn_grc_workspace:21.0.4) • GRC: Policy and Compliance management (sn_compliance:21.0.2) • Post assessment actions for Smart assessments (sn_smart_imp_auto:20.1.0) • GRC: Risk management (sn_risk:21.0.2) • Regulatory agency library (sn_reg_body_mgmt:21.0.0) • Smart assessment core (sn_smart_asmt:21.0.1)

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

Application list (continued)

Application	Dependent plugin
	• Smart assessment connected (sn_smart_asmt_conn:21.0.1) • Smart assessment designer (sn_smart_asmt_desg:21.0.3)
AI Risk and Compliance integration with Control Tower	• AI Case Management (sn_ai_case_mgmt:21.0.1) • AI Risk and Compliance management (sn_grc_ai_gov:21.0.1) • GRC: Advanced Risk (sn_risk_advanced:21.0.2)
AI Control Tower	• AI Asset Management (sn_ai_asset_mgmt: 2.0.0) • AI Control Tower Core (sn_ai_governance:4.0.2) • AI Risk and Compliance integration with Control Tower (sn_grc_ai_irm_intg:21.0.1) • Engagement dashboard for AI Control Tower (sn_ai_engagement:2.1.6) • Value dashboard for AI Control Tower (sn_ai_value:2.1.6) • AI Discovery (sn_ai_disc:1.0.4)
AI Control Tower for Now Assist 1.01	Gen AI Controller (sn_generative.ai 11.0.9)

When the dependent plug-ins are active, you can look for AI Control Tower application in the ServiceNow store, and download the application when available.

For more information about downloading any application from ServiceNow store, see Download any application from ServiceNow

Configure AI Control Tower

Configuring the AI Control Tower for Now Assist from the **Configurations** page.

Before you begin

Role required: AI steward [sn_ai_governance.ai_steward]

Procedure

1. Navigate to the **Configurations** view in the AI Control Tower.
2. Expand **Data** and select **Opt in** under **Data sharing**.
   To view the **Data** section, confirm that the generative AI Controller plugin is installed.
   The **Opt in** option is enabled by default.

> **Note:** Opting out of data sharing can only be reversed by your Account Executive or the Now Support team.

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

For more information about Data sharing, see the FAQs (common questions) section on the Data Sharing page.

3. Select **Controls** and activate the **Automatically trigger playbooks** option.
   The **Automatically trigger playbooks** option is inactive by default.
   Enabling the **Automatically trigger playbooks** and adding an AI asset to this environment automatically triggers approval requests.

   Note: When the Automatically trigger playbooks isn't enabled, the approval requests aren't generated automatically, however, the asset manager can initiate them manually.

   Verify to have the **Automatically trigger playbooks** option enabled in your production environment.

Result
The AI Control Tower is configured.

Configure multi-instance management for AI Control Tower

Configuring Multi-instance management for AI Control Tower.

Before you begin

Note: The feature is available on Yokohama release onwards.

This feature isn’t supported on Government Community Cloud (GCC) and on-premises instances.

Role required: AI steward [sn_ai_governance.ai_steward]

About this task

Note: Make sure the plugin com.glide.mif.mtls is active. If it isn’t active, install the plugin com.glide.mif.mtls by submitting a support request with Now support for MIF features.

Procedure

1. Log in to all sub-prods and select your prod (managed) instance as manager for the AI Control Tower application.
2. To verify the first step, log in to the prod instance and navigate to the managed instances tab under Multi-instance setup page to view all the sub-prods (managed) instances.

Result
Multi-instance management is configured for the AI Control Tower.

For more information about Trust configuration and Trust configuration management, see Cross-instance application trust configuration

For more information about AI asset synchronization process, see Multi-instance Setup section under Exploring the AI Control Tower configurations.

Configure third-party LLMs using AI Control Tower

Use the AI Control Tower to configure third-party LLMs based on the data routing configuration for the APJC region..

Before you begin
Role required: AI steward [sn_ai_governance_ai_steward]

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

Procedure

1. Navigate to **All** > **AI Control Tower** > **Configurations** > **Controls** > **AI model providers**.
2. Select **Edit**.

   You can configure AI model requests through these types of data routing configuration:

   a. Regional

   b. Global
3. Configure LLMs through Regional data routing.

   a. Select **Regional**.
   All allowed model providers are displayed:
   - Now LLM Service
   - AWS Claude
   - Azure OpenAI
   - Google Gemini

   b. Select the model providers according to your organizational policies.
   When Azure OpenAI is selected, you're asked to Allow Spillover, which redirects the traffic to Global data routing.

   c. Enable spillover by selecting **Yes**.

   d. Select **Save**.
   The Data routing confirmation consent request appears informing you to provide the consent.

   e. Select **Agree and proceed**.

   f. Select **Save**.
   The third-party LLMs are configured with Regional data routing configuration for the APJC region.
4. Configure LLMs through Global data routing.

   a. Select **Global**.
   All allowed model providers are displayed:
   - AWS Claude
   - Azure OpenAI
   - Google Gemini
   - Now LLM Service

   Note: Selecting global data routing redirects your data processing location to a global datacenter.

   b. Select all the model providers.

   c. Select **Save**.

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

Result

The third-party LLMs are configured with the Global data routing configuration for the APJC region.

For information about Data routing, Fallback, and Spillover, see Explore AI model providers

For information about scenarios for configuring third-party LLMs for regulated markets and AMS & EMEA regions, see Explore Third-party LLMs and regions

Using AI Control Tower

Use AI Control Tower for creating governance workflows.

Create an AI Control Tower Playbook workflow

Create a AI Control Tower Playbook workflow for Now Assist approvals to approve or reject an asset.

Before you begin

Important: The AI Control Tower Approval Playbook for the Now Assist approval isn't created manually. They're automatically triggered when an approval request is created while the approval mandate is enabled. To automatically trigger, activate the Automatically trigger playbooks feature under configuration.

Role required: AI steward [sn_ai_governance_ai_steward]

Procedure

1. Navigate to **Workspaces** > **AI Control Tower**.
2. From the AI Control Tower, open the **AI assets** view.
3. From the navigation menu of the AI assets view, navigate to Now Assist **Approvals** and open an Approval record.
4. Under the **Playbook** section of the approval record, you can see the Now Assist **Approvals** workflow.

   You can save the **Playbook** record using the **Save** option or select **Cancel approval request** to cancel the workflow during any step.

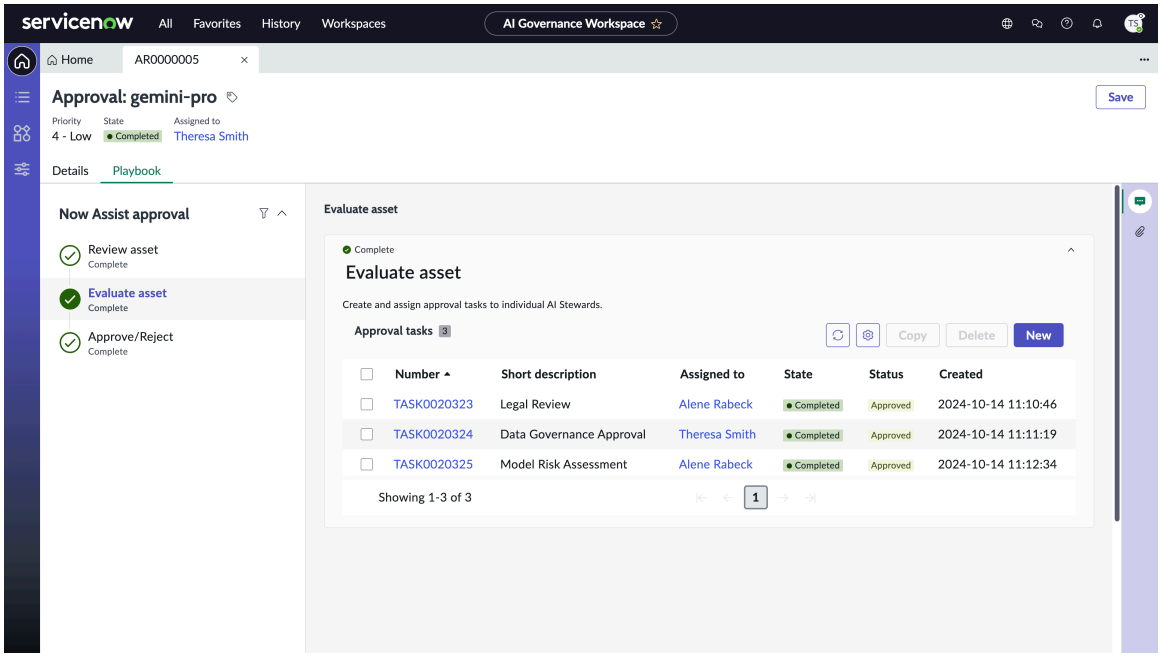
5. **Review asset**- You can view all the details of an asset in review.
6. Select **Next**.
   The approval status changes when the Review asset step is marked as complete.
7. **Evaluate asset**- AI Stewards can create and assign the approval tasks to other individual AI Stewards to evaluate the assets.
   For information on creating a Now Assist approval task for evaluating assets, see Create a Now Assist approval task.

   Note: You can **Copy** or **Delete** the tasks and create a task using the **New** option in the Approval tasks list.

8. **Approval/Reject**- AI Stewards can select the Risk score from the drop-down menu and enter the Close notes, while approving or rejecting an asset based on the evaluation of the tasks.
9. Select **Reject** or **Approve asset** to complete the workflow.

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

The **Risk score** and **Close notes** can't be modified after the **Approve/Reject** step is marked as complete.



What to do next

If you need to activate any skill from Now Assist Admin, an approval request is now required. The AI Control Tower approvals are integrated with both the builder team on the Skill Kit and Now Assist Admin. You may only use models, skills, or AI systems that have been approved. Approval is granted when the "AI steward approval required" option is enabled.

Create a Now Assist approval task

Create an Now Assist approval task to evaluate the assets. The Now Assist approval task is a sub task of an approval request where you will need approvals from other organizations or entities like Legal review, Security review, Data review, Piracy review etc.

Before you begin

Role required: admin

Procedure

1. Navigate to Now Assist **Approvals** and open an Approval record.
2. Under the **Playbook** section, confirm that the **Review asset** step is marked as complete.
3. Select **Create new task** under **Evaluate asset**.
   A new approval task is created. The task number and task type fields are auto-populated.
4. Enter the **Short description**, **Description**, **Due date**, and **Close notes**.
5. Select the **Priority**, **State**, and **Status** of the task from the drop-down menu.
6. Enter and search for **Assigned to** and **Parent approval record**.
   The **Parent approval record** gets auto-populated and appears in a read-only mode.
7. Select **Save**.
   A Now Assist Approval task of an approval request is created.

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

Creating AI assets

You can create AI assets to track and manage the life cycles of your AI systems, AI models, prompts, and datasets.

Create AI system assets

Create AI assets to track and manage the life cycles of your AI systems.

Before you begin

Role required: sn_ai_governance_ai_steward or sn_ai_asset_mgmt.ai_asset_owner

About this task

An AI system is a software artifact that provides AI and machine learning (ML) capabilities to generate outputs, such as predictions, content, recommendations, and decisions, with varying levels of autonomy. Each AI system can be associated with one or more AI models, which may also be associated with one or more prompts and datasets. These AI models, prompts, and datasets enable the AI system to process data and solve complex problems with little to no human intervention.

Procedure

1. Navigate to **Workspaces** > **AI Control Tower**.
2. From the AI Control Tower, open the AI assets view.
3. From the navigation menu of the AI assets view, navigate to either **AI asset inventory - Managed** > **AI Systems** or **AI asset inventory - Unmanaged** > **AI Systems**.
4. Select **Add AI system**.
5. In the Details section of the Add AI system form, fill in the fields.

Details fields in the Add AI system form

Field	Description
Name	Name of the AI system.
Provider	People or organization that developed the AI system.
Version	Version of the AI system.
Description	Brief description of the AI system.
Documentation	Additional information about the AI system, such as the method used to evaluate the efficacy of the AI system.
Managed by	User who is assigned to manage the AI system. This field is automatically set to the user who creates the AI system asset. **Note:** This field is editable only if you have the AI steward (sn_ai_governance_ai_steward) role. If you have the AI asset owner (sn_ai_asset_mgmt.ai_asset_owner) role, this field is read-only.
State	State of the AI system. The options are **Draft** and **Deployed**.
License details	License that you are using for the AI system.
Asset type	Type of AI that your AI system is. Select one of the following options:

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

Field	Description
	○ **Agentic AI** ○ **Classic AI** ○ **Generative AI**
Department	Department that the AI system belongs to.
Supported locations	Locations in which the AI system is supported.

6. Select **Next**.
7. In the Related assets section of the Add AI system form, associate the AI system with other AI assets that are related to it.

   a. Depending on the type of related AI asset that you want associate the AI system with, select one of the following options from the asset type menu:

   - To associate the AI system with any of its supported components or subsystems, select **Sub AI systems**.

     Note: This option is available only if you set the **Asset type** field to either **Agentic AI** or **Generative AI** in the Details section of the Add AI system form.

   - To associate the AI system with any of its integrated AI models, select **AI models**.
   - To associate the AI system with a dataset that is used to test an integrated AI model, select **Evaluation datasets**.
   - To associate the AI system with a prompt that you are providing to an integrated AI model, select **Prompts**.
   - To associate the AI system with any of its integrated AI tools, select **Tools**.

     Note: This option is available only if you set the **Asset type** field to **Agentic AI** in the Details section of the Add AI system form.

   The corresponding asset list opens.

   b. Associate the AI system with the related AI asset.
   Depending on if the related AI asset is currently available in your asset inventory, proceed with one of the following options:

   - To associate the AI system with a related AI asset that is currently available, use the following steps:

     i. Select the Add from inventory drop-down menu on the header of the corresponding asset list.

     ii. Select **Add from inventory**.

     iii. In the Add from inventory dialog box, select the check box for the related AI asset that you want to associate the AI system with.

     Note: If you want to associate the AI system with multiple related AI assets, select the check box for each asset.

     iv. Select **Add**.
     The dialog box closes and the selected AI asset appears in the corresponding asset list.

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

- To associate the AI system with a related AI asset that is not currently available, use the following steps:

  Note: This option is not available for AI tools.

  i. Select the Add from inventory drop-down menu on the header of the corresponding asset list.

  ii. Select **Create**.

  iii. Create the AI asset.

  - If you are creating an AI system asset, return to step 5.
  - If you are creating an AI model asset, follow the proceeding steps in Create AI model assets.
  - If you are creating a prompt asset, follow the proceeding steps in Create prompt assets.
  - If you are creating a dataset asset, follow the proceeding steps in Create dataset assets.

  The new AI asset appears in the corresponding asset list.

8. Select **Next**.
9. In the Use & purpose section of the Add AI system form, fill in the fields.
   This section of the Add AI system form enables you to specify the intended use and purpose of the AI system.

Use & purpose fields in the Add AI system form

Field	Description
Intended outcome of the AI system	Primary business outcome that the AI system is expected to achieve. Select one of the following options: ○ **Not Applicable** ○ **Efficiency Boost** ○ **Quality Enhancement** ○ **Decision Guidance** ○ **Automation of Tasks** ○ **Customer Experience Upgrade** ○ **Insight Generation**
Interaction type with end users	Type of interaction between end users and the AI system, including whether outputs are visible, actionable, or interactive. Select one of the following options: ○ **Not Applicable** ○ **No Direct Interaction** ○ **Background Support** ○ **Notifications & Prompts** ○ **User-Facing Recommendations** ○ **Chat-Based Interaction** ○ **Interactive Experience**

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

Field	Description
Level of human involvement	Capacity at which users can guide, review, and accept AI system activities during operation. Select one of the following options: ○ **Not Applicable** ○ **Full User Control** ○ **User-Guided with AI Support** ○ **Shared Control** ○ **AI-Initiated with User Approval** ○ **Full Automated Workflow**
System autonomy level	Extent to which the AI system can initiate actions or decisions without human intervention. Select one of the following options: ○ **None** ○ **Assistive (AI suggests)** ○ **Semi-Automated (acts with confirmation)** ○ **Condition-Based Automation** ○ **Event-Triggered Automation** ○ **Fully Automated Execution**
Type of output produced	Type of output that the AI system produces for users, systems, or downstream processes. Select any of the following options: ○ **Automated Decisions** ○ **Generated Content** ○ **Insight & Summaries** ○ **Ranking & Scores** ○ **Recommendations** ○ **Simple Alerts** ○ **System Actions**
Area where the AI system is used	Business or operational areas where the AI system is used or provides value. Select any of the following options: ○ **Customer services** ○ **External Partner Ecosystem** ○ **Finance & Accounting** ○ **HR & Workforce** ○ **Internal Operations** ○ **IT & Security** ○ **Sales & Marketing** ○ **Supply Chain**
People affected by the AI system	Users who may be directly or indirectly affected by the AI system outputs or decisions. Select any of the following options: ○ **External Partners** ○ **General Customer Base**

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

Field	Description
	○ **Internal Team** ○ **Public or Large Audiences** ○ **Specific Customer Groups**
Data used by the system	Types of data that the AI system uses to process inputs and generate outputs. Select any of the following options: ○ **Behavioral or Usage Data** ○ **Business Operational Data** ○ **Customer Interaction Data** ○ **Profile or Account Data** ○ **Public or General Info** ○ **Sensitive Business Data**
Additional use and purpose details	Additional information or context that helps clarify the specific use and purpose of the AI system.

For more information on classifying AI systems based on regulatory risk at intake by applying a configured Risk Assessment Methodology (RAM), see Assessment templates and risk assessment methodologies and Request an AI use case.

10. Select **Submit for review**.

> **Note:** If you aren't ready to submit the AI system for review, select **Save** instead. You can then edit it and submit it for review later by navigating to **AI assets** > **Lifecycle** > **Onboard** and then selecting the AI system from the list. When the AI asset record opens, continue editing the AI system details.

Result
The AI system is added to your AI asset inventory. It automatically enters the onboard stage of the AI asset life cycle and is updated with a Lifecycle status of AI steward review.

What to do next
Users who are assigned the AI steward (sn_ai_governance.ai_steward) role can start the review process to begin tracking and managing the life cycle of the AI system. For detailed instructions, see Complete AI asset lifecycle.

Create AI model assets

Create AI assets to track and manage the life cycles of your AI models.

Before you begin
Role required: sn_ai_governance_ai_steward or sn_ai_asset_mgmt.ai_asset_owner

About this task
An AI model is a program that is trained to process data and generate outputs, such as predictions, content, recommendations, and decisions, without human intervention. You can use an AI model to perform a specific task, such as image recognition, data classification, or price prediction.

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

Procedure

1. Navigate to **Workspaces** > **AI Control Tower**.
2. From the AI Control Tower, open the AI assets view.
3. From the navigation menu of the AI assets view, navigate to either **AI asset inventory - Managed** > **AI Models** or **AI asset inventory - Unmanaged** > **AI Models**.
4. Select **Add AI model**.
5. In the Details section of the Add AI model form, fill in the fields.

Details fields in the Add AI model form

Field	Description
Name	Name of the AI model.
Provider	People or organization that developed the AI model.
Version	Version of the AI model.
Supported languages	Languages that are supported by the AI model.
Deployment guidelines	Guidelines or instructions on how to deploy the AI model.
Training procedure	Method that you're using to train the AI model.
Context window	Maximum number of tokens that the AI model can process and recall when generating outputs.
Model size in MB	Size of the AI model in megabytes (MB).
Model parameters info	Internal variables that the AI model learns during training to process data and generate outputs.
Description	Brief description of the AI model.
Managed by	User who is assigned to manage the AI model. This field is automatically set to the user who creates the AI model asset. **Note:** This field is editable only if you have the AI steward (sn_ai_governance_ai_steward) role. If you have the AI asset owner (sn_ai_asset_mgmt.ai_asset_owner) role, this field is read-only.
State	State of the AI model. The options are **Draft** and **Deployed**.
License details	License that you are using for the AI model.
Base model	Base model, also known as a foundation model or pre-trained model, that the AI model was built from.
Department	Department that the AI model belongs to.
Supported locations	Locations in which the AI model is supported.

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

Field	Description
Model card	Brief document that describes important information about the AI model, including the context, intended use, training data, and limitations of the model.
Model weights info	Numerical parameters that define how the AI model learns during training so that it can generate more desired outputs.
Required infrastructure	Integrated software and hardware components that are required for developing, deploying, and managing the AI model.
Evaluation metrics report	Report that provides data for the metrics that you're using to evaluate the effectiveness of the AI model.

6. Select **Next**.
7. In the Related assets section of the Add AI model form, associate the AI model with related datasets.
   Related datasets include any collections of data that can help test or train the AI model.

   a. Depending on the type of related dataset that you want associate the AI model with, select one of the following options from the asset type menu:
      - To associate the AI model with a dataset that can help test it, select **Evaluation datasets**.
      - To associate the AI model with a dataset that can help train it, select **Training datasets**.

      The corresponding asset list opens.

   b. Associate the AI model with the related dataset.
      Depending on if the related dataset is currently available in your asset inventory, proceed with one of the following options:
      - To associate the AI model with a related dataset that is currently available, use the following steps:
        i. Select the Add from inventory drop-down menu on the header of the corresponding asset list.
        ii. Select **Add from inventory**.
        iii. In the Add from inventory dialog box, select the check box for the related dataset that you want to associate the AI model with.

        > **Note:** If you want to associate the AI model with multiple related datasets, select the check box for each dataset.

        iv. Select **Add**.
        The dialog box closes and the selected dataset appears in the corresponding asset list.
      - To associate the AI model with a related dataset that is not currently available, use the following steps:
        i. Select the Add from inventory drop-down menu on the header of the corresponding asset list.
        ii. Select **Create**.
        iii. Create the dataset by following the proceeding steps in Create dataset assets.
        The new dataset appears in the corresponding asset list.

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

8. Select **Submit for review**.

> **Note:** If you aren't ready to submit the AI model for review, select **Save** instead. You can then edit it and submit it for review later by navigating to **AI assets** > **Lifecycle** > **Onboard** and then selecting the AI model from the list. When the AI asset record opens, continue editing the AI model details.

Result
The AI model is added to your AI asset inventory. It automatically enters the onboard stage of the AI asset life cycle and is updated with a Lifecycle status of AI steward review.

What to do next
Users who are assigned the AI steward (sn_ai_governance.ai_steward) role can start the review process to begin tracking and managing the life cycle of the AI model. For detailed instructions, see Complete AI asset lifecycle.

Create prompt assets

Create AI assets to track and manage the life cycles of your prompts.

Before you begin
Role required: sn_ai_governance_ai_steward or sn_ai_asset_mgmt.ai_asset_owner

About this task
A prompt is the instructive input that you provide to AI models to elicit specific responses or outputs. The more detailed and descriptive a prompt is, the higher the quality of the response or output. Examples of prompts include instructions, questions, and commands.

Procedure
1. Navigate to **Workspaces** > **AI Control Tower**.
2. From the AI Control Tower, open the AI assets view.
3. From the navigation menu of the AI assets view, navigate to either **AI asset inventory - Managed** > **Prompts** or **AI asset inventory - Unmanaged** > **Prompts**.
4. Select **Add prompt**.
5. In the Details section of the Add prompt form, fill in the fields.

Details fields in the Add prompt form

Field	Description
Name	Name of the prompt.
Provider	People or organization that developed the prompt.
Version	Version of the prompt.
Description	Brief description of the prompt.
Documentation	Additional information about the prompt, such as the method used to evaluate the accuracy and quality of AI model responses based on the given input.
Managed by	User who is assigned to manage the prompt. This field is automatically set to the user who creates the prompt asset.

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

Field	Description
	Note: This field is editable only if you have the AI steward (sn_ai_governance_ai_steward) role. If you have the AI asset owner (sn_ai_asset_mgmt.ai_asset_owner) role, this field is read-only.
State	State of the prompt. The options are **Draft** and **Deployed**.
Prompt information	Input that you want to provide to AI models.

6. Select **Next**.
7. In the Related assets section of the Add prompt form, associate the prompt with related AI models.

   Related AI models include any AI models that you can provide the prompt to. These AI models can process inputs from the prompt and help generate appropriate outputs.

   Depending on if the related AI model is currently available in your asset inventory, proceed with one of the following options:

   - To associate the prompt with a related AI model that is currently available, use the following steps:

     a. Select the Add from inventory drop-down menu on the header of the AI models list.

     b. Select **Add from inventory**.

     c. In the Add from inventory dialog box, select the check box for the related AI model that you want to associate the prompt with.

     Note: If you want to associate the prompt with multiple related AI models, select the check box for each model.

     d. Select **Add**.
     The dialog box closes and the selected AI model appears in the AI models list.

   - To associate the prompt with a related AI model that is not currently available, use the following steps:

     a. Select the Add from inventory drop-down menu on the header of the AI models list.

     b. Select **Create**.

     c. Create the AI model by following the proceeding steps in Create AI model assets.
     The new AI model appears in the AI models list.

8. Select **Submit for review**.

   Note: If you aren't ready to submit the prompt for review, select **Save** instead. You can then edit it and submit it for review later by navigating to **AI assets** > **Lifecycle** > **Onboard** and then selecting the prompt from the list. When the AI asset record opens, continue editing the prompt details.

Result
The prompt is added to your AI asset inventory. It automatically enters the onboard stage of the AI asset life cycle and is updated with a Lifecycle status of AI steward review.

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

What to do next
Users who are assigned the AI steward (sn_ai_governance.ai_steward) role can start the review process to begin tracking and managing the life cycle of the prompt. For detailed instructions, see Complete AI asset lifecycle.

Create dataset assets

Create AI assets to track and manage the life cycles of your datasets.

Before you begin
Role required: sn_ai_governance_ai_steward or sn_ai_asset_mgmt.ai_asset_owner

About this task
A dataset is a collection of structured or unstructured data that you can use to train and test AI models. Datasets can help you improve the performance and efficacy of your AI models.

Procedure
1. Navigate to **Workspaces** > **AI Control Tower**.
2. From the AI Control Tower, open the AI assets view.
3. From the navigation menu of the AI assets view, navigate to either **AI asset inventory - Managed** > **Datasets** or **AI asset inventory - Unmanaged** > **Datasets**.
4. Select **Add dataset**.
5. On the form, fill in the fields.

Add dataset form

Field	Description
Name	Name of the dataset.
Provider	People or organization that developed the dataset.
Acceptable usage	Acceptable use for the dataset. You can specify more than one use.
Version	Version of the dataset.
Data type	Type of data that the dataset contains. You can specify more than one data type.
Source	Source of the dataset, such as internal customer data within an organization or publicly available data from a government agency.
Description	Brief description of the dataset.
Managed by	User who is assigned to manage the dataset. This field is automatically set to the user who creates the dataset asset. **Note:** This field is editable only if you have the AI steward (sn_ai_governance_ai_steward) role. If you have the AI asset owner (sn_ai_asset_mgmt.ai_asset_owner) role, this field is read-only.
State	State of the dataset. The options are **Draft** and **Deployed**.
Creation type	Method by which the dataset was created. Select one of the following options:

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

Field	Description
	○ **Curated**: The dataset was created through careful selection, organization, and refinement of data. ○ **Derived**: The dataset was created by processing, transforming, or combining data from existing datasets. ○ **Synthetic**: The dataset was created by using artificial data that mimics real-word data.
Base datasets	Base dataset that the dataset was built from.
Department	Department that the dataset belongs to.
Dataset creation date	Date and time that the dataset was created.
Dataset card	Brief document that describes important information about the dataset, including the context, intended use, limitations, and potential biases of the dataset.

6. Select **Submit for review**.

> **Note:** If you aren't ready to submit the dataset for review, select **Save** instead. You can then edit it and submit it for review later by navigating to **AI assets** > **Lifecycle** > **Onboard** and then selecting the dataset from the list. When the AI asset record opens, continue editing the dataset details.

Result
The dataset is added to your AI asset inventory. It automatically enters the onboard stage of the AI asset life cycle and is updated with a Lifecycle status of AI steward review.

What to do next
Users who are assigned the AI steward (sn_ai_governance.ai_steward) role can start the review process to begin tracking and managing the life cycle of the dataset. For detailed instructions, see Complete AI asset lifecycle.

Creating requests for AI assets

You can manage deployed AI assets effectively by creating change requests for updates to existing assets or by creating offboarding requests for retiring assets.

Create change requests for AI assets

Create a change request to modify the relationships between a deployed AI asset and its related assets.

Before you begin
Role required:

- AI steward (sn_ai_governance_ai_steward)
- AI asset owner (sn_ai_asset_mgmt.ai_asset_owner)

About this task
To modify the relationships between related AI assets, a user with the AI asset owner (sn_ai_asset_mgmt.ai_asset_owner) role can initiate the change process by creating and submitting a change request. After the request is submitted, a user with the AI steward

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

(sn_ai_governance_ai_steward) role can review and approve the request. This approval triggers the AI Control Tower application to create a new AI asset record for the asset that the change request was created for. The asset then automatically enters the onboarding stage of the asset life cycle.

You can create change requests for the following AI asset types:

- AI systems (classic, generative, and agentic)
- AI models
- Datasets

Procedure

1. Navigate to **Workspaces** > **AI Control Tower**.
2. If you have the AI asset owner (sn_ai_asset_mgmt.ai_asset_owner) role, create a change request.

   a. Initiate the request by using one of the following navigation options.

Navigation option	Procedure
AI Control Tower Home view	i. From the Home view of the AI Control Tower workspace, select the **Create request** drop-down menu. ii. Select **Create a change request**.
Change requests list	i. From the AI Control Tower workspace, open the AI assets view. ii. From the navigation menu of the AI assets view, navigate to **Requests** > **Change requests**. iii. Select **Create change request**.
AI asset record	i. From the AI Control Tower workspace, open the AI assets view. ii. From the navigation menu of the AI assets view, locate the **AI asset inventory - Managed** section and then select the subsection for the type of AI asset that you want to create a change request for. Alternatively, navigate to **Lifecycle** > **Deploy**. iii. Select an AI asset that meets the following criteria: ▪ The State is set to Deployed. ▪ The Lifecycle phase is set to Deploy. ▪ The Lifecycle status is set to Deployed. iv. Select the **Create a request** drop-down menu. v. Select **Create a change request**.

   b. On the form, fill in the fields.

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

New change request form

Field	Description
Change details	
Asset in review	The AI Asset that you want to modify related asset relationships for. **Note:** If you initiated the change request from an AI asset record, this field populates automatically.
Version	Updated version number of the AI asset.
Due date	Date and time at which the request must be completed.
Justification	Justification for creating the request.
Sub AI systems **Note:** This form section appears only if you are creating a change request for an AI system with an Asset type of either Generative AI or Agentic AI.	
Updated Sub AI systems	Related AI components or subsystems that you want to associate the AI asset with.
AI models **Note:** This form section appears only if you are creating a change request for an AI system.	
Updated models	Related AI models that you want to associate the AI asset with.
AI prompts **Note:** This form section appears only if you are creating a change request for an AI system.	
Updated prompts	Related prompts that you want to associate the AI asset with.
Evaluation dataset **Note:** This form section appears only if you are creating a change request for an AI system or AI model.	
Updated evaluation datasets	Related evaluation datasets that you want to associate the AI asset with.
Tools **Note:** This form section appears only if you are creating a change request for an agentic AI system.	
Updated tools	Related tools that you want to associate the AI asset with.
Training datasets	

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

Field	Description
Note: This form section appears only if you are creating a change request for an AI model.	
Updated training datasets	Related training datasets that you want to associate the AI asset with.
Source **Note:** This form section appears only if you are creating a change request for a dataset.	
Updated source	Related data source that you want to associate the AI asset with.
Base datasets **Note:** This form section appears only if you are creating a change request for a dataset.	
Updated base datasets	Related base dataset that you want to associate the AI asset with.
Dataset card **Note:** This form section appears only if you are creating a change request for a dataset.	
Updated dataset card	Related dataset card that you want to associate the AI asset with.

c. Select **Submit for review**.

d. In the Submit change request dialog box, select **Submit request**.
After the dialog box closes, you are automatically redirected to the new change request record.

3. If you have the AI steward (sn_ai_governance_ai_steward) role, approve or reject the request.

a. Open the change request record by using one of the following navigation options.

Navigation option	Procedure
Change requests list	**i.** From the AI Control Tower workspace, open the AI assets view. **ii.** From the navigation menu of the AI assets view, navigate to **Requests** > **Change requests**. **iii.** From the list of available change requests, select the record number for the request that you want to approve.

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

Navigation option	Procedure
AI asset record	i. From the AI Control Tower workspace, open the AI assets view. ii. From the navigation menu of the AI assets view, locate the **AI asset inventory - Managed** section and then select the subsection for the type of AI asset that the change request was created for. Alternatively, navigate to **Lifecycle** > **Deploy**. iii. Select the AI asset that the change request was created for. iv. On the AI asset record, select the **Requests** tab. v. Select **Change requests** from the request types menu. vi. From the list of available change requests, select the record number for the request that you want to approve.

The change request record opens.

b. On the **Details** tab, set the **Assigned to** field to the user who you want to assign the change request to.
You can assign the request to yourself or to any other user with the AI steward (sn_ai_governance_ai_steward) role.

c. If you are the assigned user, approve or reject the request.

- To approve the request, select **Approve**.
- To reject the request, select **Reject**.

- If you approved the request, the Status changes to Approved and the State changes to In progress.

  The AI Control Tower application then generates a change task for each asset that is impacted by the request. These change tasks appear on the **Change tasks** tab and specify the actions that must be taken on the impacted assets. Each change task is assigned to a user with the AI asset owner (sn_ai_asset_mgmt.ai_asset_owner) role. If you have any concerns or need further clarification on the request, you can create additional change tasks manually. After the assigned users complete all change tasks, the AI Control Tower application automatically creates a new AI asset record for the asset that the change request was created for. The State of the request then changes to Completed.

  Note: The **Change tasks** tab is not available if the change request was created for a dataset.

- If you rejected the request, the Status changes to Rejected and the State changes to Completed.

Create offboarding requests for AI assets

Create an offboarding request to retire AI assets that are no longer needed.

Before you begin

Role required:

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

- AI asset owner (sn_ai_asset_mgmt.ai_asset_owner)
- AI steward (sn_ai_governance_ai_steward)

About this task

To retire a deployed AI asset, a user with the AI asset owner (sn_ai_asset_mgmt.ai_asset_owner) role can initiate the offboarding process by creating and submitting an offboarding request. After the request is submitted, a user with the AI steward (sn_ai_governance_ai_steward) role can review and approves the request. This approval triggers the offboarding workflow that assigns tasks to impacted asset owners and updates the Status of the AI asset record to Retired upon completion.

Note: We have automated the ServiceNow AI model to initiate the offboarding workflows when a deprecation date is present.

You can create offboarding requests for the following AI asset types:

- AI systems (classic, generative, and agentic)
- AI models
- Datasets
- MCP servers

Procedure

1. Navigate to **Workspaces** > **AI Control Tower**.
2. If you have the AI asset owner (sn_ai_asset_mgmt.ai_asset_owner) role, create an offboarding request.

   a. Initiate the request by using one of the following navigation options.

Navigation option	Procedure
AI Control Tower Home view	i. From the Home view of the AI Control Tower workspace, select the **Create request** drop-down menu. ii. Select **Create an offboarding request**.
Offboarding requests list	i. From the AI Control Tower workspace, open the AI assets view. ii. From the navigation menu of the AI assets view, navigate to **Requests** > **Offboarding requests**. iii. Select **Create offboarding request**.
AI asset record	i. From the AI Control Tower workspace, open the AI assets view. ii. From the navigation menu of the AI assets view, locate the **AI asset inventory - Managed** section and then select the subsection for the type of AI asset that you want to retire. Alternatively, navigate to **Lifecycle** > **Deploy**. iii. Select an AI asset that meets the following criteria:

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

Navigation option	Procedure
	▪ The State is set to Deployed. ▪ The Lifecycle phase is set to Deploy. ▪ The Lifecycle status is set to Deployed. **iv.** Select the **Create a request** drop-down menu. **v.** Select **Create an offboarding request**.

b. On the form, fill in the fields.

New offboarding request form

Field	Description
Asset in review	**Note:** If you initiated the offboarding request from an AI asset record, this field populates automatically.
Due date	Date and time at which the request must be completed.
Justification	Justification for creating the request.

c. Select **Submit for review**.

d. In the Submit offboarding request dialog box, select **Submit request**.
After the dialog box closes, you are automatically redirected to the new offboarding request record.

3. If you have the AI steward (sn_ai_governance_ai_steward) role, review the request.

a. Open the offboarding request record by using one of the following navigation options.

Navigation option	Procedure
Offboarding requests list	**i.** From the AI Control Tower workspace, open the AI assets view. **ii.** From the navigation menu of the AI assets view, navigate to **Requests** > **Offboarding requests**. **iii.** From the list of available offboarding requests, select the record number for the request that you want to review.
AI asset record	**i.** From the AI Control Tower workspace, open the AI assets view. **ii.** From the navigation menu of the AI assets view, locate the **AI asset inventory - Managed** section and then select the subsection for the type of AI asset that the offboarding request was created for. Alternatively, navigate to **Lifecycle** > **Deploy**. **iii.** Select the AI asset that the offboarding request was created for. **iv.** On the AI asset record, select the **Requests** tab.

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

Navigation option	Procedure
	v. Select **Offboarding requests** from the request types menu. vi. From the list of available offboarding requests, select the record number for the request that you want to review.

The offboarding request record opens.

b. On the **Details** tab, set the **Assigned to** field to the user who you want to assign the offboarding request to.
You can assign the request to yourself or to any other user with the AI steward (sn_ai_governance_ai_steward) or AI asset owner (sn_ai_asset_mgmt.ai_asset_owner) role.

c. If you are the assigned user, select **Start review**.

The **Offboarding workflow** tab appears and then the offboarding workflow initiates. In addition, the **Impacted assets** tab automatically updates with the list of assets that are impacted by the request.

The offboarding workflow contains the Assess, Pre-offboarding tasks, and Offboard stages. You can create and delete relevant tasks in all three stages. Any tasks that are generated for impacted assets also appear in the corresponding workflow stages.

Note: Each user who is assigned to an impacted asset must complete the corresponding task by either accepting or rejecting it.

- If the user accepts the task, the Status task changes to Accepted and the State changes to Completed.
- If the user rejects the task, the Status of the task changes to Rejected and the State changes to Completed.

d. Complete all tasks in the Assess stage.

e. Select **Mark as complete** to move to the Pre-offboarding tasks stage.

f. Complete all tasks in the Pre-offboarding tasks stage.

g. Select **Mark as complete** to move to the Offboard stage.

h. Complete all tasks in the Offboard stage.

i. Select **Mark as complete**.

After all three stages are marked as complete, the Status of the request changes to Approved and the State changes to Completed. The State of the asset then changes to Retired.

View AI assets by life-cycle stage

View AI assets based on the AI asset life-cycle stage that they are currently in. Use this information to determine which AI assets require your attention.

Before you begin

Role required: sn_ai_governance.ai_steward or sn_ai_asset_mgmt.ai_asset_owner

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

Procedure

1. Navigate to **Workspaces** > **AI Control Tower**.
2. From the AI Control Tower, open the AI assets view.
3. From the navigation menu of the AI assets view, locate the **Lifecycle** section.
   You can select **Cancel** to cancel AI asset.
4. Depending on the life-cycle stage that you want to view AI assets for, select one of the following Lifecycle subsections:
   - To view the list of AI assets that are currently in the New stage, select **New**.
   - To view the list of AI assets that are currently in the assess stage, select **Assess**.
   - To view the list of AI assets that are currently in the build and test stage, select **Build and test**.
   - To view the list of AI assets that are currently in the deploy stage, select **Deploy**.
5. Select an AI asset to launch the life-cycle review process or to view and modify additional details about the asset.
   - To initiate the life-cycle review process for the asset, select **Start review** and then follow the proceeding steps in Complete AI asset lifecycle.
   - To continue a life-cycle review that is already in progress, select the **Lifecycle** tab and then follow the proceeding steps in Complete AI asset lifecycle.
   - To view or modify general details about the asset, select the **Details** tab and then navigate to **Details** > **Details**.
   - To view or modify the use and purpose of the asset, select the **Details** tab and then navigate to **Details** > **Use & purpose**.

     Note: This option is available only for AI systems.

   - To view, modify, or add associations to other AI assets that are related to this asset, select the **Details** tab and then locate the **Related assets** section of the navigation menu. Select the subsection for the type of related AI asset that you want to view, modify, or add associations for.

     For example, if you want to view all datasets that can help train an AI model, select the **Training datasets** subsection.

     Note:

     If you want to associate the asset with additional related AI assets, select one of the following options from the Add from inventory drop-down menu:

     - To associate the asset with a related AI asset that is already available in your asset inventory, select **Add from inventory**.
     - To associate the asset with a related AI asset that is not already available in your asset inventory, select **Create**. This option is not available for related AI tools.

     You can then follow the proceeding steps in the corresponding AI asset creation topic:

     - Create AI system assets
     - Create AI model assets
     - Create prompt assets
     - Create dataset assets

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

- To view evaluation metrics for the asset, select the **KPIs & metrics** tab.

  Note: This tab is available only for AI systems and AI models.

- To view, modify, or create value templates that are mapped to the asset, select the **Value template** tab.

  Note: This tab is available only for AI systems.

  For more information on value templates, see Using value templates.

- To view associated onboarding, change, and off-boarding requests for the asset, select the **Requests** tab.

  For more information on change and off-boarding requests, see Creating requests for AI assets.

- To view the evaluations that are enabled for the asset, select the **Evaluation template** tab.

  Note: This tab is available only for AI systems.

Complete AI asset lifecycle

Complete the AI asset lifecycle process starting from assessment through deployment.

Before you begin

Role required: AI steward [sn_ai_governance.ai_steward] or Asset owner [sn_ai_asset_mgtm.ai_asset_ownerAI]

Procedure

1. Navigate to **Workspaces** > **AI Control Tower**.
2. From the AI Control Tower, open the AI assets view.
3. From the AI assets view, locate and open an asset which has its **Lifecycle status** as steward review.
4. In the asset record view, initiate the lifecycle process by selecting **Start review**.
5. On the **Lifecycle** tab, select **Assess** to create tasks.

   Note: The AI steward can manually create tasks under **Lifecycle**>**Assess** and assign them to an asset owner. The asset owner or the AI steward can mark this task as complete and move on to the next stage.

6. Select **Build and test** and create tasks under the **Development plan** and **Pre-deployment assessments** phases.

   Note: An AI steward can create tasks in the **Development plan** and **Pre-deployment assessments** phases, marking the tasks as complete in both the phases using the **Complete** button.

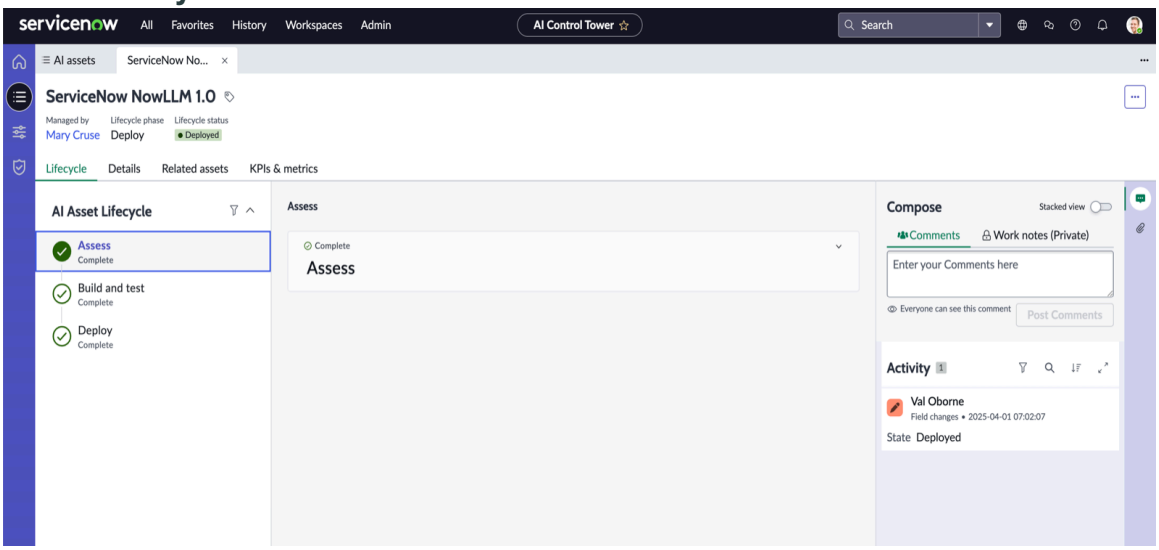
7. To create tasks and complete the AI asset Lifecycle, select **Deploy**.
   The asset owner can mark the task as complete.

Result

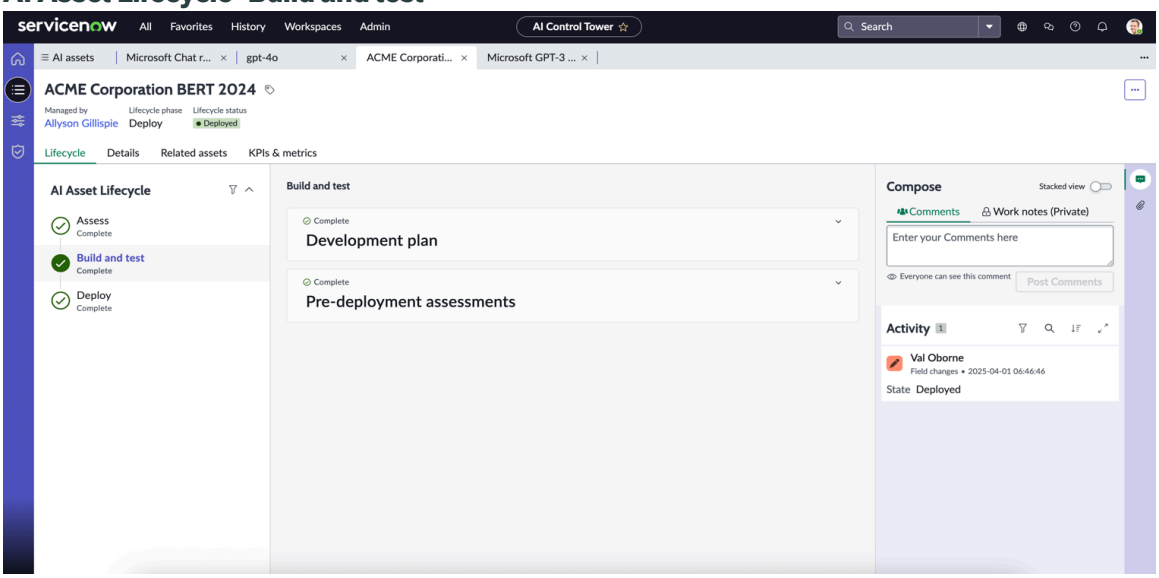
The life cycle status of an AI asset is deployed and completed.

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

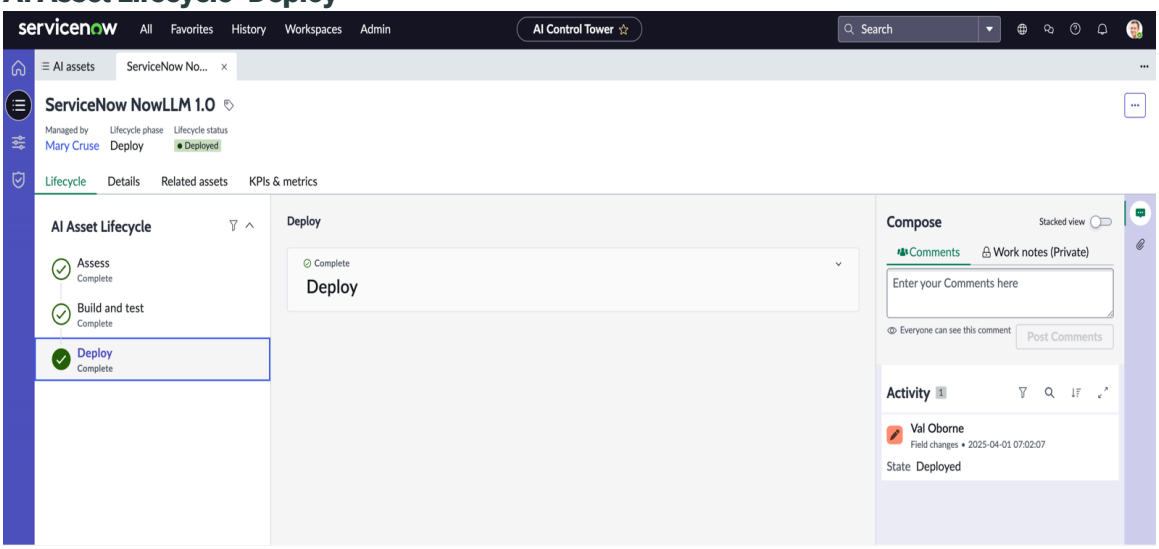
AI asset Lifecycle- Assess



AI Asset Lifecycle- Build and test



AI Asset Lifecycle- Deploy



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

To manage the approval workflow for ServiceNow® AI, see Create an AI Control Tower Playbook workflow.

Using value templates

Use value templates to define, calculate, and track the value delivered by your AI systems and models.

Value templates introduces a unified and structured framework for defining, calculating, and tracking the value delivered by your AI assets. You can define value based on personas and formulas using customizable templates.

Value templates offer the following benefits.

- Transparency: Provides visibility into how value is determined for an AI asset.
- Customization: Supports various persona-driven and formula-based definitions.

Currently, value templates are supported for skills and agents. A value template is assigned to an asset when a skill or agent is activated and deployed as an asset in the system. The following table lists the default value templates.

Default value templates

AI asset	Default value template
ServiceNow skill	Daily skill executions (Usage) x average time per execution (Time) x 50% (Acceptance rate) For example, for Incident Summarization skill, if the daily executions are 1500, average time per execution is 0.2 minutes, acceptance rate is 90%, the value as per the default value template is 1500x0.2x90/100, which is 270 minutes.
ServiceNow AI agent	Daily AI agent executions (Usage) x average time per execution (Time) x 50% (Acceptance rate) For example, for Problem investigator agent, if the daily executions are 1000, average time per execution is 1.5 minutes, acceptance rate is 50%, the value as per the default value template is 1000x1.5x50/100, which is 750 minutes.
AWS and Microsoft AI agents	Daily AI agent invocations (Usage) x 15 mins (Time) x 50% (Acceptance rate)

Add a value template

Add a value template to define, calculate, and track the value delivered by an AI asset.

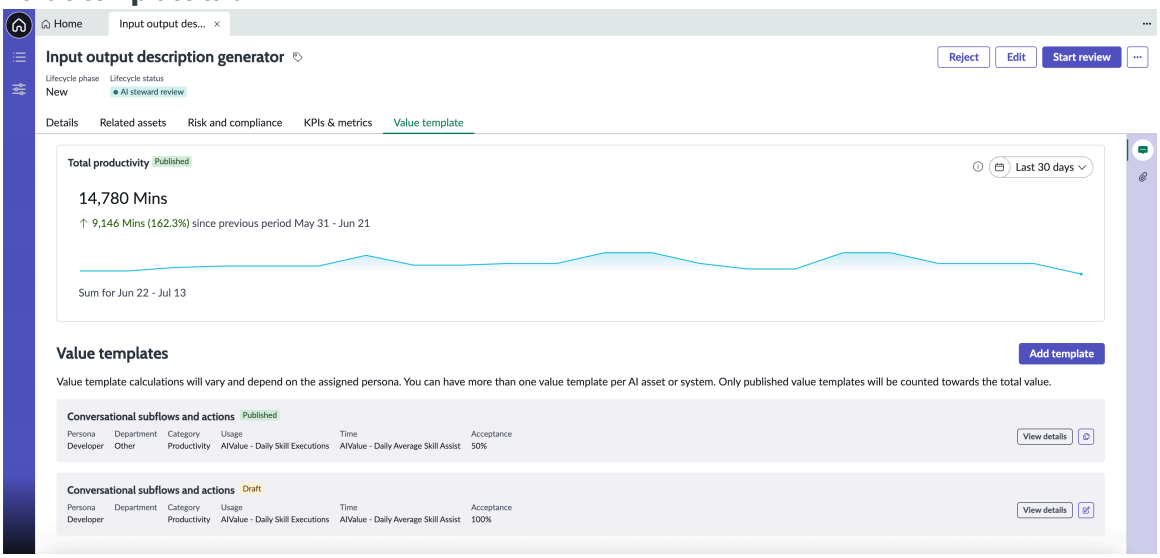
Before you begin
Role required: [sn_ai_governance_ai_steward]

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

Procedure

1. Navigate to **Workspaces** > **AI Control Tower**.
2. Select the **AI asset inventory** tab and navigate to the asset that you want to assign a value template to.
3. In the **Value template** tab, select **Add template**.

Value template tab



4. On the Add value template form, fill in the fields.

   For a description of the field values, see Add value template form.

5. Select **Save as draft**.
6. Validate and test the formula against the available dataset by selecting **Validate and calculate output**.
   You can change the values and validate the formula until you're satisfied with the output.
7. If you're satisfied with the output, publish the template by selecting **Publish**.
   When a value template is published, it's modified across all instances and subsequent calculations are based on the new template.

Edit a value template

Edit a value template to update how value is calculated for AI assets.

Before you begin

Role required: [sn_ai_governance_ai_steward]

Procedure

1. Navigate to **Workspaces** > **AI Control Tower**.
2. Select the **AI asset inventory** tab and navigate to the asset that you want to edit the value template for.
3. In the **Value template** tab, select **Duplicate**.
   A new tab with the same details as the original template opens for you to edit.
4. Edit the value template according to your requirements.
5. Select **Save as draft**.

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

6. Validate and test the formula against the available dataset by selecting **Validate and calculate output**.
   You can preview calculations from the past 30 days using the new template before publishing. This functionality enables you to evaluate the impact of the new template while preserving the historical data associated with the previous version.
7. If you're satisfied with the output, publish the template by selecting **Publish**.
   When a value template is published, it's modified across all instances and subsequent calculations are based on the new template.

Create an AI case in the AI Control Tower

Create an AI case in the AI Control Tower by providing a detailed description, such as system behavior, affected users, and any relevant data. Ensure all necessary information, such as supporting attachments, is included for prompt resolution.

Before you begin

Role required: sn_ai_case_mgmt.ai_case_analyst or sn_ai_case_mgmt.ai_case_manager

Procedure

1. Navigate to **All** > **AI Control Tower**.
2. Go to the **AI cases** dashboard.
3. On the AI cases dashboard, on the **AI cases** tab, select **Create AI case**.
4. On the form, fill in the fields.
   For a description of the field values on the Create New AI case form, see Create new AI case form.
5. Select **Save**.
   After you submit this form, the AI case team will provide the insights and support you need.

Create new AI case form

Use the Create New AI case form in the AI Control Tower to report an AI case with the necessary details.

See the following table for a description of the field values.

Create new AI case form

Field	Description
Number	Number of the request. This field is automatically set to a request number.
Name	Name of the AI case. For example, `Incorrect AI recommendation`.
Description	Description about the AI case in detail. For example, include information such as the expected recommendation, actual suggestion, context, impact, and steps to reproduce the issue, especially for cases such as Incorrect AI Recommendation.
Type	Type of the AI case. This field is automatically set to **AI case**.

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

Create new AI case form (continued)

Field	Description
Sub-type	Subtype of the case. For example, Adversarial attacks (Deliberate manipulation of AI models to produce incorrect results).
State	Workflow state of the request. This field is automatically set to **New**.
Priority	Urgency of the AI case. It enables the AI case team to triage requests effectively and responds based on how critical the request is. The options are as follows: • **Critical** • **High** • **Moderate** • **Low** • **Planning**
Requester	Name of the person who raised the request.
Requested on behalf of	Name of the person that you created the request for.
Primary entity	Entity impacted by the AI case. Only the entities identified in the impacted areas are available for selection as the primary entity.
Entity owner	User who is the owner of the entity. This field is automatically set based on the entity selected in the Impacted areas related list.
Assignment	
Assignment group	Group that is assigned to the request. **Note:** The assignment group is preconfigured to the request type during the configuration setup.
Case analyst	Analyst who will analyze and work on the AI case. The AI case analyst is a part of the Assignment group.
Watch list	Person who must be informed about the AI case.
Accountable executive	Person who is accountable for the AI case.
Primary origin	
Location	Location where the AI case occurred. For example, Japan.
Sub-location	Sub location of the AI case occurrence. For example, the sub location is Tokyo.

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

Create new AI case form (continued)

Field	Description
Impacted business unit	Business unit that is affected by the reported AI case. For example, Finance.
Impacted department	Department that is affected by the AI case. For example, Customer support.
Source	Source of the AI case creation. This field is automatically set to Manual when the case is manually created. If the case is reported from the Employee Center, the field displays the source as Employee Center.
Additional source	Mode of how the AI case is reported when the case is created manually. This field is available only when **Manual** is selected from the **Source** field and the record is saved. The choices are as follows: • Email • Phone
Schedule	
Date of occurrence	Date when the AI case occurred. For example, the case may have occurred on 18-02-2025.
Date of discovery	Date when the case was discovered. For example, the case may have occurred on 18-02-2025, but was discovered on 12-03-2025.
Reported date	Date the case is reported.
Case closure SLA	Expected date of case closure.
Investigation planned start	Planned start date to investigate the case.
Investigation actual start	Actual start date of case investigation.
Remediation planned start	Planned start date to remediate the case.
Remediation actual start	Actual start date of case remediation.
Investigation planned end	Planned end date to investigate the case.
Investigation actual end	Actual end date of case investigation.
Remediation planned end	Planned end date of case remediation.
Remediation actual end	Actual end date of case remediation.
Breach analysis	
Breach status	Status to indicate if a breach has occurred or not. The choices are as follows: • **To be determined**: This value is the default value. • **Breach detected**: If the breach is confirmed.

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

Create new AI case form (continued)

Field	Description
	• **Future potential**: If the breach hasn't occurred but may occur in the future. • **Not a breach**: When there's no breach detected.
Breach start	Date the breach started. This field only appears when **Breach detected** or **Future potential** is selected from the **Breach status** field.
Reporting status	Status of the AI case being reported to the regulators. This field is automatically set based on the reporting status of the regulations associated with the case. If an AI case has many regulations and even if one regulation is identified as reportable, then the reporting status of the case is set to **Reportable**. The default value of this field is **To be determined**.
Breach end	Date the breach ended. This field only appears when **Breach detected** or **Future potential** is selected from the **Breach status** field.
Breach significance identified	Date when the user identifies that the breach is significant. This field only appears when **Breach detected** or **Future potential** is selected from the **Breach status** field.
Root cause analysis	
Primary cause	Primary cause of the AI case occurrence. This field is automatically set to the primary cause that is selected in the Cause and consequences related list.
Primary consequence	Consequences of the primary cause for the AI case.
Overall observations	Observations made regarding the AI case.
Remediation taken	Field to indicate if remediation measures have been taken to address the AI case. The choices are as follows: • **Yes** • **No**
Overall preventive measures	Preventive measures taken toward the case.
Activity	
Work notes (Private)	Notes or information about the request.
Additional comments (Customer visible)	Additional information about the request that you want to share with your customers.

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

Create an AI issue in the AI Control Tower

Identify and manage issues related to the impacted areas for the reported AI case in the AI Control Tower.

Before you begin

Role required: sn_grc_ai_gov.ai_risk_and_compliance_analyst or sn_grc_ai_gov.ai_risk_and_compliance_manager

Procedure

1. Navigate to **All** > **AI Control Tower**.
2. Go to the **AI cases** dashboard.
3. On the AI cases dashboard, on the **AI cases** tab, select **Create issue**.
4. On the form, fill in the fields.
   For a description of the field values on the Create New AI case, see Create new AI issue form.
5. Select **Save**.

Create new AI issue form

Use the Create New Issue form to identify and manage issues related to the impacted areas for the reported AI case.

See the following table for a description of the field values.

Create New Issue form

Field	Description
Number	Number of the issue. This field is automatically set to a request number.
Name	Name of the issue. For example, `Cyber Attack on Acme`.
Issue source	Source from where the issue was created. This field is automatically set to **AI case**.
Issue type	Type of the issue. The options are as follows: • **Control design effectiveness failure**: The control was poorly designed and cannot effectively prevent or detect the intended risk. • **Control operative effectiveness failure**: The control was well-designed but failed during execution or wasn't followed correctly. • **Control doesn’t meet requirement**: The control is in place but doesn't satisfy regulatory, policy, or business requirements. • **Control doesn’t exist**: There is no control present to address a known risk or requirement. • **Non-compliance to a regulation**: A law or regulation was not followed, potentially exposing the organization to penalties. • **Non-compliance to a policy**: An internal policy was not adhered to, which could lead to risks or inefficiencies. • **Improvement or suggestion to an existing policy**: A recommendation to enhance an existing policy for better clarity, coverage, or effectiveness. • **Recommendation for a new policy**: A proposal to create a new policy to address a gap or need that currently isn’t covered.

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

Create New Issue form (continued)

Field	Description
	• **Process optimization or improvement**: Opportunities identified to improve efficiency, accuracy, or effectiveness of a business process. • **Observation**: A general note or finding that may not be an issue now but could warrant attention. • **Data breach**: Unauthorized access, disclosure, or loss of sensitive or personal data. • **Fraud**: Intentional deception for personal or organizational gain, such as misappropriation of assets. • **Misstatement**: Errors or omissions in financial or operational reporting that misrepresent facts. • **Training**: Gaps or needs identified in knowledge or skills that require attention. • **Documentation**: Issues related to missing, outdated, or inaccurate documentation. • **Risk issue**: A broad risk-related concern that may not fall under other specific categories. • **Other**: Any issue that doesn't fit into the above types but is still worth tracking and resolving.
Classification	Classification of the issue. The options are as follows: • **Compliance** • **Risk** • **Audit** • **Vendor Risk**
Location	Location where the issue occurred. For example, Japan.
State	Workflow state of the issue. This field is automatically set to **Review**. The options are as follows: • **New** • **Analyze** • **Respond** • **Review** • **Closed Complete** • **Closed Incomplete** For more information on Issue management workflow and life cycle, see View AI assets by life-cycle stage.
Substate	Substate of the issue. This field is auto-filled.
Priority	Urgency of the issue. It enables the team to triage requests effectively and respond based on how critical the request is. This field is automatically set to **Review**. The options are as follows:

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

Create New Issue form (continued)

Field	Description
	• **Critical** • **High** • **Moderate** • **Low** • **Planning**
Issue rating	Rating of the issue. The options are as follows: • **Very High** • **High** • **Moderate** • **Low** • **Very Low**
Description	Description about the issue in detail. For example, ACME experienced a cyber attack that resulted in unauthorized access to internal systems.
Assignment	
Assignment group	Group to whom the issue is assigned. For example, Risk Managers.
Assigned to	User to whom the issue is assigned.
Issue manager group	Manager group that the issue is assigned to. The options are as follows: • **Compliance Managers** • **IT Risk Managers** • **Risk Managers**
Issue manager	Manager to whom the issue is assigned.
Watch list	Person who must be informed about the issue.
Schedule	
Due date	Date when the issue is due.
Confirmed date	Confirmation date for the issue. This field is auto-filled.
Created	Date on which the issue is created. This field is automatically set to the current date and time.
Closed	Date on which the issue is closed.
Planned start date	Planned start date for the issue.
Planned end date	Planned end date for the issue.
Duration	Duration for the issue in days, hours, minutes, and seconds.
Actual start date	Actual start date for the issue.

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

Create New Issue form (continued)

Field	Description
Actual end date	Actual end date for the issue.
Actual duration	Actual duration for the issue in days, hours, minutes, and seconds.
Issue grouping	
Issue group rule	Group rule for the issue. This field is auto-filled.
Parent issue	Parent issue that is associated with the issue.
Action plan	
Recommendation	Recommendation for the issue. For example, `Enhance cybersecurity measures and provide regular security training`.
Action plan	Action plan for the issue. For example, `Upgrade the firewall to the latest version and implement multi-factor authentication for all administrative accounts.`
Activity	
Work notes (Private)	Notes or information about the issue.
Additional comments (Customer visible)	Additional information about the issue that you want to share with your customers.
Settings	
Functional domain	Functional domain that the issue belongs to. For example, an issue may belong to the AI Risk and Compliance domain.

Using the access map

Learn how to use the **Access map** in AI Control Tower.

Before you begin

Role required: AI Steward

Procedure

1. Select **access map** in the header of the **Security and privacy** tab.
   You can also access the map by navigating to **All** > **AI Security and Privacy** > **Access Map**.
2. Select an AI agent/agentic workflow using the **Access map**.
   You can also use the **Search by agent or workflow** field, and the **AI agents** and **Agentic workflows** to refine your search. Additionally, agents with access issues will have a ⚠ icon.
3. In the node map, select a node to view the connections of your AI agent/agentic workflow.
   Agents with access issues with be highlighted with a ⚠ icon in the map.
4. Select a node to get further details and configure settings.

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

Access map node details

Icon	Label	Details
	Agentic workflow	**Description** A brief description of the workflow **Domain** The domain scope of the workflow **Execution mode** How the workflow is executed (Autonomous or manual) **Created by** The creator of the workflow
	Agent	**Description** A brief description of the agent **Domain** The domain scope of the agent **Created by** The creator of the workflow **Tables accessed** Review the table accesses of the agent over the last month. Select **Operation** to filter specific access operations. **Note:** Only agents with access issues will display the **Access issues** details. **Access issues**

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

Icon	Label	Details
		Details the resources and operations causing the access issues. **Resource** The resource the agent is having issues accessing **Operation** The operation the agent attempted **Count** The number of times of the agent encounter the issue
	Tool	**Description** The description of the tool. **Domain** The domain scope of the tool. **Execution mode** The execution mode of the tool (Supervised or Autonomous) **Type** The type of tool **Target record** The record the tool modifies. **Created by** The creator of the tool **Note:** Select **Open tool record** to view additional tool details.

Create an Automation rule

Create an Automation rule to set AI assets as managed.

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

Before you begin
Role required: AI steward

Procedure
1. Navigate to **AI Control Tower** > **Configuration** > **Controls** > **Automation rules**.
2. Select **+Create new**.
3. Enter the rule according to your requirements in the Rule name field.
4. Enter the description of the rule in the Description field.
5. Select the Asset and define the conditions.
6. Select the Activate the rule check box.

Result
An automation rule is created for AI assets.

AI Control Tower reference

Reference topics for AI Control Tower.

Add value template form

Add value template form fields and descriptions.

Add value template form

Field	Description
Calculation builder	
Persona	The user role for whom the value is being calculated, for example, Agent, Requester.
Usage	The usage indicator that you want to associate the value template with.
Time value type	The type of time-based metric used in the calculation, for example, Constant or Indicator. Choose a value that aligns with your performance goals.
Time constant	A fixed time value (in hours, minutes, or days) used as a baseline or threshold in the calculation. Useful for comparisons or normalization.
Time indicator	The time indicator that you want to associate the value template with.
Acceptance rate type	The type of acceptance metric to include, such as Indicator or Constant.
Acceptance rate constant	A fixed percentage or ratio used to represent an expected or target acceptance rate in the value calculation.
Attributes	
Template name	A unique and descriptive name for the value template. This name appears in the list of

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

Add value template form (continued)

Field	Description
	templates and should clearly indicate its purpose.
Value template category	The categorization of the template based on its use case, for example, Productivity, Efficiency, Accuracy. Helps in organizing and filtering templates.
Department	The business unit or department associated with the value template. This helps in aligning the template with organizational goals.
Description	A brief explanation of what the value template does, including its purpose and how it contributes to measuring value.
Calculation output	
Instance	The instance where the value template is applied. This field is useful for managing value templates across instances. **Note:** You can validate a formula in a non-production instance before publishing it to a production instance.

AI Control Tower roles

Certain roles are installed along with the installation of the AI Control Tower.

Roles and their descriptions

Role title [name]	Description	Contains roles
AI steward [sn_ai_governance_ai_	**Note:** The organization decides on assigning the AI steward role. By adding the users to the AI stewards group, allows user to have additional permissions related to playbook. The AI steward is responsible for: • Configuring AI Control Tower • Adoption of AI governance practices. • Adoption of managing AI Control Tower and linking the AI asset Inventory. • Execution of AI Control Tower initiatives. • Understand the AI assets and AI Control Tower policies.	• sn_nowassist_admin.user • sn_ai_governance.workspace_admin • sn_aia.admin • aig_admin • sn_mcp_client.admin

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

Roles and their descriptions (continued)

Role title [name]	Description	Contains roles
	• Creating AI assets. • Completing the AI asset lifecycle. • Collaboration of cross-functional teams within the organization to confirm that the organization policies are adhered. • Creating AI Control Tower Approval Playbook for Now Assist approvals. • Configure third-party LLMs and SLMs. • Configure Multi-instance management. • Add and edit a value template. • Learning to use the access map • Approve or reject an approval request. For AI discovery: • Activate or deactivate hyperscaler connections. • Select the hyperscaler connections to discover agents and usage on-demand. For AI Gateway: • Add an MCP server via AI Agent Studio. • Set up MCP client connections.	
AI Control Tower Workspace user [sn_ai_governance_workspace_user]	The AI Control Tower Workspace user is responsible for: • Own and manage the AI assets. • Access the AI Control Tower home page. • Exclusive access to the AI portfolio tab.	None
AI asset owner/ Product owner [sn_ai_asset_mgmt.ai_asset_owner]	The AI asset owner/Product owner is responsible for: • Ensure that AI assets are represented accurately and kept up to date. • Manage AI assets like AI systems, AI models, datasets, and prompts through their asset lifecycle from intake to retirement.	None

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

Roles and their descriptions (continued)

Role title [name]	Description	Contains roles
	• Access My overview, Value, and Adoption tabs. • Creating an AI asset from the AI Control Tower home page using **Create AI Asset icon**. • Marking the deploy phase of the AI asset lifecycle task complete. If the AI asset gets deployed, then the state of the task doesn't change anything automatically in the asset table or the asset governance details record.	

AI Risk and Compliance roles

The AI Risk and Compliance application installs the essential role to perform respective day-to-day operational tasks for managing AI systems across the enterprise.

Roles and their descriptions

Role title [name]	Description	Contains roles
AI Risk and Compliance Admin [sn_grc_ai_gov.ai_risk_and_compliance_admin]	The AI Risk and Compliance Admin can perform the following tasks: • Set up risk and impact assessment frameworks. Configure risk assessment methodologies, risk contribution factors, and impact assessment templates. • Define automation rules for impact assessments to determine applicable risks and controls based on the assessment responses. • Set up and profile AI case types. • Delete AI systems. • Enable or disable Entity-Based Access for record types associated with entity properties, and configure	• sn_smart_asmt.template_manager • sn_grc_ai_gov.ai_risk_and_compliance_mana • sn_smart_asmt.assessment_admin • sn_grc_workspace.state_model_admin • sn_smart_asmt.template_contributor • sn_ai_case_mgmt.ai_case_admin • sn_reg_body_mgmt.writer • sn_risk_advanced.ara_admin • sn_rec_pg_vertical.admin • sn_grc_ent_access.admin **Note:** GRC: Entity Based Access application must be installed for this role to be available.

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

Roles and their descriptions (continued)

Role title [name]	Description	Contains roles
	the Entity-Based Access settings as needed. **Note:** GRC: Entity Based Access application must be installed to use this feature.	
AI Risk and Compliance Manager [sn_grc_ai_gov.ai_risk_and_compliance_manager]	The AI Risk and Compliance Manager can access all AI systems on the system and perform the following tasks: • Initiate impact assessments. • Manage the life cycle of an AI system. • Initiate risk assessments. • Initiate control attestations. • Write and update access to the bulk access update configuration. **Note:** GRC: Entity Based Access application must be installed to use this feature.	• sn_grc_ai_gov.ai_risk_and_compliance_analy • sn_smart_asmt.template_contributor • sn_smart_asmt.template_manager • sn_risk_advanced.risk_asmt_project_manag • sn_ai_case_mgmt.ai_case_manager • sn_grc_ent_access.bulk_access_config_adm **Note:** GRC: Entity Based Access application must be installed for this role to be available.
AI Risk and Compliance Analyst [sn_grc_ai_gov.ai_risk_and_compliance_analyst]	The AI Risk and Compliance Analyst can access all AI systems assigned to them in the system and perform the following tasks only on the assigned records: • Initiate impact assessments. • Manage the life cycle of an AI system. • Initiate risk assessments. • Initiate control attestations.	• sn_ai_case_mgmt.ai_case_analyst • sn_smart_asmt.assessment_reader • sn_smart_asmt.template_reader • sn_grc_ai_gov.ai_risk_and_compliance_busi • sn_grc_ai_gov.ai_risk_and_compliance_read • sn_grc_workspace.user • sn_grc_workspace.state_model_reader • sn_risk_advanced.ara_creator • sn_risk_advanced.ara_assessor • sn_risk_advanced.ara_approver • sn_risk_advanced.risk_asmt_project_user
AI Risk and Compliance Business User	The AI Risk and Compliance User can perform the following tasks:	• sn_grc_workspace.assessment_template_co • sn_smart_asmt.actor

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

Roles and their descriptions (continued)

Role title [name]	Description	Contains roles
[sn_grc_ai_gov.ai_risk_and_compliance_business_user]	• Create AI case on the Employee Center. • Work on the assigned tasks. • Perform control attestations.	• sn_grc_workspace.user • sn_smart_asmt.assessment_reader • sn_risk_advanced.risk_asmt_project_reader **Note:** For more information on AI Control Tower roles, see AI Control Tower roles.
AI Risk and Compliance Reader [sn_grc_ai_gov.ai_risk_and_compliance_reader]	The AI Risk and Compliance Reader can have read access to the AI systems and AI impact assessments.	• sn_grc_workspace.user • sn_grc_workspace.state_model_reader
AI System Reader [sn_grc_ai_gov.ai_risk_and_compliance_ai_system_reader]	The AI System Reader can have read access to the AI systems on AI Control Tower workspace and AI Risk and Compliance workspace.	NA
AI Case Business User [sn_ai_case_mgmt.ai_case_business_user]	The AI Case Business User can create AI case and AI inquiry on the Employee Center.	sn_grc_case_mgmt.grc_case_business_user
AI Case Analyst [sn_ai_case_mgmt.ai_case_analyst]	The AI Case Analyst can review the AI cases and AI inquiries assigned to them in the system and perform the following tasks only on the assigned records: • Identify and manage impacted and related areas such as policies, regulations, and enterprise-wide compliance risks. • Identify and manage issues related to impacted areas to eliminate the root causes.	• sn_grc_case_mgmt.grc_case_analyst • sn_ai_case_mgmt.ai_case_business_user
AI Case Manager [sn_ai_case_mgmt.ai_case_manager]	The AI Case Manager can review all the AI cases, AI inquiries, and its associated information.	• sn_ai_case_mgmt.ai_case_analyst • sn_grc_case_mgmt.grc_case_manager
AI Case Admin [sn_ai_case_mgmt.ai_case_admin]	The AI Case Admin can manage type profiles to segregate AI cases. They can set up assignment rules and delete AI cases.	• sn_grc_case_mgmt.grc_case_admin • sn_ai_case_mgmt.ai_case_manager

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

Tables installed with AI Control Tower

Certain Tables or fields are installed along with the activation of the AI Control Tower core plugin.

The following table is provided to understand the AI assessment and AI task fields.

AI Control Tower details

Asset	ai_asset	Reference
Risk classification	risk_score	choice
Status	approval_status	choice

AI asset assessment request

Sys ID	sys_id	SYS_ID
AI asset	asset	Reference
Parent Task	parent_task	Reference (ai_gov_review_task)
State	state	
Type	type	Choice
Close Notes	close_note	

AI asset governance task

Sys ID	sys_id	SYS_ID
Parent Task	parent_task	Reference
Approval	approval	Reference (sysapproval)
State	state	
Type	type	Choice

AI asset data model attributes

Additional attributes for the AI asset data model.

Attributes

AI model product model: Product Information for the AI model that is used by the AI system to generate responses without human intervention (cmdb_ai_model_product_model).

AI model product model

Attribute	Description
Model parameters info	Number of parameters for the model.
Supported languages	Languages supported.

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

AI model product model (continued)

Attribute	Description
Model size	Size of the model in MB. Mostly applicable for models developed and deployed within the organization.
Deployment guidelines	Instructions applicable for models developed and deployed within the organization.
Source	Links or details of source of the model sources example: Hugging face, Microsoft, and so on.
Training procedure	Types of training • Decision Trees • Deep Neural Networks • Linear Regression • Logistic Regression • Random Forest • Supervised Learning • Unsupervised Learning • Reinforcement Learning • Transfer Learning • Semi-Supervised Learning • Instruction Finetuning • Supervised Finetuning
Context window	Size of input sequences that the model can handle (number of tokens).

AI dataset product model: Product Information for the collection of data that is used to train and test AI models (cmdb_ai_dataset_product_model).

AI dataset product model

Attribute	Description
Data type	Describes data, example: Text, Image, Video, and Table
Source	Links or details of source of the dataset sources, example: Customer, Wikipedia, Hugging face, Crowd sourced, and so on.
Acceptable usage	Acceptable usage of the data according to license / contract example: Training, and Evaluation.

AI prompt product model: Product Information for instructions given to AI models to get a response for AI system (cmdb_ai_dataset_product_model).

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

AI prompt product model

Attribute	Description
Documentation	Links and information about requirements, design, and related information.

AI system product model: Product Information for software that provides ML / AI capability to generate outputs, such as decisions, recommendations, content, or predictions (cmdb_ai_system_product_model).

AI system product model

Attribute	Description
Documentation	Links and information about requirements, design, and related information.

Information asset

Attribute	Description
Data classification	Classification according to organization's data classification model, example: Public, confidential, and customer confidential.

AI digital asset

Attribute	Description
ServiceNow® record reference	Reference to Now Assist record.
ServiceNow® table	Now Assist table.

AI system digital asset

Attribute	Description
AI models	Reference to more than one associated models.
Evaluation Dataset	Reference to more than one associated datasets used for evaluation.
Evaluation Metrics Report	Details of evaluation results.

AI model digital asset

Attribute	Description
Base model	This AI model version was derived from an internal model developed within the organization.
Model weights info	Additional model information if available. Mostly applicable for models developed within the organization.
Required infrastructure	Documentation of infrastructure requirements for model deployment, primarily for models deployed within

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

AI model digital asset (continued)

Attribute	Description
	an organization. Example: Infrastructure stack and processing requirements.
Training dataset	Reference to one or more associated datasets used for training the model. These datasets are mostly applicable for models developed within the organization.
Evaluation dataset	Reference to one or more associated datasets used for evaluating the model. These datatsets are mostly applicable for models developed within the organization.
Evaluation metrics report	Links or details of evaluation results
License details	Link or detail to applicable licenses applied to the model.
Model card	Links to shareable model card (Internal and external model card).

AI dataset digital asset

Attribute	Description
Base datasets	This version of the AI dataset was derived from the previous version.
Dataset card	Information on number of records, distribution, and so on. Documentation for data quality and known risks and limitations.
License details	Link or detail to applicable licenses applied to the dataset, example: CommonCore, Apache 2.0,etc.

AI prompt digital asset

Attribute	Description
Prompt information	Details of the prompt.
AI model	Reference to the AI model for which the prompt is created.

Evaluation dashboard reference

Reference topics for the Evaluation dashboard.

Components installed with the Evaluation dashboard

Several types of components are part of the Evaluation tab, including scheduled jobs, tables, system properties, and flows.

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

Scheduled jobs installed

Scheduled job	Description
CE Populate Value Aggregates Chats – Daily	This scheduled script runs daily and randomly selects 1000 conversations from yesterday's conversations. After that, for each conversation, this job extracts the chat duration and classifies them as small, medium, or large. It also classifies the chats in which a Knowledge article or catalog item was invoked. For the evaluated chat, it also classifies the conversations based on the chat performance and populates that data into the Evaluation Value Aggregates table.
Evaluation Value Calcuation - Runs Only once after install	Deletes all the records on the Evaluation Value Aggregates tables, runs the calculations again and stores the aggregated value in the Evaluation Value Aggregates table. The data is from the first evaluation date.

Tables installed

Label	Name
Evaluation	[sn_na_conv_eval_evaluation]
Evaluation configurations	[sn_na_conv_eval_evaluation_configurations]
Evaluation Metrics	[sn_na_conv_eval_evaluation_metrics]
Evaluation Set	[sn_na_conv_eval_evaluation_set]
Evaluation Value Aggregates	[sn_na_conv_eval_evaluation_value_aggregates]

Remote tables installed

Table	Description
Conversation Evaluator Value Calculations [sn_na_conv_eval_st_value_calcs]	For the given query, the definition for this remote table calculates the time savings and efficiency percentage for small, medium, and large chats. Also, it returns the time savings and efficiency when a Knowledge article or catalog item was invoked.
Conversation weekly calculations [sn_na_conv_eval_weekly_cals]	For the given query, the definition for this remote table calculates the time savings and efficiency percentage for small, medium, and large chats for different weeks of the selected date range. Also, it returns the time savings and efficiency when aKnowledge article or catalog item was invoked for all the different weeks of the selected date range.

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

System properties installed

Property	Description
sn_na_conv_eval.errorBandMinRecords	Minimum number of records required to calculate the error band for upper and lower deviation. By default, the value is 30.
sn_na_conv_eval.evalWeights	Contains weights to each evaluation metric for chat evaluation. This property is used to compute total or composite scores for evaluation records.
sn_na_conv_eval.maxEvaluateCount	Maximum number of records to evaluate in a day. By default, the value is 200.
sn_na_conv_eval.total_sampled_conversations	Edit this property to control the total number of conversations that can be sampled for value calculations. By default, the value is 1000.
sn_na_conv_eval.value_chat_classifier	Edit this property to change the definition of small, medium, and large conversations. By default, the value it stores is 4, 10. Here, 4 and 10 signify the total number of inbound messages. Fewer than or equal to 4 inbound messages in the sys_cs_message table for a conversation means that it's a small conversation. More than 4 inbound messages and fewer than or equal to 10 inbound messages means that it's a medium conversation, and more than 10 inbound messages means that it's a large conversation.
sn_na_conv_eval.ce_value_calculation_weights	Value calculation weight values for each type of evaluated chat.
sn_na_conv_eval.eval_value_rerun_status	Re-state the value calculations once after the installation. This property will check the status of the Conversation Evaluator Value Rerun status. If it has run, then the script will change the value of this system property to false.

Business rules installed

Name	When	Insert	Update	Filter Conditions
Add info message for Evaluation set	after	TRUE	TRUE	stateCHANGESTOIn Progress^evaluation_type=conversation^
Scale Up labeling metric	before	TRUE	TRUE	metric_type=Labeling^metric_nameINhel
updateLabelingScoresOnEvaluation	after	TRUE	TRUE	metric_type=Labeling^raw_scoreVALCHA
Update deviation scores	before	TRUE	TRUE	metric_type=LLM Generated^scoreVALCHANGES^EQ

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

Name	When	Insert	Update	Filter Conditions
getAutoEvalCompleteScore	after	FALSE	TRUE	stateCHANGESTOComplete^total_scorel

Flows installed

Flow	Description
Execute Evaluation	Performs evaluations when conversations are completed. By default, the Execute Evaluation flow is deactivated. You can use the nightly scheduled job Execute Evaluations to evaluate the chats. If you want to evaluate the chats on chat completion, activate the Execute Evaluation flow.
Execute Batch Evaluation	Performs batch evaluations, evaluating up to 100 completed virtual agent conversations. Flow is triggered when the Evaluation set is created or updated and the Evaluation Type is Conversation.

Flow actions installed

Flow action	Description
Randomize conversations	Performs randomization of conversations and returns 100 conversations randomly from a given query.
invokeApiDefinition	Invokes OneExtend Capability in the large language model (LLM).
Chat Classifier Eval	Gives the title, category, and whether the evaluation should be executed.
buildTranscript	Builds the transcript from a conversation.
evalExecuteCondition	Checks if the transcript is good enough to be evaluated.

Script includes installed

Script includes	Description
evalExecuteCondition	Use this script include to update the evaluation condition.
evalUtils	Primary Utility function for the Evaluator.

Related topics

Remote tables

Evaluation metrics and calculations

Metrics against which conversations are evaluated and calculation of adjusted scores.

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

Metrics

The Select metric list shows all the metrics against which each conversation is evaluated for the selected date range. You can filter the evaluation trend based on each metric. The following metrics are available:

Metric	Description
Request Completion	Measures the virtual agent's ability to complete user requests by accurately identifying the user's intent and gathering all required information (slot filling).
Intent Accuracy	Shows the virtual agent's ability to comprehend user requests, resulting in relevant responses.
Slot Filling	Shows the virtual agent's capability to interpret user responses and extract structured answers to the required questions.
Smooth Flowing Conversation (Deadlock avoidance)	Checks if the virtual agent responds dynamically, successfully moving the conversation forward without repetition.
Context Retention	Shows if the virtual agent succeeds in retaining and using information provided during the conversation, including request interpretation and slot filling.
Truthfulness (Hallucination Prevention)	Shows if the virtual agent generated genuine responses grounded in conversation, excluding fabrication or memory and comprehension failures.
Conciseness (Redundancy Avoidance)	Checks the virtual agent's ability to avoid superfluous or verbose and generic responses, which doesn't contribute to the core intent of the conversation.
Coherence	Checks for clear logical flow, structure, and organization of the virtual agent's responses.
User Satisfaction	Weighted average of all the other metrics on which the conversation was evaluated.

Note: All the metrics are rated on a scale of 3 or 5, and are finally scaled up to 5.

Calculations

Calculation of deviations and Adjusted Score:

To align the auto-evaluation scores with human judgment over time, a deviation is calculated and used to produce an adjusted score on metric level.

- Upper Deviation

  Condition: If the number of human-labeled scores that are higher than the auto-evaluated scores in the last 6 months is more than 30.

  Calculation: The top 90% of these cases are taken and the difference (delta) between the human score and the auto-evaluated score is averaged. This delta is the Upper Deviation.

- Lower Deviation

  Condition: If the number of human-labeled scores that are lower than the auto-evaluated scores in the last 6 months is more than 30.

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

Calculation: The top 90% of these cases are taken and the difference (delta) between the human score and the auto-evaluated score is averaged. This delta is the Lower Deviation.

- Adjusted Score

  The final Adjusted Score is calculated based on the availability of the deviations.

  - If at least 30 distinct evaluations of both upper and lower deviations are labeled for a given metric, Error band is calculated as SUM(Avg labeling score – LLM score)/Distinct evaluations. This error band is added to Auto-Eval score to get Adjusted Score.
  - If neither deviation is available, then Adjusted Score = Auto-Eval Score

Calculation of Auto eval user satisfaction score, Human user satisfaction score, and Upper and Lower deviation on Evaluation level:

- Auto eval user satisfaction score: For a given evaluation, get all the scores for each metric that are LLM generated and calculate SUM(metric score * metric weight)/SUM(metric weights).
- Human user satisfaction score: For a given evaluation, if at least one metric is labeled, it's considered to calculate the human user satisfaction score. If labeled, the labeling score is used, or else LLM score is used. Calculated as SUM(metric score * metric weight)/SUM(metric weights).
- Gap: Gap is calculated as (Human user satisfaction score – Auto eval satisfaction score).
- Upper Deviation: If the Gap is positive and there are more than 30 records, the error band is calculated at top 90% by SUM(Positive Gap) / Distinct evaluations. This error band is added to the Auto eval user satisfaction score.
- Lower Deviation: If the Gap is negative and there are more than 30 records, the error band is calculated at top 90% by SUM(Negative Gap) / Distinct evaluations. This error band is added to the Auto eval user satisfaction score.
- Adjusted user satisfaction score is calculated as SUM(Gap) / Distinct evaluations.

Note:

- The evaluator provides aggregated score per chat, even if there are multiple different requests made by user.
- Performance Analytics indicators are used to calculate the average score over time. If you run batch jobs on historical data, then by the definition of Performance Analytics indicators, these evaluations are counted on the evaluation date in aggregated scores and not counted for scores on the actual chat date.

Evaluation flow

The workflow for evaluation execution, which performs evaluations when conversations are completed.

Conversations are evaluated using the following logic:

1. Conversation capture:

   All end-user interactions with the virtual agent are logged in the Conversation table [sys_cs_conversation]. When a user ends the conversation, the record's state is updated to Complete.

2. Automated flow evaluation trigger:

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

Flow name: Execute Evaluation.

Trigger condition:

- Table: Conversation table [sys_cs_conversation]
- State: Complete
- Device type: Web Client, Slack, Teams, Bot to Bot, Messenger

Sequence of execution:

Action 0: Check evaluations count for today

- Perform a query on evaluation table and to get record count.
- If record count is less than Max Number of evaluations per day, continue to Action 1, else end flow.

Action 1: evalExecuteCondition

- Invokes the `evalExecuteCondition.executeEvaluation` script Include with conversation reference.
- Generates a random number (1–100). Proceeds only if ≤10 (10% random sampling).
- Outcome: Returns true or false for further processing.

Action 2: Conditional Branch

- If true: Proceed to the next action.
- If false: Evaluation stops.

Action 3: Lookup Interaction Table:

Matches the conversation's channel metadata with the interaction table to fetch related records.

Action 4: Application Scope Filter:

If the interaction's application scope doesn't include `hr`, continue.

Action 5: buildTranscript:

Detailed Transcript Construction:

- Tags: [User]: For user messages, [Virtual Agent]: For virtual agent messages.
- For any referenced Knowledge article:
  - Pulls the complete article body to replace genius result, tagged with `[Virtual Agent]: Help articles for user query:` and delimited by Article_Start/Article_End.
  - If the Knowledge article is in HR scope/inaccessible, skip evaluation.
  - If the Knowledge article content is >10,000 words: Truncate at 10,000.
  - Attached files (PDF/Word/Txt): Use genius result instead.
- For referenced Catalog Items:

  Extracts name, short description, description, annotated as `[Virtual Agent]: Please choose one of the below options:` with citation number.

- If the first message is to the live agent, or the live agent is invoked within the first 120 words: skip evaluation.

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

Outputs:

- ExecuteEvaluation (true/false)
- Chat Transcript
- Knowledge articles or catalog items referred
- Sys_id of first live agent invocation (if any)
- List of skills to invoke (all evaluation skills for Evaluation dashboard)
- Additional evaluation logs

Action 6: Conditional Branch:

If ExecuteEvaluation is true: Continue to Action 7.

Action 7: Chat Classifier Eval

- Builds the initial transcript from *sys_cs_message*.
- Uses Chat topic classifier to determine:
  - Should the conversation be evaluated? (ExecuteEvaluation: true/false)
  - Topic Name
  - Category (IT/HR)
- If ExecuteEvaluation is true: Proceed to Action 6.

Action 8: Create or Update Evaluation Record:

Create a record on Evaluation [sn_na_conv_eval_evaluation] table with:

- Document Conversation: Conversation reference
- State: Processing
- Topic, Category, Knowledge article or catalog references, first live agent sys_id, type, user who initiated, message log

Action 9: For each skill:

Repeats for each skill flagged in Action 6.

Action 10: invokeApiDefinition

- Inputs: Skill name, conversation, transcript, evaluation id
- Calls Now Assist Skill API asynchronously.
- Post processing available in sys_generative_ai_response_validator, performs the following parsing:
  - Score
  - Reason for Score
  - Examples for the reasoning
- Parsed data is created on the Evaluation Metrics [sn_na_conv_eval_evaluation_metrics] table (Score, Reasons, Examples, and the entire reasoning for scoring [Scratchpad]).

Action 11: Waits 7 seconds before continuing to the next skill.

Special behavior and edge case handling:

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

- Sampling: Only 10% of conversations (randomly chosen) are evaluated.
- Channel Filter: Only Web, Slack, Teams, Bot to Bot, Messenger.
- Application Scope: Excludes records with `_hr_` in the scope.
- Knowledge article controls: No evaluation for HR or inaccessible. Knowledge articles, limits on Knowledge article size, and file handling.
- First live agent invocation: Excludes conversations routed to the live agent at the start or within 120 words.
- The Request Completion skill is added as part of a business rule where the score is tagged as the lowest between Slot filling and Intent.
- The reason on the record is added as follows:

```
if (Slot filling score > Intent score) {
Intent reason is used
} else if (Slot filling score < Intent score) {
Slot filling reason is used
} else {
Both are used
}
```

Evaluation flow for batch evaluations

Batch evaluation enables Eval admins to evaluate up to 100 completed virtual agent conversations at once, based on a saved query.

Flow name: Execute Batch Evaluation.

The flow creates evaluation records and invokes Now Assist skills for each eligible conversation, mirroring the single-conversation evaluation logic, but at scale. It enforces HR scope exclusions, topic/category validation, transcript construction rules, early live-agent exclusions, and asynchronous scoring through skills.

Batch evaluations are performed using the following logic:

Trigger

- Table: Evaluation set [sn_na_conv_eval_evaluation_set]
- Condition: State changes to In Progress and Evaluation type = Conversation

Inputs

- Evaluation Set record with:
  - Query filter: A query that targets conversations to be evaluated (for example, sys_cs_conversation filters).
  - Evaluation type: Conversation
  - State: In Progress (to start)
- LLM/Skills: Chat Topic Classifier, plus the evaluation skills listed after this.

High-level behavior

- Reads the query filter and randomly samples up to 100 conversations.
- Skips already-evaluated conversations.
- Excludes HR-scoped interactions.
- Uses Chat Topic Classifier to validate evaluation eligibility and extracts Topic and Category.

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

- Builds a transcript with controlled inclusion of Knowledge articles and catalog sources, and applies early live agent exclusions.
- Creates an Evaluation record and asynchronously invokes all selected evaluation skills, writing scores and rationale to metrics.

Sequence of execution:

Action 1: If the query filter isn't empty

- Purpose: Guard clause.
- Logic: Look up the Evaluation Set record and check the query filter field.
- If the query filter is present: Proceed to Action 2.
- If empty: Stop and optionally log `No query provided`.

Action 2: Randomize conversations

- Purpose: Select a bounded, random sample of conversations from the provided query.
- Logic:
  - Execute the query to get matching conversation records.
  - Randomly select up to 100 conversations.
    - If >100 matches, cap at 100.
    - If <100, select all.
  - Validate the query; if invalid, return false and an empty or partial array.
- Outputs:
  - success: true/false
  - conversation_ids: array of sys_ids (max 100)
- If success = true: Proceed to Action 3; otherwise, stop and log the validation error.

Action 3: Look up the evaluation table to check prior evaluation

- Purpose: Avoid duplicate evaluations.
- Logic: For each conversation sys_id, check sn_na_conv_eval_evaluation for existing records indicating that it's already evaluated or is in progress (implementation choice: state not in canceled/failed).
- If not previously evaluated: Proceed to Action 4 for that conversation.
- If already evaluated: Skip this conversation, optionally log `Already evaluated`.

Action 4: Look up the interaction record

- Purpose: Enforce HR scope exclusion.
- Logic: Resolve the interaction related to the conversation. If its application scope contains `hr`, skip the conversation.
- If the scope doesn't contain `hr`: Proceed to Action 5.

Action 5: buildTranscript

- Purpose: Construct the final, minute-level transcript and determine downstream skill set and guardrails.
- Steps:

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

- Aggregate all conversation messages.
- Tag user messages as `[User]:` and virtual agent messages as `[Virtual Agent]:`.
- Knowledge articles:
  - If genius results reference Knowledge articles, query the Knowledge article and replace the genius snippet with the entire article body.
  - Annotate with `[Virtual Agent]: Help articles for user query:` and wrap content within Article_Start and Article_End.
  - Constraints:
    - If the KB is HR-scoped or inaccessible, don't evaluate (skip conversation).
    - Truncate the article body to a maximum of 10,000 words.
    - If the KB content source is attached files (PDF/Word/Txt), fall back to the genius result instead of full file content.
- Catalog Items:
  - If genius results reference catalog items, query sc_cat_item and build a string: catalog name, short description, description.
  - Annotate with `[Virtual Agent]: Please choose one of the below options:` and include citation order.
- Live Agent Exclusions:
  - If the first user message requests a live agent, skip evaluation.
  - If a live agent is invoked within the first 120 words, skip evaluation.

- Outputs:
  - ExecuteEvaluation: true/false (post-guardrail outcome)
  - Chat transcript
  - Knowledge articles referred
  - Catalog items referred
  - First live agent occurrence: Sys_id of the conversation message (if present)
  - Skills to invoke:
    - Coherence Chat Evaluation
    - Conciseness Chat Eval
    - Context Retention
    - Inadequate Slot Filling Chat Eval
    - Intent Accuracy Chat Eval
    - Smooth Flowing Conversation Chat Eval
    - Truthfulness Hallucination Chat Eval
  - Additional logs
- If ExecuteEvaluation = true: Proceed to Action 7; otherwise, skip the conversation.

Action 6: If Block

- Purpose: Branch to record creation.
- Logic: If ExecuteEvaluation from Action 6 is true, go to Action 8.

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

Action 7: Chat Classifier Eval

- Purpose: Validate whether the conversation should be evaluated and extract high-level labels.
- Logic:
  - Build a lightweight transcript from sys_cs_message for classification input.
  - Invoke Chat topic classifier skill with the transcript.
  - Receive:
    - Execute evaluation: true/false
    - Topic Name
    - Category: IT or HR
- If Execute evaluation = true: Proceed to Action 6.
- If false: Skip conversation and log the classifier decision.

Action 8: Create or Update evaluation record

- Purpose: Persist an evaluation entry for this conversation.
- Table: sn_na_conv_eval_evaluation
- Field population:
  - Document conversation: Conversation reference
  - State: processing
  - Topic: from Action 5
  - Category: from Action 5
  - KB Referred: from Action 6
  - Catalog Referred: from Action 6
  - First live agent occurrence: from Action 6
  - Type: chat summarization
  - User: initiating user for the conversation
  - Message log: Additional logs from Action 6
- On success: Proceed to Action 9.

Action 9: For Loop over skills

- Purpose: Execute each selected evaluation skill.
- For each skill in the list from Action 6:
  - Action 10: invokeApiDefinition
    - Inputs: Skill Name, Conversation, Transcript, Evaluation Id
    - Behavior:
      - Invoke the Now Assist skill asynchronously.
      - The post-processor writes results into sys_generative_ai_response_validator.
      - Extract JSON response fields:

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

- Score
- Reason for Score
- Examples supporting the reasoning

- Create child metric records in sn_na_conv_eval_evaluation_metrics linked to the parent evaluation.

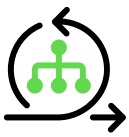
- Action 11: Wait

  Pause seven seconds before proceeding to the next skill to manage rate limits or throttling.

Knowledge Graph

The Knowledge Graph application, uses the structured and unstructured data from ServiceNow records, knowledge bases, and external sources to enhance the performance of Now Assist Virtual Agent, AI agents, and generative AI skills.

Get started

Explore  Explore Knowledge Graph	Configure  Configure Knowledge Graph	Use Knowledge Graph  Create and edit Knowledge Graph schema
Using Enterprise Graph  Learn about the Enterprise Graph schema.	Natural Language queries in Knowledge Graph  View example use cases of natural language queries in Knowledge graph.	Reference  Additional information to configure Knowledge Graph

Exploring Knowledge Graph

Knowledge Graph provides a connected representation of data that maps entities and their relationships, adding context and meaning to information to enable intelligent search, insights, and AI-driven experiences.

Knowledge Graph overview

ServiceNow Knowledge Graph application enhances the ServiceNow AI Platform® by creating a semantic layer that connects data, relationships, and context across the enterprise. It structures

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

information as a graph of entities and connections, bringing context and meaning to the available raw data. By leveraging AI, it powers personalized and intelligent experiences across the Now Assist ecosystem including Now Assist Virtual Agent, AI Agents, and AI Search to deliver more relevant insights and automation.

By linking data and embedding semantic meaning, ServiceNow Knowledge Graph transforms raw data into actionable knowledge, fueling faster resolutions, richer insights, and more personalized AI-powered experiences.

ServiceNow Knowledge Graph application:

1. Natural language queries: Allow users to query structured ServiceNow data conversationally—for example, "Who is my manager?", "What's the status of my incident?", or "What assets are assigned to me." For more examples, see Natural language queries use cases and examples
2. User context: Provide Now Assist in Virtual Agent and AI Agents with contextual information such as role, department, or location to deliver personalized and relevant responses.
3. Slot-filling: Simplify user interactions by automatically populating known information in forms or chat requests, reducing effort and friction.
4. Smarter AI Agents: Empower agents to retrieve insights, facts, and relationships directly from the knowledge graph to automate actions and improve accuracy.
5. Integration with Workflow Data Fabric: Use an integrated data layer that gives real-time, secure, and unified access to data from any source without duplicating it, enabling faster, data-rich workflows across the platform.

Knowledge Graph components

ServiceNow Knowledge Graph is a graphical representation of real-world entities (tables) and their relationships, stored in a database. Each graphical representation is called as a Knowledge Graph Schema that consist of relevant entities called nodes. The relationship between each entity is called an edge.

To create, edit and manage Knowledge Graph schemas for your specific use cases, ServiceNow Knowledge Graph has introduced, Knowledge Graph designer.

Users can use these schema to customize and extend Knowledge Graph integrations with Now Assist Virtual Agent and AI agents, and Now Assist panel. Servicenow also provides some Prebuilt integrations, that connect the Knowledge Graph with Now Assist Virtual Agent and AI agents, and other ServiceNow applications, enabling seamless AI-powered workflows and personalized experiences.

Knowledge Graph Designer

Knowledge Graph Designer is a dedicated, no-code UI where Knowledge Graph administrators (kg_admins) can effortlessly:

- Design and manage Knowledge Graph schemas, and configure the related nodes (tables), properties (columns), and their relationships.
- Analyze results of the Knowledge Graph APIs integrated in downstream products by auditing the schema using natural language queries and achieved responses.

The Knowledge Graph Designer streamlines the entire process, from schema creation and data ingestion to performance monitoring and results analysis. The new approach confirms a scalable, flexible, and intuitive way for Knowledge modeling.

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

Prebuilt Integrations

By unifying Knowledge across platforms and integrating with Now Assist, AI Search, AI agents, and skill kit, the prebuilt integrations of Knowledge Graph helps the customer drive productivity, enhance decision-making, and unlock the full potential of enterprise data, while maintaining robust data governance and permission controls.

In this release, the available prebuilt integrations are:

1. Integration with Now Assist Virtual Agent, Now Assist panel and AI agents for User Context: Helps users with personalized responses.
2. Integration with Now Assist Virtual Agent and Now Assist panel for Slot filling: Helps pre-fill the slots for Virtual Agent topics using the Natural Language Querying feature of Knowledge Graph.
3. Integration with Now Assist Virtual Agent and Now Assist panel for Employee schema: Helps requesters and fulfillers with personalized responses on people queries and Natural Language queries. Also supports people citation card. By default the user NLQ graph is connected which is used for people queries but you also have sample graph schema for other employee queries. To see more sample graphs for employee queries refer, KB article.
4. Integration with AI agents as a tool: Used to retrieve results in Natural Language and perform follow up tasks that are assigned to the AI agents.

Knowledge Graph users

Knowledge Graph users

User role	KG Functionality	Description
Knowledge Graph Admin (kg_admin)	Knowledge Graph Designer	The Knowledge Graph Admin can create and manage Knowledge Graph schemas.
Requester	Prebuilt integration with Virtual Agent and Agentic AI	Helps requesters with personalized answers, natural language queries, and fewer conversation turns with pre-filled slots for LLM topics and skills.
Fulfiller	Prebuilt integration with Now Assist panel and Agentic AI	Helps fulfillers with personalized answers, natural language queries, and fewer conversation turns with pre-filled slots for LLM topics and skills.

Knowledge Graph benefits

Knowledge Graph benefits

Benefit	Feature	Users
Enhance user experience (with prebuilt integration for downstream products like Now Assist Virtual Agent and Agentic AI)	Provides accurate data with minimal user effort.	Requester
Simple and easy to use	Creates a complex data model called Knowledge Graph schema with numerous entities and	kg_admin

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

Knowledge Graph benefits (continued)

Benefit	Feature	Users
	their relation within a few steps.	
Easy to manage	Editing Knowledge Graph schemas to add new nodes or edges is simple.	kg_admin
Customizable Knowledge Graph schema	Provides an option to copy the ServiceNow Knowledge Graph schemas for customization.	kg_admin
Test a Knowledge Graph schema	Provides an option to test a Knowledge Graph schema by running a query.	kg_admin
Use Workflow Data Fabric to retrieve information without saving or copying them from the source entity. This ensures efficiency and security.	Leverage Workflow Data Fabric tables in Knowledge Graph to retrieve data.	Requester
Create a Knowledge Graph schema with Workflow Data Fabric tables.	To create and manage Workflow Data Fabric tables, see Managing data fabric tables	kg_admin

What to explore next

To learn more about configuring and using Knowledge Graph, see:

- Configuring Knowledge Graph
- Using Knowledge Graph Designer
- Reference for Knowledge Graph

Access Knowledge Graph Schema

Use Knowledge Graph Designer to create, edit, and manage Knowledge Graph schema.

Before you begin

Role required: kg_admin

Procedure

Navigate to **All** > **Knowledge Graph** > **Knowledge Graph Designer**.
The Knowledge Graph Designer landing page displays a list of all the Knowledge Graph schemas available for users.

ServiceNow some provides prebuilt schema that are non-editable. Currently, the following prebuilt schemas are available:

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

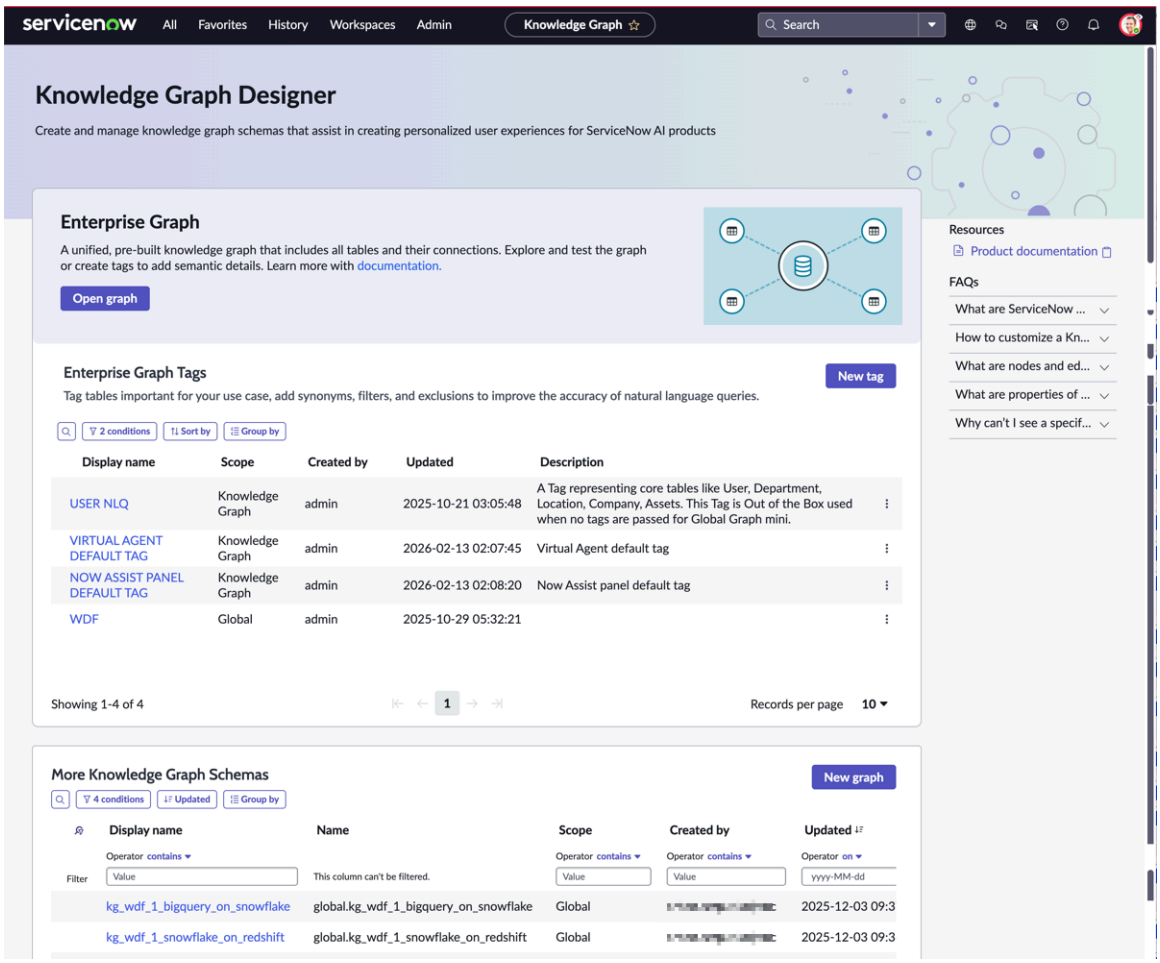
- User graph schema: Used to provide logged in user's details to Virtual Agent for personalised response. The profile section is used by AI agents for additional user context.
- Natural Language Query Scheme: Helps the requester with personalized responses on people queries and Natural Language queries. Also supports people citation card. By default the user NLQ graph is connected which is used for people queries but you also have sample graph schema for other employee queries. For more details on other prebuilt schema, see KB article.

ServiceNow® has also introduced Enterprise Graph, a pre-configured Knowledge Graph schema that maps all instance tables and their connections, enabling natural language queries from data across all tables. For more details, see Using Enterprise graph schema.

Along with Enterprise Graph, you can also create, edit and manage tags with Knowledge Graph Designer. You can use these tags to mark the key tables, that are important for answering natural language questions for specific use cases. For details, see Tagging in Knowledge Graph Designer

Knowledge Graph supports Admin to create Knowledge Graph schema with WorkFlow Data Fabric tables. This helps accessing data from different system in real time without moving it.

You can create a Knowledge Graph schema from the landing page.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Leverage Knowledge Graph prebuilt integration with Now Assist Virtual Agent and Now Assist Panel

The prebuilt integrations of Knowledge Graph can help ServiceNow users to drive productivity, enhance decision-making, and unlock the full potential of enterprise data- while maintaining robust data governance and permission controls.

In this release, the available prebuilt integrations with Now Assist Virtual Agent and Now Assist Panel are:

1. Integration with Now Assist for User Context: Helps requesters and fulfillers with personalized responses.
2. Integration with Now Assist for Slot filling: Helps requesters and fulfillers in pre-filling the slots for LLM topics and skills execution using Natural Language Querying of Knowledge Graph.
3. Integration with Now Assist for Natural Language Query graph: Helps requesters and fulfillers with personalized responses on people queries and Natural Language queries. Also supports people citation card.

Note: To enable Knowledge Graph for Now Assist Virtual Agent, ensure that *sn_vad_genai.knowledge_graph.enabled* and *sn_ais_assist.enable_knowledge_graph_nlq* system properties are set to true. See Add a Knowledge Graph schema to a chat assistant.

For more information see Add a Knowledge Graph schema to a chat assistant.

Note:

- To enable Knowledge Graph functionality in the Now Assist Virtual Agent and conversational catalog, make sure to fill the Slot Filling Schema.
- To use Knowledge Graph for Question Answering from VA, fill out the Natural Language Query Schema in the Now Assist Panel.
- You can only have one active Knowledge Graph schema at a time for either Slot Filling or Natural Language Query. Multiple active Knowledge Graph schemas are not supported.

Integration with Now Assist for User Context

For the users of Now Assist, Knowledge Graph integrates the context from the prebuilt User Profile Schemas that provide personalized responses.

By leveraging relationships between users, teams, and content, products like AI Search and Now Assist can provide relevant, permission-aware answers instead of generic results.

With Knowledge Graph, responses are dynamically tailored based on:

- Who the user is: Role, department, and location
- Who they collaborate with: Manager, reportee
- What assets do they have

Here's an example use case:

- An employee uses Now Assist Virtual Agent for information on parental leave policy. They enter the query in the Virtual Agent window `What is my parental leave policy?`
- Virtual Agent receives the user information like. The employee is based in the country: USA, state: California, City: Santa Clara, from the Knowledge Graph User Profile Schema.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- This additional user profile context is used to personalize the synthesized response to the exact location of the employee
- Therefore, instead of getting a link to the parental leave policy document or a generic response, the employee gets a tailored contextualized answer:

```
Your company offers a generous parental leave policy to its
employees in California. As of January 2022, the company
increased paid time off for workers who give birth to a maximum
of 24 weeks from the previous 18. In addition to the company's
internal policies, California state law provides further
protections. The California Family Rights Act (CFRA) offers
eligible employees up to 12 weeks of unpaid, job-protected
leave to care for their own serious health condition or that
of a family member, or to bond with a new child. This is
complemented by the Pregnancy Disability Leave (PDL) law, which
provides up to four months of unpaid, job-protected leave for
employees disabled by pregnancy, childbirth, or related medical
conditions.
```

Integration with Now Assist for Slot filling

Knowledge Graph enhances the Now Assist user experience and makes the process seamless and efficient by reducing the slot-filling questions asked during conversations.

Here’s an example use case:

An employee uses Virtual Agent to request a laptop replacement. Virtual Agent uses the assigned Knowledge Graph schema to find information and resolve the query with minimal user inputs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. The user uses Virtual Agent to query `Need assistance in laptop replacement.`
2. Virtual Agent processes the query and generates the following prompts required for this request:
   - Topic: New laptop request
   - Name
   - Location
   - Department
   - Laptop model
   - Address
3. Virtual Agent first tries to gather this information using Knowledge Graph.
4. The Knowledge Graph schema leverages LLM to retrieve the data from all the relevant entities, called nodes using the relationship between these nodes, called edges, and provides the following output:
   - Topic: New laptop request
   - Name: John Doe
   - Location: Santa Clara
   - Department: Marketing
   - Badge Template: User input needed
   - Address: 123 Street, CA, USA
5. Virtual Agent requests verification of the output and add details for the missing fields.
6. The user can edit and verify the provided information before confirming. Once confirmed, the request is processed with minimal effort from the user.
7. Virtual Agent processes the user input and completes the user query.

Knowledge Graph leverages the existing information available in the internal databases and auto-populates it to reduce the efforts while making the entire experience seamless.

Integration with Now Assist for people queries

Now Assist can now provide users with information about people in your organization.

If you ask Virtual Agent about a person,the Information about that person you're searching for appears in the synthesized response, along with an inline people citation.

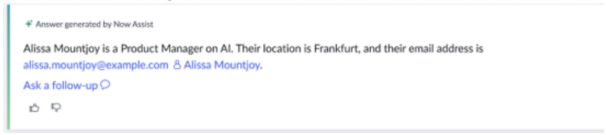
Inline citations appear at the end of the relevant synthesized response sentence. Selecting an inline citation results in a popover containing either a link to an article or source, or a description and action to start the action.

Note: Shared files only appear if Knowledge Graph admin has activated the *sn_kg_conn_user_shared_files* record in Knowledge Graph related data map [sn_kg_related_data_map_list] table. To active it, set the **Active** field to **True**.

Here's an example:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example of synthesized response for people on the portal's search results page



Example of synthesized response with people inline citations in chat



Selecting the person's name presents a popover. The information in the popover can include the following information:

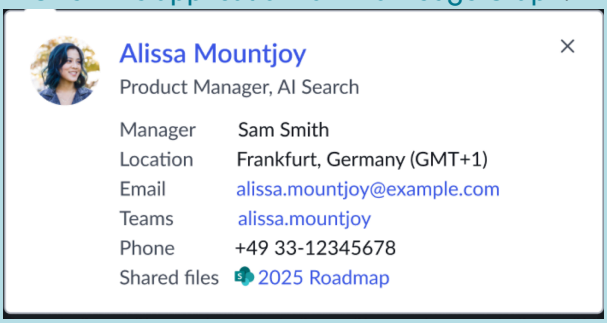
- Manager
- Location
- Email
- Teams
- Phone

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Shared files

Important:

- Shared Microsoft SharePoint files between you and the person found, appear only on the people popover.
- The shared files only appear after you have completed the prompt to **Log in**, and signed in successfully. If you do not have a valid token, you will be prompted to sign in and re-directed to Microsoft login page.
- If you have not configured Microsoft OneDrive application, see Configure Microsoft OneDrive application for Knowledge Graph.



Integration with AI Search

Knowledge Graph is also integrated with AI Search. You can enable Now Assist Multi-Content Response Genius Results to get answers from Knowledge Graph on AI Search.

Ensure that you have already configured Knowledge Graph with an assistant before enabling it on AI search.

To configure and enable the AI Search see Now Assist Q&A Genius Results.

Leverage Knowledge Graph prebuild integration with AI agents

The prebuilt integrations of Knowledge Graph, can improve live agent productivity by assisting AI agents in performing tasks and answering user queries.

In this release, the available prebuilt integrations with AI agents are:

1. Integration with Now Assist AI agents for User Context: Helps users with personalized responses.
2. Integration with AI agents as a tool: Used to perform specific tasks that are assigned to the AI agents.

Integration with AI agents for User Context

For the users of Now Assist, AI agents integrates the context from the prebuilt User Profile schema to fetch relevant data and provide personalized responses.

By leveraging relationships between users, teams, and content, Now Assist AI agents can fetch relevant, permission-aware user information and provide answers to reduce slot-filing requirement.

Here's an example use case of how Knowledge Graph is used within AI agents for user context.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- A user needs assistance in writing their resume, so they use ServiceNow AI agents called Resume builder.
- The user ads their task: `Build my resume`.
- AI agents uses Knowledge Graph to fetch the following information from the user profile schema:
  - First name
  - Last name
  - Location
  - Email ID
  - Phone number
  - Occupation
- AI agents reverts with the available information and request for the missing information to proceed building the resume.
- Once the agent as all the required information, it provides user with output.

Integration with AI agents as a tool

Knowledge Graph can now be used as a tool within AI agents. Users can choose Knowledge Graph to perform tasks, while creating an Agent.

You can define the flow action to use Knowledge Graph as a reusable operation in automating the ServiceNow AI Platform features without having to write code.

See Add a Knowledge Graph to an AI agent to add Knowledge Graph to an AI agents in AI Agent Studio.

Example use case:

- A user wants to view and retire all the assets assigned to another user. To execute this task, user will use Asset Manager agent.
- Asset Manager uses Knowledge Graph to fetch the asset information related to the user.
- After processing the query, the agent provides the following output:`User has the following assets:1.<asset1> 2.<asset2> 3.<productID>- <asset 3> <purchase date>. The email address for the user is <emailId> and username <username1>.`
- User then proceeds with the task of retiring the assets: `Proceed with retiring the assets associated with Username1.`
- The next tool, used execute this task, will take the table name and sysId displayed in Knowledge Graph output and proceed with the task.

Example: Knowledge Graph API output in AI agents

When a user ask a query `What is my Manager's name`?, AI agents reaches out to Knowledge Graph to fetch this information. Here's an example output of the query:

```
{
    "Manager": [
        {
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
            "sys_user": {
                "user_name": {
                    "value": "abel.tuter",
                    "displayValue": "abel.tuter"
                },
                "sysId": "62826bf03710200044e0bfc8bcbe5df1",
                "name": {
                    "value": "Abel Tuter",
                    "displayValue": "Abel Tuter"
                },
                "first_name": {
                    "value": "Abel",
                    "displayValue": "Abel"
                },
                "last_name": {
                    "value": "Tuter",
                    "displayValue": "Tuter"
                }
            }
        }
    ]
}
```


In this example, the table that is called is `sys_user` and column referred is `user_name`. The output also displays SysId , example: `"sysId": "62826bf03710200044e0bfc8bcbe5df1"`.

Natural language queries use cases and examples

View example use cases of some natural language queries used in Knowledge Graph and Enterprise Graph.

Natural Language Query support in Knowledge Graph and Enterprise Graph allows users to interact with structured data using human-like language. This feature enhances the user experience by enabling more intuitive and flexible search capabilities within the ServiceNow platform.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Supported Queries

Knowledge Graph and Enterprise Graph, understand and process queries that reference specific tables, columns, choice values, and conditions. It also handles queries involving system-generated identifiers (sys-ids), person names, dates, and contextual pronouns. Additionally, it supports aggregate and sorting queries, enabling users to perform simple statistical analysis and order results.

Here are some examples of the supported query types:

Queries with specific reference

Queries having reference to all table names, column names, or choice values and conditions required are supported. Ensure that the conditions are specified exactly.

Here are some example queries and description of how they works in Knowledge Graph.

Example use cases

Example	Description
Show all stories which are ready for acceptance.	Stories refers to stories table and ready for acceptance refer to choice.
Find details of CI associated with P1 priority incident.	CI refers to look at configuration table and find those CI which are liked to incident with priority 1.
List all change requests with high risk.	Change request refers to change_request table, risk refers to column, and high refers to choice.
Give details of CI named "AppleMacBookPro15".	CI refers to configuration item table and named refers to name column. The exact name of CI is also specified.
Show all incidents which are open.	Incident refers to the table; "in progress," "new," and "on hold" are choice value in state column options that correspond to open.
Share all case task assigned to "Dev – Knowledge Graph" group.	Case task is reference to table, group is a reference to assignment group column, and exact condition to look for is specified exactly.
Find Configuration items where manufacturer = Lenovo.	Configuration item is reference to cmdb_ci table, manufacturer is column, and the exact name of manufacturer is specified.
Show me assets assigned to <Person Name>.	Assets refer to asset table, assigned to is column and person name is well specified.
Find details of CI associated with P1 priority incident.	CI refers to look at configuration table and P1 indicated the choice value of the priority column.

Queries with Sys-IDs, Person Names, and Date References

Queries having person names or references to dates such as yesterday, last month and contextual pronouns to the right person. Queries using system-generated identifiers (sys-ids) also works and can identify the table based on sys-ids.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Here are some example queries and description of how they works in Knowledge Graph.

Example use cases

Example	Description
Find details of request items linked to REQ0010144.	The request item refers to sc_req_item table and REQ sys-id refers to to sc_request table.
Show details for INC0000044.	INC as sys-id refers to Incident table.
Get status for CHG000567.	CHG as sys-id refers to change request table and the status refers to the status column for that record.
Find department for person who requested RITM0010144.	The RITM sys-id refers to sc_request_item and department refers to department column from cmn_department which is linked to the person who requested RITM0010144.
Find details of problem linked to INC0000047.	This query requests the details of problem record linked to incident INC0000047.
Show incidents assigned to Fred Luddy.	Finds Incident table with user in assigned to column.
List tasks created by Beth Anglin.	Task refers to task table and User name refers to created by column.
Find approvals pending for Joe’s reportees.	Find approvals from the sysapproval_approver table for approval that are pending for the user's reportees.
Display incidents created more than 3 days ago and are not resolved yet.	Uses Incident table, created date, and state column to fetch the details.
Find certificates expiring in the next 30 days.	Certificates refer to table name, the next 30 days reference can be converted to date related condition and can be understood from the date reference in query.
Find case task created in last 30 minutes.	Case task refers to table, created refers to created column and time related condition can be understood from time reference in query.
What are the Problems open with my group?	Problem refers to problem table, group refers to column, ‘my’ reference will be resolved to the person asking the query.
Show my open incidents.	Uses Incident table and 'my' reference is the person asking the query.
What is the status of my request?	Uses request table and 'my' reference is the person asking the query.
What are their pending approvals?	Uses `sysapproval_approver` table and 'their' reference is the person who was referred to in previous turn of conversation.

Aggregate or Sorting Queries

These queries allow users to perform simple statistics and sorting directly in the Virtual Agent or Now Assist panel.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Here are some example queries and description of how they works in Knowledge Graph.

Example use cases

Example	Description
Count of incidents grouped by priority.	Counts the records in incident table grouped by priority column.
Count of problems grouped by stage.	Counts the records in problem table grouped by stage column.
Count case task assigned to group "Dev – Knowledge Graph".	Counts the records in case task table that are assigned to assignment group = Dev – Knowledge Graph.
List in progress hr cases by ordered by created date.	Lists the HR cases from HR cases table ordered by the created date field.
Show top 5 users by number of incident assigned to them.	Lists the top 5 user by number of incidents from the incident table assigned to them.

Unsupported Query Categories

ServiceNow's Knowledge Graph does not support the following types of queries:

Queries Missing References or with Misspellings

Queries that lack references to table, columns, choice values, or conditions, or contain misspellings, are not supported.

Here are a few examples of unsupported queries that will work if rephrased, as suggested below:

Example unsupported queries

Example	Rephrase to:
Find incidents about SAP login issue.	Find incidents where short description = SAP login issue.
Find CI where manufacturer is Orrange. (Incorrect spelling)	Find CI where manufacturer is Orange.
Find CI in SanDiego city.	Find CI in San Diego city.
Locate changes related to patch upgrade.	Locate change request where description contains patch upgrade.
Change request for Abel Tuter.	Change request assigned to Abel Tuter or Change request opened by Abel Tuter.
Problem by Mark Andrew.	Problem updated by Mark Andrew or Problem closed by Mark Andrew.

Queries on Journal Fields or Document IDs

Queries involving journal entries, document IDs, or glidelist references are not supported.

Here are a few examples of unsupported queries:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Show latest comments added to INC0000044.
- What comments or work notes were added in the approval of change request CHG0000046.
- Show tasks with 'urgent' in comments.
- Find incident where Joe commented about email resolution steps.

Keyword Queries or Queries Requiring KB Resolution

Keyword-based queries or those requiring resolution from the Knowledge Base (KB) are handled by AI Search, not the Knowledge Graph.

Here are a few examples of unsupported queries:

- VPN issue
- Password reset
- Incident SLA
- How can I request a windsurf license?
- How to install global protect for PC?

Workflow Data Fabric integration with Knowledge Graph

Knowledge Graph supports Workflow Data Fabric to fetch third party information securely without copying the data.

The integration of Workflow Data Fabric (WDF) with the ServiceNow Knowledge Graph (KG) allows organizations to seamlessly connect external and internal data sources. This integration extends the functionalities of the Knowledge Graph to external data, ensuring real-time access and enabling advanced insights, workflow automation, and personalization using external data.

The key feature of Workflow data fabric integration with Knowledge graph are:

- Real-Time Access: Zero Copy connectors from WDF and the semantic layer of the Knowledge Graph ensure that external data is accessible in real-time.
- Deep Data Connection: WDF tables are automatically included in the Enterprise Graph, enhancing the connectivity and utility of external data within the ServiceNow platform.
- Natural Language Query Support: Users can ask natural language queries (NLQ) in Virtual Agents and AI Agents, leveraging both external (WDF) and internal data.

The benefits of using Workflow data fabric integration in Knowledge Graph are:

- Requestors and Fulfillers: Can ask NLQ questions and receive insights from external (WDF) tables integrated with the KG schema.
- Admins: Can create a KG schema with WDF tables and link it to the Now Assist Panel, AI Agents, or Now Assist Virtual Agents.

Admins can add WDF tables to the Knowledge Graph just like any other internal table. However, only WDF tables configured with a primary key are supported.

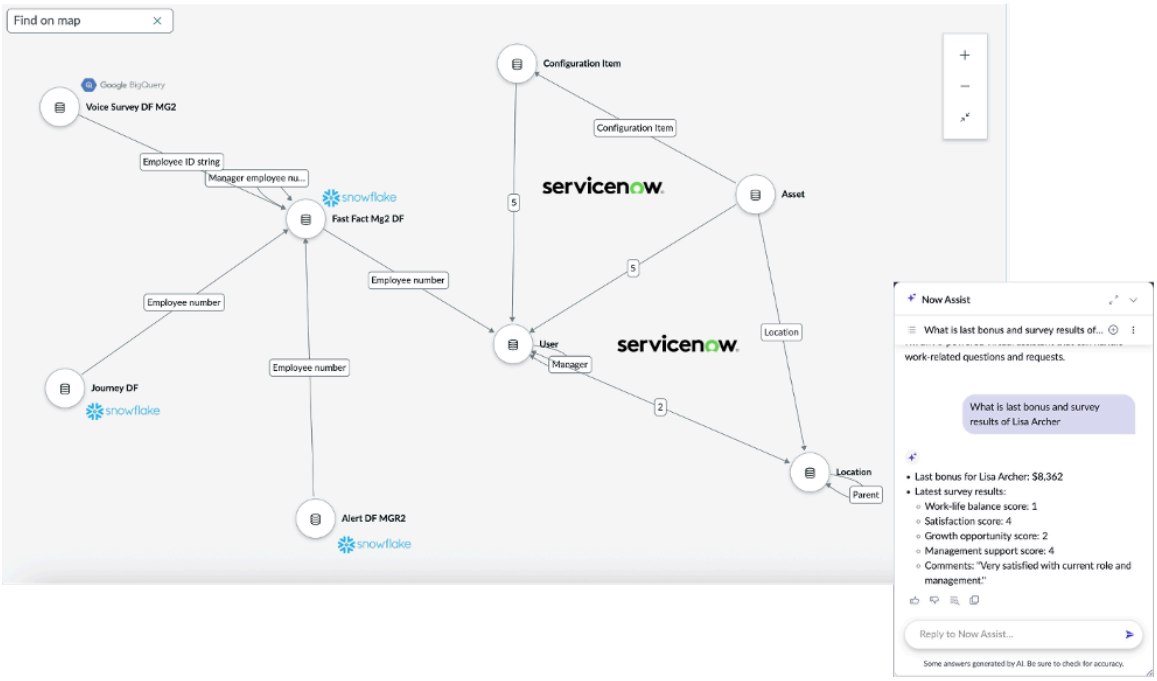
Here are some integration best practices:

- Primary Key Requirement: Always add a primary key to your WDF table during configuration. This is mandatory for supporting WDF tables in the Knowledge Graph.
- Meaningful Naming: Use clear and meaningful names when creating WDF tables. Add synonyms to your tables in the KG schema as needed to improve the accuracy of natural language queries.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Here is some example use cases:

Snowflake and Google BigQuery Integration: WDF tables sourced from Snowflake and Google BigQuery can be added to the Knowledge Graph schema, allowing users to query and analyze data from these external sources within ServiceNow.



Improving Natural Language Queries with Tag configuration

AI instructions add more business context to natural language queries. They guide users and promote responses that more closely align with their needs and expectations.

AI instructions

AI instructions help clarify how the knowledge graph interprets user queries and accesses data. When a user submits a query, relevant instructions at all applicable levels like node, property, and edge are considered.

Referencing data within tables directly in instructions is not a general guideline. Instead, instructions should be generalized and context-driven.

You can add instructions at three levels:

- Node (entity/table)
- Property (column/attribute)
- Edge (relationship)

The Always Include flag

The option to set an **Always include** flag enables users to indicate that a particular instruction should be considered unconditionally. This is useful for enforcing critical business filters, such as always excluding certain asset statuses unless requested otherwise.

Instructions:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Use **Always include** for business-critical filters that must apply regardless of user query phrasing.
- Omit **Always include** for context-sensitive instructions that should only apply when relevant.
- Verify that flagged instructions do not conflict with each other or produce duplicate results.

Examples of Good AI Instructions:

Effective AI instructions are clear, precise, and contextually aware. The following examples illustrate the recommended approach:

- "Always add WHERE alm_asset.install_status <> 7 to filter out retired assets by default." This instruction respects business context automatically. Retired assets are excluded unless the user explicitly requests them.
- "When generating a query involving sc_cat_item, don't include state='published' unless explicitly mentioned." This avoids unintended filtering caused by ambiguous instructions.
- "For life-cycle-related queries, use sam_sw_product_lifecycle_report as the primary/default table."- This aligns query logic with performance and business requirements.

Identifying good vs. bad instructions

Good Instruction	Bad Instruction
Consider normalized_company only when user asks for normalised manufacturer.	Check company (no indication of which property or relationship to use).
Always add WHERE alm_asset.install_status <> 7 to filter out retired assets by default.	Filter by manufacturer (no column or condition specified).

Poor instructions may result in hallucinations. The AI generates incorrect data or misinterprets the query intent. This can occur due to:

- Missing context or incomplete references.
- Ambiguous column or table names.
- Contradictory guidelines across instruction levels.

AI instructions must adhere to the same basic guidelines as user queries. Reference structured elements, avoid unsupported query types, and follow prescribed business logic. What applies to good queries also applies to effective AI instructions.

Support for Synonyms and Data Filters

Synonyms:

Knowledge Graph supports configuring synonyms for tables, columns, edges, and reverse edges to enhance flexibility and user-friendliness. Synonyms enable users to define alternate names for these entities, accommodating different terminologies.

- Each entity can have up to five synonyms.
- Synonyms promote broad coverage for commonly used alternate terms.
- Maintaining a manageable synonym set reduces ambiguity in query resolution.

Data Filters:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Data filters are applied deterministically as post-processing rules to refine query results. After initial Cypher generation, filters are systematically applied to promote consistency and reliability of outcomes.

- Filters are applied after Cypher generation, not during.
- Dynamic filters enable real-time adaptation based on user context.
- Improves precision and relevance of returned data. Does not modify the core graph query.

CI relationships and Knowledge Graph

CI Relationships (Rel CI) enable the Knowledge Graph to answer natural language questions about service dependencies and infrastructure topology by storing typed parent-child relationships between CMDB configuration items.

The *CMDB_REL_CI* table stores relationships between configuration items (CIs) in the ServiceNow CMDB. Each relationship connects a parent CI to a child CI through a defined relationship type, enabling the Knowledge Graph to understand and traverse the topology of your IT environment.

Rel CI support in the Knowledge Graph allows users to ask natural language questions about how services, servers, databases, and other CIs relate to one another without writing queries or navigating CMDB tables directly.

How CI relationship data is stored

Each record in the Rel CI table represents a directional relationship between two CIs. Relationships are two-way and are described by a relationship type that consists of a parent descriptor and a child descriptor, separated by double colons:

`<parent descriptor>::<child descriptor>`

This means every relationship can be read in two directions:

- Parent → child: read using the parent descriptor
- Child → parent: read using the child descriptor

For example, a record in the Rel CI table has Bond Trading (*cmdb_ci_service*) as the parent, lnux100 (*cmdb_ci_linux_server*) as the child, and a relationship type of `Depends on::Used by`. This relationship is read as:

- Bond Trading depends on lnux100
- lnux100 is used by Bond Trading

Note: The direction of the relationship affects how you phrase queries to the Knowledge Graph. Parent-to-child queries (for example, "depends on") return more reliable results than child-to-parent queries (for example, "used by").

Knowledge Graph support for CI relationships

The Knowledge Graph can answer questions about CI relationships when the query clearly specifies all three of the following:

- The class of the parent CI (for example, service)
- The relationship direction — either the parent descriptor or child descriptor (for example, depends on)
- The class of the child CI (for example, Linux server)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Class hierarchy inheritance

When you define a relationship between two CI classes, the Knowledge Graph automatically extends that relationship to all classes higher up in the CI class hierarchy. This means users can query at a more general class level and still get results across all matching subclasses.

For example, a relationship defined between service and linux server also applies to server, which is a parent class of Linux server in the hierarchy. Querying for servers rather than Linux servers returns results across all server subclasses — including Linux server, windows server, UNIX server, and others.

Note: If a query returns fewer results than expected, try using a broader parent class in the hierarchy (for example, "server" instead of "Linux server") to include all inherited CI types.

Supported query patterns

The following table shows examples of queries the Knowledge Graph can answer using Rel CI data. Each query specifies a parent class, a relationship descriptor, and a child class.

Supported query patterns

Scenario	Example Query
Service that depends on a specific Linux server	Which services depend on 'lnux100' Linux server?
Servers that a specific service depends on	'Bond Trading' service depends on which UNIX server?
All server types a service depends on (using hierarchy)	'Bond Trading' service depends on which server?
Computers connected to a database	Which databases are connected by computers?
Multi-hop relationship across three CI types	What database runs on UNIX server that connects to 'nc6500-a01' network gear?

Unsupported query patterns

The following query types are not currently supported. Use the recommended alternatives to get the results you need.

Unsupported query patterns

Limitation	Unsupported Query	Recommended Alternative
Negation of a relationship	Which business capabilities have no related business applications?	Rephrase to ask for what does exist rather than what does not.
Unspecified relationship type	Show me services related to Linux servers.	Show me services depending on Linux servers.
Skipping steps in a multi-hop path	Show me servers in New York.	Show me servers in racks present in datacenters located in New York.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Enabling Rel CI for Knowledge Graph

Rel CI support for the Knowledge Graph is inactive by default. Set both of the following system properties to true to enable it:

System properties for Rel CI support

System Property	Purpose
sn_kg.description_generation.enable_cmdb_rel_ci	Enables description generation for Rel CI data
sn_kg.query.enable_cmdb_rel_ci	Enables Knowledge Graph querying against Rel CI data

Note: After enabling both properties, allow time for Rel CI data to be fully indexed before running queries. Results may be incomplete until indexing is complete.

Configuring Knowledge Graph

Plan and configure Knowledge Graph implementation.

Configuration overview

Knowledge Graph is a ServiceNow AI Platform feature that is available with activation of any ServiceNow generative AI application.

On activation, the following Knowledge Graph plugins are installed by default:

- Graph QL plugin
- Knowledge Graph application

To implement the Knowledge Graph adoption, you must configure Knowledge Graph for Virtual Agent.

For pre-built integration with Now Assist Virtual Agent for Slot filling, configure Knowledge Graph in Virtual Agent. Refer Configuring assistants overview for admin guided setup.

Knowledge Graph now supports Workflow Data Fabric that enables users to retrieve data from different systems without moving them. This ensures efficiency and security when working with sensitive information.To configure and integrate the Workflow Data Fabric tables for Knowledge Graph, see Managing data fabric tables

Configure LLM for Knowledge Graph

Choose which large language model (LLM) service provider for Knowledge Graph.

Before you begin

LLMs are part of the foundation for Knowledge Graph. Different LLMs can provide slightly different performance and responses. You can select which LLM to use for your Knowledge Graph schema.

Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **sys_one_extend_capability.list** > **TextToResult**.
2. Open the TextToResult record and select **OneExtend Definition Configs** (related list).
3. For one of the following options, change the value of the default field to **True**.
   - **TextToResult (NowLLM)**
   - **TextToResult (Claude Large)**
   - **TextToResult (GPTLarge)**
   - **TextToResult (Gemini)**

   Once you select the default LLM, ensure that the Default field value for another LLM is set to **False**.
4. Select the LLM to open the record and update the default field to **True**.

Configure Microsoft OneDrive application for Knowledge Graph

Use Microsoft SharePoint for fetching user-specific external data, such as shared files, from external services through a Knowledge Graph API.

Before you begin

Knowledge Graph uses Microsoft OneDrive application for authentication of Microsoft SharePoint required to fetch external data such as shared files used for people citation in Virtual Agent. Use the below process to setup the necessary authentication, used for people citation in Virtual Agent.

To view the share files, configure the Microsoft OneDrive application with Knowledge Graph.

To complete the configuration, you must:

- Create Microsoft OneDrive application.
- Setup Knowledge Graph application on Microsoft OneDrive tenant.
- Update Microsoft OneDrive permissions for delegated access.
- Create a new client secret and complete the authentication process.

Role required: admin

Procedure

1. Go to portal.azure.com and select **view** to Manage Microsoft Entra ID.
2. To register the app, do the following:
   - Select **Manage** > **App registration** and select **New Application**.
   - Add a Name for the application
   - Select **Accounts in this organizational directory only** single-tenant apps for use.
   - Select Register.

     Note: Copy and save the Application ID and Tenant ID for future reference.
3. Once the new application is registered, add the client credentials:
   - Select **Add a certificate or secret**.
   - Select **New client secret** and add Description and Expiry duration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Select **Add**.
- Ensure that you copy and save the New client secret value that is created.

4. Select **Overview** from the left navigation pane to add the redirect URL.
   - Select **Add a redirect URI**.
   - Select **Add a platform** from the platform configurations section.
   - Add `https://<instanceURL>/oauth_redirect.do` in **Redirect URIs** and **Front-channel logout URL** fields on the configure web page.

     Note: Replace the <instanceURL> with your instance path. Example: *abc.service-now.com*.

   - Select both the Access token and ID tokens option from the **Select the token you would like to be issued for authorization endpoint** section.
   - Select **Configure**.
5. Select **Manage** > **App permissions** from the left navigation pane, to add permissions.
   - Select **Add a permission**.
   - Select **Microsoft Graph**.
   - Select **Delegated permission**.
   - Add and select *offline.access* and *Sites.Read.All* in Select permissions section.
   - Select **Add permissions**.
   - Ensure that the **Admin consent required** field is set to Yes for the newly added Microsoft Graph.
6. Go to your ServiceNow instance to change the Application registry settings:
   - Select **All** > **System OAuth** > **Application registry**
   - Select **KG MS OneDrive Delegated Connector**.
   - Add the copied Application ID in the **Client ID** field.
   - Add the Client secret in the **Client secret** field.
   - Add the Tenant ID in the placeholder for [tenantId] in the Authorization URL and Token URL field.
   - Ensure that the Redirect URL is correct.
   - Select **Update**.

Using Knowledge Graph Designer

Use Knowledge Graph Designer to create customized Knowledge Graph schemas that consist of nodes and edges.

Knowledge Graph schemas are the customized Knowledge graphs that consist of nodes and edges.

Like tables, nodes also denote and store details about entities such as user, location, department. You can add nodes to the Knowledge Graph schema so that the system can reach out to all the relevant tables to fetch the data. You can also edit node properties such as data source, start node, end node, and connected nodes.

Knowledge Graph enables you to select the node columns that can be queried.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The relationship or connections between these nodes are referred to as edges. You can select and edit the connecting edges and available edges for each node.

With Knowledge Graph Designer, you can create, edit, duplicate, or delete a Knowledge Graph schema.

Apart from ServiceNow tables, you can also create Knowledge Graph schema using WorkFlow Data Fabric tables to retrieve data from different systems in real time without moving it.

There's an option to test a Knowledge Graph schema by generating and running a query.

View Knowledge Graph schemas

By default, ServiceNow published Knowledge Graph schemas are available with the product. View the prebuilt schemas provided by ServiceNow, or the schemas created by your organization admin. There are currently two prebuilt schemas available that are read-only but can be cloned and edited:

1. User profile schema
2. Employee schema

Create a Knowledge Graph schema

Create customized Knowledge Graph schema that will be used by Virtual Agent, AI Agents and Now Assist Panel.

Before you begin

Role required: kg_admin

Procedure

1. Navigate to **All** > **Knowledge Graph** > **Knowledge Graph Designer**.
   The UI displays a list of all the Knowledge Graph schemas on the landing page.
2. Start creating a Knowledge Graph schema by selecting **Create New**.

   Note: You can create Knowledge Graph schema using ServiceNow tables or the external Workflow Data Fabric tables.

3. On the form, fill in the fields.

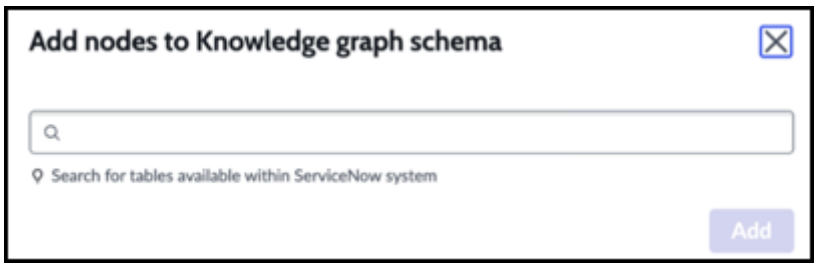
Edit Knowledge Graph Schema details form

Field	Description
Display Name	Display name for the Knowledge Graph schema.
Name	Optional name for the Knowledge Graph schema.
Scope	Scope used when creating the Knowledge Graph schema. This field is a read-only.
Description	Knowledge Graph schema overview to provide additional information to users.

4. Select **Create**.
   The Add nodes to the Knowledge graph schema window is displayed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

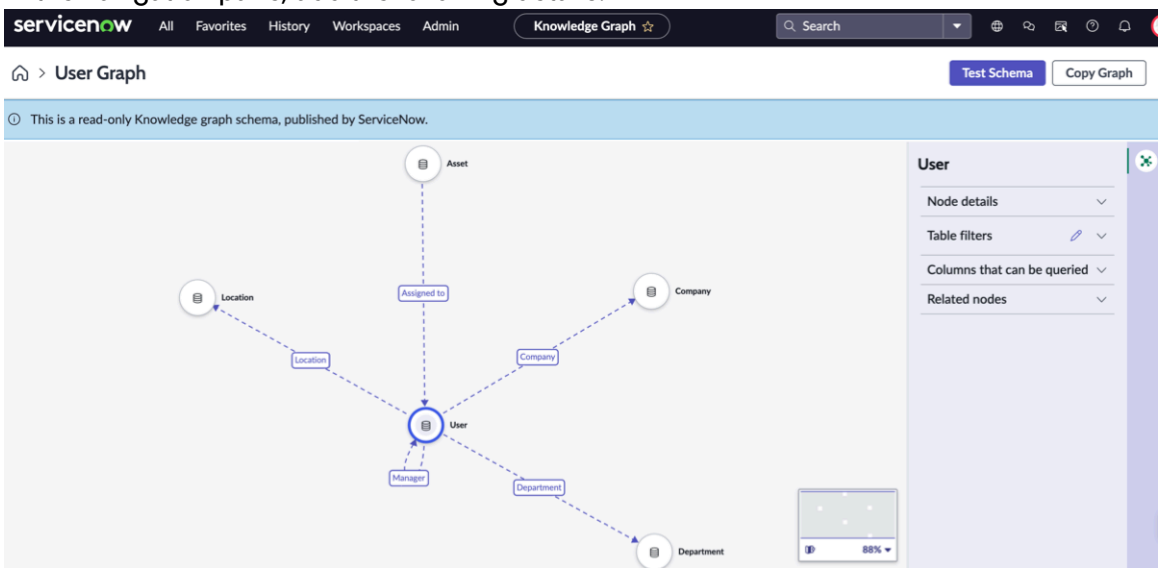
5. Enter or search for the nodes that you want to add to the Knowledge Graph schema and select **Add**.



You can search and select the Workflow Data Fabric tables, if integrated.

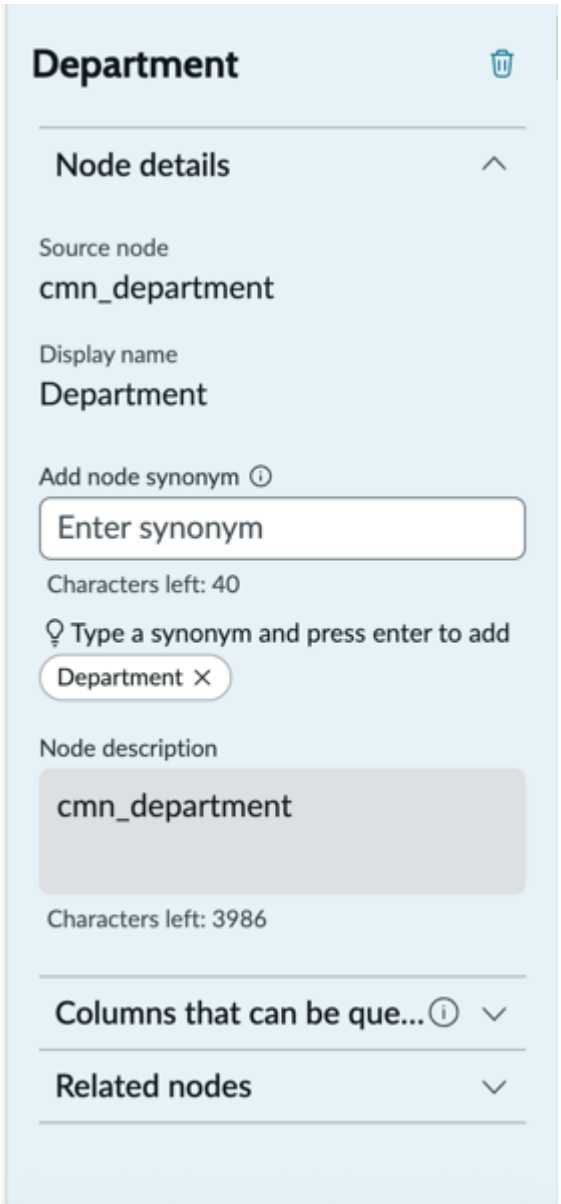
The Knowledge Graph schema is created and displayed on the Knowledge Graph canvas.

6. In the navigation pane, add the following details:



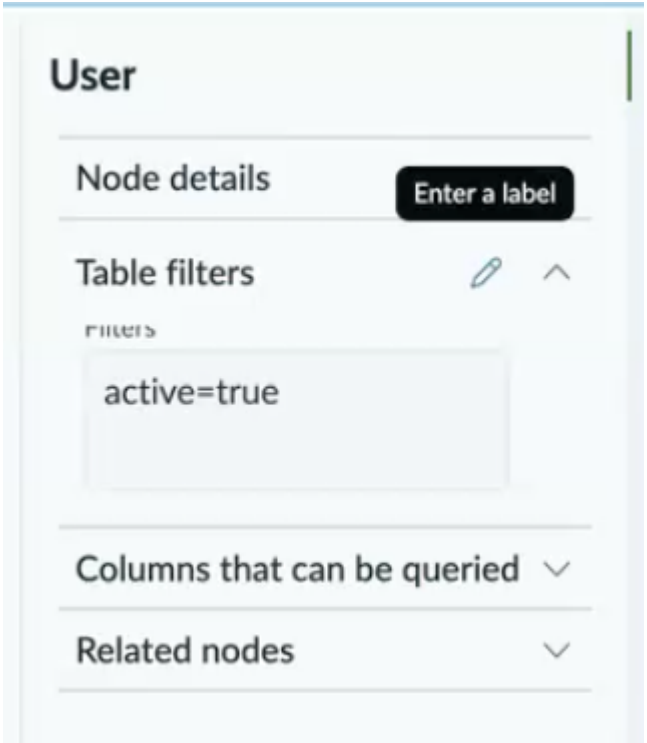
7. In the Node details section, you can add or edit the following fields.
   - Add node synonym
   - Node description

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



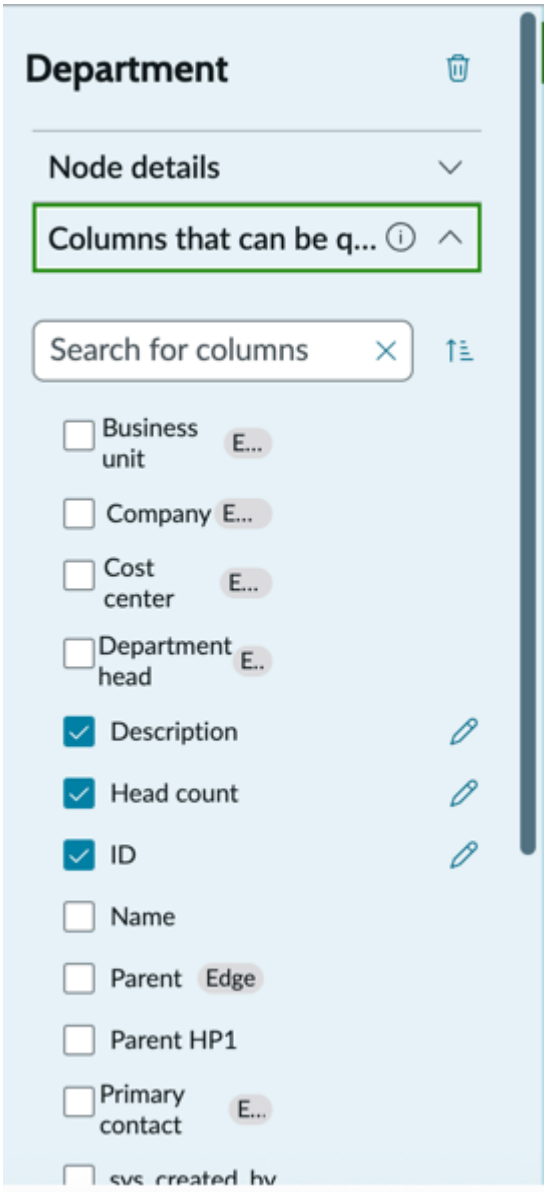
8. In the **Table filters** section, view and add filters to set rules that control which records are shown in query results.
   To add table filters, provide the following information:
   - Field
   - Operator
   - Value
   - Add Condition set: optional field to add an alternative filter condition.
   - Apply filters to child tables: optional check box to add the same filter conditions for the child table of the selected table in the graph.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



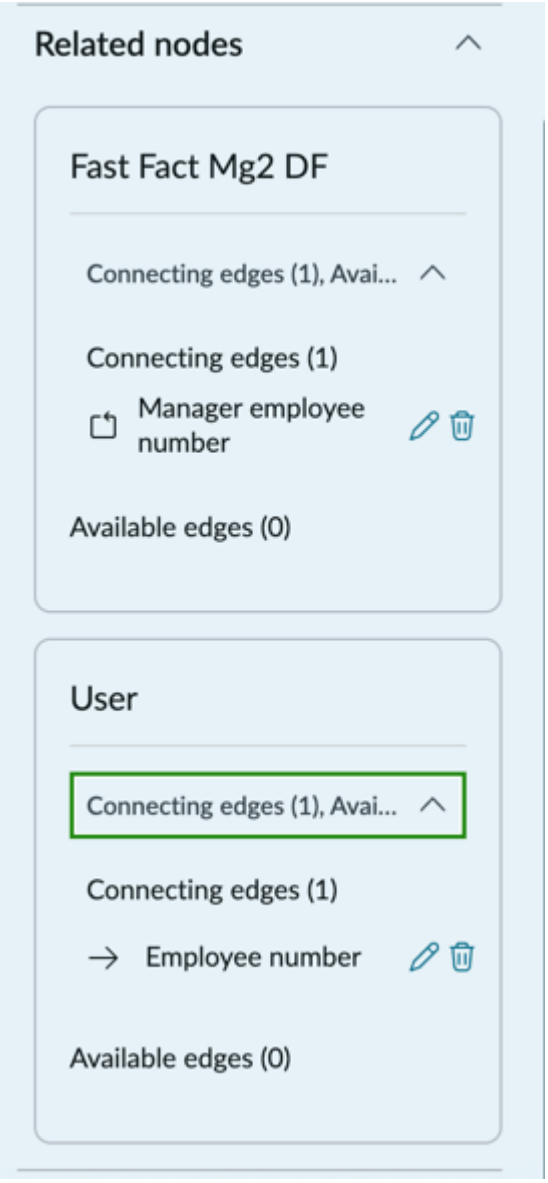
9. In the Columns that can be queried section, search for and select the desired columns and select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

10. Add, delete, or edit edges in the Related nodes section and select **Save**.



Edit a Knowledge Graph schema

Edit a Knowledge Graph schema to customize or update the existing schema details.

Before you begin

Role required: kg_admin

Procedure

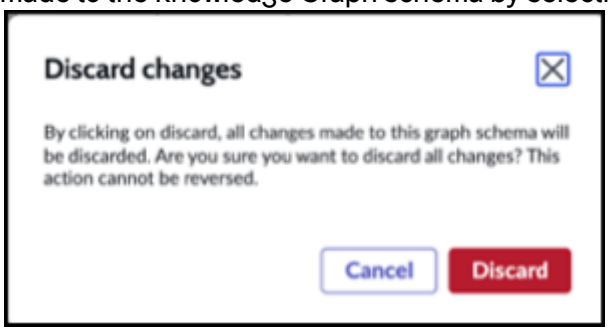
1. Navigate to **All** > **Knowledge Graph** > **Knowledge Graph Designer**.
   The UI displays a list of all the Knowledge Graph schemas on the landing page.
2. From the list of existing Knowledge Graph schemas, select a Knowledge Graph schema to edit.
   The Knowledge Graph schema opens in the Knowledge Graph canvas.
3. Select the edit icon ( ).
4. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Edit Knowledge Graph Schema details form

Field	Description
Display Name	Display name of the Knowledge Graph schema.
Name	Name of the Knowledge Graph schema.
Scope	Scope used to create the Knowledge Graph schema.
Description	Knowledge Graph schema overview for users.

5. To save the changes, select **Save**.
6. Delete all the changes made to the Knowledge Graph schema by selecting **Discard**.



Manage nodes in a Knowledge Graph schema

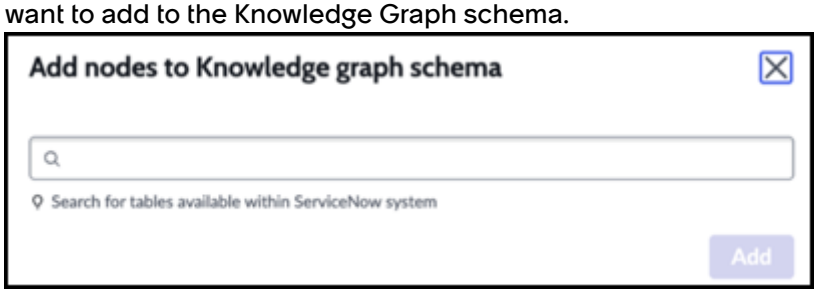
Add or delete a node from an existing Knowledge Graph schema.

Before you begin

Role required: kg_admin

Procedure

1. Navigate to **All** > **Knowledge Graph** > **Knowledge Graph Designer**.
   The UI displays a list of all the Knowledge Graph schemas on the landing page.
2. Select a Knowledge Graph schema from the list.
   The Knowledge graph schema opens in the Knowledge Graph canvas page.
3. From the toolbar, select the **Add nodes** option.
   The Edit Knowledge Graph Schema details form is displayed.
4. In the Add nodes to Knowledge graph schema window, enter or search for the nodes that you want to add to the Knowledge Graph schema.



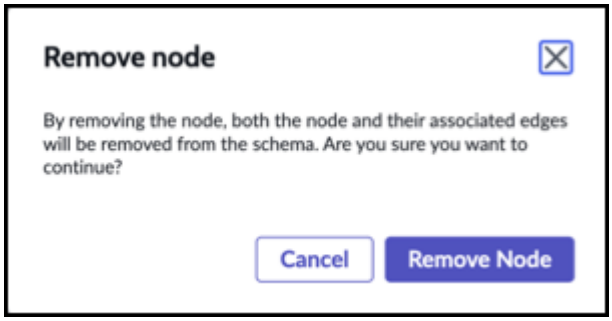
   You can search and select the Workflow Data Fabric tables, if integrated.

5. Select **Add**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The nodes are added to the Knowledge Graph schema and displayed on Knowledge Graph canvas.

6. Select the node to update the node details in the side navigation pane.
7. Select **Save** to save your changes.
8. Select **Remove node** to delete a node from the schema.



Add or delete edges to a Knowledge Graph schema

Add, edit, or delete edges that connect the nodes to customize a Knowledge Graph schema.

Before you begin

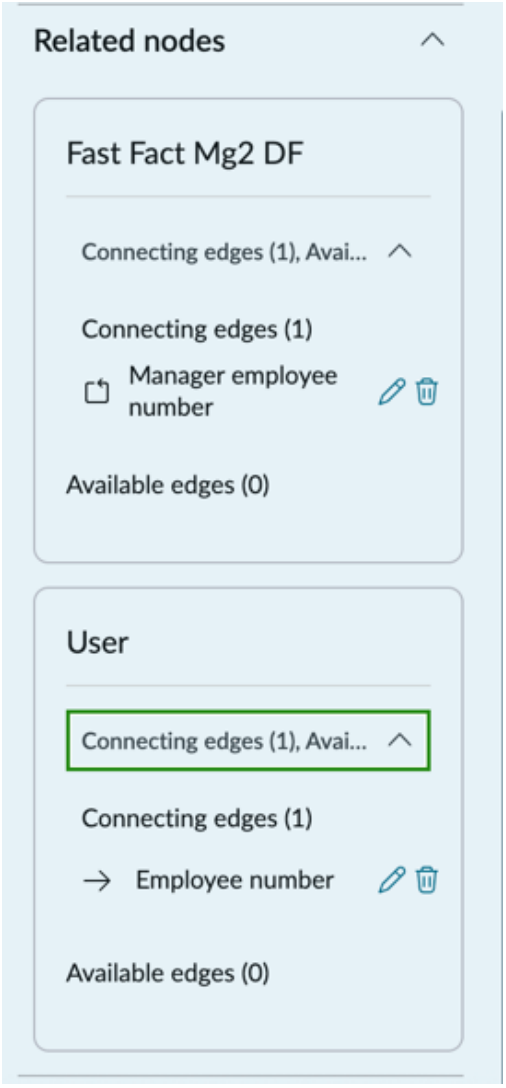
Role required: kg_admin

Procedure

1. Navigate to **All** > **Knowledge Graph** > **Knowledge Graph Designer**.
   The UI displays a list of all the Knowledge Graph schemas on the landing page.
2. Select a Knowledge Graph schema from the list.
   The Knowledge Graph schema opens in the canvas.
3. From the Knowledge Graph schema, select a node to edit.
4. In the side navigation pane, select **Related nodes**.

   Each related node displays a list of connecting edges and available edges.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



5. To establish a connection to an available edge, select the plus icon ( ) against the edge. The available edge is moved to **Connecting edges**. Select an edge to see its parent node and grandparent nodes in the hierarchy.
6. To reconfigure a connecting edge, select the edit icon (and update the following:
   - Edge type
   - Edge description
7. To delete an edge and move it to available edges, select the remove edge (
8. To delete the edge, select **Remove edge** when prompted.
9. To save your changes, select **Save** or to undo all your changes, select **Discard**.

Test a Knowledge Graph schema

Promote functionality by entering a query and testing the Knowledge Graph schema.

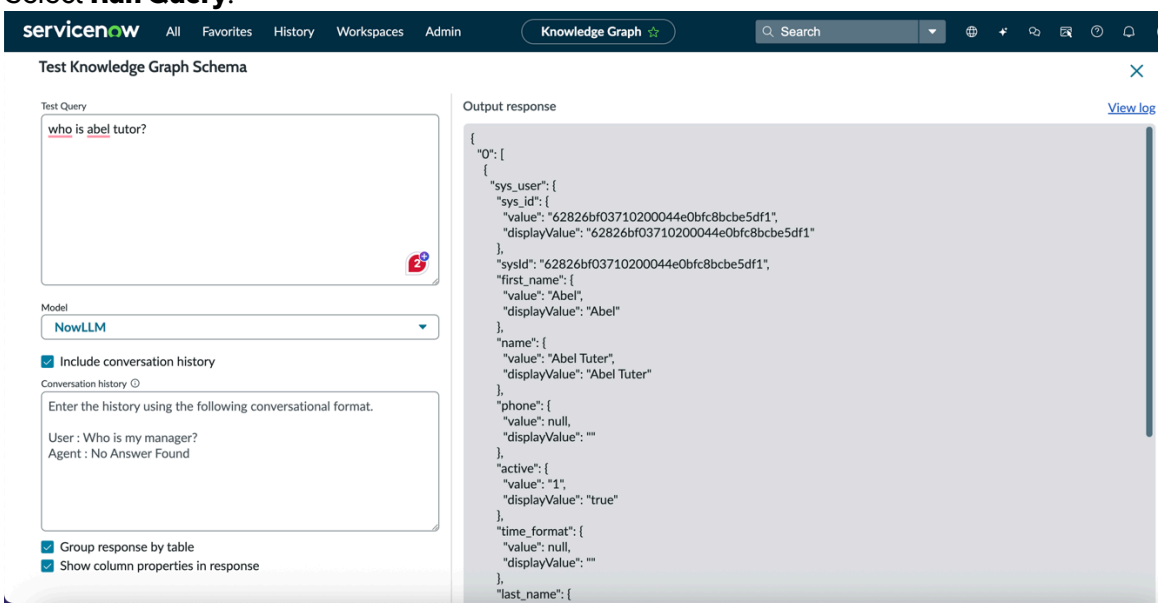
Before you begin

Role required: kg_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Knowledge Graph** > **Knowledge Graph Designer**.
   The UI displays a list of all the Knowledge Graph schema on the landing page.
2. From the list of Knowledge Graph schemas, select the Knowledge Graph that you want to test.
   The Knowledge Graph schema opens in the canvas.
3. From the toolbar, select **Test**.
4. In the Query section of the Test Knowledge Graph Schema window, enter your question.
5. Select an LLM model from the following:
   - **NowLLM**
   - **Claude**
   - **Gemini**
   - **GPT**
6. **Optional:** Select **Include conversation history** check box to add the any previous chats.
7. **Optional:** Select **Group response by table** to sort the results.
8. **Optional:** Select **Show column properties in response** to view the associated properties.
9. Select **Run Query**.



Analyze Knowledge Graph logs for debugging

Review Knowledge Graph logs and history to analyze performance and diagnose issues.

Before you begin

Ensure that you do not make any changes to the production instance.

Role required: admin

Procedure

1. Navigate to **All** > **sys_properties.list** > **sn_kg.log.level** or **All** > **System Properties** > **All Properties** > **sn_kg.log.level**.
2. Open *sn_kg.log.level* and in the value field add `debug`.
   The default value is **Err** but you can use any of the following values:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- emerg
- alert
- crit
- err
- warning
- notice
- info
- debug

3. Select **Update**.
   The logs will now be added to *syslog*.

Create a copy of a Knowledge Graph schema

Create a copy of a Knowledge Graph schema and duplicate it for further customization.

Before you begin

Verify that you have not selected a ServiceNow prebuilt Knowledge Graph schema that is non-editable.

Verify that the scope selected when creating a Knowledge Graph schema is the same as the scope of the existing Knowledge Graph schema.

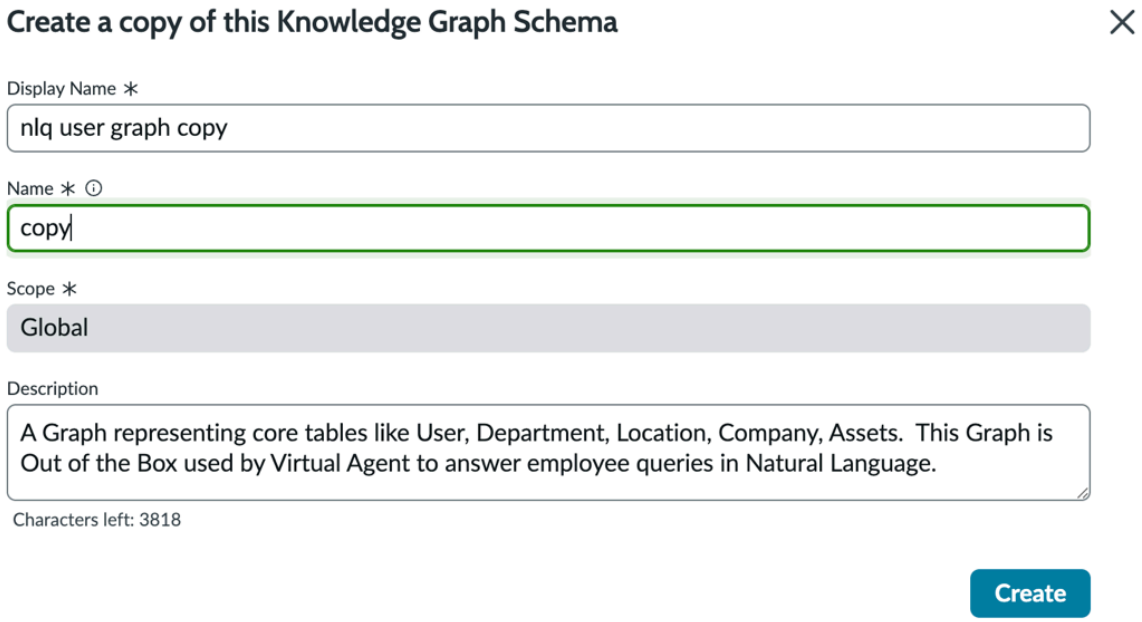
Role required: kg_admin

Procedure

1. Navigate to **All** > **Knowledge Graph** > **Knowledge Graph Designer**.
   The UI displays a list of all the Knowledge Graph schemas on the landing page.
2. Select a Knowledge Graph schema.
3. Select the more icon ( ).
4. Select **Copy Graph**.
   Verify that the scope of the existing Knowledge Graph schema is the same as the new schema.
5. On the form, fill in the fields.

Field	Value
Display Name	Existing name of the Knowledge Graph schema.
Name	Name of the cloned Knowledge Graph schema.
Scope	Scope under which you want to create the Knowledge Graph schema.
Description	Knowledge Graph schema overview for users.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



6. Select **Create**.

Using Enterprise graph schema

Use Enterprise Graph for accurate natural language query responses, across the entire database.

Enterprise graph is a pre-configured Knowledge Graph schema that maps all instance tables and their connections, enabling natural language queries for data across all tables.

To see some examples of the Natural Language query responses, refer Natural language queries use cases and examples.

Enterprise graph schema simplifies Knowledge Graph setup by providing a preconfigured schema, eliminating the need for custom schema creation in Knowledge Graph designer.

Admins can choose Enterprise graph as the Knowledge Graph schema when using AI agents, Now Assist Virtual Agent, or Now Assist panel admin and add tags to enhance accuracy.

By mapping all tables, the Enterprise graph schema expands query capabilities to cover the entire database, whereas a custom or out-of-the-box (OOTB) schema limits queries to only the tables included in its specific schema.

Benefits of Enterprise Graph on Prebuilt integration

Now Assist Virtual Agent- With Enterprise Graph integration, Now Assist Virtual Agent can respond to a wide range of questions about their enterprise data from requester, helping to deflect queries.

Now Assist panel- With Enterprise Graph enabled, Now Assist panel allows fulfillers to obtain insights from the structured data in instance by asking natural language questions, which boosts their productivity.

AI agents- With Enterprise Graph enabled, AI agents can access relevant information from structured data in ServiceNow instance directly from the Knowledge Graph, giving them essential context for tasks.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Modes of Enterprise Graph

Enterprise Graph has two modes:

- Enterprise Graph- regular mode
- Enterprise Graph (Small)

Regular mode is used for use cases with large number of tables (more than 50 tables), to:

- Searches across all tables, but gives more priority to tables which are included in the tag added at time of configuring Enterprise Graph in the consuming app (AI agent or VA).
- Useful for scenarios where answer is expected from large number of tables.

Enterprise Graph (Small) is a version of Enterprise Graph that limits the search to only the tables tagged with specific labels. This means it can only answer queries related to those tables. It's useful for:

- Searches only within tables which are included in the tag added at time of configuring Enterprise Graph (Small) in consuming app (AI agent or VA). Tagging is mandatory in this mode.
- Useful for latency-sensitive use cases with limit of 50 tables

Tags in Knowledge Graph

Tags are lists of key tables that are important for answering natural language questions in a specific use case. They are required to be used with Enterprise Graph and Enterprise Graph (Small).

Tagging in Knowledge Graph Designer

Use Knowledge graph tags to mark the key tables, that are important for answering natural language questions for specific use cases.

Use Knowledge Graph Designer to create and edit tags from the **Enterprise Graph Tags** section.

Tagging in different Enterprise Graph modes

Name	Description
Enterprise Graph- regular mode	Provides hint to Enterprise Graph on which tables to prioritize when retrieving information, to improve the accuracy of results. It is recommended to only tag up to 30 tables for regular mode as hints to prioritize while searching.
Enterprise Graph (Small)	Provides a list of tables to run the search operation. Enterprise Graph (Small) will only search on the tables provided by tags. It is mandatory to add a tag when Enterprise Graph (Small) is used with any of the prebuilt integration with AI agent, Virtual Agent or Now Assist Panel. You can tag a maximum of 50 tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tagging in different Enterprise Graph modes (continued)

Name	Description
	Note that to enable querying on columns of table which are references to another table, also include referenced table in tags.

Creating tags

Tags are created by kg_admins using the Knowledge Graph Tags section and has to be linked to schema when selecting the Enterprise Graph or Enterprise Graph (Small) as the Knowledge Graph schema for your use case in Virtual Agent, Now Assist Panel or AI agent.

Knowledge graph ships the following prebuilt tags:

1. VIRTUAL AGENT DEFAULT TAG: A prebuilt tag listing common to core, ITSM, and HR tables for integrating Enterprise Graph with the Virtual Agent on the employee portal. Admins can edit this tag to include additional tables relevant for answering Virtual Agent queries.
2. NOW ASSIST PANEL DEFAULT TAG: A prebuilt tag listing common to core, ITSM and HR tables which can be leveraged as general tag on Now Assist panel. Admins can edit this tag to add additional important table for answering questions on Now Assist panel.

Follow these steps to create or edit tags:

- Go to the Knowledge Graph Designer to view the existing tags and their associated tables. You can use an existing tag, edit one to fit your needs, or create a new tag for your particular use case.
- Navigate to the admin center for consuming apps.
  - For the Now Assist Virtual Agent or Now Assist Panel, go to your virtual agent's setup, select Information Sources, and open the Knowledge Graph.
  - For AI agent, proceed to its setup and open the Knowledge Graph tool for the AI agent.
- Select Enterprise graph or Enterprise Graph (Small) as the KG schema and select the required tag from Tag drop down to link it to graph schema.

Note: In this release, the Enterprise Graph option is hidden by default for Now Assist Virtual Agent and Now Assist Panel, and only Enterprise Graph (Small) can be selected as schema option. To use Enterprise Graph, Admins can enable it as a schema option by setting the property `sn_kg.va.enable_global_graph_all_modes` to true. The role required is Admin.

Recommended tables for different use cases

- Now Assist Virtual Agent: For Now Assist Virtual Agent include the tables necessary to answer the expected user queries. For example if you want to deploy a new virtual agent on employee center portal, you should add a tag having tables relevant for answering employee questions.
- Now Assist Panel: Now Assist Panel operates as a unified assistant across multiple workspaces, hence it is essential to tag tables as follows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- An individual tag for each workspace that encompasses key tables relevant to that workspace. This ensures the Now Assist panel can deliver responses specific to queries related to the workspace tables when accessed within each workspace.
- Create a general tag that includes the most important tables required to answer users questions on Now Assist panel, allowing the Now Assist panel to provide relevant answers when used outside of any workspace.
- For more information on adding workspace tags in Now Assist Panel, see Add a Knowledge Graph schema to a chat assistant.

- AI agent: For an AI agent use case, include key tables relevant to the agent's purpose. For example, for asset manager agents User, Asset, and other crucial tables should be included in the tag. For more information, see Add a Knowledge Graph to an AI agent.

Create a new Knowledge Graph tag

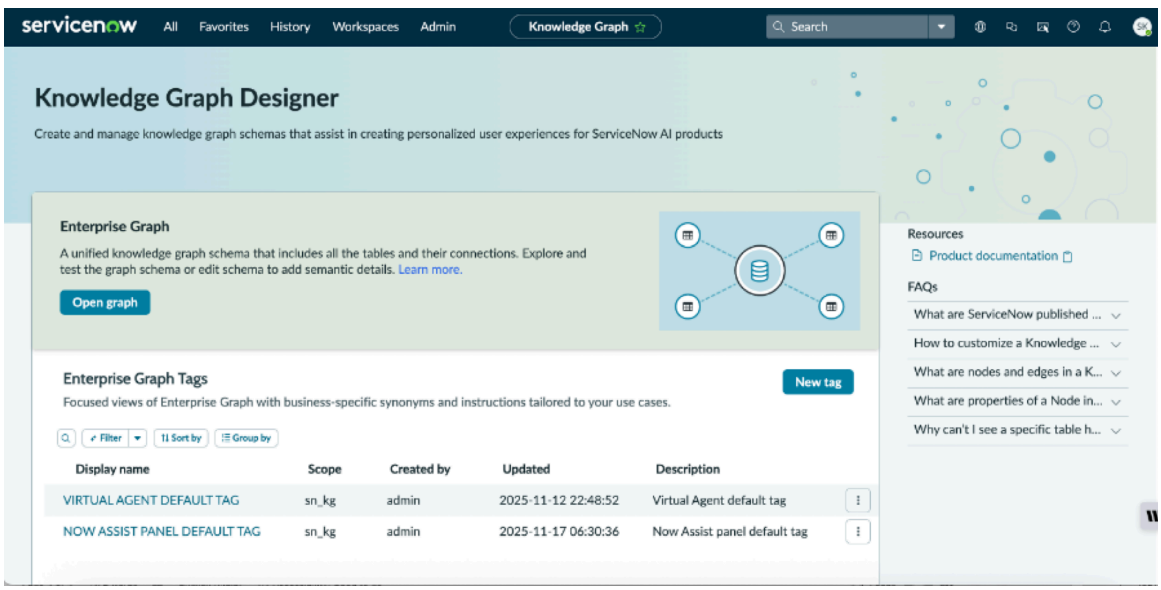
Create new Knowledge Graph tags for Now Assist Virtual Agent, AI agent or Now Assist Panel Enterprise Graph using Knowledge Graph Designer to improve accuracy of natural language queries.

Before you begin

Role required: admin

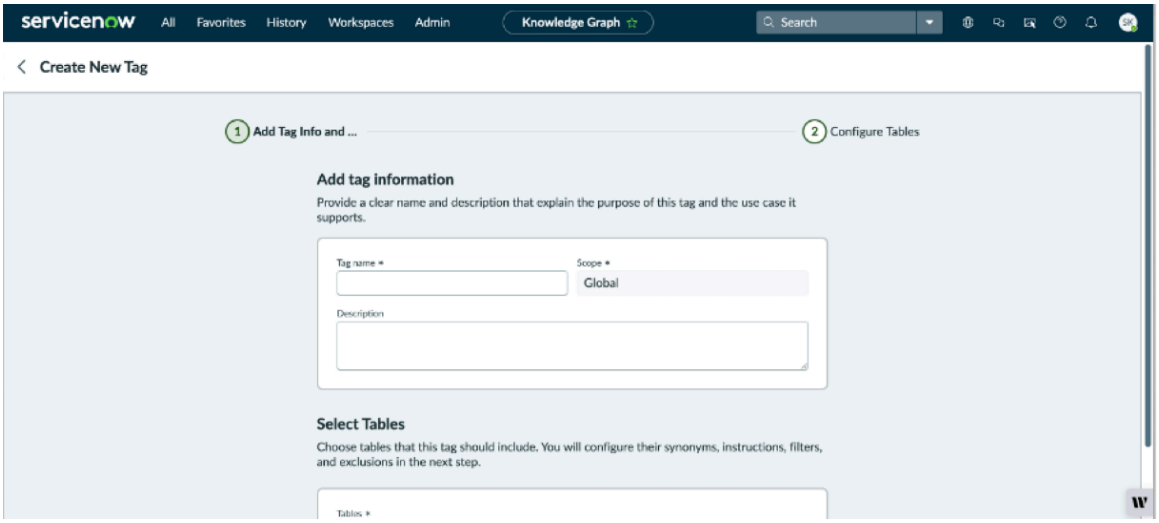
Procedure

1. Navigate to **All** > **Knowledge Graph Designer** > **Enterprise Graph Tags** > **New tag**.
   You are navigated to the Create New Tag record page.



2. On the Create New Tag record page, add the following information:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



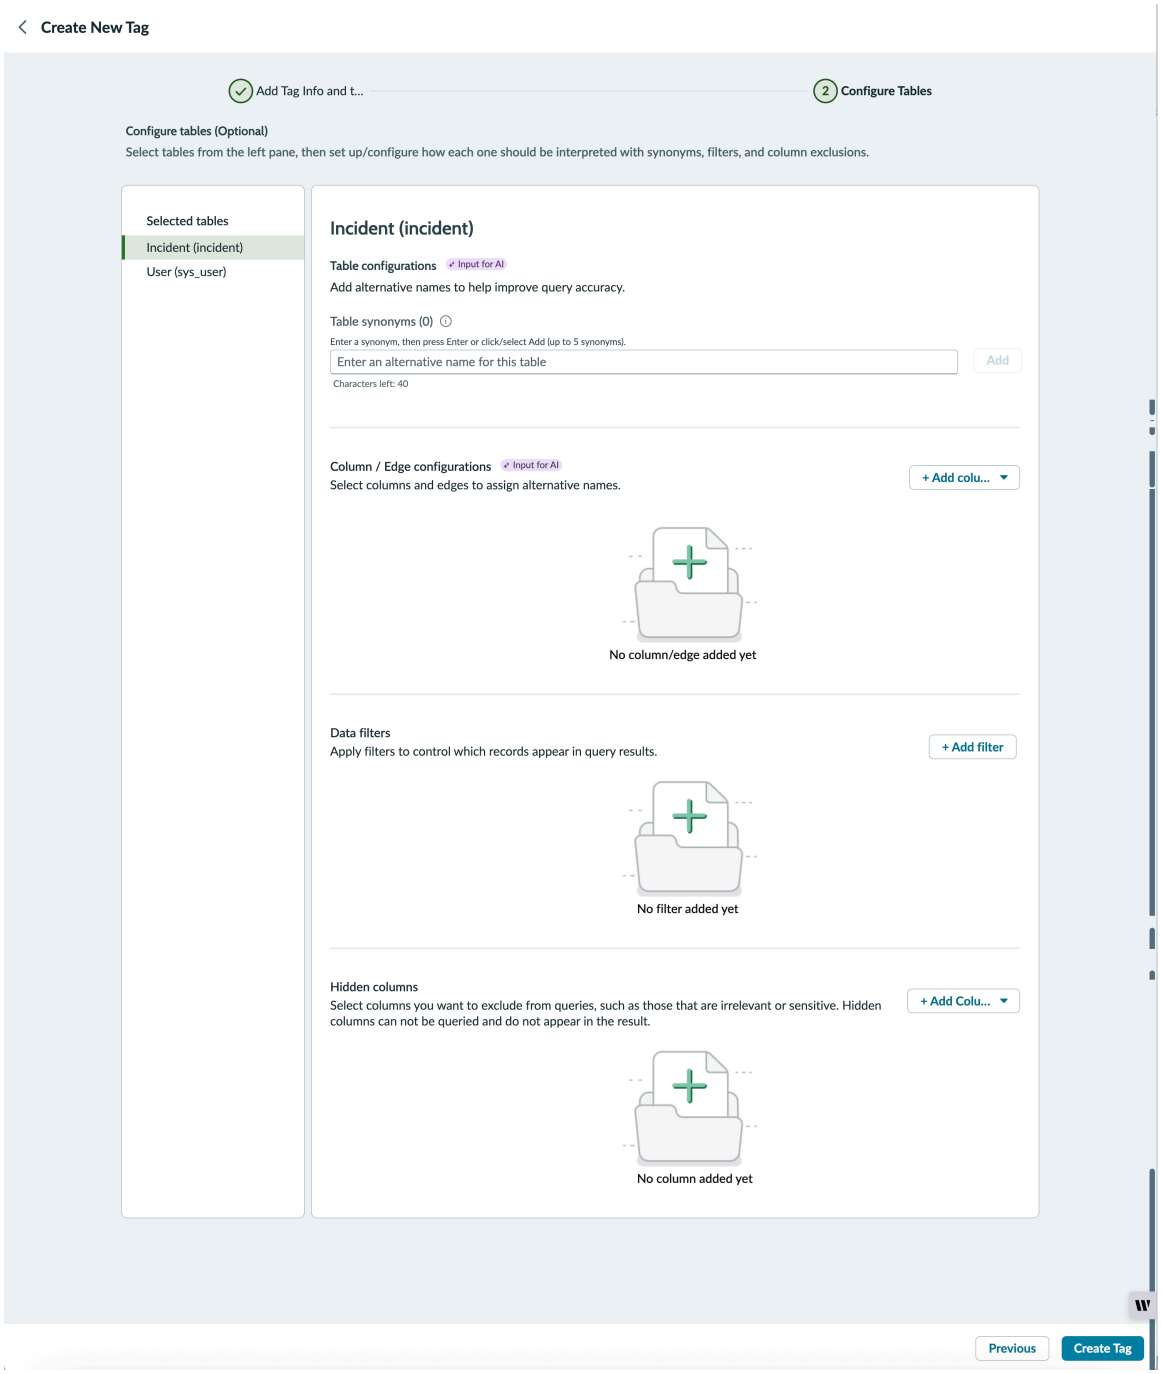
- Tag Name: Name for the tag.
- Description
- Scope: The default selection in Global. Read-only field.
- Select Tables: The tables or nodes that you want to add to the tag.

3. Select **Next** to go to Configure tables page.

   The details that are added to the Configure tables page will help to further improve the accuracy of the natural language queries.

   The selected tables will be displayed on the left pane of the Configure tables page.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. Select a table from the left pane to add tables configurations, column and edge configurations, data filters and hidden columns.
5. Add alternative names in the **Table synonyms** field and select **Add**.
   - Use the commonly used table names as synonyms to help AI recognize it in natural language queries.
   - Example: you can add Tickets as a synonym for the Incident table.
   - You can add up to five synonyms only.
6. To delete a synonym, select ⊗ icon next to the synonym.
7. Select **Add Column** and choose the columns and edges that you want to add synonyms to.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Add commonly used column and edge names as synonyms to help the AI match user terms to this field.
- Example: you can add Opened On as a synonym for Created Date.
- You can add up to five synonyms only.

Use the
The selected columns and edges are added to the **Column / Edge configurations** section.

8. To delete a column from the section, select  next to the column.
9. Select **Add filter** option to add filters for managing which records appear in the query results and add the following details:
   - Field
   - Operator
   - Value

   Example: To exclude inactive users from search results, apply `Active = True` filter.
10. Select **Apply filters to child tables** check box if you want to add the same filter conditions for the child table of the selected table in the tag.
    If enabled, these filters will also apply to child tables that inherit data from this table.
11. Select irrelevant or insensitive columns from the **Hidden column** drop-down to exclude them from queries.
    Hidden columns can not be queried and do not appear in the result.

    For example, add `Employee number` as hidden columns for `sys_user` table to hide the employee number from query results.
12. Select **Create tag** and repeat the same process for each of the selected columns.

Managing Knowledge Graph tags

Edit or delete Knowledge Graph tags for Now Assist Virtual Agent, AI agent or Now Assist Panel Enterprise Graph use case.

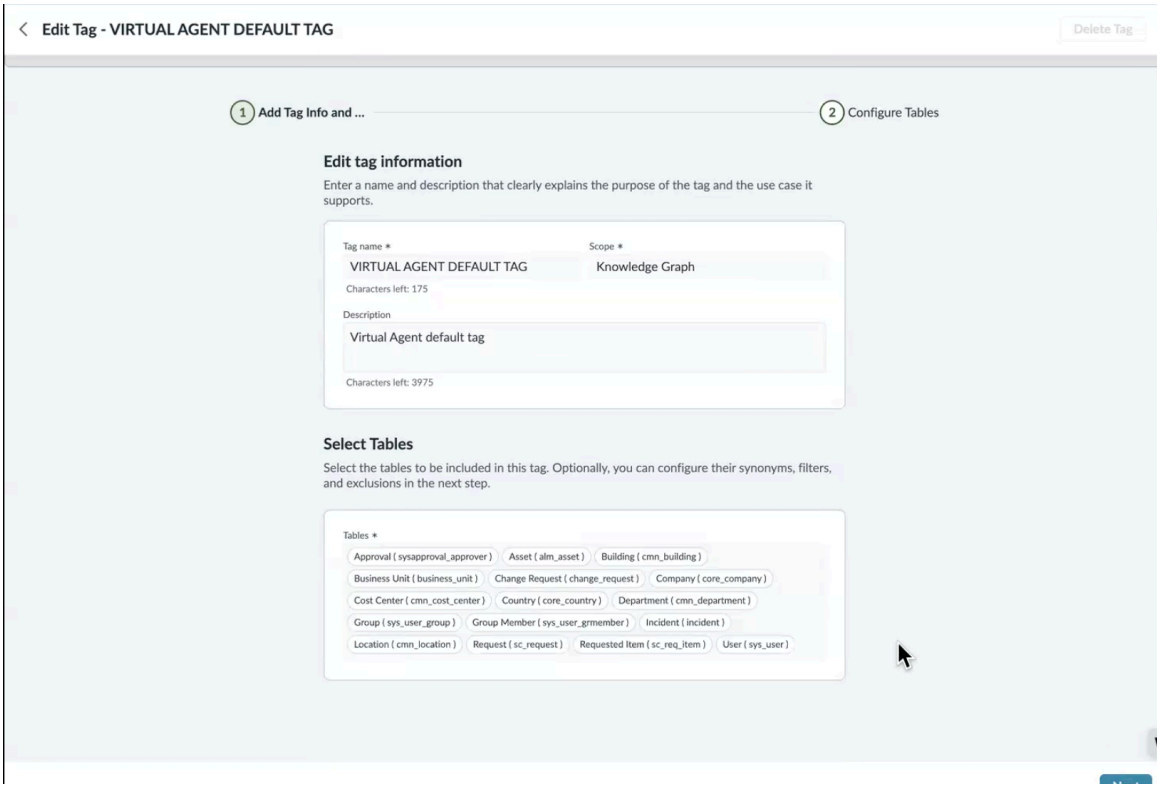
Before you begin

Role required: admin

Procedure

1. To edit an existing tag, go to **All** > **Knowledge Graph Designer** > **Enterprise Graph Tags**.
2. Select a tag to open it in Knowledge Graph canvas.
   You can see all the tables that are part of the tag on the Knowledge Graph canvas.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



3. Select **Edit tag** option to open the Edit Tags window or select a node to edit.
   If you select a node you will see the Tag configuration section on the right pane.
4. You can also select the Edit tag configuration ( ) icon from the **Tag configuration** section to open the Edit Tags window.
5. To delete a tag, select the ⋮ icon next to a tag on Enterprise Graph tags section on home page.
6. Select **Delete** from the drop-down.
   You can also select the edit option here, to go to the Edit tags window.
7. When prompted, confirm the deletion process.

Setting up Enterprise graph in sub-production instance

Use these steps to setup Enterprise graph in sub-production instance manually.

Before you begin

Unlike production instances, the Enterprise Graph setup does not start automatically in the sub-production instance.

You have to manually enable the system property:
`sn_kg.description_generation.enable.non_production_envs`

When you enable this system property, the Enterprise Graph setup starts in the sub-production instances and can take several days to complete.

Until this setup is finished, the Enterprise Graph schema won't work effectively for answering queries.

This step creates the necessary descriptions for the Enterprise Graph to respond to queries.

Role required: kg_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Initial setup for Enterprise Graph schema in production instance

Setup and use Enterprise Graph Schema, a unified knowledge graph schema, that captures all the ServiceNow and third party tables and their connections.

Before you begin

When you install Enterprise Graph, it automatically begins configuring instance tables. This may take several days and until the setup is complete, the schema won't respond effectively to queries.

If you have already completed the setup in sub-production instance, you can import the generated descriptions to production instance and fasten the process.

Assuming that you have completed the setup in sub-production instance, do the following to complete the initial setup in production instance:

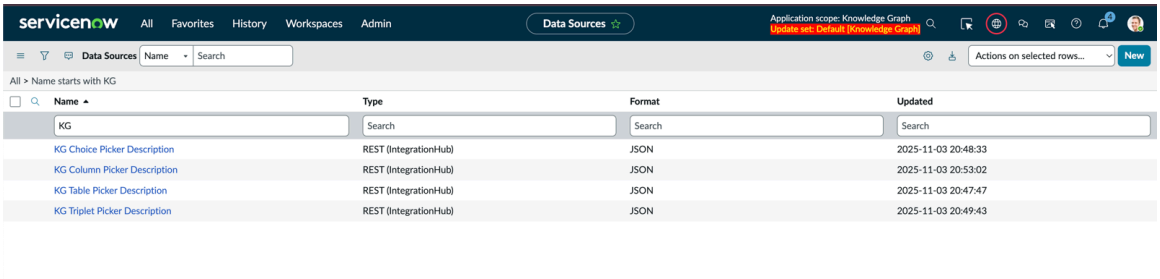
- Create an integration user, in the sub-production instance. For more information, see https://www.servicenow.com/docs/bundle/zurich-platform-administration/page/administer/users-and-groups/task/t_CreateAUser.html
- Verify if the descriptions are generated in sub-production instance.
- Use the following procedure to:
  - Import the data into production.
  - Run the associated Transform Maps.
- Remove tracker records from the production instance.

Role required: admin

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connection & Credential Aliases**.
2. Open the record with name **Description_Connector**.
3. Select **Create New Connection & Credential** from the related links section.
4. Enter the required details in the **Create Connection & Credential** form.
5. Add the user name and password that was used, while creating the new user.
   See https://www.servicenow.com/docs/bundle/zurich-platform-administration/page/administer/users-and-groups/task/t_CreateAUser.html for more information.
6. Navigate to **All** > **System Import Sets** > **Administration** > **Data Sources**.
7. In the name field search bar, enter **KG** to view the related tables.
   The following tables are displayed:
   - KG Choice Picker Description.
   - KG Column Picker Description.
   - KG Table Picker Description.
   - KG Triplet Picker Description.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



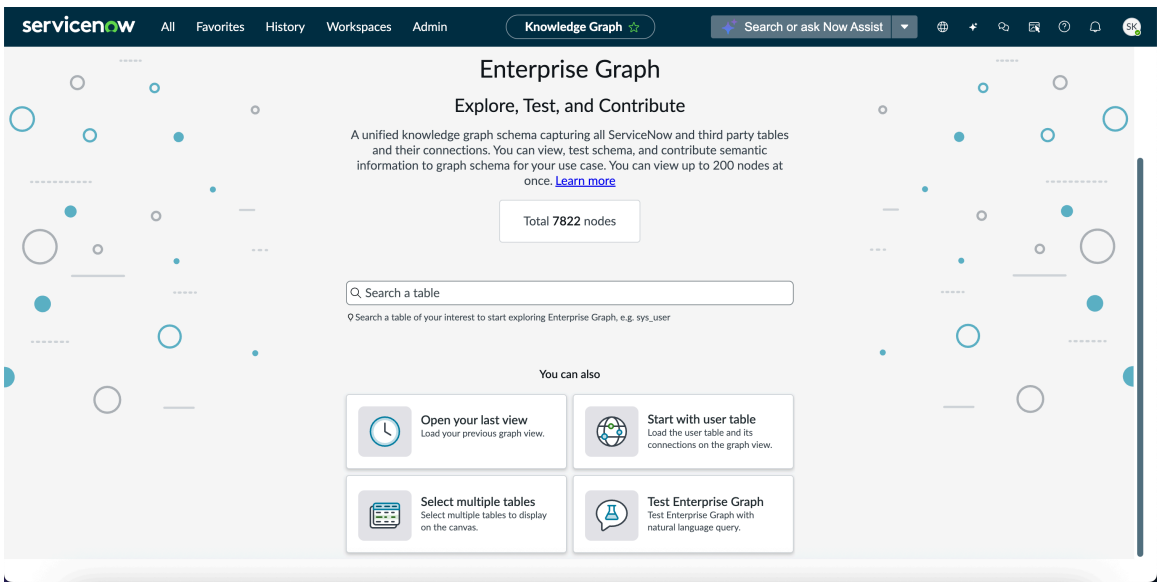
8. Open one of the displayed table and select **Load All Records** from the Related links section. Ensure the data is loaded.
9. Navigate to **System Import Sets** > **Advanced** > **Import Sets**, once the data is loaded. For detailed information, see import set documentation https://www.servicenow.com/docs/bundle/yokohama-integrate-applications/page/administer/import-sets/reference/import-sets-landing-page.html
10. Open the import set and select **Transform** from the Related links section. For detailed information, see Transform map documentation https://www.servicenow.com/docs/bundle/yokohama-integrate-applications/page/script/server-scripting/concept/c_CreatingNewTransformMaps.html
11. Select **Transform** to complete the transformation.
12. Navigate to the selected table and verify if the data is loaded.
13. Repeat the import set and transformation steps for other three tables.

Access Enterprise Graph schema

View and manage Enterprise Graph schema, select tables and test the schema with Knowledge Graph Designer.

You can view and manage the Enterprise Graph schema, go to **All** > **Knowledge Graph Designer** > **Open Enterprise Graph**.

Use Enterprise Graph to view and test the schema. You can also contribute synonyms to the schema for your use case. Synonyms will help improve the accuracy of Enterprise Graph for answering queries.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

On the Enterprise Graph start page, use the search bar to search and select a table you want to view or manage. Alternately, you can use the **Open your last view** or **Start with user tables** to view the user and the related table.

The Enterprise Graph start page displays the total number of nodes. You can use the **Select multiple tables** option to select tables that you want to view in the schema.

Within **Select multiple tables** option, you can select an application, tag or individual tables. Once you select the node, you will see selected nodes and their associated tables on the schema.

You can also use the **Test Enterprise Graph** option to run a query on the Enterprise Graph schema.

View and manage Enterprise Graph schema

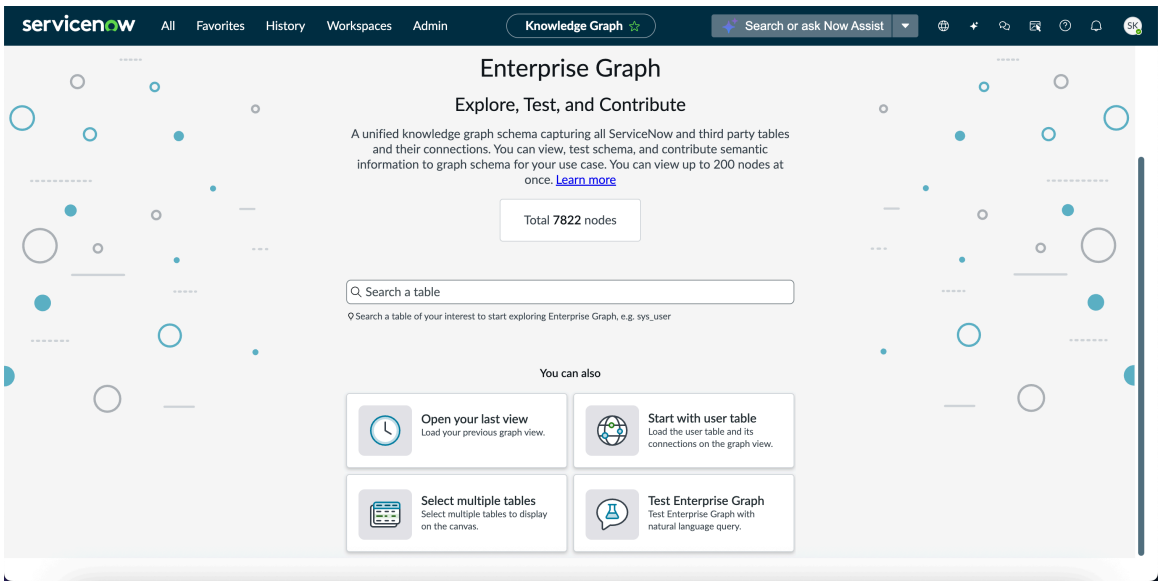
Use Knowledge Graph Designer to view and manage Enterprise Graph schema

Before you begin

Role required: admin

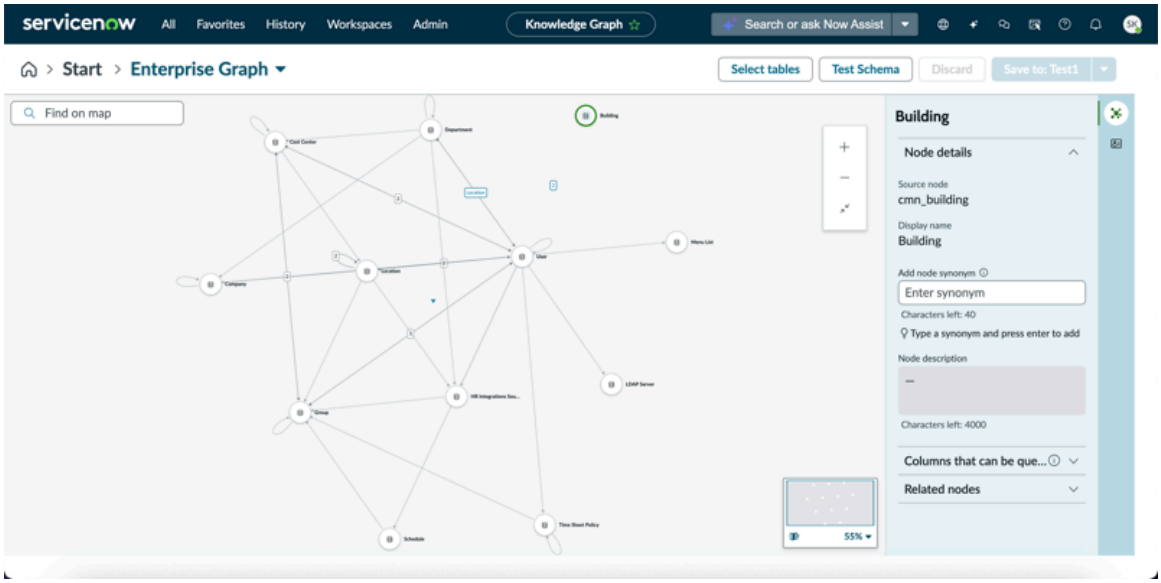
Procedure

1. Navigate to **All** > **Knowledge Graph** > **Knowledge Graph Designer** > **Enterprise Graph**.
   You are navigated to the Enterprise Graph start page.



2. Do one of the following to go to the Enterprise Graph canvas:
   - Search and select a table from the search bar.
   - Select **Open your last view**.
   - Select **Start with user tables**.
   - Select**Select multiple tables**.

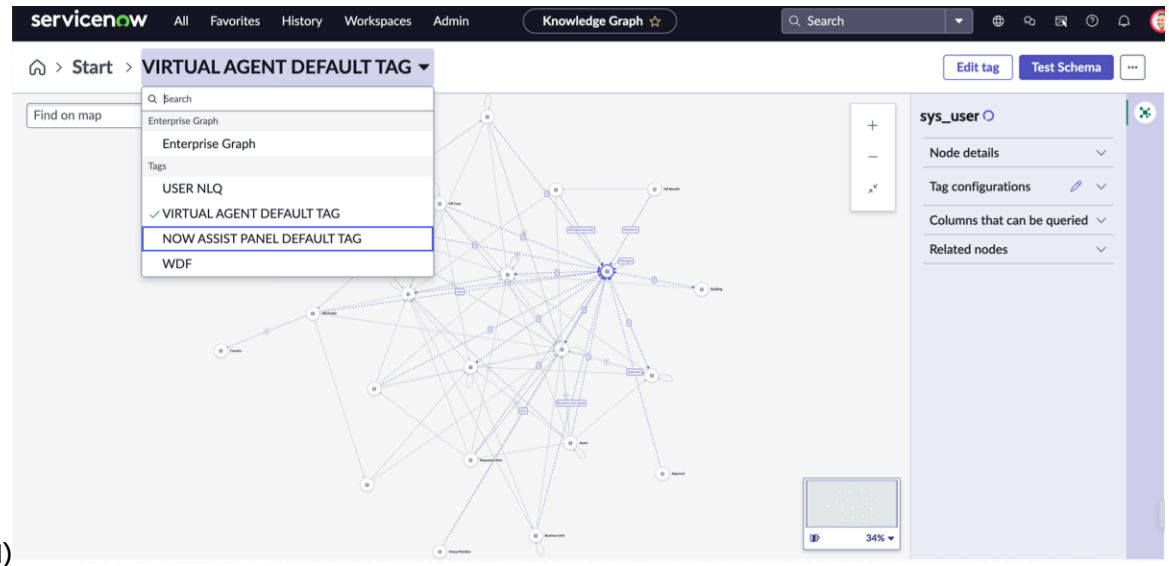
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



The Enterprise graph schema opens in the Knowledge Graph canvas page.

3. **Optional:** You can select a tag from the drop-down next to Enterprise graph to see a specific tag.

(Optional)



4. From the toolbar, select **Nodes details** option to view details or add a synonym for the node. The other fields are read-only and cannot be edited. In the tag view of Enterprise Graph, you can also see tag configuration when you click on a node.
5. Use the icon in the **Tag configurations** section to go to the Edit tag configurations page and update the following details:
   - Node synonyms
   - Column with synonyms
   - Data filters
   - Hidden columns
6. In the **Columns that can be queried** section, search and view the columns of the tables.
7. From the toolbar, select **Related nodes** to view the related nodes.
8. Select **Save** to save the synonyms that you have added.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can save to the most recent contribution or choose a another contribution from the drop-down, to save your changes.

9. Select **Discard** to delete your changes.

Select multiple tables

Choose one or multiple nodes to view in the Enterprise graph canvas.

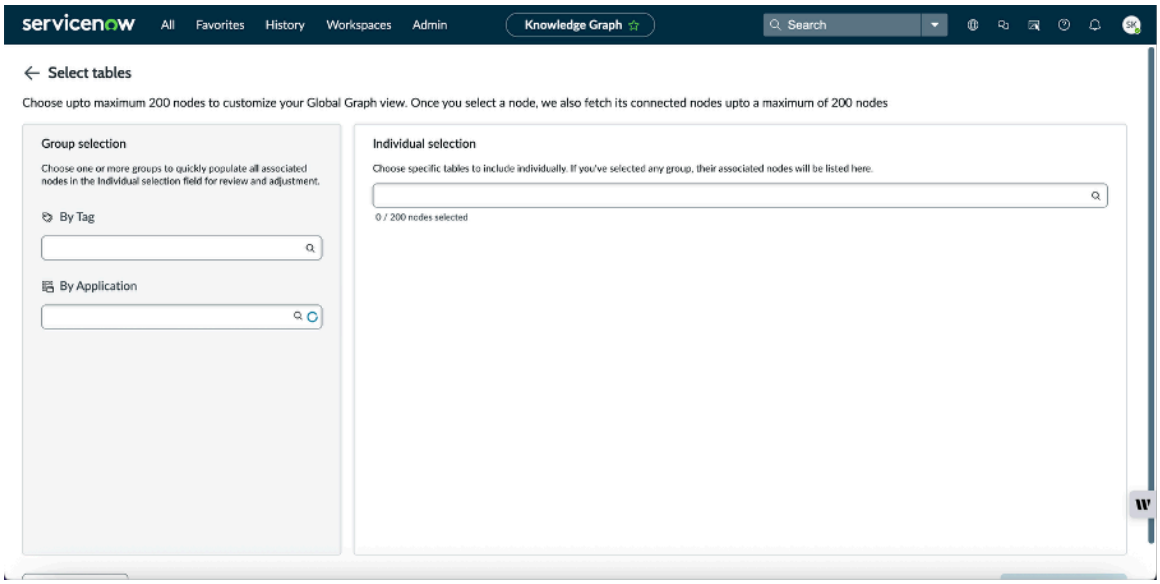
Before you begin

You can select an entire group or individual tables to view in Enterprise Graph canvas. All the connected nodes or associated tables also gets populated, once you select the node. However, you can only choose up to 200 nodes at once.

Role required: admin

Procedure

1. Navigate to **All** > **Knowledge Graph Designer** > **Open Enterprise Graph** > **Select multiple tables**.
   Choose up to 200 nodes to customize your view. Once you select a node, all the related edges will also be selected.



2. Select a group using the **By Tag** or **By Application** search bar.
   You can search and choose multiple groups and all the associated nodes. The selection will be listed in the Individual selection section.
3. Select an individual node using the **Individual Selection** search bar.
4. Select **Apply** to view these nodes and its related tables on Enterprise Graph canvas.

Test an Enterprise Graph schema

Promote functionality by entering a query and testing the Enterprise Graph schema.

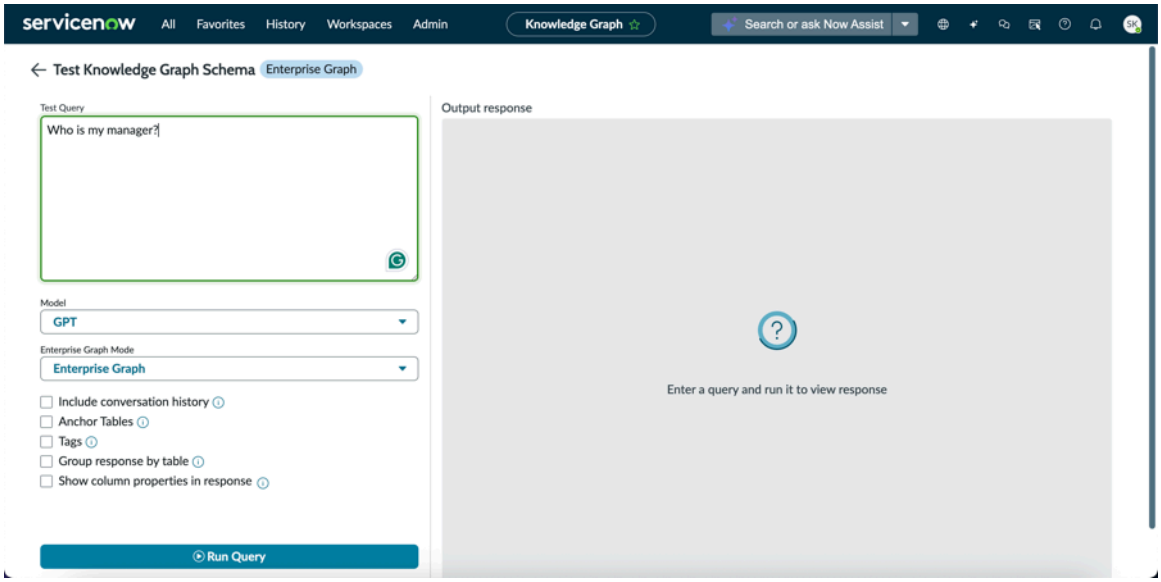
Before you begin

Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

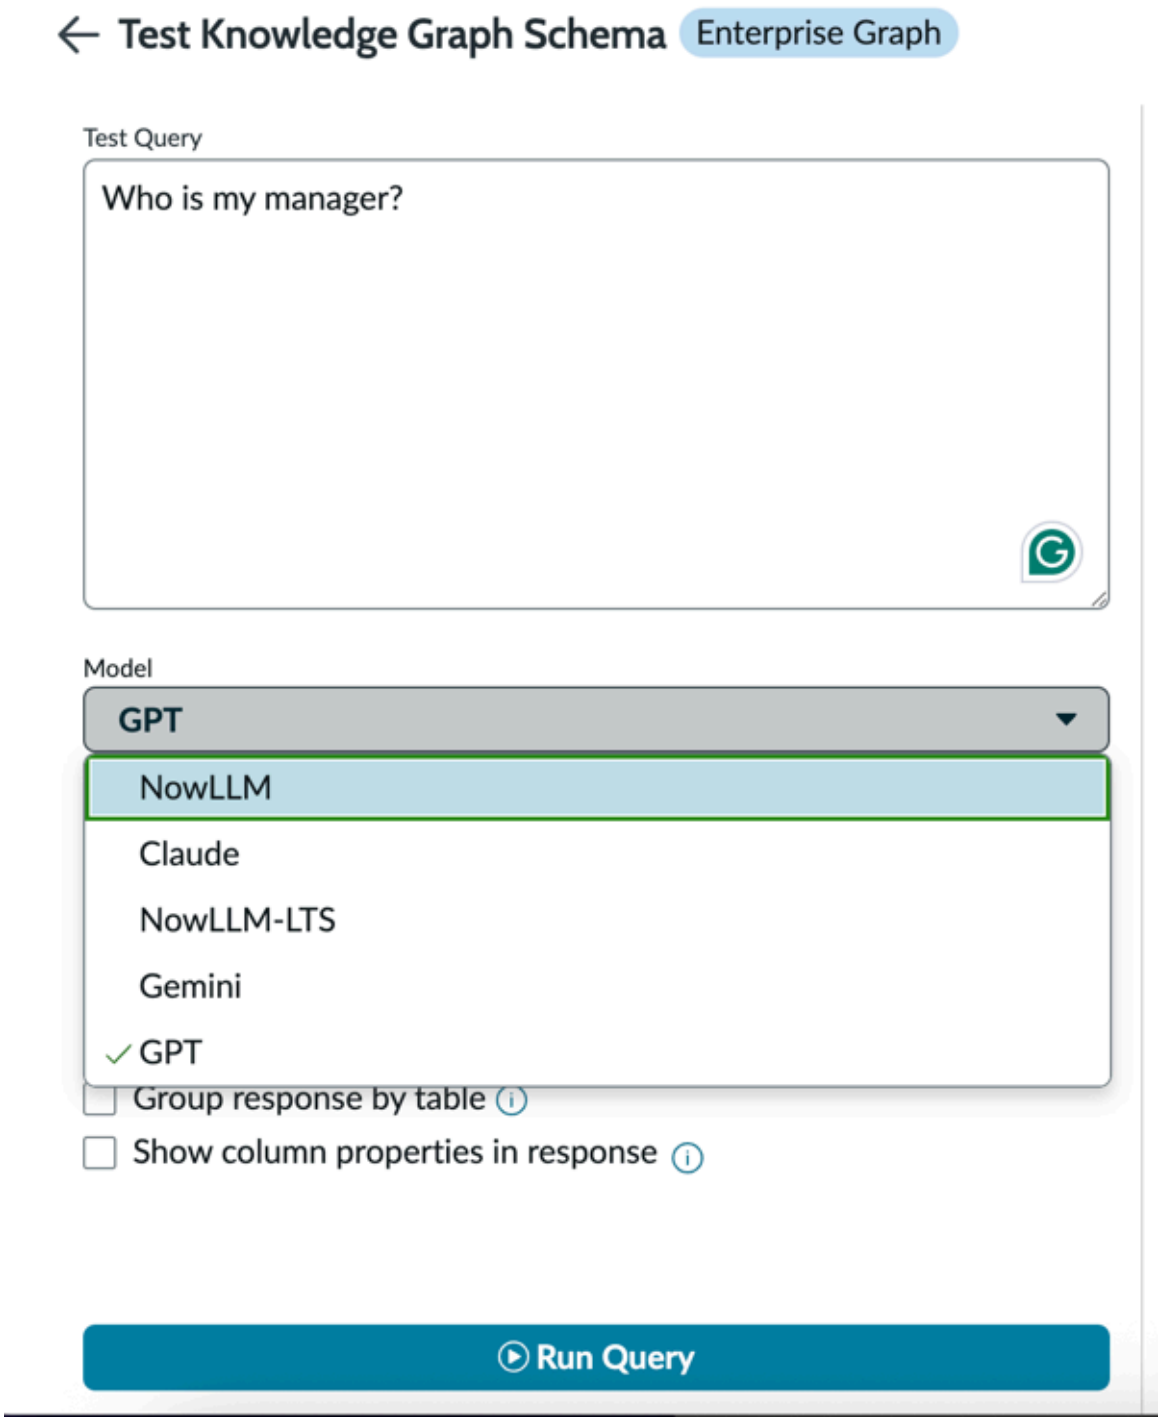
Procedure

1. Navigate to **All** > **Knowledge Graph** > **Knowledge Graph Designer** > **Enterprise Graph**.
2. Select **Open your last view** to load and view the recently opened Enterprise Graph schema.
   The Enterprise graph schema opens in the Knowledge Graph canvas page.
3. From the toolbar, select **Test**.
   You are navigated to Test Knowledge Graph Schema page.



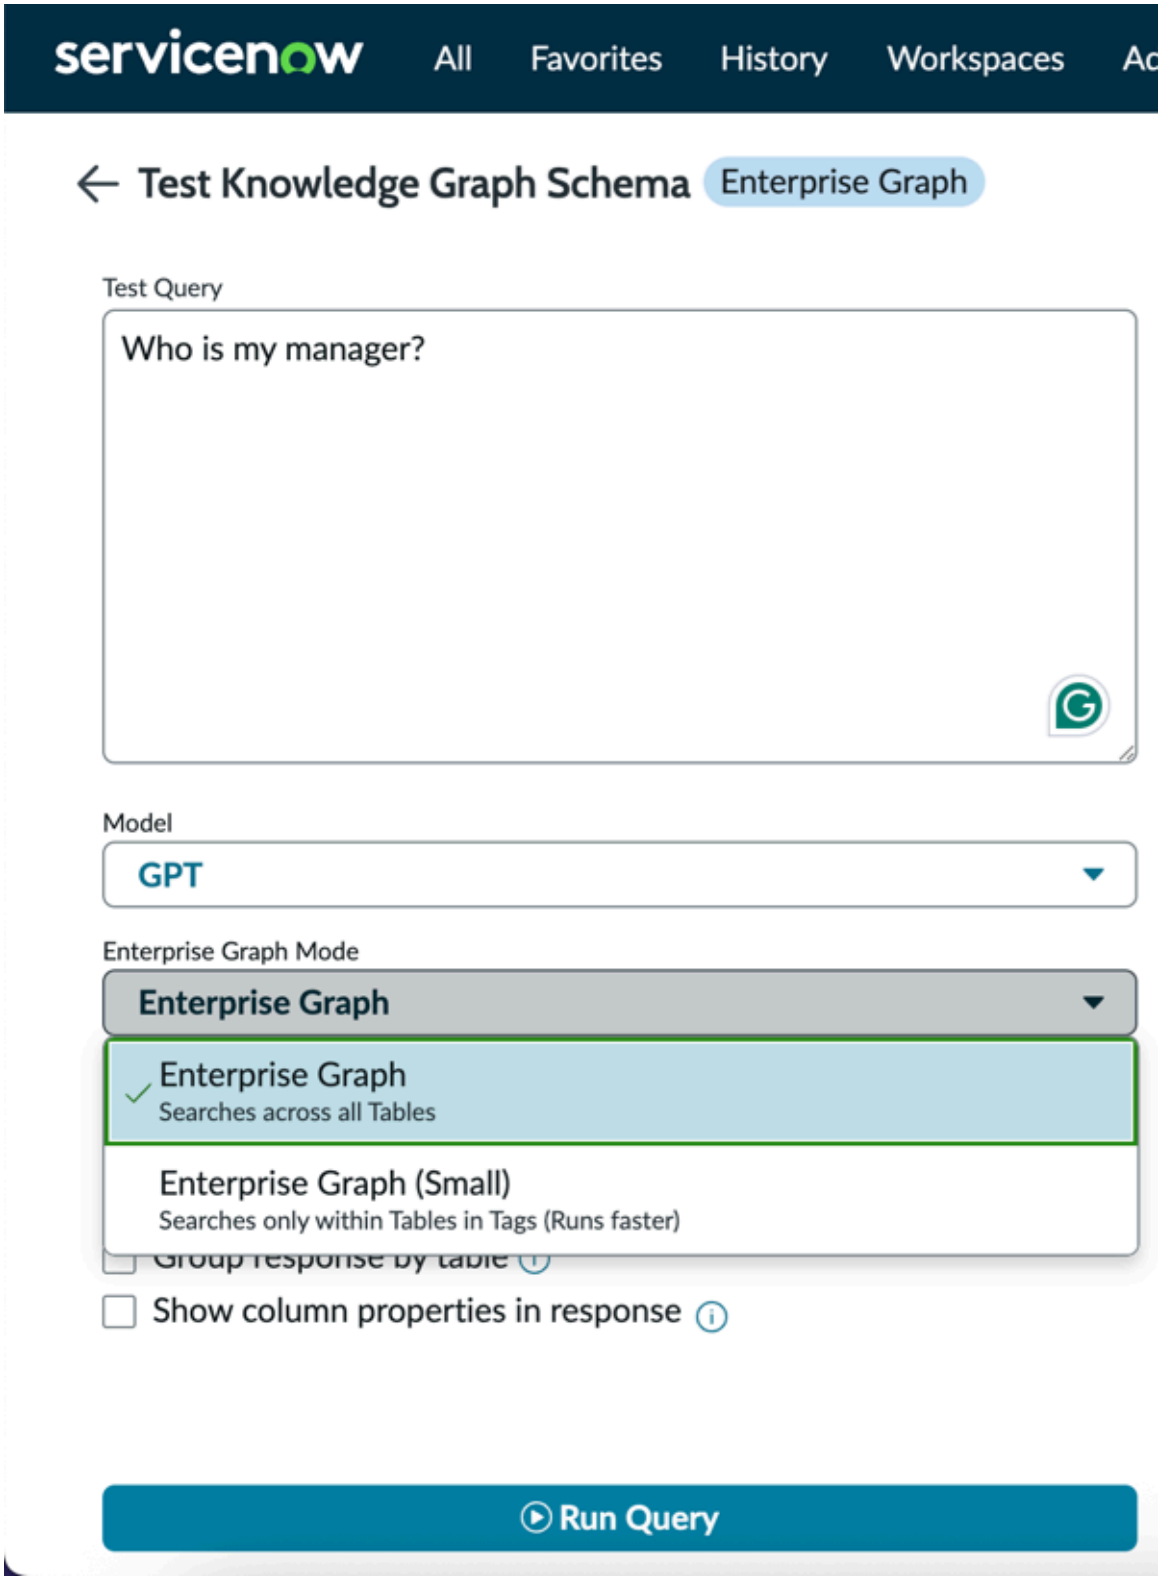
4. In the Query section of the Test Knowledge Graph Schema window, enter your question.
5. Select an LLM model from the following:
   - **NowLLM**
   - **Claude**
   - **NowLLM-LTS**
   - **Gemini**
   - **GPT**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



6. Select an Enterprise Graph mode.
   The options are:
   - Enterprise Graph
   - Enterprise Graph (Small)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



7. **Optional:** Select **Include conversation history** check box to add the any previous chats.
8. **Optional:** Select **Anchor Tables** check box to add the Anchor tables.
9. **Optional:** Select **Tags** check box to add the associated tags.
   In Enterprise Graph (Small) mode, tags are mandatory.
10. **Optional:** Select **Group response by table** to sort the results.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

11. **Optional:** Select **Show column properties in response** to view the column properties along with associated data.
12. Select **Run Query**.

Reference for Knowledge Graph

Knowledge Graph uses the following terminologies.

Terminologies

Term	Description
Knowledge Graph Schema	Structured representation of real-world entities and their relationships, stored in a graph database.
Knowledge Graph Canvas	The editor used to create and edit Knowledge Graph Schemas. You can also use it to add or delete nodes and edges.
Node	An entity or object that represents a data point such as a person, place, product, or concept.
Data Source	The source table that is linked to a node.
Node Properties	The attributes or columns of the source table linked to a node.
Connected Node	The nodes that have a connection with the selected node.
Edge	Represent the relationships or interactions between nodes.
Connecting Edges	The edge that connects two nodes.
Available Edges	Available connection between selected nodes.
Start Node	The start data point for creating an edge or relationship.
End Node	The end data point for creating an edge or relationship.

Knowledge Graph roles

Certain roles are required to use Knowledge Graph functionality.

Knowledge Graph Admin [kg_admin]

This user can create and update Knowledge Graph schema using Knowledge Graph designer.

Contains Roles

List of roles contained within the role.

Knowledge Graph Admin [kg_admin].

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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.

MCP Server Console

MCP Server Console enables secure and governed access to functionality on a ServiceNow instance for AI applications with Model Context Protocol (MCP) servers. MCP servers extend ServiceNow AI Platform® functionality into any external MCP client and employee experience over the Model Context Protocol.

https://player.vimeo.com/video/1154347099?h=fd6c7893de&badge=0&autopause=0&player_id=0&app_id=58479

In this video, learn about the Model Context Protocol and how to create and use MCP servers with MCP Server Console.

Get started

Explore  Learn about MCP Server Console.	Configure  Create an MCP server and configure its tools.
Connect  Call an MCP server from an MCP client.	Reference  Get details about components installed with MCP Server Console.

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.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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.

For more information, see the Now Assist documentation.

Exploring MCP Server Console

Learn about how you can use Model Context Protocol (MCP) servers to allow AI applications to access data and perform actions on a ServiceNow instance.

MCP Server Console overview

The Model Context Protocol defines a standard method of communication between large language models (LLMs) and external systems, such as a ServiceNow instance. AI applications connect to an external system from an MCP client, such as an AI agent, using an MCP server. The server tells the client which external resources it can access and how to access them. Then, from an MCP client, users can request information from the server and automate functionality using the available tools and data that the server returns. For general information about the Model Context Protocol and MCP terminology, see the Model Context Protocol documentation on the modelcontextprotocol.io website.

With MCP Server Console, you can create MCP servers on an instance and configure the tools that they expose or use a preconfigured Quickstart Server. Tools define which functionality and data a server exposes to clients and the actions that can be performed on an instance by clients. You can create tools based on existing capabilities, such as Now Assist skills.

Then any AI application, such as Microsoft Copilot or Claude, can call a server from a client using OAuth 2.0 authentication and get a list of available tools. With these tools, employees using MCP clients can prompt the server for information from the instance or to perform an action on the instance, such as to summarize a list of recently closed incidents or cases. You can also use the ServiceNow Model Context Protocol Client to let an AI application on one instance access tools from an MCP server on another instance.

MCP Server Console users

User	Description
AI administrator	AI administrators can create MCP servers and configure which tools they expose to MCP clients. They also create OAuth inbound integrations on the instance for each client and configure clients with the server and authentication details.

MCP Server Console workflow

The following infographic shows the workflow for AI administrators to get started managing MCP servers with MCP Server Console.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. As an AI administrator, you identify which functionality on an instance to access from an external MCP client and employee experience.
2. You can use the preconfigured Quickstart Server or create an MCP server with the appropriate tools to use the desired functionality.

   You can create one or more servers that expose different tools for different use cases, such as for HR or IT workflows, or for different clients.
3. If needed, you can create additional tools from Now Assist skills and configure which fields are exposed to clients and then add the tools to servers.
4. You create OAuth inbound integrations for each client to secure access to servers on an instance.
5. You configure clients to connect to a server using the server and OAuth inbound integration details.

   Then the client calls the server, which validates the OAuth token and responds with the list of tools available for use with the server.
6. Employees use clients, such as AI agents, to prompt the server for data from the instance or to perform an action on the instance.

Note: With AI Gateway in AI Control Tower, AI administrators can monitor MCP server access and view metrics for servers and their tools. For more information, see AI Gateway.

MCP Server Console benefits

Benefit	Feature	Users
Integrate with any AI application and MCP client using a standard protocol.	Create a Model Context Protocol server	AI administrator
Control which tools and fields are exposed to MCP clients.	Create a tool for a Model Context Protocol server	AI administrator
Securely access functionality from a ServiceNow instance in any external employee experience.	Configure an MCP client to connect to an MCP server	AI administrator

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Quickstart Server in MCP Server Console

MCP Server Console includes a preconfigured Quickstart Server to help you get started learning about MCP servers and connecting MCP clients to a ServiceNow instance. The Quickstart Server includes tools for looking up and summarizing incident and case records.

Quickstart Server tools

Name	Description
Case summarization	Get a summary of details in a case record generated by Now Assist.
Incident summarization	Get a summary of details in an incident record generated by Now Assist.
Look up Case Records	Look up case records to enable summarizing them with the Case Summarization tool.
Look up Incident Records	Look up incident records to enable summarizing them with Incident summarization tool.

What to explore next

To learn more about configuring and using MCP Server Console, see:

- Configuring MCP Server Console
- Configure an MCP client to connect to an MCP server
- MCP Server Console reference

Related topics

Model Context Protocol Client

Configuring MCP Server Console

Create a Model Context Protocol (MCP) server and configure the tools and inputs it exposes to MCP clients.

Configuration overview

1. Create a Model Context Protocol server

   An AI administrator creates a server and adds tools to the server.

2. Create a tool for a Model Context Protocol server

   If additional tools are needed, the AI administrator identifies which functionality to expose and creates tools based on Now Assist skills. From the tools, they configure which fields are exposed to clients as tool inputs.

After configuring a server, the AI administrator creates an OAuth inbound integration for each client and configures clients to connect to the server using the server and authentication details. For more information, see Connecting to an MCP server from an MCP client.

Create a Model Context Protocol server

Create a Model Context Protocol (MCP) server and configure which tools it exposes to MCP clients.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: sn_mcp_server.admin or admin

About this task
Depending on your requirements, you can create one or more servers that expose different tools for different use cases, such as for HR or IT workflows, or for different clients. You can also use the preconfigured Quickstart Server.

Procedure
1. Navigate to **All** > **MCP Server Console**.
2. From the Configuration tab, select **Servers**.
3. Select **New**.
4. On the form, fill in the fields.

Server form

Field	Description
Label	An internal name for the server.
Short description	A description of the capabilities and purpose of the server.

5. From the Tools section, add the tools that the server exposes.
   a. Select **Add tools**.
   b. Search for tools from the list and select the tools to expose.
   If you need to create a tool, select **Create tool**. For more information, see Create a tool for a Model Context Protocol server.
   c. Select **Add**.
6. Select **Create** to make the server active.

> **Note:** To save your changes but not make the server active yet, select **Save as draft** instead. To activate it later, you can select **Activate** from the server record or the Servers page.

After creating the server, the Server URL field is populated with the endpoint that you use to integrate with clients: `https://<instance>.service-now.com/sncapps/mcp-server/mcp/<server-name>`.

Trouble?
If you need to turn off access to an existing server, select **Deactivate** from the server record or the Servers page.

What to do next
Configure clients to connect to the server. For more information, see Connecting to an MCP server from an MCP client.

Create a tool for a Model Context Protocol server

Create a tool from a Now Assist skill to expose it to Model Context Protocol (MCP) clients from an MCP server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

If you aren't using the Quickstart Server, create a server to which you can add tools. For more information, see Create a Model Context Protocol server.

Role required: sn_mcp_server.tools_admin, sn_mcp_server.admin, or admin

About this task

Each server must include at least one tool.

- Tools define which functionality and data a server exposes to clients and the actions that can be performed on an instance by clients. Tools are based on existing capabilities, such as Now Assist skills, including custom skills created with Now Assist Skill Kit. For a list of Now Assist skills that can be used as tools, see Now Assist skill support in MCP Server Console.
- Tools include inputs that correspond to the fields of the existing capability. Inputs that are enabled for a tool are exposed to clients.

Procedure

1. Navigate to **All** > **MCP Server Console**.
2. From the Configuration tab, select **Tools**.
3. Select **New**.
4. On the form, fill in the fields.

Tool form

Field	Description
Now Assist skill	A Now Assist skill to use as a tool with servers.
Label	An internal name for the tool.
Servers	One or more servers from which you want to expose the tool.
Description	A description of what the tool does and the inputs it exposes to clients. The description can be accessed from clients and used to determine when to call this tool.

In the Tool inputs section, the fields associated with the capability are added.

5. **Optional:** Turn off inputs from the tool that you don't want to expose.
   a. In the Tool inputs section, locate the tool input.
   b. From the Enabled column, select the toggle to turn off the input.

   Note: Some tool inputs are required and can't be turned off.

6. Select **Create**.

Trouble?

Note: If you add inputs to a Now Assist skill after a tool has been created, you must create another tool to include the additional inputs and replace the existing tool.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

Configure clients to connect to the server and use the tool. For more information, see Connecting to an MCP server from an MCP client.

Related topics

Now Assist skills

Now Assist Skill Kit

Connecting to an MCP server from an MCP client

Connect to a Model Context Protocol (MCP) server from an MCP client by creating an OAuth inbound integration and configuring the client with the server details.

With the Quickstart Server or after creating a server, you can complete the following steps to connect to a server from a client:

1. Create an OAuth inbound integration for an MCP client
2. Configure an MCP client to connect to an MCP server

The process to configure a client to connect to a server is dependent on the client used. For more information, refer to the documentation for your AI application and client.

Create an OAuth inbound integration for an MCP client

Secure access to Model Context Protocol (MCP) servers on an instance by creating an OAuth inbound integration for each MCP client.

Before you begin

Role required: oauth_admin, mi_admin, admin

About this task

For each client that you want to access servers on an instance, create an OAuth inbound integration in Machine Identity Console. To create the OAuth integration, you need a redirect URL from the client. For more information, refer to the documentation for your AI application and client.

Procedure

1. Navigate to **All** > **MCP Server Console**.
2. From the Configuration tab, select **Servers**.
3. From the OAuth setup required banner, select **Set up OAuth**.
   Alternatively, you can navigate to **All** > **Machine Identity Console** and select the **Inbound integrations** tab.

   Note: In the list of existing inbound integrations, you might see integrations created with the same names as servers (including underscores). These are integrations for monitoring servers from AI Control Tower and shouldn't be used to integrate with clients.

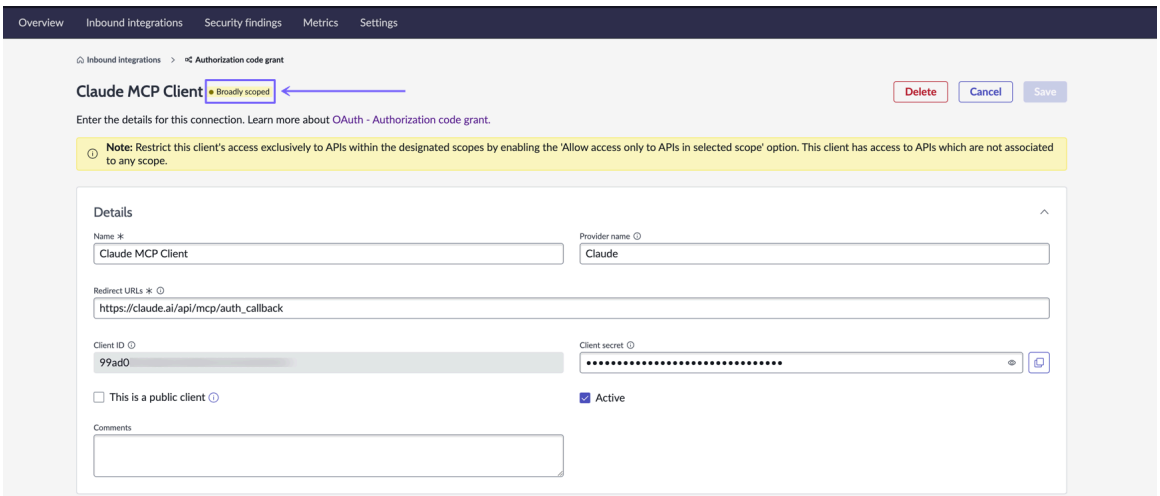
4. Select **New integration**.
5. Select **OAuth - Authorization code grant**.
6. On the form, fill in the required fields.
   For more information about this form, see Configure an OAuth authorization code grant.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Inbound integration form

Field	Value
Details section	
Name	Enter a name for the OAuth integration.
Redirect URLs	Enter the redirect URL for a client. The authorization code is sent to this URL after authentication. To get the redirect URL, refer to the documentation for your AI application and client. To connect to the ServiceNow MCP client on another instance, use the following redirect URL: `https://<client-instance>.service-now.com/oauth_redirect.do`. For more information, see the Model Context Protocol Client documentation.
Auth scope section	
Allow access only to APIs in selected scope	Clear the check box to make the OAuth integration broadly scoped.
Advanced options section	
Token Format	Select **JWT**.

7. Select **Save**.
   The OAuth inbound integration is created as broadly scoped with a client ID and client secret that you use when configuring the client to connect to servers on the instance.



What to do next

Configure the client to use the client ID and client secret to authenticate with servers on the instance.

Configure an MCP client to connect to an MCP server

Configure a Model Context Protocol (MCP) client to connect to an MCP server and prompt the server to perform a task.

Before you begin

Role required: none

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

The process to configure a client to connect to a server is dependent on the client used. The following procedure is a high-level overview of the workflow to configure a client to call a server. For more information, refer to the documentation for your AI application and client. For an example that demonstrates how to connect from a server on one instance to the ServiceNow Model Context Protocol Client on another instance, see the example following this procedure.

Procedure

1. Configure a client with the required server details as determined by the client.

Server details for MCP client configuration

Server details	Value
Server URL	`https://<server-instance>.service-now.com/sncapps/mcp-server/mcp/<server-name>` To connect to the preconfigured Quickstart Server, use `https://<server-instance>.service-now.com/sncapps/mcp-server/mcp/sn_mcp_server_default`.
Host	`<server-instance>.service-now.com`
Base URL	`/sncapps/mcp-server`
Scope	mcp_server
Authentication type	OAuth 2.0
Identity Provider	Generic OAuth 2
Authorization URL	`https://<server-instance>.service-now.com/oauth_auth.do`
Token URL	`https://<server-instance>.service-now.com/oauth_token.do`
Token Revocation URL	`https://<server-instance>.service-now.com/oauth_revoke.do`
Refresh URL	`https://<server-instance>.service-now.com/oauth_auth.do`
Redirect URL	`https://<server-instance>.service-now.com/oauth/callback`
Client ID	The client ID from the OAuth inbound integration on the server instance.
Client secret	The client secret from the OAuth inbound integration on the server instance.

After configuring these details, the client calls the server with the `Authorization: Bearer <token>` header. If the token is validated by the server, the client receives the list of tools available for use.

2. From the client, you can view the list of tools available to determine how to prompt the server.
3. Enter a prompt for the information you need or for the tool to perform an action on the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, if the Look up Incident Records tool is available, you could enter "Get all open incidents." With the Case summarization tool, you could enter "Summarize all cases closed this week."

The server runs the relevant tools and returns the result to the client as JSON data. The client presents the response as formatted text.

Example: Connecting to an MCP server from ServiceNow Model Context Protocol Client

This example demonstrates how to connect to a server from an AI agent on another instance using the ServiceNow Model Context Protocol Client. First, you configure the client to call the preconfigured Quickstart Server. From an AI agent, you access the Quickstart Server's list of tools and add individual tools to the agent. Lastly, you test the agent in AI Agent Studio by providing a prompt and seeing the agent's response. For more information, see the Model Context Protocol Client documentation.

Role required: sn_mcp_client.admin

1. On the server instance, create an OAuth inbound integration for the ServiceNow Model Context Protocol Client.

   For more information, see Create an OAuth inbound integration for an MCP client.

2. On the client instance, navigate to **All** > **AI Agent Studio** > **Settings**.
3. Select **Manage MCP Servers**.
4. Select **New**.
5. Add the Quickstart Server.

   For more information about this step, see Add an MCP Server with OAuth 2.1.

   a. On the Add MCP server form, fill in the fields.

Add MCP server form

Field	Value
Name	Quickstart Server
Authentication Type	OAuth 2.1
MCP server URL	`https://<server-instance>.service-now.com/sncapps/mcp-server/mcp/sn_mcp_server_default`

   b. Select **Next**.

   c. On the form, fill in the fields.

Add MCP server form

Field	Value
Client registration type	Manual Registration
Grant type	Authorization Code

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
Token authentication method	Client Secret Post
Client ID	The client ID from the OAuth inbound integration for the Model Context Protocol Client on the server instance.
Client secret	The client secret from the OAuth inbound integration for the Model Context Protocol Client on the server instance.
Authorization URL	`https://<server-instance>.service-now.com/oauth_auth.do`
Token URL	`https://<server-instance>.service-now.com/oauth_token.do`
Token Revocation URL	`https://<server-instance>.service-now.com/oauth_revoke.do`

d. Select **Save**.

6. Verify the OAuth configuration.
   a. Select **Authenticate**.
   b. Select **Allow** to allow the client to connect to the server.
7. Add tools from the Quickstart Server to an AI agent.

   For more information about this step, see Add an MCP server to an AI agent.

   a. In AI Agent Studio, select the **Create and manage** tab.
   b. From the AI agents tab, select an existing agent or create one.

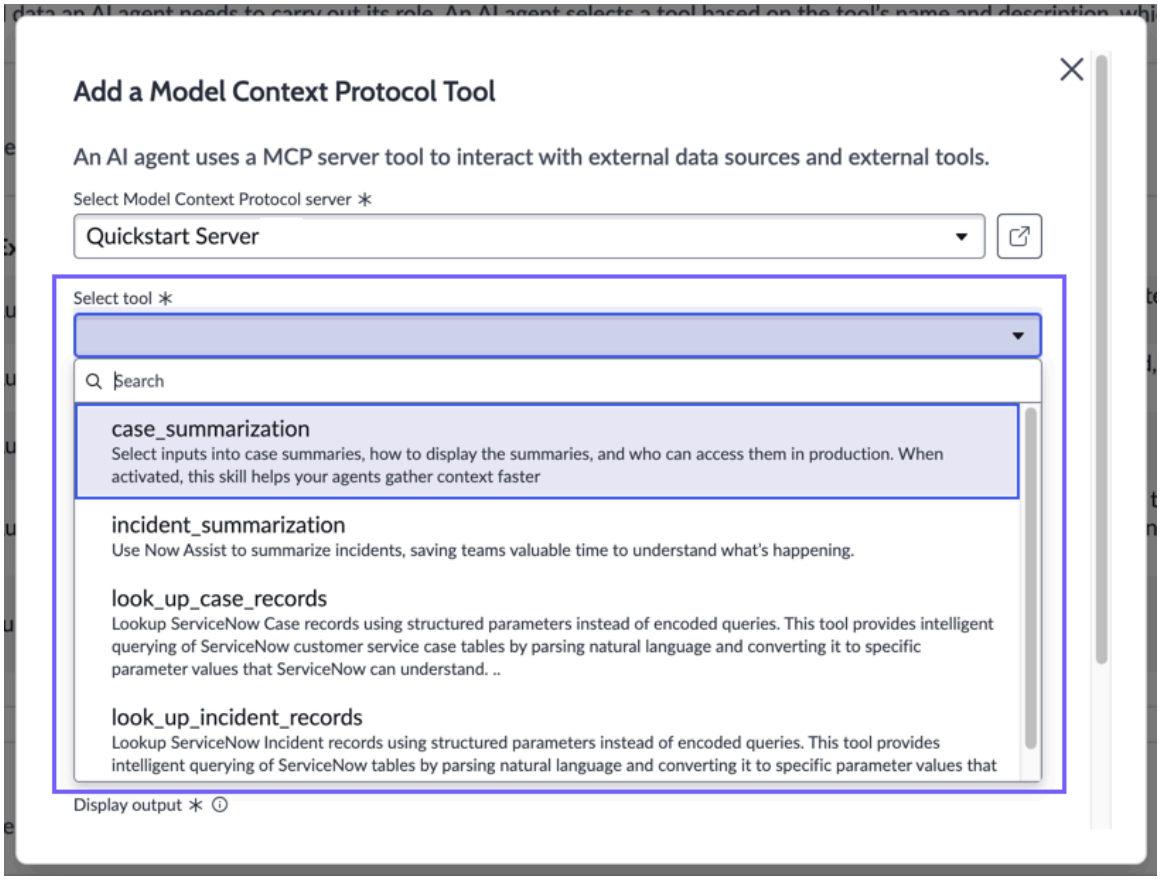
   For information about creating an agent, see Create an AI agent.

   c. Select **Add tools and information**.
   d. Select **Add tool** > **MCP server tool**.
   e. On the form, fill in the fields.

Add a Model Context Protocol Tool form

Field	Value
Select Model Context Protocol server	Quickstart Server
Select tool	Select the tools from the Quickstart Server that you want to use with this AI agent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



f. Select **Add**.

g. Select **Save and continue**.

8. Test the AI agent.

   For more information about this step, see Test an AI agent.

   a. In AI Agent Studio, select the **Testing** tab.

   b. Select **Start manual test**.

   c. In the Choose a test type field, select **AI agent or workflow**.

   d. Select the AI agent you configured and its version.

   e. In the Task field, enter a prompt to get information or to perform an action on the instance.

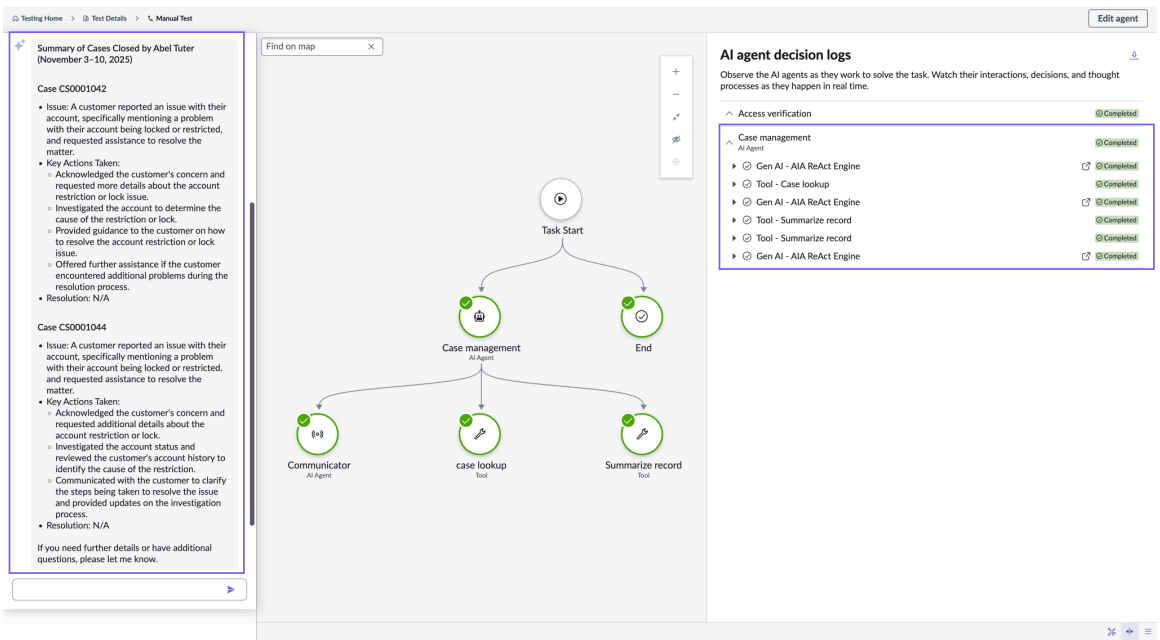
      The prompt should be based on which tools are available. For example, if you added the Look up Case Records and Case summarization tools, you can enter "Summarize all cases closed this week."

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



f. Select **Continue to Test Chat Response**.

The AI agent calls the server, and the server runs the tools requested based on the prompt. The server returns the information to the agent as JSON data, and the agent presents it as formatted text. In this example, the agent returns summaries of the cases closed by Abel Tuter in the past week.



MCP Server Console reference

Reference topics include information about MCP Server Console roles, tables, and more.

Now Assist skill support in MCP Server Console

MCP Server Console supports creating tools for Model Context Protocol (MCP) servers from a subset of Now Assist skills.

The following Now Assist skills are preconfigured as tools MCP Server Console:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Case summarization
- Incident summarization

In addition, you can create tools from the following Now Assist skills or from custom skills created with Now Assist Skill Kit. Only Now Assist skills that don't rely on internal information as inputs, such as sys_ids, are available to be created as tools. Some Now Assist skills don't currently meet the criteria required to be used as tools. For more information, contact your ServiceNow representative.

- Analyze the CMDB Search request
- Change request risk explanation
- Code Assist AutoComplete
- Code Assist Edit
- Code Assist Generation
- Common control objective creation
- Correlation insights generation
- Customer Service Summary
- Docs summarization in Collaborative Work Management
- Duration to Days convertor
- EAP doc summarization
- Error resolution recommendation
- Feedback summarization
- Flow generation
- Invoice Inquiry solution generator skill
- Multi feedback summarization
- Playbook Generation
- Playbook generation with images
- Policy Based HR Case Evaluator
- Problem investigator skill
- Project doc summarization
- Recommend preferred solution for VIT
- Refine content
- Regulatory alert impacted citations
- Regulatory alert impacted control objectives
- Regulatory alert impacted controls
- Relevant invoice finder
- Resume Skill Extraction
- Skill keyword extractor
- Spoke Generation
- Subflows and actions
- Supplier recommendation

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Survey Filler Answer Suggestion skill
- Text2Test Sentinel
- Text2Test Skill
- Vulnerable item deduplication

Related topics

Create a tool for a Model Context Protocol server

Now Assist skills

Now Assist Skill Kit

Components installed with MCP Server Console

Several types of components are installed with activation of the MCP Server Console application, including tables and user roles.

Roles installed

Role	Description	Contains roles
MCP Server administrator [sn_mcp_server.admin]	Performs create, read, update, and delete (CRUD) operations on all MCP Server Console tables.	sn_mcp_server.tools_admin
MCP Server tools administrator [sn_mcp_server.tools_admin]	Performs create, read, update, and delete (CRUD) operations for tools.	• canvas_user • sn_nowassist_admin.user • web_service_admin
MCP Server viewer [sn_mcp_server.viewer]	Invokes tools and views all MCP Server Console tables (read-only access).	• one_extend_viewer • canvas_user • sn_nowassist_admin.user • web_service_admin

Tables installed

Table	Description
AI Skill Tool Definition [sn_mcp_ai_skill_tool_definition]	Tools based on Now Assist skills. Extends the Tool Definition table.
AI Skill Tool Input [sn_mcp_ai_skill_tool_input]	Inputs for tools based on Now Assist skills. Extends the Tool Input table.
Provisioned MCP Server	MCP servers deployed on the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
[sn_mcp_server_registry]	
Scripted REST Tool Definition [sn_mcp_scripted_rest_tool_definition]	Tools based on Scripted REST APIs that enable other tools. Extends the Tool Definition table.
Scripted REST Tool Input [sn_mcp_scripted_rest_input]	Inputs for tools based on Scripted REST APIs. Extends the Tool Input table.
Tool Definition [sn_mcp_tool_definition]	Tools defined for use by MCP servers.
Tool Input [sn_mcp_tool_input]	Inputs defined for use by MCP server tools.

Now Assist Readiness Evaluation

The Now Assist Readiness Evaluation app helps prepare your organization for implementing agentic AI for ITSM and CSM, Now Assist for five products, or both.

Choose one of these tiles to get started.

Explore	Configure	Use
		
Learn about the Now Assist Readiness Evaluation app.	Run scheduled jobs and configure the app through the Now Assist Readiness Evaluation guided setup.	Learn about how to use the Now Assist Readiness Evaluation dashboards and assessments to prepare for implementing agentic AI, Now Assist, or both.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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.

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.

Related topics

Now Assist

Now Assist AI agents

Exploring Now Assist Readiness Evaluation

The Now Assist Readiness Evaluation app automates assessment processes, evaluates data readiness impacting implementation, and provides actionable insights to promote adopting Now Assist quickly. The app enables you to assess whether updates, installations, or customizations of your instance could affect implementation. The assessments provide direct hyperlinks to improve any issues found.

Now Assist Readiness Evaluation overview

The Now Assist Readiness Evaluation app is a solution designed to simplify and automate the agentic AI and Now Assist implementation assessment process. The app helps to determine whether your organization's instance is ready to implement generative and agentic AI features in Now Assist. Previously, manually assessing your organization's instance readiness was time consuming and took significant effort. Now, you can use the Now Assist Readiness Evaluation app to automate gathering, processing, and analyzing instance data so that you can review your

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

instance readiness more quickly. The app provides results within seconds, helping to reduce manual effort and promoting rapid and reliable assessments. The app includes assessment information for Now Assist by separating agentic AI solutions and generative AI solutions by tabs. Agentic AI solution information appears in the agentic_AI  tab and generative AI solution information appears in the  tab.

The following list details what agentic solutions for Now Assist are evaluated and assessed after the necessary scheduled jobs have run.

Agentic AI for Now Assist for IT Service Management (ITSM)

Evaluates whether customizations on ITSM-related parent tables (such as Incident and Task) may interfere with AI agent behavior, task execution, or decision-making. Two use cases are available to select: **Categorize Incidents** and **Generate Change Request Plans**. The assessment identifies and flags high-risk customizations to help promote seamless agent operations and stable instance performance. Each finding includes interactive hyperlinks to the impacted records, enabling you to review, validate, and address potential issues without navigating away from the dashboard.

Agentic AI for Now Assist for Customer Service Management (CSM)

Assesses whether customizations on CSM-specific tables (such as Case and Interaction) impact AI agent functionality. The assessment evaluates data completeness, field modifications, and scripting issues that may hinder agent-driven tasks and offers targeted estimated remediation to maintain agent efficacy. Each assessment finding includes interactive hyperlinks to the affected records, enabling you to navigate, validate, and resolve issues directly from the dashboard.

The following list details what generative AI Now Assist products are evaluated and assessed after the necessary scheduled jobs have run.

Now Assist for AI Search

This feature provides insights into the status and configuration of AI Search, such as active or inactive profiles, widget versions, and estimated remediation efforts. The assessment also evaluates Knowledge Management by assessing article counts, categories, group restrictions, and top searches or articles viewed. Additionally, the assessment highlights areas for improvement, offering actionable feedback to enhance knowledge quality and hygiene.

Now Assist for Virtual Agent (VA)

This feature checks if Virtual Agent is live, tracks active topics and catalog items, and identifies barriers to conversational functionality. The assessment highlights opportunities to convert non-conversational items to conversational ones and tracks automation in catalog items, along with estimated remediation efforts.

Now Assist for IT Service Management (ITSM)

This feature reviews customizations and configurations in the Incident and Task tables, such as fields, UI actions, and scripts. The assessment highlights changes made, tracks data completeness, and provides key metrics to improve accuracy and consistency in ITSM processes. The assessment also specifies the estimated remediation effort required and reports the language used in the instance.

Now Assist for Customer Service Management (CSM)

This feature reviews customizations in the Case and Task tables, tracks changes to fields, and monitors data completeness for key fields. The assessment provides insights into past cases and field-filling percentages, helping improve data accuracy in Customer Service Management. The assessment also reports the estimated remediation effort needed and the language used in the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist for HR Service Delivery (HRSD)

This feature analyzes customizations in the HR Core Case and Task tables, tracking field-level changes and promoting data completeness for critical fields. The assessment provides insights into historical HR Core cases, including field population percentages, to enhance data accuracy in HRSD processes. Additionally, the feature identifies estimated remediation efforts required and reports the languages used within the instance.

Note: Additional configuration is needed to run the HRSD assessment successfully. For more information, see Configure the Now Assist Readiness Evaluation guided setup.

Now Assist Readiness Evaluation users

Users

User	Description
Admin	Admins use the Now Assist Readiness Evaluation app to prepare for implementing generative and agentic AI features in Now Assist.

Now Assist Readiness Evaluation benefits

Now Assist Readiness Evaluation benefits

Benefit	Feature	Users
Run scheduled jobs for applicable agentic and generative AI Now Assist products from a single assessment job.	Run the GenAI/ AgenticAI Assessment scheduled job	Admins
Review an overall dashboard of your issues before implementing agentic and generative AI Now Assist.	Using Now Assist Readiness Evaluation dashboard	Admins
View your overall go or no-go status to implement Now Assist generative AI skills or agentic AI agents, and view the go or no-go status' issues categorized by percentage.	Assessing readiness status	Admins
Review summary information of overall implementation readiness.	Reviewing your Now Assist assessment and Reviewing your Agentic AI assessment	Admins
Improve the gaps for implementation by using the hyperlinks found in the assessments to fix issues. Direct hyperlinks to records and tables make it easier to validate, investigate, or act without switching contexts.	Reviewing your Now Assist assessment and Reviewing your Agentic AI assessment	Admins
Receive an estimated remediation effort of development work days to fix issues before implementation.	Reviewing your Now Assist assessment and Reviewing	Admins

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist Readiness Evaluation benefits (continued)

Benefit	Feature	Users
	your Agentic AI assessment	

What to explore next

To learn more about configuring and using the Now Assist Readiness Evaluation app, see:

- Configuring Now Assist Readiness Evaluation
- Using Now Assist Readiness Evaluation

Configuring Now Assist Readiness Evaluation

Run the scheduled jobs and complete the Now Assist Readiness Evaluation guided setup configuration steps before viewing the generative AI and agentic AI assessment results.

Configuration overview

Before configuring the Now Assist Readiness Evaluation app, confirm that the Now Assist Readiness Evaluation app is installed. Navigate to **All** > **System Definition** > **Plugins**, and search for either the app name, `Now Assist Readiness Evaluation`, or app ID, `sn_assess`, to confirm.

1. Run the GenAI/AgenticAI Assessment scheduled job

   Run the GenAI/AgenticAI Assessment scheduled job so that agentic and generative AI assessment results populate in **Workspaces** > **Now Assist Readiness Evaluation**.

2. Configure the Now Assist Readiness Evaluation guided setup

   Complete the configuration for the five Now Assist jobs. Additional configuration is required to work with the Now Assist for HRSD assessment.

Run the GenAI/AgenticAI Assessment scheduled job

Before you can view the agentic AI or Now Assist assessment results, you must first run the GenAI/AgenticAI Assessment scheduled job.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Scheduled Jobs**.
2. Set the Scheduled Jobs filter to `Application is Now Assist Readiness Evaluation`, and then select **Run**.

   Tip: If the Application column isn't displayed in your ServiceNow instance, personalize the column settings using the Update Personalized List icon from the list header.

3. Select **GenAI/Agentic AI Assessment** from the Name column.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Select **Execute Now**.
   This action begins the scheduled job for the agentic and generative AI assessments found in **Workspaces** > **Now Assist Readiness Evaluation**.

What to do next

If you must rerun all the scheduled jobs, or specifically, the agentic AI scheduled jobs for ITSM and CSM, repeat these steps.

Note: You must complete the additional steps in the Now Assist Readiness Evaluation guided setup if you want to work with the Now Assist for HRSD product. If you don't complete the additional guided setup steps for the Now Assist for HRSD product, the assessment continuously fails. For more information about completing the Now Assist Readiness Evaluation guided setup, see Configure the Now Assist Readiness Evaluation guided setup.

Configure the Now Assist Readiness Evaluation guided setup

Before you can review the Now Assist assessments in **Workspaces** > **Now Assist Readiness Evaluation**, you must first complete the configuration for the five Now Assist jobs.

Before you begin

Before completing this guided setup, you must first run the generative AI and agentic AI scheduled jobs. For more information, see Run the GenAI/AgenticAI Assessment scheduled job.

Role required: admin

Procedure

1. Navigate to **All** > **System Definition** > **Plugins**.
2. In the search bar, search for `Now Assist Readiness Evaluation`.
3. Select the Installed tab to view the installed store application **Now Assist Readiness Evaluation** App id: sn_assess.
4. Select **Now Assist Readiness Evaluation**.
5. In the Get started area, select **Configure**.
6. Review the Want to know how it works? screen, and then select **Continue**.
   You're redirected to the Agentic AI Assessment Guided Setup.

   Tip: You can also access the Agentic AI Assessment Guided Setup screen by navigating directly to **All** > **Now Assist Readiness Evaluation** > **Guided Setup**.

   Three sections appear: Scheduled Jobs, Restricted Caller Access Privileges, and System Properties. The Restricted Caller Access Privileges and System Properties sections are only applicable if you want to work with Now Assist for HR Service Delivery (HRSD).
7. Select **Start** in the Scheduled Jobs section.
   You're redirected to the five scheduled Now Assist generative AI jobs.

   Note: You're unable to run agentic AI scheduled jobs from this guided setup. To run agentic AI scheduled jobs, you must do so through the GenAI/AgenticAI Assessment scheduled job. For more information on running the agentic AI scheduled job, see Run the GenAI/AgenticAI Assessment scheduled job.

8. Select the link to the specific scheduled job to have that job open in a new browser tab.
   For example, if you select the Now Assist Assessment - AI Search link, that scheduled job appears in a new browser tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

9. Select **Execute Now** to begin that job.
   The job begins and you're redirected back to the guided setup.
10. Select **Next** in the guided setup.
11. Repeat steps 8-10 until you have run the scheduled jobs for the five Now Assist assessments.
    You have completed the Now Assist assessment jobs. If you ran the HRSD job, continue to the following optional steps.
12. **Optional:** If you have run the job for HRSD, select the **Change The State to Allowed for the following records** link on the Restricted Caller Access Privileges screen.

    Note: The following sub-steps can only be completed when in the Human Resources: Core application scope. For more information on changing application scopes, see Select an application from the application picker.

    a. Change the Status column from `Requested` to `Allowed` on the Restricted Caller Access Privileges table, and then go back to the guided setup.
    b. Select **Next** on the Restricted Caller Access Privileges screen.
    c. Select the **Configure System Properties** link on the System Properties screen to review the system properties.
    d. Select **Mark Complete**.
13. Select **Close**.

Result
You have completed configuration for the Now Assist assessments.

What to do next
You can navigate to **Workspaces** > **Now Assist Readiness Evaluation** to review your assessment results.

Using Now Assist Readiness Evaluation

The Now Assist Readiness Evaluation app helps you find actionable items in your implementation preparation and provides direct hyperlinks to improve upon those gaps.

Overview of Now Assist Readiness Evaluation

Use the Now Assist Readiness Evaluation app to help you prepare to launch generative AI, agentic AI, or both in Now Assist for your organization. The following list is a high-level overview of the documentation topics in this section that you can refer to on how to use and interpret the Now Assist Readiness Evaluation results.

- Using Now Assist Readiness Evaluation dashboard

  View an overall dashboard of your assessments' findings divided up into percentages and bar chart visuals.

- Assessing readiness status

  View an overall dashboard of readiness and issue categorization before implementing Now Assist.

- Reviewing your Now Assist assessment

  View the Now Assist assessment report results for the five available Now Assist products.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Reviewing your Agentic AI assessment

  View the agentic AI assessment report results for IT Service Management (ITSM) and Customer Service Management (CSM).

Using Now Assist Readiness Evaluation dashboard

The Now Assist Readiness Evaluation dashboard comes with two dashboard tabs: **Agentic AI - Assessment** and **Now Assist Assessment**.

Navigate to **Workspaces** > **Now Assist Readiness Evaluation** to access the Now Assist Readiness Evaluation dashboard. The Now Assist Readiness Evaluation dashboard is the first tab you see.

The Now Assist Readiness Evaluation dashboard is a powerful product feature that accelerates generative and agentic AI readiness for Now Assist by delivering automated, data-driven insights.

Note: The **Remediation properties** tab may also appear on the Now Assist Readiness Evaluation dashboard. This tab only appears if the *sn_assess.effort_visibility* system property is set to true. For more information about this system property, see Now Assist Readiness Evaluation system properties.

After reviewing the dashboard and making the suggested assessment changes, rerun your scheduled jobs to see how your results improve.

Agentic AI - Assessment dashboard tab

The Now Assist Readiness Evaluation Dashboard's Agentic AI - Assessment tab helps determine agentic AI readiness for IT Service Management (ITSM) and Customer Service Management (CSM) by delivering automated, data-driven insights.

Note: For results to appear in this dashboard, ensure that you have completed the configuration steps first. For more information on configuration, see Configuring Now Assist Readiness Evaluation.

The following list describes details of the Agentic AI - Assessment dashboard tab.

Total Issues

This area of the dashboard shows the total issues identified.

Issues - Trends

This area of the dashboard shows the trends over time in a bar chart for the products that need the most attention before implementation.

Total Estimated Effort (Days)

This area of the dashboard shows the total estimated work effort in development days.

Estimated Effort - Categorized

This area of the dashboard shows the estimated work effort needed to correct the issues for each product, depending on the issue category, in a bar chart.

Issues - Categorized

This area of the dashboard shows the percentage of each issue category for each product in a bar chart.

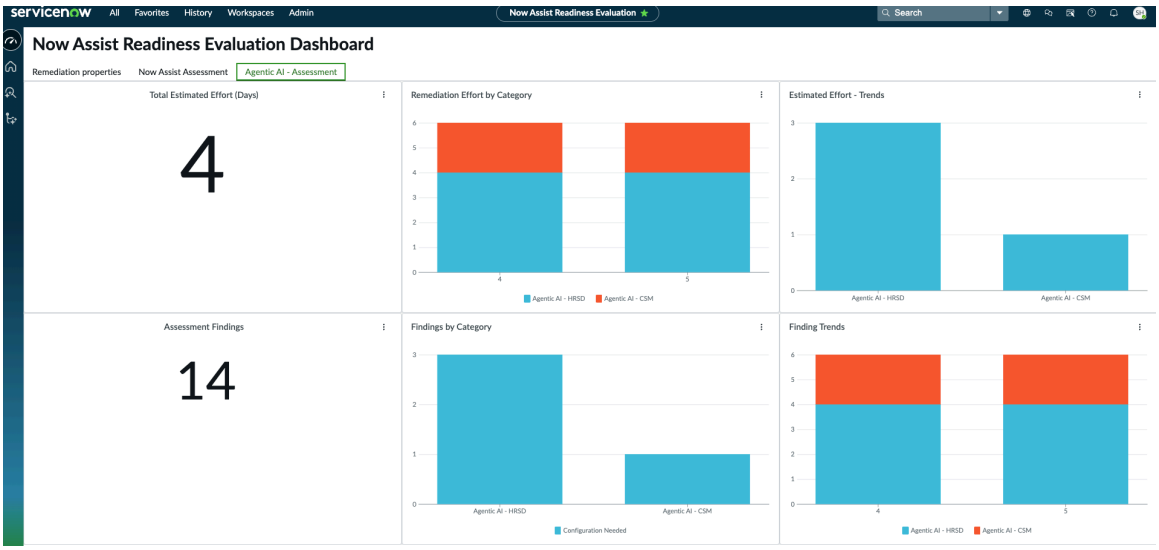
Estimated Effort - Trends

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This area of the dashboard shows the trends over time in a bar chart for each products' estimated work effort.

You can select each widget on the dashboard to have that table data open in a separate tab.

Example of the Now Assist Readiness Evaluation Dashboard's Agentic AI - Assessment



Now Assist assessment dashboard tab

The Now Assist Readiness Evaluation Dashboard's Now Assist dashboard tab helps determine generative AI readiness for Now Assist by delivering automated, data-driven insights.

Note: For results to appear in this dashboard, verify that you have completed the configuration steps first. For more information on configuration, see Configuring Now Assist Readiness Evaluation.

The following list describes details of the Now Assist Assessment dashboard tab.

Total Estimated Effort (Days)

This area of the dashboard shows the total estimated work effort in development days.

Remediation Effort by Category

This area of the dashboard shows the estimated work effort needed to correct the findings for each product, depending on the findings category, in a bar chart.

Estimated Effort - Trends

This area of the dashboard shows the trends over time in a bar chart for each products' estimated work effort.

Assessment Findings

This area of the dashboard shows the number of assessment items that need attention.

Findings by Category

This area of the dashboard shows the percentage of each assessment finding category for each product in a bar chart.

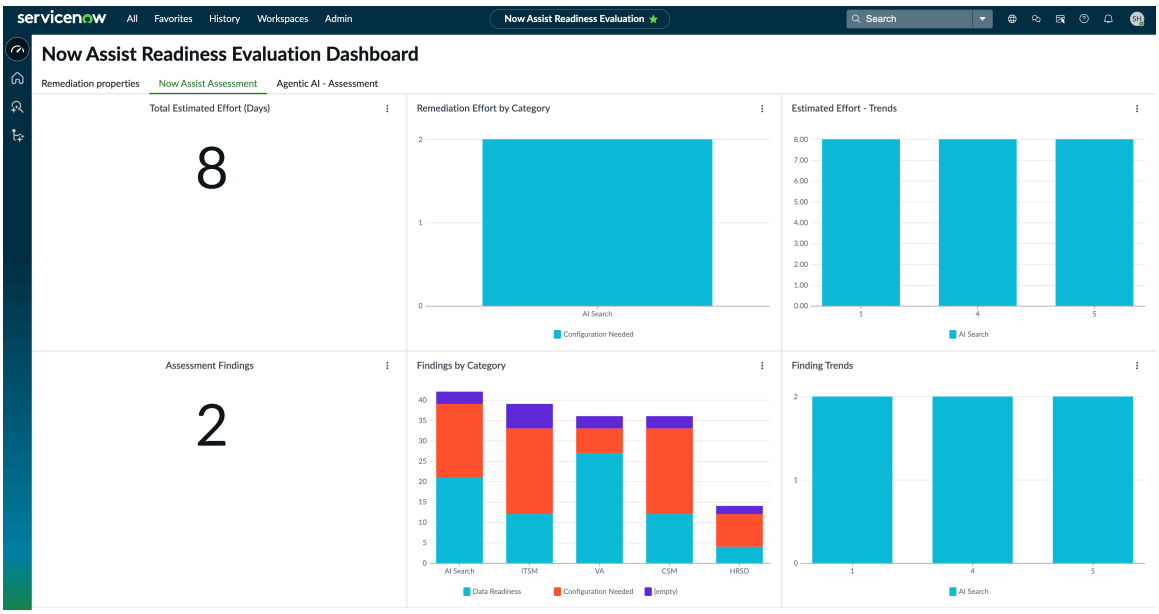
Finding Trends

This area of the dashboard shows the trends over time in a bar chart for the products that need the most attention before implementation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can select each widget on the dashboard to have that table data open in a separate tab.

Example of the Now Assist Readiness Evaluation Dashboard's Now Assist Assessment



Assessing readiness status

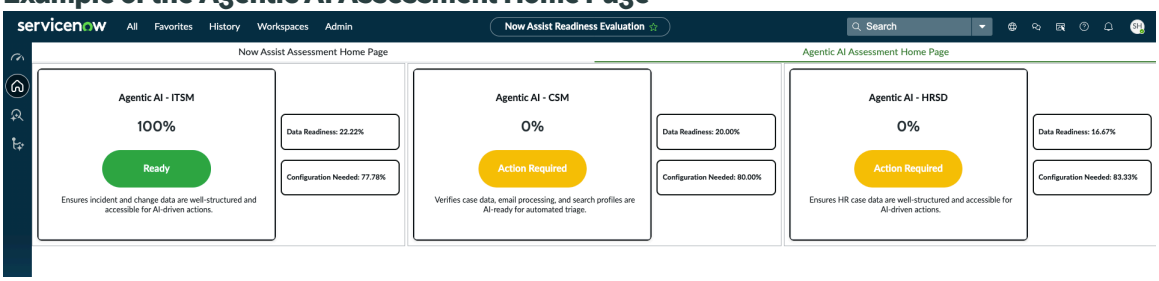
The Now Assist Readiness Evaluation app shows you a high-level overview of your go or no-go status for your organization's readiness to implement agentic AI, generative AI, or both in Now Assist.

Navigate to **Workspaces** > **Now Assist Readiness Evaluation** and select the Home icon to review your assessments' home page.

An overall ready or action-required status is determined by percentage rules. A green **Ready** button appears if over 75% of your instance is ready to implement agentic AI, generative AI, or both in Now Assist. A yellow **Action Required** button appears if 75% or less of your instance is determined to be ready to implement agentic AI, generative AI, or both in Now Assist. Each finding is broken down and automatically tagged as a product, data, or configuration issue. Having issues categorized lets you know whether you should open a Now Support case, cleanse data, or adjust platform settings.

The agentic AI Assessment Home Page tab details the Ready or Action Required for Now Assist agentic AI in ITSM and CSM.

Example of the Agentic AI Assessment Home Page

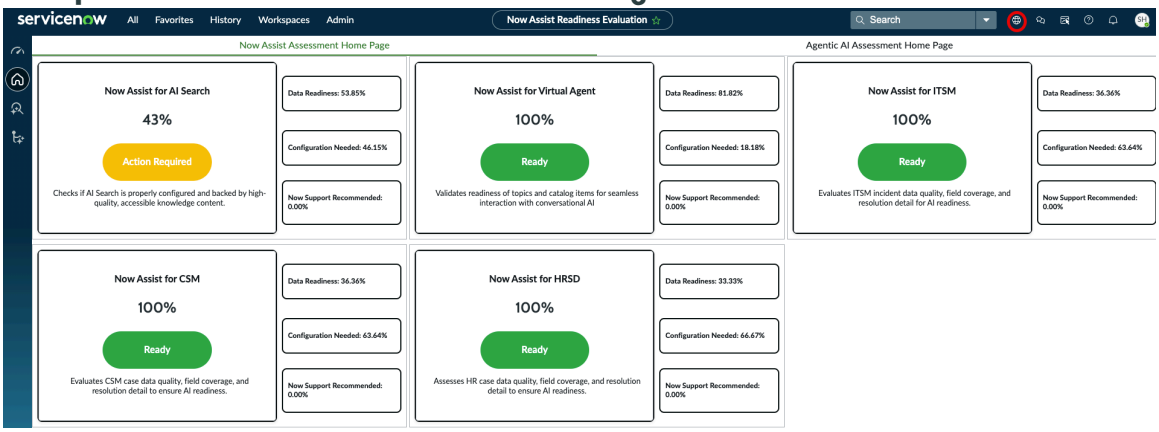


The Now Assist Assessment Home Page tab details the go or no-go status for Now Assist products and features, including:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Now Assist for AI Search
- Now Assist for Virtual Agent
- Now Assist for IT Service Management (ITSM)
- Now Assist for Customer Service Management (CSM)
- Now Assist for HR Service Delivery (HRSD)

Example of the Now Assist Assessment Home Page



Select the assessment feature's tile on the home page to be redirected to that assessment tab's report results.

To impact these percentages, rerun your scheduled jobs after making your suggested assessment fixes.

Reviewing your Now Assist assessment

This automated assessment process evaluates potential implementation impacts and provides actionable insights for Now Assist products.

Navigate to **Workspaces** > **Now Assist Readiness Evaluation** and select the assessment icon  to review your Now Assist assessment.

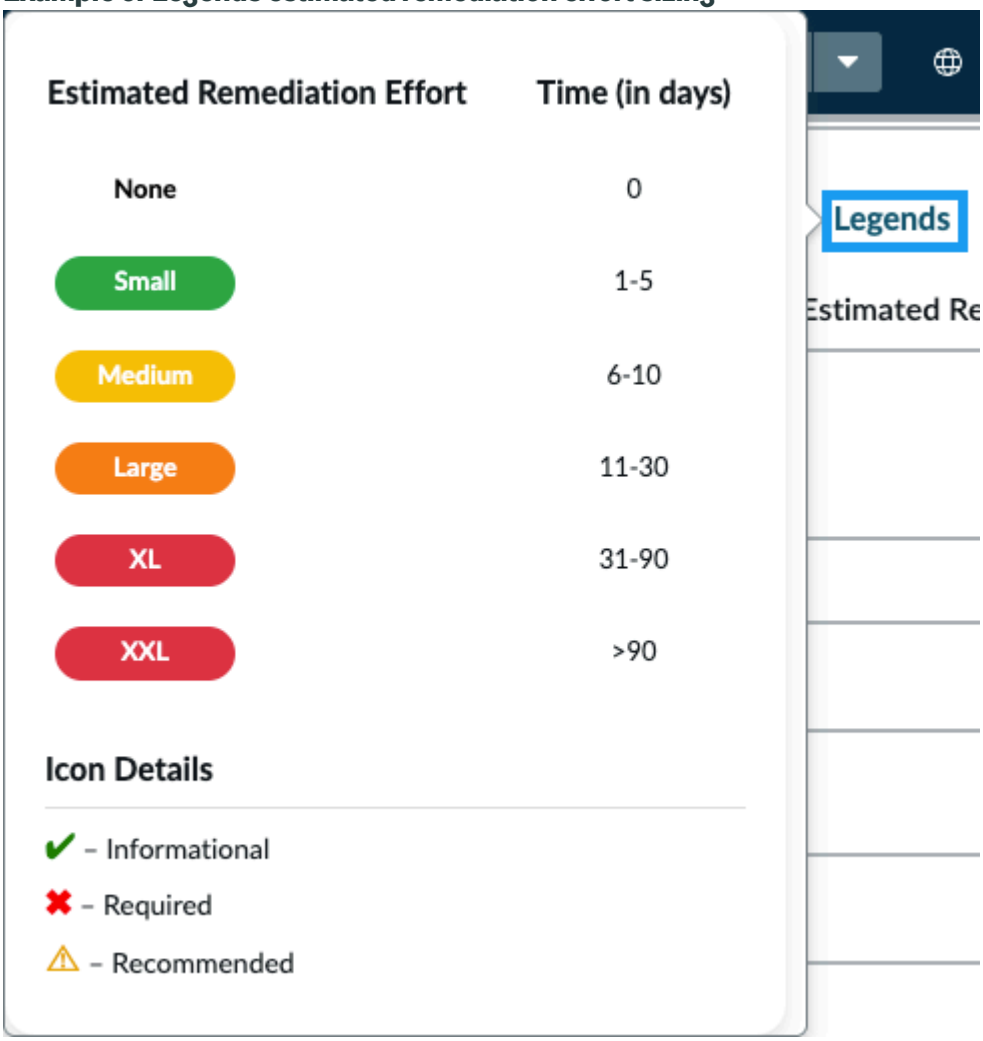
Note: Before viewing your Now Assist assessments, you must first have run the scheduled jobs to get the assessment results. If you haven't run your scheduled jobs, you are prompted to do so on an individual job level. For more information on running scheduled jobs, see Run the GenAI/AgenticAI Assessment scheduled job and Configure the Now Assist Readiness Evaluation guided setup.

The summary tab provides a consolidated summary of features, including Now Assist capabilities across AI Search, Virtual Agent (VA), IT Service Management (ITSM), Customer Service Management (CSM), and HR Service Delivery (HRSD), as well as their corresponding estimated remediation efforts. The results shown are estimates. You should evaluate results provided by Now Assist Readiness Evaluation for accuracy and appropriateness for your use case.

To understand the estimated remediation effort, select the **Legends** button. The legend shows the estimated remediation effort using sizing verbiage and colors, along with an estimated timeline of remediation in development days. The legend also includes descriptions of the icons found next to the assessment questions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example of Legends estimated remediation effort sizing



After reviewing the summary, select the individual assessment links to be redirected to that specific assessment in a new browser tab. The questions for each Now Assist assessment differs. For example, the questions and answers on the Now Assist for AI Search assessment are different than the questions and answers on the Now Assist for Virtual Agent assessment.

Note: The assessments gather information from tables. To view any of the tables' available questions, answers, assessments, or assessment run successes or failures, search for the applicable table in the navigation filter:

- Questions [sn_assess_question] table
- Answers [sn_assess_answer] table
- Assessment [sn_assess_assessment] table
- Assessment Run [sn_assess_assessment_run] table

For more information about how to search for tables in the navigation filter, see Navigate directly to a table.

The assessments include visual indicators to guide you on non-blocker and blocker results. The following list details the descriptions of the assessment's visual indicators:

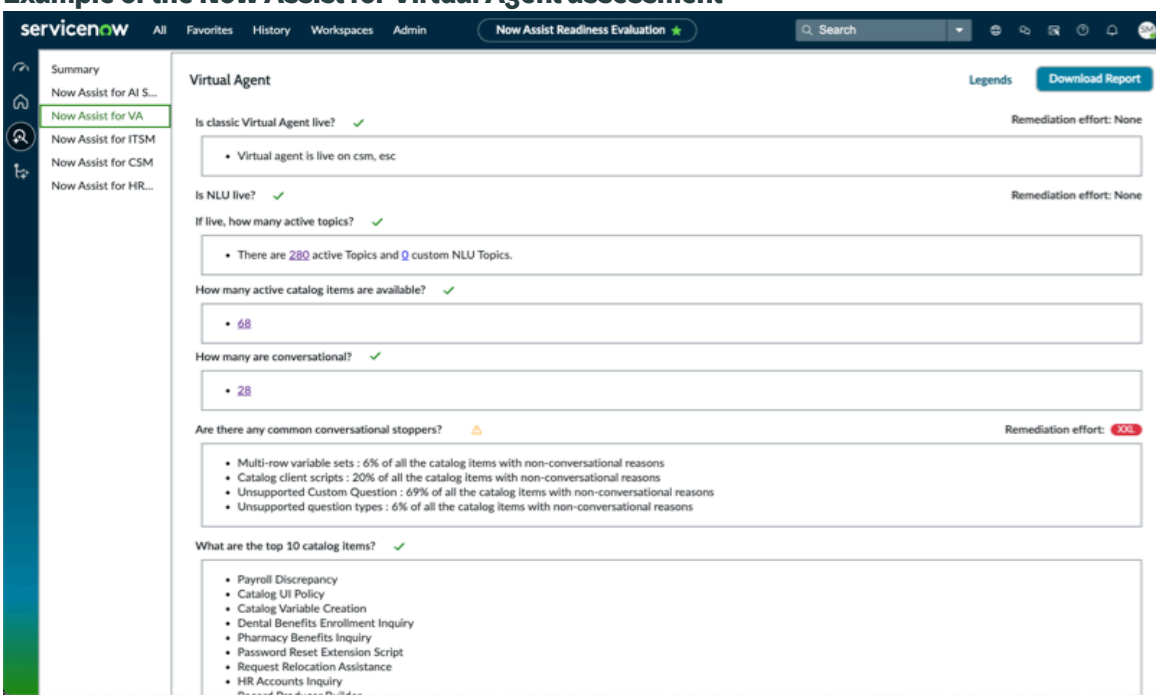
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- ✓: Informational area and no further action needed.
- ✗: Further action is needed to resolve findings before implementation. These required areas include estimated remediation efforts.
- ⚠: Recommended area warning, but the finding should be evaluated before implementation. These recommendations are observations that should be reviewed prior to implementation but no estimated remediation effort is provided.

You can select **Download Report** to export the entire Now Assist assessment report in PDF format from the Summary tab. The exported report still includes hyperlinks to the specific individual assessments. Selecting **Download Report** on an individual assessment downloads only that report in PDF format. After selecting a hyperlink in the PDF, that record or table opens in a separate browser tab.

Note: You don't have to select **Download Report** to work directly with the problematic records or tables found in the assessment. Each assessment includes hyperlinks to the records or tables that need attention. Selecting a hyperlink opens the record or table in a separate browser tab.

Example of the Now Assist for Virtual Agent assessment



After reviewing the assessment and making the suggested assessment changes, rerun your scheduled jobs to see how your results improve. For more information on rerunning a specific job or all jobs, see Run the GenAI/AgenticAI Assessment scheduled job.

Reviewing your Agentic AI assessment

This automated assessment process evaluates potential implementation impacts and provides actionable insights for agentic AI in Now Assist for IT Service Management (ITSM) and Now Assist for Customer Service Management (CSM).

Navigate to **Workspaces** > **Now Assist Readiness Evaluation** and select the agentic AI assessment icon to review your agentic AI assessment.

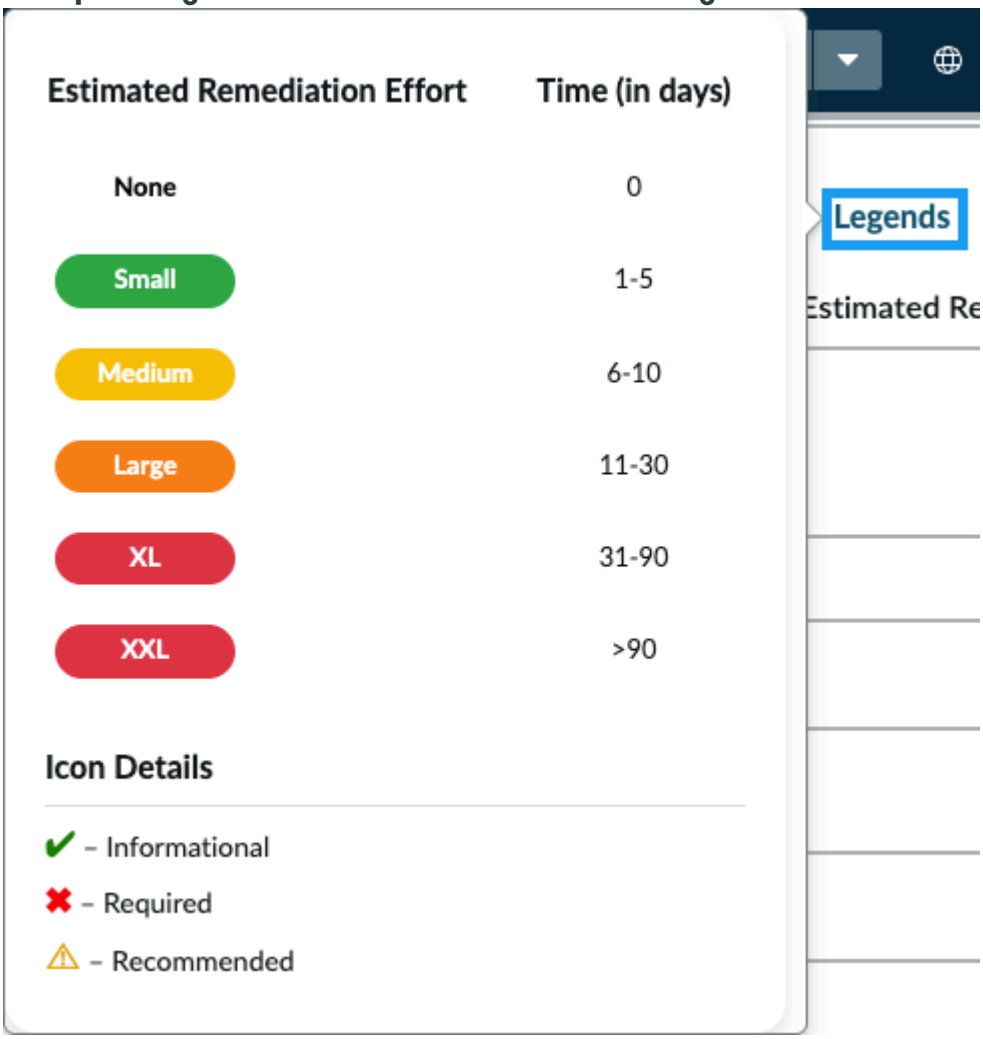
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Before viewing your agentic AI assessments, you must first have run the scheduled job to get the assessment results. If you haven't run your scheduled jobs, you are prompted to do so on an individual job level. For more information on running scheduled jobs, see Run the GenAI/AgenticAI Assessment scheduled job.

The summary tab provides a consolidated summary of overall findings along with summarization and resolution notes for ITSM and CSM. The results shown are estimates. You should evaluate results provided by Now Assist Readiness Evaluation for accuracy and appropriateness for your use case.

To better understand the estimated remediation effort, select the **Legends** button. The legend shows the estimated remediation effort using sizing verbiage and colors, along with an estimated timeline of remediation in development days. The legend also includes descriptions of the icons found next to the assessment questions.

Example of Legends estimated remediation effort sizing



After reviewing the summary, select the individual assessment tabs to review those specific assessments. The questions for each assessment differ. For example, the questions and answers on the **Agentic AI - ITSM** assessment are different than the questions and answers on the **Agentic AI - CSM** assessment.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: The assessments gather information from tables. To view any of the tables' available questions, answers, assessments, or assessment run successes or failures, search for the applicable table in the navigation filter:

- Questions [sn_assess_question] table
- Answers [sn_assess_answer] table
- Assessment [sn_assess_assessment] table
- Assessment Run [sn_assess_assessment_run] table

For more information about how to search for tables in the navigation filter, see Navigate directly to a table.

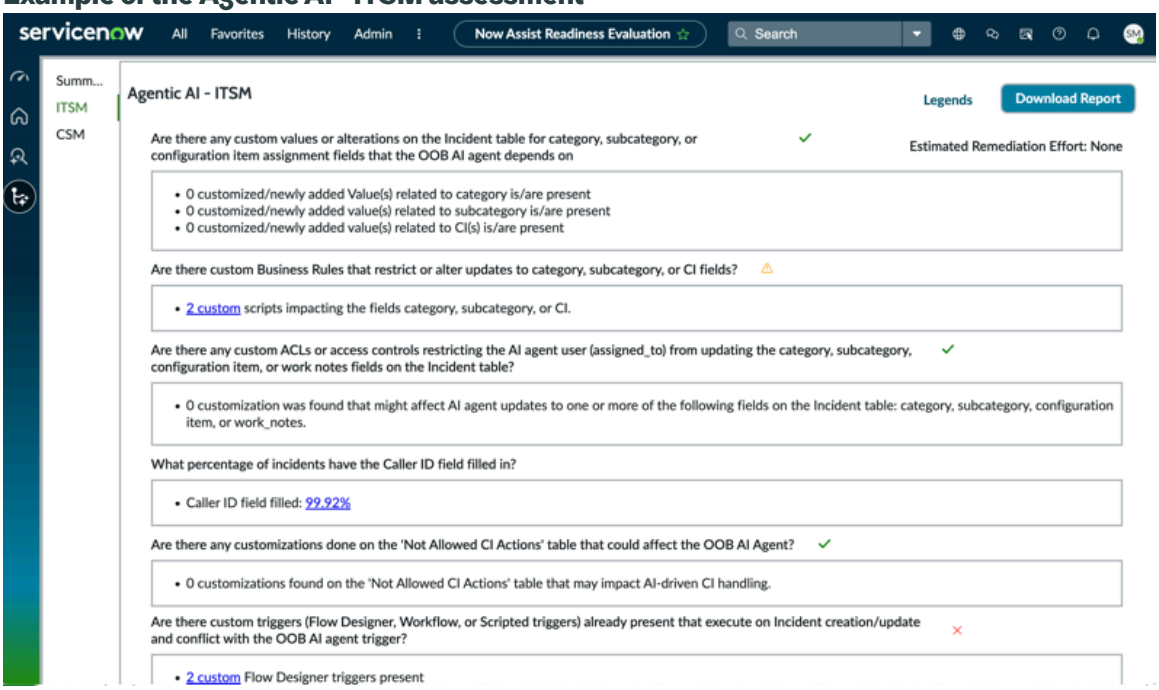
The assessments include visual indicators to guide you on non-blocker and blocker results. The following list details the descriptions of the assessment's visual indicators:

- ✓: Informational area and no further action needed.
- ✗: Further action is needed to resolve findings before implementation. These required areas include estimated remediation efforts.
- ⚠: Recommended area warning, but the finding should be evaluated before implementation. These recommendations are observations that should be reviewed prior to implementation but no estimated remediation effort is provided.

You can select **Download Report** to export the entire agentic assessment report in PDF format from the Summary tab. The exported report still includes hyperlinks to the specific individual assessments. Selecting **Download Report** on an individual assessment downloads only that report in PDF format. After selecting a hyperlink in the PDF, that record or table opens in a separate browser tab.

Note: You don't have to select **Download Report** to work directly with the problematic records or tables found in the assessment. Each assessment includes hyperlinks to the records or tables that need attention. Selecting a hyperlink opens the record or table in a separate browser tab.

Example of the Agentic AI - ITSM assessment



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

After reviewing the assessment and making the suggested assessment changes, rerun your scheduled jobs to see how your results improve. For more information on rerunning a specific job or all jobs, see Run the GenAI/AgenticAI Assessment scheduled job.

Now Assist Readiness Evaluation reference

Now Assist Readiness Evaluation admins and implementers can work with select system properties for additional customization.

Now Assist Readiness Evaluation system properties

Use system properties to customize your readiness assessment results. Access these properties from the System Property [sys_properties] table.

Now Assist Readiness Evaluation system properties

Property	Description
sn_assess.assessment_limit	Reduce performance issues when a large volume of data is assessed. Adjust this system property, if needed, so that the job can run successfully. The default value is `10000`.
sn_assess.effort_visibility	Customize the estimated remediation effort for select efforts. This property controls multiple sections associated with remediation effort. The default value is `false`. When this system property is set to `true`, the **Remediation properties** tab appears on the Now Assist Readiness Evaluation dashboard. You customize the estimated remediation efforts via the **Remediation properties** tab. All input values represent estimated remediation efforts in days. After making changes to the estimated remediation efforts, select **Save**, and then re-run the assessments to see the updated remediation efforts reflected.
sn_assess.task_limit	Decide the maximum number of records to process for the ITSM assessment. The default value is 50. Be cautious of setting a higher limit or leaving the value empty since that would fetch all active records which could impact performance.

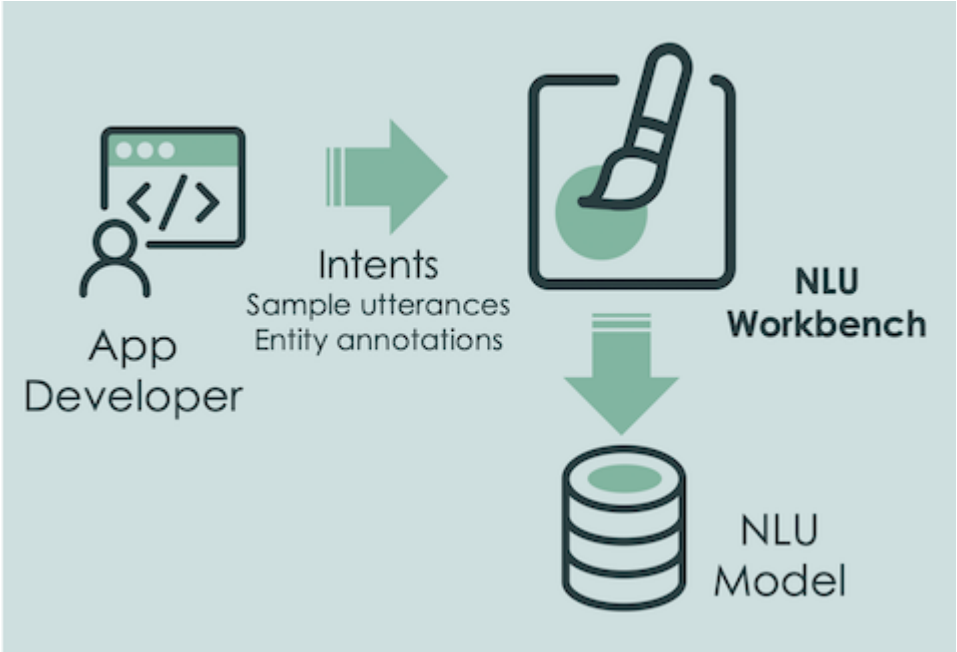
Natural Language Understanding

ServiceNow® Natural Language Understanding (NLU) provides an NLU Workbench and an NLU inference service that you can use to enable the system to learn and respond to human-expressed intent. By entering natural language examples into the system, you help it understand word meanings and contexts so it can infer user or system actions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Overview of Natural Language Understanding

User input flow in the NLU model build process



This image shows you the user input flow in the NLU model build process.

NLU terminology

In NLU parlance, these terms identify the key language components the system uses to classify, parse, and otherwise process natural language content.

Intent

Something a user wants to do or what you want your application to handle, such as granting access.

Utterance

A natural language example of a user intent. For example, a text string in an incident's short description, a chat entry, or an email subject line. Utterances are used to build and train intents and should therefore not include several or ambiguous meanings or intents.

Entity

The object of, or context for, an action. For example: a laptop, a user role, or a priority level.

System entity

These are predefined in an instance and have highly reusable meanings, such as date, time, and location.

User defined entity

These are created in the system by users and can be built from words in the utterances they create.

Common Entity

A context commonly used and extracted via a pre-defined entity model, such as currency, organization, people, or quantity.

Vocabulary

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

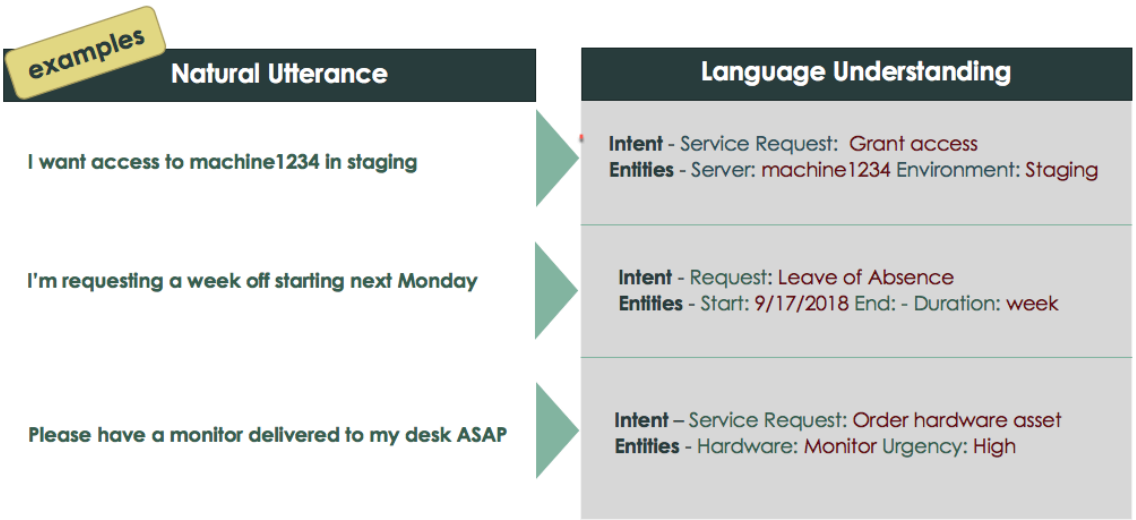
Vocabulary is used to define or overwrite word meanings. For example, you can assign the synonym “Microsoft” to the acronym “MS”.

NLU Model

A collection of utterance examples and their associated intents and entities that the system uses as a reference to infer intents and entities in a new utterance. The NLU Workbench comes with pre-built NLU models for specific business units, such as an ITSM model. You can also create custom models.

This image illustrates how Natural Language Understanding processes and renders utterance examples into intents and entities in the system.

NLU processes and renders utterance examples into intents and entities

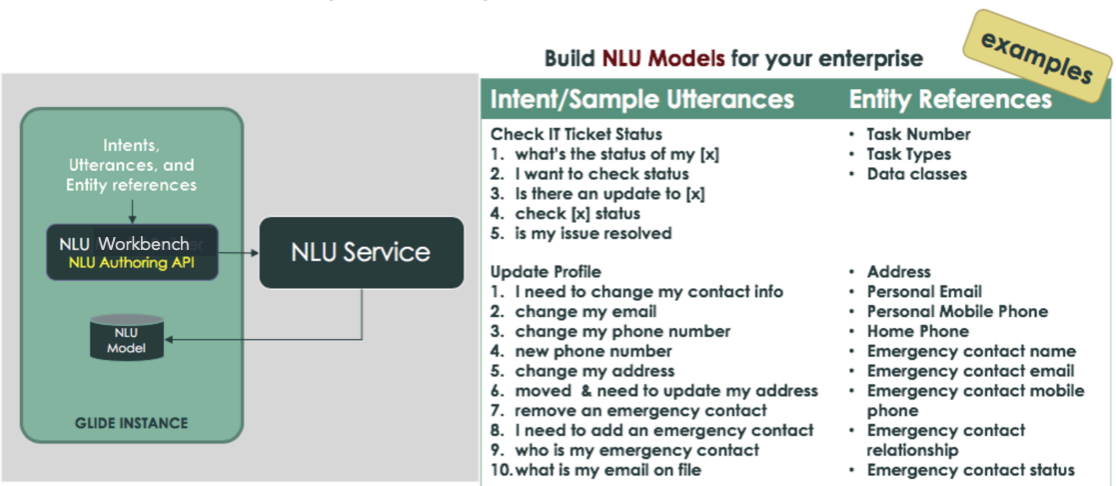


NLU Workbench

Use the NLU Workbench to create morphological representations of human language. These models enable you to create intents and entities expressed in natural language utterances. Any ServiceNow application can invoke an NLU model to get an inference of intents and entities in a given utterance.

Using the nlu_admin role, you build your models in the NLU Workbench, where you create, train, test, and publish them iteratively.

Overview of NLU Authoring API helping administrators build their models



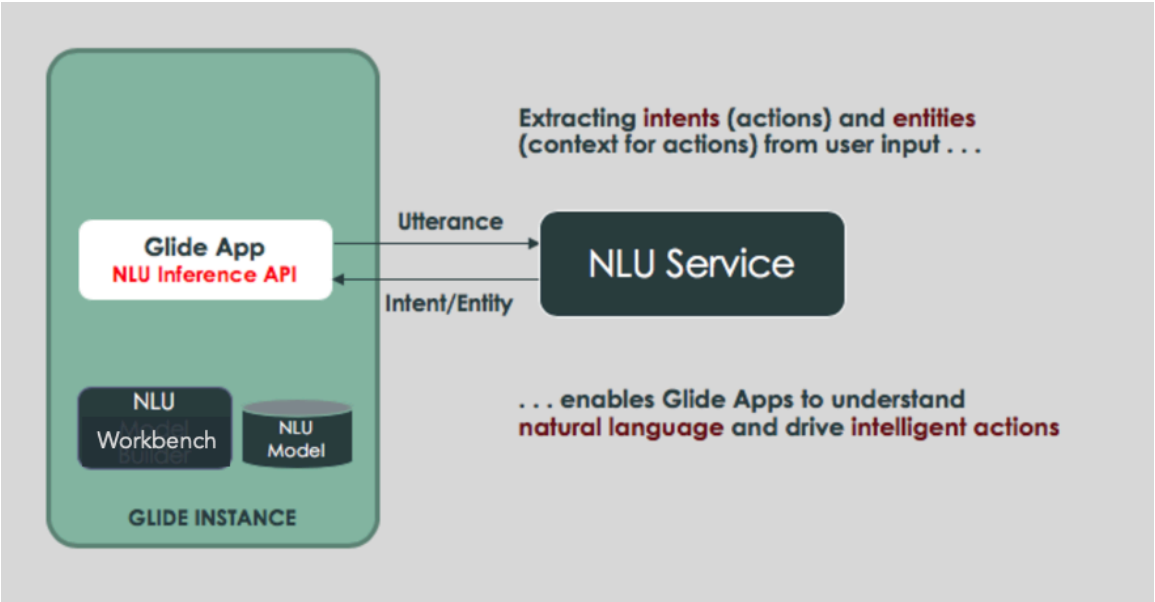
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For information on how to build and use an NLU model, see: Create an NLU model.

NLU inference service

Natural Language Understanding provides an NLU inference service that helps the system to understand natural language and drive intelligent actions. This service trains and predicts intents and entities for a given user utterance in your model so that its text translates into machine-understandable formats, such as APIs and parameters.

Overview of how the system uses an NLU inference API to extract intents and entities

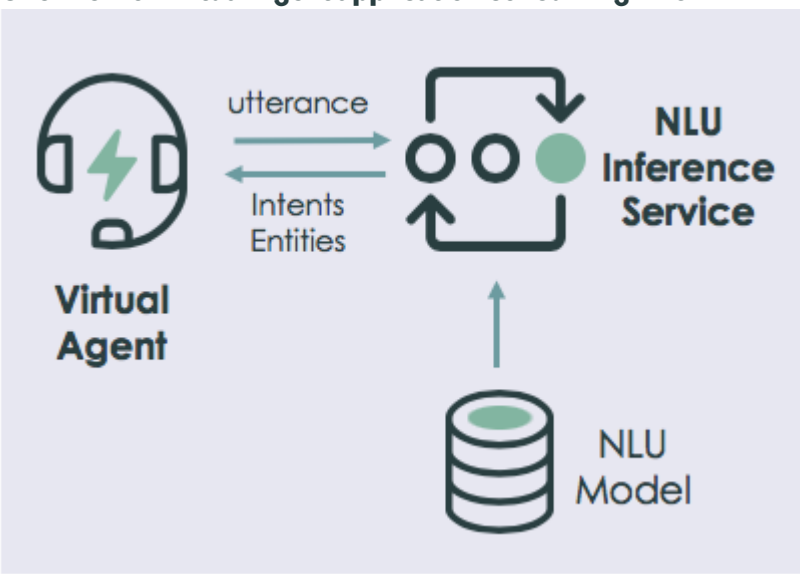


Here, the system uses an inference API to train NLU algorithms by using sample record data to identify intents and entities that are strong candidates for accurate prediction.

NLU model consumption

Other ServiceNow® applications consume NLU model output, such as Virtual Agent.

Overview of Virtual Agent application consuming NLU



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, Virtual Agent administrators can configure a Virtual Agent Designer conversation flow to consume NLU models so that agent chatbots can better understand user statements in the conversation. For more information on how Virtual Agent consumes NLU models, see: Natural Language Understanding (NLU) topic discovery in Virtual Agent.

Get started

Explore  Learn about NLU concepts and features.	Use  Create, test, translate, and publish your NLU models.	Integrate  Learn how Virtual Agent administrators can update NLU models from within Virtual Agent Designer.
	Reference  Learn about using models in different languages for use in other applications	

Troubleshoot and get help

- Virtual Agent & NLU community page
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Natural Language Understanding

ServiceNow® Natural Language Understanding (NLU) provides an NLU Workbench and an NLU inference service that you can use to enable the system to learn and respond to human-expressed intent. By entering natural language examples into the system, you help it understand word meanings and contexts so it can infer user or system actions.

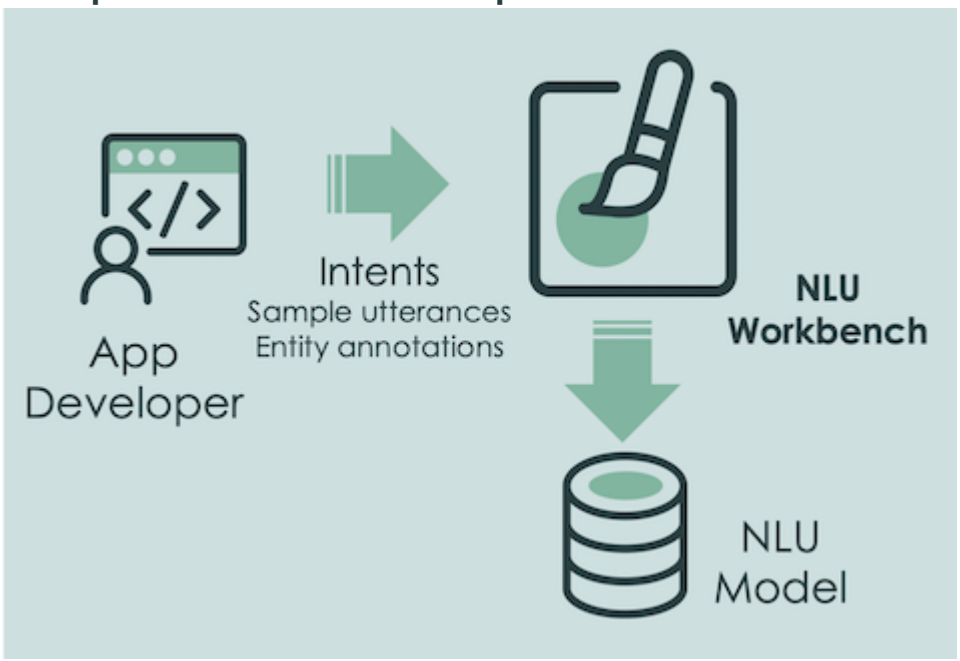
Natural Language Understanding overview

Natural Language Understanding (NLU) is a component of artificial intelligence natural-language processing that deals with machine-reading comprehension. Predictive Intelligence uses NLU so the system can understand word meanings and word contexts to infer user or system actions.

Natural Language Understanding workflow

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

User input flow in the NLU model build process



This image shows you the user input flow in the NLU model build process.

Natural Language Understanding benefits

Benefit	Feature	Users
Use the NLU Workbench to create models for Virtual Agent and AI Search.	NLU Workbench	NLU Administrator
Build the model's content by adding intents, entities, vocabulary, and test set utterances. Your NLU model content determines how the model responds to user inputs.	Model Management	NLU Administrator
Add a language to an existing NLU model by translating it. Use one of several translation options to add a secondary model in a supported language.	Model Management	NLU Administrator
Virtual Agent administrators can access and update their NLU models from within the Virtual Agent Designer user interface.	Virtual Agent Integration	Virtual Agent Administrator

Activate the NLU Workbench

Activate the following plugins to activate the NLU Workbench.

Before you begin

Role required: admin

About this task

Activate the following plugins if they aren't already active in your instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugin	Description
NLU Workbench - Core com.glide.nlu	Installs the required tables for persisting NLU models that are created using the NLU Workbench.
NLU Workbench com.snc.nlu_studio	Enables the creation of Natural Language Understanding (NLU) models. These models can understand the intent (action) and entities (details about the action) for a given user utterance. Any ServiceNow application can invoke an NLU model. Requires the NLU Workbench - Core plugin and the NLU Common Model plugin.
Predictive Intelligence com.glide.platform_ml	Enables the NLU Service APIs used for model creation and inference. Enables the creation of machine learning solutions that are trained on data in your instance. Provides frameworks for classification, similarity, and clustering. A trained solution can be invoked by any application by using a prediction API.
NLU Common Model com.glide.nlu.model	Packages all language NLU common models. Also includes commonly used pattern entities that can be imported and used in any NLU model in the NLU Workbench. Commonly used patterns like email, phone, and ServiceNow specific pattern entities such as INT, RITM are made available. Requires the NLU Workbench - Core plugin.
NLU Active Learning-Properties com.glide.nlu.active_learning_properties	Enables the nlu_admin to configure the system properties for the Expert Feedback Loop application. If you don't use the Expert Feedback Loop application, this plugin is not used and can safely be ignored. For more information about the system properties for the Expert Feedback Loop, see the NLU Expert Feedback Loop documentation.

The following two plugins are for apps associated with a for-fee subscription and are available on the ServiceNow Store. Installing these apps adds additional features to the NLU Workbench. Contact your account manager if you are interested in these apps.

Plugin	Description
Intent Discovery sn_nlu_discovery	Delivers the Intent Discovery feature, which identifies user intents by analyzing incident/case data. Use this application to help identify which intents to model and build for Virtual Agent conversations to attain maximum deflection. For more information, see Install Intent Discovery.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugin	Description
NLU Workbench - Advanced Features sn_nlu_workbench	Delivers the Model Performance, Multi-model Batch Testing, Cross-model Conflict Review, and Expert Feedback Loop features. For more information, see Install NLU Workbench - Advanced Features.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the following plugins using the filter criteria and search bar: NLU Workbench - Core (com.glide.nlu), NLU Workbench (com.snc.nlu_studio), Predictive Intelligence (com.glide.platform_ml), and NLU Common Model (com.glide.nlu.model).

   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.

NLU Workbench roles

NLU Workbench is installed with these roles.

To learn more about managing per-user subscriptions, see Managing per-user subscriptions in Subscription Management and contact your account representative.

NLU user [nlu_user]

Users who have read access to models in NLU Workbench.

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 installed with NLU Workbench - Core. This role does not have access to NLU Model Performance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

NLU editor [nlu_editor]

Users who can edit NLU models assigned to them in NLU Workbench.

Contains Roles

List of roles contained within the role.

- import_admin
- sn_ace.ace_user
- nlu_user

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 installed with NLU Workbench - Core.

NLU administrator [nlu_admin]

Users who can use NLU Workbench to manage NLU models.

Contains Roles

List of roles contained within the role.

- sn_ace.ace_user
- sn_nlu_workbench.nlu_feedback_admin
- nlu_editor
- nlu_user

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 installed with NLU Workbench - Core.

NLU feedback admin [sn_nlu_workbench.nlu_feedback_admin]

Data labelling (NLU feedback) Admin role - to manage data labelling across models also ability to optimize models.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Contains Roles

List of roles contained within the role.

- sn_ace.ace_user
- platform_ml_write
- platform_ml_create
- ml_labeler

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 installed with NLU Workbench - Advanced Features. For more information, see NLU Workbench - Advanced Features.

This role is used in Expert Feedback Loop modules. For more information, see NLU Expert Feedback Loop or Issue Auto Resolution Tuning in NLU.

NLU Workbench properties

Refer to these system properties for the Natural Language Understanding (NLU) application.

NLU Workbench properties and their usage

To access your system properties, use the admin or nlu_admin role and the following path in the application navigator: **All** > **NLU Workbench** > **Settings**.

Model Settings

Label and Name	Default value	Plugin	Recommen usage
Maximum number of utterances per intent *glide.nlu.utterances_per_intent.value_limit*	200	NLU Workbench	Use fewer than 200 utterances per intent t keep your model well balanced in terms of intent size.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Model Settings (continued)

Label and Name	Default value	Plugin	Recommen usage
			Note: Value must b greate than 5 and le than o equal 300.
Maximum number of records in a Table vocabulary source *glide.platform_ml.api.max_nlu_lookupsource_records*	100,000	NLU Workbench	Keep the value unde 100,000.
Maximum number of values in a List vocabulary source *glide.nlu.static_lookup.value_limit*	1,000	NLU Workbench	Keep the value unde 1,000.
Enable pre-built vocabulary for software names *glide.mlpredictor.option.nlu.@LookupSources:software*	enabled	NLU Workbench	Enable pre-built vocabulary so the system ca recognize software names.
Enable pre-built vocabulary for hardware names *glide.mlpredictor.option.nlu.@LookupSources:hardware*	enabled	NLU Workbench	Enable pre-built vocabulary so the system ca recognize hardware names.

Advanced Settings

Label and Name	Default value	Plugin	Recommended usage
Maximum number of records for Intent Discovery classification *sn_nlu_discovery.intent_discovery_max_classification_limit*	300,000	Intent Discovery	Keep the number of records less than 500,000.
Minimum number of records for Intent Discovery classification *sn_nlu_discovery.intent_discovery_min_classification_limit*	10,000	Intent Discovery	Use at least 10,000 records to get high quality results.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced Settings (continued)

Label and Name	Default value	Plugin	Recommended usage
Minimum number of records for NLU performance analysis *sn_nlu_workbench.glide.nlu.performance.min_clustering_records*	5,000	NLU Workbench - Advanced Features	Use at least 5,000 records to get high quality results.
NLU Conflict Detection - Moderate Threshold *sn_nlu_workbench.glide.nlu.conflict.moderate_threshold*	.85	NLU Workbench - Advanced Features	Must be a decimal between 0 and 1. Keep this threshold less than the Critical Threshold.
NLU Conflict Detection - Critical Threshold *sn_nlu_workbench.glide.nlu.conflict.critical_threshold*	.95	NLU Workbench - Advanced Features	Must be a decimal between 0 and 1. Keep this threshold greater than the Moderate Threshold.
The maximum number of rows in a batch test import file *sn_nlu_workbench.glide.nlu.batch_test.max_import_rows*	10,000	NLU Workbench - Advanced Features	Make sure your batch test import file has no more than 10,000 rows.
The maximum number of utterances to display for feedback in the expert feedback loop *glide.mlpredictor.option.nlu.activeLearning.label_candidate_table.max_respon*	300	NLU Workbench - Advanced Features	Pull no more than 300 utterances from your users' Virtual Agent chat logs to display for feedback in the Expert Feedback Loop application.The minimum umber of utterances a user should review before tuning the model
The minimum number of utterances a user should review before tuning the model *sn_nlu_workbench.glide.nlu.optimize.min_labeled_data*	100	NLU Workbench - Advanced Features	Provide and save feedback for at least 100 utterances from your users' Virtual Agent chat logs so you can execute the Tune Model feature in the Expert Feedback Loop application.
The maximum number of records to fetch from Virtual Agent chat logs *glide.mlpredictor.option.nlu.activeLearning.va_chat_logs.max_row_limit - 3000*	3,000	NLU Workbench - Advanced Features	If there is high NLU usage, increasing the default value to a maximum of 50,000 records will increase the data available for the active learning job to filter up on and display in the Expert Feedback Loop application to give feedback on.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced Settings (continued)

Label and Name	Default value	Plugin	Recommended usage
Size limit on Label Candidate Table (used for pruning the table) *glide.mlpredictor.option.nlu.activeLearning.label_candidate_table.max_data_s - 10000*	10,000	NLU Workbench - Advanced Features	The recommended usage for this property is the same as the property above.
Size limit on Labeled Data Table (used for pruning the table) *glide.mlpredictor.option.nlu.activeLearning.label_table.max_data_size - 10000*	10,000	NLU Workbench - Advanced Features	The recommended usage for this property is the same as the property above.
Enable this property to unblock your instance during NLU model training. The training will be scheduled for an off-peak time, and we will notify you when it's done. *glide.mlpredictor.scheduled.nlu.model.training*	False	NLU Workbench - Advanced Features	False

To get more feedback data from Virtual Agent (VA) chat logs, refer to the **Procuring additional VA feedback data on demand** section in the Expert Feedback Loop documentation.

NLU language support

The NLU Workbench application provides support for creating NLU models in different languages for use in other applications, such as Virtual Agent.

The platform supports NLU for 17 languages. Ten of these languages available to your models have both intent and entity support; the remaining 7 languages have intent-only support. (Japanese entity support includes character annotation.)

Note: Install the language plugins for languages you want to use in NLU. Installing and activating language plugins ensures that the languages are available in your instance. For more information see Activate a language.

Available in Tokyo, Utah, and Vancouver	Intent and entity support (in Utah and Vancouver)	Intent-only support (in Utah and Vancouver)
• Brazilian Portuguese • Chinese (simplified) • Danish • Dutch • English • Finnish • French • French Canadian	• Brazilian Portuguese • Dutch • English • French • French Canadian • German • Italian • Japanese	• Chinese (simplified) • Danish • Finnish • Korean • Norwegian • Polish • Swedish

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Available in Tokyo, Utah, and Vancouver	Intent and entity support (in Utah and Vancouver)	Intent-only support (in Utah and Vancouver)
• German • Italian • Japanese • Korean • Norwegian • Polish • Portuguese • Spanish • Swedish	• Portuguese • Spanish	

With intent and entity support, NLU can understand sophisticated utterances such as intent-entity relationships, system entities, and user-defined entities. NLU relays this information to Virtual Agent, and the user is usually taken directly to the conversation topic that offers resolution.

With intent-only support, the focus is on intent recognition. With Virtual Agent using NLU, users are directed to the desired conversation topic, where qualifying follow-up questions may be asked before being taken to a topic that offers resolution.

Utterances for all languages are case insensitive during intent prediction.

Below is a detailed list of the currently available NLU languages and the intents, entities, vocabulary, features, and applications they support.

Supported languages and features

	Intents	System-Defined Entities	User-Defined Entities	Vocabulary	Vocabulary Sources	Prebuilt SW/HW Vocabulary	Intent Discovery	Conflict Review	Batch Testing	NLU Performance	Fast Training	Expert Feedback Loop	Model Use
English	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	VA, Search
French	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓		VA, Search
Spanish	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓		VA, Search
German	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓		VA, Search
Japanese	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓		VA, Search
Portuguese	✓	✓	✓	✓	✓				✓	✓	✓		VA
Swedish	✓			✓					✓	✓	✓		VA
Italian	✓		✓	✓	✓				✓	✓	✓		VA
Chinese (simplified)	✓			✓					✓	✓	✓		VA
Brazilian Portuguese	✓	✓	✓	✓	✓				✓	✓	✓		VA
Dutch	✓		✓	✓	✓				✓	✓	✓		VA
French Canadian	✓		✓	✓	✓				✓	✓	✓		VA
Polish	✓			✓					✓	✓	✓		VA
Korean	✓			✓					✓	✓	✓		VA
Danish	✓			✓					✓	✓	✓		VA
Finnish	✓			✓					✓	✓	✓		VA
Norwegian	✓			✓					✓	✓	✓		VA

NLU Service updates

Refer to this documentation so you are up to date with changes to the NLU Service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Service update summary

The NLU Service helps the system to understand natural language and drive intelligent actions. This service trains and predicts intents and entities for a given user utterance in your NLU model so it can understand human-expressed natural language, whether spoken or written. The source of this documentation is KB0953693.

This service is updated independently of your instance upgrade, and without any action required by you. These updates are done on a bi-monthly basis (once every two months) to improve the quality of NLU model training and predictions. Major updates are aligned with family releases such as Rome, San Diego, Tokyo, etc. Minor updates are automatically updated so you are using the latest version when you retrain an NLU model. While most of these updates don't impact your existing use of NLU, there may be some changes you need to be aware of.

May 2023 NLU Service update

- Introduced dialog acts to enable natural mid-conversation in Virtual Agent (VA) and improve conversation fluidity. Affirm, negate, and modify dialog acts are supported in English and enabled by default for all new VA topics.
- Migrated all languages to use new language models, boosting average intent prediction quality by 10% across all languages.
- Enabled customers to manage and edit irrelevant utterances for their models to improve irrelevance detection.
- Removed the requirement for a model to have 2 or more intents in the model, making it easier for end-to-end topic testing in VA.

March 2023 NLU Service update

- Improved intent/entity detection through better handling of common words in vocabulary sources.
- Improved latency and memory utilization for system entity (NER) detection.
- Updated version support so that customers will need to use newer versions of the NLU Service and cannot point to n-2 releases older than their current glide version.

January 2023 NLU Service update

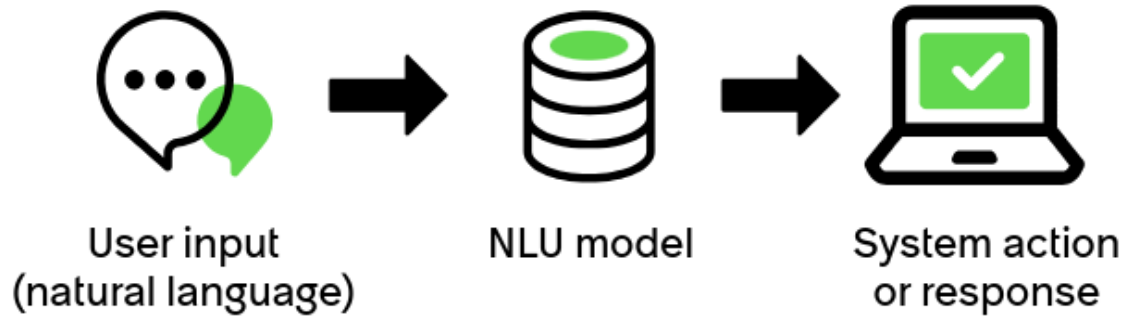
- Created the new DATE-TIME system entity (English only) for use in Virtual Agent.
- Added vocabulary source and entity support (simple, mapped, and open ended) for the Dutch and Italian languages, and system entity support for Portuguese and Brazilian Portuguese.
- Upgraded the ServiceNow Language Model for the Danish, Swedish, Finnish, and Norwegian languages, improving their average prediction quality by 17% from Tokyo.
- Improved the handling of punctuation in entities and special characters in vocabulary sources.
- Incorporated feedback provided on intent predictions by admin users in the NLU Workbench to improve model training data for Virtual Agent.

NLU models

Use NLU models to apply ServiceNow Natural Language Understanding on your instances. Create, manage, test, and publish NLU models with the NLU Workbench.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Overview of how user input is transformed to a system action



A model is a collection of utterances, intents, entities, and vocabulary that the system uses to respond to natural language inputs from your users. The model takes the natural language input from your users and matches it to actions to be performed by the system.

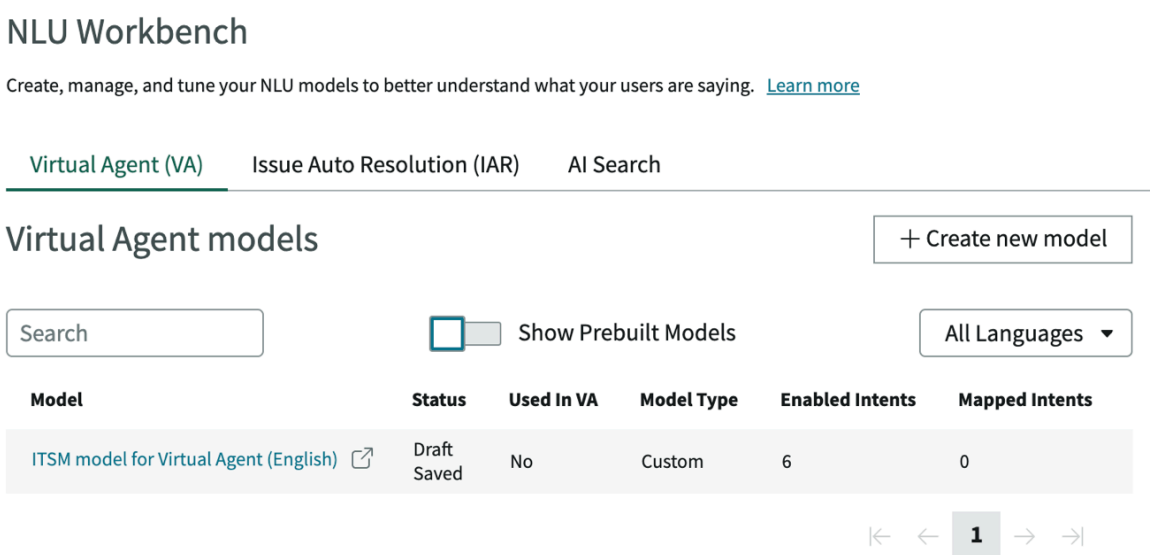
The utterances in your model data are examples of what your users might ask for. These examples are used to train your model to recognize which system actions to perform in response. These system actions are called intents.

Usage, roles, and navigation

Use the nlu_admin role to access the NLU Workbench.

Navigate to **NLU Workbench** > **Models**.

There are three tabs on the NLU Workbench homepage: Virtual Agent, Issue Auto Resolution, and AI Search. The Virtual Agent tab opens by default. Select the correct tab for your application.



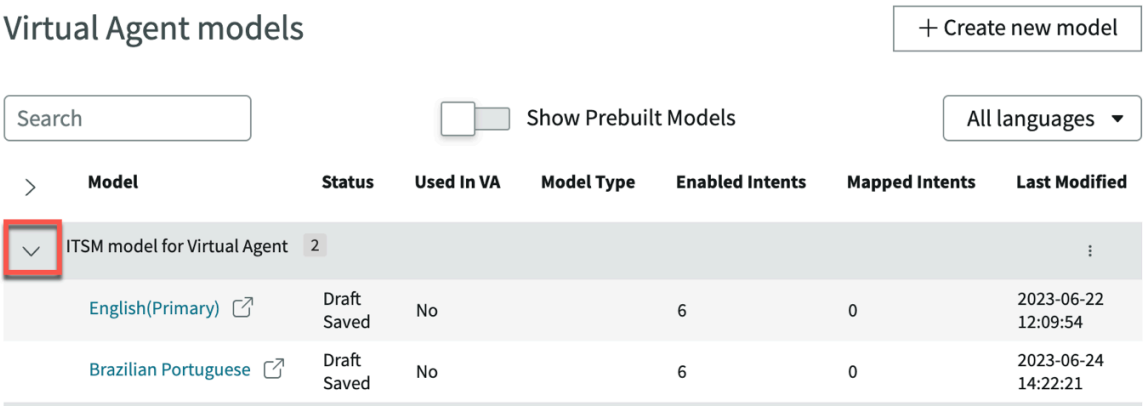
You can use the **Search** field to look for models by name.

The default setting of the **All Languages** filter displays all models regardless of their language. You can filter models by language using the list in **All Languages**. To learn more about the available languages, see NLU language support.

If you have a multilingual grouped model, its row in the list of models has an arrow on the left side. Select the arrow to expand the group so that all language versions of the model are

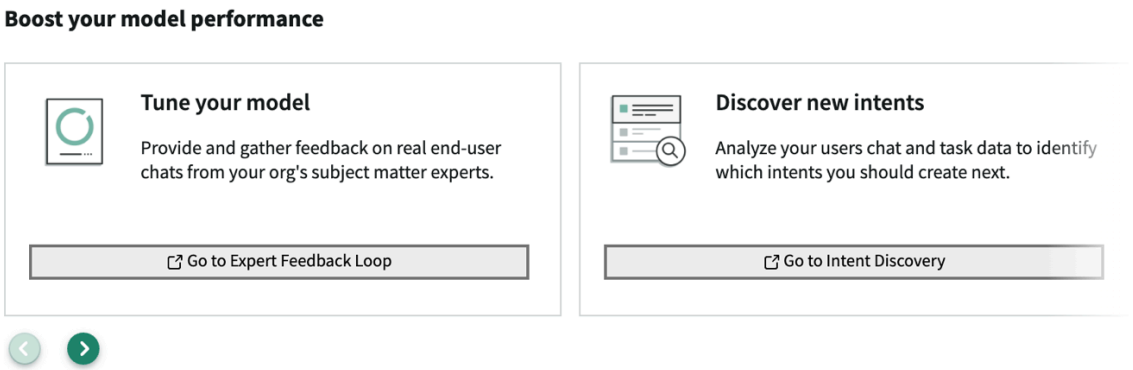
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

displayed. For more information on model languages and grouping, see Multilingual model management.



The **Used In VA** column indicates whether a model has already been linked to another application such as Virtual Agent. For AI Search, this column indicates whether the model has already been linked to a Genius Result. Similarly, the **Mapped Intents** column displays a count of intents that have been mapped or linked to Virtual Agent topics.

Scroll down the NLU Workbench homepage to view the **Boost your model performance** section. This section displays cards you can use to access functions that are available to your model type.



Model applications

The output of your NLU models can be consumed by the ServiceNow applications Virtual Agent, Issue Auto Resolution, and AI Search. The NLU Workbench homepage organizes each application's models by tab. In the list of models, different columns are displayed depending on the application.

NLU Models for Virtual Agent

For more information on Virtual Agent, see Virtual Agent.

- The tab for Virtual Agent displays a list of models. Select a model name to open an overview page for that model.
- The **Show Prebuilt Models** toggle is off by default. Switch this toggle on to display read-only models that can be copied and used as the basis for your custom models.
- The **Boost your model performance** section displays a total of 5 cards for functions available to Virtual Agent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

NLU Models for IAR

Fewer functions and columns are displayed for IAR than for the other applications in NLU Workbench. This is because the IAR model is prebuilt and is configured in the IAR Admin console.

- When you select the name of your IAR model in NLU Workbench, you are taken to IAR Tuning rather than to a model overview. For more information, see Issue Auto Resolution Tuning in NLU.
- The **Show Prebuilt Models** toggle is not displayed.
- The **Create new model** button is not displayed because the prebuilt model is used directly.
- The section **Boost your model performance** is not displayed.

For more information on setting up IAR, see Using Issue Auto Resolution.

NLU Models for AI Search

- There is one prebuilt model for AI Search. You can copy this model as the basis for a custom model.
- The **Show Prebuilt Models** toggle is on by default.
- The **Create new model** button is available.
- The **Boost your model performance** section displays two cards for functions available to AI Search.

For an introduction to managing AI Search in the NLU Workbench, watch the following video. A video on how to manage AI Search in the NLU Workbench.

For more information, see the AI Search landing page.

What you can do next

To learn more about model content and the model life-cycle, see Model management.

Model management

Manage your NLU model's life cycle in the NLU Workbench. Model management phases guide you through the iterative process of building, testing, and publishing your model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Bringing your NLU model from creation to deployment requires multiple steps, separated into phases. You can return to earlier phases when you want to adjust and maintain your model.

The phases available for your model depend on the model's application. The system will display a phase, button, or function only when it applies to your model's application.

To use the model management phases, ensure you have all the necessary NLU plugins. For more information, see the following resources:

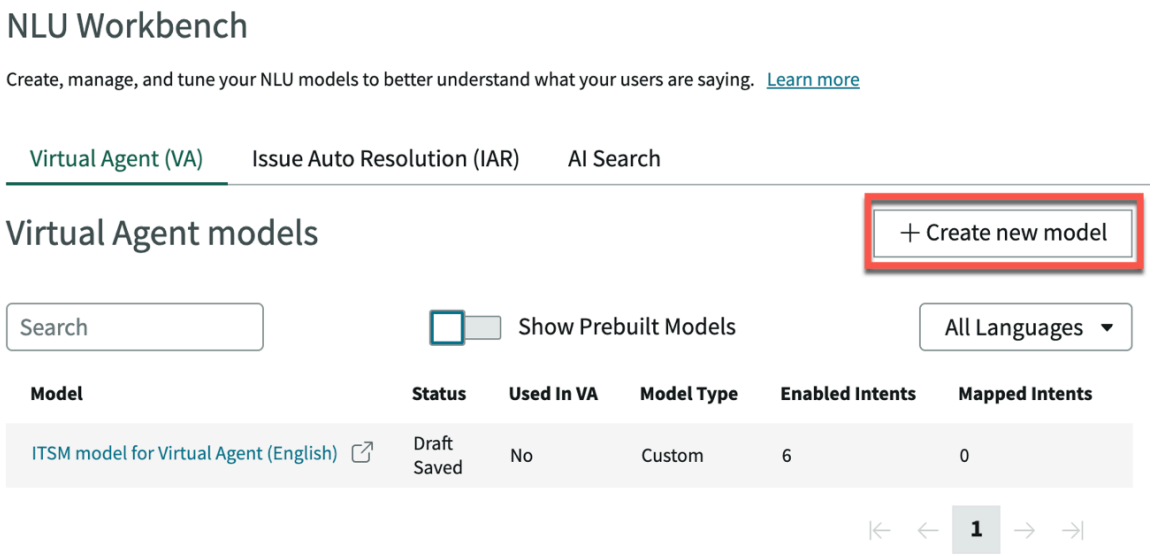
- Activate the NLU Workbench
- Install NLU Workbench - Advanced Features
- Install Intent Discovery

Note: NLU Workbench - Advanced Features and Intent Discovery are available from the ServiceNow Store. NLU model testing and performance monitoring require NLU Workbench - Advanced Features.

Create a model

To create a model for Virtual Agent or AI Search, navigate to **NLU Workbench** > **Models**. The Virtual Agent tab opens by default. Select the appropriate tab for the model you want to create.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



You can choose different ways to create a model:

- **Use prebuilt model**: Copy one of the included read-only models, and add content specific to your business.
- **Import data from CSV**: Upload a CSV file that contains training utterances and matched intents.
- **Start from blank**: Go through the process of setting up a new model from scratch.

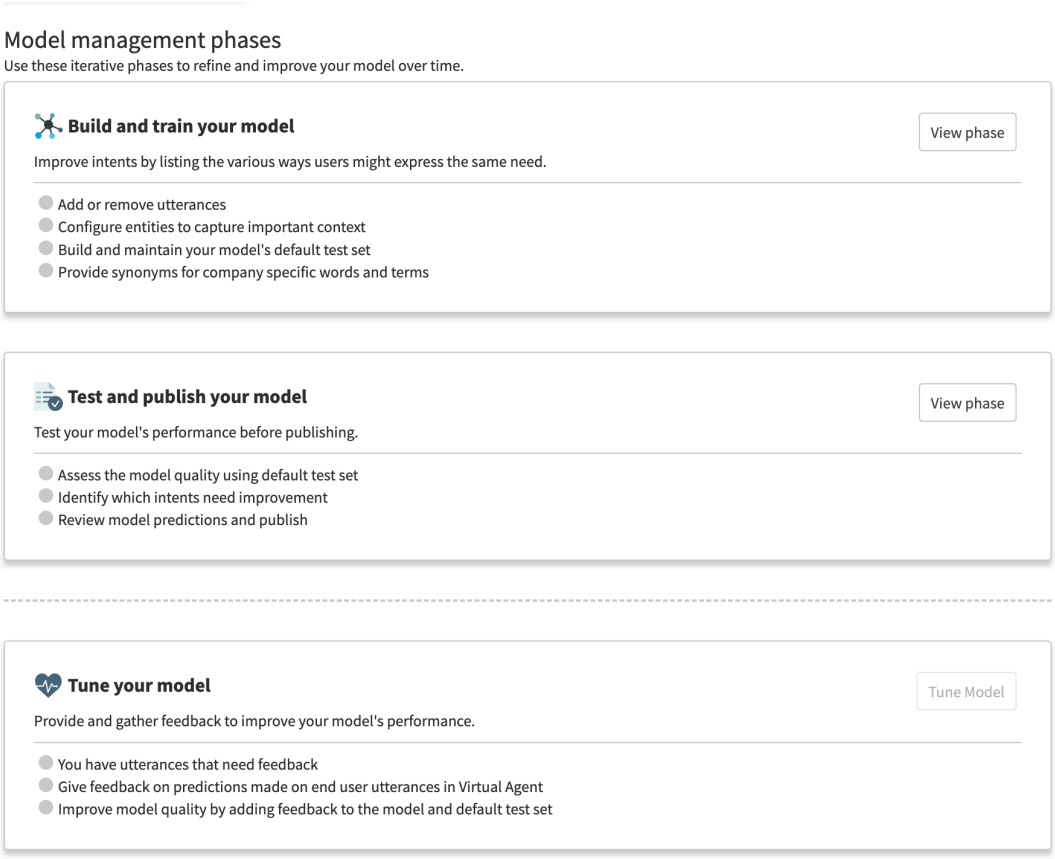
To get started, see Creating models.

Model management phases

After creating a model, access its management phases by navigating to **NLU Workbench** > **Models**. Select the tab for your model's application, then the name of the model to open the **Model details** page on the model overview.

There are three phases on a Virtual Agent model's overview page: Build and train your model, Test and publish your model, and Tune your model. These phases guide you as you build and improve your model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Build and train your model

Build the model by adding and managing content:

- Intents: Add more intents to broaden the range of user requests that your model can understand.
- Entities: Add more entities so that your model can extract more contextual details from your users' requests.
- Vocabulary: Add vocabulary to enable the model to better understand words and phrases that are specific to your business, such as industry terms and acronyms.
- Test set: Add test utterances and their expected intents to your model's default test set.

To learn more, see Build and train your model.

Train your model using utterances that the model is likely to encounter from your users. To learn more, see Train and try your NLU model.

Test and publish your model

Test your model to gauge the performance and identify areas for improvement.

Note: Model testing requires the NLU Workbench - Advanced Features store application. For more information, see NLU Workbench - Advanced Features.

For more information on testing and thresholds, see Test and publish your model.

When you're satisfied with the results of testing, publish your model to make it available for use by other applications. For more information, see Publish your NLU model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tune your model

If NLU Workbench - Advanced Features is installed, and your model is created for Virtual Agent, the **Tune your model** phase is enabled. With this phase, you can use **Expert Feedback Loop** to incorporate actual user utterances into your model.

For more information, see Tune your model.

If your model is created for Issue Auto Resolution, you will be taken to IAR Tuning by selecting the name of your model in the IAR tab of the NLU Workbench homepage. For more information, see Issue Auto Resolution Tuning in NLU.

Model settings

Use the **Settings** page of the model overview to change the name and description of the model. You can also modify the confidence threshold of the model. The confidence threshold determines how confident the model must be to predict an intent.

For more information, see NLU model settings.

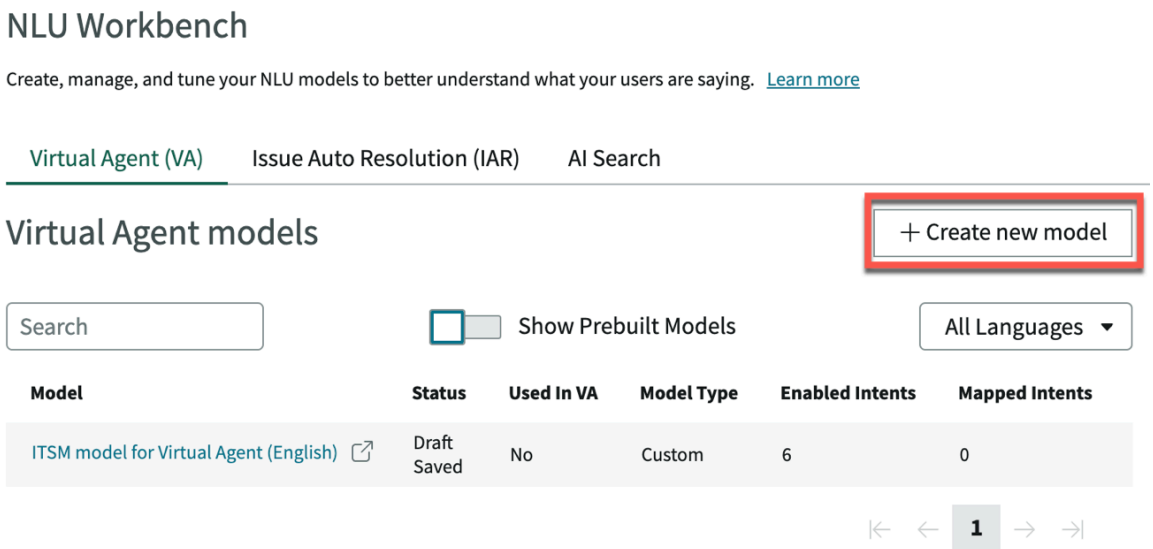
Creating models

Creating models is the first step to taking advantage of Natural Language Understanding (NLU) in your instances. Create models for Virtual Agent and AI Search in the NLU Workbench.

You create NLU models for consumption by the ServiceNow applications Virtual Agent or AI Search.

Note: For Issue Auto Resolution, a prebuilt model is provided for you to configure.

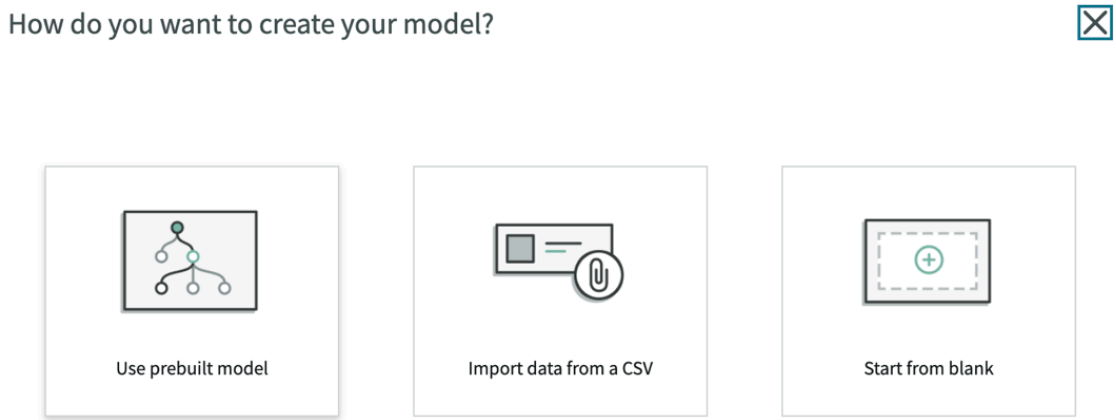
To start creating models, set your scope to the application scope you want for your new model. Then navigate to **NLU Workbench** > **Models**. The Virtual Agent tab opens by default. Select the appropriate tab for the model you want to create.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

New model creation

When you select the **Create new model** button, a modal opens to display your model creation options. Start by selecting one of the



icons:

- Create an NLU model using a pre-built model: Copy a prebuilt model and its contents as a starting point for your new model.
- Create an NLU model from a CSV file: Upload a CSV file containing a list of intents and corresponding utterances.
- Create an NLU model from blank: Build a model from scratch and add intents and utterances as you go.

After creating, add content to your model. The intents, utterances, entities, and vocabulary you add helps improve the model's ability to interpret natural language. See Build and train your model.

Duplicating, exporting, and updating

After creating a model, you have options to use that model across other models and instances. With the NLU Workbench, you can perform the following actions with your models:

- Duplicate an NLU model: Copy a model to create a model with the same content.
- Export an NLU model: Export a model as a CSV file containing the associated utterances and intents. Share the model or use it to create one.
- Add an NLU model to an update set: Add a model and its artifacts to an update set to transfer the model across instances.

Create an NLU model from blank

Create a Natural Language Understanding (NLU) model from scratch. Start with an empty model to have full control over the model's content.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- You can create NLU models for Virtual Agent and AI Search.
- Role required: nlu_admin or admin

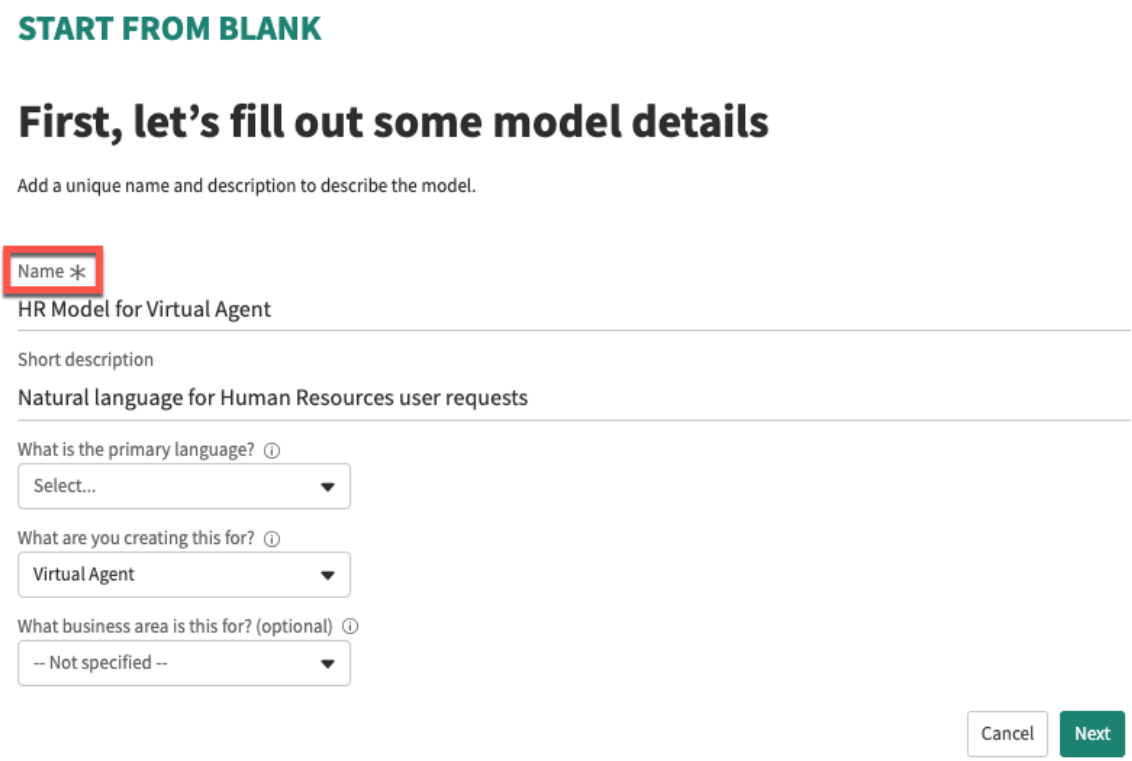
About this task

In this example procedure, you're building an NLU model to help Virtual Agent understand human-expressed intent regarding Human Resources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Set your scope to the application scope you want for your model.
2. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
3. Select the tab for the type of model you want to create, such as AI Search.
4. Select the **Create new model** button.
5. In the **Add some details** window, select the **Start from blank** button.
6. In the define details page, enter a unique **Name** and **Short description**.

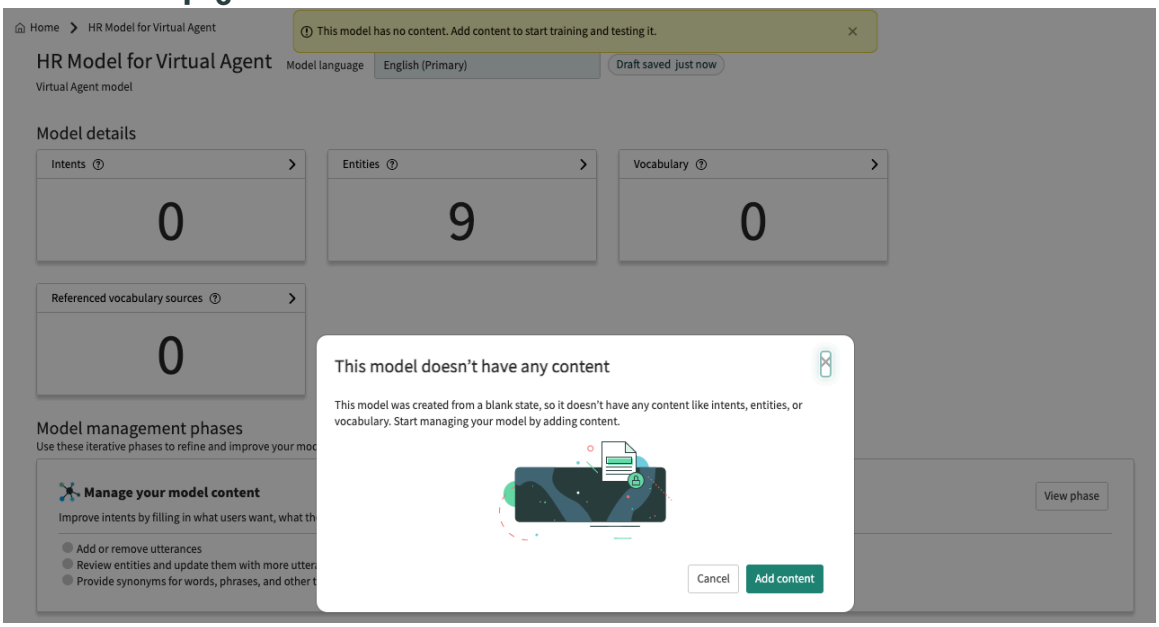


   In this example scenario, you enter `HR Model for Virtual Agent` for the name and `Natural language for Human Resources user requests` for the description.
7. Select the language and purpose from the drop-down lists.
   In this example scenario, select `English` and `Virtual Agent`.
8. **Optional:** Select a business area for your model.
9. Select **Next**.
   Your model starts building. When complete, select **View model** to open the **Model details** page for your new model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

Model details page



Your new model contains no content. Select **Add content** to begin adding intents, entities, and vocabulary. For more information, see Build and train your model.

Add test utterances and intents to build the model's default test set. For more information, see Test set creation and management.

Create an NLU model from a CSV file

Upload a CSV or XLSX (Excel Workbook) file containing utterances and their intents to create a Natural Language Understanding (NLU) model. Use this method to quickly create models from your data or other exported models.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- You can create NLU models for Virtual Agent and AI Search.
- Role required: admin or nlu_admin

About this task

In this example procedure, you're building an NLU model to help Virtual Agent understand user requests regarding calendars.

In a CSV file, you've listed the intents and utterances in two columns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Sample CSV setup

Intent	Utterance
schedule	when is my next meeting
schedule	what time is my next meeting
schedule	schedule a meeting for 11am
schedule	cancel my 4:30 meeting
holiday	when is the next holiday
holiday	do we get Black Friday off
timeoff	request friday off
timeoff	how many PTO days do i have

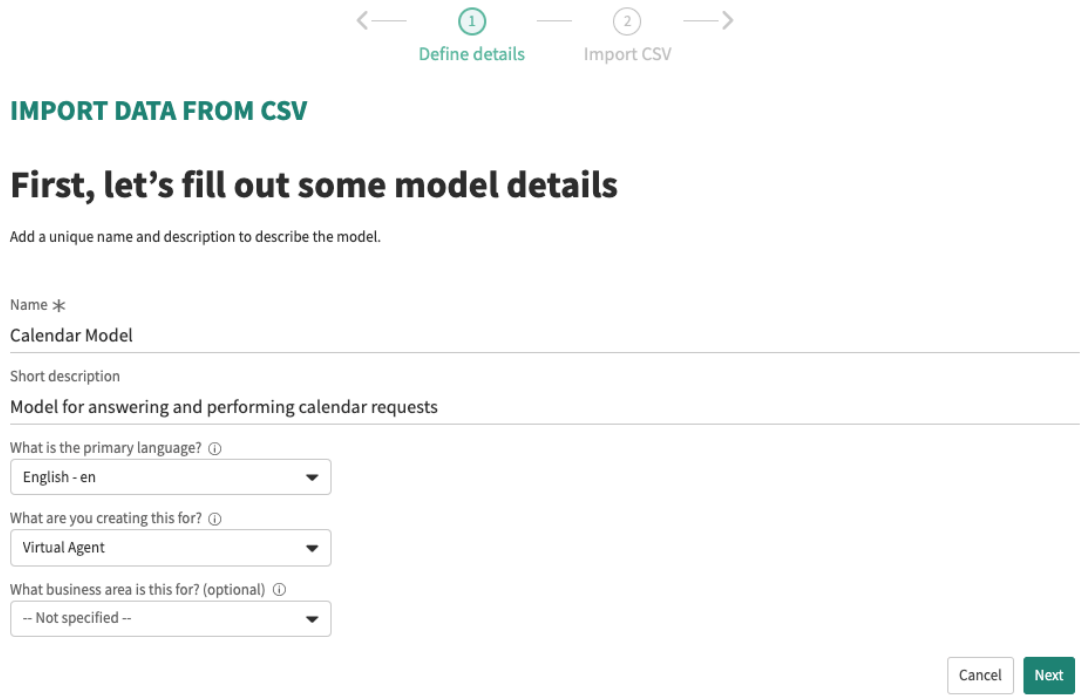
Note the following for creating NLU models by CSV import:

- A model needs at least 1 intent with a minimum of 5 training utterances in each intent. For optimum performance, aim to have 15 training utterances per intent.
- Utterances should not contain a comma.
- Importing with a CSV file does not preserve entities. Make sure to annotate utterances as needed after importing.

Procedure

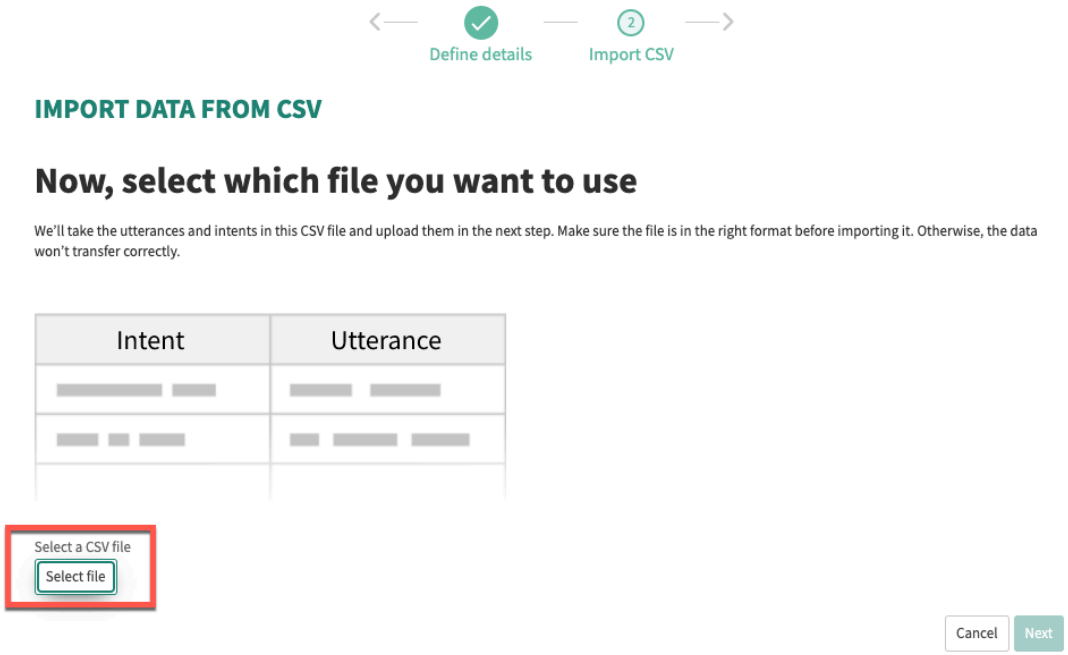
1. Set your scope to the application scope you want for your new model.
2. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
3. Select the tab for the type of model you want to create, such as AI Search.
4. Select the **Create new model** button.
5. In the **How do you want to create your model?** window, select **Import data from a CSV**.
6. In the **Add some details** window, add the **Name** and **Short description** for the model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



In this example scenario, you enter `Calendar Model` for the name and `Model for answering and performing calendar requests` for the short description.

7. Select the language and purpose from the drop-down lists.
   In this example scenario, you select `English` and `Virtual Agent`.
8. Click **Next**.
9. On the **Import CSV screen**, click **Select file**.



10. Choose the CSV or XLSX (Excel Workbook) file from the pop-up.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

11. Select **Next**.
    Your model starts building. After completion, select **View model** to open the model details page.

What to do next

Add intents and training utterances to continue building the model. Add entities and vocabulary to help the model understand inputs from your users. For more information, see Build and train your model.

Add test utterances and intents to build the model's default test set. For more information, see Test set creation and management.

Create an NLU model using a pre-built model

Use the prebuilt Natural Language Understanding (NLU) models in the NLU Workbench to create a new one.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- You can create NLU models for Virtual Agent and AI Search.
- Role required: nlu_admin or admin.

About this task

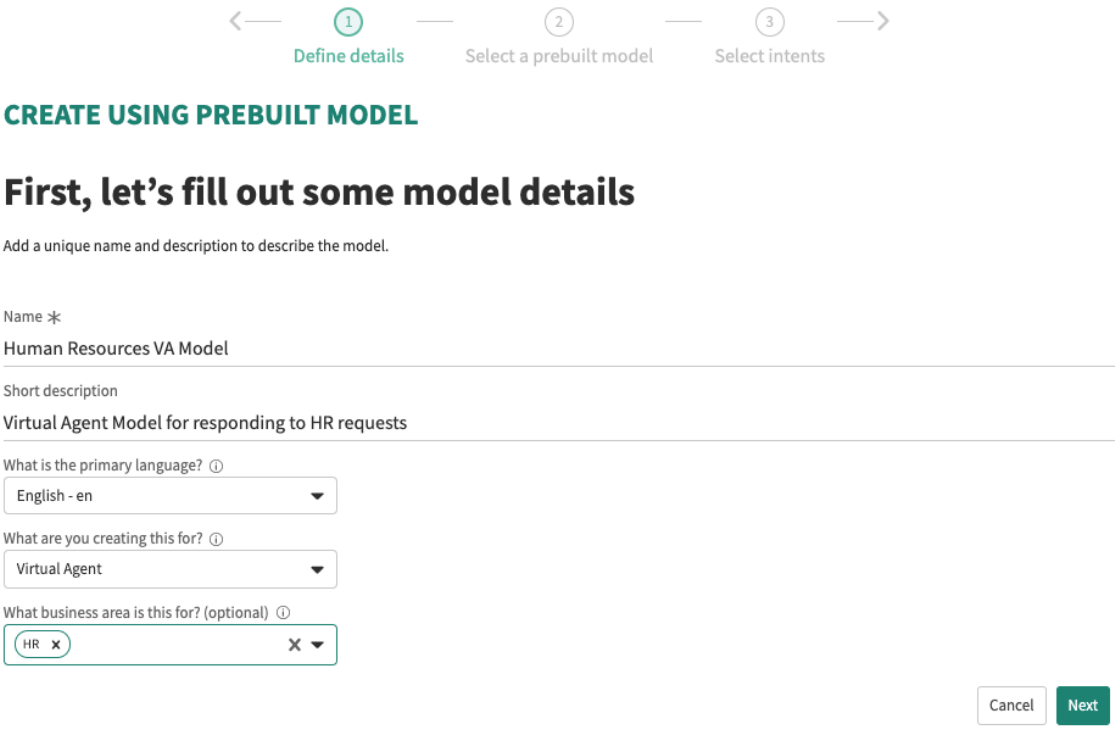
In this example procedure, you're building an NLU model based on the prebuilt HR NLU for VA model. You base your model on the prebuilt model, then edit or add to the model content to meet your business requirements.

Your new model contains intents, entities, and vocabulary from the prebuilt model. The new model also contains an empty default test set, ready to be populated with test utterances.

Procedure

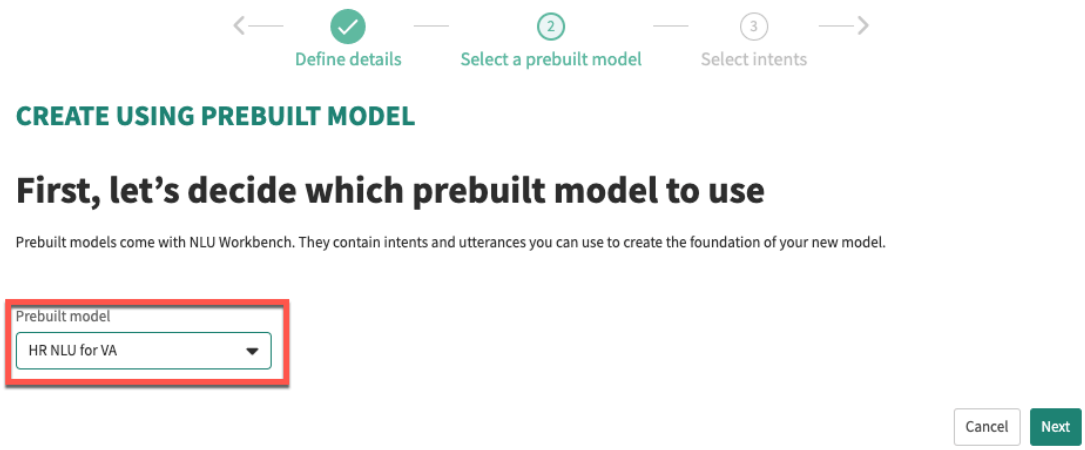
1. Set your scope to the application scope you want for your new model.
2. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
3. Select the tab for the type of model you want to create, such as AI Search.
4. Select the **Create new model** button.
5. In the **How do you want to create your model?** window, select **Use prebuilt model**.
6. On the **Add some details** window, fill in the **Name** and **Short description** for the model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



In this example scenario, you enter `Human Resources VA Model` for the name and `Virtual Agent Model for responding to HR Requests` for the short description.

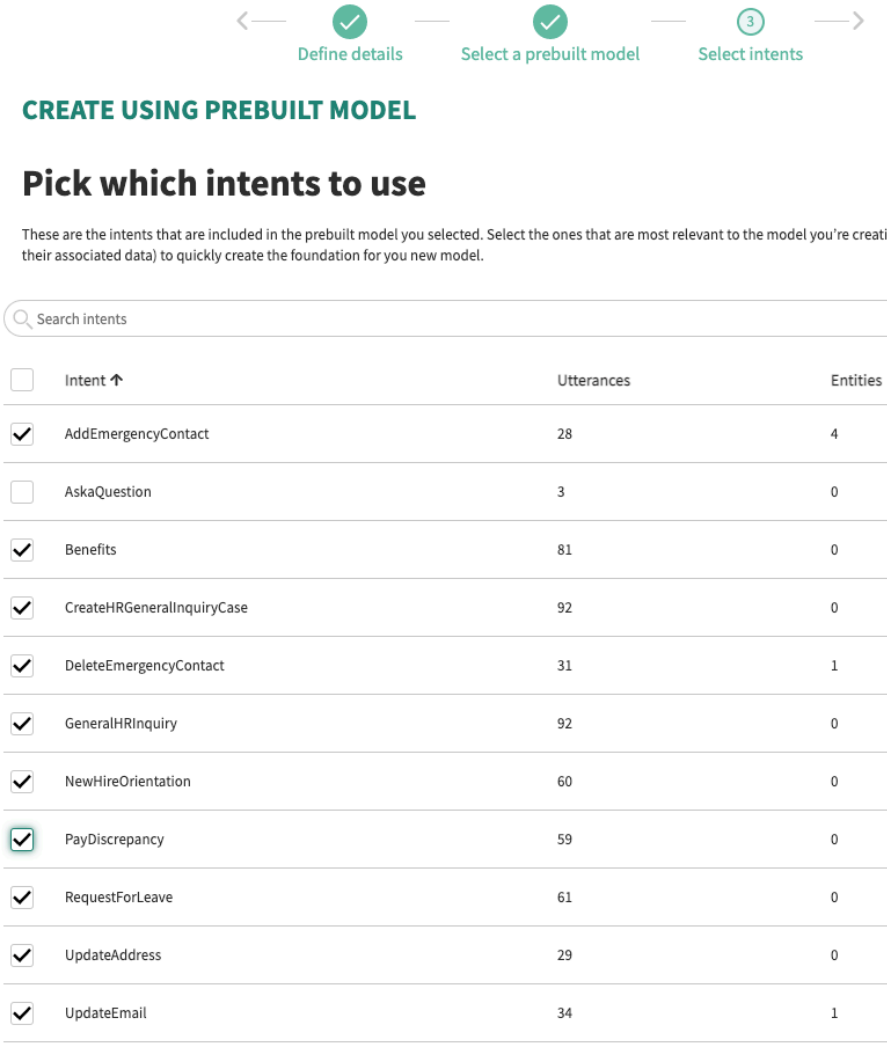
7. Select the language and purpose from the drop-down lists.
   In this example scenario, you select `English` and `Virtual Agent`.
8. **Optional:** Select the business area for the model.
   In this example scenario, you select HR.
9. Click **Next**.
10. Select the prebuilt model from the drop-down list.



In this example scenario, you select `HR NLU for VA`.

11. Select **Next**.
12. On the **Select intents** screen, select the prebuilt intents to add to your model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



In this example scenario, you select 10 intents from the list.

Note: You can select all the prebuilt intents by checking the **Intent** box at the top of the list.

13. Select **Next**.
    Your model starts building. After completion, select **View model** to open the **Model details** page for your new model.

What to do next

Add more intents and training utterances to continue building the model. Update entities and vocabulary to help the model understand inputs from your users. For more information, see Build and train your model.

Add test utterances and intents to build the model's default test set. For more information, see Test set creation and management.

Duplicate an NLU model

Duplicate an existing Natural Language Understanding (NLU) model to create a new one. Duplicating a model copies the settings and contents of the original model, including its default test set.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Identify an existing NLU model to copy from. The source model should have the language and purpose you want for the new model. If a suitable model does not exist, create and train one.
- Role required: nlu_admin or admin.

About this task

In this example procedure, you've already created a model that's titled NLU for Access Requests. You want to customize it for another department. You duplicate the model to create a new, separate model. To customize the new model, you then add content related to the other department.

Procedure

1. Set your scope to the scope of the model you want to create.
   After creation, these model settings cannot be changed: language, purpose, scope.
2. Navigate to **All** > **NLU Workbench** > **Models**.
3. On the far right column of the model list, select the **More options** menu for the model you want to duplicate.



4. Select **Duplicate this model**.
5. In the **Duplicate this Model** window, enter a name and description for the new model.



   In this example scenario, you enter `NLU for Access Requests Copy1` for the description.

6. Click **Duplicate**.
   The system duplicates the model. The **Model details** page loads for your new model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

You can customize the new model by adding new content to it. Train, test, and publish the new model. For more information, see Build and train your model.

To duplicate an entire model group, or one of the models in a group, see Multilingual model management.

To transfer a model to a different instance, see Add an NLU model to an update set.

Export an NLU model

Export a Natural Language Understanding (NLU) model to create a CSV file of the intents and utterances. You can then use the CSV file to edit, share, and import.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Role required: nlu_editor, nlu_admin, or admin

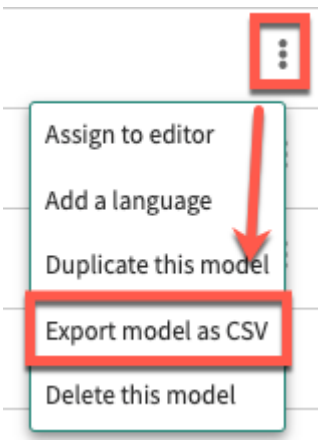
About this task

Exporting an NLU model creates a CSV file. The file contains a table of the utterances from the model and the matched intents. The data comes from the **Utterances** tab for each intent in the model.

Note: The file doesn't contain the sources of the utterances. Also the file doesn't transfer entities or vocabulary associated with the model. To export all model data, see Add an NLU model to an update set.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, and find your model in the list.
3. Scroll to the far right column of your model's row, and select the more options menu of the model you want to export.



4. In the more options menu, click **Export model as CSV**.
   The model CSV file downloads to your browser.

What to do next

You can use the CSV file to share the model or edit the utterances. You can also create a model by importing the CSV file. For more information, see Create an NLU model from a CSV file.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Add an NLU model to an update set

Use update sets to move your Natural Language Understanding (NLU) models from one instance to another. Update sets include all records needed for your model to function on the target instance.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- The following instructions apply to NLU models for Virtual Agent and AI Search.
- Role required: admin

About this task

Use update sets to transfer a model from a source instance to a target instance. The target instance must already have the same scope as the model's scope in the source instance.

Adding a model to an update set includes the following resources:

- Model content (intents, utterances, entities, annotations, vocabulary, and the default test set)
- Associated vocabulary sources
- Corresponding latest active ML solution
- ML model artifacts
- ML solution and definition (last three runs, provided one of them was successful)

For more information on update sets, see System update sets.

If you want to add the model to a different update set, you must make that update set the current one. See Create and select an update set as the current set.

For optimum portability, add your NLU model to a new, dedicated update set rather than to a system Default update set. To do this, a new update set must be manually created in the model's scope. An error is displayed if the current update set is the scope's Default. This error contains a link for manually creating a new update set.

If the model contains records from multiple scopes, such as vocabulary tables, its update set must have a parent-child (batch) structure. Follow the procedure from step 5 to create a parent-child update set.

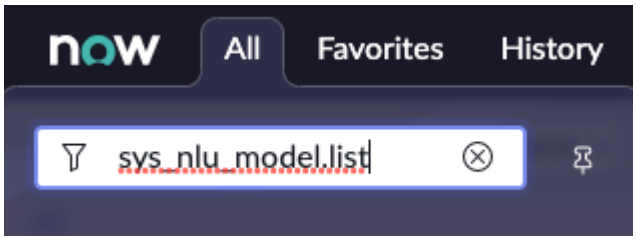
When models (either Global or scoped) are moved using update sets, their training and publishing state remains the same after transfer. So a model that is trained and published before it is added to an update set does not require retraining or republishing on the target instance.

For information about parent-child update sets, see Working with batched update sets. The following video provides a demonstration of this process. A video walkthrough of migrating NLU models by update sets in the Vancouver release.

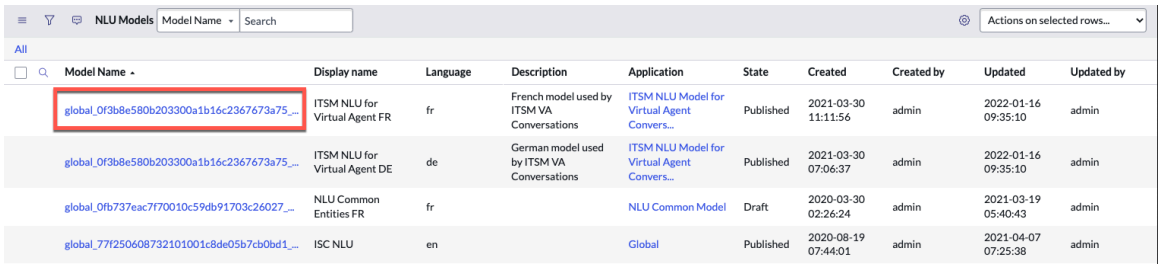
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Select **All** and enter `sys_nlu_model.list` into the navigator.

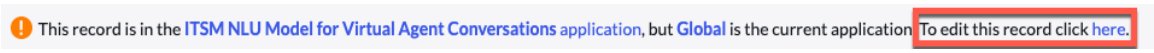


2. Select the **Model Name** of the model.

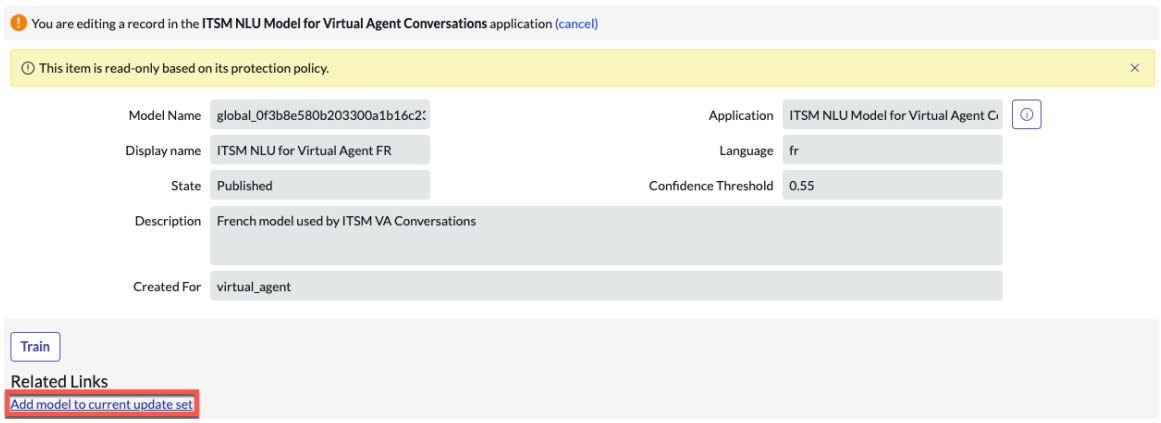


The **Display name** is the name you gave the model when creating it. The **Model name** is given to the model by the system.

3. **Optional:** If Global is your current application, follow the prompt at the top to edit the record.

(Optional) 

4. In the **Related Links** section, click **Add model to current update set.**

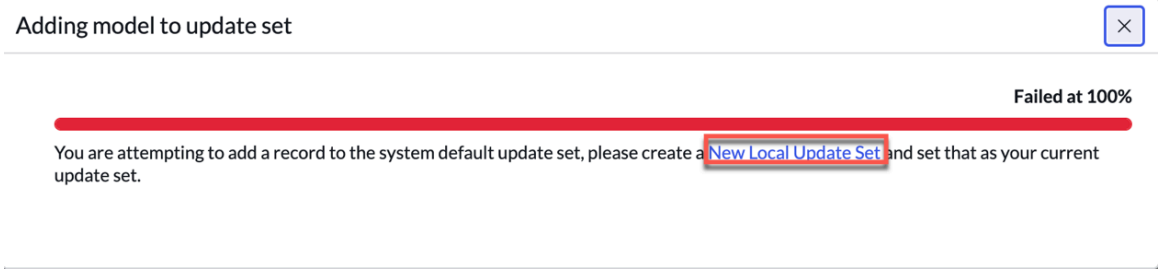


If the model's records are all in one scope, the system adds the model to the current update set. Once completed, the record can be found in the sys_update_set table.

However, the system displays an error if the designated scope is the Default scope. Continue with the following steps.

5. If the system displays the error **You are attempting to add a record to the system default update set**, continue with the remaining steps in this procedure.

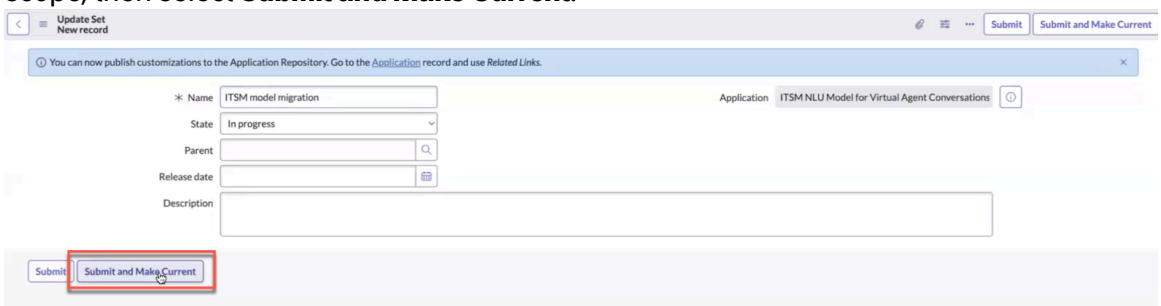
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



In the error banner, select the **New Local Update Set** link to create a new update set which is not Default. In the case of a parent-child update set, the parent is first created in the model's scope. (Any child update sets are created in a later step).

New Local Update Set opens a new record in the sys_update_set table.

6. For the new update set record, provide a name, review other values including Application scope, then select **Submit and Make Current**.



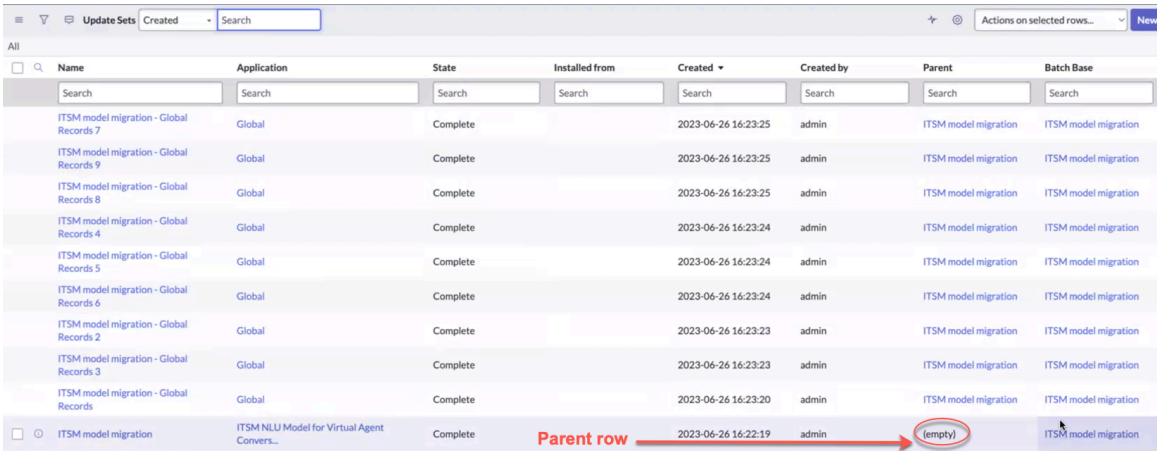
A new update set is created, and the screen re-opens to the model's record in the sys_nlu_model table.

7. In the model's record in sys_nlu_model, select **Add model to current update set** under **Related Links**.
When the model's records are in multiple scopes, this step creates child update sets that have Global scope.

> **Note:** If the model's records are all in one scope, the system adds the model to the current update set. Once completed, the record can be found in the sys_update_set table.

8. If the model is in a non-Global scope: In the Update Sets table sys_update_set, locate the parent update set and its child update sets (refresh the page if necessary).
   - The parent update set has the value **(empty)** in the Parent column.
   - Child update sets have the name of the parent in the Parent column, and **Global** in the Application column.
   - Both parent and child update sets contain the same value in the Batch Base column.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



9. Open the parent's record in sys_update_set and set the **State** field to `Complete`.
   A Confirmation dialog box pops up. Select `Yes` to confirm. This sets the parent and all child update sets to Complete.
10. To migrate the update set file, select **Export Update Set Batch to XML** in the parent's record in sys_update_set.
    The parent and child update sets are packaged together and exported.

What to do next

In the target instance, navigate to **Retrieved Update Sets** and select **Import Update Sets from XML**. When unpackaged, the parent and all children will be listed in the Retrieved Update Sets table. Open the parent and select **Preview Update Set Batch**.

If a referenced record in the source instance is not present on the target instance, you may encounter errors while applying the update set. For example, the origin field on sys_nlu_intent could refer to an intent from a different model that doesn't exist on the target instance. You can click **Accept remote update** on the failed records to commit the update set anyway.

Delete an NLU model

Delete a Natural Language Understanding (NLU) model permanently.

Before you begin

Role required: admin or nlu_admin

About this task

Deleting an NLU model removes the model and its contents, including its default test set, from the system. You cannot delete a model if it meets any of the following criteria:

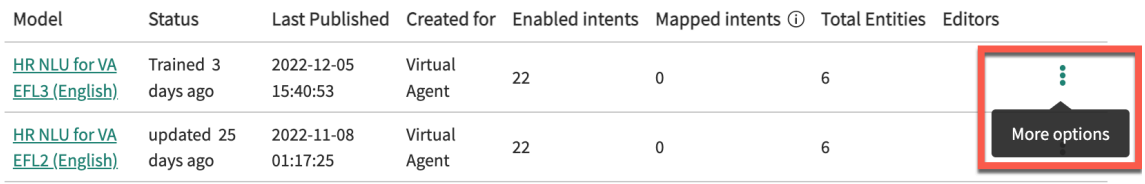
- The model is pre-built
- The model contains at least one intent that is mapped to a Virtual Agent topic

> ⚠ **Warning:** Deleting a model cannot be undone.

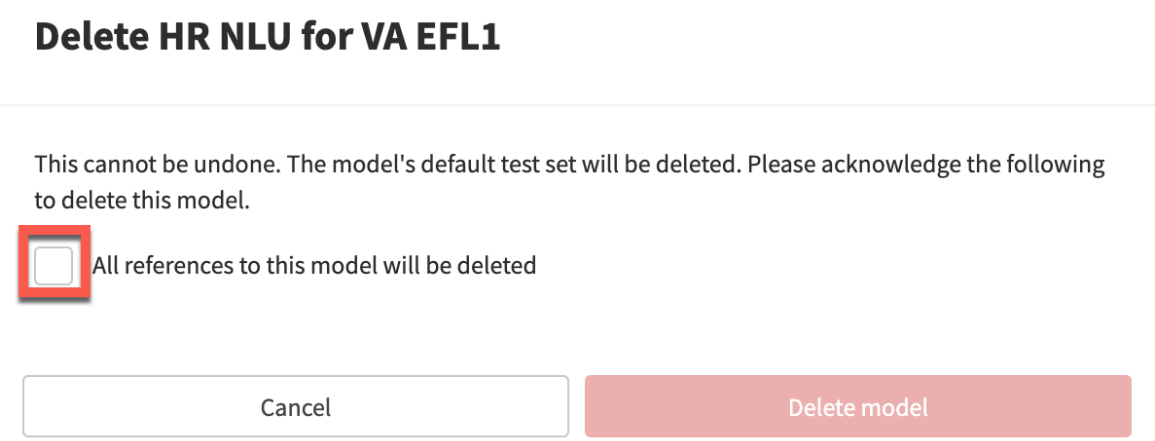
Procedure

1. Set your scope to the scope of the model you want to delete.
2. Navigate to **All** > **NLU Workbench** > **Models**.
3. On the far right column of the model list, select the more the **More options** menu for the model you want to delete.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. Select the option to **Delete this model**.
5. Select the check box to acknowledge that all references to this model will be deleted.



6. Select **Delete model**.
   The system deletes your model and reloads the page.

Build and train your model

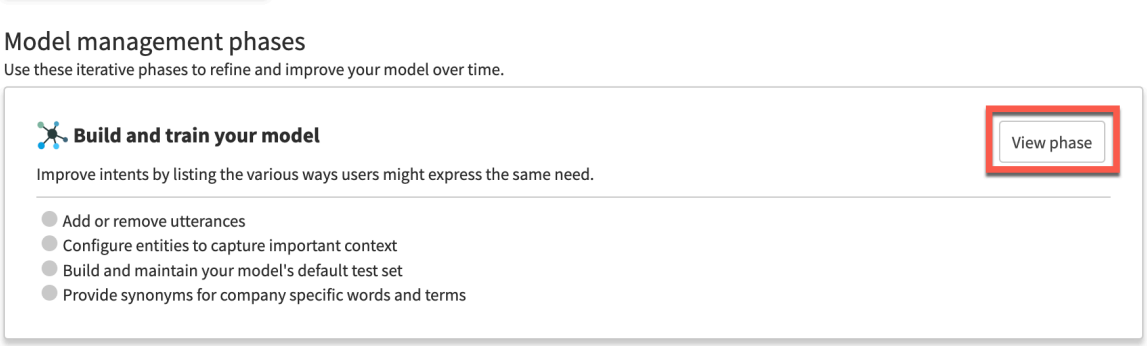
After creating a model, build the model's content by adding intents, entities, vocabulary, and test set utterances. Your NLU model content determines how the model responds to user inputs.

Models are made up of the following content:

- Intents: An action the user wants to do or wants the application to do.
- Entities: Object or context for an action.
- Vocabulary: Add vocabulary to help your model understand the range of words in your users' utterances.
- Test set: To assess model performance, add test utterances and the intents that you expect to be predicted for those utterances.

To access the model content, navigate to **NLU Workbench** > **Models**. The Virtual Agent tab opens by default. Select the tab for your model's application and then select the name of the model to open the **Model details** page. In the **Build and train your model** card, select **View phase**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Intents

When your model receives user input, it uses an intent to perform a system action. For example, a user types in `I have a critical issue with a slow laptop`. The model matches the utterance input to the intent #TroubleshootSlowComputer. If the intent is linked to a Virtual Agent topic, it triggers further action.

Intents contain training utterances, or examples of user inputs that would trigger the system action. Provide realistic utterances that the model might encounter from your users. The quality of training utterances affects the accuracy of your model.

For more information, see NLU intents.

Entities

Your intents use entities to provide additional context for the model when receiving inputs. In the computer example, the laptop is the entity, or object of, the action.

NLU entities fall into two categories: system-defined and user-defined. System entities such as DATE, TIME, and LOCATION are available by default in your instance. You can create your own user-defined entities to provide additional associations and meaning for your business requirements.

All entities are reusable across other NLU models. However, you must add them to a training utterance for each model to use them.

For more information, see NLU entities.

Vocabulary

Your users' input may contain a wide variety of words and phrases. Also, your model may not understand some terms used in specialized domains or business areas.

To improve your model's ability to understand a wide range of user input, you can define synonyms by creating vocabulary items.

For example, your model includes an entity for the term computer. When a user types in `I need a new computer`, the model knows how to respond. However, if a user enters `laptop` or `workstation`, the model might fail to predict the intent. You can add vocabulary to the model to train it to understand synonyms and variations.

You can also use tables and lists as vocabulary sources. Your models can look up the vocabulary sources when predicting intents.

For more information, see NLU vocabulary.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Test set

Your model contains a default test set that you can use to evaluate the model's performance. Initially the test set is empty, ready to be populated with your content. Add test utterances and their expected intents to build the test set.

For more information, see Test set creation and management.

Test panel

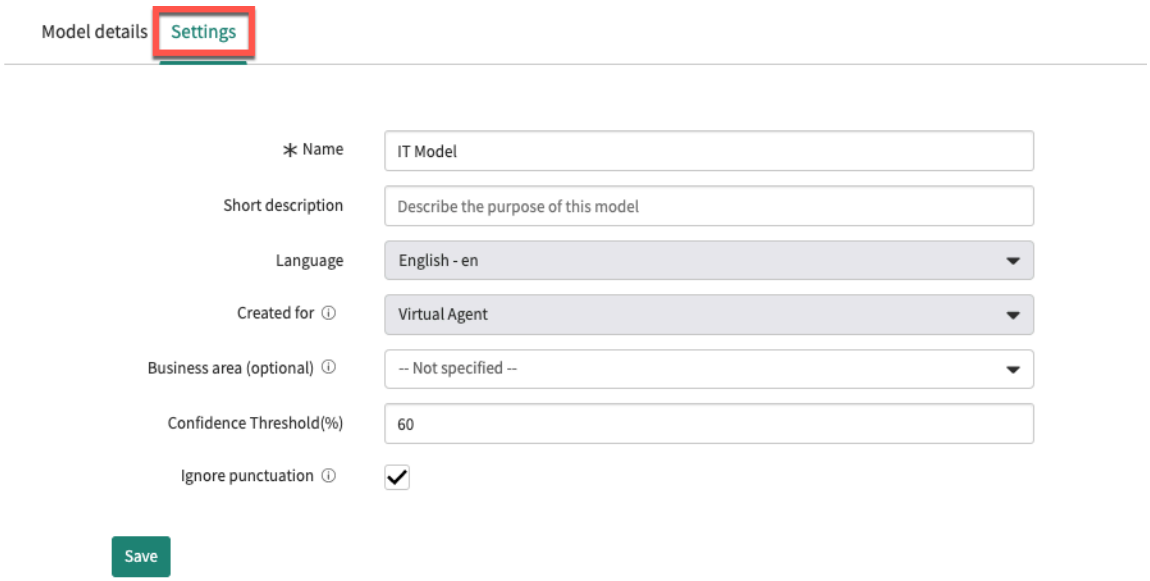
Access the test panel by clicking **Train model** or **Try model** in the **Build and train your model** phase. Training incorporates new content into your model. With **Try model**, you can manually enter individual utterances to test what intents the model predicts for them.

For more information, see Train and try your NLU model.

You can also use the test panel to provide feedback on your model's predictions. Your feedback helps improve intent prediction. See Test panel feedback.

Settings

Use the **Settings** tab to edit the name, short description, and confidence threshold of the model. You can't change the language or purpose of the model.



For more information on the confidence threshold, see Test and publish your model. For more information on Settings, see NLU model settings.

NLU intents

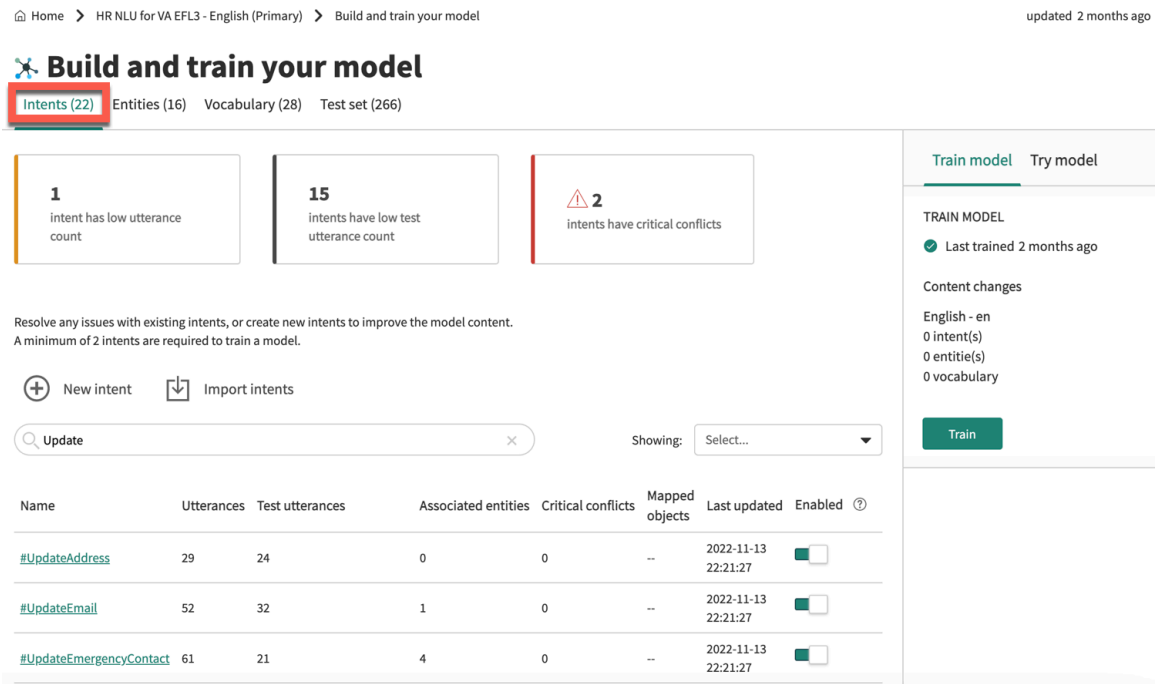
Intents drive your models' responses by matching a system action to user inputs. Models with good intents help Virtual Agent and Search respond to your users accurately.

Think of intents as the core that drives the comprehension of natural language. Intents help the models translate utterances from your users into actions that the system can perform. Intents perform best if your utterances, entities, and vocabulary work together to support your model. Using vocabulary and realistic utterance examples can help the system to be more accurate when predicting intents.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When creating your Natural Language Understanding (NLU) model, you add intents to match user inputs with system actions. The more intents that a model has, the more actions it can take when it receives utterances from your users. Model intents can perform various actions, from creating hardware service requests to adding users to groups.

To access your model's intents, navigate to a model's overview page. In the **Build and train your model** phase, select **View phase**. The **Intents** tab shows by default.



A model can contain up to 750 intents. However, models with over 300 intents or 4500 utterances, whichever comes first, take longer to train, test, and publish.

Note: As you create intents, keep in mind that they can sometimes impact each other. For example, you build and test an intent that works on one or two utterances. But when you test it in a larger intent environment, it might behave differently. To reduce the likelihood of such an event, you may want to create at least five intents in a model before you start proper testing.

To start adding content to your model, see Create an NLU intent.

The **Enabled** column shows whether or not the intent is active in predictions. An NLU admin can deactivate an individual intent but keep it in the model.

However, if an intent is mapped to a published Virtual Agent topic:

- You cannot deactivate the intent.
- You cannot delete the intent.

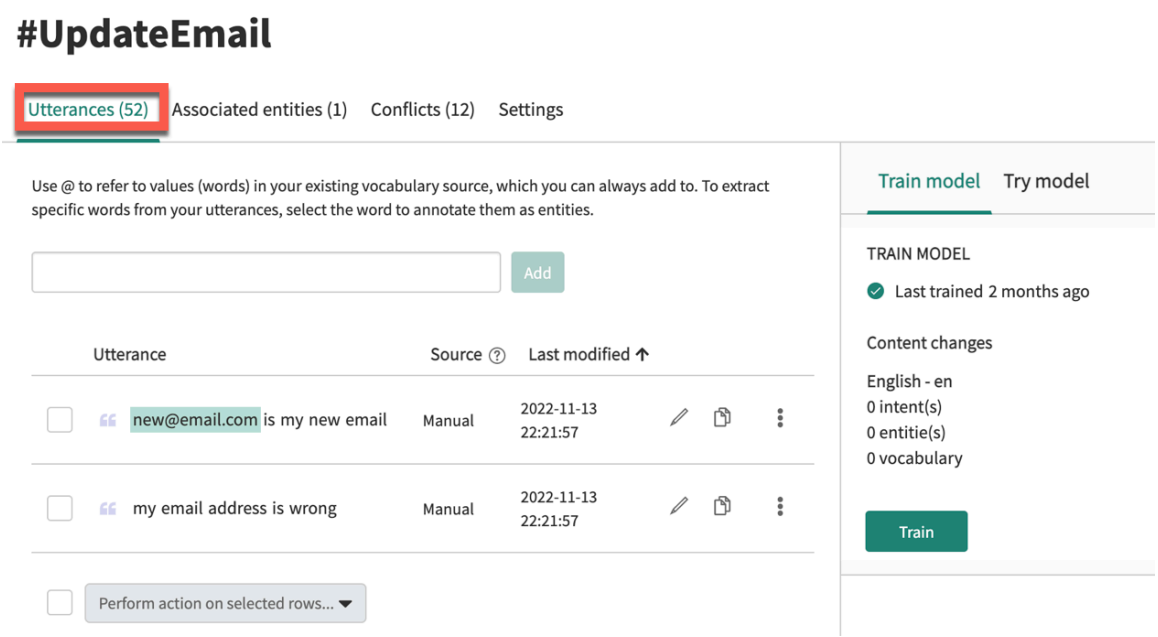
After changing the **Enabled** status of an intent, retrain the model.

Utterances

Intents contain training utterances, which are examples of inputs that the model may see from your users. Each intent in a model has its own utterances. When trained, the model learns to recognize similar utterances from your users and then respond with the matching intent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Select an intent to access the intent details page. The **Utterances** tab lists all the utterances currently in the intent.



After adding utterances, you can edit, copy, move, or delete the utterances using the icons in the right column. You can move or delete multiple utterances by selecting the box on the left first and using the **Perform action on selected rows** button.

Here are some things to consider when adding utterances to intents:

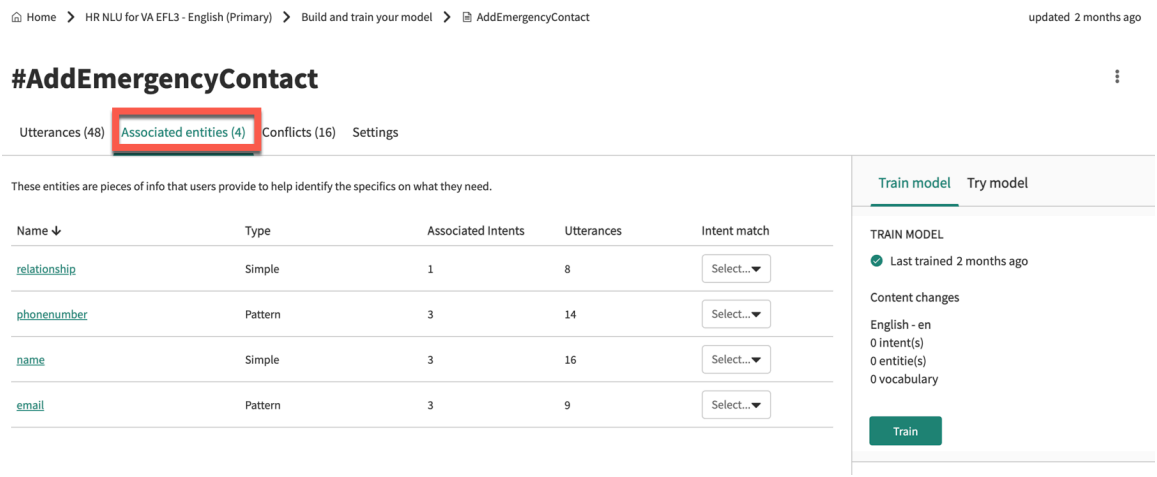
- A model must have at least 1 intent, with a minimum of 5 utterances in each intent.
- An intent needs at least 5 utterances to begin training.
- The system currently supports utterances up to 25 words or 200 characters. Utterances that exceed that limit fail to return an intent prediction.
- The system currently supports up to 20,000 utterances in a single model.
- Models with more than 4,500 utterances take longer to train, test, and publish.

Provide vocabulary for any words or phrases that are relevant to your organization or domain when adding utterances to your model. The vocabulary helps with intent prediction for words and phrases that your users are likely to use. You can use the @ symbol when adding an utterance to call on a vocabulary source. For more context and examples, see NLU vocabulary.

Associated entities

Your model uses entities to provide additional context and meaning when predicting user input. You add entities to the training utterances of your intent to provide the system with more information to perform the intended action.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

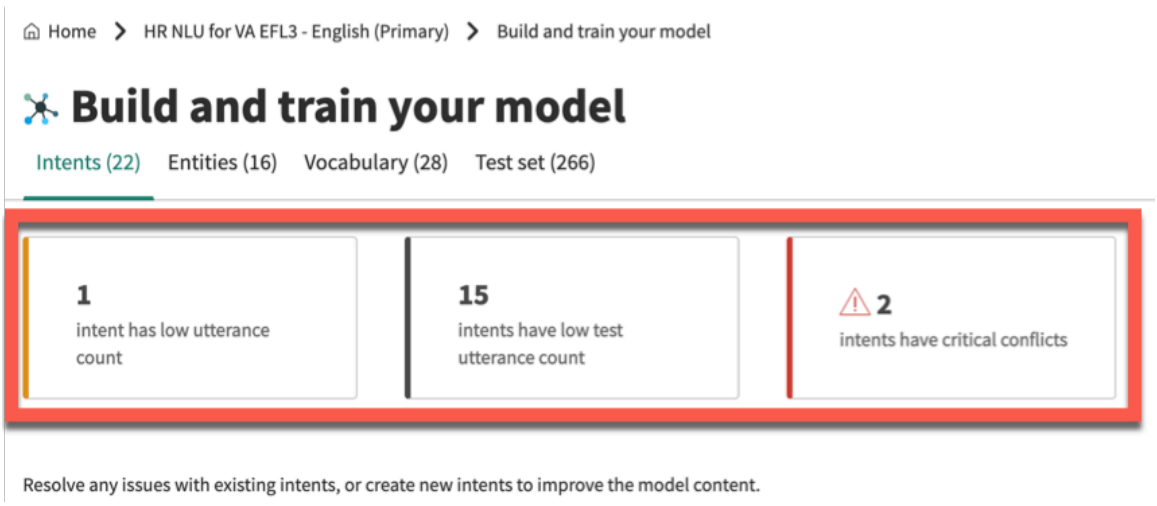


For more information, see NLU entities.

Intent issues

Building large models increases the chance that intents overlap, conflict, or fail to contain enough training utterances. For example, the utterance examples in one intent may end up identical to the examples in another intent. If your intents conflict, the model may not know which intent to predict when receiving user input.

If your model has issues or conflicts, the **Intents** page displays cards showing the number of intents affected.



Select a card to see a filtered list of intents with that issue. Resolving issues ensures that your intents meet the requirements and work as intended.

For more information, see Resolve intent issues.

Create an NLU intent

Create an intent for your Natural Language Understanding (NLU) model. Intents provide your model with a system action to perform when it receives user input.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Workbench - Advanced Features plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Ensure that you are in the same application scope as your model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- You can create intents for Virtual Agent and AI Search models in NLU Workbench.
- Role required: admin or nlu_admin

About this task

This procedure shows you how to create an intent. To reuse intents from other models, see Reusing intents from prebuilt NLU models.

Here's an example of how intents can interact with the vocabulary in their training utterances.

- Intent: #AddMembersToDistributionList
- Utterance A: "Please add Carlos Santana to the uxinfodev list"
- Utterance B: "I'm mistakenly removed from the arlo-drury-directreports group"
- Result: The system doesn't recognize uxinfodev or arlo-drury-directreports and can't use these words to predict the intent.
- Solution: Add uxinfodev and arlo-drury-directreports as vocabulary items and add synonyms to them. The synonyms you provide help add more context to the utterance and the intent in which they reside. Your intent prediction confidence can be even higher if you also mark them as entities.

Note: Training utterances and utterances from your users have a limit of 25 words or 200 characters. Utterances that exceed that limit fail to return an intent prediction.

Do not include unrealistic terms such as "OrderLaptop" or "sfsdfasdfas" in training utterances. Utterances should be correct and natural examples in the model's language.

In the following example procedure, you're creating an intent and adding utterances that users might say when requesting information on payment. You've already created an NLU model that you've titled HR Model for Virtual Agent and you're creating an intent in that model.

Procedure

1. Set your application scope to your model's scope.
2. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
3. Select the tab corresponding to your model's application, then select the name of the model you want to add an intent to.
   In this model example, you select your HR Model for Virtual Agent model.
4. In the **Build and train your model** card, select **View phase**.
5. Select **Intents**.
6. Select **New Intent**.
7. In the **Create an intent** window, add a name and description for the intent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



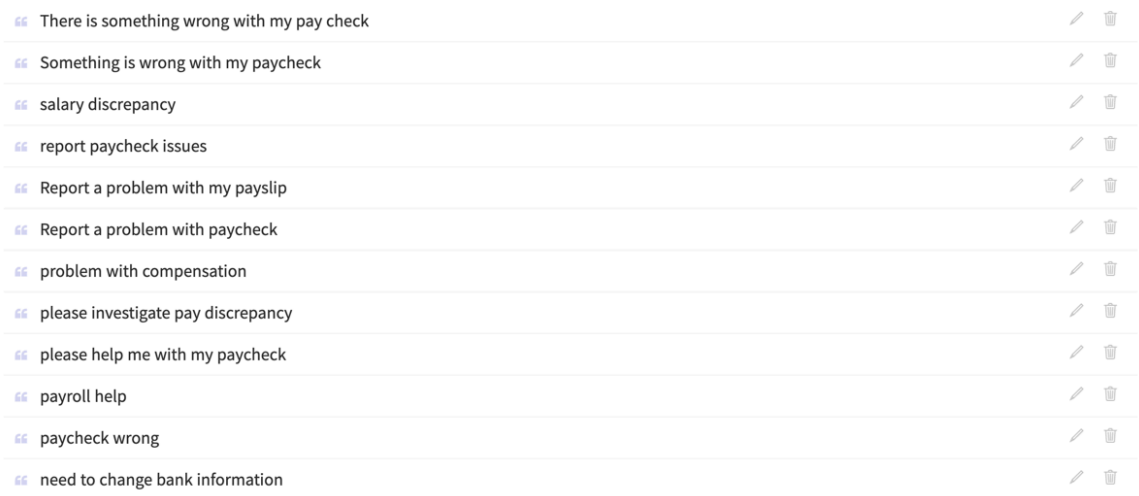
In this example scenario, you enter `PayDiscrepancy` for the name. You can also add a description.

Note: When you create an intent, the system adds a hashtag to the intent name.

8. Select **Add intent**.
   The #PayDiscrepancy intent screen appears, including sections for its utterances, associated entities, and settings. The intent draft status is also shown in the upper right corner of the Intent screen.
9. In the **Utterances** tab, enter training utterance examples that are relevant to the intent.

   Note: The utterance examples that you provide must be unique and contain fewer than 25 words or 200 characters. Aim to add at least 15 utterances with as much variety between them as possible. You must add at least 5 utterances to begin training the model.

   In this scenario, you add the following utterances into the field and select **Add**



As you build and retrain your model iteratively, you can check how your updated intent affects the model's predictions. See Train and try your NLU model.

What to do next

Train your model to save your updates. For issues with intents, see Resolve intent issues.

To improve your utterances, add entities to provide context. See NLU entities.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The available Intent Discovery feature can help identify intents that would be possible to add, based on your historical data.

Reusing intents from prebuilt NLU models

Reuse Natural Language Understanding (NLU) intents by importing them from a prebuilt NLU model to other models. Reusing intents saves time when building your models.

Prebuilt Virtual Agent NLU models provide language understanding needed for chatbot conversation flows in ITSM, CSM, and HR topics. Each NLU intent in these models maps to a single Virtual Agent conversation topic created in Virtual Agent Designer accordingly.

The prebuilt models are read-only, but contain validated intents that you can reuse in your NLU models. For example, the prebuilt HR model for Virtual Agent contains several intents regarding HR requests. Adding the intents to one of your models makes that model ready to respond to user inputs regarding HR.

Prebuilt Virtual Agent model content

Each of the prebuilt Virtual Agent models uses language that's pertinent to their respective business unit (BU). The NLU that processes this language, built from a word corpus of 3 million words, is context-aware of general linguistic patterns and both ServiceNow and user-defined vocabularies.

General Patterns	ServiceNow Vocabulary	User Defined Vocabulary
• Date • Time • Duration • Location • Phone number • Phone extension • Email • Email domain • Website address • Computer drive • Executable file • IP address • Bytes • Proper Names	• Task number • Incident number • Change number • Case number • KB article number • Problem number	• Standard phrases • List – ServiceNow Table: Field Examples • Change request: Type • Task: Impact • Task: Urgency • List – user defined • Regular expressions

The prebuilt models map to common conversation topics used by Virtual Agent chatbots across each of the three BUs. Those topics correspond with intents that perform BU-specific actions.



For more NLU and Virtual Agent context, refer to the following product documentation:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- For ITSM context, see ITSM Virtual Agent.
- For HR context, see Virtual Agent for HR Service Delivery.
- For CSM context, see Customer Service Virtual Agent conversations.

Virtual Agent and NLU plugins and roles

To access the prebuilt Virtual Agent models in the NLU Workbench, you must install and activate these plugins using the admin role.

Plugins	Descriptions
• CSM Virtual Agent conversations [com.sn_csm.virtualagent] • Human Resources Scoped App: Virtual Agent Conversations [com.sn_hr_virtual_agent] • ITSM Virtual Agent conversations [com.snc.itsm.virtualagent]	Enables Pre-built Virtual Agent topics, topic blocks, and ServiceNow NLU models for the Customer Service Management, HR Service Delivery, and IT Service Management applications. • Requires the Glide Virtual Agent (com.glide.cs.chatbot) plugin to be activated. • Automatically activates the NLU Workbench (com.snc.nlu_studio) plugin.

Creating NLU models that reuse predefined Virtual Agent intents

The prebuilt Virtual Agent models are set to read-only and can't be edited. However, the intents in these models can be imported into your new model, alongside any new intents you've created in your model. You can also copy the prebuilt model and use the intents as a foundation for a new model.

To reuse the intents from a prebuilt Virtual Agent model, choose the prebuilt model when importing intents. See Import an NLU intent.

NLU models and their intents and entities are associated to an application scope. The scope can't be changed later, so verify your target application scope before you begin.

Import an NLU intent

As you create intents for your Natural Language Understanding (NLU) model, you can also import and reuse intents from other models in the same application scope. Reusing intents saves time when building new models.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- You can import intents for Virtual Agent and AI Search models.
- Verify that the application scope of your source model and target model are the same.
- Role required: nlu_editor, nlu_admin, or admin

About this task

When you import an intent, the following associated items are also imported:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- training utterances, with their entity annotations
- entities

The following items are not included when importing an intent:

- test utterances in the default test set
- regular and pattern vocabulary

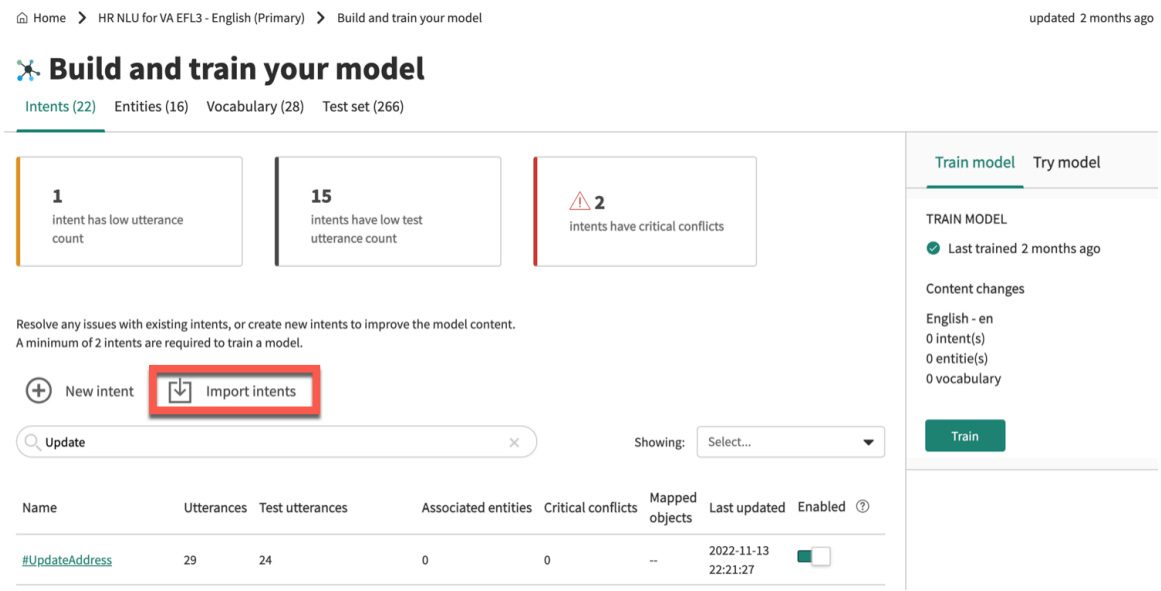
Note: You cannot import intents in the following situations:

- When the model already has an intent with the same name as the intent you are trying to import.
- When the model already has an entity with the same name but different attributes as what you are trying to import.

In this example scenario, you are building the content of your NLU model and you want to enhance the model by reusing intents from another model.

Procedure

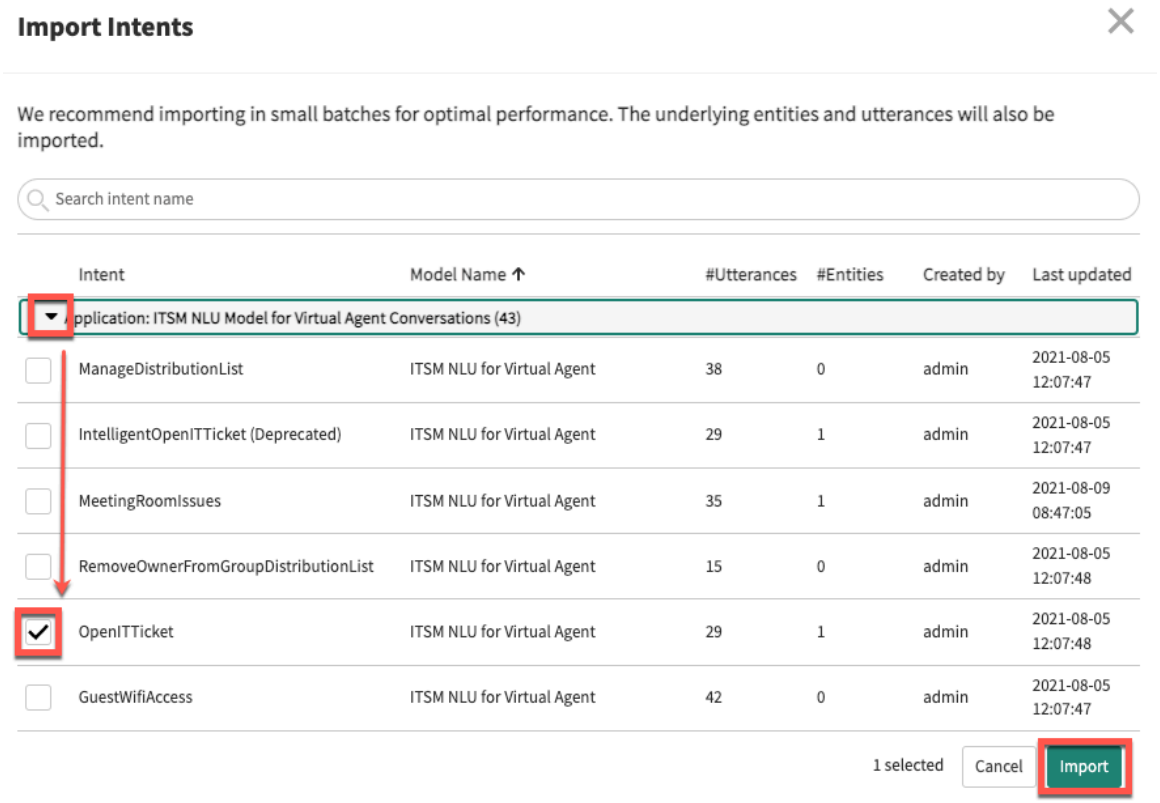
1. Ensure that you are in the target application scope, and navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, then select the name of the model you want to add an intent to.
3. In the model overview page, find the **Build and train your model** phase and select **View phase**.
4. In the **Intents** tab, select **Import intents**.



   The **Import intents** screen appears, showing a list of NLU models from which you can select intents to import.

5. Locate the intents you want to import, and select their check boxes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



In this scenario, you select the **OpenITTicket** intent in the ITSM NLU Model for Virtual Agent Conversations model.

6. Select **Import**.

   Result: The **#OpenITTicket** intent appears in the intents list of your target model. By importing this particular intent, you have added 29 additional training utterances to your model.

What to do next

Review the annotations of the newly imported training utterances. Add vocabulary annotation if needed.

Train your model to save your updates. For issues with intents, see Resolve intent issues.

If there are any secondary language models associated to your target model, consider importing and translating the new content. See Import primary model content to a secondary model.

For information about prebuilt Virtual Agent models that you can import from, see Reusing intents from prebuilt NLU models.

Resolve intent issues

Use the issue cards to identify intents that have conflicts, need reviewing, or need more utterances. Resolving intent issues ensures the intents in your Natural Language Understanding (NLU) models work properly.

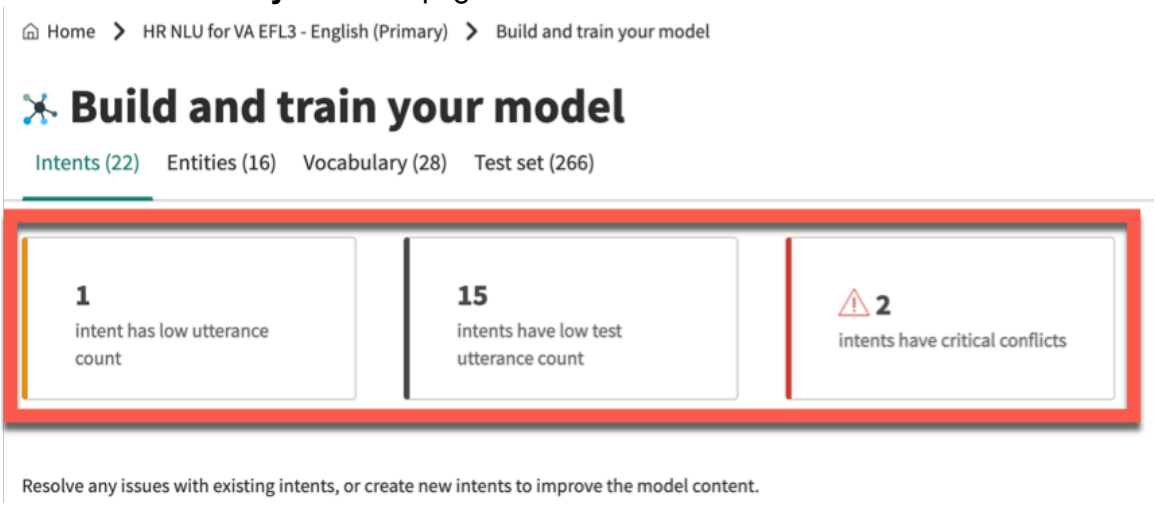
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Workbench - Advanced Features plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- You can resolve conflicts for Virtual Agent and AI Search models.
- Role required: nlu_admin or admin. When assigned to a model, the nlu_editor can resolve intent issues for that model.

About this task

When a trained model contains intents with issues, the issue cards are displayed on the intents tab of **Build and train your model** page. Select each card to filter the list of intents for that issue.



The cards display different issues:

- **intents have low utterance count**: Intent doesn't contain the required minimum of 5 training utterances. This card also displays when the intent contains less than the recommended count of 15 utterances.
- **intents have low test utterance count**: The default test set does not have enough test utterances for the intent. This card displays when the count of test utterances for an intent is below the recommended 15.
- **intents have critical conflicts**: Intent contains training utterances that overlap with utterances in another intent.

  Note: Utterances which are marked as `Not relevant`may also conflict with utterances assigned to intents. For more information, see Irrelevance detection in NLU.

- **intents need review**: Intent was translated but must be reviewed by an nlu_admin or nlu_editor.

The issue cards are hidden under the following circumstances:

- The model is a pre-built model
- The model is in a draft state and not trained
- The model contains no conflicts
- A successful conflict report doesn't exist for the latest trained version
- The NLU Workbench - Advanced Features plugin isn't installed

Resolve issues in the following ways:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Issue	Resolution
intents have low utterance count	Add more training utterances to the intent. Intents need at least 5 utterances. The recommended count is at least 15 utterances.
intents have low test utterance count	Add more test utterances for the intent to the model's default test set. The recommended count is at least 15.
intents have critical conflicts	Remove or edit any utterances so that each utterance applies to only one intent.
intents need review	Have an nlu_admin or nlu_editor review the intent. For more information, see Enable or disable a secondary model intent.

In this example procedure, you are resolving critical conflicts in one of your intents. You built a Virtual Agent model for handling HR requests, but two of your intents, GeneralHRInquiry and CreateHRGeneralInquiryCase included similar utterances.

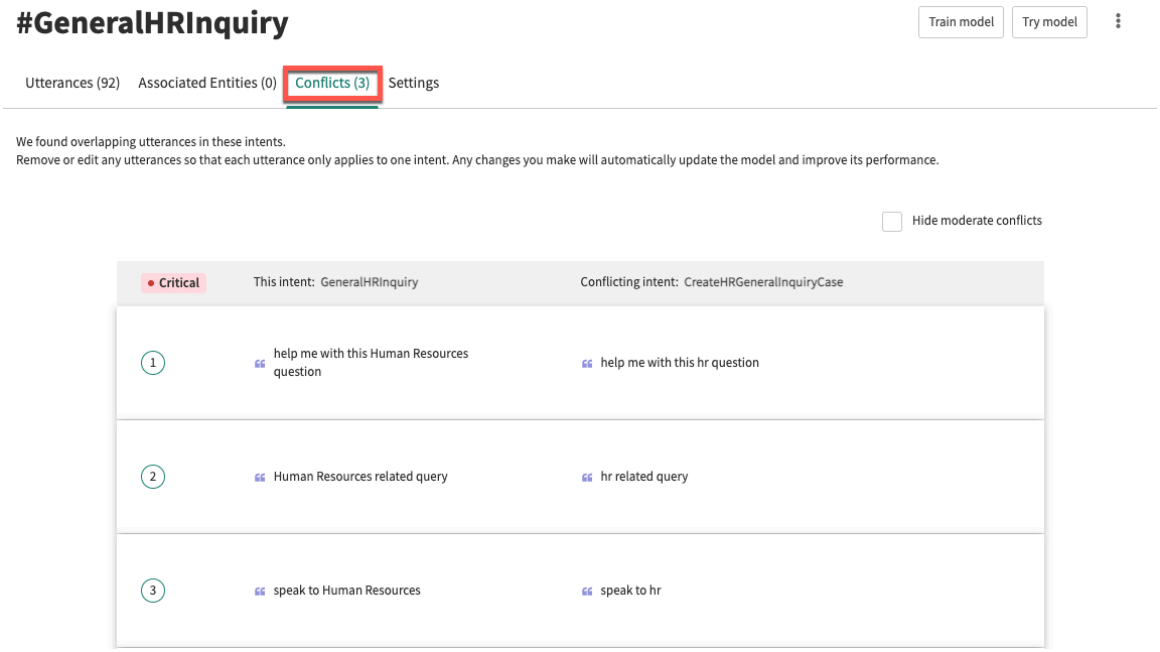
Procedure

1. Ensure that you are in the same application scope as your model, and navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab corresponding to your model's application, then select the name of your model.
3. In the **Build and train your model** card, click **View phase**.
4. On the intents tab, click the **intents have critical conflicts** card.
   The intents list filters and shows the intents with critical conflicts.
5. Click the intent name you want to resolve.
   In this example, you click the GeneralHRInquiry intent.
6. Click the **Conflicts** tab.

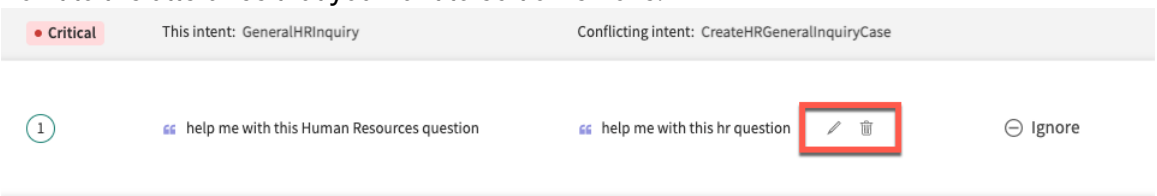
   The conflicts tab appears and lists the conflicting utterances. The utterances for the current intent show on the left. The conflicting utterances for the other intent appear on the right.

   Note: You can hide moderate conflicts to focus on the critical conflicts. You can also ignore conflicts. However, be sure to fix all conflicts for best performance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



7. Point to the utterance that you want to edit or remove.



In this example procedure, click the trash icon to remove the utterance from the CreateHRGeneralInquiryCase intent.

The utterance is removed from the intent.

Note: Any changes you make automatically save.

What to do next

Utterances designated as `Irrelevant` or `Not relevant`may also conflict with training utterances. Irrelevant utterances are displayed as though they are in their own intent named **NO_INTENT**. However, these utterances cannot be edited or deleted in Conflicts. Copy the utterance to the Irrelevance Detection module to edit or delete. For more information, see Irrelevance detection in NLU.

Continue to resolve all conflicts. When finished, train your model to see a list of updated conflicts, if any. Resolve all intent issues prior to testing and publishing your model.

NLU entities

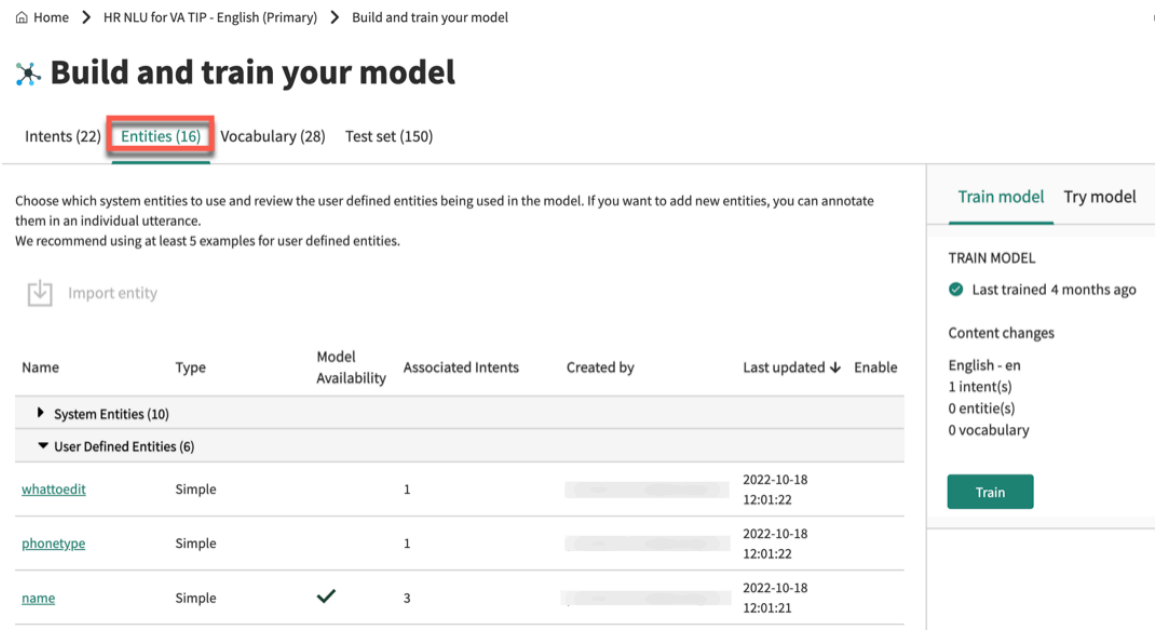
Entities provide your model with additional context when receiving user input. Add entities to your utterances and intents to improve the predictions of your Natural Language Understanding (NLU) model.

Think of entities as the object of the action that the user wants to perform. The model interprets the utterance by matching it with an intent, but also uses entities to gather more details on the request.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, if a user types in `I have a critical issue with a slow laptop`, the system predicts a match to the intent #TroubleshootSlowComputer. From this utterance, the model also identifies the following entities:

- HARDWARE (entity) - laptop (value)
- urgency (entity) - High (value)



NLU entities fall into two categories: system and user-defined. System entities such as DATE, TIME, and LOCATION are available by default in your instance. You can create your own user-defined entities to provide context relevant to your business.

There are five types of user-defined entities. Select one of the following links to learn more.

- Create a simple entity
- Create a mapped entity
- Create a pattern entity
- Create a system-derived entity
- Create an open-ended entity

When you create entities, annotate them on utterances to provide examples to help your model learn. By annotating entities, you provide your model with linguistic associations and meaningful context for the system vocabulary. Annotations strengthen the relevance of entities and help your model perform the correct action in response to your users' inputs.

You add entities to your utterances when you are creating the intents. The entities then become associated to that intent, giving you the **Associated Intents** number.

Regular expressions

Regular expressions (regex) help your model establish patterns that improve that model's ability to locate, match, and manage text. Use regular expressions with pattern entities to help your model understand formats such as email addresses, phone numbers, and incident numbers.

To learn more, see Using regular expressions in entities.

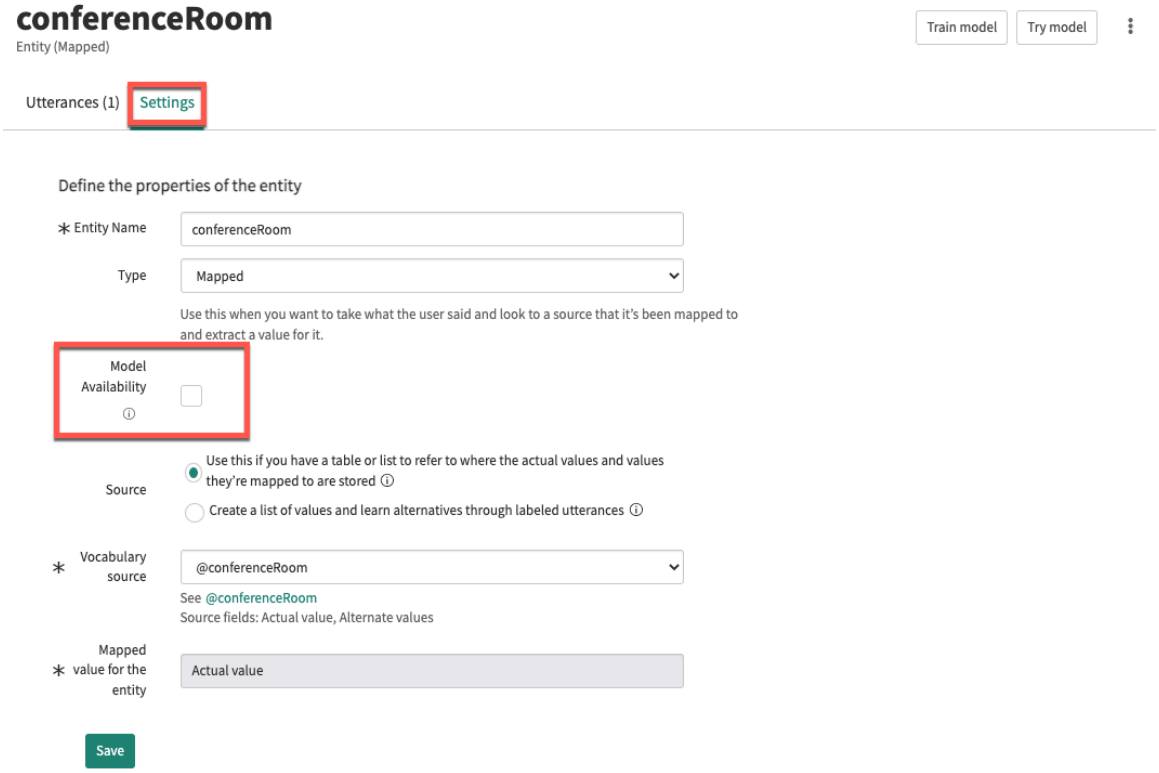
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Model availability

When you create an entity, you can choose to make the entity available for reuse by other intents in the model. If you didn't select the **Model availability** box when creating the entity, you can edit the entity afterwards.

On the Model details page, select **Entities**. Select the name of the entity to bring up the entity details page. Then, select the **Settings** tab.

Entity settings page



Create a simple entity

Create one or more simple entities from words in your utterance examples. An entity is an object of, or context for, an action.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model for Virtual Agent or AI Search.
- Create or use an existing intent.
- Role required: nlu_editor, nlu_admin, or admin. The nlu_editor must be assigned to the model.

About this task

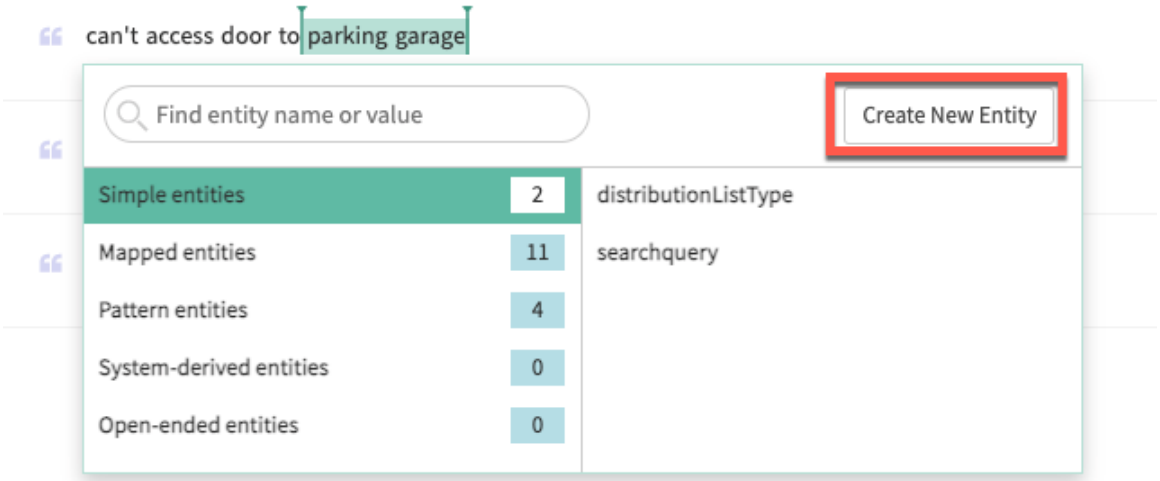
Simple entities are words or phrases whose value can be extracted by your model. Simple entities are identified based on the context in which the entity is used in an utterance.

For the following example procedure, you've already created an intent that's titled SubmitAccessRequest and you're creating a simple entity for the type of access the user is requesting.

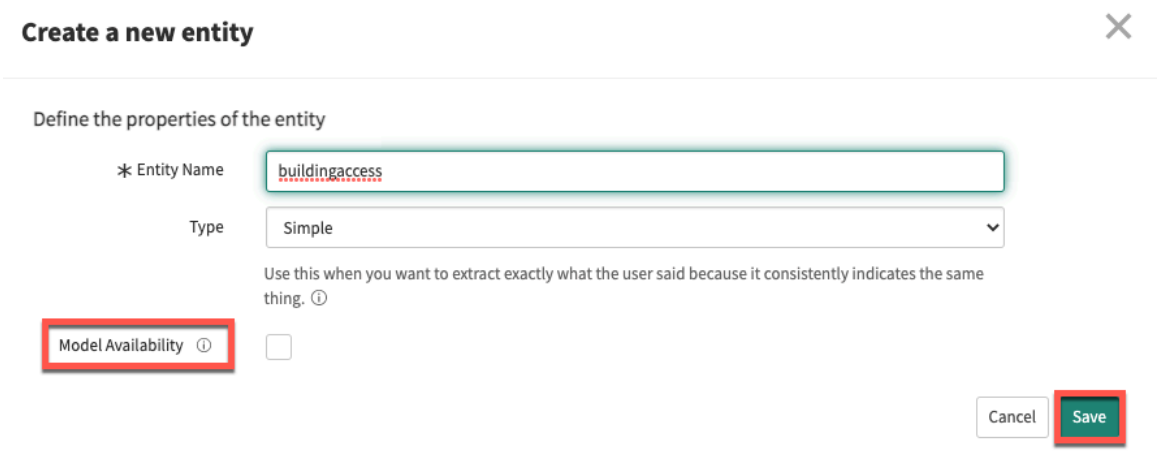
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab corresponding to your model's application, then the name of the model you want to add an entity to.
   For this example, you choose the NLU for Access Requests model.
3. In the **Build and train your model** card, click **View phase**.
4. In the intents tab, click the name of the intent.
   For this example, you click SubmitAccessRequest.
5. Click a word in one of the utterances to bring up the entity window.
   For this example, click the word parking garage.



6. Click **Create New Entity**.
7. In the Create a new entity screen, enter a name and select the **Simple** entity type.
   For this example, you enter `buildingaccess` for the **Entity Name**.



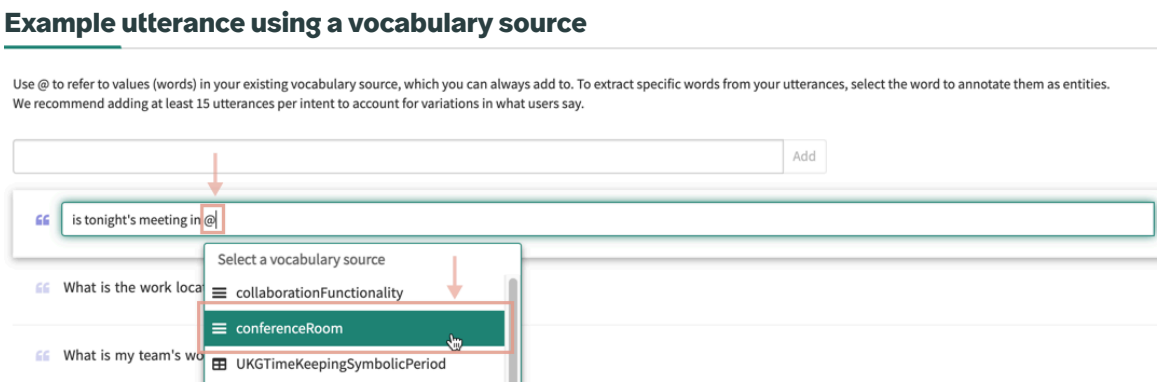
Note: Select **Model availability** to make the entity available to every intent in the model. If you check the box, the entity shows in the **Associated Entities** tab.

8. Click **Save**.
   The entity saves and the words in the utterance remain highlighted.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

Your utterances can reference a vocabulary source by using the @ handle. If you have a list of values that are defined in a vocabulary source, you can annotate the @ handle as a simple entity to extract it rather than repeating the utterance for all of the values. The referenced vocabulary source can be a table or a list. For example, the following image shows how you invoke a vocabulary source that lists various conference room names.



For more information, see NLU vocabulary.

Create a mapped entity

Create an entity mapped to a vocabulary source, or to a list of values you manually create for the entity. Mapped entities can help provide multiple values the model can use as context when interpreting utterances.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model for Virtual Agent or AI Search.
- Create or use an existing intent.
- Role required: nlu_editor, nlu_admin, or admin. The nlu_editor must be assigned to the model.

About this task

Mapped entities take the words of the utterance and extract value based on a designated source. The model uses the source when predicting the intent.

When you create a mapped entity, you have the following three options for the source.

- Manual list of values: Use this option to manually enter a list of values for the entity. For example, you could create a mapped entity named priority and map it to the word urgent in an utterance, then manually build a list for it with values of High, Medium, and Low.
- Table vocabulary source: Use this option if you have a ServiceNow table that has the values you're looking for. Mapping an entity to a table vocabulary source enables the entity to reference multiple values from the table. For example, use a @Location vocabulary source, where @Location has values for cities and countries.
- List vocabulary source: Use this option if you don't have a ServiceNow table that has the values you're looking for. For example, use a @mouse vocabulary source, where @mouse has values for various models of hand-held computer devices.

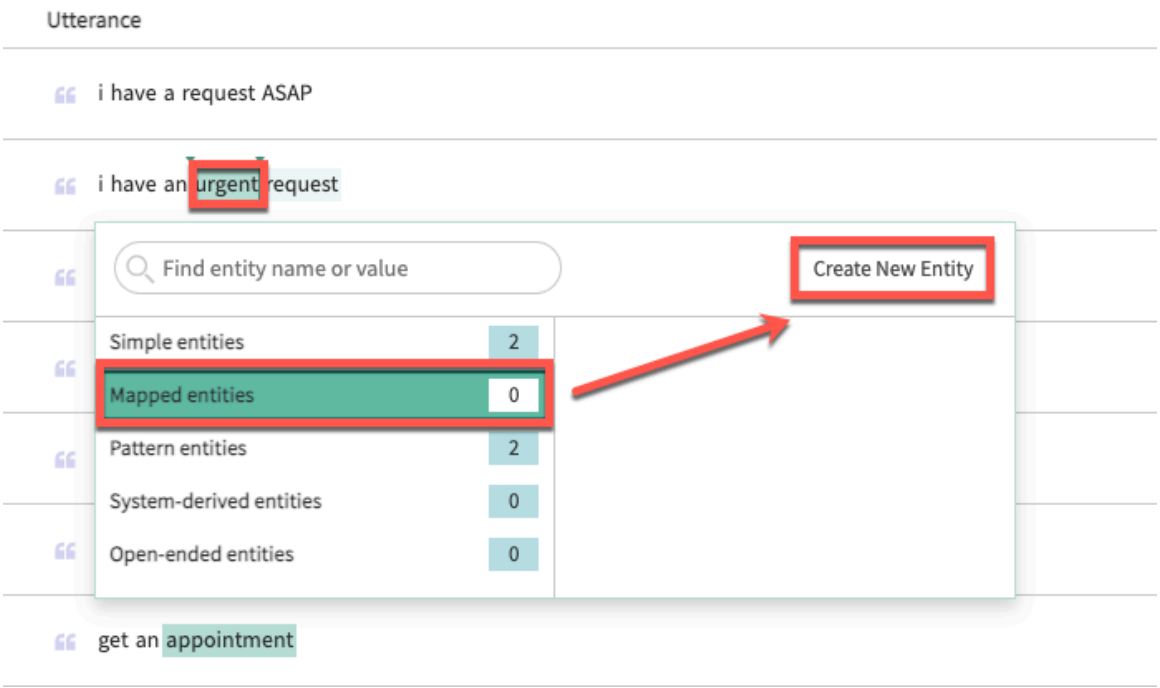
In this example procedure, you create a mapped entity for urgency.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, then the name of your model.
3. On the model details page, select the **Intents** tab.
4. In the Intents section of the model, select the name of an intent.
   For this example procedure, you select **#SubmitRequest**.
5. In the Utterances tab, select a word in an utterance

   In this scenario, you select the word urgent in the utterance I have an urgent request.
6. Select **Mapped entities**.
7. Select **Create New Entity.**



8. On the form, configure the fields.

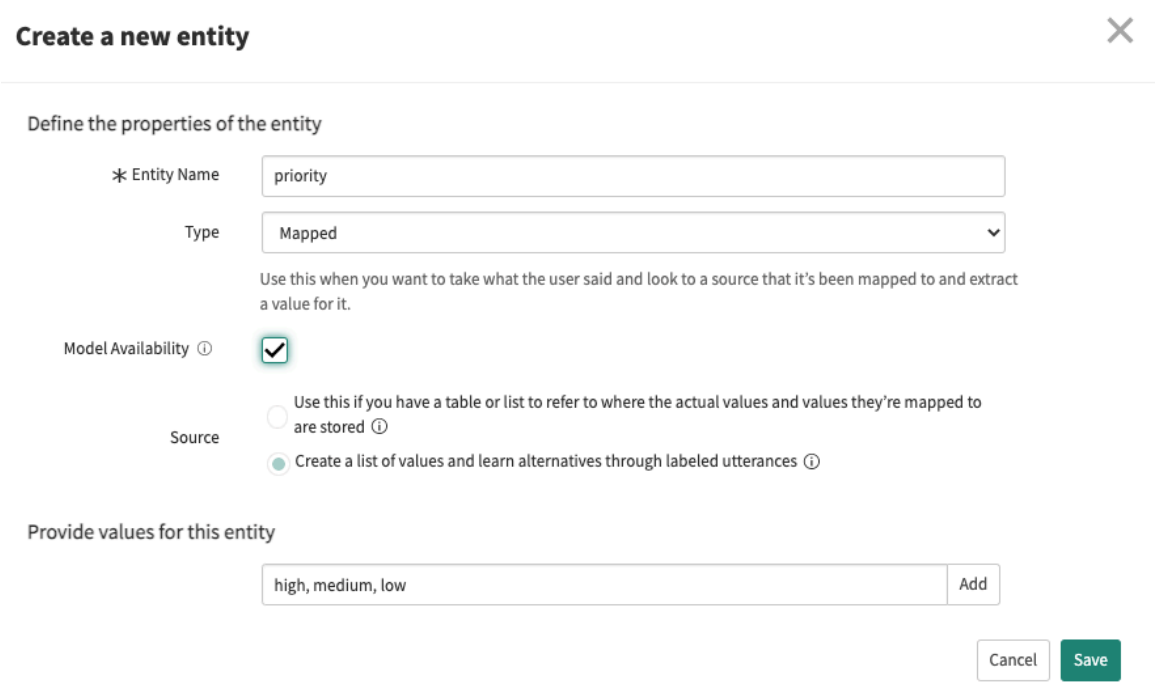
Field	Description
Entity Name	Name for the entity.
Type	Type of entity.
Model availability	Select this option if you want this entity to be included in all intents in your model.
Source	Source of the entity values.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Provide values for this entity	Values used to provide context for the model.

For this example procedure, use the following configurations:

- **Entity Name**: `priority`
- **Type**: Mapped
- **Model availability**: Select the check box
- **Source: Use this if you have a table or list to refer to where the actual values and values they're mapped to are stored**
- **Mapped value for the entity**: `high, medium, low`.



9. Click **Save**.

Result: Your mapped entity saves. The entity appears on the **Associated entities** tab. Now the model can leverage machine learning and use the values provided to identify possible values.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

You can create a mapped entity using a vocabulary source to use the values in the source as the mapped entity.

Related topics

Create a table vocabulary source

Create a list vocabulary source

Create a pattern entity

Create a pattern entity from a word or phrase with repeatable patterns, such as email addresses and phone numbers. These patterns help the system to recognize similar utterances based on the patterns.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model for Virtual Agent or AI Search.
- Create or use an existing intent.
- Role required: nlu_editor, nlu_admin, or admin. The nlu_editor must be assigned to the model.

About this task

Pattern entities help your model identify, contextualize, and govern similar types of content. Inputs such as phone numbers and email address follow the same pattern. By creating a pattern entity, you can train the model to recognize all variations of those inputs.

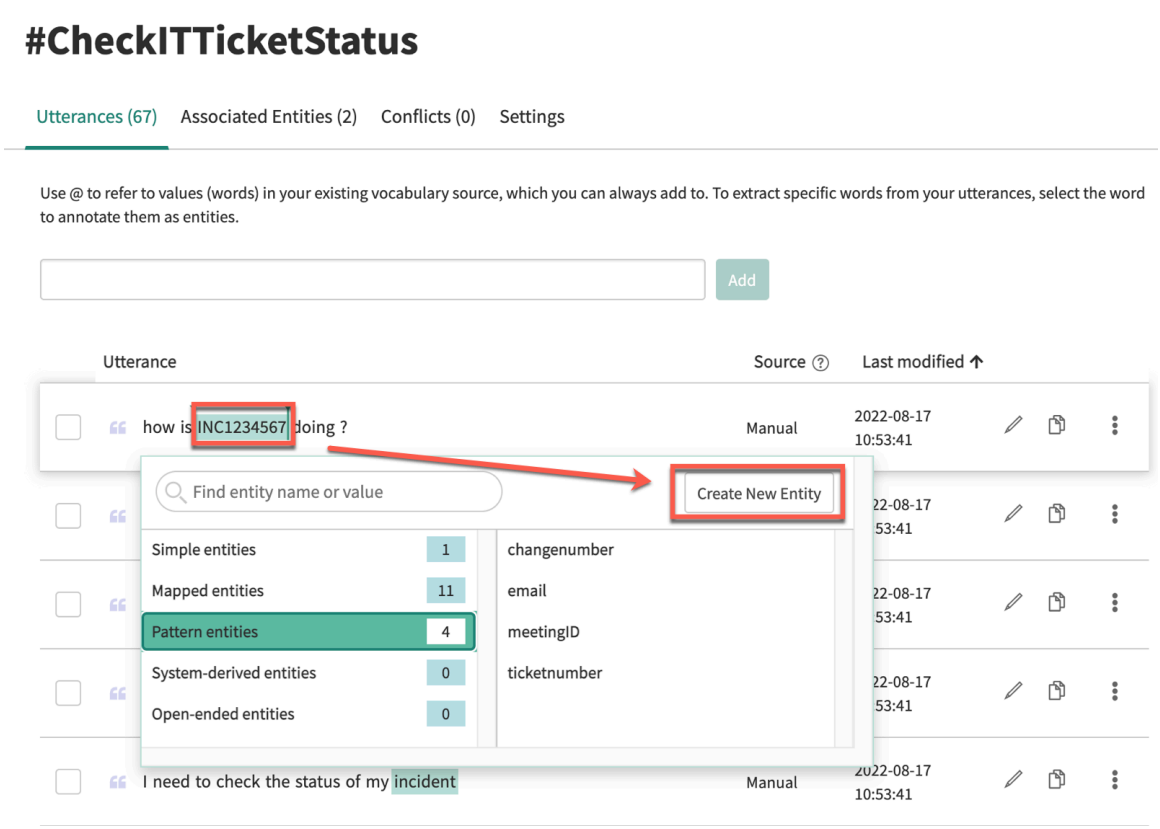
Note: Pattern entities use regular expressions (regex). The regex field value is a Java regular expression. For more information on how to use regex, see Using regular expressions in entities.

In this example scenario, you've created an intent that's titled #CheckITTicketStatus. In this example procedure, you're creating a pattern entity for incident record numbers.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, then the name of the model.
3. On the model details page, select **Intents**.
4. Select the name of an intent.
   In this example scenario, you select **#CheckITTicketStatus**.
5. Select one of the words of an utterance.
   In this scenario, you click **INC1234567**.

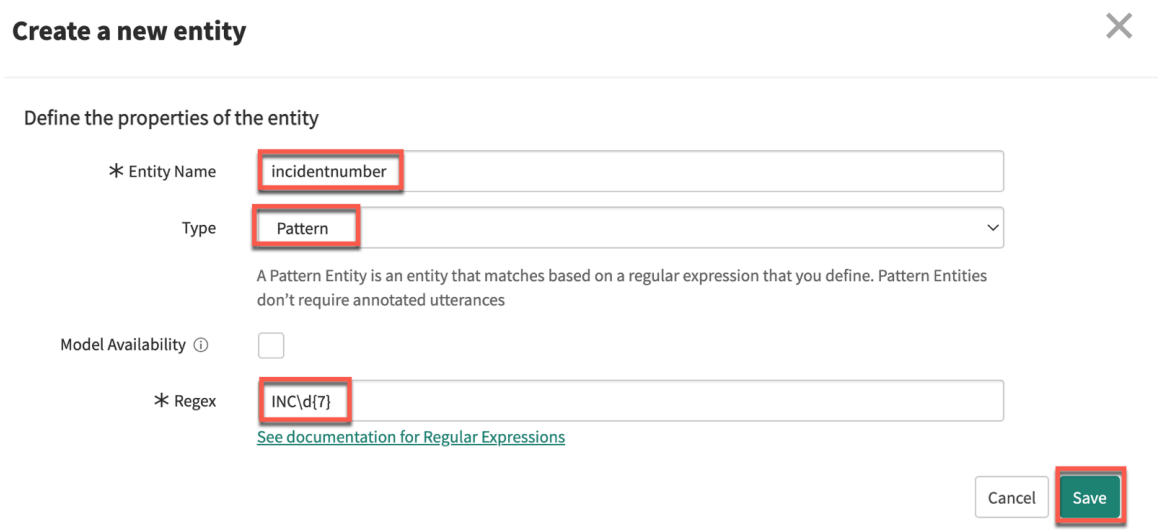
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



6. Select **Create New Entity**.
7. In the Create a new entity window, configure the fields.

   For this example procedure, use these configurations:

   - **Entity Name**: Enter `incidentnumber`
   - **Type**: Select **Pattern**
   - **Model Availability**: Select this box if you want this entity to be included in all intents in your model
   - **Regex**: Enter `INC\d{7}`



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. Click **Save**.
   The pattern entity saves and appears in the **Associated Entities** tab. The model can use the pattern entity to interpret variations of the annotated utterance. For example, the model can match the #CheckITTicketStatus intent if a user enters any incident record number.

Create a system-derived entity

Create a custom entity that's derived from a default system entity such as date, time, duration, or location.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model for Virtual Agent or AI Search.
- Create or use an existing intent.
- Role required: nlu_editor, nlu_admin, or admin. The nlu_editor must be assigned to the model.

About this task

System entities are pre-built in your ServiceNow instance by default, and appear on the Entities section of the Model screen. These entities, such as DATE, TIME, and DATE_TIME, are enabled by default. You can disable and re-enable them if needed by clicking **Enable**.

System-derived entities extend your system entities, providing more context. For example, your model already understands date formats thanks to the system entity DATE. However, you can create system-derived entities, such as startdate and enddate, to extract more information from user utterances regarding dates.

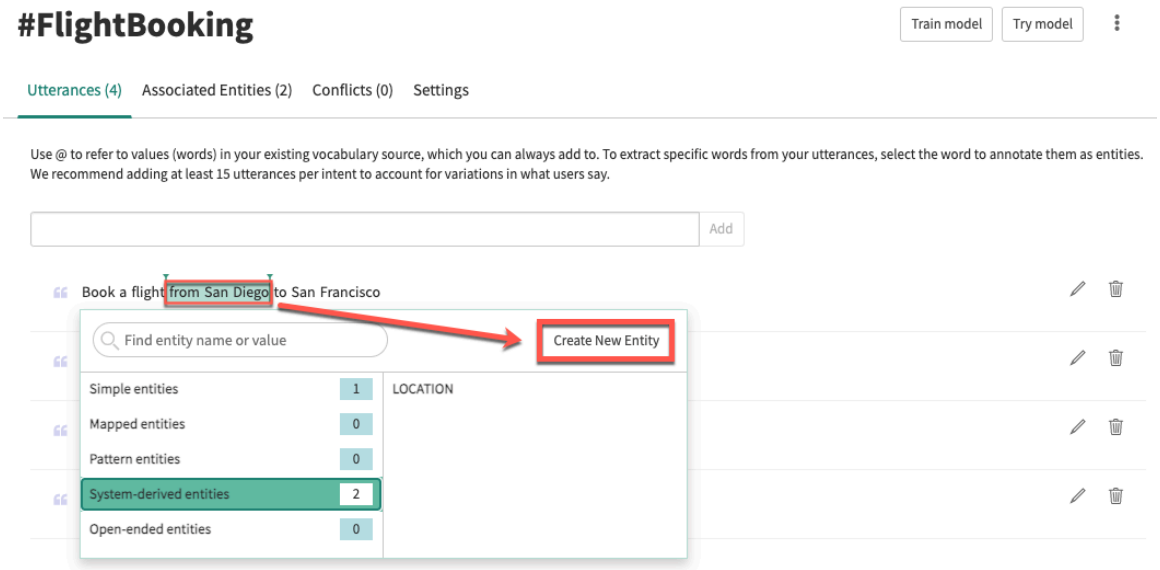
For the following example scenario, you're creating a model for booking flights, cars, hotels, and events. You want your intent #FlightBooking to interpret utterances with your users flight requests. Your model contains the system entity LOCATION, but flight plans usually involve two locations.

In this example procedure, you create two system-derived entities to gather the departure and arrival locations of the flight.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, then the name of your model.
3. In the model details page, select **Intents**.
4. Select the name of the intent.
   For this example, you select the **#FlightBooking** intent.
5. On the **Utterances** tab, select a word or phrase from one of the utterances.
   In this example scenario, you have an example utterances `book a flight from San Diego to San Francisco`. For this step, click **from San Diego**.

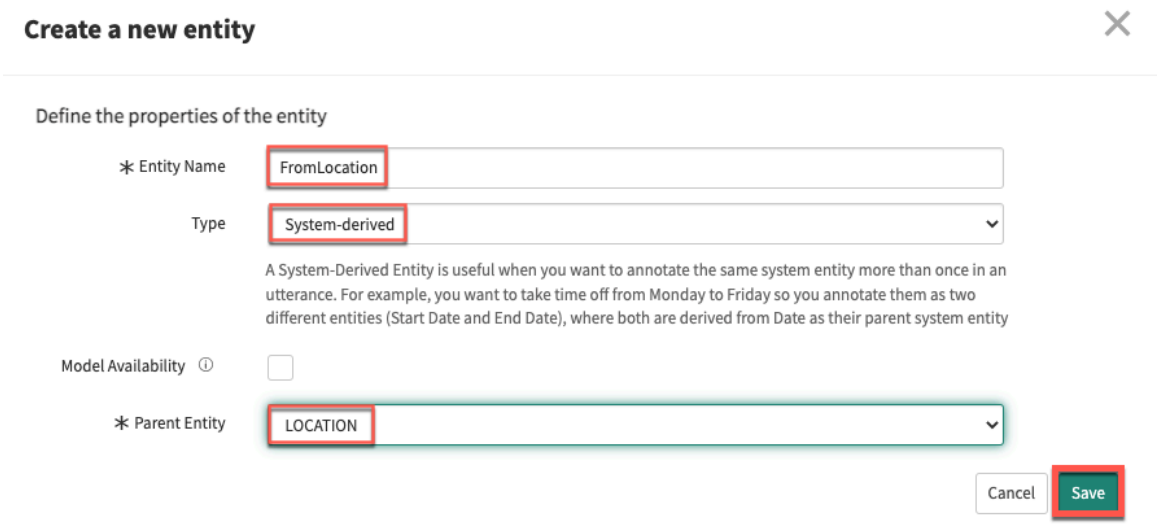
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



6. Select **Create New Entity**.
7. In the Create a new entity screen, configure the fields.

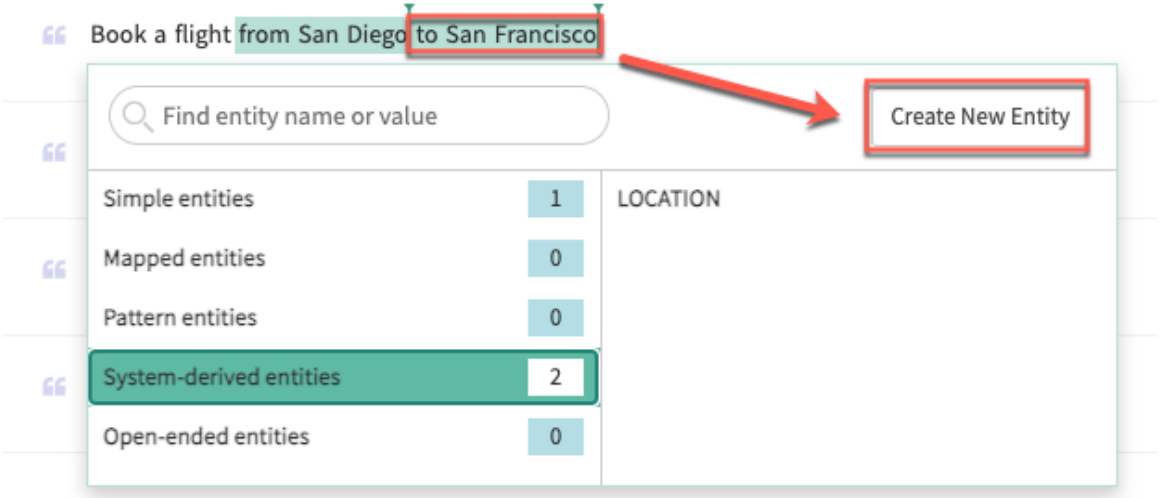
   For this example, use these configurations.

   - **Entity Name**: Enter `FromLocation`
   - **Type**: Select **System-derived**
   - **Model Availability**: Select this check box if you want this entity to be included in all intents in your model
   - Parent Entity: Select **LOCATION**



8. Select **Save**.
9. On the same utterance, select another word or phrase.
   In this example scenario, you click **to San Francisco**.

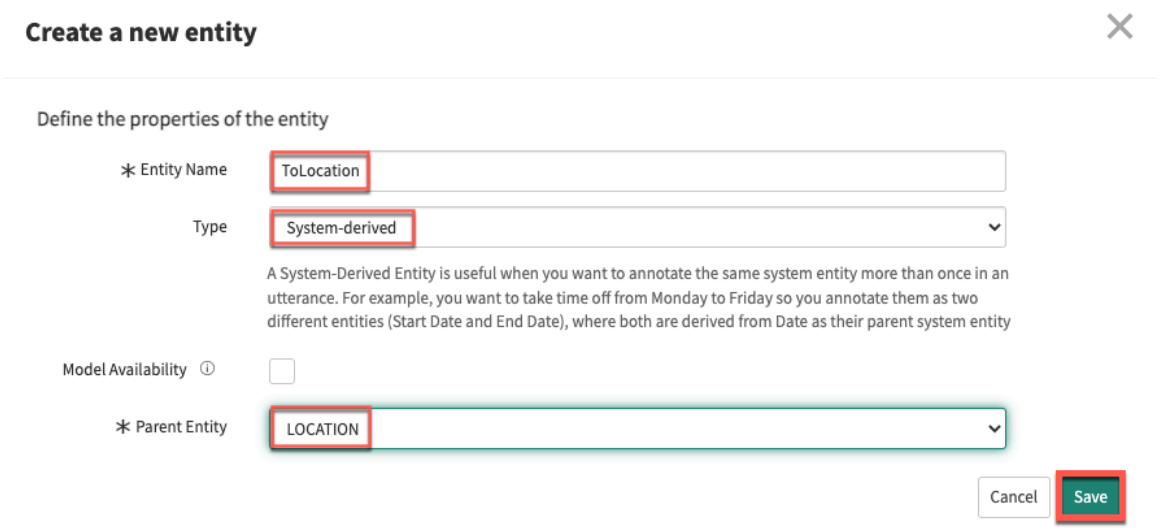
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



10. Select **Create New Entity**.

For this example, use these configurations.

- **Entity Name**: Enter `ToLocation`
- **Type**: Select **System-derived**
- **Model Availability**: Select this check box if you want this entity to be included in all intents in your model
- **Parent Entity**: Select **LOCATION**



11. Select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result: Your two system-derived entities appear in the entity window.



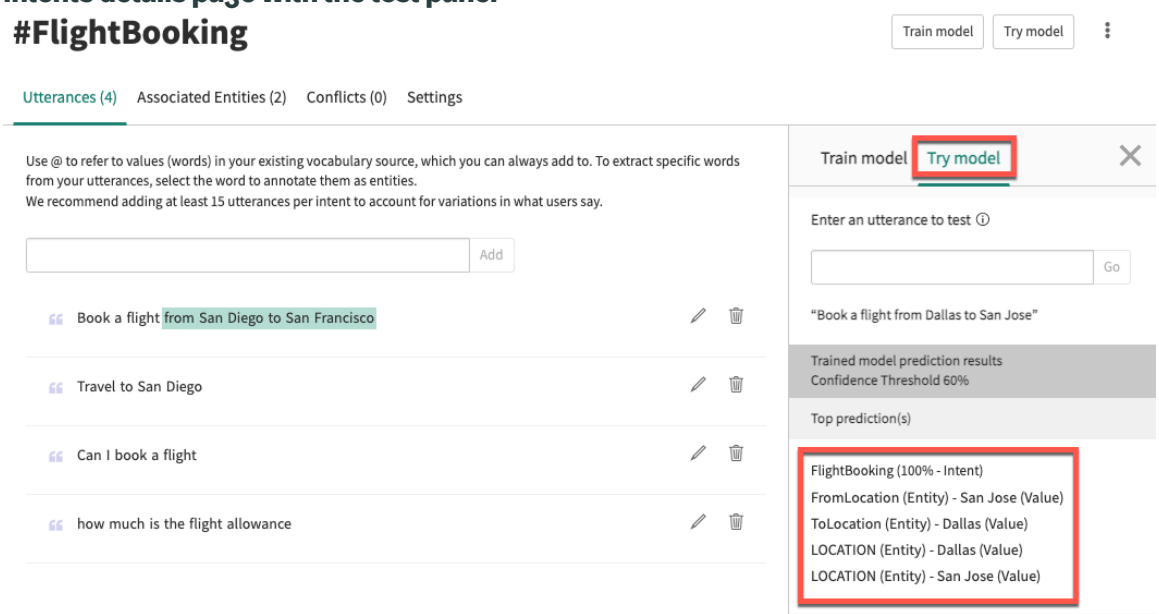
What to do next

Train your model to save the entities. You can try your model to see if it recognizes and interprets variations of the entities that you created.

For this example, you can test your model with a different departure and arrival location.

1. Select **Try model**.
2. Enter `book a flight from Dallas to San Jose`.
3. Select **Go**.

Intents details page with the test panel



The model predicts the intent, showing the entities it used to determine the values. Notice how it uses both the system entity and system-derived entity when predicting.

Create an open-ended entity

Use an open-ended entity when you want to improve intent prediction accuracy. Open-ended entities help your model focus on the context of the utterances.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model for Virtual Agent or AI Search.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Create or use an existing intent.
- Role required: nlu_editor, nlu_admin, or admin. The nlu_editor must be assigned to the model.

About this task

Open-ended entities tell the model to focus on the context of the entity rather than the entity itself. When you mark a word or phrase as open-ended, the system skips the entity and predicts the intent from the context that precedes or follows the entity in the utterance.

For example, in the utterance `I want to order an iPhone`, you annotate the words "an iPhone" as an open-ended entity. The model focuses on the context, predicting the user wants to order something. Since there are numerous things the user could want to order, naming all of them would be an unbearable task for the model author.

Using an open-ended entity instead of a simple entity helps the model focus on the rest of the utterance and not the entity. In the iPhone example, the entity itself is less relevant; so you want the system to ignore it.

In other scenarios you should use a simple entity, as there could be multiple intents where you shouldn't have the system ignore the entity.

Note: You can't annotate a vocabulary source (referenced by @vocab_source in an utterance) as an open ended entity. You can only annotate a vocabulary source as a simple entity or a mapped entity. For example, if the utterance is “I want to order a laptop”, then the word “laptop” can be annotated as an open ended entity. However, if the utterance is “I want to order @laptop” where @laptop refers to a table vocabulary source or a list vocabulary source, it can't be annotated as an open ended entity.

For this example scenario, you've created an NLU model with an intent for your users to order company merchandise.

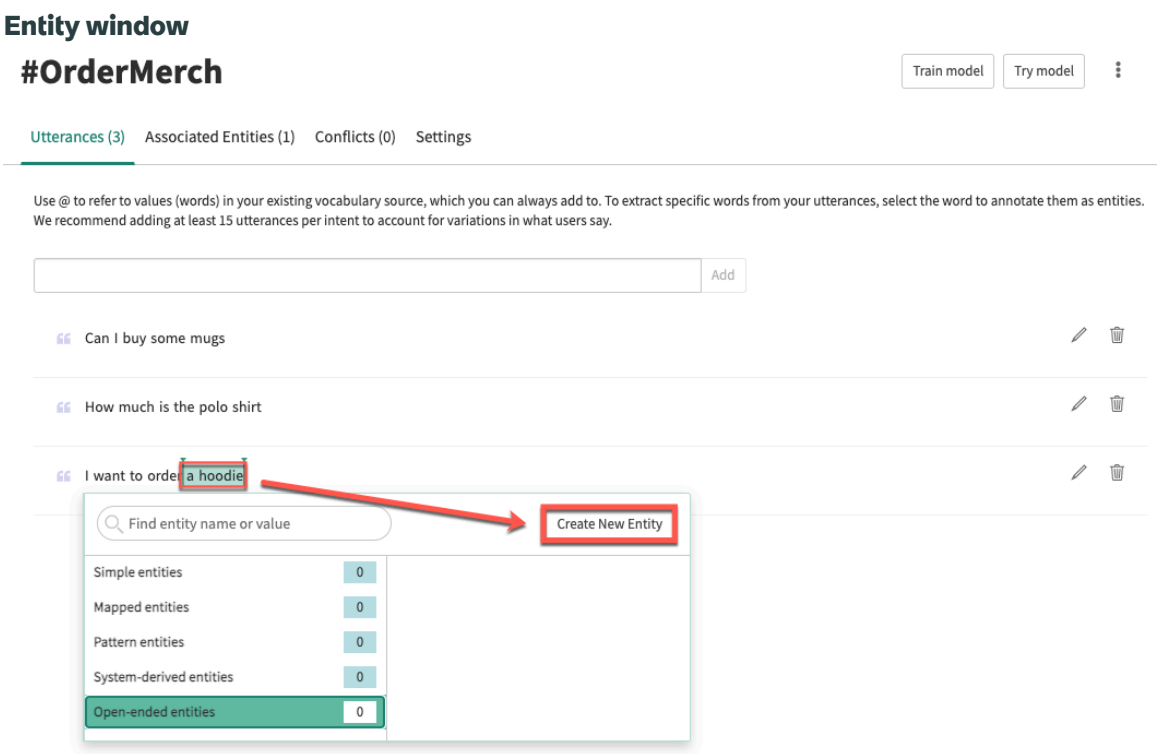
In the following example procedure, you create an entity from one of your utterances so the system can recognize it as open-ended and reusable in other NLU models in your instance.

Note: You can use only one open-ended entity per intent.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, then select the name of your model.
3. On the Model details page, click **Intents**.
4. Select the name of the intent you want to add the entity to.
   For this example, you select the #OrderMerch intent.
5. In the Utterances tab, select a word or phrase of one of the utterances to bring up the entities window.
   For this example, you select a hoodie.

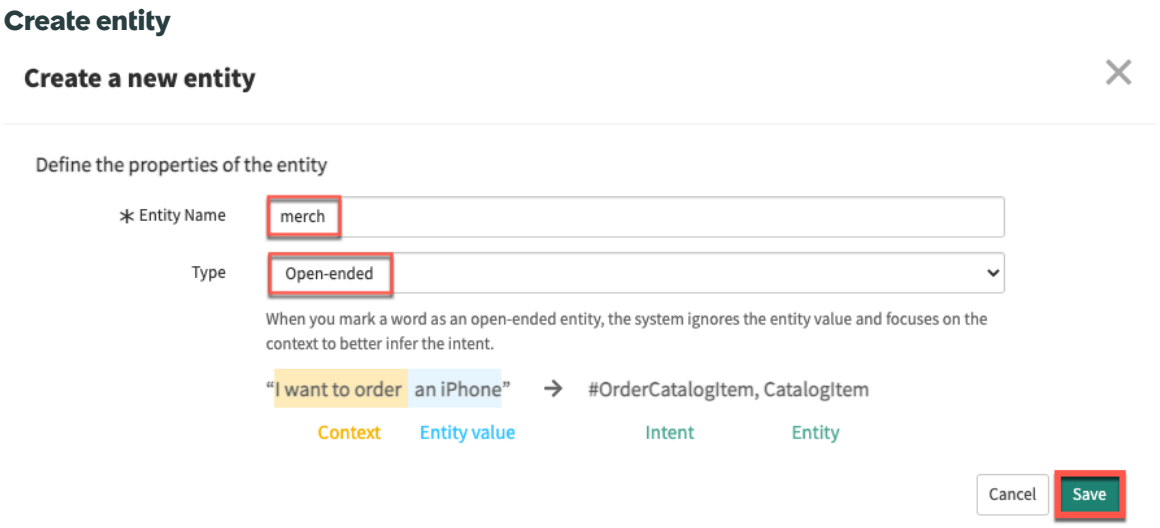
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



6. Select **Create New Entity**
7. On the Create a new entity screen, configure the fields.

   For this example, use the following configurations:

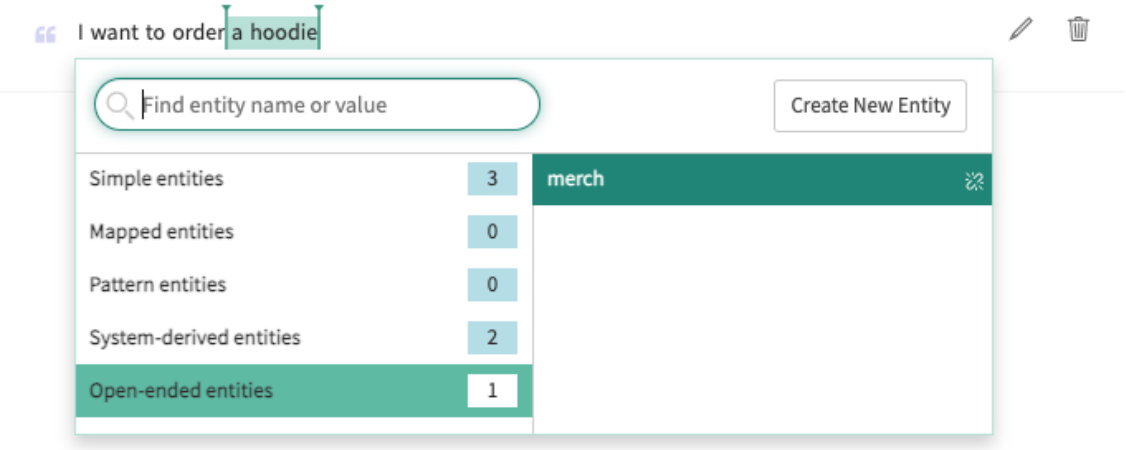
   - **Entity Name**: `merch`
   - **Type**: Select **Open-Ended**



8. Select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The merch open-ended entity is annotated in the Utterances section of your model's Intent screen. When you point to its name, you can see that it persists as a new entity in the annotation details. This entity is reusable in all other NLU models in your instance.

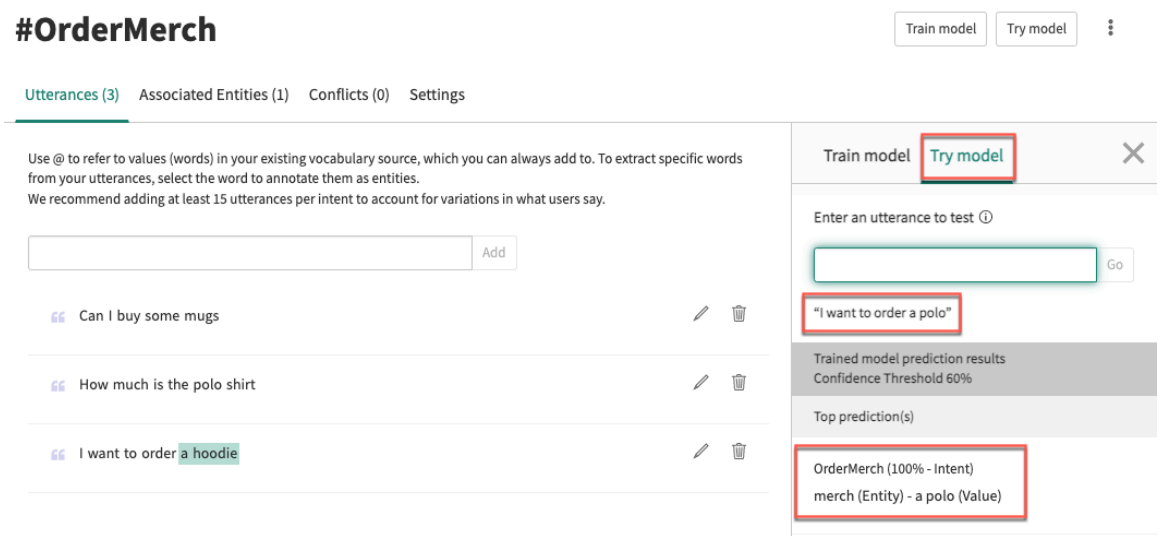


What to do next

Train your model to save the entities. You can try your model to see if it interprets the utterance based on the context of the entity, rather than the entity itself.

For this example, you can test your model with a different merchandise item.

1. Select **Try Model**.
2. Enter `I want to order a polo`.
3. Select **Go**.



The model predicts the intent and shows that it used the merch entity for the a polo value.

Import entities

Reuse entities that you have created across your other Natural Language Understanding (NLU) models. Importing entities saves time and helps improve the intents in your model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model.
- Create or use existing intents and entities.
- Role required: nlu_editor, nlu_admin, or admin. The nlu_editor must be assigned to the model.

About this task

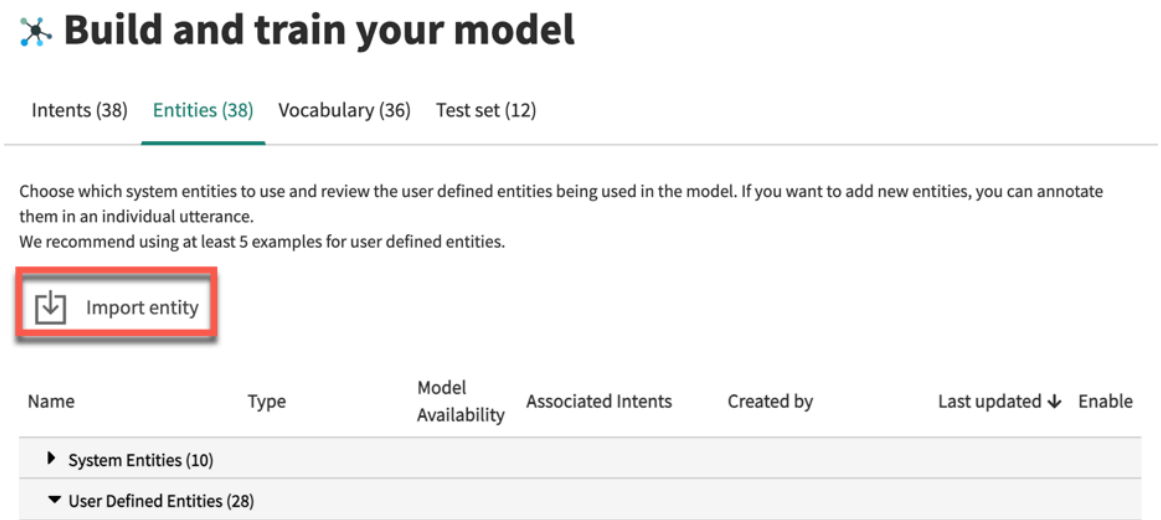
New entities are usually created inside an individual intent. To help save time and improve your intents, you can also reuse entities from your other models by importing them. After importing entities into a model, annotate the model's utterances using the new entities.

Note: All models include several system entities by default.

In this example procedure, you are importing several entities to a new model.

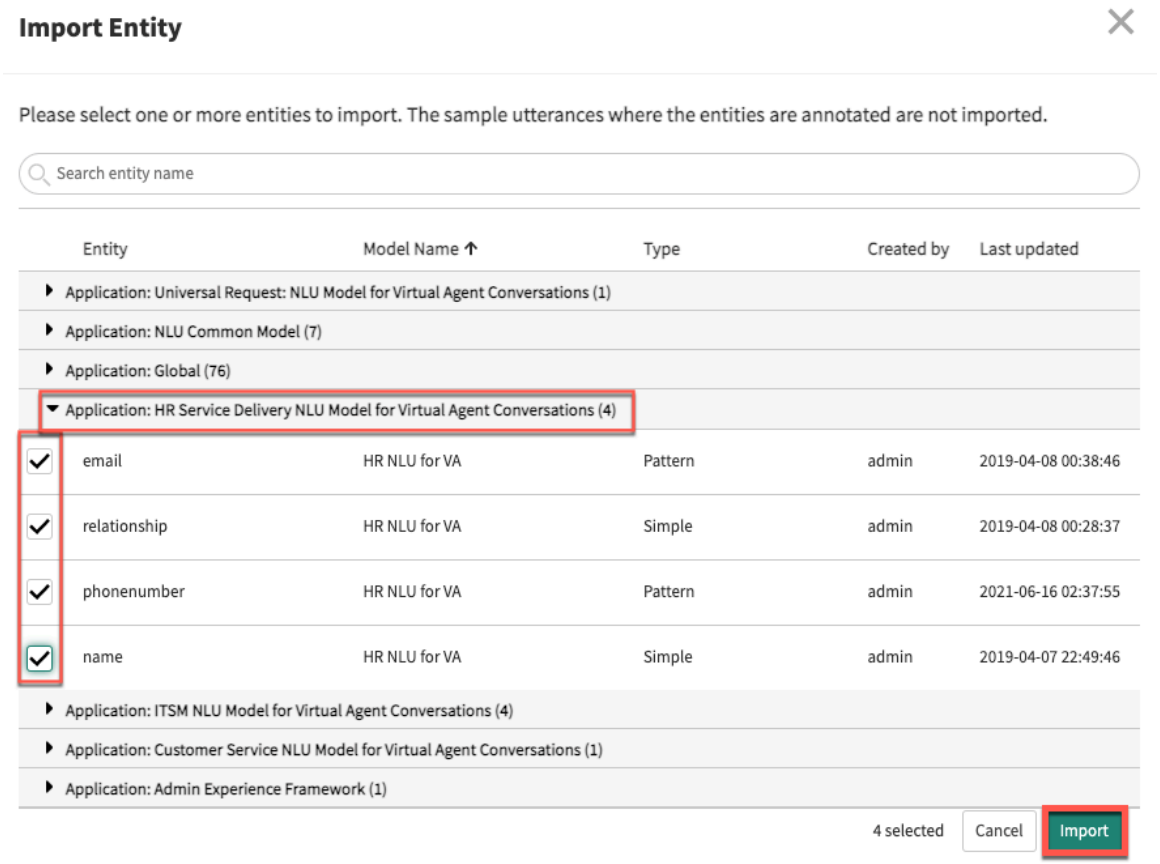
Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab corresponding to your model's application, then select the name of the model you want to add entities to.
3. In the **Build and train your model** phase, select the **Entities** tab.
4. Select **Import entity**.



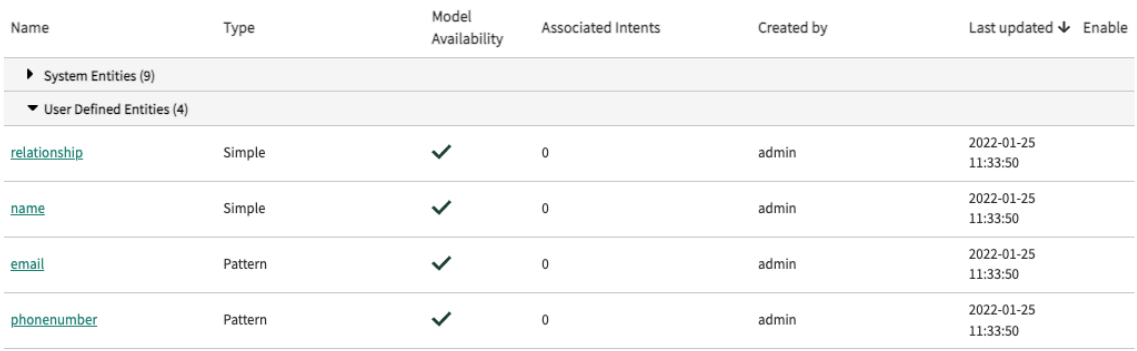
5. In the Import Entity window, select the entities you want to import.
   In this example, you click the drop-down arrow for Application: HR Service Delivery NLU Model for Virtual Agent Conversations. Select the four entities.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Note: Importing entities does not import any utterances that are annotated with those entities. You must annotate your current model's utterances using the new imported entities, then retrain the model.

6. Click **Import**.
   The selected entities import to the model. The entities appear under **User Defined Entities** in the Entities tab. By default, **Model Availability** will be enabled so that all intents in the model can use the entity.



What to do next

Use the imported entities to annotate the utterances of the model. Train your model to save your changes.

Related topics

Import an NLU intent

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Using regular expressions in entities

Learn how to use regular expressions in your NLU entities to establish patterns that help the system locate, match, and manage text.

Pattern entities use regular expressions (regex) to match any pattern of text, such as the format of an email address, a phone number, or an incident or case ID.

Note: ServiceNow uses and supports Java regex exclusively and not other vendor options, such as Perl regex.

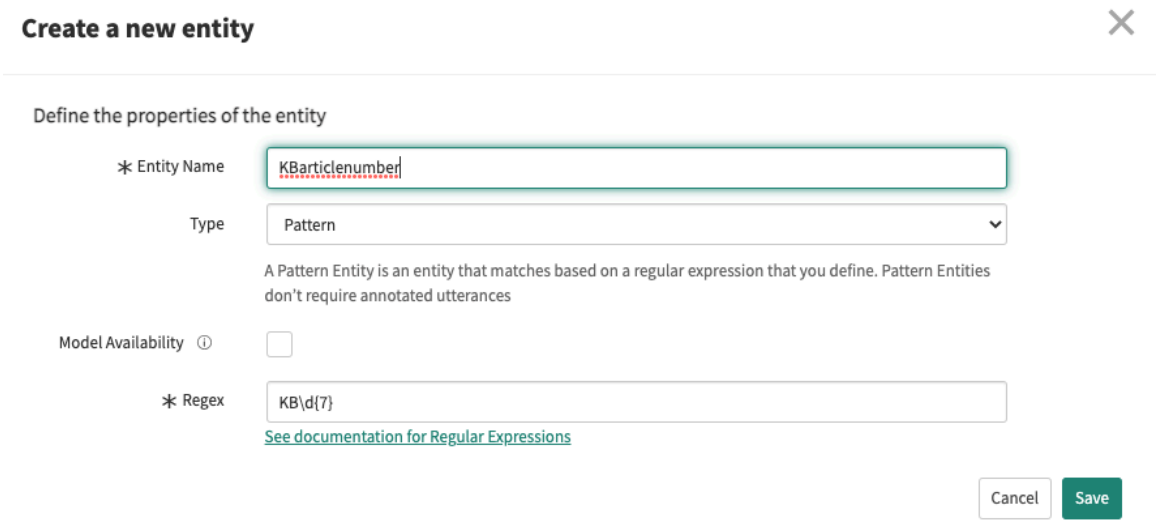
Regex examples

For a full example of regex code used in a pattern entity, see Create a pattern entity.

The following four examples provide the format and regex code for different pattern entities.

Knowledge base article

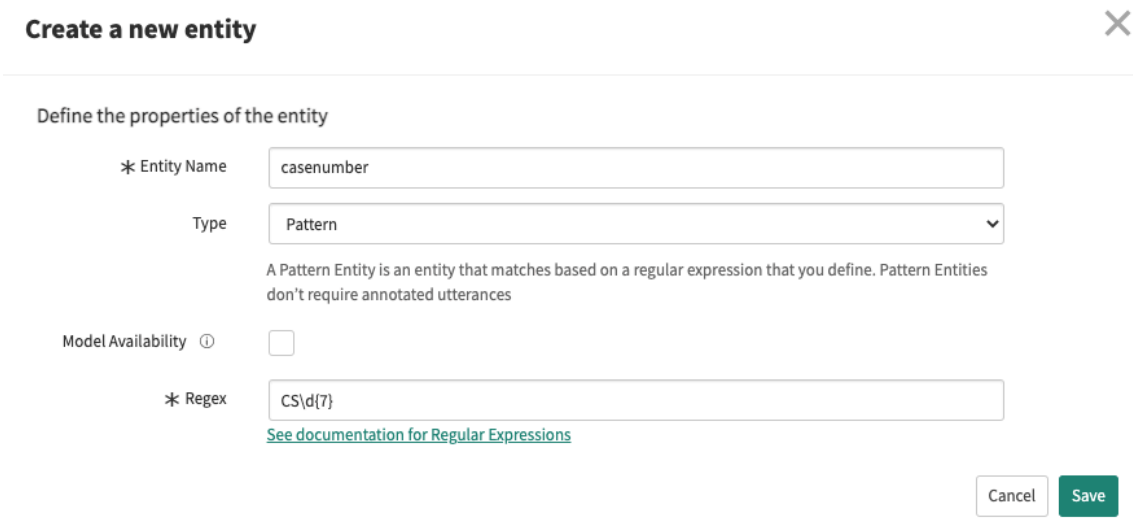
ServiceNow Knowledge Base (KB) articles use the format KB1234567. The regex code is KB\d{7}, where KB = knowledge base record and d{7} = 7 digits.



Case number

Case tasks use the format CS1234567. The regex code is CS\d{7} where CS = case record and d{7} = 7 digits.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Email address

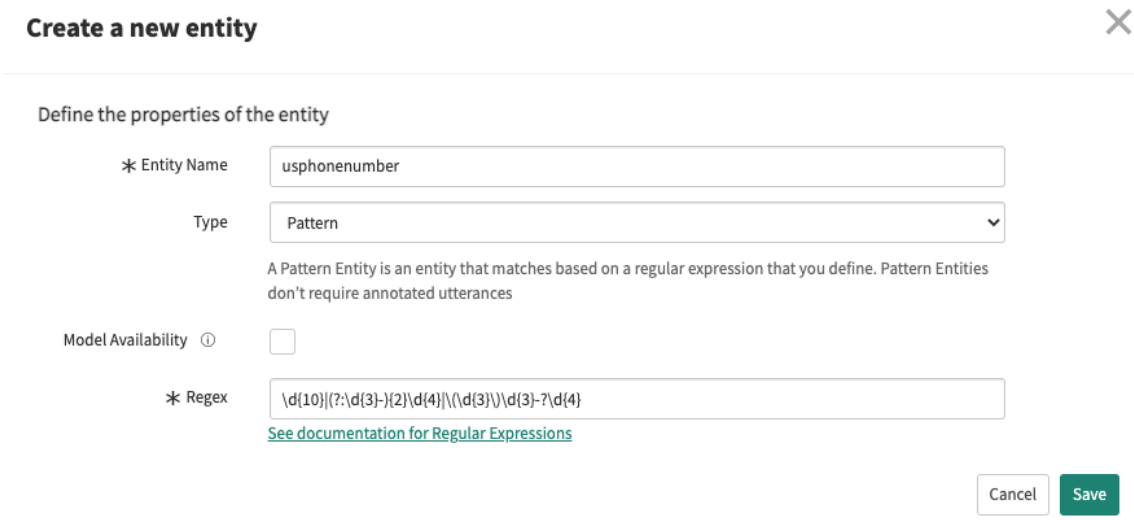
Email addresses use the format name@servicenow.com. The regex code is \b[a-zA-Z0-9&*/_.+-]+@[a-zA-Z0-9-]+\.[a-zA-Z0-9-.]+\b, which supports these formats: name@servicenow.com and name@servicenow.co.in.



Phone number

United States phone numbers use the format 555-123-4567. The regex code is \d{10}|(?:\d{3}-){2}\d{4}|\(\d{3}\)\d{3}-?\d{4}, which supports these formats: 5108882062,510-888-2062, and (510)888-2062.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Regex resources

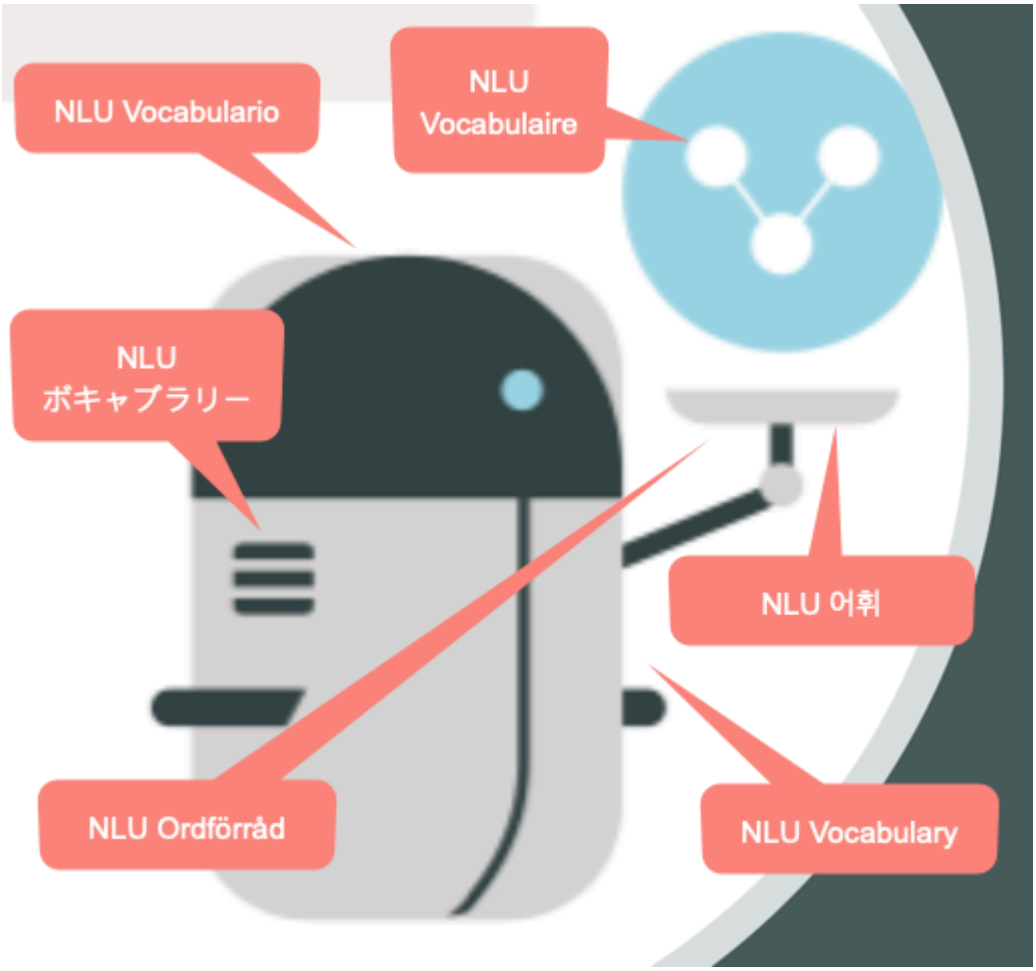
For further information on regular expressions, see the following external links.

- Java Regular Expressions
- Java Regular Expression Tester
- Pattern (Java Platform SE7)
- Java regex match abbreviations

NLU vocabulary

Use NLU vocabulary items to help the system recognize the various ways your users express their requests. Use vocabulary sources to help the system recognize objects in tables or lists, such as names of conference rooms or catalog items.

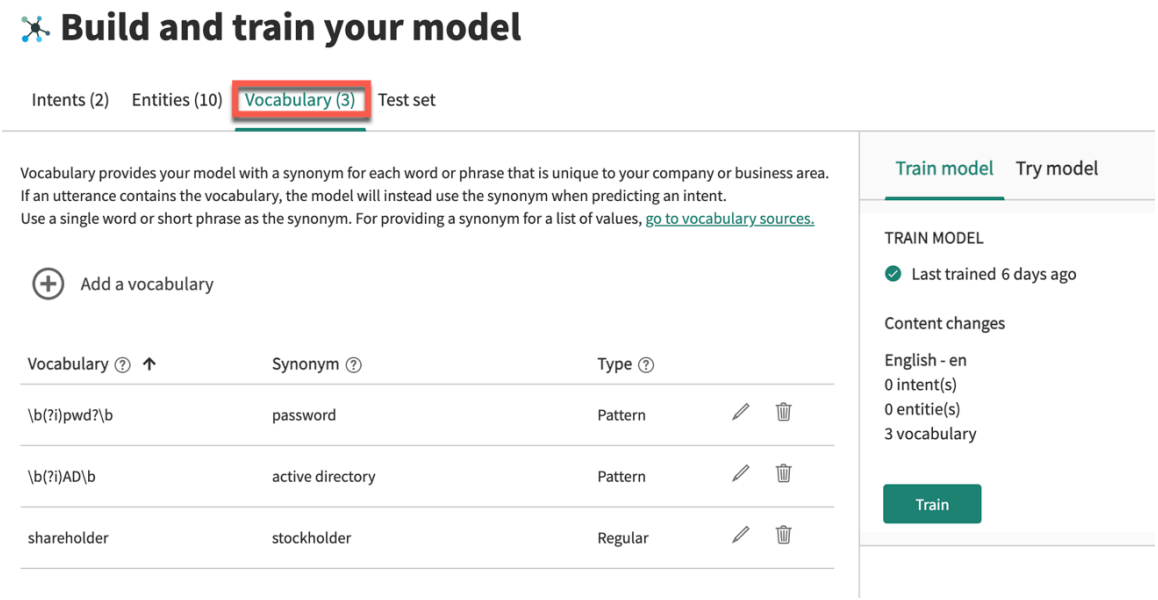
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Vocabulary usage and context

Vocabulary helps your model with the various words and phrases that it may encounter from your users. Since humans are not computers, they may come up with different ways to say the same thing. For example, if one user inputs an acronym rather than the full phrase, the system might be able to predict the correct intent by using the context in the utterance. However, by defining the acronym with a vocabulary item, you raise the model's confidence level and capability to predict intents correctly.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Vocabulary items cover the varying words or phrases that might appear in utterances. The vocabulary items are mapped to the synonyms that you provide, for intent prediction. This helps the system recognize the various ways your users express their requests, while at the same time ensuring the consistency, confidence, and accuracy of predictions.

Vocabulary item types

NLU provides the following vocabulary types that you can use to create and configure a vocabulary item.

Type	Definition
Regular	A word or phrase that is not commonly known, such as a business or industry specific term or acronym. Regular vocabulary is case insensitive, so all case variations will be captured by the vocabulary you create. See Create a regular vocabulary item.
Pattern	A regular expression (regex) that can capture specific formats such as email addresses and phone numbers. See Create a pattern vocabulary item.

Vocabulary sources

You can also use vocabulary sources to cover various objects that your users might request. Vocabulary items and vocabulary sources differ in their usage:

- Use a vocabulary item for an individual word, phrase, or pattern that can easily be mapped to a single synonym for the model to use.
- Use a vocabulary source to reference a ServiceNow table or list so that the values can all be replaced by the single synonym you define. Vocabulary sources can be reused across all your NLU models.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use vocabulary sources when your user utterances refer to objects in a set. The sources can be referenced from a list of values from a ServiceNow table or a static list that you define, such as a list of catalog items or conference rooms. Once you create and save these sources, you can use the @ symbol to specify them in training utterances. You can also use these sources as entity values.

For example, say you already have a list of all the conference rooms for your office. You can create a vocabulary source to reference that list rather than typing each conference room name into your intent's training utterances.

To create vocabulary sources, refer to Create a table vocabulary source and Create a list vocabulary source.

Vocabulary usage in relation to an intent

Here's an example of how an intent can interact with the vocabulary in its utterance examples.

Note: In NLU vocabulary, the synonym replaces the vocabulary that appears in the utterance.

Intent: OrderSoftware

Utterance	Issue and Solution
"I need to access sfcrm"	• Issue: the system doesn't recognize the acronym sfcrm, so can't accurately predict the intent. • Solution: Add sfcrm as a regular vocabulary item, and provide a synonym such as CRM software.
"I need to install Word"	• Issue: The term word is very common and does not necessarily indicate a software product. The term install may provide helpful context, but the confidence score may be too low to predict the OrderSoftware intent. • Solution: Create a pattern vocabulary item with a regex for capitalized Word, so that the system can recognize it as a software product. This added constraint makes it more likely that the OrderSoftware intent will be predicted. **Note:** To extract the specific software name to use in a Virtual Agent conversation topic, annotate it as an entity in the utterance.

Use a single word or short phrase as a synonym for best results. You can map multiple vocabulary items to one synonym. Do not map one vocabulary item to multiple synonyms.

For more context and examples, see Create an NLU intent.

A Regex example for a pattern vocabulary item

Let's say that you want to use a vocabulary item to identify the acronym IT, and map it to the synonym information technology.

Regular vocabulary items are case-insensitive by default. This means that a regular vocabulary item would match both the acronym IT and the common word it. So you decide to use a pattern vocabulary item.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The regex pattern that you can use is `\bIT\b`. The `\b` is a word boundary marker. This ensures that the pattern would not match ITSM or JIT, for example. The default case sensitivity of pattern vocabulary items means that `\bIT\b` would not match the common word it.

Regex details: You can turn off case sensitivity in a pattern by using `(?i)`. You can end that mode by using a minus sign as in `(?-i)`. For example, `(?i)te(?-i)st` should match both test and TEst, but not teST or TEST.

For more information, see Using regular expressions in entities.

Related topics

Create an NLU intent

Create a regular vocabulary item

Add a word or phrase that your users might use, and match that vocabulary item to a synonym. Your model uses the synonym during intent prediction.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model for Virtual Agent or AI Search.
- Role required: nlu_editor, nlu_admin, or admin. The editor must be assigned to the model.

About this task

Regular vocabulary items provide the model with a synonym for words or phrases your users might use in an utterance. The model uses the synonym to replace the vocabulary when predicting the intent. Use a single word for the synonym when possible.

Regular vocabulary items are case-insensitive by default. If you need to create a case-sensitive vocabulary item, use a pattern vocabulary item. For more information, see Create a pattern vocabulary item.

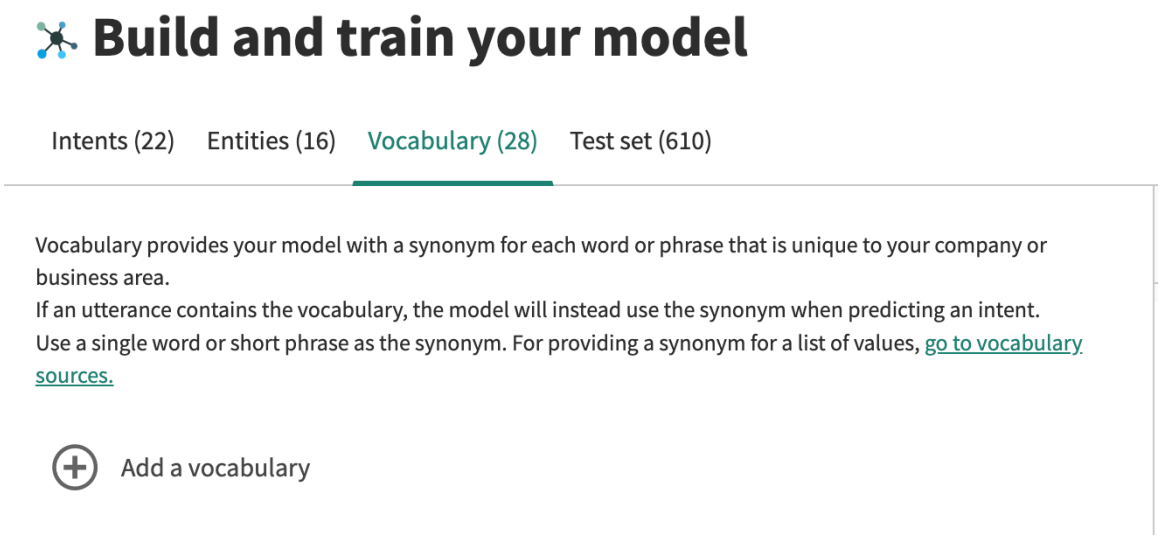
Note: Choose a synonym that is a commonly-occurring word in the same language as your model.

In this example scenario, you are adding a vocabulary item for the word credentials. Say that your users may use credentials to refer to their password. By creating a vocabulary item, you can make sure that the system correctly predicts the intent for an utterance such as reset my credentials.

Procedure

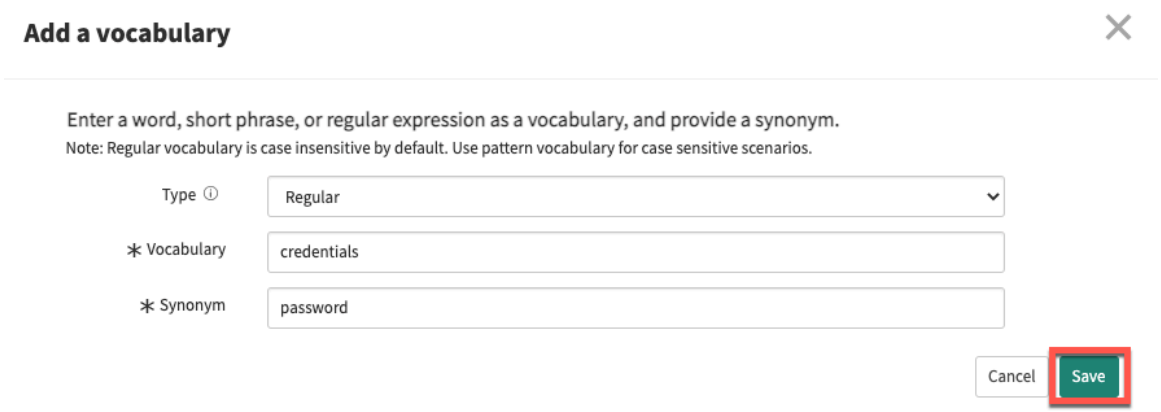
1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab corresponding to your model's application, then select the name of your model.
3. On the Model details tab of the model overview, select the **Vocabulary** card.
4. In the Vocabulary tab, select **Add a vocabulary**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



5. In the **Add a vocabulary** window, select **Regular** as the Type.
6. Add a vocabulary word or phrase that your users might use, and then add the synonym that the model should use for intent prediction.

   In this example procedure, add credentials as the vocabulary and password as the synonym.



7. Select **Save**.

What to do next

To deploy your new vocabulary item, train and publish your model again.

Add more vocabulary items to improve model coverage and accuracy.

Create a pattern vocabulary item

Use regular expression (regex) encoding to establish a pattern format for vocabulary items such as email addresses, phone numbers, and record naming conventions. You can create your own patterns for the vocabulary data in your instance.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model for Virtual Agent or AI Search.
- Role required: nlu_editor, nlu_admin, or admin. The editor must be assigned to the model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

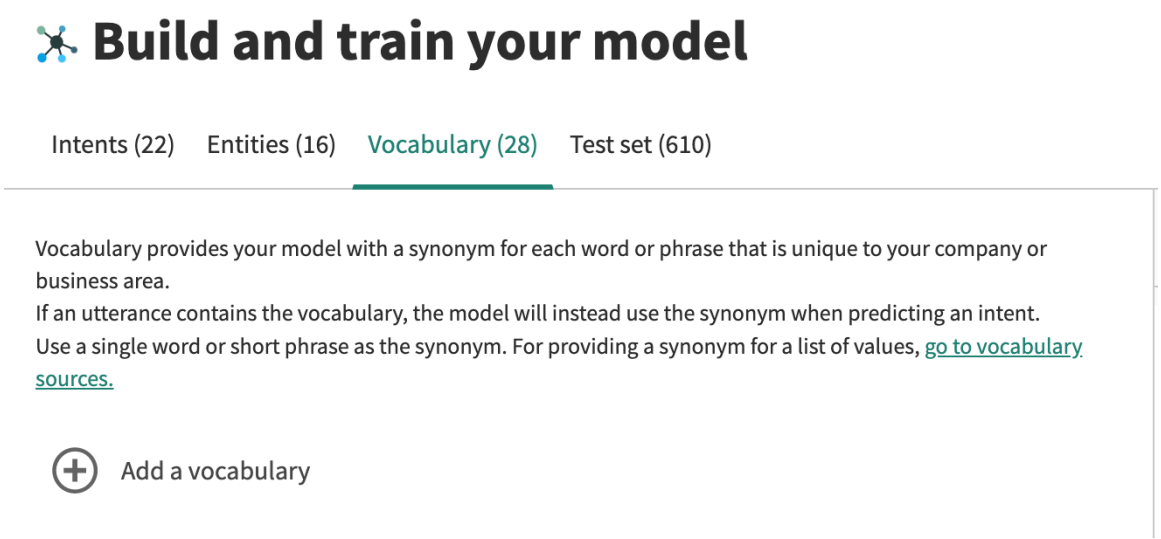
Pattern vocabulary items use regular expressions (regex) to find variations of words and phrases in your users' utterances. The pattern then maps these variations to the synonym you provide. The synonym is used during intent prediction.

In this example scenario, you've created an NLU Virtual Agent model for requests related to authentication. To refer to multi-factor authentication, your users often use "MFA", "mfa", or "mfa authentication". You can provide a pattern to cover these variations. They are replaced with the matching synonym multi-factor authentication when predicting an intent.

See Using regular expressions in entities for further details on regular expressions.

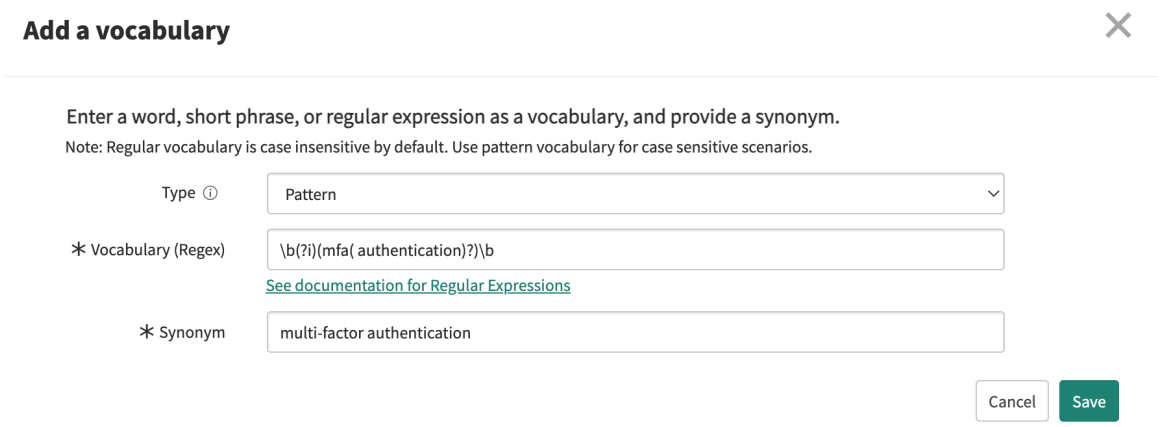
Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab corresponding to your model's application, then select the name of your model.
3. On the Model details tab, select the **Vocabulary** card.
4. Click **Add a vocabulary**.



5. In the **Type** field of the **Add a vocabulary** window, select **Pattern**.
6. Add a regular expression that covers words or phrases in your users' utterances, and a synonym the model uses to predict the intent.
   In this example, the regular expression covers several variations of "mfa authentication".

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



7. Click **Save**.

   Result: Your pattern vocabulary item appears in the Vocabulary section of your Model screen. Each time a user's utterance contains one of the spelling variations, the system recognizes the meaning as multi-factor authentication, in this example.

Related topics

NLU vocabulary

Using regular expressions in entities

Create a list vocabulary source

Create a list of words or phrases to act as a vocabulary source. The values in the list source are replaced by the synonym if they are detected in a user utterance.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Role required: admin, nlu_admin, or nlu_editor

About this task

Creating a list vocabulary source enables your models to interpret all the alternate and actual values of the list if they appear in a user utterance. The model interprets those values as the synonym you provide when predicting an intent.

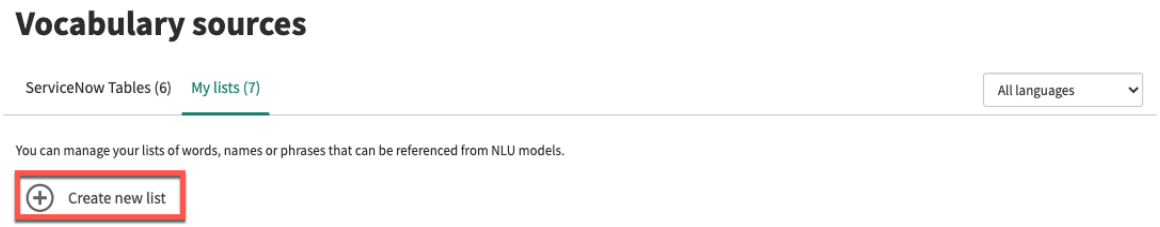
Note: During intent prediction, the synonym you provide replaces the value. During entity prediction, the actual value is used as the entity if the actual value or one of the alternate values are detected in the utterance.

In this example procedure, you're creating a list vocabulary source for your company's meeting rooms. You add the names of the rooms and alternatives your users might call them. Once you create the list, you add it to an utterance in an intent for the model. The model interprets the intent (for example, #BookMeetingRoom) and uses the name the user entered as an entity (for example, Everest).

Procedure

1. Navigate to **All** > **NLU Workbench** > **Vocabulary Sources**.
2. Click the **My lists** tab.
3. Click **Create new list**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

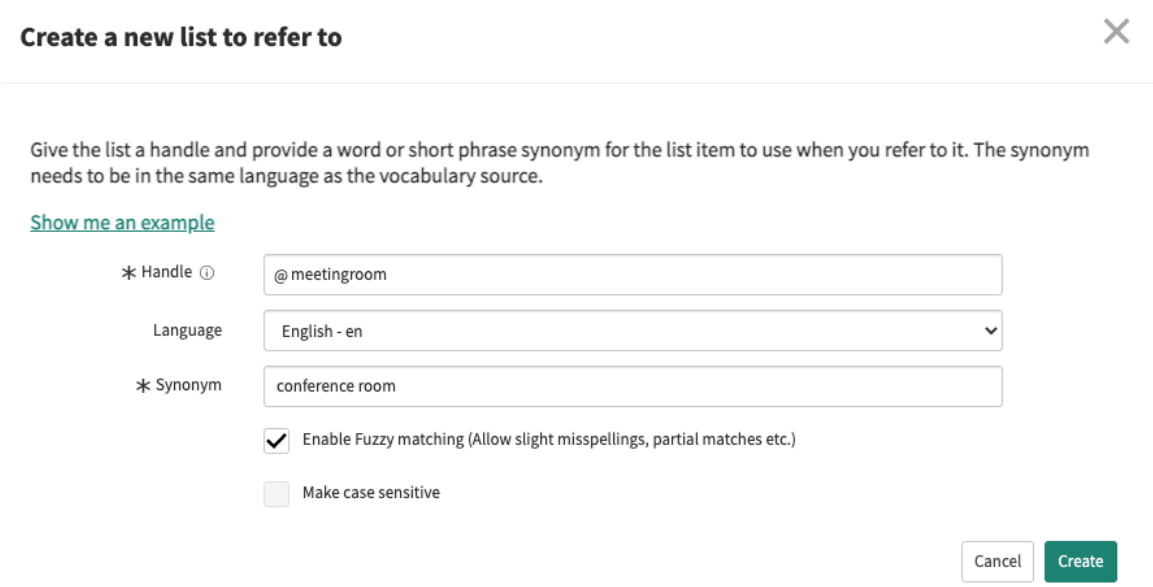


4. In the **Create a new list to refer to** window, configure the fields.

Field	Description
Handle	Name for the vocabulary source. Used to refer to in an utterance.
Language	Language of the vocabulary source. Synonym must be in the same language.
Synonym	Word or phrase the model uses during intent prediction. Choose a commonly-occurring word in the same language as your model.
Enable Fuzzy Matching	Check this box if you want the items on the list to match when a user utterance has slight misspellings.
Make case sensitive	Check this box to make the values in the list case sensitive. Utterances with wrong cases won't match.

For this example, use the following configurations:

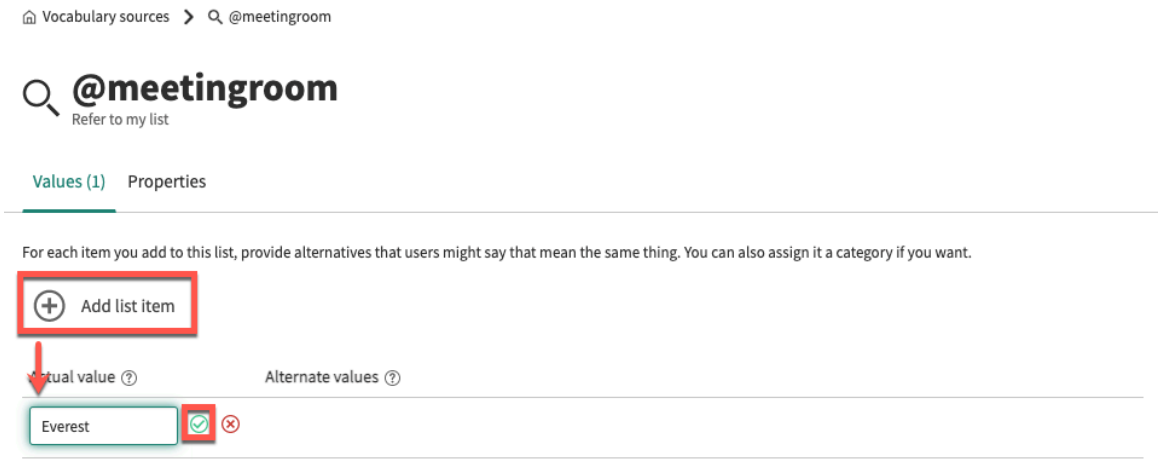
- **Handle**: `@meetingroom`
- **Language**: `English - en`
- **Synonym**: `meeting room`
- **Enable Fuzzy Matching**: Select the box.
- **Make case sensitive**: Leave the box clear.



5. Click **Create**.
   Your list vocabulary source draft appears in the My lists section of the Vocabulary sources screen.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Click the name of the list vocabulary source.
7. Click **Add list item**.
8. Enter a value for the list and click the green check mark.



In this example, you enter `Everest`.

9. Double-click the area under **Alternate values** to add alternatives separated by a comma.



10. Train the model to make the list vocabulary source available.

What to do next

Add the rest of the break room names and alternatives.

You must retrain the model after updating a list vocabulary source. For more information, see Train and try your NLU model.

Then you can use the list vocabulary source when annotating a training utterance. Use the @ symbol with the handle to refer to this vocabulary source.

Create a table vocabulary source

Use the values from a ServiceNow table as a vocabulary source. Your Natural Language Understanding (NLU) models use your provided synonym to interpret utterances that contain values from the chosen source fields of the table.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Role required: admin or nlu_admin

About this task

Add a ServiceNow table to be used as a vocabulary source. Select one or more source fields of the table. Then, provide a synonym to be used by the model when interpreting utterances that contain values from the table. Note that the **Reference** field is not supported as a source field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When you create and sync a table vocabulary source, the values from the table in your instance are extracted, and a vocabulary source is created in the NLU Service. For more information on the NLU Service, see NLU Service Updates.

Note: You should not create multiple vocabulary sources that reference the same table and fields. Create only one vocabulary source (and provide a synonym) for the particular table and fields you want to reference in the model. Otherwise, it would cause confusion in the model and interfere with prediction quality.

In this example procedure, you want your models to be able to interpret utterances with locations. The ServiceNow location (cmn_location) table already contains information regarding important locations. You set up the location table to be used as a vocabulary source.

Procedure

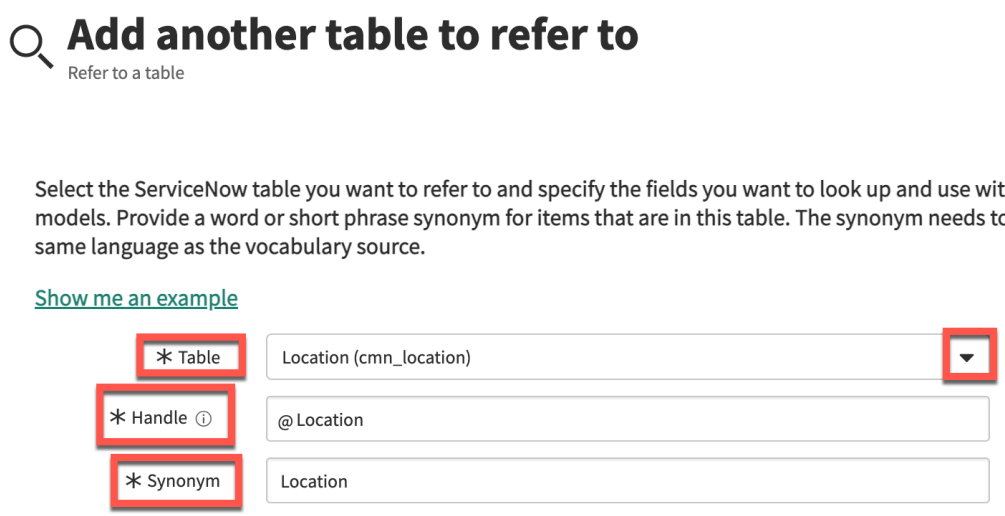
1. Navigate to **All** > **NLU Workbench** > **Vocabulary Sources**.
2. Click **ServiceNow Tables**.
3. Click **Add another table**.
4. On the Add another table to refer to page, configure the fields.

   a. Select the table, handle, and synonym for the table.

   In this example, use the following configurations:

   - **Table**: Select **Location (cmn_location)**.
   - **Handle**: Use the system-generated handle, @Location.
   - **Synonym**: Use the system-generated synonym, Location. You can update this value to ensure that it's a synonym for each individual item in the vocabulary source.

     Note: Choose a synonym that is a commonly-occurring word in the same language as your vocabulary source.



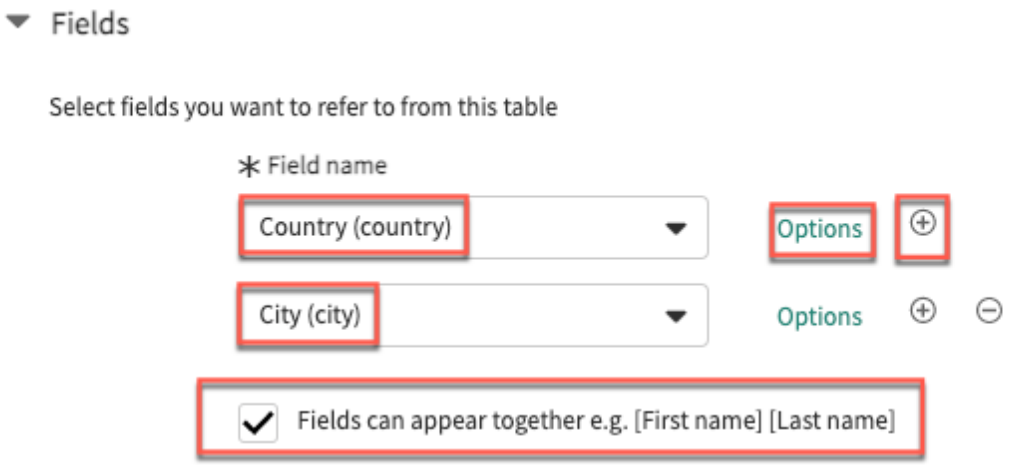
   b. Select the fields you want to refer to from the source table.

   Note: To add multiple fields, select the plus icon.

   In this example, use the following configurations:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Field name**: Select the **Country (country)** and **City (city)** fields.
- **Options**: Click **Options**, then select the **Use this field to look up values** check box for both of the **Field name** values. Checking the box makes it so either a country name or a city name can be used in an utterance for finding a record in the vocabulary source. You can also use multiple comma-separate values in this field, for example, if your city column has multiple names, such as NYC, New York, and New York City.
- **Fields can appear together**: Leave the check box empty. This field is typically used in an NLU Search model, where users can enter words next to each other in an utterance to find a record. For example, you can enter `Pierre Development` while searching for an employee named Pierre who works in the Development team, where **Name** and **Team** are two fields in the table.



c. Configure the advanced options.

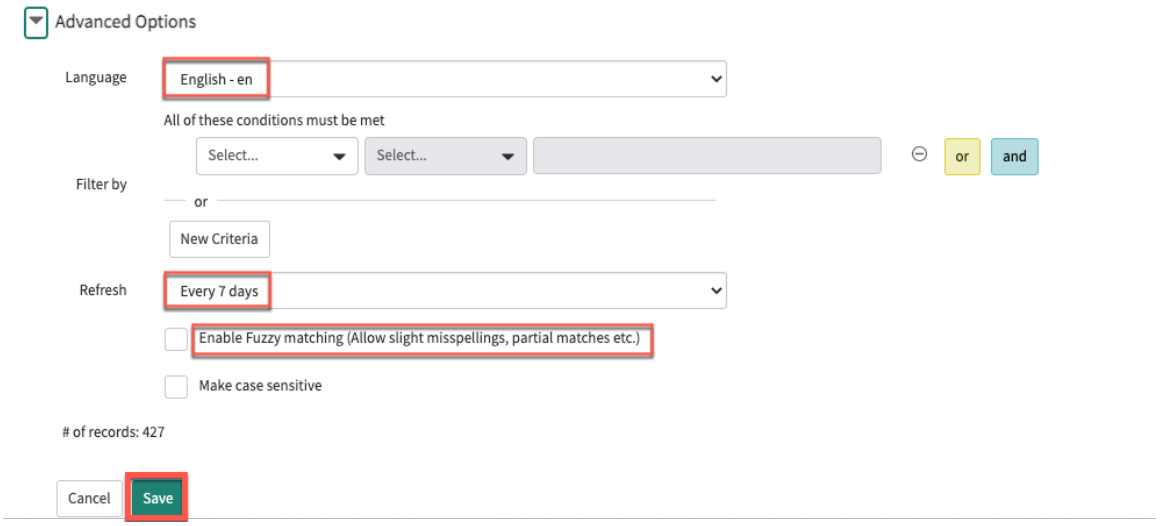
In this example, use the following configurations:

- **Language**: Select **English - en**.
- **Filter by**: You can use the condition builder to filter the values in the source table for the vocabulary source. Leave empty for this scenario.
- **Refresh**: Select **Every 7 days**. By selecting this, you set the system to get new values from the table every seven days. There are different refresh options you can choose based on how often the source table is updated.
- **Enable Fuzzy matching**: Checking this box make it so a record can be matched even if an utterance contains a slightly misspelled word or part of a word. For example, when searching for a city such as Kansas City, you might enter Kansas City or just Kansas, and the system will still be able to match with the correct location record. Fuzzy matching may

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

return false matches in some cases. Use fuzzy matching sparingly and test out your model before using it.

- **Make case sensitive**: If you don't enable fuzzy matching, you can check this box to make the values in the source table case sensitive. If an utterance doesn't use the correct cases, the model doesn't predict the intent.



5. Click **Save**.

   Result: Your **@Location** table vocabulary source appears on the Vocabulary sources screen, where it begins to sync with its source table.

6. If the sync doesn't start immediately, select the **Sync Lookup** icon on the far right to sync the source table data with the NLU Service.
   Result: When the sync completes, you can use the vocabulary source in your models. You don't need to retrain the model after updating a table vocabulary source.

What to do next

Add a vocabulary source to a model by annotating an utterance. Use the @ symbol when adding an example utterance to an intent.

Sync a table vocabulary source

Synchronize your table vocabulary sources to obtain the latest changes to the ServiceNow source table. Synchronizing your vocabulary sources ensures your NLU models have the latest values when predicting intents.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Role required: admin or nlu_admin

About this task

When you reference a vocabulary source in an utterance, it pulls the values at the time the model is trained. However, if the table values change over time, the model still references the values from the last training session.

You can select a schedule to automatically refresh the vocabulary values used by NLU. This schedule can be edited later. For more information, see Create a table vocabulary source.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

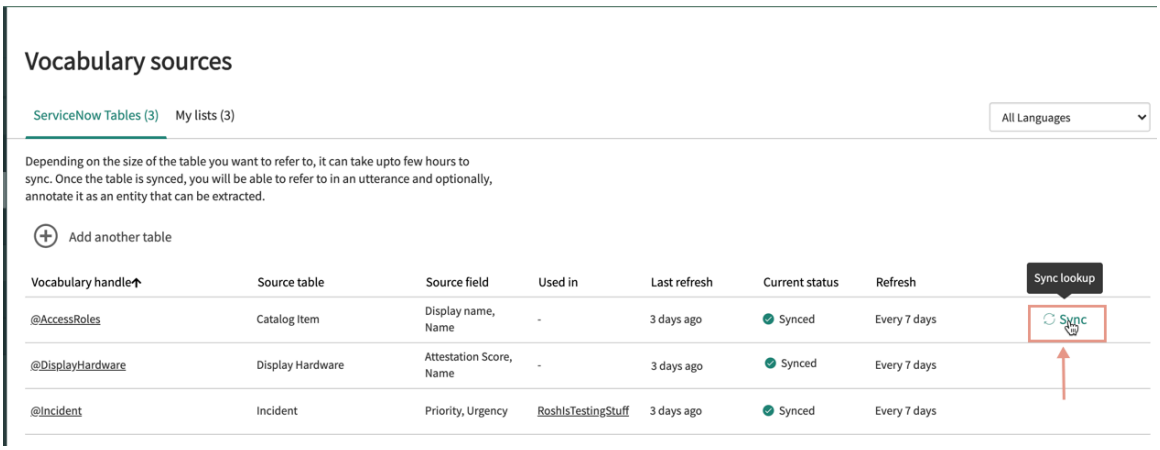
You can also manually sync a table vocabulary source, such as before model training.

Note: If an utterance references a table vocabulary source that has never been synchronized, the model fails to train. Check the vocabulary source's current status, and sync manually if **Never synced**.

In the following example scenario, you're manually syncing the @AccessRoles vocabulary source to the Catalog Item table.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Vocabulary Sources**.
2. In the ServiceNow Tables tab, point just to the right of the **Refresh** column to invoke the **Sync lookup** icon.

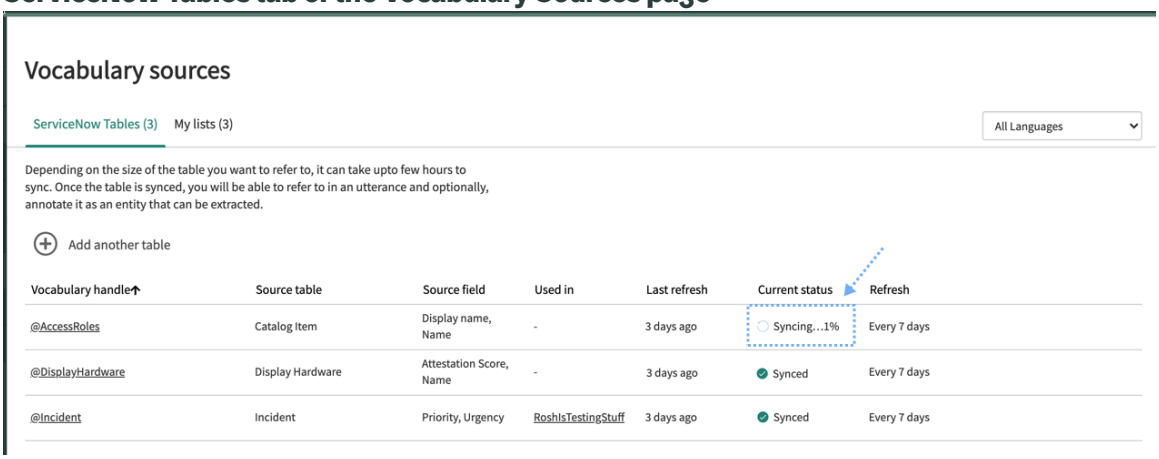


3. Select **Sync**.
   The vocabulary source starts synchronization.

What to do next

The syncing operation can take some time, depending on the source table size.

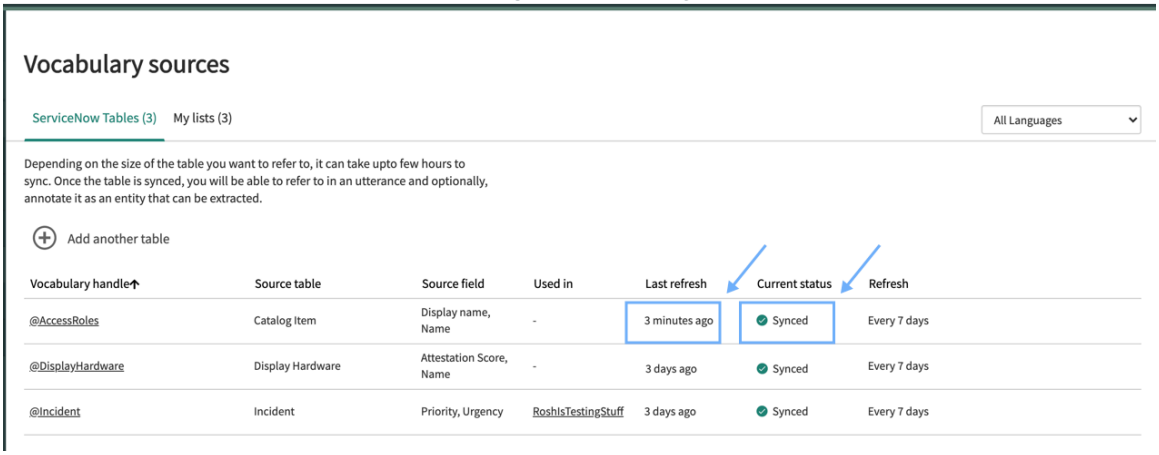
ServiceNow Tables tab of the Vocabulary Sources page



The values in the Last refresh and Current status columns reflect the current status of the vocabulary source.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ServiceNow Tables tab of the Vocabulary Sources page



Proceed to train your model. For more information on model training, see Train and try your NLU model.

For information about errors when syncing a table vocabulary source, see article KB1588239 in the Now Support Knowledge Base.

Related topics

Create a table vocabulary source

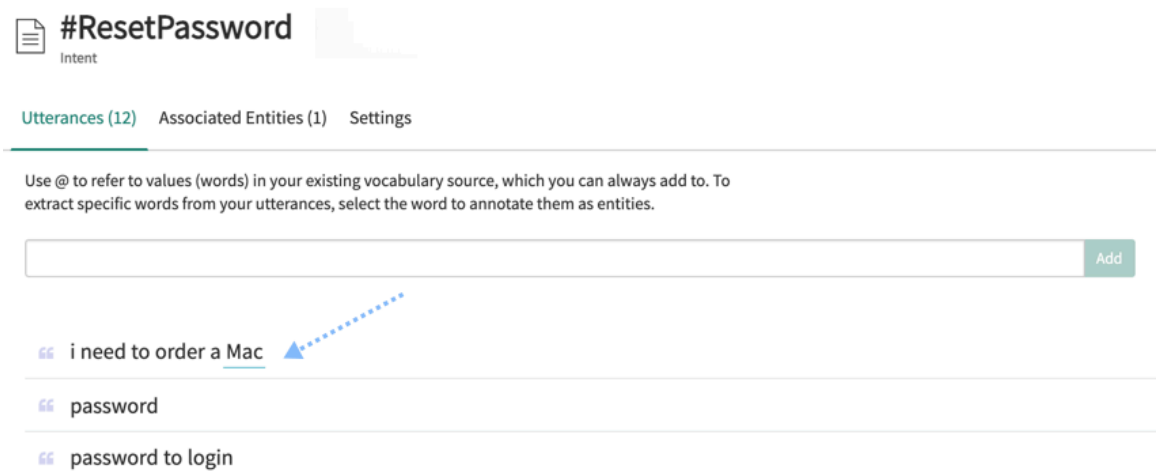
Pre-built vocabulary

Use ServiceNow® pre-built vocabulary for software and hardware terms so the system recognizes their multiple variations in utterances.

Your Natural Language Understanding models contain pre-built vocabulary settings you can use when adding an example utterance. The vocabulary provides definitions for software and hardware terms whether they're expressed in slang or professional usage. The vocabulary can also recognize product misspellings.

For example, for one of your example utterances, you enter `I need to order a Mac`. When the system recognizes a pre-built vocabulary item, the term has a blue line under it.

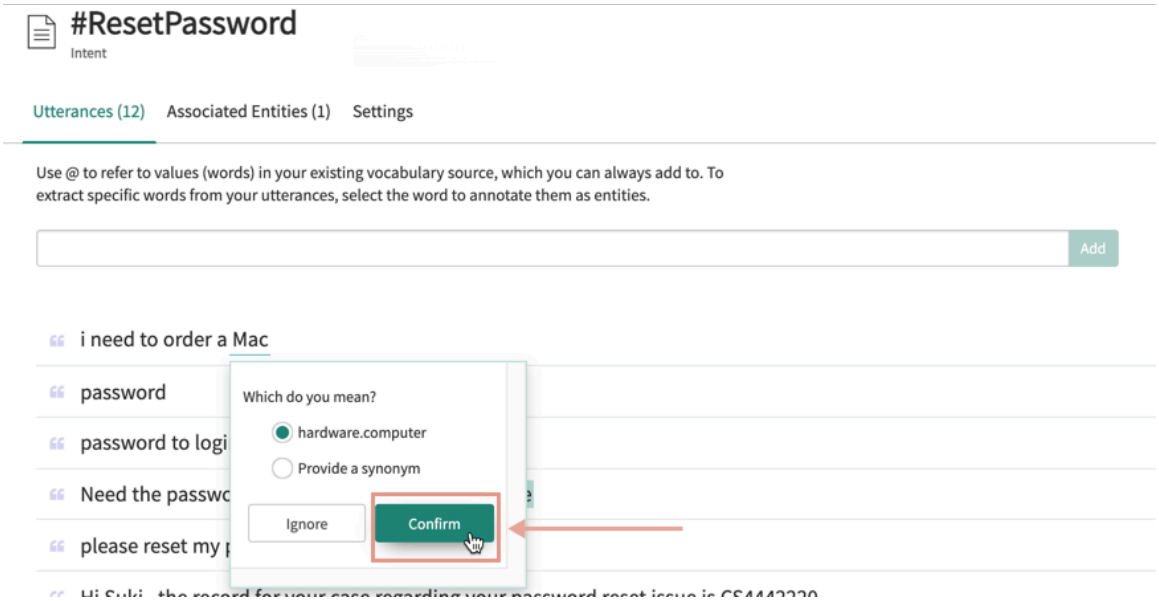
Utterance tab of the Intent details page



When you click the word, a window appears with two options to choose for the word:

- A pre-built suggested definition for the word
- An option to add a synonym

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



If you select the first option and click **Confirm**, the system uses the pre-built suggested definition and the blue line disappears.

If you select the second option, enter a synonym for the word, and click **Confirm**, the word and the synonym are added to the model vocabulary.

If you select either of the two options and click **Ignore**, the blue line disappears and the word remains as it was previously.

Test set creation and management

Use the default test set of your NLU model to test the model's performance and accuracy. Manage your test set over time by building or updating its content in the NLU Workbench.

Embedded video: Test set management in NLU Workbench

Note: To test your model, install the ServiceNow® Store application NLU Workbench - Advanced Features. For more information, see Install NLU Workbench - Advanced Features.

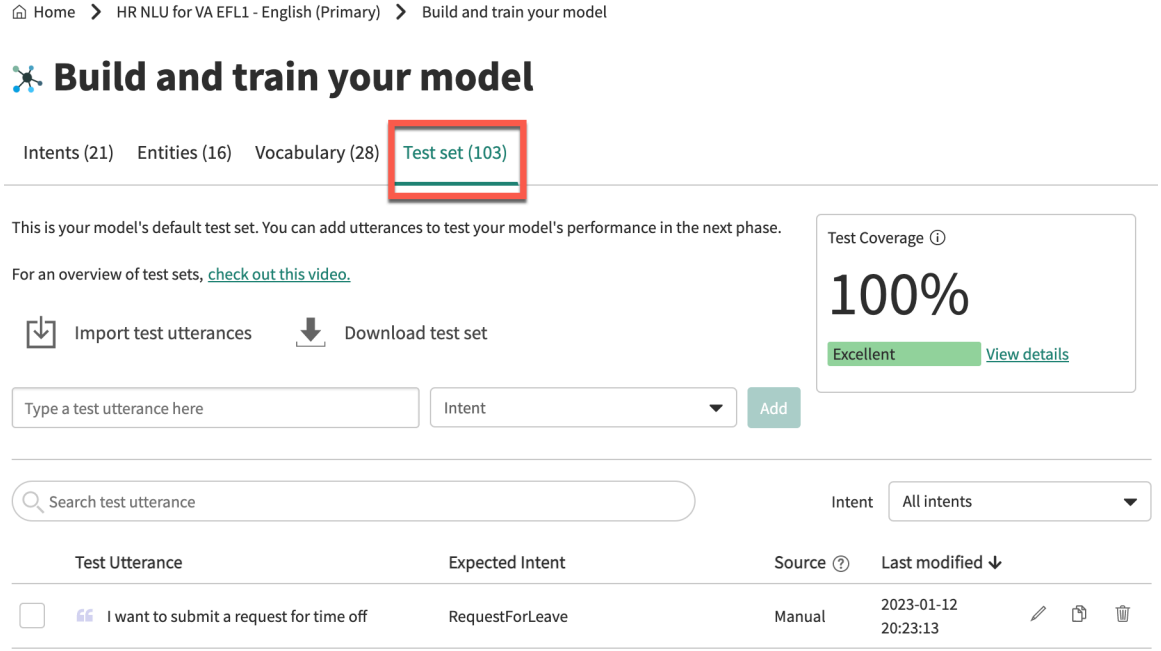
When you create an NLU model for Virtual Agent or AI Search, a default test set is created and associated to the model. You can use the default test set to evaluate the model's performance. Initially the test set is empty, ready to be populated with your content.

Access your default test set

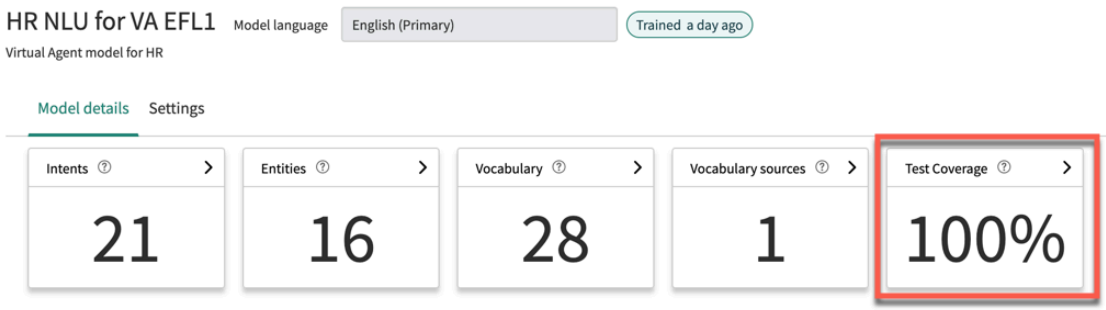
Access your default test set with one of the following methods.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Navigate to **All** > **NLU Workbench** > **Models**. Select the tab for your model's application, then the name of your model from the list. On the model's overview page, find the **Build and Train your model** card and select its **View phase** button. Then select the **Test set** tab.



- Navigate to **All** > **NLU Workbench** > **Models**. Select the tab for your model's application, then the name of your model from the list. On the model's overview page, select the **Test Coverage** tile.



- Navigate to **All** > **Multi-model Batch Testing** > **Test sets** tab. Find the name of your model. Default test sets are labeled as **Default**.



Add content to your default test set

Add utterances and their expected intents to build and manage your test set over time. You can add content to the default test set with the following methods:

- Add test utterances and their expected intents manually. From the model's overview page navigate to **Build and train your model** > **Test set** tab. Type your input into the Type a test utterance here field, select an appropriate intent, then select the **Add** button.

  These test utterances are assigned a source of Manual.

- Import test utterances and their expected intents from a CSV file or from other models. To import content to a default test set, from the model's overview page navigate to **Build and train your model** > **Test set** tab. Select **Import test utterances**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Imported test utterances are assigned a source of Manual.

- The Expert Feedback Loop feature lets you add actual user utterances from Virtual Agent chat logs to the test set.

  These test utterances are assigned a source of Expert Feedback. For more information, see NLU Expert Feedback Loop.

Test Coverage

The Test Coverage score is the percentage of a model's enabled intents that have test utterances in the default test set. Before testing your model, ensure that there is at least 60% coverage. The higher the Test Coverage score, the more accurate the performance testing results.

Your test coverage needs to be at least 60%, with at least 5 test utterances per intent, in order for the system to provide an optimal confidence threshold during batch testing. For more information about the confidence threshold, see NLU model settings.

Aim to have about 10 percent of a model's test utterances marked as "not relevant", meaning that there is no intent associated. This helps assess how the model handles irrelevant utterances which should not have any intent predicted. For more information about irrelevant utterances, see Irrelevance detection in NLU.

Use the test set

To use the default test set from the **Test and publish your model** phase, see Test and publish your model.

To use the test set in Multi-model Batch Testing, see Multi-model Batch Testing.

Characteristics of default test sets

When an instance is upgraded, default test sets are created for any existing models that don't already have them.

When you copy a model using **Duplicate this model**, the original's default test set is copied into the new model. For more information, see Duplicate an NLU model.

The utterances in the test set shouldn't be the same as the utterances in the training set.

Default test sets can't be deleted separately from their models.

Test set utterances should be in the same language as their model.

Test sets are available for Virtual Agent or AI Search models.

Downloading or moving default test sets

Default test sets can be downloaded or moved as follows.

- Default test sets can be separately downloaded in CSV format. To download the test set, from the model's overview page navigate to **Build and train your model** > **Test set** tab. Select **Download test set**.

  Note: Test sets that are downloaded from **Download test set** contain test utterances and their expected intents, but not the sources.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Default test sets can be moved with update sets. When you add an NLU model to an update set, its default test set is added, including test utterances, expected intents, and sources. For more information, see Add an NLU model to an update set.
- When using the **Export model as CSV** function in the **All existing models** table, the default test set is not included. For more information, see Export an NLU model.

Train and try your NLU model

Train and try your model iteratively so that its intents and entities are validated, compiled, and saved to your model.

Before you begin

- Make sure that the NLU Workbench - Core plugin, NLU Workbench plugin, and Predictive Intelligence plugin are all installed and activated.
- Create an NLU model. For more information, see Creating models.
- Create one or more NLU intents and their associated entities for your model. For more information, see NLU intents.
- If any utterance references a table vocabulary source, ensure that the source has been synchronized so that its values are available to your model. For more information see Sync a table vocabulary source.
- Role required: nlu_editor, nlu_admin, or admin. The NLU editor must be assigned to the model.

About this task

Training your model saves any changes you made to the content, and checks for conflicts or errors. Training also makes a model available for publishing.

After training, you can try your model by manually entering individual utterances to see what intents are predicted.

Note: To run a test of your model against a list of test utterances, see Test and publish your model.

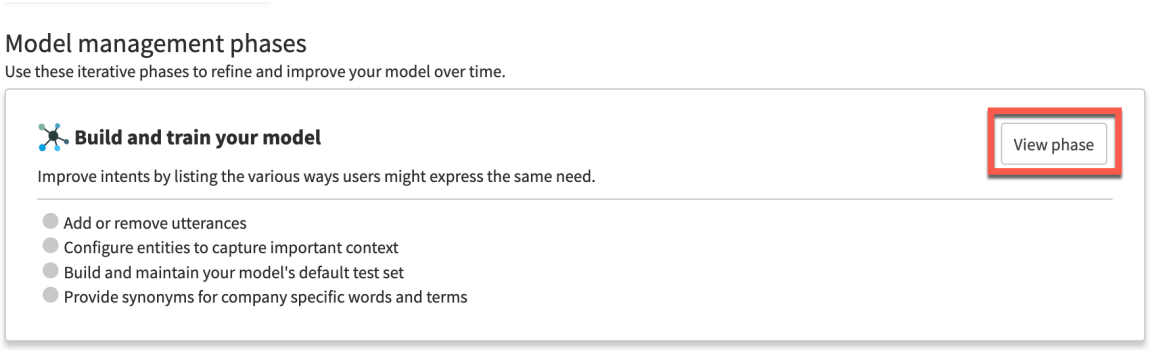
The mid-conversation responses of Dialog Acts can't be tried or tested in NLU Workbench.

In this example scenario, you've already built sufficient model content by adding intents, utterances, entities, and their associated annotations. Following the example procedure, you first train your NLU model. Then you try your model by manually entering utterances so you can check the prediction results and confidence scores.

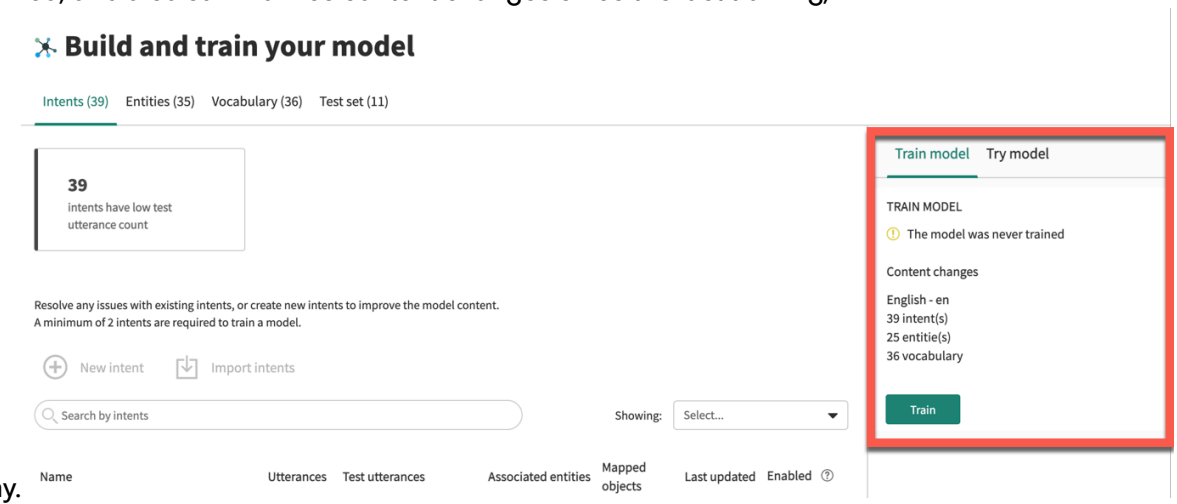
Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, then select the name of your model.
3. In the **Model details** tab of the model overview, make sure there is enough content in Intents, Entities, and Vocabulary.
4. On the **Build and train your model** card on the model overview, select **View phase**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



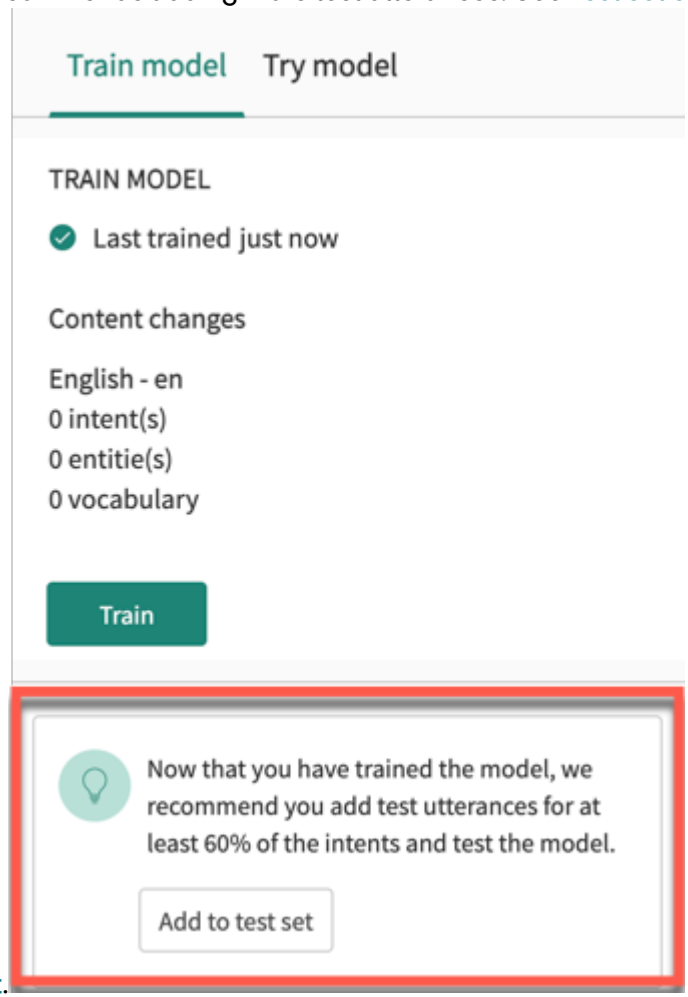
5. When the **Build and train your model** phase opens, ensure that the **Train model** tab is selected.
   Result: The **Train model** tab displays the last time the model was trained, and also summarizes content changes since the last training, if any.



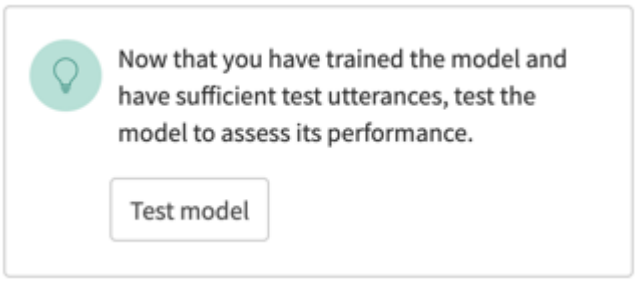
6. Select the **Train** button.
   Result: The system displays a progress bar during training. When finished, the system displays one of two recommendations:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- When less than 60% of the model's intents are covered in the default test set, the system recommends adding more test utterances. See Test set creation and management.



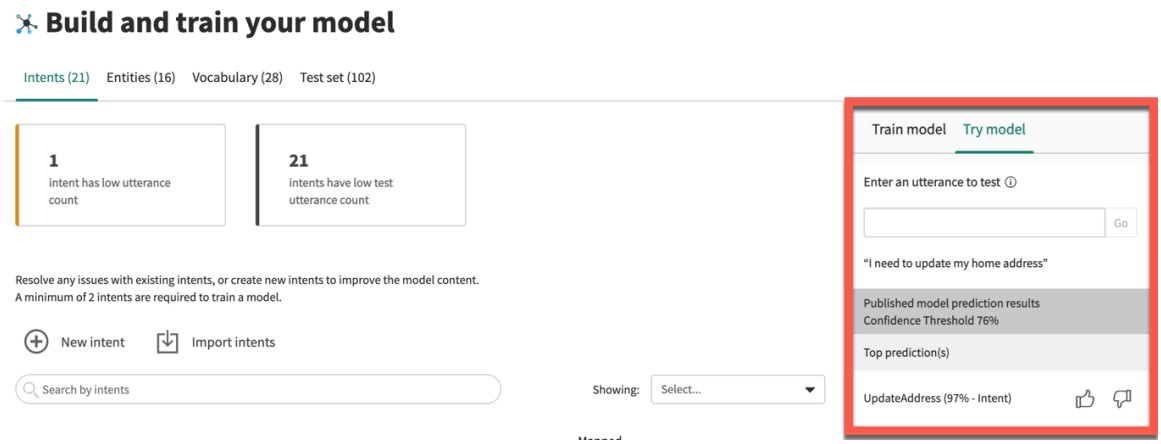
- When over 60% of the model's intents are covered in the default test set, the system recommends proceeding to testing. See Test and publish your model.



7. To manually try individual utterances, select the **Try model** tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. In the text field under **Enter an utterance to test**, type an utterance and select **Go**.



Result

In this example, you entered `I need to update my home address` as the utterance to try.

1. The system displays the model's confidence threshold, which is 76% in this example.
2. Under **Top prediction(s)**, the system displays all intents that were predicted with a confidence score greater than the threshold.
3. In the example, the intent UpdateAddress is predicted with a confidence score of 97%, which is greater than the threshold of 76%.

The **Try model** results also display thumbs-up and thumbs-down icons for you to provide feedback. For more information, see Test panel feedback.

What to do next

- Continue trying various utterances to check that your updates to model content are effective. See Compare draft and published versions of your NLU model.
- To test your model against a list of test utterances, use its default test set in the Test and publish your model phase, or navigate to Multi-model Batch Testing.
- To adjust the model's confidence threshold, use the **Settings** tab on the model's overview page. For more information, see NLU model settings.
- If you're satisfied with the results of your testing, Publish your NLU model.

Test panel feedback

When testing your NLU model on the Try model section of the test panel, use this feature to provide feedback on the model's intent predictions.

Summary context

When a model is trained and tested for an utterance and the model returns an intent prediction, you can provide a thumbs up or thumbs down rating on the predicted intent it returns. Marking a different intent prediction as correct adds the utterance to the corrected intent. All other feedback is captured for continual learning. The system then incorporates your feedback to optimize the model predictions. This feature requires the nlu_admin role to access and test the model. NLU editors can also access the test panel if an NLU admin assigns them to it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Providing prediction feedback

The ratings you provide help the system to match an intent to an utterance. These ratings are essential for the system to continuously learn, evolve, and improve the accuracy of the intent predictions based on user input. They also enable you to notify the system if the intent prediction is correct or not.

The following scenarios below show examples of how to interact with your model test panel and provide prediction feedback to the system. In all scenarios, you use these four steps:

1. In the **Build and train your model** phase of your model, select **Try model** to open the test panel.
2. In the test panel's **Enter an utterance to test** field, enter a brief utterance that's similar to a training utterance in one of the intents.
3. Click **GO**.

   Result: The system returns its predictions for your test utterance in the **Top Predictions(s)** section of the test panel.

4. Click the **Thumbs Up** icon or the **Thumbs Down** icon.

   If you want the system to know it has predicted the correct intent for your utterance, select the **Thumbs Up** icon.

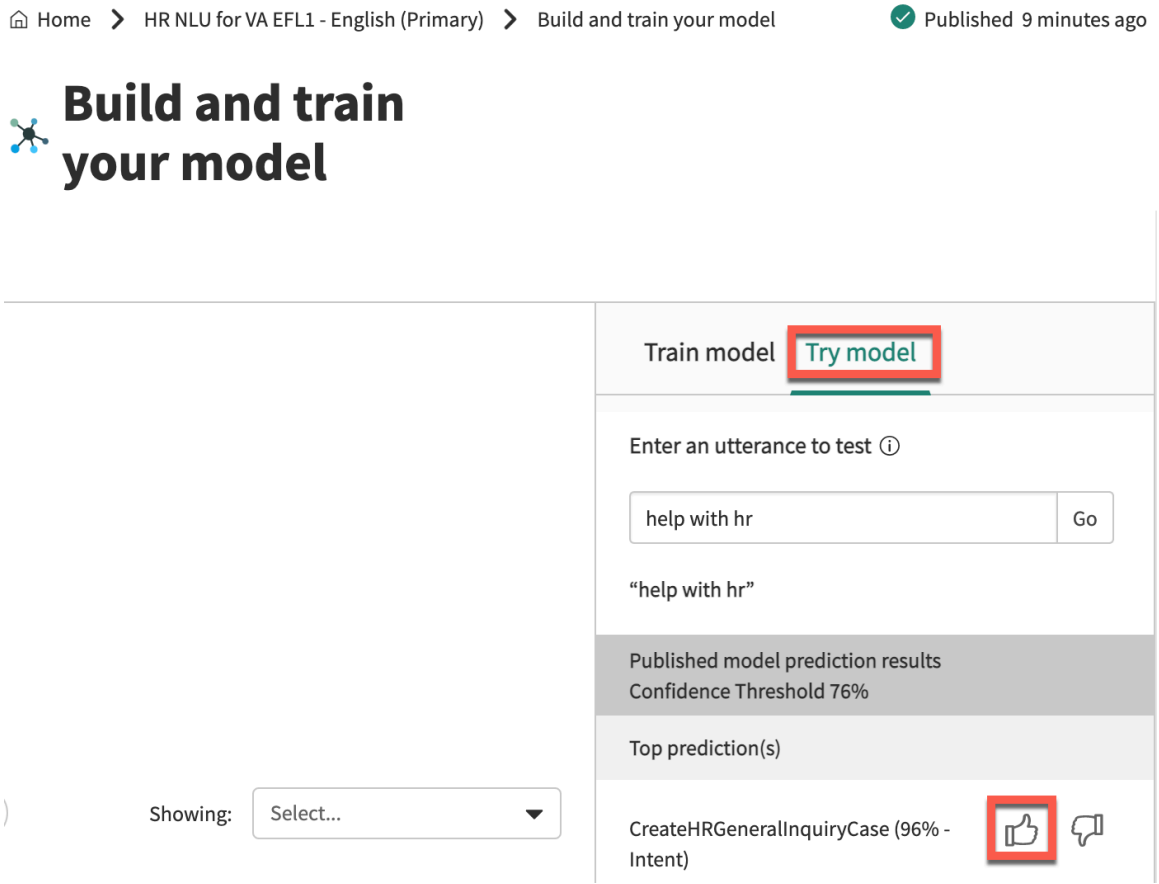
   In all other cases, select the **Thumbs Down** icon, which opens the **Provide feedback to improve this prediction** section. Here you can choose an intent other than the top predicted intent.

Scenario 1: On the Try model section of the test panel, you enter `help with hr` as the utterance. When the top prediction results appear, you're confident that the predicted intent is the correct match to your utterance. So in this case, you click the **Thumbs Up** icon.

Results:

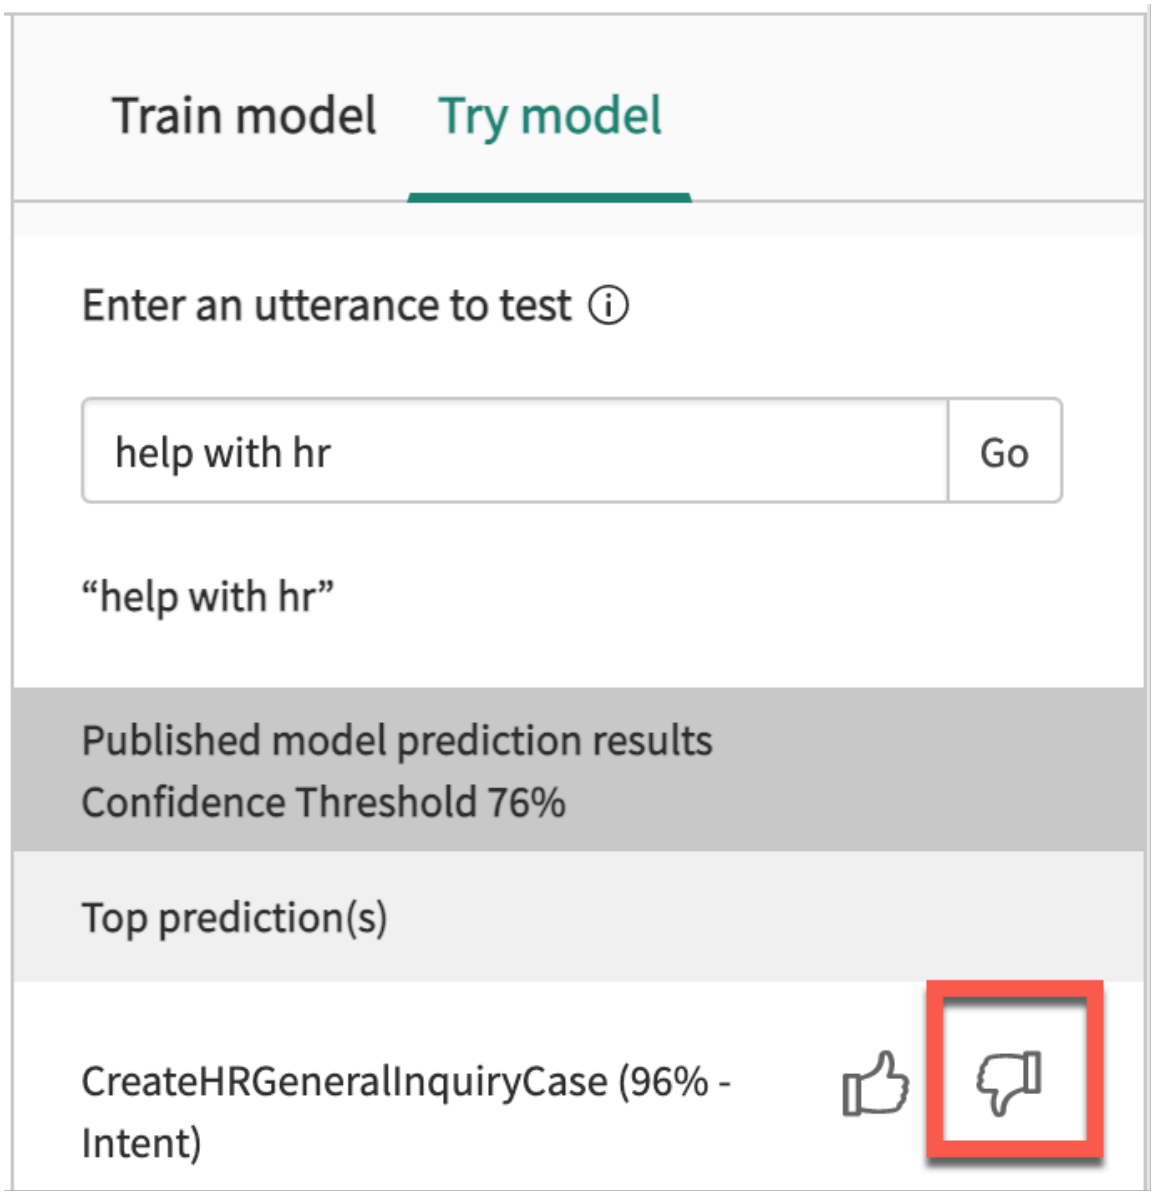
- The system predicted the correct intent, which in this case is **#CreateHRGeneralInquiryCase**.
- Your feedback notifies the system that it has matched the correct intent to your test utterance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Scenario 2: In a separate model on a separate instance, a different user enters the same `help with hr` utterance. The system responds with the top prediction results for the intent, but the user isn't sure if it's the correct intent or not. So this user clicks the **Thumbs Down** icon, as shown in the image below.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

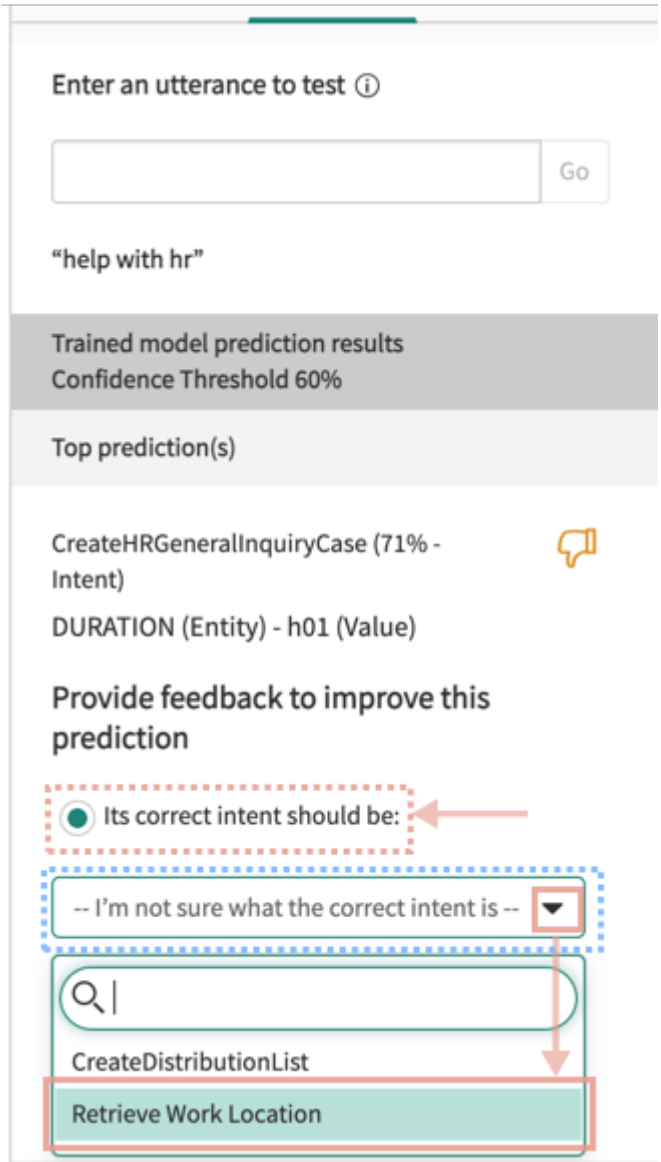


Result: The panel expands to show the **Provide feedback to improve this prediction** section where users can submit feedback that may help to improve the intent prediction.

There are two options here:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

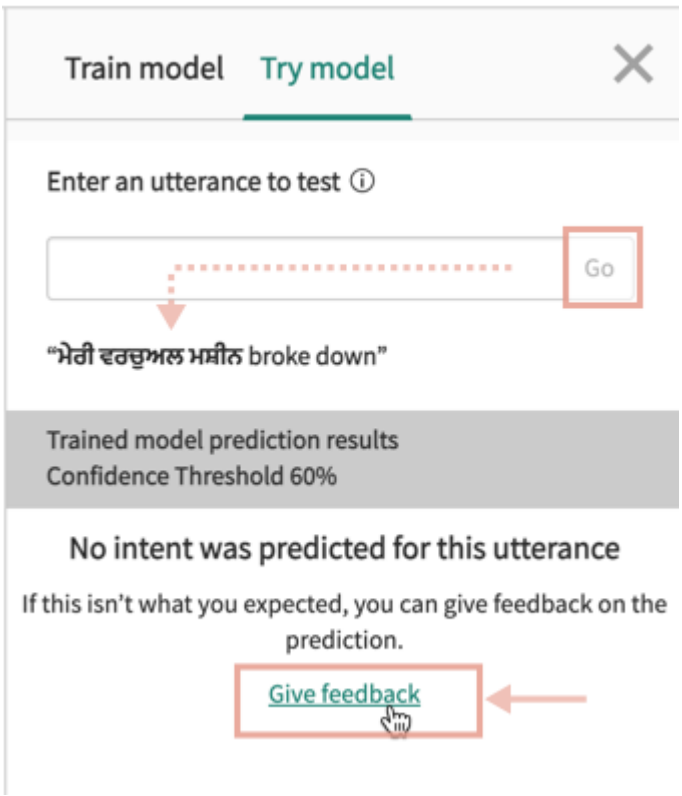
- If users click the **Its correct intent should be:** button, a list appears where they can choose a more appropriate intent for the test utterance. In this example scenario, a user selects the **Retrieve Work Location** intent, as shown in the image below.



- If you click the **I'm not sure what the correct intent is** prompt, instead of returning a top prediction, the system shows the next best intent predictions available.

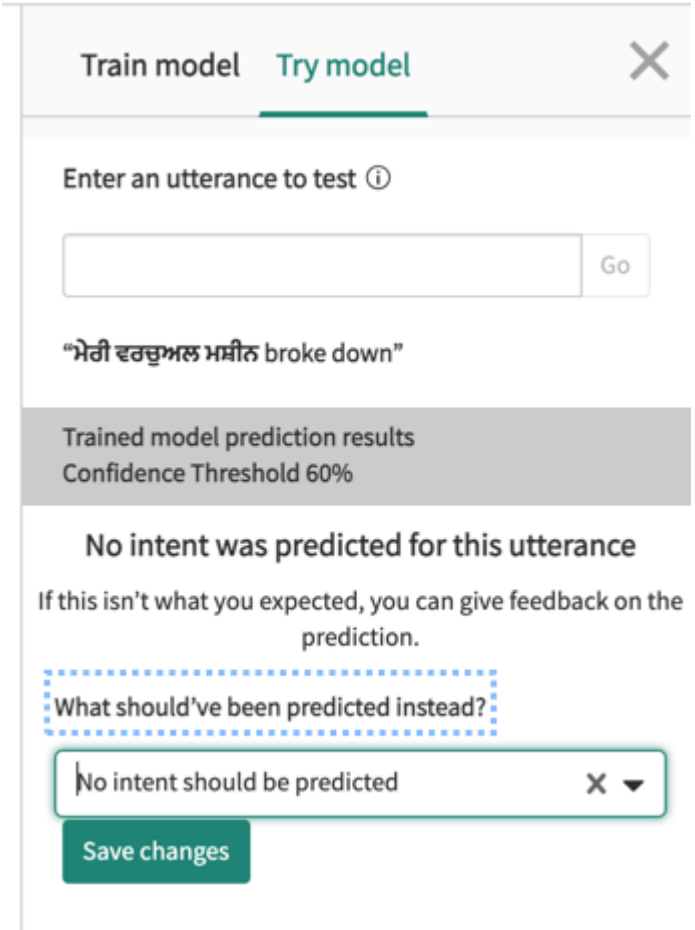
Scenario 3: In a separate model on a separate instance, another user submits an utterance that uses gibberish, or uses a language that's different from the language the model uses. For example, a user mistakenly submits an utterance comprised of both non-English and English languages, as shown in the image below.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



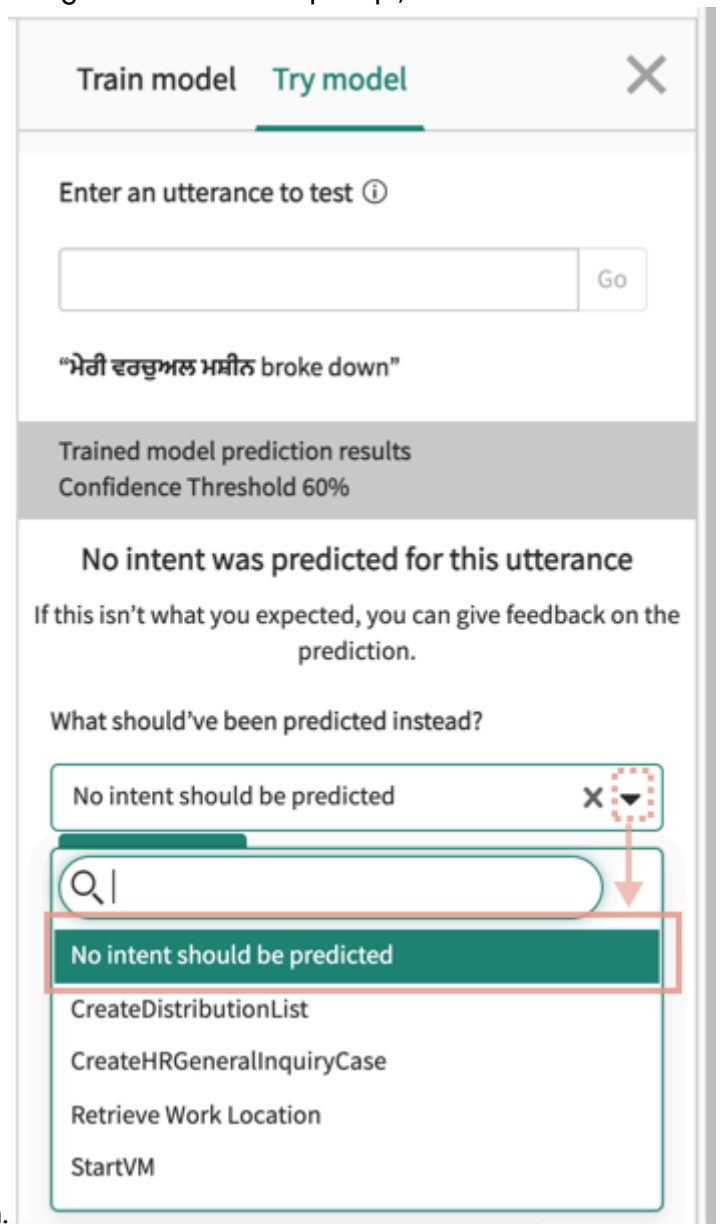
Result: The system doesn't return a prediction because the utterance uses two different languages together. Since no intent was predicted, the user clicks the **Give feedback** option which expands the Try model section to show other intent alternatives.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

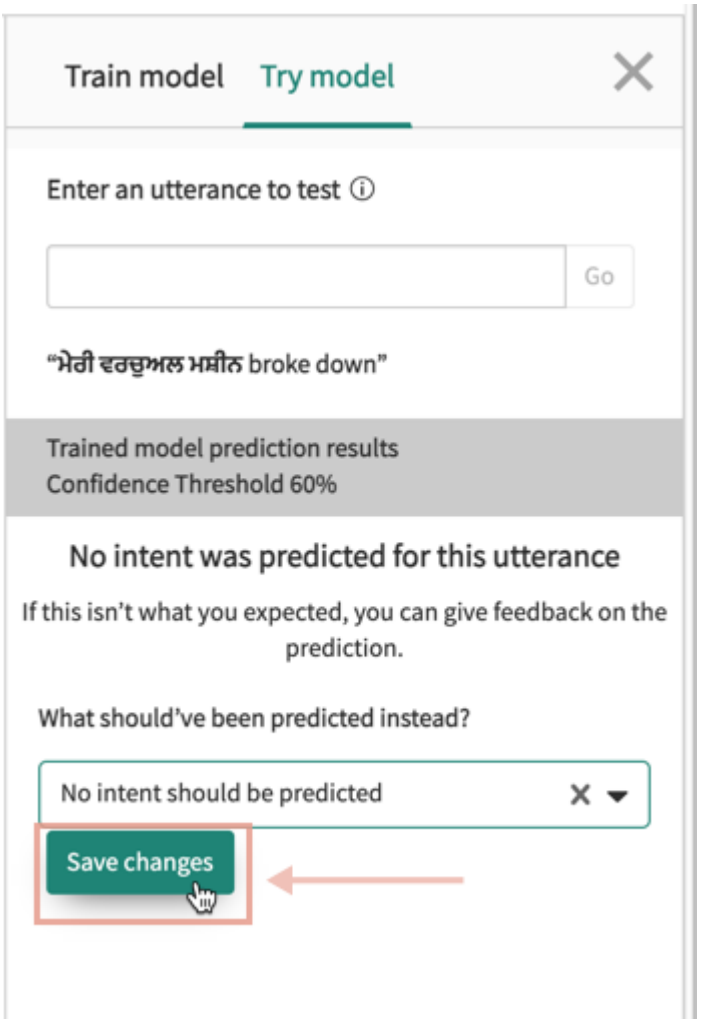
So instead of choosing an intent from the prompt, this user selects the **No intent should be predicted** option.



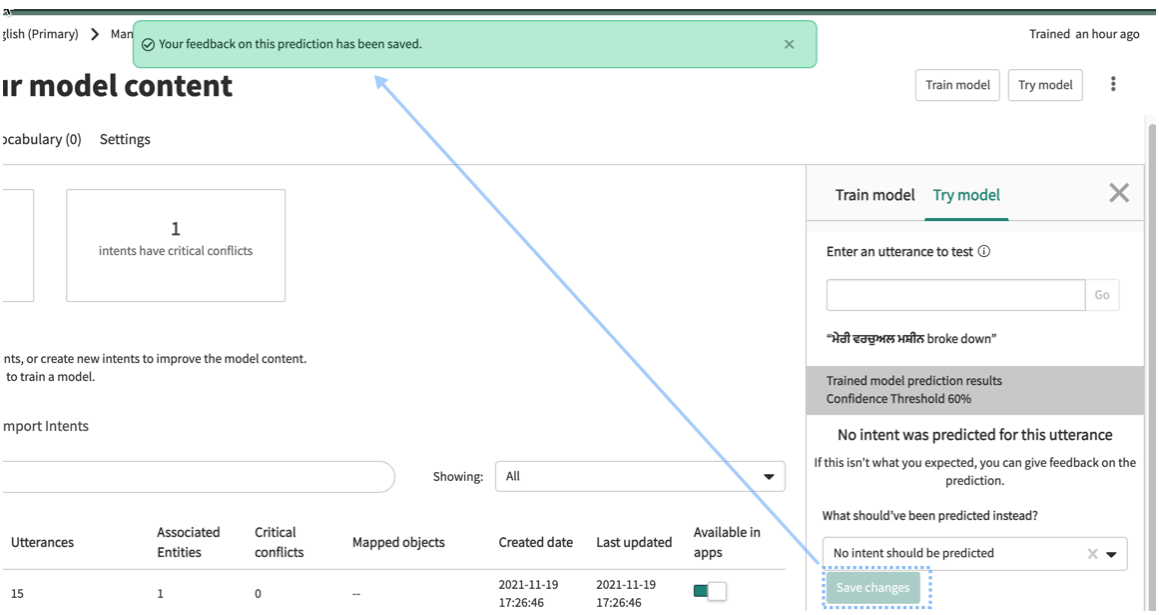
Note: When you choose and save **No intent should be predicted**, the utterance is removed from all intents which it is a part of.

Scenario 4: Along with choosing from a list of your model's intents for a prediction, you can also directly notify the system that the utterance is irrelevant to the model. To do this, you click the **Exclude this model's predictions for this utterance** button, then click **Save changes**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Result: A banner appears at the top of the screen confirming the user feedback for the prediction is saved, as shown in the image below.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Accessing your feedback records

Your feedback data is stored in the **ml_labeled_data** table, which is also used by other ServiceNow products. This table can also house multiple sources, such as Virtual Agent chat logs that can be used for future predictions.

Test and publish your model

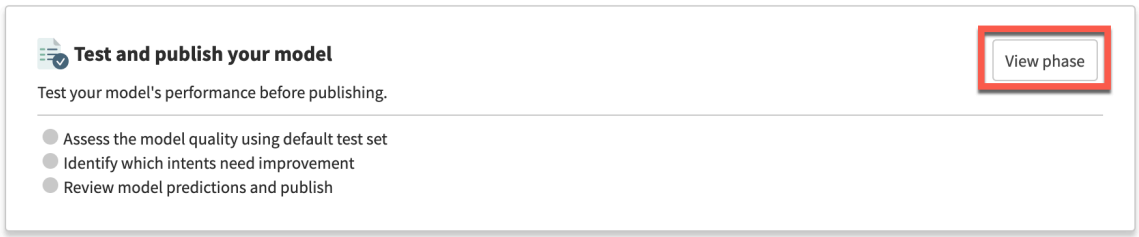
Assess the performance of your NLU model to identify areas for improvement. Then publish your model to make it available to other applications such as Virtual Agent.

Summary usage

Test your Virtual Agent or AI Search model against its default test set to see how the model responds. Test results provide information you can use to improve your model.

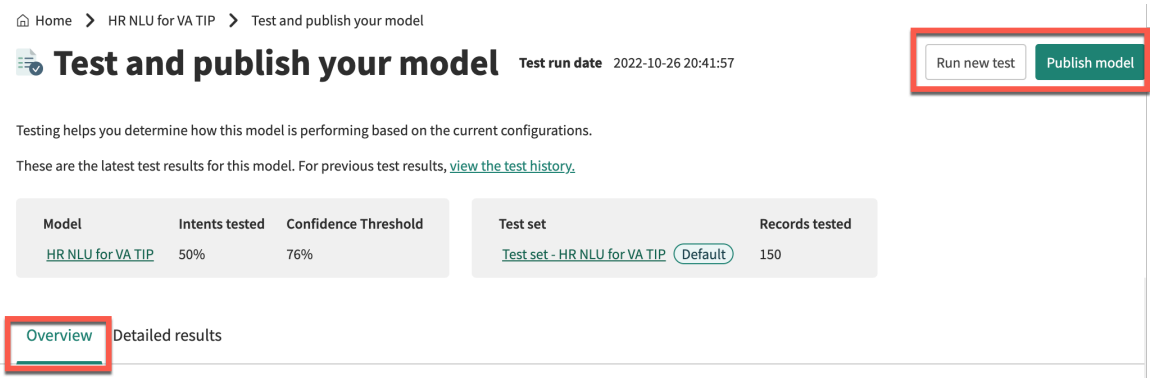
Note: Testing your model requires the Multi-model Batch Testing feature, available with the NLU Workbench - Advanced Features application from ServiceNow® Store. For more information, see Install NLU Workbench - Advanced Features.

To test your model, navigate to **NLU Workbench** > **Models**. Select the tab for your model's application, then select the name of the model. In the **Test and publish your model** card, select **View phase.**



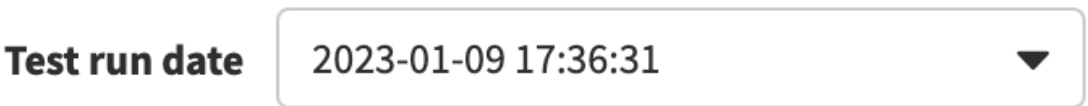
Overview of testing and publishing your model

The **Test and publish your model** phase opens in the **Overview** page by default. Buttons for **Run new test** and **Publish model** are located here.



Overview provides information about a previous test run, with bar charts summarizing the test results.

If you have earlier test runs, you can view those by selecting from the **Test run date** list.

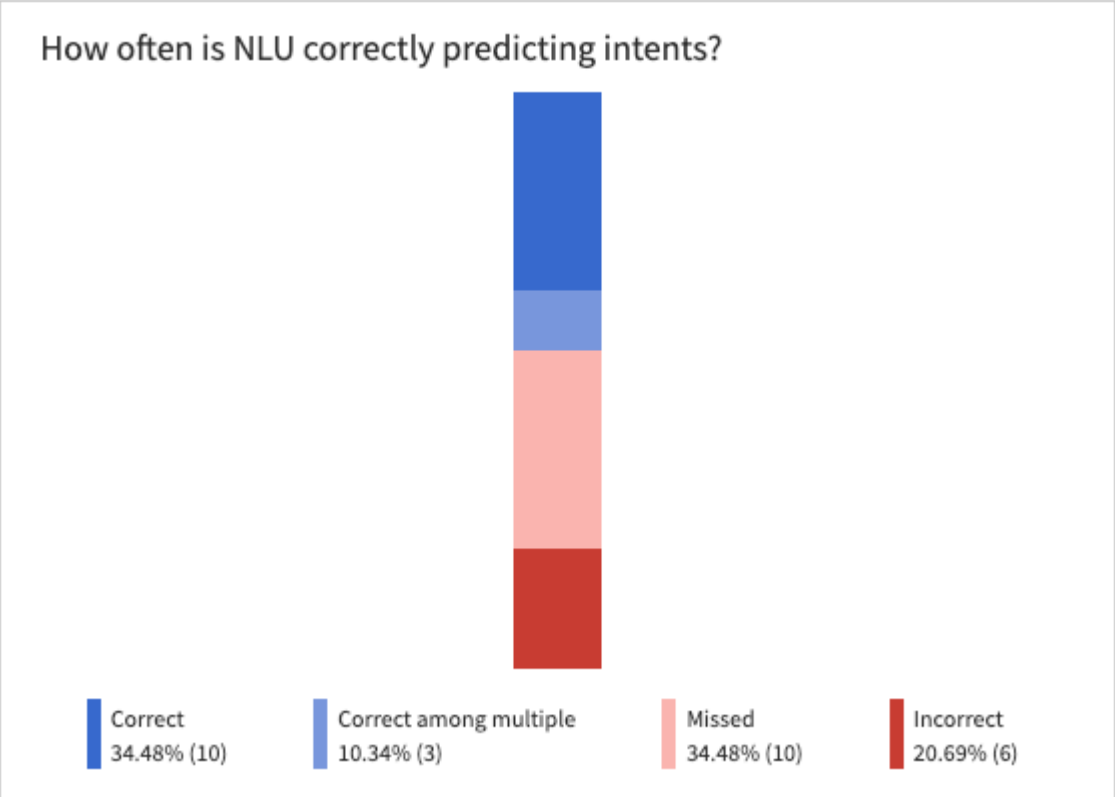


© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To drill down into the test results table, select the **Detailed results** tab. Each test utterance is listed in **Detailed results**, with its prediction.

Understanding test results

The test results show how your model responded to the utterances in the test set.



The bar chart shows the prediction percentages for correct, correct among multiple, missed, and incorrect:

Percentage	Description
Correct	The percentage of utterances for which your model correctly predicted the intent. When the model predicts no intent for utterances marked as `Not relevant`, that result is counted as Correct.
Correct among multiple	For utterances that had more than one intent predicted. The percentage of utterances for which the model correctly predicted the intent or intents, but also predicted intents that did not belong to the utterance.
Missed	The percentage of utterances for which your model did not predict an intent, even though there was an expected intent.
Incorrect	The percentage of utterances for which your model predicted an intent that was not correct.

Testing can affect the model's confidence threshold. The confidence threshold determines how confident a model must be to predict an intent for an utterance. For more information on confidence thresholds, see NLU model settings.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For information about utterances which should not have any intent predicted, see Irrelevance detection in NLU.

Publish model

The **Publish model** button makes the current version of the model available to other applications such as Virtual Agent.

> **Note:** If the model has not been trained, the **Publish model** button is unavailable. Return to the **Build and train your model** phase to train the model before publishing.

For more information on publishing your model, see Publish your NLU model.

Multi-model Batch Testing

In the **Test and publish your model** phase, you test your model against its default test set. With Multi-model Batch Testing, you can test against other test sets, test multiple models at once, and see your test results. To use Multi-model Batch Testing, navigate to **NLU Workbench** > **NLU Advanced Features** > **Multi-model Batch Testing**.

For more information, see Multi-model Batch Testing.

For more information about test sets, see:

- Create a test set
- Test set creation and management

For information about the process of testing, see Test your model.

Test your model

Test your Natural Language Understanding (NLU) model against its default test set. Testing helps determine how your model is performing with the current content.

Before you begin

- Make sure that the NLU Model Builder - Core plugin, NLU Model Builder plugin, NLU Workbench - Advanced Features plugin and Predictive Intelligence plugin are all installed and activated.
- Have a trained model for Virtual Agent or AI Search. For more information, see Build and train your model.
- Have a test set for testing models. For more information, see Test set creation and management.
- Role required: nlu_editor, nlu_admin, or admin. The editor must be assigned to the model.

About this task

Run a test of your model using its default test set, which contains test utterances and their expected intents.

- If an expected intent in your test set does not correspond to any intent in the model, the utterances with those intents are not used for testing. They are not included in the test results.
- The mid-conversation responses of Dialog Acts can't be tried or tested in NLU Workbench.
- When the model returns no prediction for utterances which are marked as `Not relevant`, that result will be counted as **Correct**.
- If your test set does not cover at least 60% of intents in the model, the system will not recommend a confidence threshold. However, you can still run the test.

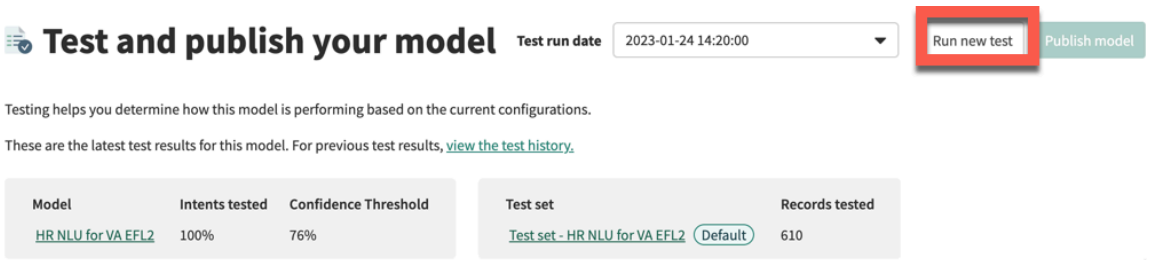
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In this example scenario, you've trained your model and want to assess the performance.

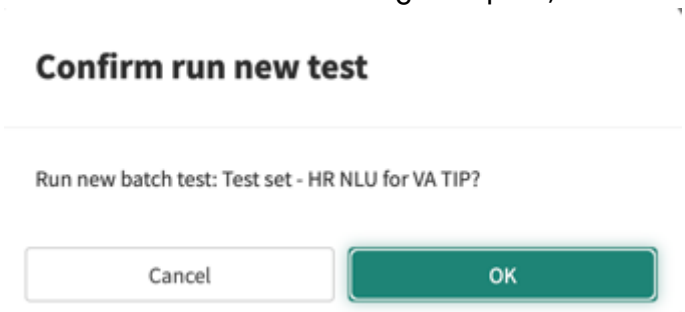
Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, then the name of the model you want to test.
3. On the model overview's **Model details** tab, locate the card for the **Test and publish your model** phase, then click **View phase**.
4. Select **Run new test**.

   Note: Other phases are unavailable during a test.



5. The **Confirm run new test** dialog box opens, so select **OK**.



The test begins and the **Testing is in progress...** page loads.

Result

When the test is finished, the **Test and publish your model** page reloads. The **Test run date** field reflects the date and time of this test.

The **Overview** tab displays a chart of the test results. It also displays a list of the top 5 incorrect intents and the top 5 missed intents.

The **Detailed results** tab lists all of the test utterances and their prediction outcomes.

You can see previous test results by clicking **view the test history** on **Test and publish your model**, or by navigating to **Batch Testing** > **Test results**.

What to do next

Use the results to edit and improve your model's content. When you are satisfied with the results, publish your model to make it available to consuming applications such as Virtual Agent.

Publish your NLU model

Publish your Natural Language Understanding (NLU) model to activate it and make it available for use in other applications that consume NLU.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Make sure that the - NLU Workbench - Core plugin, NLU Workbench plugin, and Predictive Intelligence plugin are all installed and activated.
- Have a trained and tested NLU model.
- Role required: admin or nlu_admin

About this task

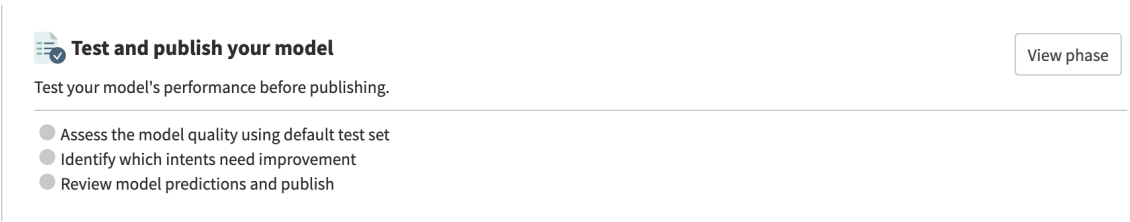
For this procedure, you've already trained, tested, and refined your NLU model and are satisfied with the results.

If your model is already published, you can publish it again. However, you must train the model again before republishing.

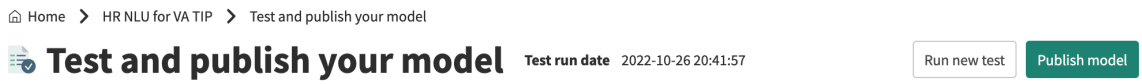
Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab corresponding to your model's application, then select the name of your model.
3. On the Model overview page, locate **Test and publish your model**, then select **View phase**.

   Note: If the model hasn't been built or trained yet, the **Test and publish your model** phase is not available.



4. On the **Test and publish your model** screen, select **Run new test** to assess the model using its default test set.
   Testing can be skipped, but model performance may not be optimal. If testing is skipped, a confirmation message displays, asking if you want to continue without testing. For more information on testing, see Test and publish your model.
5. Click **Publish model**.



   Note: If the model hasn't been tested yet, the **Publish model** button is white. If the model has already been tested, the button is green. If the last trained model is already published, the **Publish model** button is unavailable. Return to **Build and train your model** to train again.

Result: The most recent version of your NLU model is published. The model is active and available for use in other ServiceNow applications, such as Virtual Agent. Publishing also replaces any older versions of the model that are currently in use by those applications.

Compare draft and published versions of your NLU model

Compare a draft trained Natural Language Understanding (NLU) model to its most recent published version. Test and review the changes to make sure that your draft model will have increased performance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Workbench - Advanced Features plugin and Predictive Intelligence plugin are installed and activated.
- Role required: nlu_editor, nlu_admin, or admin. The editor must be assigned to the model.

About this task

In this example scenario, you're training and trying a published NLU model in the NLU Workbench iteratively with the goal of improving its prediction confidence scores.

When you try an utterance against an NLU model:

- If the model is trained and never published, the Test Model panel shows only trained model results.
- If the model is trained and published, the Test Model panel shows only published model results.
- If you've made changes to a published model and trained it, the Test Model panel shows both the trained model results and the published model results for comparison.

In this example procedure, you've cloned the model from a pre-built read-only HR model. You've cloned the model to create your own business-specific version of it while leveraging the existing intents from the pre-built model.

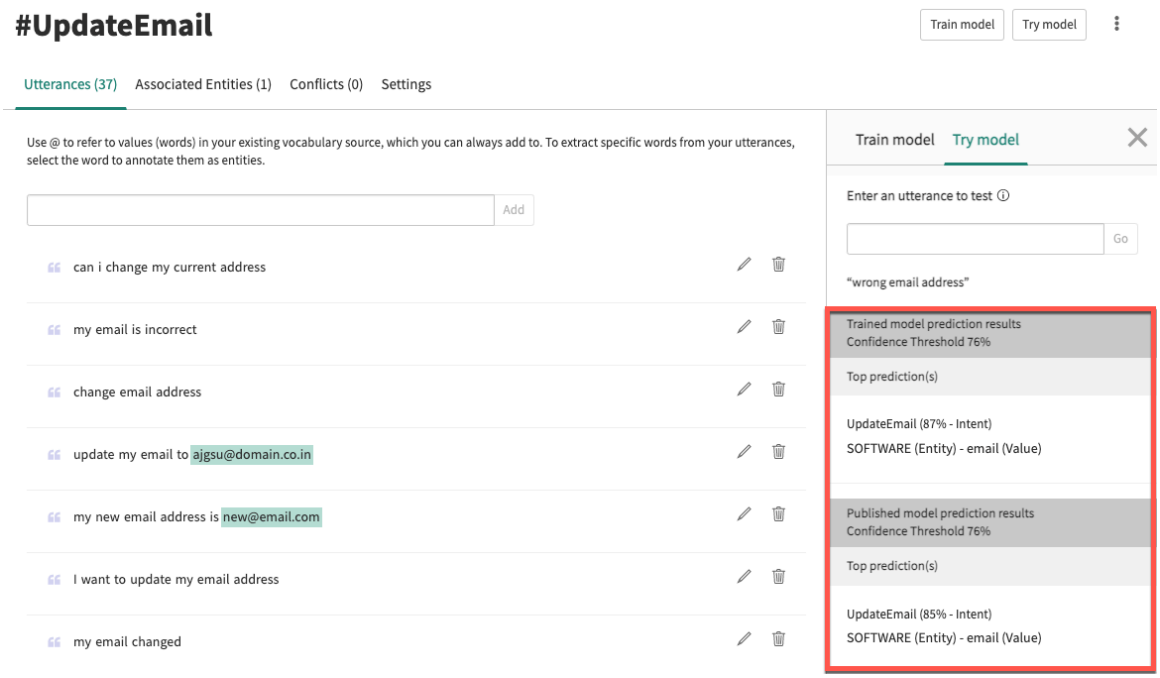
Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab corresponding to your model's application, then select the name of your published model.
3. From the model's overview page, locate the **Build and train your model** card and click its **View phase**.
4. Make a change to the intents, utterances, entities, or vocabulary.
   In this example scenario, you add a few more training utterances to the #UpdateEmail intent.
5. Train and try the changed model so you can see its prediction scores compared to the published version's scores.
   a. On the **Train model** tab, click the **Train** button.
   b. When training is finished, the system displays The model has been successfully trained.
   c. In the **Try model** tab, enter this utterance: `wrong email address`.
   d. Click **Go**.

   The panel displays predictions results for both the published model and the trained model. Compare the results of the two versions of the model, before and after your changes. In this

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

example, the confidence score increased by a small margin. By making significant changes to the model content, the confidence score or even the intent predictions may change.



What to do next

Use the information in the test panel to see if the changes you make will improve the model's performance. Once you are satisfied with your changes, test your model before publishing. Then, publish your model to replace the current published version.

Irrelevance detection in NLU

Keep Virtual Agent chats focused with Irrelevance detection. Use the Irrelevance detection feature to train your NLU model to avoid making predictions for utterances that are not relevant.

Summary information

The Irrelevance detection feature improves the prediction accuracy of NLU models by training them to ignore certain utterances. These utterances from your users may not apply to any intent, so should not get any prediction.

To ensure that models do not predict an intent when they are not supposed to, you can mark utterances as `Not relevant`. These marked utterances are included as part of model training. When the published model encounters similar utterances from your users, no intent is matched or predicted.

The Irrelevance detection table is where you can manage utterances marked as `Not relevant`.

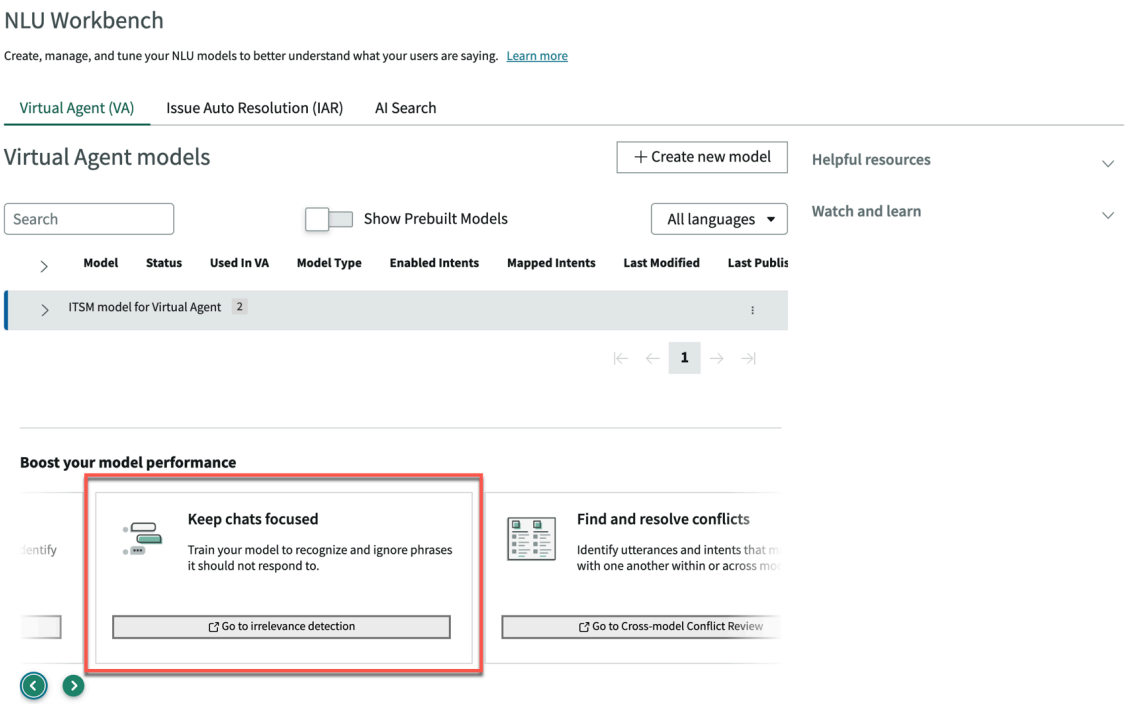
Roles, Usage, and Navigation

Use the nlu_admin or admin role to access Irrelevance Detection. The nlu_editor role is also able to access Irrelevance Detection, but must be assigned to a model to edit that model's contents.

Irrelevance Detection is available for Virtual Agent models only.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Navigate to **All** > **NLU Workbench** > **Models**. The Virtual Agent tab opens by default.
2. Scroll down the list of Virtual Agent models to the **Boost your model performance** section.
3. Scroll horizontally to locate the card **Keep chats focused**, and select its button **Go to irrelevance detection**.



Note: The instance URL for this feature is <instance-name>.servicenow.com/now/nlu-workbench/irrelevant-utterances.

Adding utterances to Irrelevance detection

There are several methods for adding utterances to **Irrelevance detection**:

- The Virtual Agent chat log: In the Expert Feedback Loop in NLU Workbench, when you review an utterance collected from the Virtual Agent chat log, you can mark it as `Not relevant`. The system will ask whether it should be irrelevant to a particular model, or to all models. After adding, these utterances may display as **NO_INTENT** in Expert Feedback Loop.

  For more information on Expert Feedback Loop in NLU, see NLU Expert Feedback Loop.

  These utterances have a Source of **VA Chat Logs** in the Irrelevance detection table.

- Manual input: In Irrelevance detection, type your utterance in the **Type utterances here** field, then select **Add**. The system will ask whether it should be irrelevant to a particular model, or to all models.

  These utterances have a Source of **Manual** in the Irrelevance detection table.

- Importing: When you use a CSV or XLSX (Excel Workbook) file to import training utterances and their intents, you can indicate irrelevant utterances by leaving the Intent column empty. These utterances may display as **NO_INTENT** in areas such as Expert Feedback Loop and Conflict Review.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information on importing utterances and intents, see Create an NLU model from a CSV file.

These utterances have a Source of **Manual** in the Irrelevance detection table.

Behavior of irrelevant utterances

There are two types of irrelevant utterances: those associated to one specific model, or those irrelevant to any model. A model can have a maximum of 200 irrelevant utterances associated to it.

Following are details of how these two types and the 200-count maximum interact.

When a model is submitted for training, at most 200 irrelevant utterances are submitted with it. First, irrelevant utterances directly associated to that model are submitted. Then, utterances which are designated as not relevant to any model are submitted. The total of these types does not exceed 200.

If a model has 200 irrelevant utterances associated to it, and a new irrelevant utterance is added, then the model's oldest utterance is dropped. The new irrelevant utterance could be of either type (not relevant to the specific model, or not relevant to any model).

A model cannot have more irrelevant utterances than normal training utterances.

Conflict review

If an utterance is marked as irrelevant, and there is a similar utterance in an intent, the model does not make a prediction for that utterance. In other words, irrelevant utterances take precedence over training utterances.

Because irrelevant utterances impact the model's predictions, they are displayed as conflicts when they overlap with training utterances.

There are two locations where conflicts with irrelevant utterances are highlighted:

- The Cross-model Conflict Review module. For more information about this module, see Cross-model Conflict Review.

  Note: The Cross-model Conflict Review module is available with the NLU Workbench - Advanced Features application from ServiceNow® Store. For more information, see NLU Workbench - Advanced Features.

- The Conflicts tab of an intent. For more information on conflicts in an intent, see Resolve intent issues.

For the purpose of reviewing conflicts, irrelevant utterances are displayed as though they are in their own intent, named **NO_INTENT**.

Note that irrelevant utterances cannot be edited or deleted on the conflict page. Copy the irrelevant utterance from the conflict page to the Irrelevance detection page to modify or delete the utterance.

More information

- In test sets, aim to include about 10% of test utterances as irrelevant. This helps to assess how your model handles utterances that should not have an intent predicted. When you import training or test utterances from a CSV or XLSX (Excel Workbook) file, you can indicate that an utterance is irrelevant by leaving the Intent column empty.

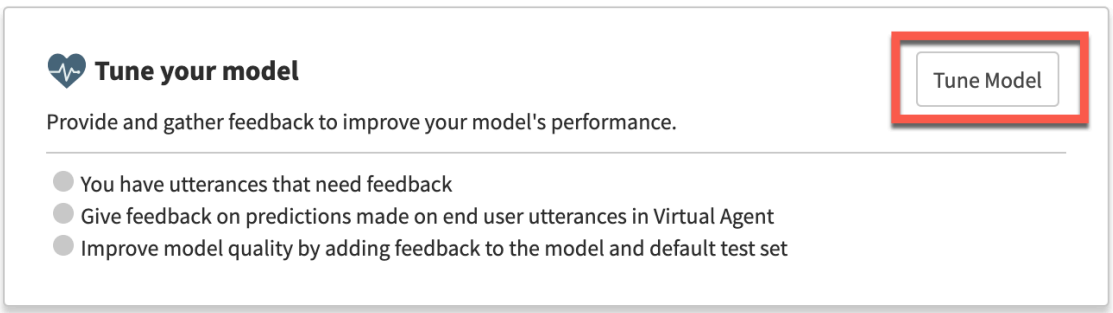
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- When testing models against test sets, results are considered Correct if no intent is predicted for an irrelevant utterance.
- Utterances that were marked as `Not relevant` can be re-assigned later. For example, if a new intent is created, existing irrelevant utterances can be manually assigned to the new intent. They then become part of the normal training utterances.

  To re-assign an utterance in Irrelevance detection, expand the list in the **Corrected intent** column and select the appropriate intent. Be sure to select the **Save feedback** button after re-assignment. Also be sure to retrain the model to incorporate these updates into the model.
- Irrelevant utterances are not associated to particular intents within a model. They are either associated to one model, or marked as not relevant to any model.
- The utterances which are marked as not relevant to any model are submitted as part of the training data for each model. In other words, these apply to all models.
- Model training is necessary to incorporate `Not relevant`utterances. Training any model adds newly marked utterances to all models.
- Utterances can be deleted or edited in the Irrelevance detection table.
- Irrelevant utterances should have content that is different from utterances associated to an intent.

Tune your model

On your model's overview in NLU Workbench, open the **Tune your model** phase to review and incorporate user utterances from the Expert Feedback Loop.



Summary usage

The **Tune your model** phase on your model's overview takes you to the NLU Expert Feedback Loop feature, which collects data from Virtual Agent chat logs and provides it to you for feedback.

Note: The ServiceNow® Store application NLU Workbench - Advanced Features must be installed to view **Tune your model**. The Virtual Agent plugin is also required.

When the Expert Feedback Loop has feedback waiting for your review, **You have utterances that need feedback** is highlighted on the card.

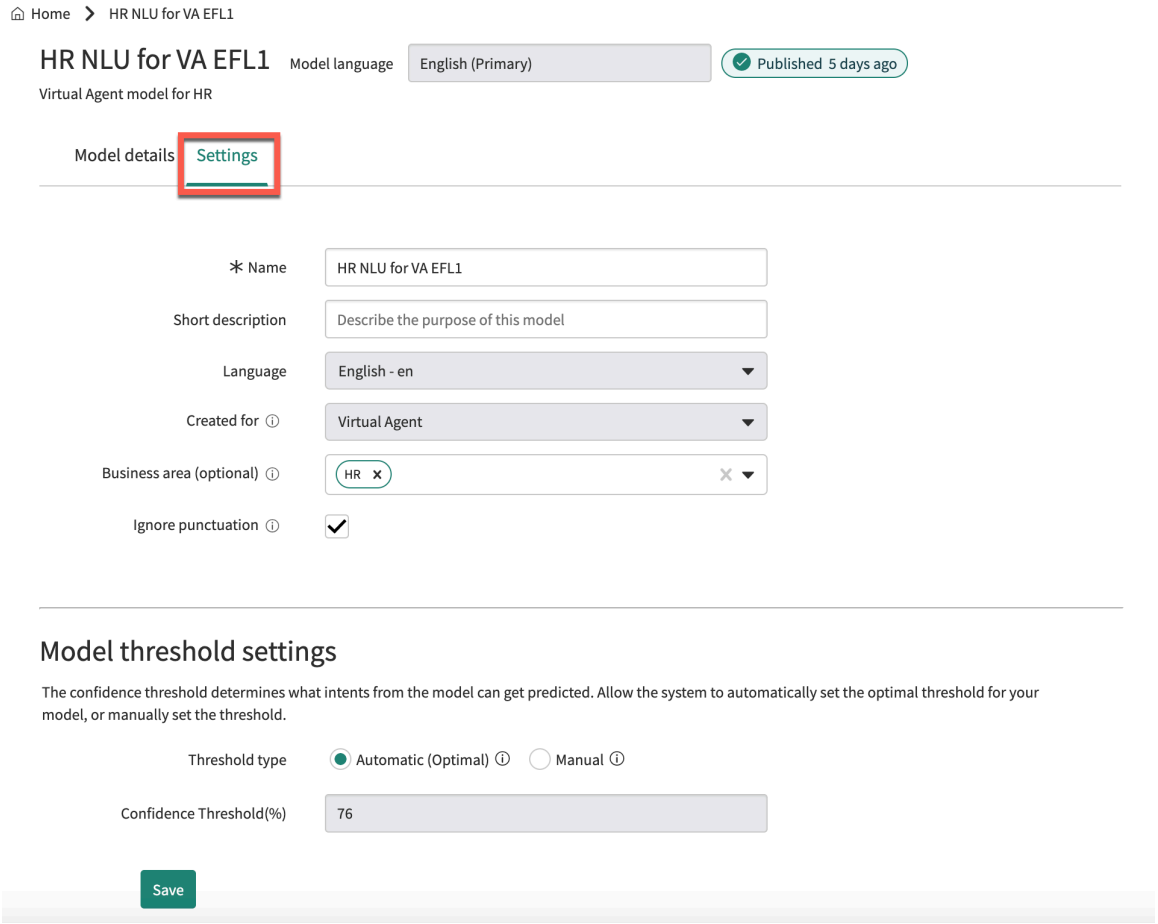
For more information, see NLU Expert Feedback Loop

NLU model settings

Change your NLU model's name, description, or confidence threshold on the **Settings** page of the model overview.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Access the model's settings by navigating to **All** > **NLU Workbench** > **Models**. Select the tab for your model's application, then your model's name. On the model's overview, select the **Model settings** tab.



Model settings

In the upper section of the model settings page, you can change the model's name, short description, and business area. You cannot change the model's language, purpose, or scope. To make a model with a different language, purpose, or scope, see Creating models.

By default, the **Ignore punctuation** check box is active. Ignoring punctuation makes it so that there is less variance between predicted intents and confidence scores for utterances with slightly different punctuation. For best results, keep the check box active.

Model threshold settings

Here you can adjust how the confidence threshold works in your model.

A threshold is a confidence score represented by a percentage. The confidence threshold of a model determines what intents from that model will be predicted for a given utterance. For example, if the model threshold is 65%, then an intent will be predicted for an utterance only when the intent has a confidence score that is at least 65%. Setting a threshold that is too low may increase the false positives by predicting intents that should not be a match for an utterance. On the other hand, a model threshold that is too high may filter out intents that you do want to get predicted. Finding the ideal threshold improves your model's ability to predict intents correctly.

There are two types of model threshold settings:

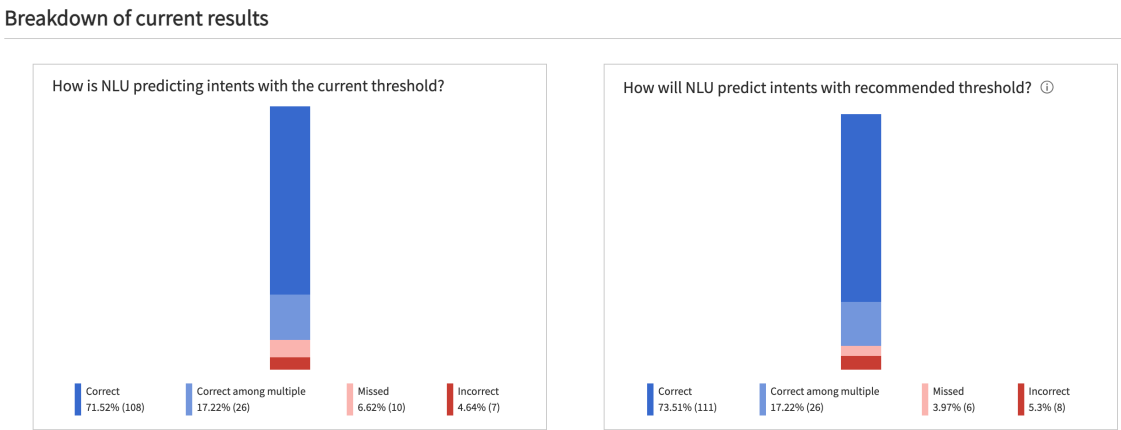
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- **Automatic** - Allow the system to choose the optimal confidence threshold for your model. The value is updated dynamically based on test results. This happens in the **Test and publish your model** phase, where your model's default test set is used.
- **Manual** - You can manually set the confidence threshold. The system may also recommend a better threshold for the model during testing. You can choose to accept recommendations.

Prebuilt models come with a tuned threshold. The confidence threshold on prebuilt models was chosen specifically for that model.

Test results include a model threshold recommendation only if they meet the following requirements:

- The test set has a Test Coverage score of at least 60%, with at least 5 test utterances per intent. For more information, see Test set creation and management.
- The test set has at least 100 utterances.
- The model is not a prebuilt model.
- The recommended threshold would have better results than the current threshold.



Test results with a recommended threshold contain a second graphic. The second graphic shows the prediction percentages with the recommended thresholds applied.

Applying the threshold recommendation may improve the prediction percentages of your model. Select **Apply recommendations** to change the threshold. The system automatically retrains the model, and the test results show the prediction percentages with the new threshold.

Multilingual model management

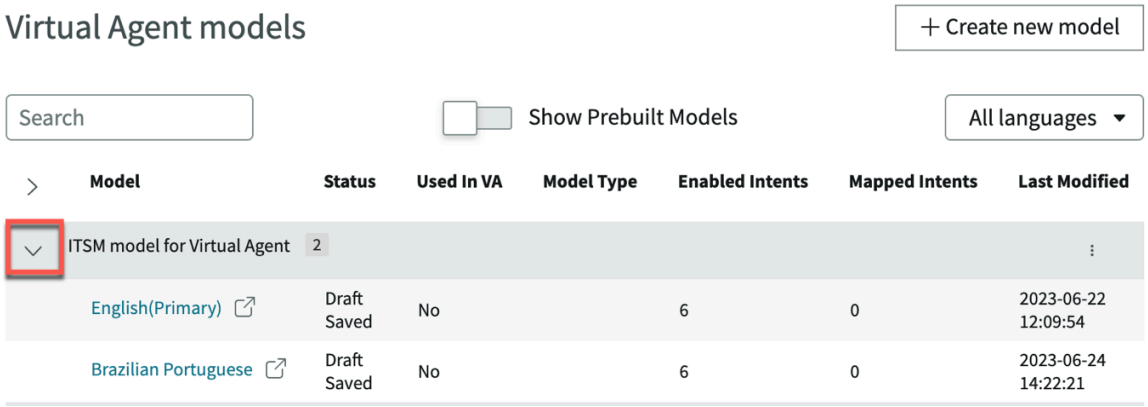
Use multilingual Natural Language Understanding (NLU) models for the system to understand user input in several languages. The NLU Workbench helps you manage and maintain a consistent structure for content across languages to provide a unified experience.

Primary and secondary languages

A primary language is the source language you choose when creating a model. These models are considered primary models. Primary models can then be translated into different languages. Those translated models are referred to as secondary models. The languages in which they are translated are referred to as secondary languages.

The NLU Workbench home displays primary and secondary language models nested under the model name. Select the arrow to the left of the model name to expand the language group.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Multilingual model management provides a way for you to group, oversee, and update your NLU models. The designation for the model language works as follows:

- Primary models have a language you assign to them during model creation and listed as English (Primary). The language of the primary model is the source language for the translations that follow later in the secondary models.
- Secondary models are translated copies of the primary model. Each secondary model uses a different language, such as Brazilian Portuguese, Polish, or Finnish.
- Any supported language can be the primary language for a primary model or the secondary language for a secondary model.
- Within a model group, you can't have two secondary models that use the same language.

For more information on the languages available in NLU, see NLU language support.

Implementing language grouping

Grouping languages is optional but helps with organizing various language versions of models. There are two methods for ensuring that model languages are grouped:

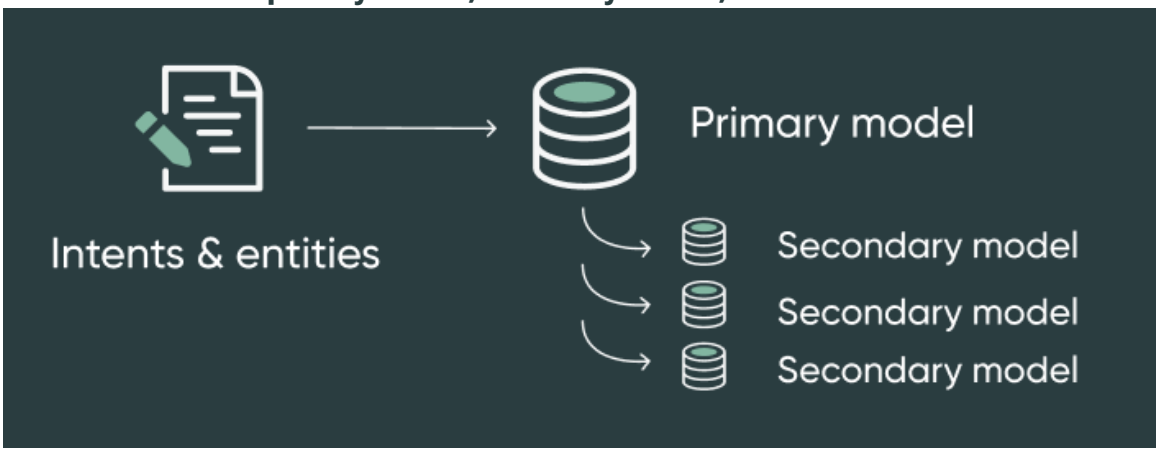
- You can add a language to a primary model in the model's More options menu, which automatically creates a secondary language model with translation options. For more information on adding and translating a language, see Translate a multilingual model.
- You can access the instance model migration page directly by visiting `<instance_name>.service-now.com/$nlu-studio.do#/model-migration`. For more information on language grouping using this method, see Model language grouping.

Primary and secondary model interactions

To ensure consistency within a model group, the names of intents and entities in all secondary models are the same as the content in the primary model. New intents can only be created in a primary model. Creating an intent in a primary model adds the intent to all the secondary models within the model group, but without any utterances in them. New intents are disabled by default. New entities can also be created in secondary models for languages that support entities. Adding an entity in a primary model creates it across all secondary models within the model group, if valid.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Interaction between primary models, secondary models, and content



When you add an intent to a primary model, the intent is added to all of its secondary models. Every intent in a secondary model is mapped to its corresponding intent in the primary model. The mapping of intents ensures that any application that uses these intents can access all the secondary intents through their corresponding primary intents.

When you delete an intent or entity in a primary model, its corresponding intents and entities are also deleted in its secondary models. Therefore, the secondary models must always follow the status of the primary model content. Although you can't delete intents in secondary models, you can disable them.

For more information on intent interactions, see Import primary model content to a secondary model.

You can't delete an entity created in a secondary model if it's a copy of a corresponding entity created in the primary model. However, you can add or delete an entity in a secondary model if it doesn't have a corresponding entity in the primary model.

Whenever you translate a model or add intents and entities to secondary models, the model must be reviewed. Secondary models marked with **Needs review** must be manually reviewed before publishing.

Viewing secondary models that need reviewing

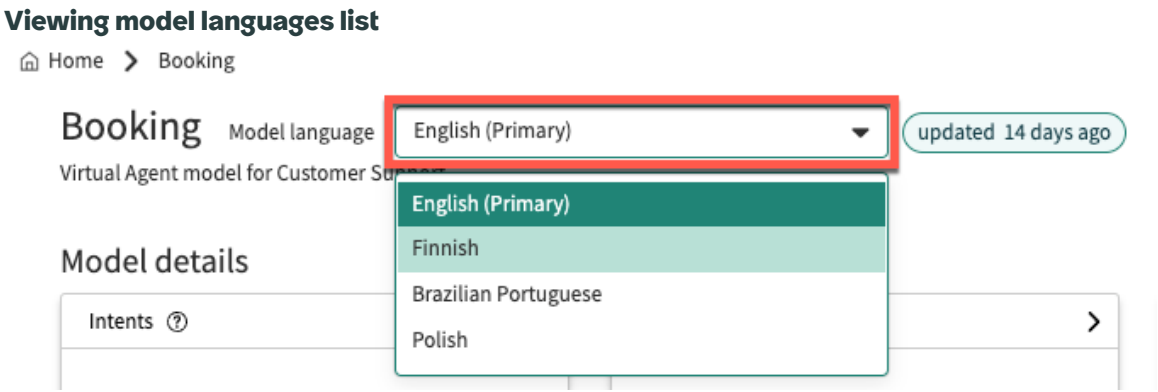


If every intent in a secondary model is disabled, then the **Train** and **Try** buttons are also disabled in the model. However, even if only one intent is enabled in the model, then you can train and test the model. Ensure that some intents are enabled in your secondary models to translate them.

Note: When training and testing, prediction scores for similar utterances across primary and secondary models can be different. The context comes across differently among languages due to inherent structural variations.

When looking at a model, you can navigate between languages in the model group. On a Model details page, use the **Model language** drop-down list to navigate to other models.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Duplicating grouped models and model groups

Using the nlu_admin role, you can duplicate primary models, secondary models, and entire model groups. You can duplicate just the primary model or a set of secondary models from among the model group.

Following is a list of system behaviors you may encounter when you duplicate these models.

- If you duplicate a secondary model that's in a model group, the duplicated version becomes a separate primary model that is outside that model group.
- All intents are duplicated. The duplicated intent maintains the same **Enabled** status as the original intent.
- When duplicating a primary model, you can duplicate a set of secondary models, or all of the secondary models along with it. This action creates a model group comprised of duplicated versions with the respective original models marked as the source models.
- When duplicating a model group, you can choose an existing secondary model to be the primary model for the duplicated group.
- If you select a secondary model in an existing model group as a new primary model while duplicating the group, all the disabled entities are enabled for the duplicated version of the secondary model. The secondary model becomes the primary model in the new duplicated model group.
- If you duplicate a primary model without any secondary model, the duplicated version becomes a separate primary model.

For more information, see Duplicate an NLU model.

Model language grouping

Language grouping makes it easier to manage your multilingual Natural Language Understanding (NLU) models. You can review existing language groups and designate new language groups.

Language grouping

Language grouping organizes your NLU models into primary and secondary models, and indicates the primary language and primary model for the group.

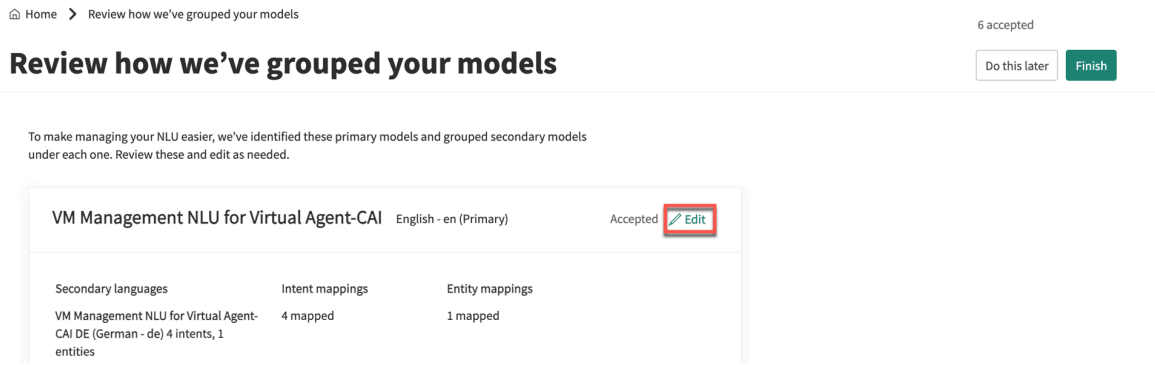
If you have added a new language to a model, a language group is created automatically, so you may already have existing language groups. For more information on adding and translating a new language model, see Translate a multilingual model.

To view current language groups and inspect possible new language groups, navigate to:

`<instance_name>.service-now.com/$nlu-studio.do#/model-migration`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

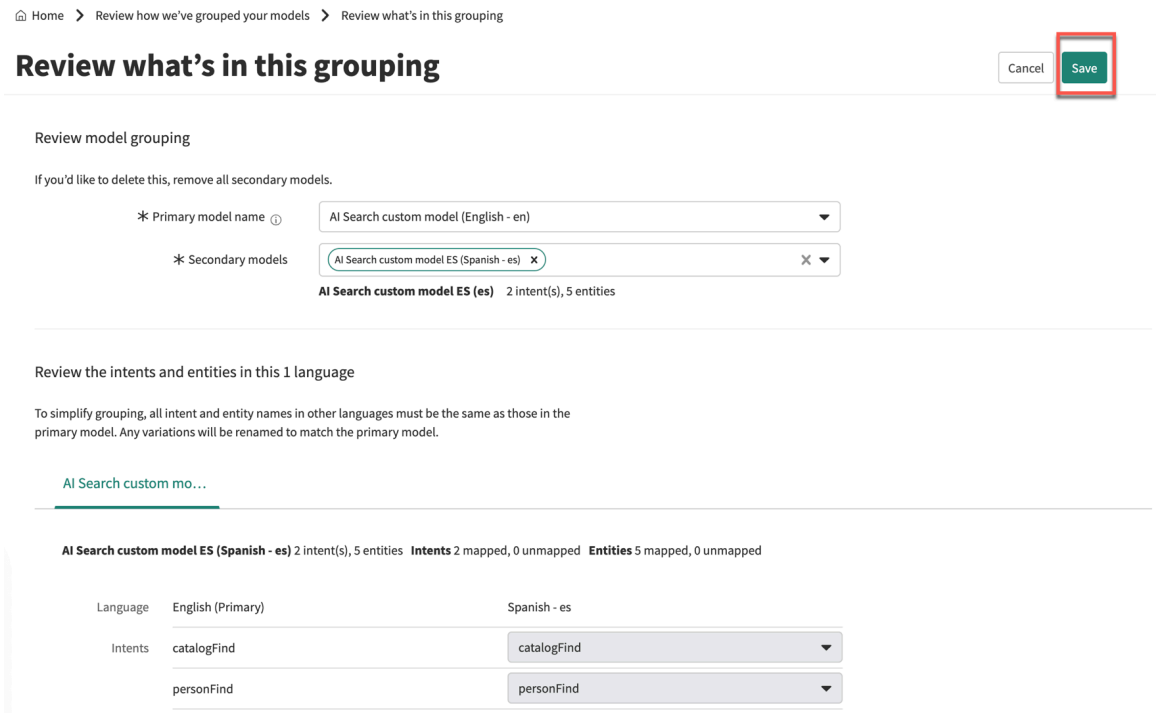
The page titled **Review how we've grouped your models** opens.



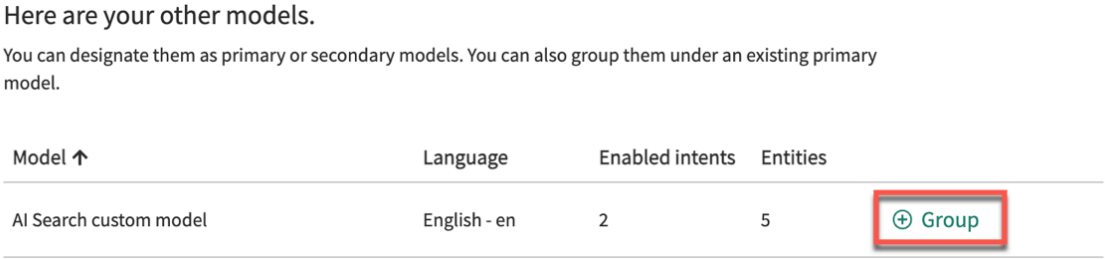
Reviewing your grouped models

There are two things you can do on the **Review how we've grouped your model** screen when using the nlu_admin role.

- You can review and edit current model language groups in the upper section. For existing language groups, any intent and entity mappings from Virtual Agent (VA) are displayed. To update VA mappings or to make one of the secondary languages the primary language in the group, select the **Edit** button. The page titled **Review what's in this grouping** opens. When you complete your edits, select the **Save** button.

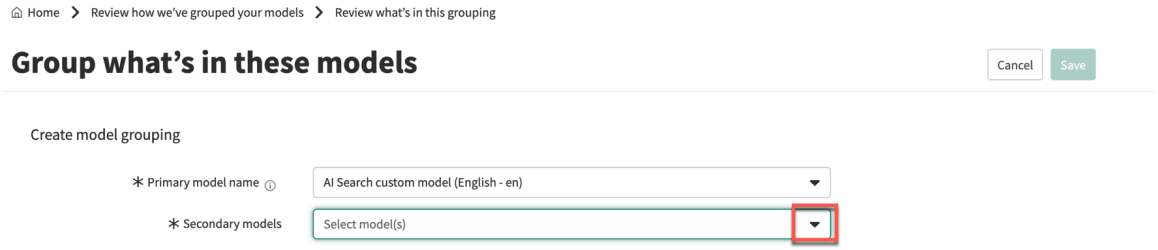


- You can manually group models in the **Here are your other models** section of the screen. Scroll down the page to view this section, if necessary. To set up a language grouping for a model, expand its **Group** button.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Group** function opens the page titled **Group what's in these models**. You can expand the list of possible Secondary models, and then select the **Save** button. Any new groups are displayed in the upper section of **Review how we've grouped your models**, after refreshing the page.



Note: The time it takes the system to group your models depends on the number and size of your models.

Translate a multilingual model

Add a language to an existing NLU model by translating it. Use one of several translation options to add a secondary model in a supported language.

Before you begin

- Activate the Localization Framework (com.glide.localization_framework.installer) plugin. See Localization Framework.
- In your target application scope, create a primary language model, or use an existing primary model.
- Multilingualism is available for Virtual Agent and AI Search models.
- Role required: nlu_admin or admin.

Note: The nlu_editor role cannot initiate a model translation, because the nlu_editor does not have permission to create a new model. See Assign an NLU editor to a model.

About this task

Translating a model creates a secondary model of the language you choose, in the same application scope as the primary model.

The secondary model contains the same content as the primary. However, the intents of the secondary model are disabled and must be reviewed prior to publishing the secondary model.

To translate a primary model and create a secondary model, you must choose a translation mode. The modes appear in the user interface as individual cards that you use in Step 4 of this procedure. Activated cards appear surrounded by a thin green border. Only activated cards can be selected as your translation mode. For the Use AI and Use a third party modes, selecting the **Add** button on an activated card starts the translation.

The No translation mode is active by default and available for all languages. However, the Use AI and Use a third party modes are activated only when the localization framework (LF) settings for the secondary language are in place.

Note: Before you add a new model in a secondary language, you must ensure that the Localization Framework (LF) settings for the secondary language are completed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Translation mode	Description
Use AI	Translate using machine-intelligence providers such as Google. Selecting the card creates a translated secondary model. Machine translations may be incorrect or too literal, so review the translated content for accuracy. For more information, see Dynamic Translation. To activate the Use AI mode, Localization Framework must be configured and activated.
Use a third party	Request translation by a professional translation team. When you select this mode, an LF task is created per the LF language group settings for the secondary language. When the translations are completed, the translated model is created and ready for review. For more information, see Localization Framework and Translation modes.
No translation	Default. Manually translate your primary language into a secondary language. Selecting this mode creates a draft secondary model with the intent names the same as that of the primary model. These intents are disabled by default, and you must enter your utterances manually. Names of intents and names of entities imported from the primary model aren't editable. Names of entities that you create directly in the secondary model are editable. This mode is also helpful if you want to draft and save the primary model first, and then translate your utterances later.

Multilingual models are supported in 17 languages.

Note: Some of these languages support only intents, not entities. If the language is intent-only, your primary model entities will not be translated or copied. For more information, see NLU language support.

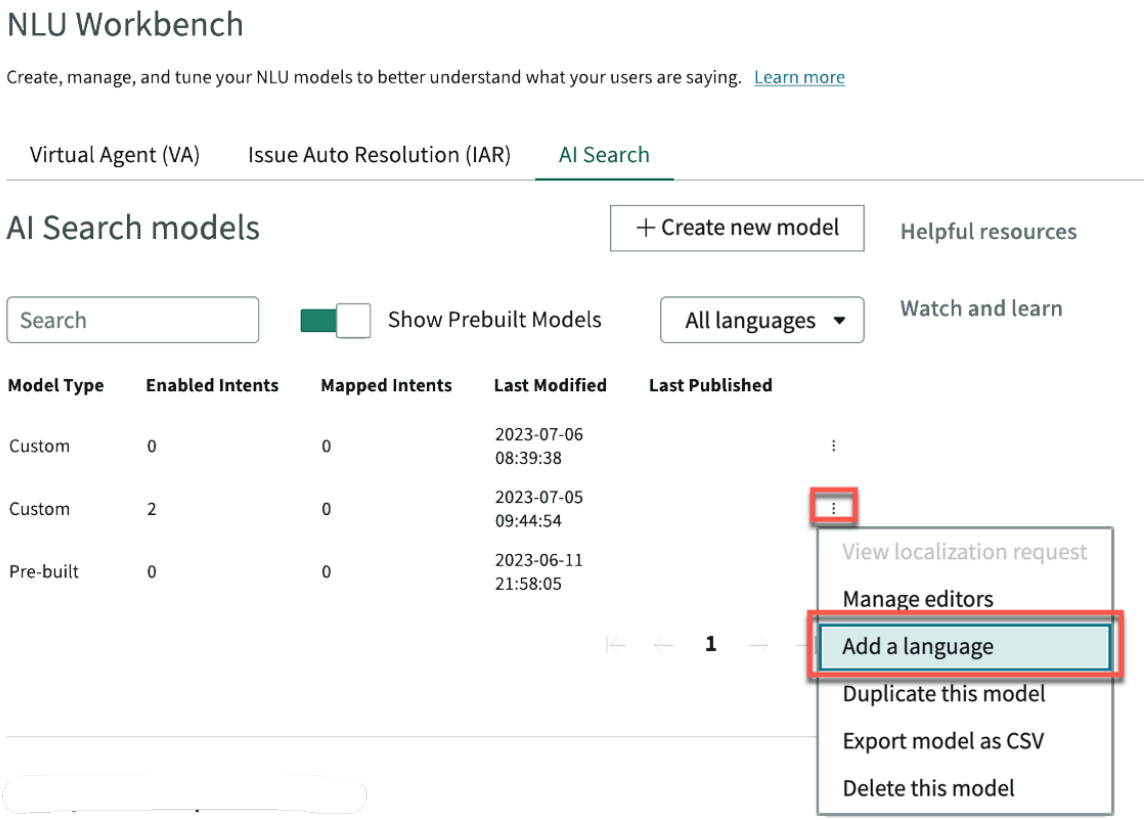
In this example scenario, your primary model uses the English language and you're adding a version in the French language, without translating utterances.

Procedure

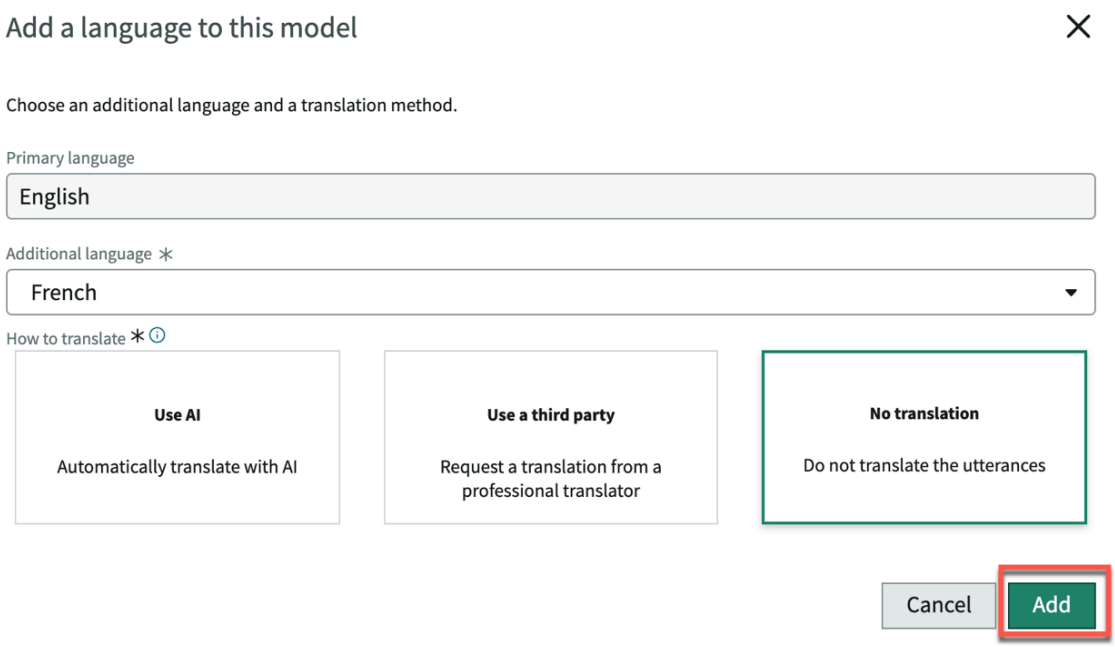
1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default. If appropriate, select the AI Search tab.
2. Identify a primary model to which you want to assign a secondary language.
3. In the row of the primary model, select the more options menu, then select **Add a language**.

Note: You may need to scroll to the right to find the more options menu.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. In the **Add a language to this model** window, choose the desired language for the secondary model and the translation method.



5. Select the **Add** button.
   The secondary model is created and the **Build and train your model** phase of the secondary model loads.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

If you chose to enable the model without translation, you must manually enter training utterances for each intent. If you choose to use software or a third-party translator, you must review the translations before you can continue working on the model.

Entering training utterances and reviewing secondary models can be delegated to users who have the nlu_editor role. For more information, see Assign an NLU editor to a model and Resolve intent issues.

Enable or disable a secondary model intent

Enable and disable intents in your Natural Language Understanding (NLU) models to make them active or inactive. Disable intents while editors or admins edit, review, or update its content and translations.

Before you begin

Role required: nlu_admin or nlu_editor

About this task

NLU admins and NLU editors can enable or disable intents in secondary models. If an intent is disabled, the model doesn't use it to make predictions. However, it also can't be accessed by any ServiceNow application using the model, such as Virtual Agent or Search.

An intent's state is represented in the **Enabled** column.



Enabling of intents works in the following ways:

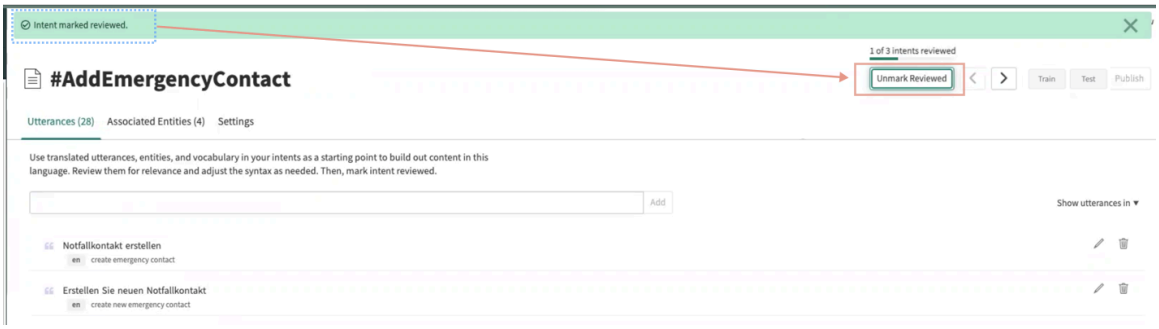
- Intents created in primary models are active by default.
- Intents from primary models are automatically created, but not enabled, in their secondary models (pending import).
- Disabling an intent in the primary model disables the corresponding intent in all secondary models.
- If an intent is disabled in the primary model, the intent can't be re-enabled in a secondary model.
- Enabling a disabled intent in the primary model does not enable the corresponding intents in the secondary models.
- Disabling intents in secondary models has no effect on the primary model or other secondary models.
- Secondary model intents that are marked with Needs review are disabled by default.
- If you disable a primary model intent that is mapped to a Virtual Agent (VA) topic, a confirmation message appears.
- If you disable a secondary model intent that is mapped to a VA topic, a confirmation message appears.
- NLU admins can enable and disable intents in any model.
- NLU editors can enable and disable intents only for models they are assigned to.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Disabling intents gives editors time to review the intent translations and update them if needed. When you're satisfied with the content, enable the intent to make it accessible to other NLU models and ServiceNow applications.

In this review example, you have a list of disabled intents in the Needs review state. The goal of this task is to review the translated content for a secondary model. When you complete your review, or if the content is fine as it is, you click Mark as reviewed. This moves the intent into the Reviewed state. You can also undo the Mark as reviewed state for an intent by clicking the **Unmark Reviewed** button, but only if you remain on the Intent screen. If you leave the screen prematurely, the **Unmark Reviewed** button disappears, and you won't be able to retrieve it.

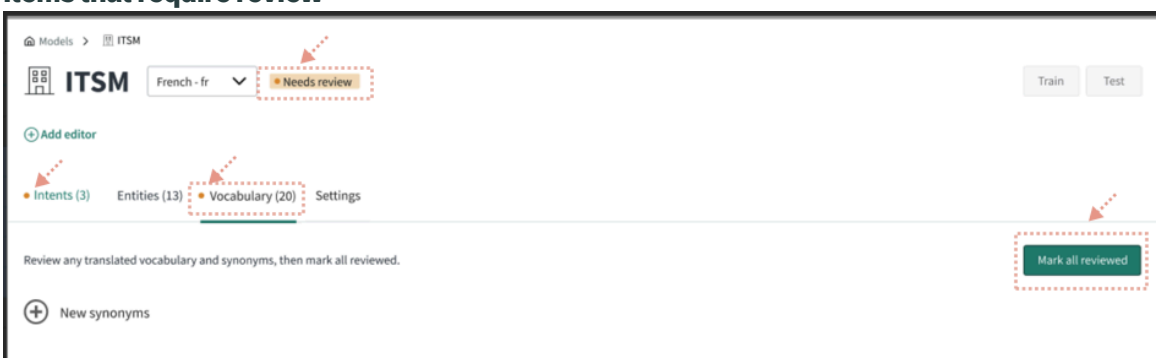
Switch between the Reviewed and Unmark Reviewed states for an intent



The editorial review also includes the Vocabulary section on the Model screen, where vocabulary that has been translated from the primary model to the secondary model is marked with an orange dot and the Needs review state. You can edit or delete the vocabulary items, even if they are in review. After you review the vocabulary and update the translations if needed, you can mark all of them at once by clicking the Mark all reviewed button. This action makes the orange dot disappear from the Vocabulary section of the Model screen.

Note: The **Train** and **Try** buttons on the Model screen are disabled until the vocabulary is reviewed, and the intents in the Needs review state are either reviewed or disabled.

Items that require review



In this example scenario, you have applied the Use Software translation mode to translate a secondary model into the Japanese language. Because the translation hasn't been reviewed yet, the system marks the intents and the model with the Needs review state. When a model is in the Needs review state, it can't be trained, tested, or published. When it moves out of the Needs review state, you can then train and test the model.

Note: An NLU editor can't create or publish a model. Only the NLU admin has permission to do this.

In this example procedure, you're reviewing the intents in a model one at a time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

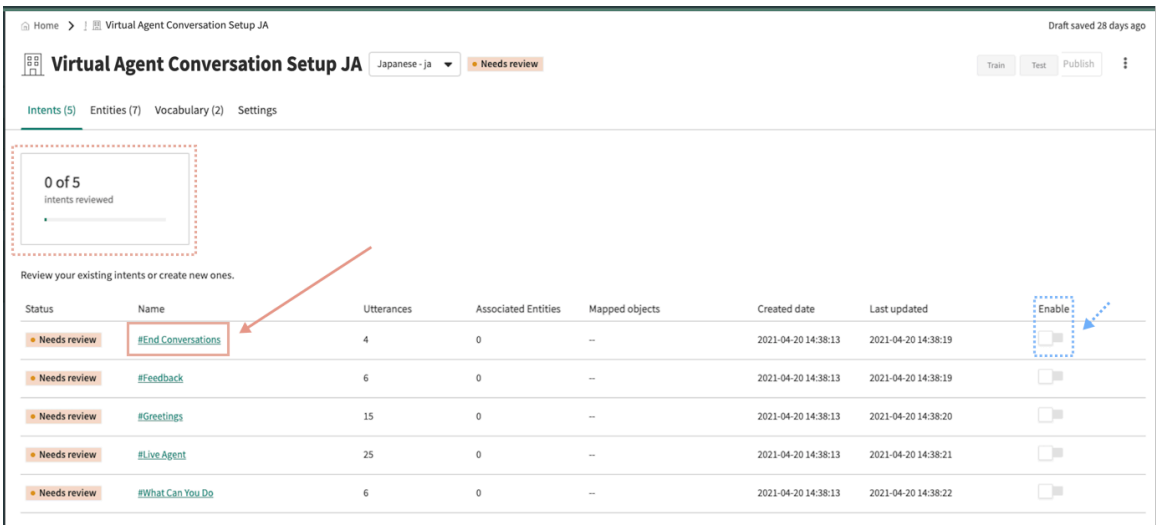
Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
2. Locate one of your secondary models that has already been translated.

   In this example scenario, you locate the Virtual Agent Conversation Setup JA model, as shown in the image below.

   In the Intents section of the model, there's a list of five intents that are marked with the Needs review state. Above the list of intents, there's also a box that shows a running count of the intents that need review.

The needs review state



3. Click an intent Name so you can access the intent content.
   In this scenario, you click the **#End Conversations** intent name.
4. Review the secondary language translations of the model utterances, entities, and vocabulary, then update the content as needed.
   When you finish your review of the intent, the Mark as reviewed button appears.
5. Click **Mark as reviewed**.
   Result: The system updates the **#End Conversations** intent from the Mark as reviewed state to the Reviewed state. In addition, the reviewed intents box changes its running count from 0 of 5 to 1 of 5.
6. If you're satisfied with your intent review, click the **Enable** button so the intent is available and can be accessed by any application, such as Virtual Agent (VA) or Search.

What to do next

Repeat Steps 1 through 6 for the four remaining intents in the list. As each intent review completes and is marked with the Reviewed state, the running count in the intents reviewed box increases toward 5 of 5.

Assign an NLU editor to a model

Assign an editor to review your Natural Language Understanding (NLU) model translations and edit model content. Delegate the maintenance, testing, and optimization of model content to an editor.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Create a model, or use an existing one. For translation review, create both primary and secondary (translated) models. For more information on multilingual model groups, see Multilingual model management.
- Assign the nlu_editor role to users. See Assign a role to a user.
- Role required: nlu_admin. The nlu_editor can't assign another editor to models.

About this task

The nlu_admin can delegate model content work to an nlu_editor. Assign language specialists as editors to review the consistency between primary and secondary (translated) models, or further localize the translation by editing intents and utterances. Delegate the maintenance of an existing model's content, without affecting the model's published status.

The nlu_editor can do the following, when assigned to a model:

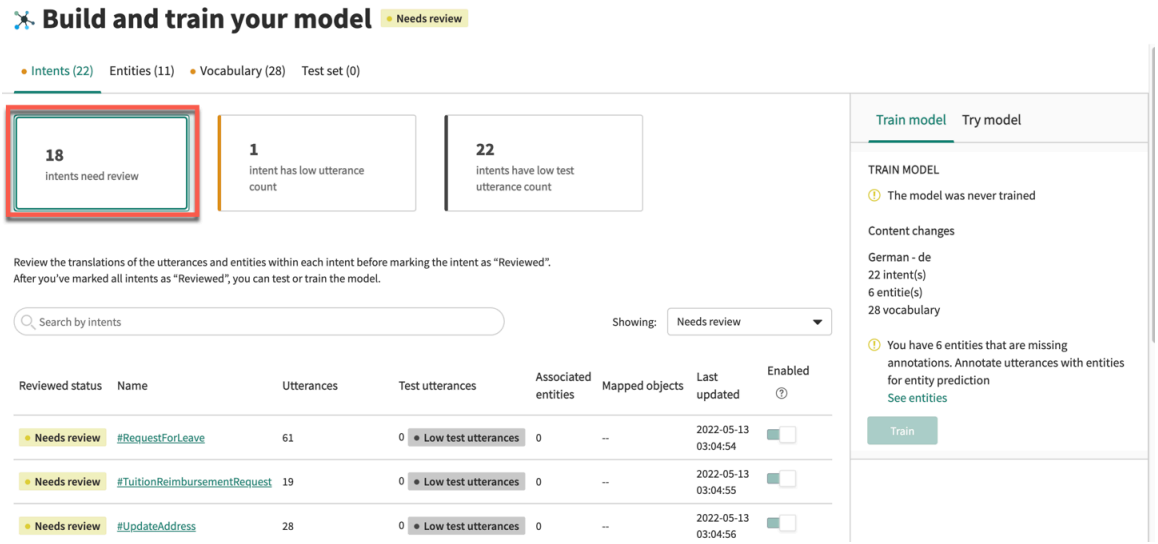
- Train and test the model.
- Export an NLU model as CSV.
- When a model or intent is in the **Needs review** state, the editor can mark it as **Reviewed**.
- Modify the model's confidence threshold.
- Edit the model name and description in **Settings**.
- Review, edit, and approve translated utterances.
- Add, edit, and delete training utterances.
- Annotate training utterances (add references to NLU vocabulary and NLU entities).
- Update entity properties (except the entity name in a secondary model).
- Add, edit, or delete the model's vocabulary.
- View the **Vocabulary sources** module (read only).
- Add, activate, and de-activate NLU intents in primary and secondary models.
- Resolve conflicting intents; also access the Cross-model Conflict Review module.
- View the Irrelevance detection in NLU module.
- Add, edit, or delete a test utterance in a test set; also access the Multi-model Batch Testing module.

Note: The editor must be assigned to a model and in the model's application scope to take these actions, unless otherwise indicated.

The nlu_editor can't do the following:

- Create, add, publish, or delete a model.
- Duplicate a model.
- Translate a model (initiate a translated secondary model).
- Import primary model content to a secondary (translated) model.
- Sync or edit vocabulary sources.
- View model performance.
- Provide feedback in NLU Expert Feedback Loop.
- Manage other editors, such as assigning an editor to a model or removing an editor from a model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



First, an admin assigns the nlu_editor role to a user. Next, the NLU admin assigns the editor to a model.

When you assign an editor to a model, you must do so at the individual model level. You can't assign an editor at the model group level.

The following models can't be assigned to an editor:

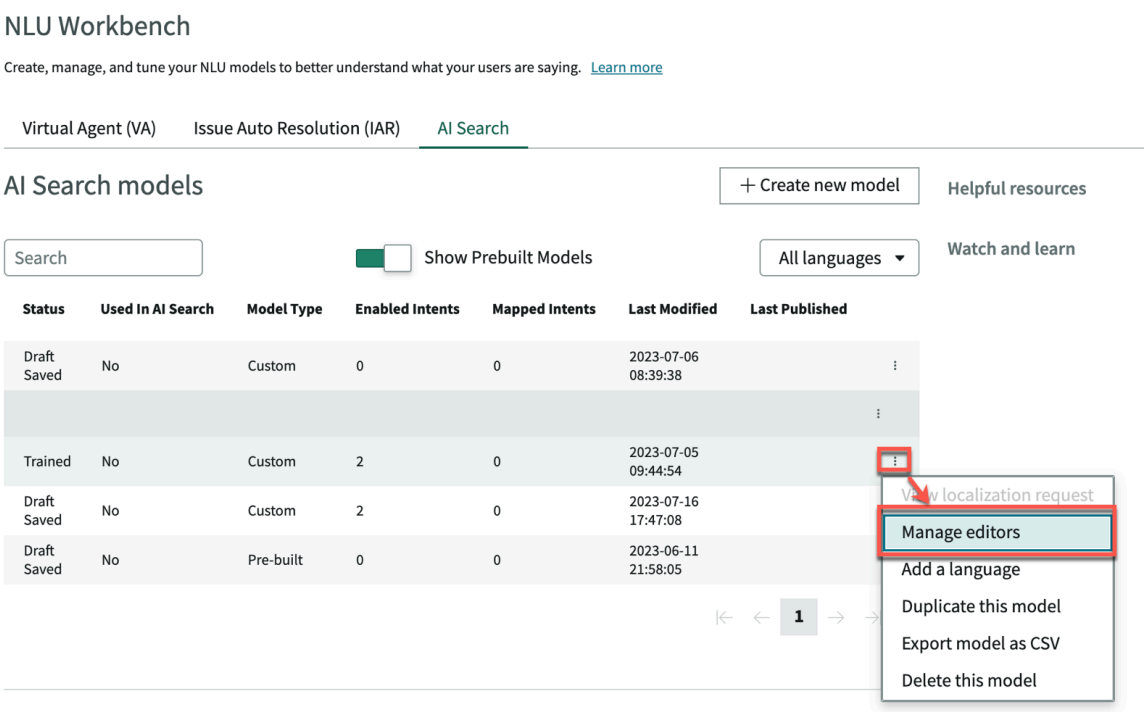
- Prebuilt models
- Models that have a translation in progress

In this example procedure, you add an NLU editor to a secondary model that needs reviewing.

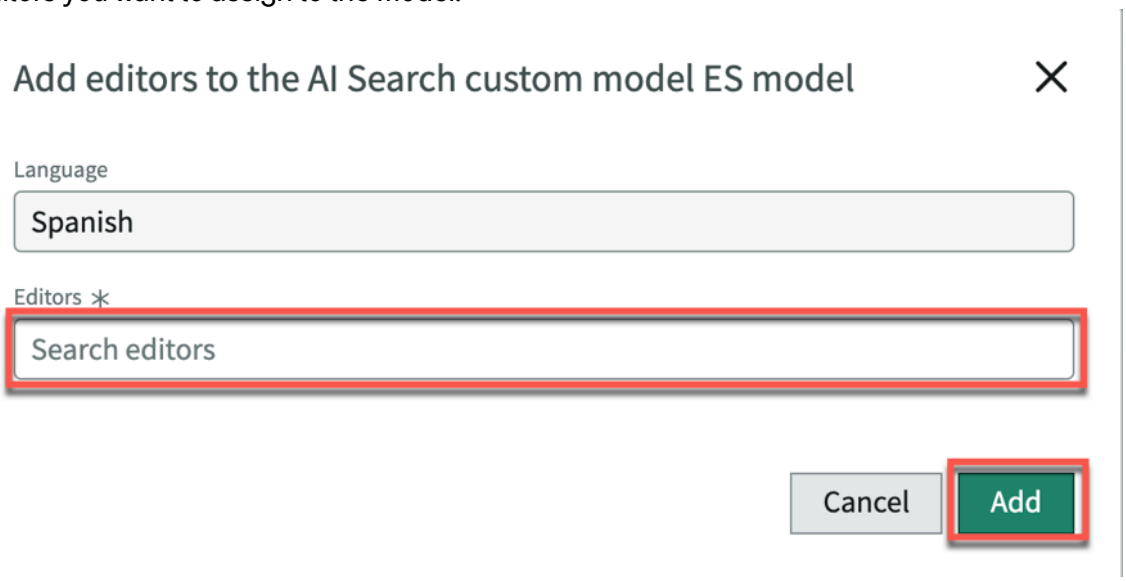
Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default. Select the appropriate tab for your model.
2. Locate the name of your NLU model and scroll to the right end of its row.
   If your model is in a multilingual group, expand the group to find the individual model. Editors cannot be assigned to groups.
3. Select the model's **More options** menu, then select **Manage editors** in the list.

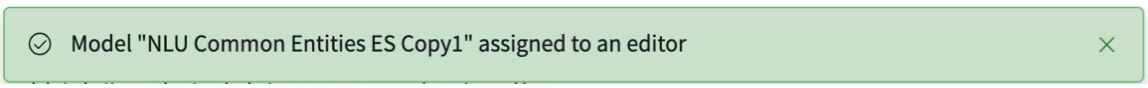
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. In the **Add editors** window, use the drop-down list and search bar to add the names of the editors you want to assign to the model.



5. Select **Add**.
   The editor is assigned to the model. A banner acknowledgement displays at the top of the screen. To add more editors, or remove editors, repeat this procedure.



What to do next

Assign more editors to the model, if necessary. You can assign up to four editors to a model at one time. To add more editors, or remove editors, repeat the above procedure.
As an NLU admin, ensure that your editors have what they need to properly review and edit models.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Import primary model content to a secondary model

When content in your primary NLU model is updated, you can import the updates directly to secondary models.

Before you begin

- Activate the Localization Framework (com.glide.localization_framework.installer) plugin. See Localization Framework.
- For Virtual Agent or AI Search models, create or use an existing primary model that has at least one secondary model.
- Role required: nlu_admin or admin.

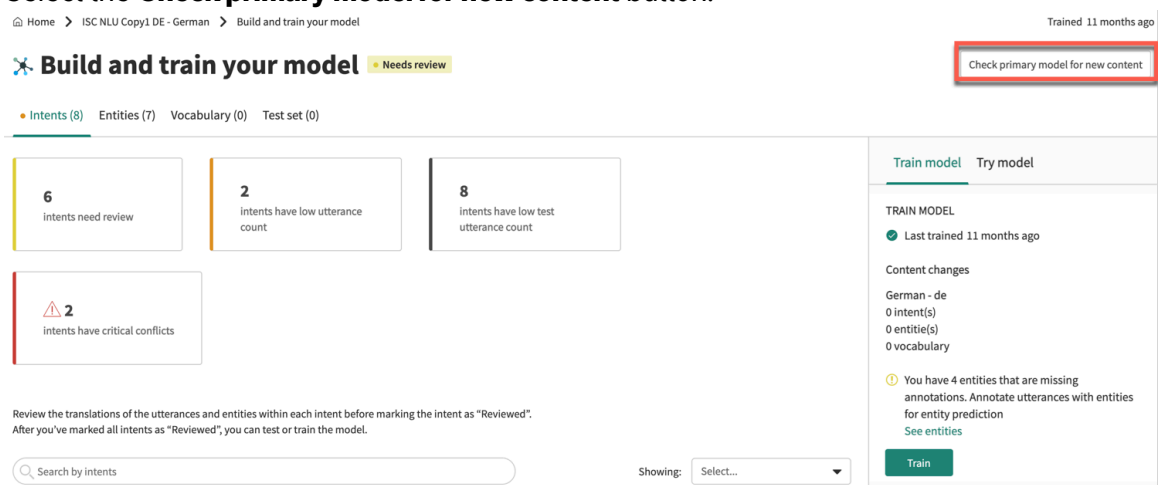
About this task

You can import content from a primary model to a secondary model after the initial translation. If you modify the primary model, those changes can be imported to your secondary models without the need to translate the entire primary model once again.

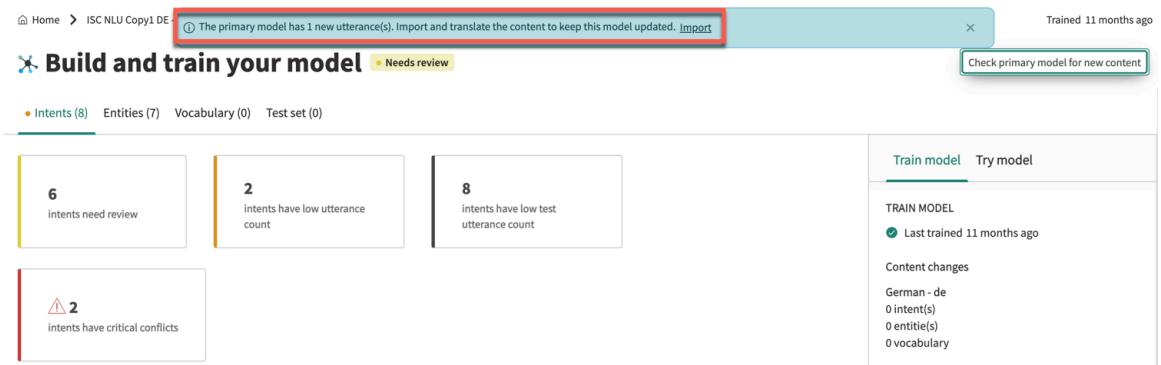
You can import primary model content to only one secondary model at a time.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default. Select the appropriate tab for your model.
2. Access a secondary model that has been translated under the primary model, and on **Build and train your model**, select **View phase**.
   By default you are taken to the **Intents** tab.
3. Select the **Check primary model for new content** button.



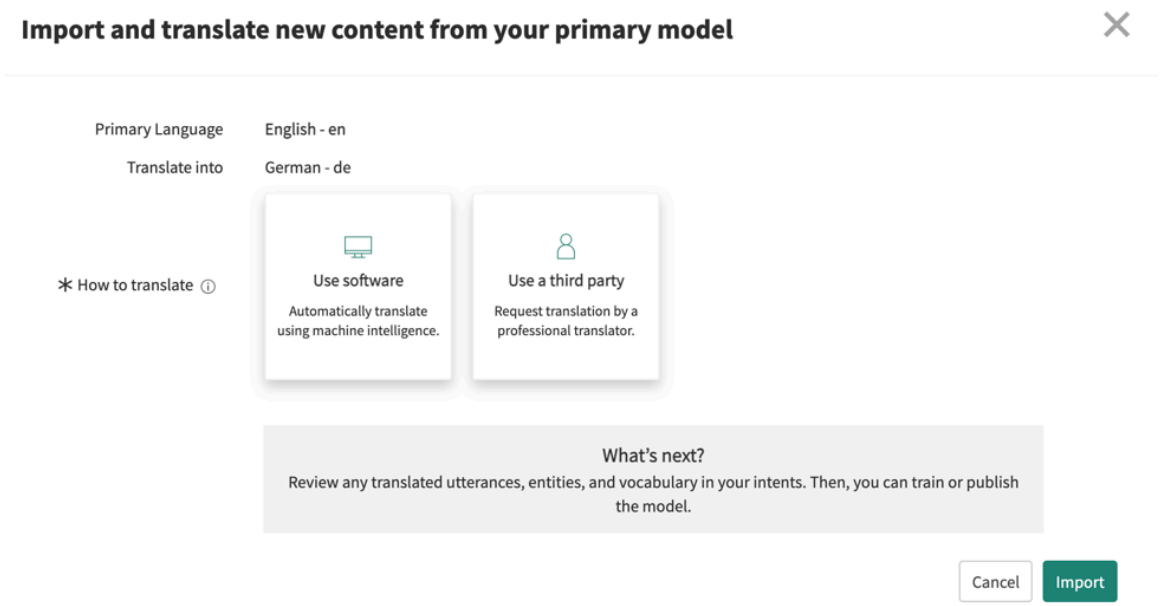
The banner Import and translate the content to keep this model updated appears at the top of the screen, along with a summary of the updates in the primary model.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: By default, newly imported intents are disabled in secondary models. Activate any imported intents manually.

4. In the banner, select **Import**.
5. In the **Import and translate new content from your primary model** window, choose a translation method and select **Import**.



The new content begins translating. When finished, the translated content must be reviewed.

What to do next

Make sure that each of your secondary models receive the new content from the primary model. Review all the new, translated content before publishing the secondary models.

Virtual Agent and NLU Workbench integration

Virtual Agent administrators can access and update their NLU models from within the Virtual Agent Designer user interface.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Note: If you have Now Assist in Virtual Agent, you can continue to use your existing NLU topics and migrate them into new LLM topics using the topic migration feature within Virtual Agent Designer. For more information on topic migration, see Migrate NLU topics to LLM topics.

Integration setup tasks, roles, and details

As Virtual Agent administrators create and configure their conversation topics, they must first create their NLU model and its associated intents in the NLU Workbench. This action requires they use the NLU Workbench and the admin or nlu_admin role.

In addition, they must also complete the following tasks in Virtual Agent **General Settings**.

- Enable NLU.
- Select the NLU service provider.
- If using language-specific NLU models, enable the languages for those models.

Virtual Agent administrators must also apply their NLU model to a conversation topic by completing the following tasks in Virtual Agent Designer:

- In **Topic Properties**, select the NLU model, the NLU intent, and the topic switching behavior.
- For input controls used in the topic flow, set the NLU properties for entity extraction.

Optionally, admins can activate Dialog Acts to enable Virtual Agent to respond flexibly when users make a modification in mid-conversation. Currently available response types are Modify, Affirm, and Negate, based on the last 5 exchanges in the conversation. Dialog Acts can be configured for English only, in Topic Properties. For more information see Dialog Acts for Virtual Agent.

After the NLU model is complete and associated with a Virtual Agent conversation topic, administrators with the virtual_agent_admin or admin role can do the following from within the Virtual Agent Designer user interface.

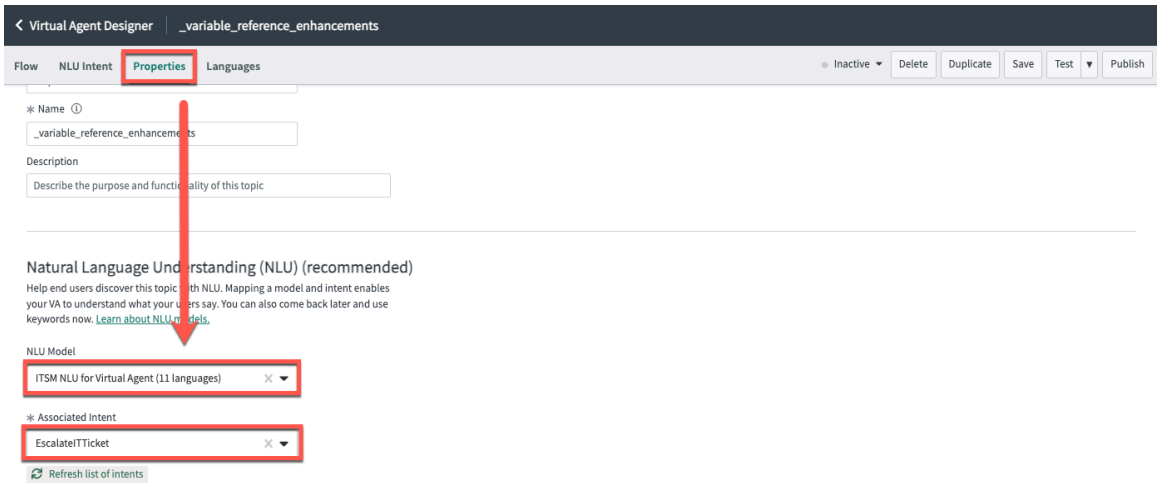
- Update NLU intent utterances.
- Train, test, and publish the NLU model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see Natural Language Understanding (NLU) topic discovery in Virtual Agent.

Publishing topics from Virtual Agent

When your model is published in the NLU Workbench, it's ready to use in Virtual Agent Designer. When editing a topic, click the **Properties** tab to select a model and intent to map to that topic.



When you click **Publish**, the model and intent are mapped to that topic and published seamlessly.

Publishing a topic with a mapped intent fails for the following reasons:

- Model isn't trained, or training is in progress.
- The last trained model is already published with a VA topic.
- The intent is not enabled in the model.

NLU Workbench - Advanced Features

NLU Workbench - Advanced Features expands the functionality of NLU Workbench to help you manage and improve your models.

Summary usage

The NLU Workbench - Advanced Features application includes the following features:

- Multi-model Batch Testing: Test large groups of utterances against your NLU models to see how the model predicts intents.
- Cross-model Conflict Review: Identify conflicting intents within or across models to improve model performance.
- NLU Model Performance: Use NLU Model Performance to see how well your models predicted intents in Virtual Agent based on end-user confirmation.
- NLU Expert Feedback Loop for VA: Provide feedback on Virtual Agent chat log utterances to help the system continuously learn and to better predict user input.
- Irrelevance detection in NLU: train your Virtual Agent models to ignore utterances that are not relevant.
- Issue Auto Resolution Tuning in NLU (Expert Feedback Loop for IAR): provide feedback to your ITSM Issue Auto Resolution model to improve its predictions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

After the application is installed and activated, Model Performance, Expert Feedback Loop for VA, and Expert Feedback Loop for IAR appear under **All** > **NLU Workbench**. Multi-model Batch Testing and Cross-model Conflict Review appear under **All** > **NLU Workbench** > **NLU Advanced Features**.

Installation

NLU Workbench - Advanced Features is available from the ServiceNow Store. For instructions on how to purchase and download, see Install NLU Workbench - Advanced Features.

Note: The available Intent Discovery ServiceNow Store application is installed separately from NLU Workbench - Advanced Features.

Related features

Install NLU Workbench - Advanced Features

You can install the NLU Workbench - Advanced Features application (com.snc.nlu.workbench.advanced) if you have the admin role.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the NLU Workbench - Advanced Features application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.
- NLU Workbench - Advanced Features requires the following plugins. Ensure that these plugins are activated before you install NLU Workbench - Advanced Features.

  Required ServiceNow plugins

  Predictive Intelligence (com.glide.platform_ml)

  Enables various Predictive Intelligence and Machine Learning capabilities for training models. See Predictive Intelligence.

  NLU Workbench - Core (com.glide.nlu)

  Adds NLU Model capabilities. See Activate the NLU Workbench

  NLU Workbench (com.snc.nlu_studio)

  Create and train NLU models. See Activate the NLU Workbench.

Role required: admin

About this task

Tables are installed with NLU Workbench - Advanced Features.

For more information, see Components installed with NLU Workbench - Advanced Features.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the NLU Workbench - Advanced Features application (com.snc.nlu.workbench.advanced) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you cannot find the application, you might have to request it from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Visit the ServiceNow Store 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. In the Application installation dialog box, review the application dependencies.

   Dependent plugins and applications are listed if they will be installed, are currently installed, or need to be installed. If any plugins or applications need to be installed, you must install them before you can install NLU Workbench - Advanced Features.

4. Select **Install**.

What to do next

The following available ServiceNow Store application is installed separately from NLU Workbench - Advanced Features:

Intent Discovery

Discover user intents from requests, incidents, or cases to help maximize deflection with Virtual Agent and NLU. For more information, see Install Intent Discovery.

Components installed with NLU Workbench - Advanced Features

Several types of components are installed with activation of the com.snc.nlu.workbench.advanced plugin, including tables.

Note: The Application Files table lists the components that are installed with this application. For instructions on how to access this table, see Find components installed with an application.

Tables installed

Table	Description
NLU Performance Ignored Clusters [nlu_performance_ignored_clusters]	Contains information on the ignored cluster. Filled when user clicks **Ignore**.
NLU Batch Test Result [nlu_batch_test_result]	Contains the parsed results of a batch test including each utterance and predicted intent.
NLU Conflict Execution [nlu_conflict_execution]	Stores the information related to each execution of a conflict analysis run for a given definition.
NLU Batch Test Run Definition [nlu_batch_test_run_definition]	Contains base information related to the batch test execution.
NLU Batch Test Run Execution [nlu_batch_test_run_execution]	Contains information related to each execution of the batch test.
NLU Conflict Result [nlu_conflict_result]	Stores the complete results of a conflict analysis execution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
NLU Batch Test Utterance [nlu_batch_test_utterance]	Contains the utterances used for a test set. Populated by the import set.
NLU Batch Test Set [nlu_batch_test_set]	Includes the information for the batch test. Referenced by the utterance record and run definition.
NLU Performance Utterance Trace [nlu_performance_utterance_trace]	Contains the information on the utterances added to an intent.
NLU Conflict Definition [nlu_conflict_definition]	Contains the details of the NLU model or pair of models used for the conflict.
NLU Model Analysis Definition [nlu_analysis_definition]	Supports IAR model tuning in NLU Workbench.
NLU Model Analysis Execution [nlu_analysis_execution]	Stores records for each instance of IAR model tuning in NLU Workbench.
ML Label User Feedback [ml_label_user_feedback]	Stores utterances that are marked Needs review.

Role installed

NLU Workbench - Advanced Features

Role	Description
NLU Feedback Admin [nlu_feedback_admin]	Data labelling (NLU feedback) Admin role - to manage data labelling across models also ability to optimize models

Multi-model Batch Testing

Test multiple Natural Language Understanding (NLU) models against a large set of utterances to evaluate the performance of the models. Add test sets, test multiple models, and see test results.

Summary usage

Use Multi-model Batch Testing to create and upload test sets comprised of utterances and their expected intents. You can then run tests against your NLU models.

Multi-model Batch Testing works with models for all supported NLU languages. See NLU language support.

Installation

Multi-model Batch Testing is part of the NLU Workbench - Advanced Features app available on the ServiceNow® Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To use Multi-model Batch Testing, ensure that the NLU Workbench - Advanced Features (com.snc.nlu.workbench.advanced) plugin is active on your instance. For more information, see Install NLU Workbench - Advanced Features and Activate the NLU Workbench.

Test sets

Test sets are lists of utterances and matched intents. Create a test set by using a table in a CSV or XLSX (Excel workbook) file. The table should contain two columns: one for utterances, and one for the expected intent. Your test set can include up to 10,000 rows.

To get the most out of testing your NLU models, your test sets should include utterances that the model is likely to encounter from your users. Test utterances should be in the same language as the model to be tested. The test set should also include utterances with no expected intents. Including utterances with no expected intent helps assess your model's ability to detect utterances which are irrelevant and shouldn't have any intent predicted.

By including these types of utterances, the test better assesses the model's ability to perceive intents and respond to your users. If your test set does not cover at least 60% of the intents of the models, you can still run the test but the recommended threshold may not be optimal.

Note: Certain test utterances are skipped during the test if their expected intent does not match any intents in the models.

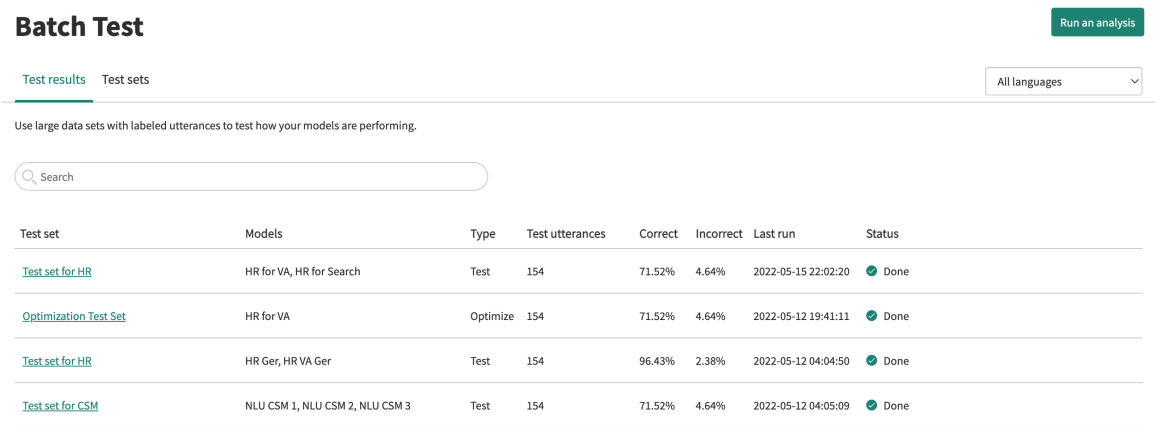
To create a test set, see Create a test set.

After you have a test set, you can test trained NLU models. To begin testing, see Run a multi-model batch test.

After running a test, your results appear on the **Test results** page.

Test results

The **Test results** page lists your completed and in-progress tests. At a glance, the results page shows the models tested against, the number of utterances, and prediction percentages.



To see the details of a test result, click the name of the test set.

The **Overview** page shows summary information about the results and includes a graphic with a breakdown of predictions.

The **Intents that need attention (Current model)** shows the top 5 missed and incorrect intents. Click the intent name to drill down into the test utterances that were predicted incorrectly. Use this information to improve the model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The **Detailed results** tab lists information about each utterance that was tested. From here, you can see the prediction outcome and confidence per model for each utterance. Filter the results by using the search bar or interacting with the filter tools and column headers.

You can also export the test results to a CSV file by clicking **Export**. The file includes the same columns as the detailed results page.

For more information on understanding your test results, see Test and publish your model.

Create a test set

To create or add to an NLU test set, you can upload a file of test utterances matched with correct intents. Use the test set to assess the performance of your model.

Before you begin

- Ensure that the NLU Workbench - Core plugin, NLU Workbench plugin, NLU Workbench - Advanced Features plugin and Predictive Intelligence plugin are all installed and activated
- You can use test sets with NLU models for Virtual Agent and AI Search.
- Role required: nlu_editor, nlu_admin, or admin. The editor must be assigned to the model.

About this task

Your CSV or XLSX (Excel Workbook) file should contain a table that pairs your test utterances with the intents that you expect for them. Your file can contain up to 10,000 utterances. Ensure that the file has columns titled "Utterance" and "Expected intent".

Note: For test sets in languages other than English, you must add the glide.import.csv.charset system property with the value UTF-8. See Import sets properties.

For the most accurate test results, include utterances that the model is likely to encounter from your users. Ensure that you have test utterances covering all the intents in your model.

Aim to include about 10% of test utterances with no expected intents. Including utterances with no expected intent helps assess your model's ability to detect irrelevant utterances that should not have any intent predicted.

Note: To indicate that a test utterance in your file has no expected intent, the value for "Expected intent" should be empty.

Example test set table

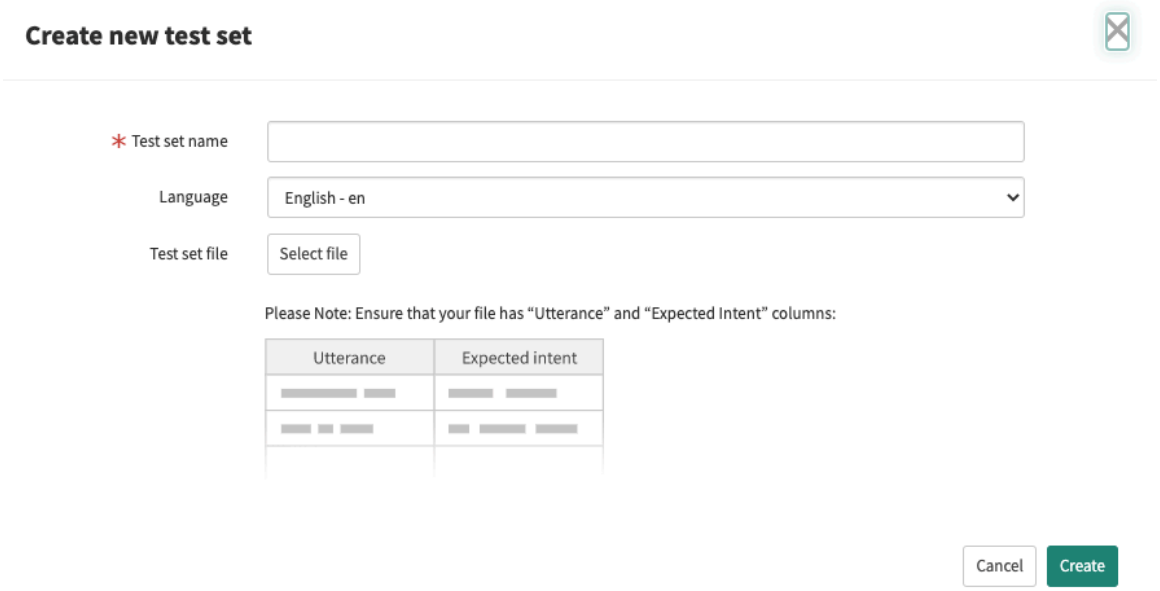
Utterance	Expected intent
Let me have a burger	Order
I want to pay	Payment
Get me something sweet	Order
Is the restaurant open	
Something wrong with my payment	Order, Payment
Total cost	Payment

Procedure

1. Navigate to **All** > **NLU Workbench** > **NLU Advanced Features** > **Multi-model Batch Testing**.
2. Click **Test sets**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Click **Create test set**.



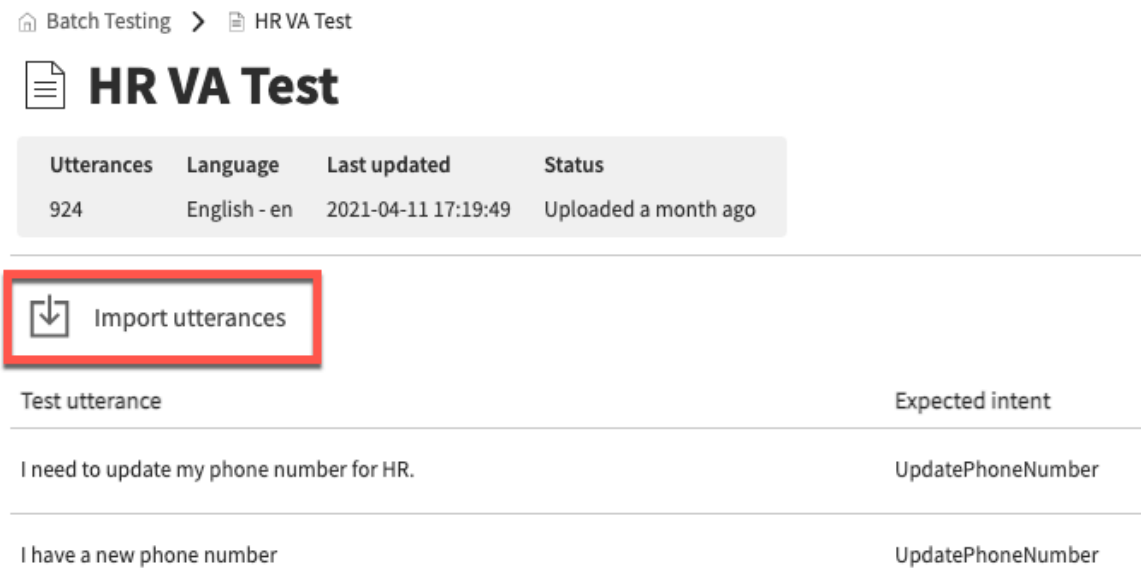
4. Choose a name for the test set.
5. Choose a language.
6. Click **Select file** and choose a CSV or XLSX (Excel Workbook) file.
7. Click **Create**.
   Your test set appears in the list.

What to do next

Use the test set to run a test on your models. To learn how, see Test your model or Run a multi-model batch test.

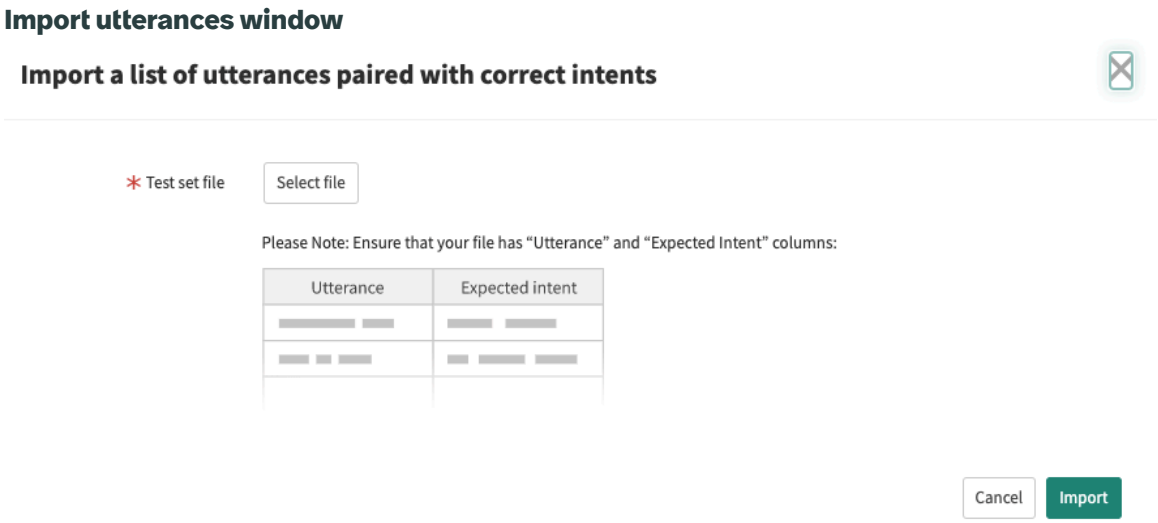
After you have created a test set, you can add more utterances to it. When viewing a test set, click **Import utterances**.

Test set page with import utterances button



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Select a CSV or XLSX (Excel Workbook) file with your additional utterances to import.



Click **Import**. The system adds your utterances to the test set. After importing, rerun any tests that use the test set.

Run a multi-model batch test

Test multiple Natural Language Understanding (NLU) models against a test set. Evaluate the quality of your models and refine them to improve intent prediction.

Before you begin

- Make sure that the NLU Workbench - Core plugin, NLU Workbench plugin, NLU Workbench - Advanced Features plugin and Predictive Intelligence plugin are all installed and activated.
- Have one or more trained models for Virtual Agent or AI Search.
- Have a test set containing test utterances with expected intents. See Create a test set or Test set creation and management.
- Role required: nlu_admin or admin. When assigned to a model, the nlu_editor role can run tests and modify test utterances for that model.

About this task

In **Multi-model Batch Testing**, you can use a test set that is not the model's default test set.

Also, you can test up to ten models at one time using **Multi-model Batch Testing**. However, tests with fewer models run more quickly.

When testing multiple models, your test set must cover at least 25% of the total intents of all the models. Use test sets that contain utterances the models are likely to encounter in Virtual Agent or AI Search.

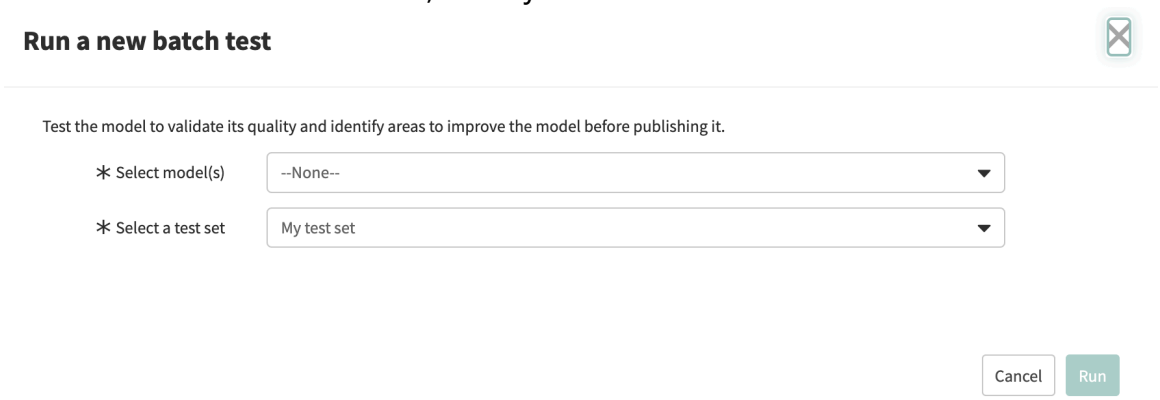
Note: If an expected intent in your test set doesn't match any intent in your models, that expected intent and its test utterances are skipped. They aren't counted or displayed in test results.

To test a single model against its default test set, use the **Test and publish your model** phase on the model's overview page. For more information, see Test and publish your model.

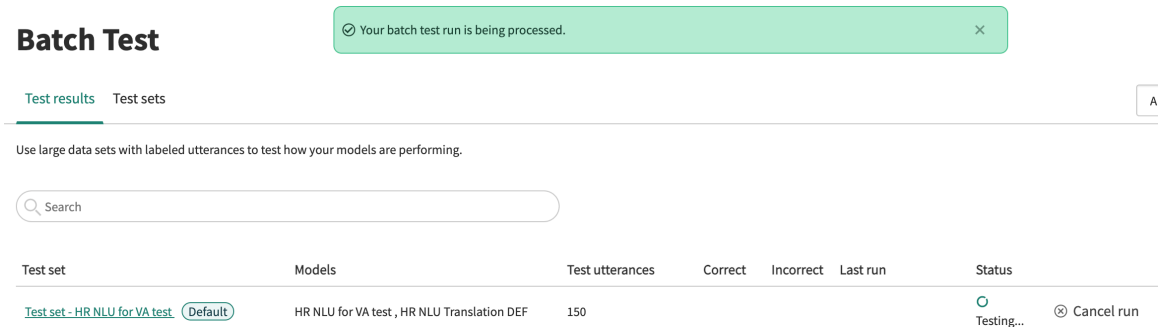
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **NLU Workbench** > **NLU Advanced Features** > **Multi-model Batch Testing**.
2. Click **Run a test**.
3. In the **Run new batch test** window, select your models from the list.



4. Select a test set from the list.
5. Click **Run**.



Your new test shows in the **Test results** list with a **Testing...** Status.

What to do next

When the batch test is finished, its Status changes to **Done**. You can click the name of the test set to view the test results.

Use the results to adjust and improve your models. Then, run the test again to assess the performance.

Batch testing may affect confidence threshold recommendations. For more information, see NLU model settings.

Cross-model Conflict Review

Identify conflicting intents within or across models so you can take corrective actions, resolve such conflicts, and improve your NLU model performance.

Summary usage

As the number of intents within a model increases, two intents may overlap in scope. This may occur when training utterances in one intent are almost identical to utterances in another intent. There may also be conflicting intents across models and even applications.

Utterances may also be marked as `Not relevant`, meaning that no intent should be predicted. When these irrelevant utterances are too similar to utterances assigned to an intent, they are displayed in Conflict Review. For more information, see Irrelevance detection in NLU.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To address and fix these issues, Cross-model Conflict Review runs an analysis on your models. Use the analysis to identify and resolve these issues prior to model publication and deployment.

Installation

Cross-model Conflict Review is part of the NLU Workbench - Advanced Features app available on the ServiceNow® Store.

To use Cross-model Conflict Review, ensure that the NLU Workbench - Advanced Features (com.snc.nlu.workbench.advanced) plugin is active on your instance. For more information, see Install NLU Workbench - Advanced Features and Activate the NLU Workbench.

Roles

To access Cross-model Conflict Review, use the nlu_admin or admin role. When assigned to a model, the nlu_editor can resolve conflicts in that model.

Running the analysis

The Conflict Review screen shows a list view of all conflict reviews created in your instance. When a review is completed, it's added to a running list of reviews. In this example scenario, you're creating the first review in your instance, so when it's completed, it's shown in the count as 1 of 1 reviews. As more conflicts arise over time, you will see multiple reviews in the list.

Conflict reviews are analyzed on either one or two NLU models. When you run an analysis on a single model, the system shows intents and utterances that are only in that model. When you run an analysis on 2 models, the system shows intents and utterances that are in both models.

Conflict reviews always run on the last trained version of the model(s) they analyze.

Conflict reviews have two types: Critical and Moderate. The standard approach is to begin with the critical ones.

When a conflict is detected, you can use one of following actions to resolve the conflict:

- Ignore the conflict
- Delete an identical or nearly identical utterance from one of the intents
- Edit the utterances to make them more distinct from each other

In this example scenario, you're resolving a conflict where two different intents contain the exact same utterance.

1. Navigate to **All** > **NLU Workbench** > **NLU Advanced Features** > **Cross-model Conflict Review**.

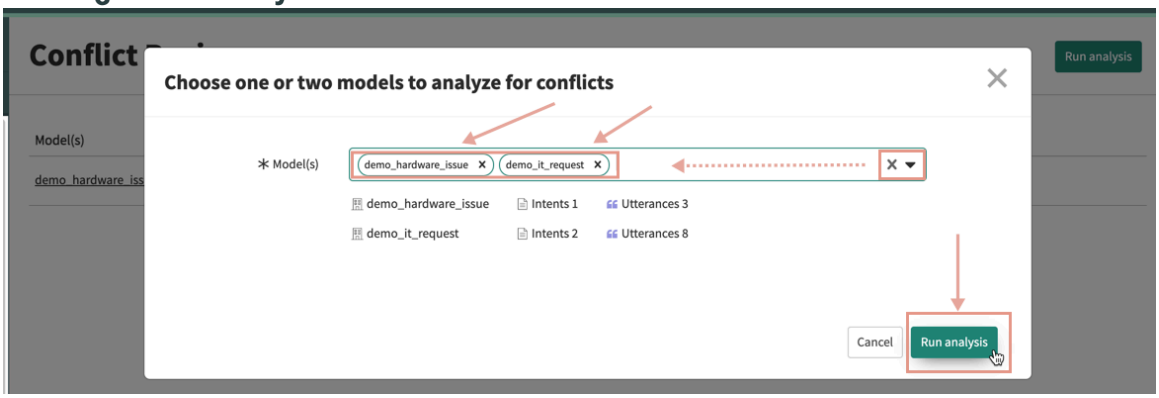
2. Select **Run analysis**.

3. In the **Model(s)** field on the Choose one or two models to analyze for conflicts screen, select two NLU models for the analysis. In this example scenario, you choose the demo_hardware_issue and demo_it_request models.

4. Select **Run Analysis**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

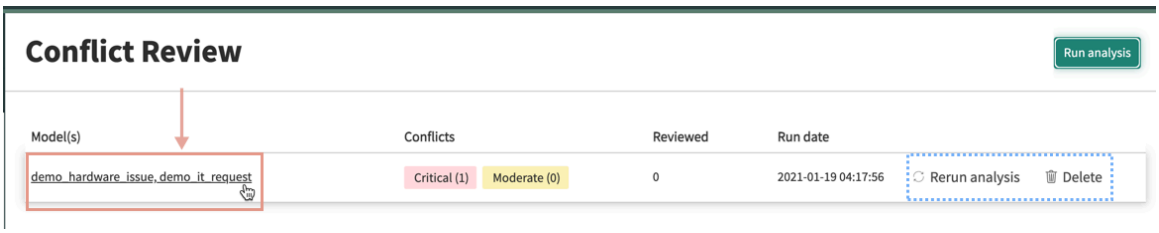
Running conflict analysis



The Conflict Review screen refreshes to show the analysis, including the two models you selected for analysis, the counts of Critical and Moderate conflicts under review, the number of reviews that have been completed, and the run date for the analysis. If you point to the far right column on the screen you will see options to rerun the analysis, or delete it and start all over.

5. In the **Model(s)** column, select your two paired models so you can drill down into the review.

Conflict review

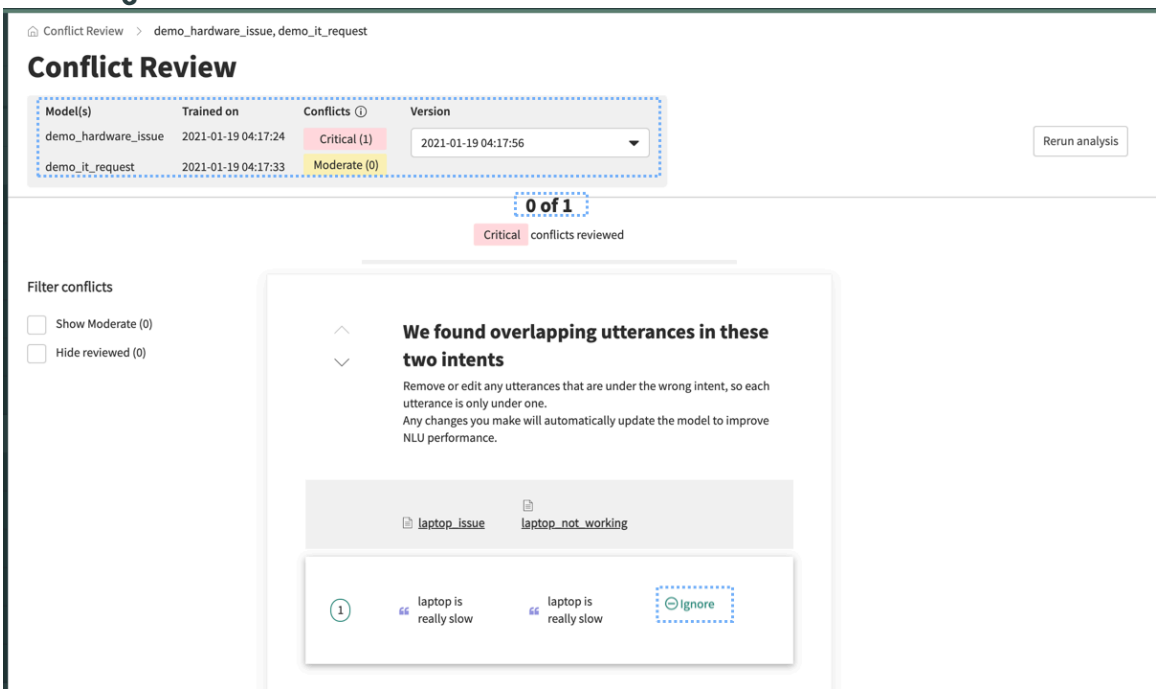


The screen refreshes to show the details of the conflict review. Note the summary of the two models you chose for the analysis, their latest training dates, the types of conflicts they hold, and the version time stamp of the analysis. Note also the 0 of 1 count, which indicates that this is the first conflict review created in this instance. As the system detects more conflicts over time, and those conflicts are reviewed, the count increases.

If you were to determine the utterances in the image below are fine as they are, then you should click **Ignore**. Clicking Ignore tells the system you've completed your review, so it marks it as reviewed and moves on to the next conflict review. However, in this scenario, you don't ignore the conflict, because intents that share the same utterance are a conflict worthy of review.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Reviewing conflict details



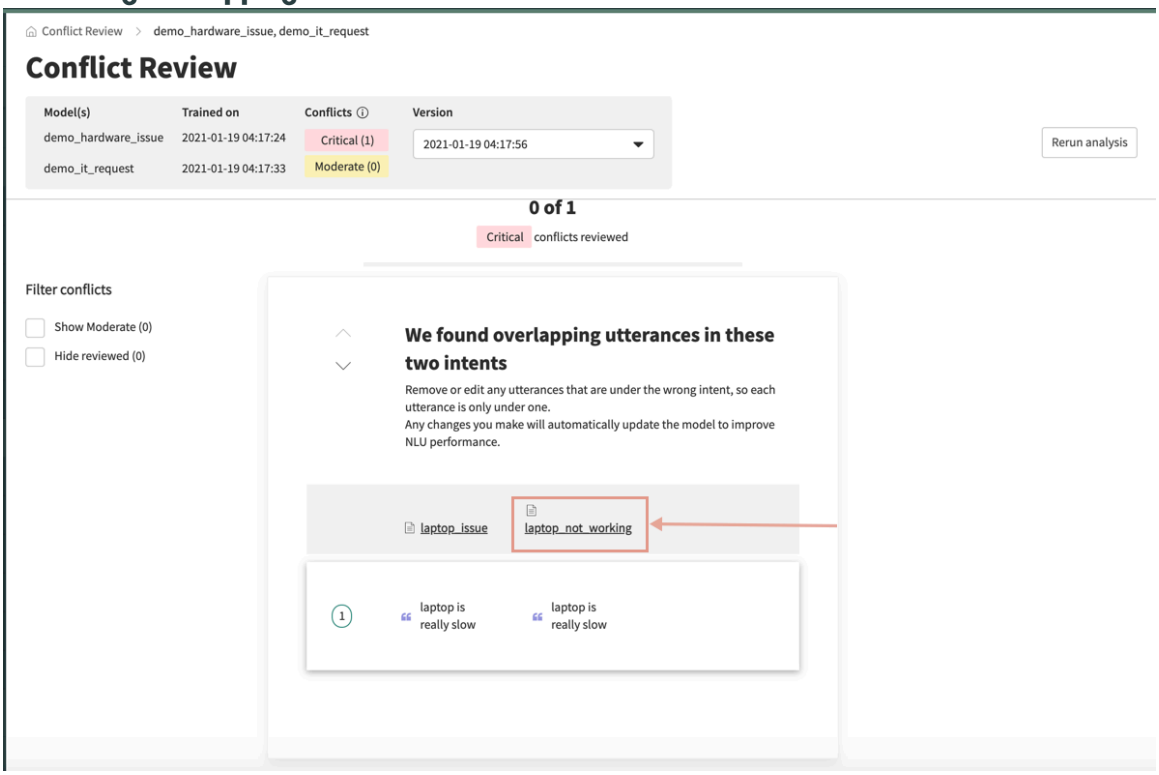
To determine how you will resolve this conflict, consider the 2 intent names and the identical utterances they share. Consider which intent is more likely to use the laptop is really slow utterance. If you compare the 2 intent names closely, you might realize that a laptop that's really slow isn't the same as a laptop that doesn't work. However, a laptop that's really slow is indeed a laptop issue. So in this example scenario, you decide to dig deeper into the intents to scan the context of their other utterances.

Note also that when an intent uses unique utterances, it helps the system to more accurately predict which utterances belong to it. Hence, you will need to edit or delete the utterance from 1 of the 2 intents. In this example scenario, you decide to dig into the #laptop_not_working intent.

6. Select **laptop_not_working**.

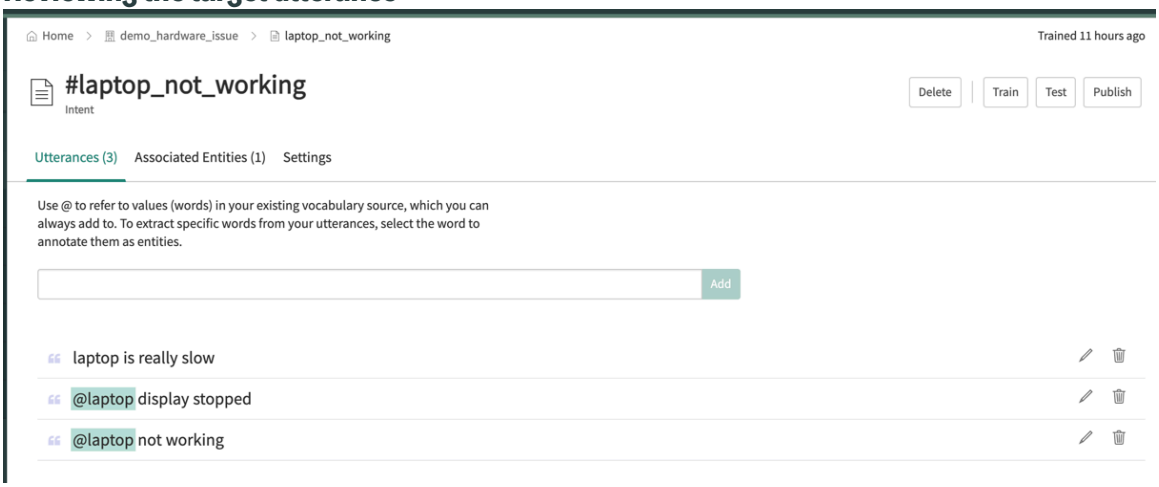
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Reviewing overlapping utterances



The #laptop_not_working Intent screen appears, showing its current 3 utterances. You make the decision to delete the laptop is really slow utterance from the #laptop_not_working intent.

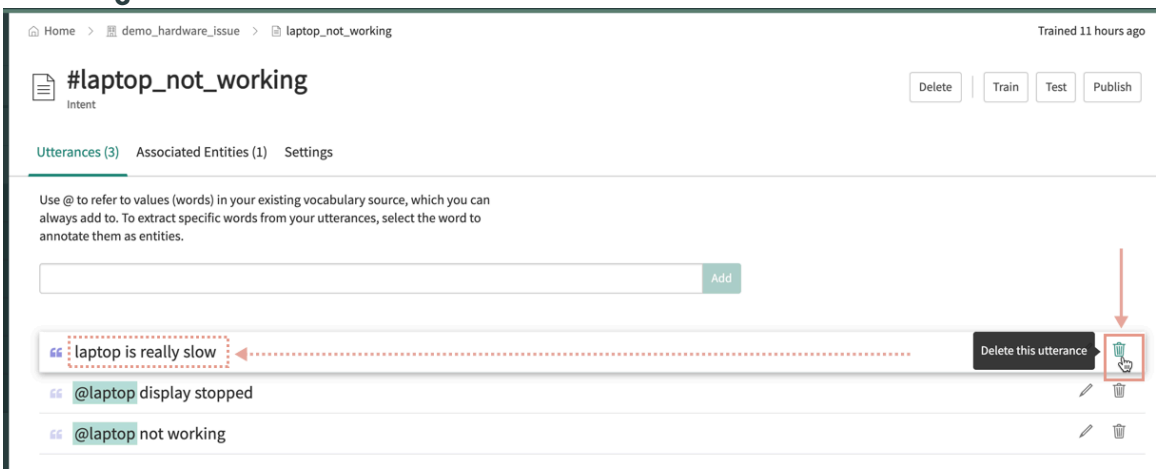
Reviewing the target utterance



7. Click the **Delete this utterance** trash can icon.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

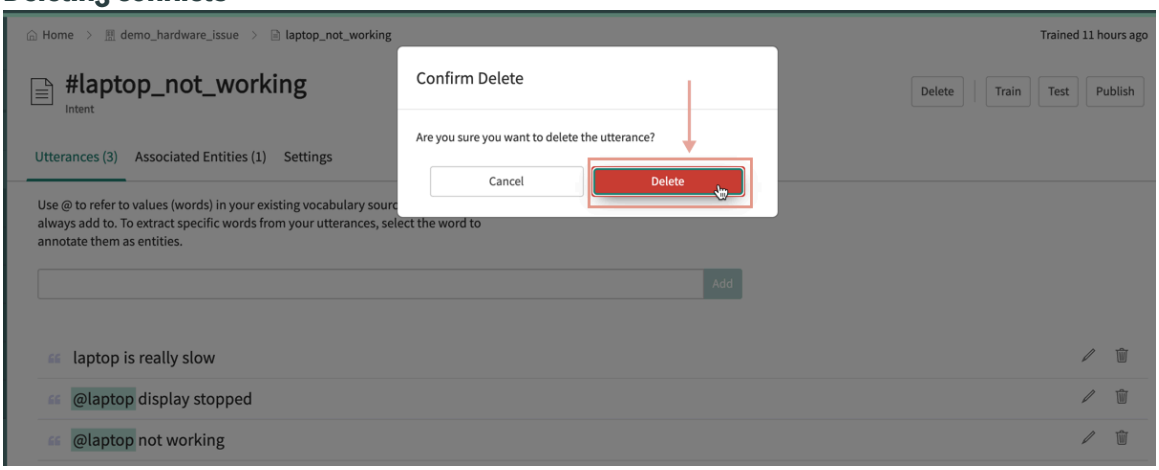
Resolving conflicts



The Confirm Delete screen appears.

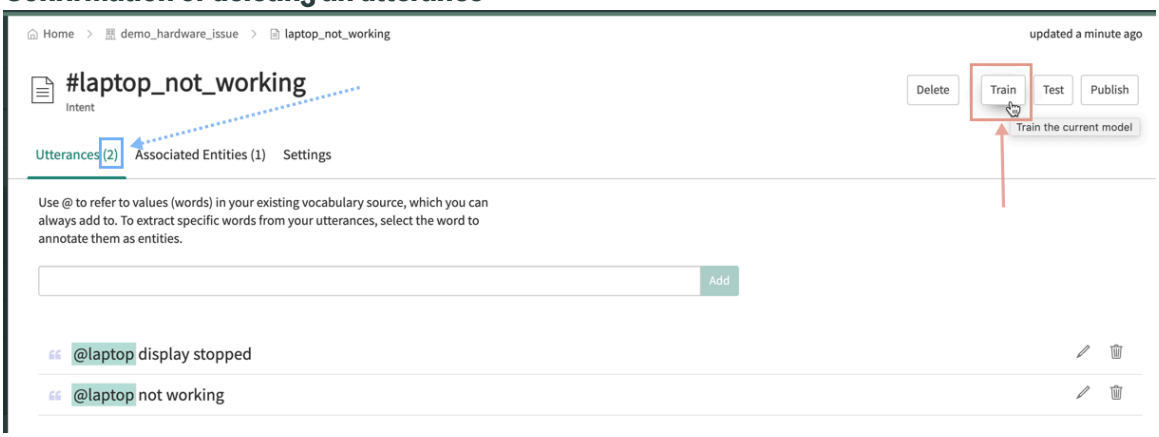
8. Select **Delete**.

Deleting conflicts



The Confirm Delete screen disappears, and the Utterances count drops from 3 to 2 because you've deleted the laptop is really slow utterance from the intent.

Confirmation of deleting an utterance



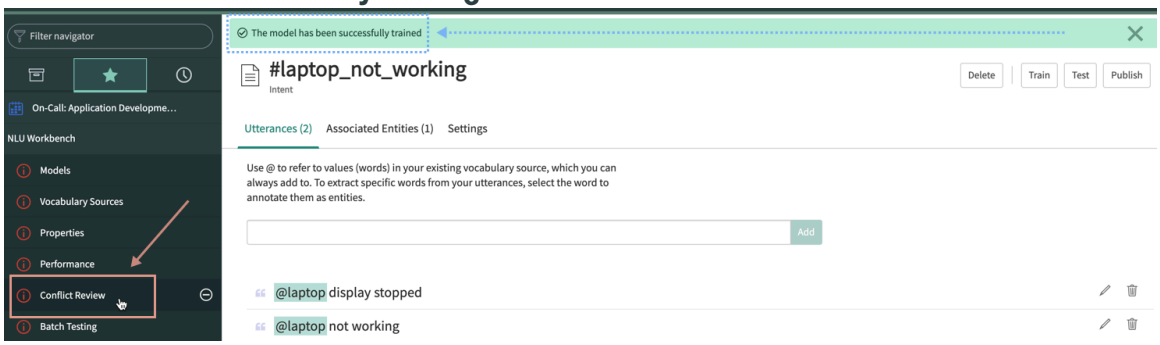
9. Select **Train**.

A banner appears on the Intent screen, confirming the model is successfully trained.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

10. Select **Conflict Review** in the navigator.

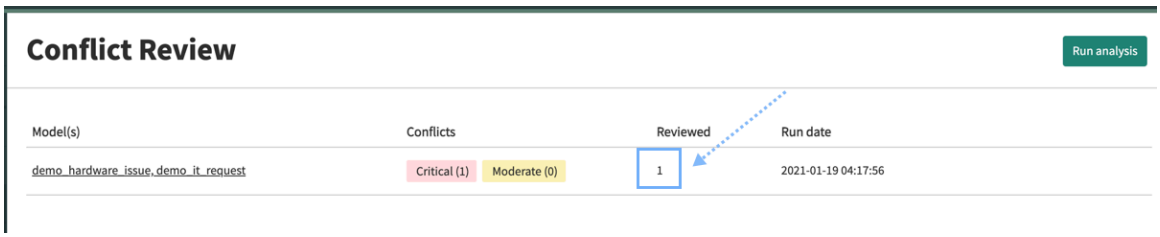
Confirmation of successfully training a model



Result:

The Conflict Review list screen appears, showing your conflict review analysis is complete, and that it's been reviewed.

Confirmation that a conflict has been reviewed



NLU Model Performance

Use NLU Model Performance to see how well your models predicted intents in Virtual Agent (VA) based on end-user confirmation.

Summary usage

The NLU Model Performance application provides an analysis and report of how well your NLU models predict VA users' intents from their utterances. As NLU models for VA are updated, published, and deployed, you can use NLU Model Performance to review the efficacy of the intent predictions they make. If these predictions are skipped, it's because they're unable to predict an intent with a high enough confidence score for the model. To improve VA topic prediction, the system groups unsupported VA utterances into clusters for an analysis and then generates a report that identifies the outcome of the analysis.

To access this application, use the nlu_admin or admin role and navigate to **All** > **NLU Workbench** > **Performance**.

Note: To use NLU Model Performance, you must first have at least 5,000 VA utterances in your instance. With the nlu_admin role, you can modify this limit by resetting the *sn_nlu_workbench.glide.nlu.performance.min_clustering_records* property.

Installation

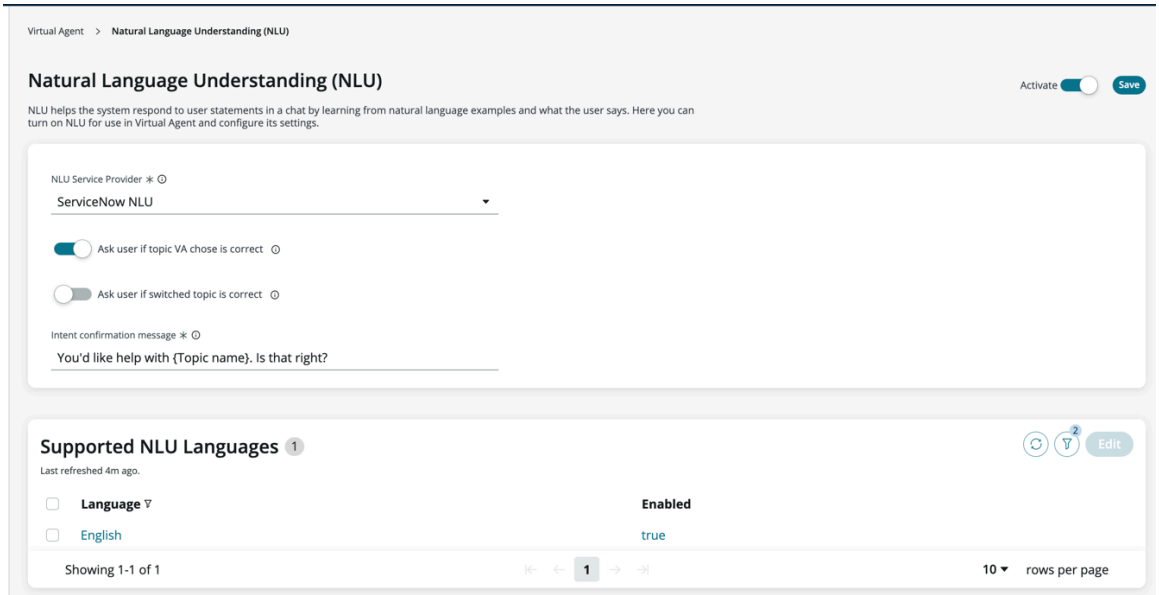
NLU Model Performance is an application available on the ServiceNow® Store. To use this application, ensure that the NLU Workbench - Advanced Features (sn_nlu_workbench) plugin is active on your instance. For more information, see Install NLU Workbench - Advanced Features and Activate the NLU Workbench. Because NLU Model Performance relies on data coming in

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

from VA, ensure that the Glide Virtual Agent (com.glide.cs.chatbot) plugin is also active. See Activate Virtual Agent.

Configuring Conversational Interfaces settings for ServiceNow NLU

To help with tracking NLU performance, you must first configure the Conversational Interfaces (CI) settings in Virtual Agent. To access these settings, navigate to the top of the **NLU Model Performance for Virtual Agent** landing page and click **Conversational Interfaces Settings**. This action takes you to the CI settings page in Virtual Agent, as shown in the image below.



To configure the settings, you need to use the virtual_agent_admin or admin role. For configuration guidance, see Implement NLU in Virtual Agent and Enable NLU languages in Virtual Agent settings.

Reviewing the user interface (UI)

The NLU Model Performance UI is comprised of the listed parts below, which are highlighted in the following image that shows the NLU Model Performance default landing page.

- The **Performance** tab and its colored chart which shows data describing scenarios that occur when users interact with VA. This chart is segmented in shades of three colors, each representing a user scenario. The rows in the chart that have text and percentages are referred to as the legend. The three scenarios in the first row of the legend have two outcome scenarios beneath them in the second row. The bars underneath the legend correspond to the colors of the scenarios above them. These bars range in size based on the percentage value for each scenario.
  - For example, the dark blue bar shows you that 10% of this batch of VA users confirmed that the system presented them the correct topic to use for their VA chatbot conversation, and that 58 % of the users (represented by the dark red bar) confirmed that the topic presented to them was incorrect. Residing between these two scenarios is a scenario showing that 32% of the users (represented by the dark brown bar) didn't make any confirmation at all.
  - It's a good idea to toggle the **Show all bars** filter back and forth at any time to see or hide the colored bars and their associated scenarios in the chart. Note also that when you first access NLU Model Performance, the default view of the chart shows only the first row of scenarios and the first row of bars.
  - There are also other views you can see depending on which scenario you're viewing. For example, when you first access NLU Model Performance and you click a scenario in the first

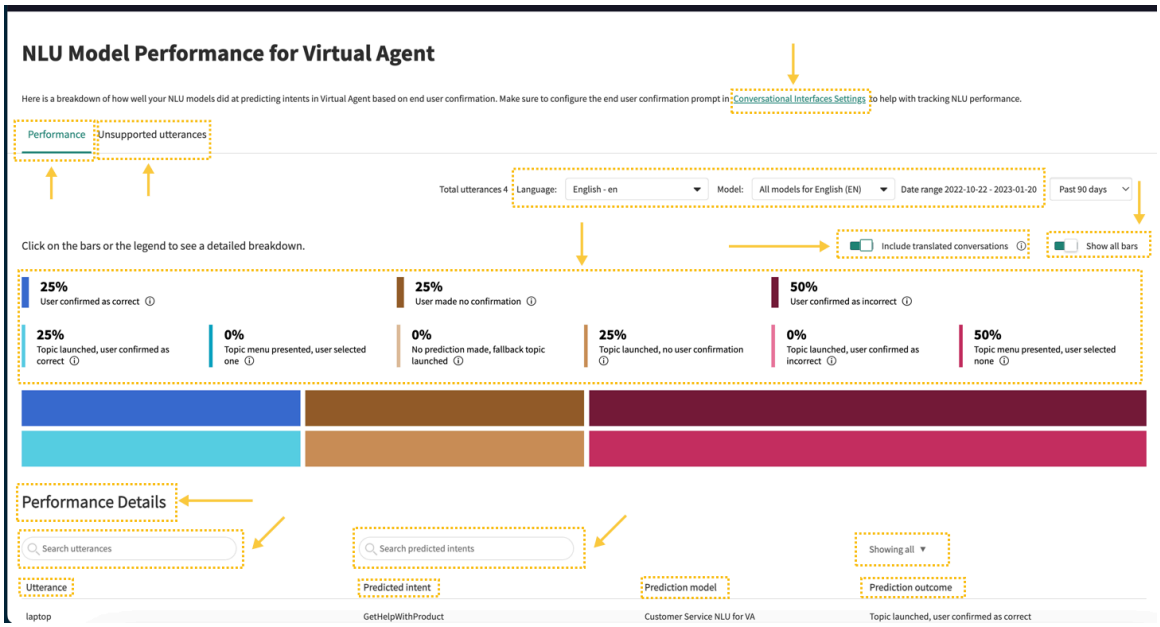
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

row of the legend, the system hides the bars of the other two scenarios. This helps to keep the UI less cluttered because it partially isolates the scenario you're focusing on from the rest of the scenarios. These actions don't change any data, they simply show you different views of the legend and bars within the chart.

- The **Unsupported utterances** tab, which shows utterances that are grouped into clusters for use in the performance analysis. This section of the UI is where you can navigate to and return from the Expert Feedback Loop application, and also where you can run the analysis.

  Unsupported utterances are different from irrelevant utterances. For more information, see Irrelevance detection in NLU.

- The **Performance Details** section located underneath the bar chart. This section has these four columns: **Utterance**, **Predicted intent**, **Prediction model**, and **Prediction outcome**. The details shown in this section interact with the legend data and bar data above them.



Reviewing the VA user scenarios

In the legend section of the UI, the text and percentages are accompanied by Information icons. Point to an icon to invoke the definition for its user scenario. Refer to the table below to see the definitions.

Scenario Definitions

Scenario	Definition
User confirmed as correct	The correct topic was presented to the end user and the end user has confirmed it is correct.
Topic launched, user confirmed as correct	The topic was launched and the user confirmed it is the topic they needed.
Topic menu presented, user selected one	Multiple topics were shown as a menu to the user and the user chose one of the topics to address their need.
User made no confirmation	The topic was launched but the user did not confirm whether or not it met their need.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

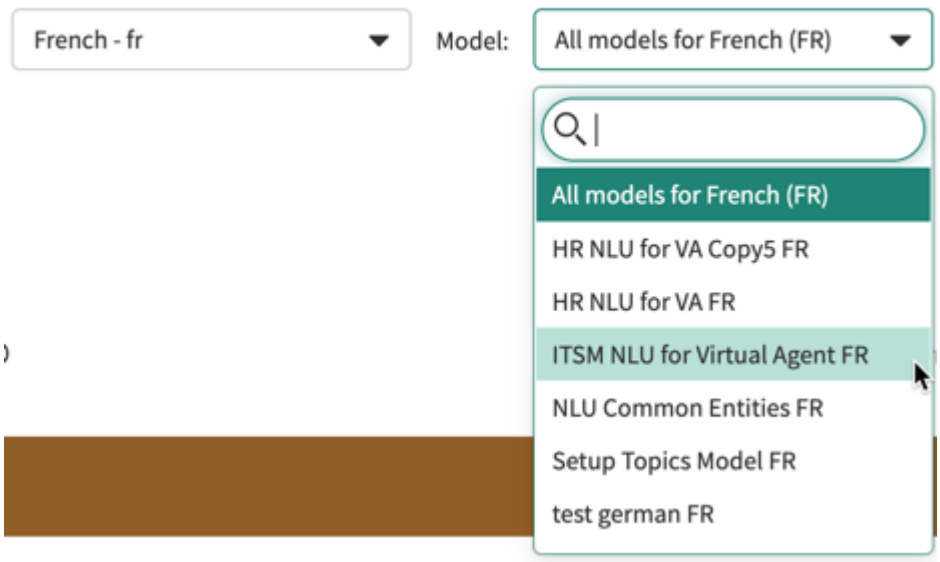
Scenario Definitions (continued)

Scenario	Definition
No prediction made, fallback topic launched	No predictions were made and the fallback topic was launched.
Topic launched, no user confirmation	The topic was launched but the user did not confirm whether it was correct.
User confirmed as incorrect	The topic(s) presented did not address the end user's needs, and the user has confirmed it is not correct.
Topic launched, user confirmed as incorrect	The topic was launched and the user confirmed it was not correct.
Topic menu presented, user selected none	Multiple topics were shown as a menu to the user and the user decided none of the topics were relevant and did not choose any topic.

Using the Language and Model filters

On the Performance tab, you can see the Language and Model filters. Next to them, you can also see the most recent Date range values you set above the legend and bars of the chart. Click the **Language** filter to see all languages that are available in NLU. Click the **Model** filter to see all prediction models that are available in your instance. The default setting for the **Language** filter is **All languages**, and the default setting for the **Model** filter is **All models**.

The Language and Model filters interact with each other. For example, if you choose `French - fr` in the **Language** filter, the **Model** filter automatically shows all prediction models in your instance that use the French language, as shown in the following image.



There are other interactions. For example, from the default view of the NLU Model Performance landing page:

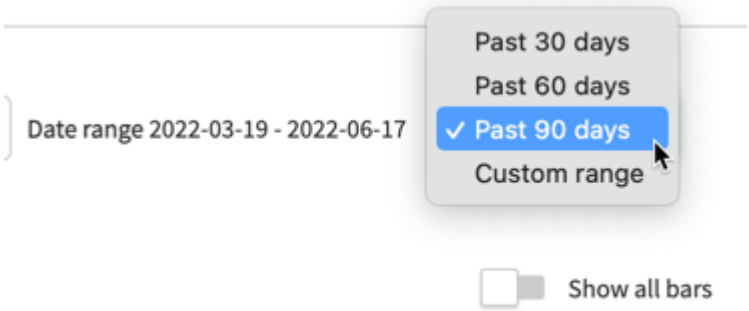
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If you select a specific model from the **Model** filter, the **Language** filter value is updated to display the language of the selected model.
- If you select a specific language from the **Language** filter, the **Model** filter only displays models of that selected language.

Underneath the Language and Model filters, you can see the **Include translated conversations** switch. To include the performance of VA conversations that were translated to your primary language using dynamic language translation, toggle the switch to the right so it changes color from grey to green. When you do this, the bars in the charts legend may also change position and colors. You can manage dynamic language translation in the **Conversational Interfaces Settings**.

Setting the Date range

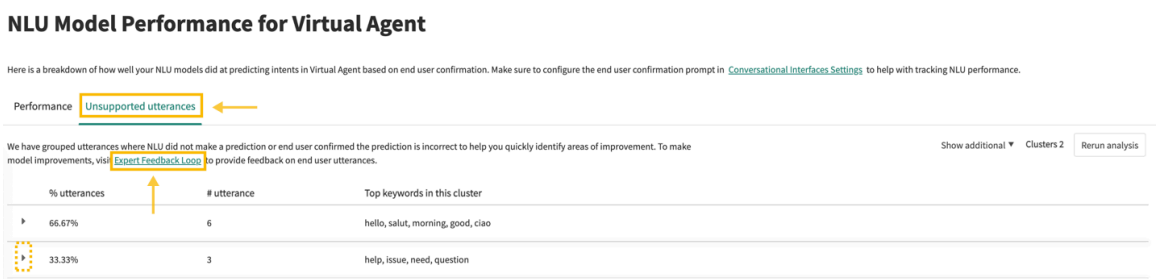
On the Performance tab, use the **Date range** to define how far back you want the system to pull VA data from. Choose **Past 30 days**, **Past 60 days**, **Past 90 days**, or **Custom range**. The further back you go in the date range, the more data you will have in your analysis.



Running an analysis

To run a performance analysis, click the **Unsupported utterances** tab. This section of the UI shows rows of expandable clusters containing VA utterances where NLU didn't make a topic prediction, or where the VA end-user confirmed that the predicted topic was incorrect. The next thing you want to do is to click **Expert Feedback Loop**. This action takes you to the NLU Expert Feedback Loop application where you review and provide feedback on the utterances that were pulled in from VA.

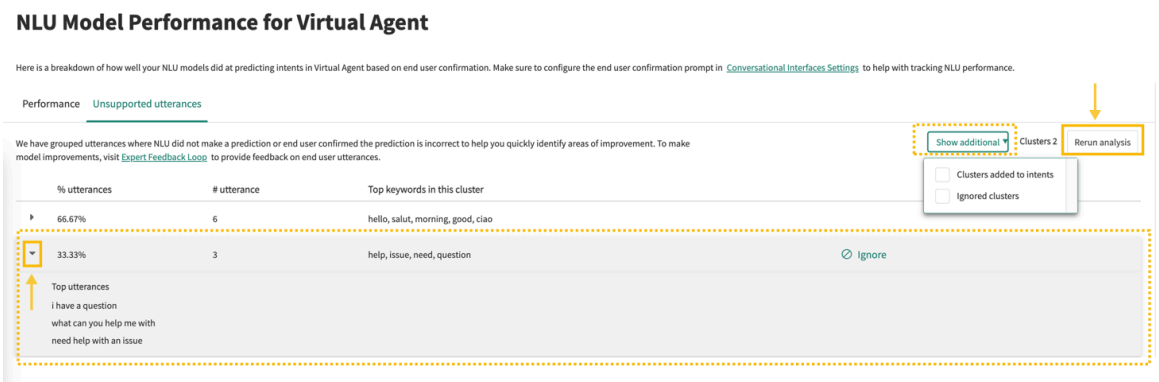
Reviewing unsupported utterances



When you leave the Expert Feedback Loop application and return to NLU Model Performance, click the **caret** icon in any cluster to open it. Within the cluster you can then see the top most representative VA utterances, as shown in the following image.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Reviewing the top representative utterances in a cluster



Click **Run analysis** or **Rerun analysis**, whichever is available. Each time you run an analysis, the system pulls the most recent VA utterances into the analysis.

Related topics

- Virtual Agent
- Natural Language Understanding in Virtual Agent

NLU Expert Feedback Loop

Provide feedback on Virtual Agent chat log utterances to help the system continuously learn and to better predict user input.

Embedded video: NLU Expert Feedback Loop, presented by the ServiceNow Virtual Agent Academy

Summary usage

The Expert Feedback Loop takes data from your instance and provides it to you for feedback. This data comes from your users' Virtual Agent (VA) chat logs and includes utterances from those logs. Using the nlu_admin role, navigate to **All** > **NLU Workbench** > **Expert Feedback Loop** and mark each utterance in a model by confirming whether the utterance is correct (match) or incorrect (mismatch) to an intent. This feedback helps the model to continuously improve the model performance.

Installation

Expert Feedback Loop is part of the NLU Workbench - Advanced Features app available on the ServiceNow® Store.

To use Expert Feedback Loop, ensure that the NLU Workbench - Advanced Features (sn_nlu_workbench) plugin is active on your instance. For more information, see Install NLU Workbench - Advanced Features and Activate the NLU Workbench.

Because Expert Feedback Loop relies on the data from VA chat logs, make sure that the Glide Virtual Agent (com.glide.cs.chatbot) plugin is also active. See Activate Virtual Agent.

Importing expert feedback loop data between instances

Before you start an import, make sure you have access to the data in your instance and have enough data to proceed with your feedback.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If you're working in a sub-production instance, you must import the feedback data from the **open_nlu_predict_intent_feedback** table in your production instance to your sub-production instance. For guidance on how to import, see Importing from another ServiceNow instance.

Here are a few data scenarios regarding the system behavior for the NLU models that house the feedback loop utterances.

- If you move a model to a different instance, the feedback data persists.
- If you upgrade the instance, the feedback data persists.
- If you clone an instance, the data does not persist, so you need to follow the procedure referenced above to import the data onto the cloned instance.

Feedback context and access

The Expert Feedback Loop provides a mechanism for you to improve NLU models deployed to VA through feedback that you provide on a select subset of utterances. For each utterance, you are asked to confirm the predicted intent or to provide the correct intent that the utterance belongs to. After this feedback is complete, the data is used to tune the model performance, resulting in an improved model that can now be redeployed to gather more end user data. This is an iterative cycle that allows the system to continuously collect data, learn from the feedback you provide, and use that feedback to further improve your NLU models.

Once every 30 days, the system pulls up to 300 utterance samples from VA chat logs to the Expert Feedback Loop. The utterances are selected for feedback based on how well they represent all the utterances in the logs. Every utterance sampled from VA chat logs has a predicted intent picked by the system.

You can also set the number of utterances you want to be pulled from VA chat logs by adjusting the setting for the *glide.mlpredictor.option.nlu.activeLearning.label_candidate_table.max_respons* system property.

Further information on the process of pulling utterances from VA chat logs can be found in the sections Procuring additional VA feedback data on demand and Using the Active Learning job, below.

As you review the utterances, decide whether each utterance belongs to its predicted intent or should be moved to a different intent. If you aren't sure about the correct intent, the utterance can be marked to revisit for further review. After you have marked at least 100 utterances with your feedback, the system uses all the marked utterances to tune and improve the model.

Reviewing your VA chat log utterances

Now that you're in the loop, you can see in the following image a list of NLU utterances that support the intents in the VA models. Each page of the list shows approximately 20 utterances per page. Your goal is to review each utterance and mark its relationship status to a given VA intent. You mark each utterance with either the **NLU_Match**, **Mismatch**, or **Unsure** value.

Marking an utterance as **NLU_Match** means that you agree with the NLU prediction for that utterance. Marking an utterance as **Mismatch** means that the utterance belongs to a different intent, and if you select **Mismatch**, a dropdown of intents appears so you can select the correct intent. If you aren't sure which intent is correct, you can mark it as **Unsure**.

Utterances that are marked and saved with the **NLU_Match** value or the **Mismatch** value are moved to the **Completed reviews** section. Utterances that are marked and saved with the **Unsure** value are moved to the **Needs further review** section.

Here's a scenario for the basic steps you use to complete your utterance reviews:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

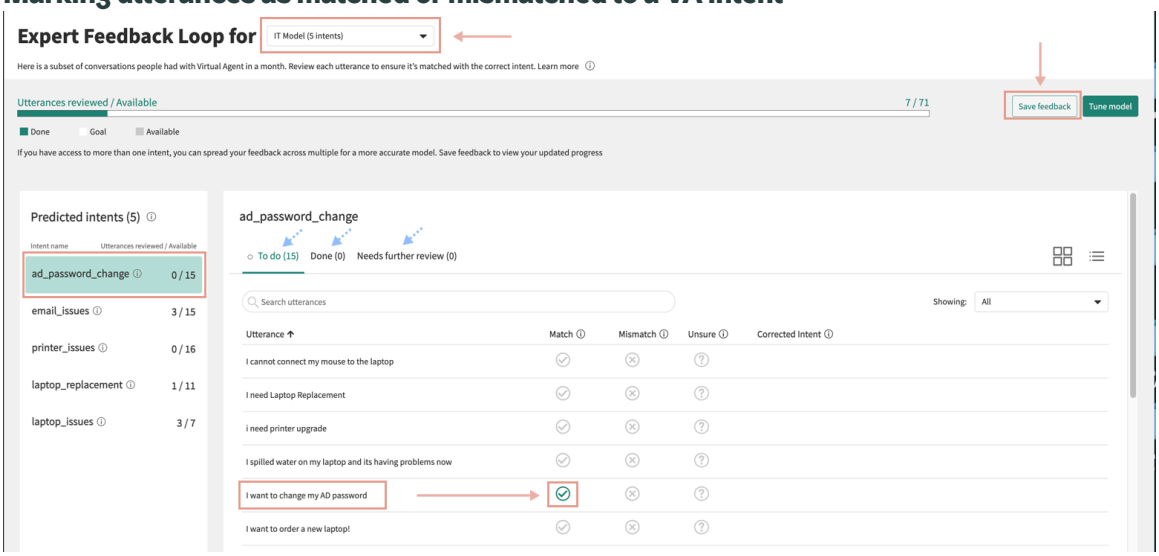
1. At the top of the Export Feedback Loop screen, select a model you want to review in the **Expert Feedback Loop for** prompt. Within this feature, in some ways it doesn't matter which model you choose because these models serve mainly to house the utterances you are marking in your reviews. So in this example scenario, you decide to select the **IT Model** model, which has five intents and many utterances within each of those intents. The number of utterances you need to review are shown next to the name of the predicted intent, as shown in the following image. The intents in these models are listed in the **Predicted intents** column on your screen. Although you can review and mark utterances in any of the predicted intents, it's a better idea to complete the reviews consecutively in the order the intents first appear. For example, begin with the **ad_password_change** intent and then onward to other predicted intents.
2. Click the **ad_password_change** intent so it loads its utterances into the **To do** section for your review.
3. Review the 15 utterances in the intent. In this scenario, you've started by correctly marking the I want to change my AD password utterance as a match to the **ad_password_change** intent.
4. Click **Save feedback**.

   Result: The system moves the utterance you marked over to the **Done** section of the screen. At the same time, the count of **To do** utterances drops from 15 to 14, while the count on the **Done** section rises from 0 to 1. If you had also marked another utterance with the **Unsure** value and saved that change, the count of utterances in the **Needs further review** section would also rise from 0 to 1.

   Note: If you don't know which intent best matches the utterance, mark it with the **Unsure** value. This action moves the utterance over to the **Needs further review** section, giving you time to mark other utterances in the **To do** section that may be easier to match or mismatch to an intent. You can always return to the **Needs further review** section to mark utterances that are more difficult to pin down to a specific intent.

5. Repeat steps 1 through 4 as you move through the remaining intents in the **Predicted intents** column.

Marking utterances as matched or mismatched to a VA intent



The utterances in the **ad_password_change** intent are easy to mark as a match to its predicted intent because there's at least one utterance within the intent that has the word password in it.

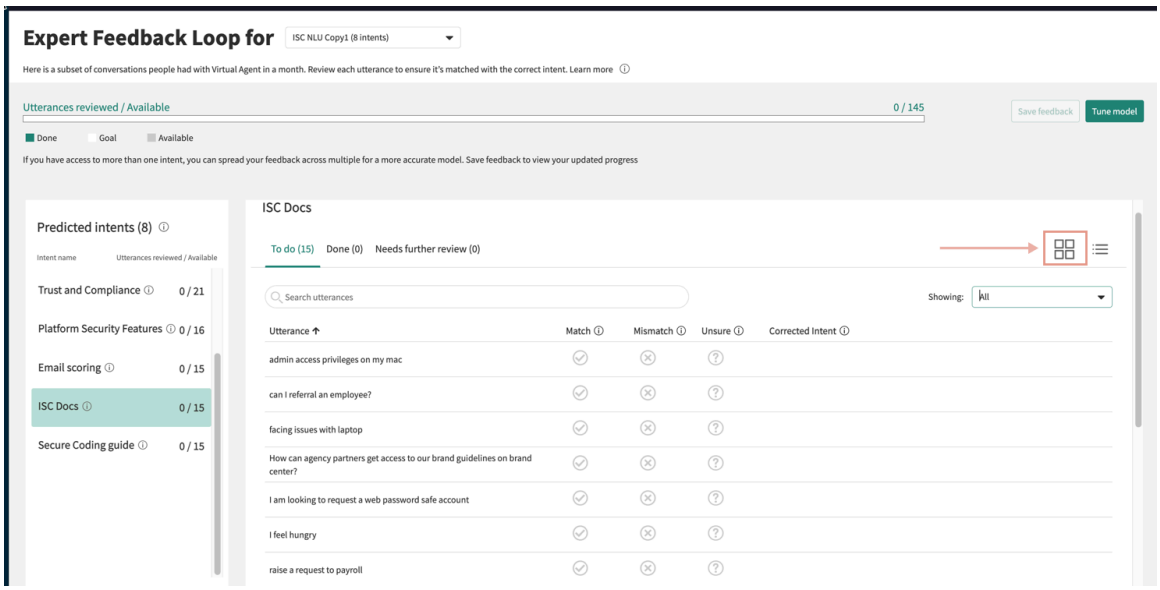
When you have finished reviewing the utterances in the predicted intent and you click **Save feedback**, the screen refreshes to highlight the next predicted intent in the model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

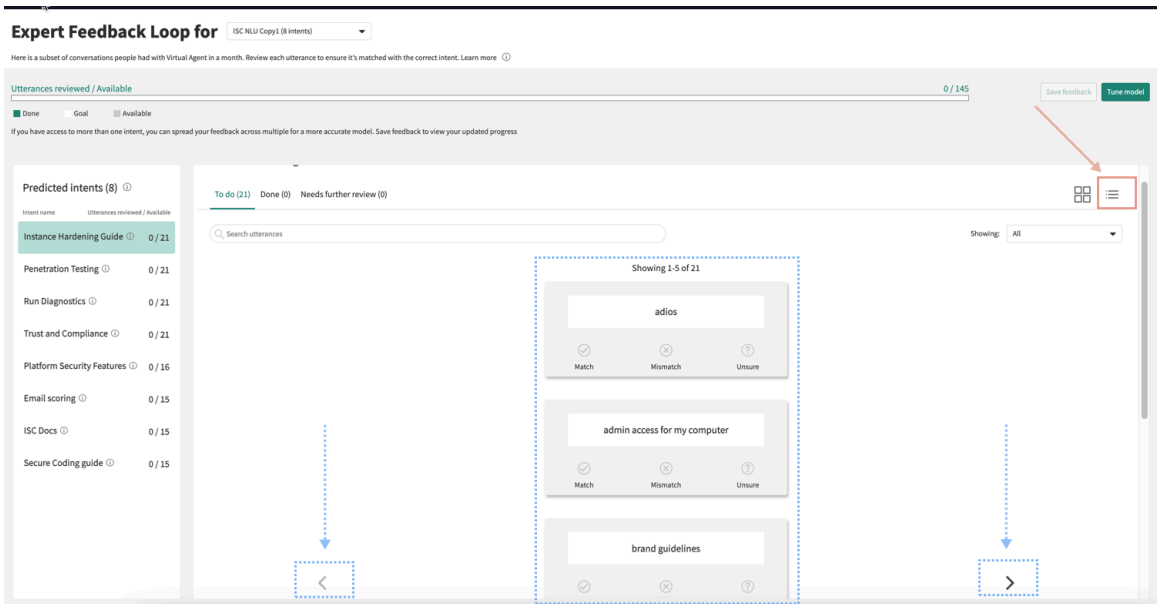
If you want to see your utterances grouped by their value, click the **Showing** filter and select one of the options in its prompt. The values are **All**, **NLU_Match**, and **Unsure**.

Using the card view option to review your utterances

Your Expert Feedback Loop utterances appear on the **To do** section in a list view by default, as shown in many of the images you've seen so far above. However, you can also choose a different view that shows each utterance framed within a card. If you choose this option, each utterance appears on the page in groups of five. While still within the list view, click the **card** icon.



Result: The card view replaces the list view. Note the left and right arrows on either side of the page of utterance cards. When you click the right arrow, it turns the page to the next set of five cards. When you click the left arrow, it turns the page to the previous set of five cards. If you want to switch back to the list view instead of the card view, click the **list** icon as shown in the upper right hand corner of the image below.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Changing an utterance value

If you mark an utterance but then change your mind on the value you chose, you can resolve the issue by clicking on **MisMatch** and selecting another intent. Click **Save feedback** to save the change.

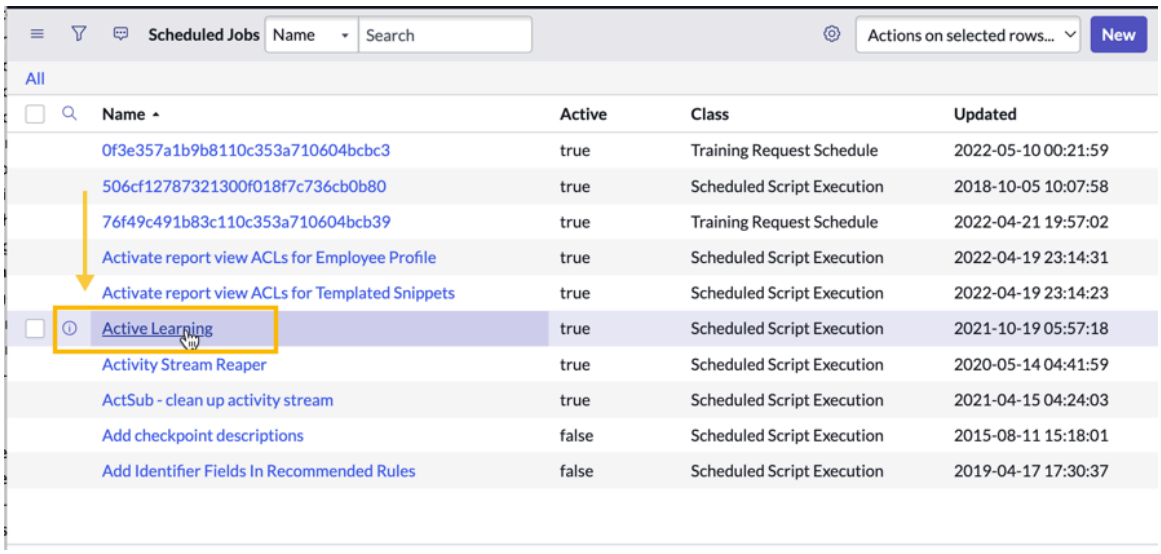
Unsaved feedback

If you finish and log out of a session of marking your utterances, but forget to save your changes, select **Unsaved changes** from the **Showing** prompt. This action displays all the utterances you have given feedback for but haven't saved yet.

Procuring additional VA feedback data on demand

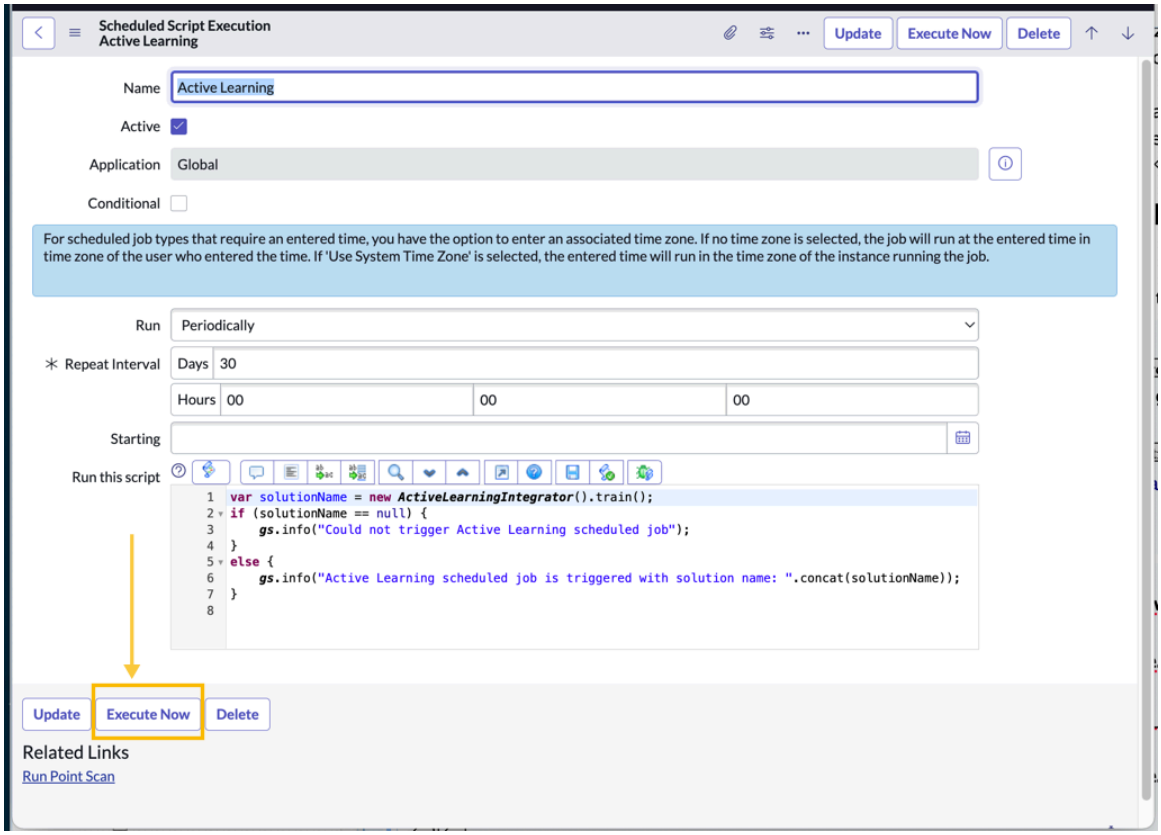
To get more data from the VA chat logs to use in the Expert Feedback Loop, take the following steps.

1. Use the nlu_admin role and navigate to **All** > **System Definition** > **Scheduled Jobs** > **Active Learning**.
2. Click **Active Learning**.



3. Click **Execute Now**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. Increase or set the values in the following four NLU system properties.

- *glide.mlpredictor.option.nlu.activeLearning.va_chat_logs.max_row_limit - 3000*
- *glide.mlpredictor.option.nlu.activeLearning.label_candidate_table.max_data_ - 10000*
- *glide.mlpredictor.option.nlu.activeLearning.label_table.max_data_size - 10000*
- *glide.mlpredictor.option.nlu.activeLearning.label_candidate_table.max_respo*

To see how these properties work within the context of other NLU properties, see NLU Workbench properties.

Reviewing uncategorized utterances

As part of the feed from VA chat logs to the Expert Feedback Loop, the system collects and displays in a list any utterances in your instance that aren't part of a VA intent. You access these utterances by selecting **Uncategorized** in the **Expert Feedback Loop for** prompt at the top of your screen.

When the screen refreshes, these utterances are shown in the **Utterance** column of the **Uncategorized utterances** section of your screen.

Important: It is extremely important to provide feedback on this set of utterances because the system is suggesting there is no associated intent for these utterances. By confirming this lack of association or by associating these utterances with an existing intent, you are helping the model to learn correctly.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

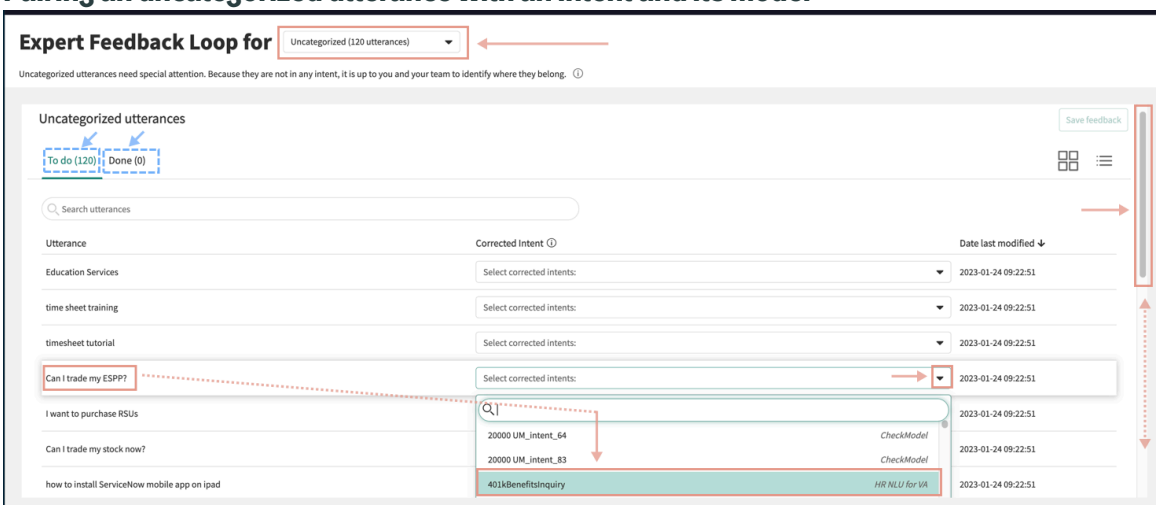
Here are the steps you use to connect an uncategorized utterance to an intent:

1. In the **Utterance** column, select an utterance from the list.
2. In the **Corrected Intent** column, search and select an intent and its model from the prompt that you think is the best match for the utterance.

For example, a good pairing for this connection is the Can I trade my ESPP? utterance from the **Utterance** column and the 401kBenefitsInquiry intent from the **Corrected Intent** column, as shown in the image below.

Note: The **Corrected Intent** column enables you to search for and use all intents across all models in your instance.

Pairing an uncategorized utterance with an intent and its model



Regarding the image above, let's cover a few items that help describe how the uncategorized utterances page works, as follows. The **To do** section collects utterances that must be connected to an intent. On this particular day we see there are 120 such utterances. Note that the **Done** section on the screen has no utterances because you've not taken any action yet.

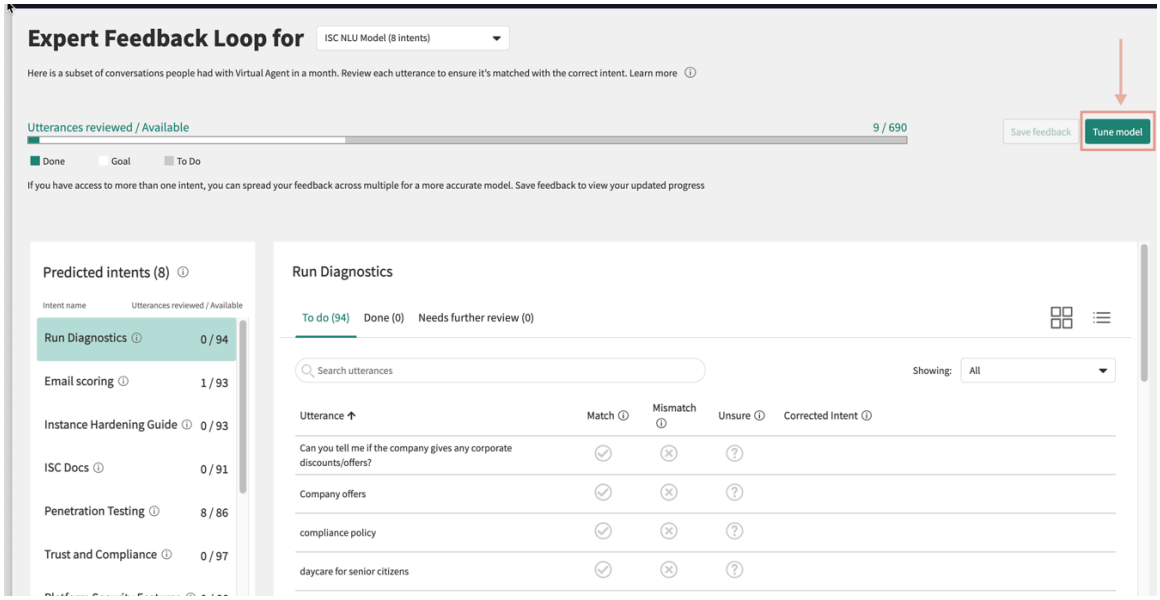
Note also that there's a vertical bar to the far right of the screen where you can scroll up and down to navigate the list of utterances. At the bottom of this scrolling screen, there are arrows that point left and right so you can to turn to the next or previous page of the full utterance list for the day.

Note also that when you do take action to make a connection, the **Save feedback** button, normally active by default, deactivates because the system recognizes your action and automatically saves your change.

Expert Feedback Loop data in the Tune Model phase

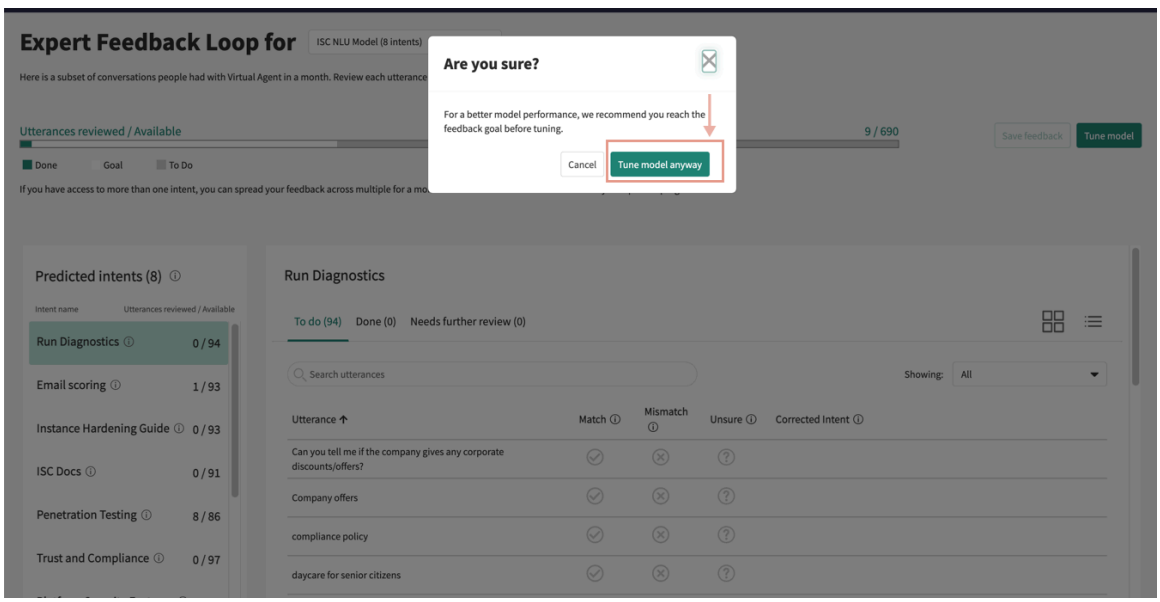
The **Tune model** button in the Expert Feedback Loop is always enabled and supports the experience of moving utterances from the Expert Feedback Loop **Done** tab to the model and its test set. Use the nlu_admin role to click this button any time you need to tune or retune your model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



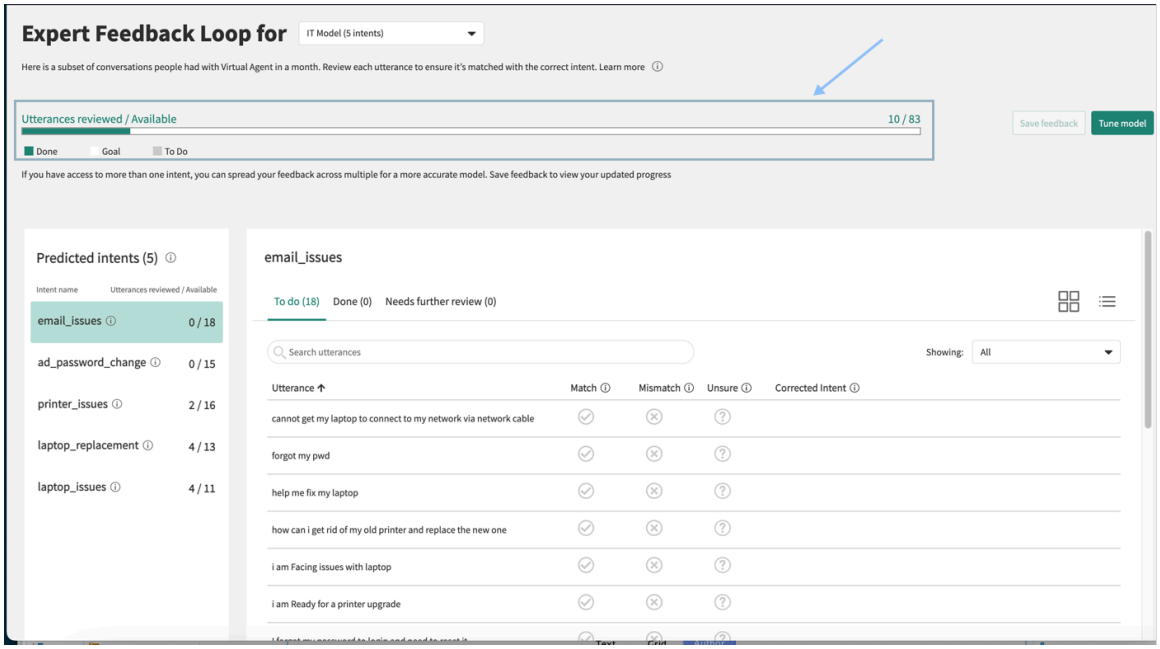
After reviewing utterances in the Expert Feedback Loop, you can push a portion of the feedback data to the default test set of your model. These utterances are then directly added to the **Test utterances** tab of your model. This helps you to continuously maintain and update your test set with real end user utterances. The system tracks the source of the test utterances for visibility into whether they came from the Expert Feedback Loop or from another source.

If you click the **Tune model** button before you have marked and saved at least 100 utterances, the screen refreshes to remind you that your model performance increases in quality when you have passed the 100 utterance feedback goal, as referenced in the text in the image below.



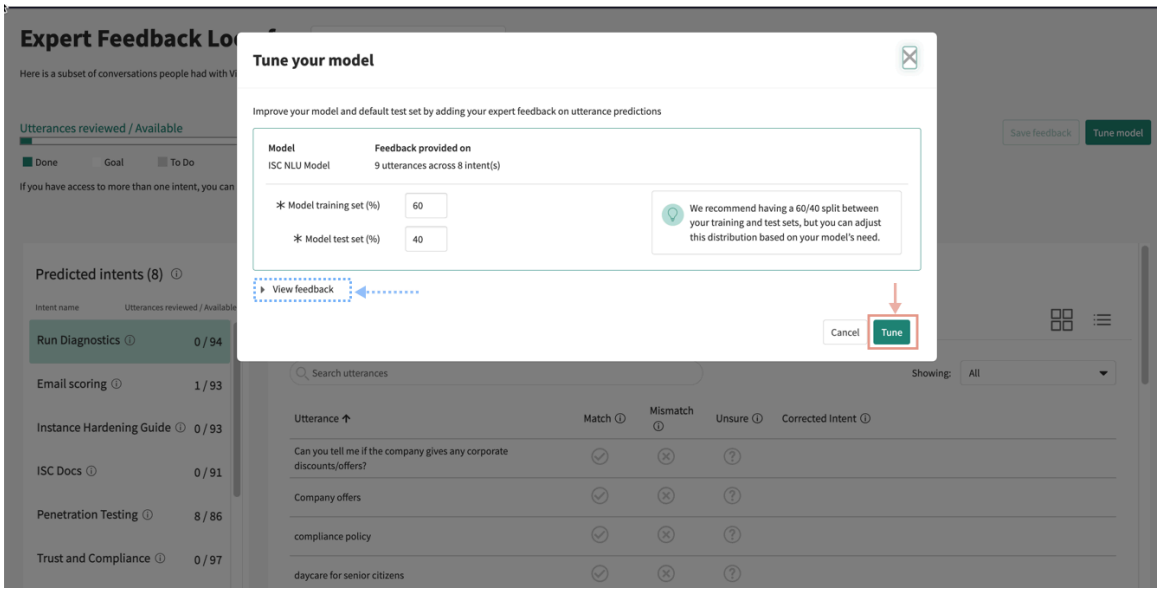
As you keep marking and saving your feedback data, you can see the ratio of reviewed utterances (in green), and those that are still available for review (in white) on the progress bar image below.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



On the **Tune your model** image below, you can see the default 60/40 percentage split for your model training set and your model test set. You can adjust these default values if needed by selecting your own chosen numbers in each of the two **(%)** boxes. Once you're satisfied with your percentage split, click the **Tune** button to tune your model based on the percentage values you chose.

Note: If you click the **View feedback** caret, it takes you to an Expert Feedback Loop screen where you can continue reviewing, marking, and saving your feedback utterances.



Using the Active Learning job

Scheduled jobs, also known as batch jobs or batch scheduling, are automated pieces of work that are performed at a specific time or on a recurring schedule. Many jobs run on schedules but we're focusing now on the Active Learning job as an example.

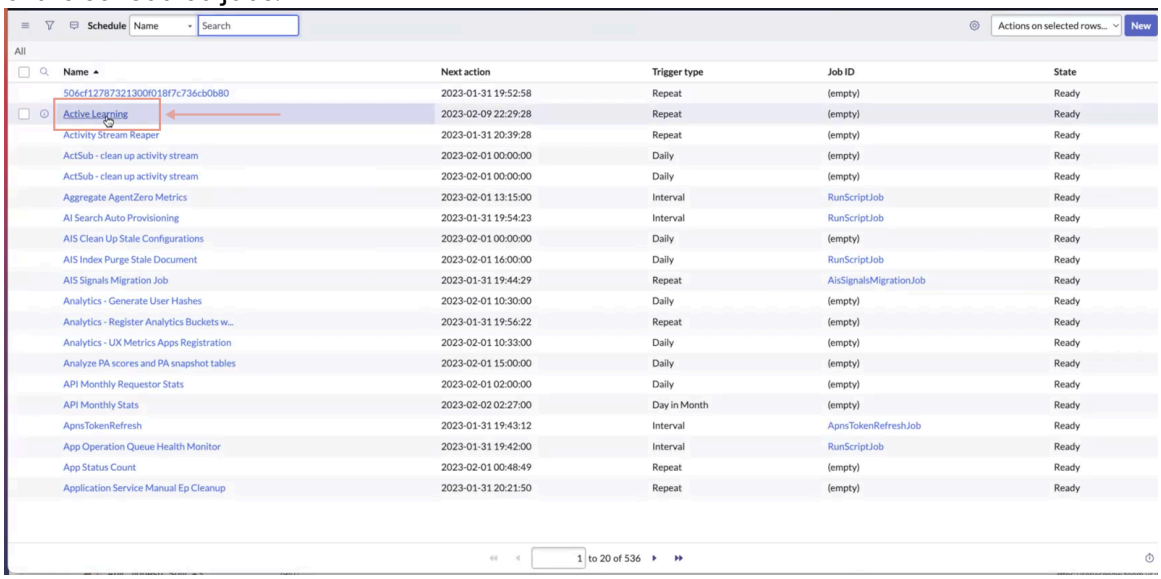
Here are three things you can do when using the Active Learning job.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

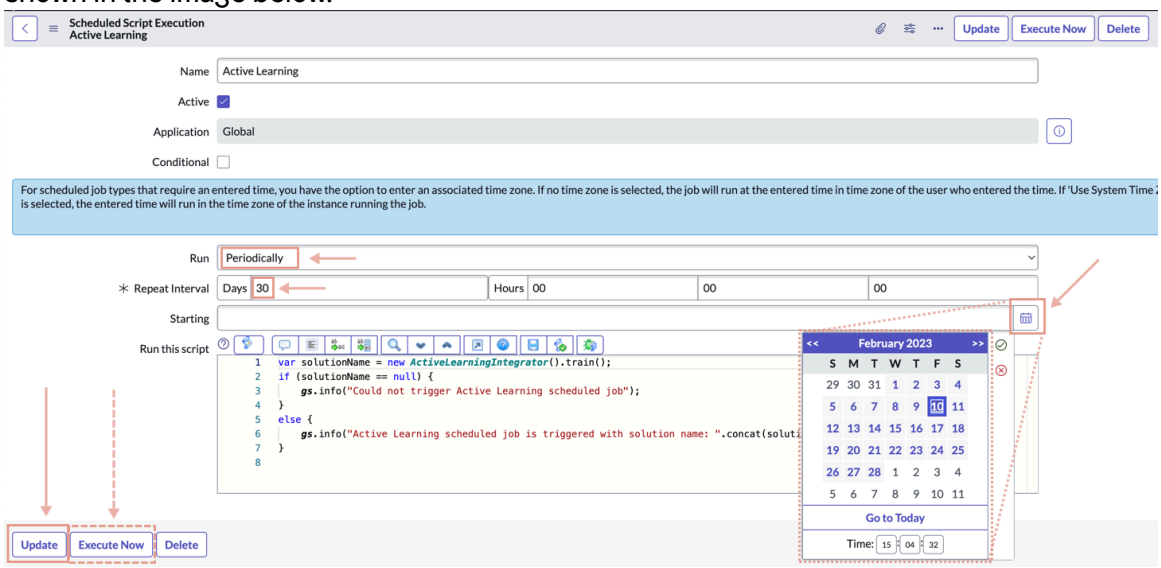
- Change the repeat frequency interval with which the Active Learning job will be set to run.
- Check when the next scheduled run is set for the Active Learning job.
- Execute the Active Learning job whenever you want (on demand).

Here are some steps to use to get started.

1. Using the nlu_admin role, navigate to the **All** field and type **sysauto_script.list**, then press the **return** key on your keyboard. The screen refreshes to show the Schedule page, which lists all of the scheduled jobs.



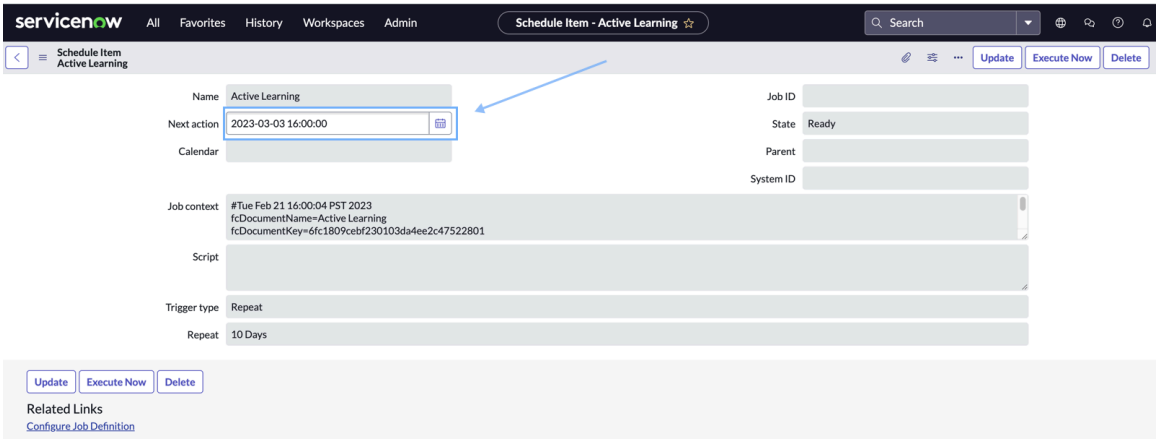
2. On the Schedule page, click **Active Learning**. A record for the Active Learning job appears, as shown in the image below.



3. On the Active Learning record, configure the following fields.
   - In the **Run** field, select one of the eight Repeat Interval options for the job. The options are: Daily, Weekly, Monthly, Periodically, One, On Demand, Business Calendar:Entry Start, and Business Calendar:Entry End.
   - In the **Repeat Interval** field, enter the number of days you want between now and the next repeat interval for the Active Learning job.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- In the **Starting** field, click the Calendar icon and choose the day and time for the next Active Learning job interval.
- Click the **Update** button to save your configurations or click the **Execute Now** button to start the Active Learning job run.
- To verify when the next Active Learning job runs, navigate to the **All** field and type **sys_trigger.list**, then press the **return** key on your keyboard. The Schedule page appears. Click **Active Learning**. The Schedule Item/Active Learning record appears and populates the **Next action** field with the date and time for the next Active Learning job run.



Detailed information about NLU's Active Learning job can be found in this KB article on the Support portal: KB1633901.

Issue Auto Resolution Tuning in NLU

Use the NLU Workbench homepage to support Issue Auto Resolution (IAR) tuning in NLU.

A video introducing Issue Auto Resolution tuning in the NLU Workbench

Summary usage and roles

Use the nlu_admin or admin role to access IAR Tuning in the NLU Workbench. IAR Tuning in the NLU Workbench requires at least the nlu_feedback_admin role. Note that the nlu_admin role contains the nlu_feedback_admin role. Also, the virtual_agent_admin role contains the nlu_admin role.

If you click on the IAR model name, you will be taken straight to tuning, in the product. You will not be taken to a model overview page, so this behavior differs from Virtual Agent or AI Search models in the NLU Workbench.

The IAR Tuning workflow

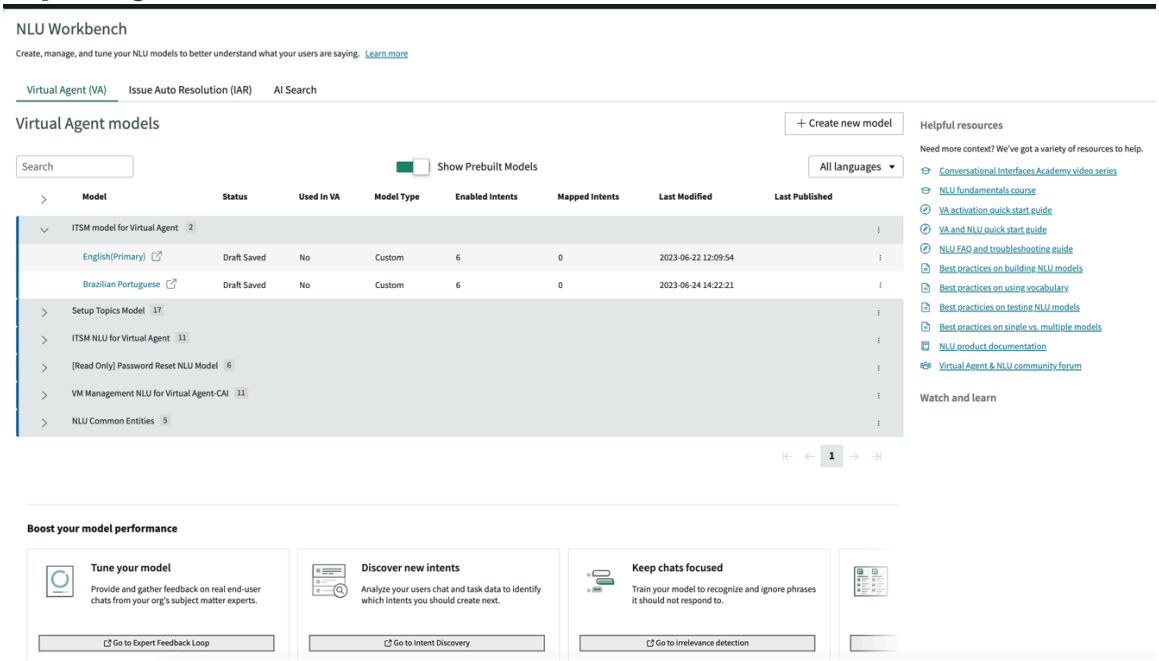
IAR admins begin their model tuning journey in the IAR Admin Console, and then land in the NLU Workbench to tune their ITSM model. If they haven't trained the ITSM model in the console yet, the workflow sends them to the Expert Feedback Loop documentation under the **Boost your model performance** section of the NLU Workbench.

How IAR models differ from NLU models

Unlike the Virtual Agent and AI Search tabs, the IAR tab doesn't use a Create new model button. The IAR-ITSM model that IAR admins use is a prebuilt model. IAR models can't be moved using update sets.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Exploring the NLU Workbench



At the top of the NLU Workbench page are three tabs that group Virtual Agent, Issue Auto Resolution, and AI Search models separately. Below those tabs are a list of models colored grey. In the Model column of the list, if you click the caret to the left of the model name, the model changes color from grey to white and opens to show the model's languages; status; usage; model type; number of enabled intents and mapped intents and the date when the model was last modified or last published.

Issue Auto Resolution tuning options

When you are tuning your Issue Auto Resolution model in NLU Workbench, you can adjust the output for several goals: precision, automation, or a balance of the two. Compare how your choice of tuning options affect match rate and coverage, before committing.

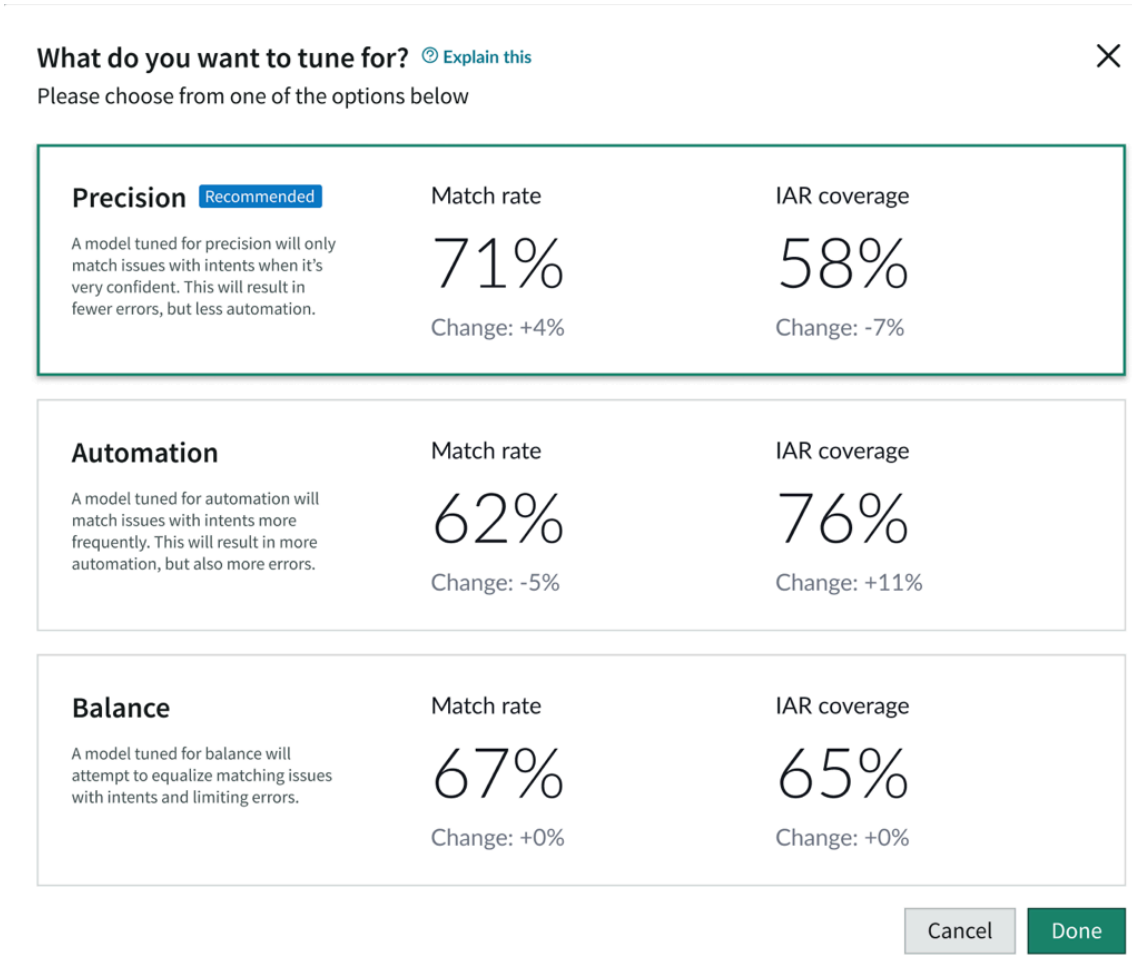
Summary usage

By default, Issue Auto Resolution tuning in NLU Workbench optimizes for precision. Depending on your business requirements, you can also tune the model for other objectives. In the **Analyze** step of Issue Auto Resolution tuning, the tuning goals list enables you to adjust for **Precision**, **Automation**, or **Balance**. As you select one of these options, the projected Match rate and IAR coverage percentages change accordingly, so you can compare possible outcomes.

To access the **Analyze** step of IAR tuning, use the nlu_admin role and navigate as follows.

1. Navigate to **All** > **NLU Workbench** > **Models**.
2. Select the Issue Auto Resolution tab, then select the model name. The tuning experience opens to step 1 (Feedback) initially.
3. Provide feedback, then select the **Analyze** button. Step 2 (Analyze) opens.
4. In the section **Here are your tuning options and projected results**, using the list **You can tune for precision, automation, or balance**, select options to see projected scenarios. You can also select the link **Learn about tuning goals** to open the following window.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Precision

When tuned for precision, the IAR model makes predictions only when its confidence is relatively high. This results in lower error rates, but also in fewer incidents resolved.

Precision is the recommended tuning option for the IAR ITSM model, so this option is selected by default.

Automation

When tuned for automation, the IAR model makes predictions at a lower confidence threshold. This results in more predictions, so more incidents are resolved. However, higher error rates are possible.

Balance

When tuned for balance, the IAR model attempts to strike a balance between precision and automation.

Match rate

The match rate is defined as the number of Incidents where the intent was predicted correctly, divided by the number of predictions for that intent. This ratio is averaged across all intents except for **NO_INTENT**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IAR Coverage

Coverage is defined as the percentage of Incidents that would be resolved because the model was able to make predictions above its confidence threshold. The predictions may contain some errors.

Using tuning options

Select several different tuning options to compare projected results. Depending on the option you select, the system presents scenarios for projected Match rates and IAR coverage rates. Also the system displays how much these rates change according to your selection.

Review further information in the **Here's a detailed breakdown** section of Analyze. Here you can drill down into results that are specific for each intent in the model.

Note that the intents are grouped into mapped and unmapped intents, depending on whether they have been mapped to Virtual Agent topics. After providing feedback in IAR Tuning, you may wish to activate some intent-to-topic mappings. To do so, expand **See unmapped intents**, then select the **Map more intents** button. This opens the IAR Admin Console.

When you have decided the optimum tuning option for your requirements, select the **Save choice** button in the **Learn about tuning goals** window. Then, select the **Tune and publish model** button to advance to the next step.

Intent Discovery

Use the Intent Discovery application to help identify opportunities for incident deflection. For example, you can use it to identify which Virtual Agent conversations to activate next.

Summary usage

For applications that consume NLU, such as Virtual Agent and AI Search, Intent Discovery helps you to better understand which prebuilt intents you can benefit from, and which custom intents would be useful to create.

Intent Discovery provides an analysis that you run on historic incident data or other task data. You can also group the run's remaining records into different clusters so you can manually add utterances to NLU intents. In addition, you can use specific clusters to create new intents in a model.

In this example scenario, you're using Intent Discovery to identify the top intents in your instance, and how much coverage they can provide across your historic incident records.

Installation

Intent Discovery is available from the ServiceNow Store. For more information, see Install Intent Discovery.

After Intent Discovery is installed and activated, it appears under **All** > **NLU Workbench** > **NLU Advanced Features**.

Note: Although organized under NLU Advanced Features in the navigation pane, Intent Discovery is a separate application that is not included when installing NLU Workbench - Advanced Features.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Intent Discovery report details

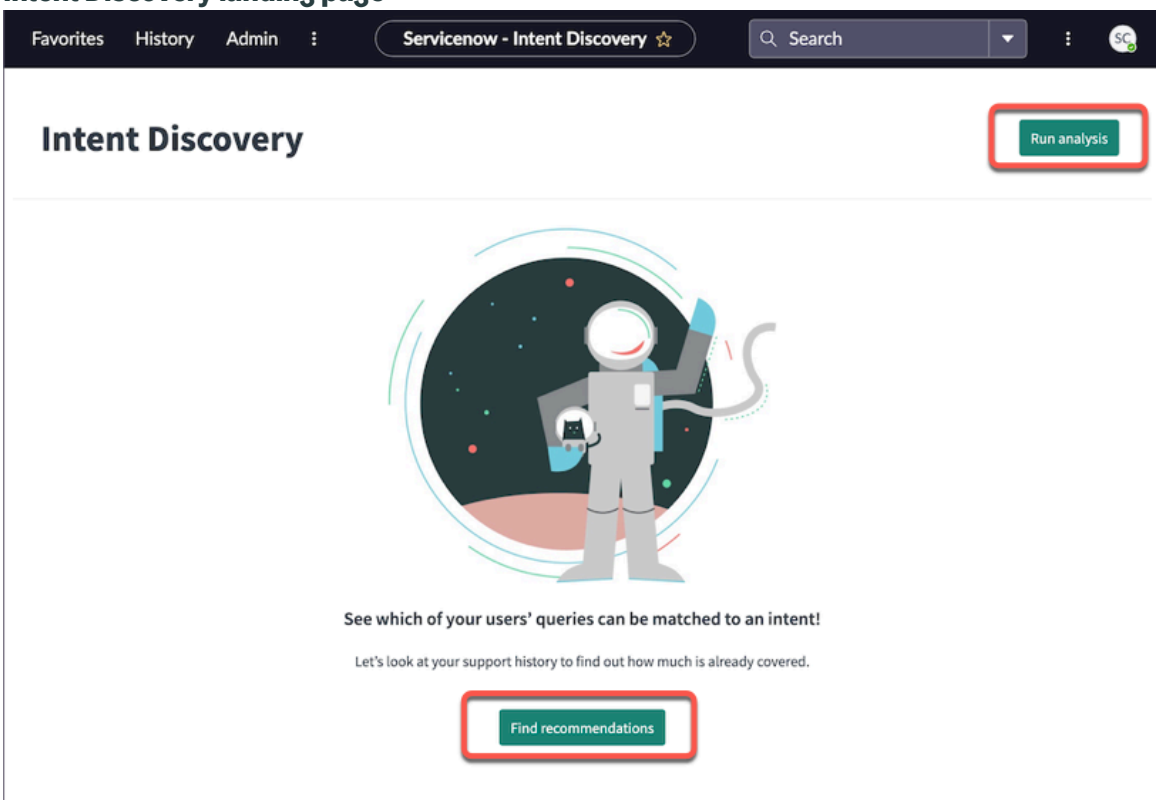
- When **Taxonomy** is selected, the generated report contains intent recommendations against the selected taxonomy. A taxonomy is a prebuilt library of intents in a specific domain. While you don't have access to the underlying intents, when you run Intent Discovery against a specific taxonomy, data that maps to any intent in the taxonomy will be identified.
- **Unmatched records** are the utterances which couldn't match to any intent in the taxonomy.
- **Recommended intents** are the intents which are found from utterances that data was run on.
- The percentage of **Unmatched records (clustered)** are the records that aren't classified (records that don't belong to any of the recommended intents).
- The percentage of unmatched records and the number of recommended intents don't need to match. It's a coincidence if they match.

Creating an Intent Discovery report

1. Using the admin or nlu_admin role, navigate to **All** > **NLU Workbench** > **NLU Advanced Features** > **Intent Discovery**.

2. Select either **Run analysis** or **Find recommendations**.

Intent Discovery landing page



Running an analysis on the report

1. For this example report, you configure the following fields on the Intent Discovery > Create new screen.

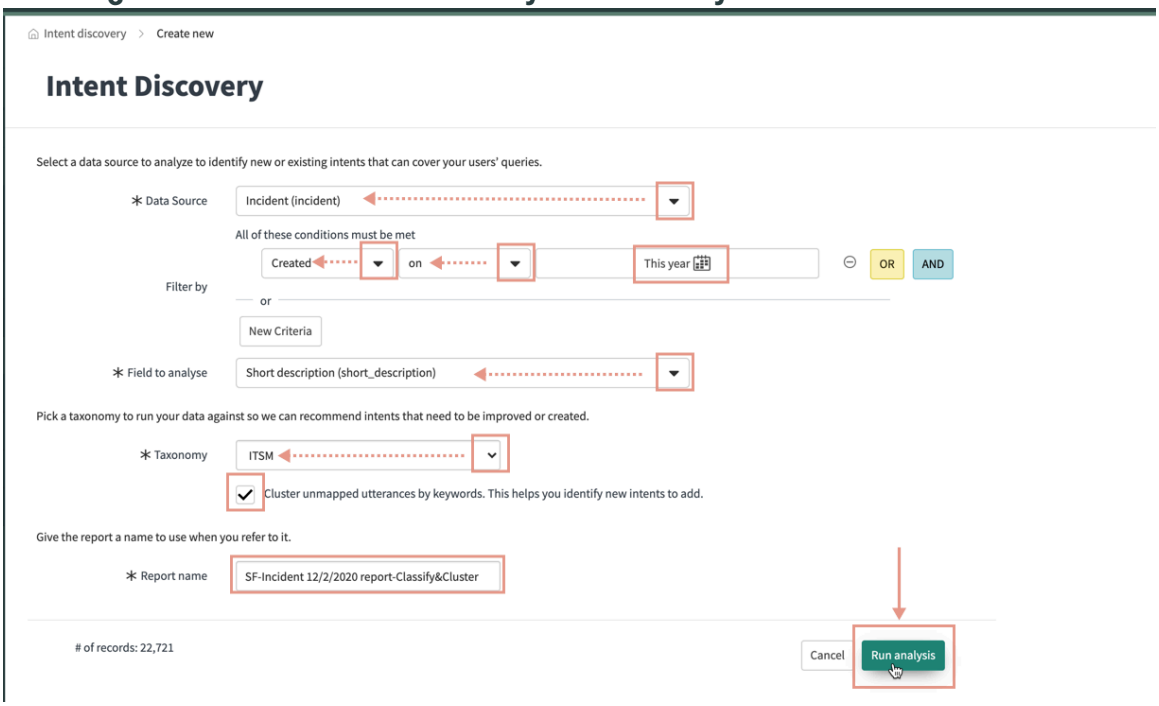
- Data Source: Select the **Incident (incident)** table.
- Filter by: **[Created] [on] [This quarter]**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Field to analyze: **Short description (short_description)**. You choose Short description because it's a highly used string field that references words that can help the system identify an intent.
- Taxonomy: Select **ITSM**. This field tells the system to run classification processing on your ITSM incident records. It has 3 options: Classification, ITSM, or blank, which defaults as Classification.
- Cluster unmapped utterances by keywords... : Select the **check box**. When you check this box, the system groups your incident records that weren't classified into clusters.
- Report name: The field automatically defaults to **Incident <month/day/year>**. You can edit the name if you prefer. In this example scenario, you enter `Incident 12/16/2020 - SF Test`.

2. Select **Run analysis**.

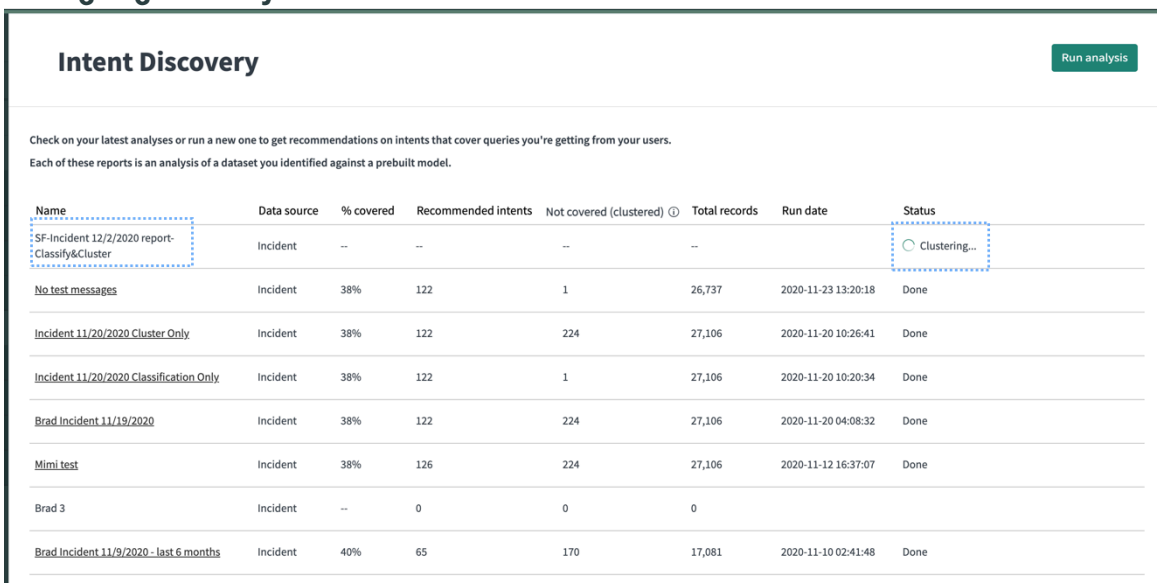
Selecting data sources in Intent Discovery for a run analysis



Result: Your report appears on the Intent Discovery screen, showing its status as the analysis begins. The subsequent status values appear in the following order during the analysis: Preparing to run, Work in progress, Clustering, and Done. This can take from 5 minutes to 30 minutes to complete. The fewer the records you have in a cluster, the less time it takes. Turning clustering off can also speed up the process.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

An ongoing run analysis

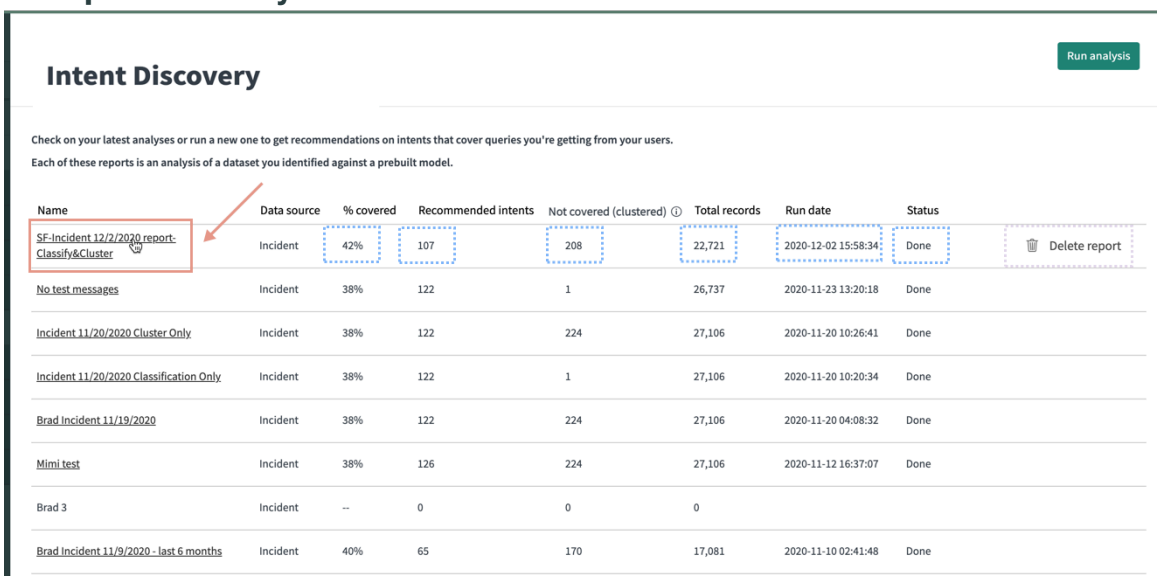


When the analysis is complete, the column values on the screen appear, with the **Status** column value set to **Done**, as shown in the image below.

Note: If you want to delete the report and start over, point to the right of the Status column to invoke the **Delete report** icon.

3. Select the **Name** of your report.

A completed run analysis



Result: The screen refreshes, showing the analyzed incident records and the remaining incident records that were not classified.

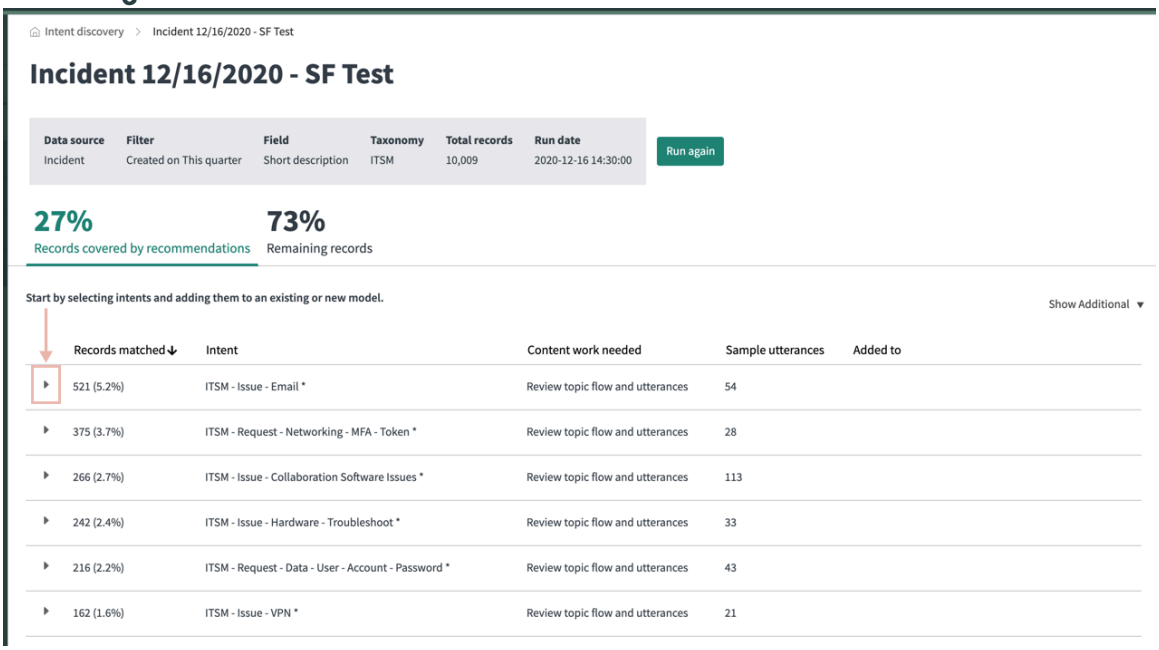
Importing recommended intents to new or existing custom models

Before importing intents to an NLU model, ensure that you are in the same application scope as the model. For more information, see Select an application from the application picker.

1. On the Records covered by recommendations section of the screen, select the caret icon on a recommended intent you want to add to a custom model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

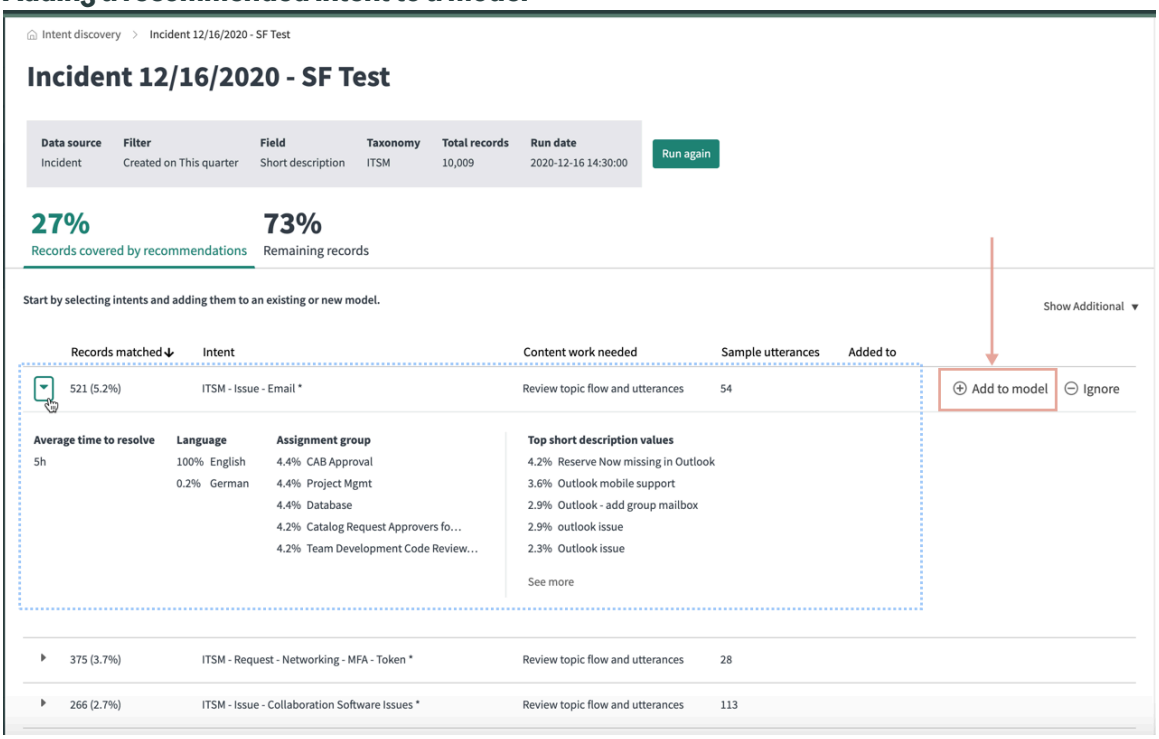
Reviewing a recommended intent



Result: The details of the recommended intent appear so you can review them, as shown in the image below.

2. Select **Add to Model**.

Adding a recommended intent to a model



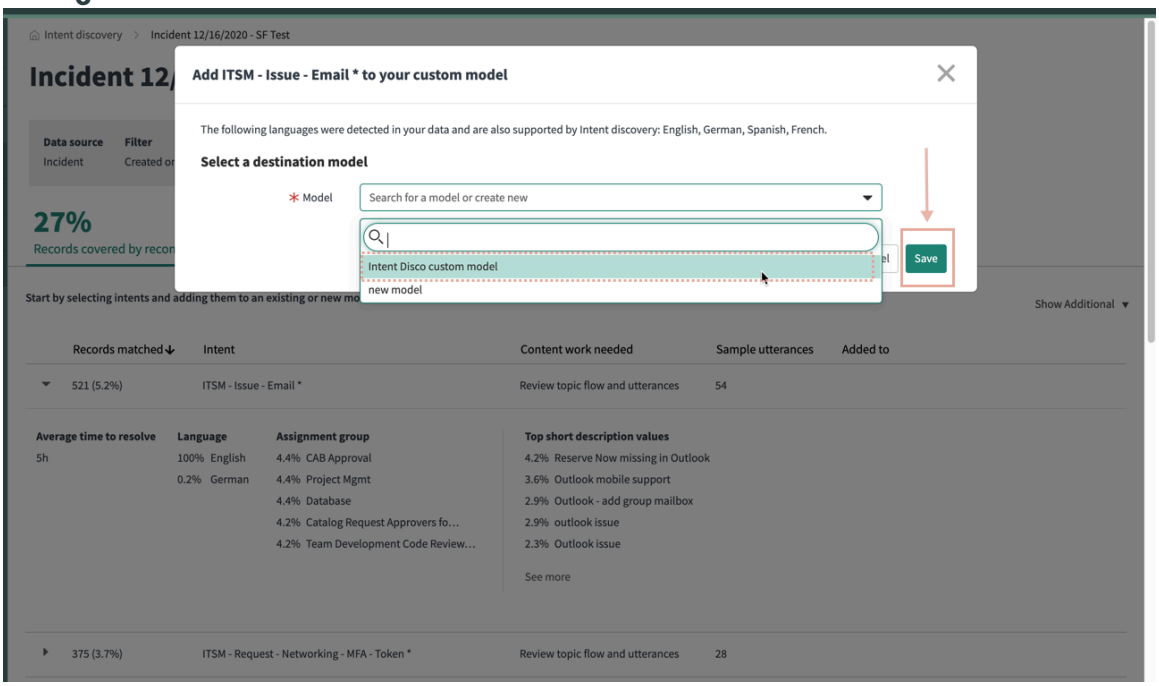
3. On the Select a destination model screen that appears, choose a model you want to add the recommended intent to. If you can't find an appropriate model, create a new one, return to the report, and add the new model.

Note: The model you choose must have the same application scope as your current scope.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

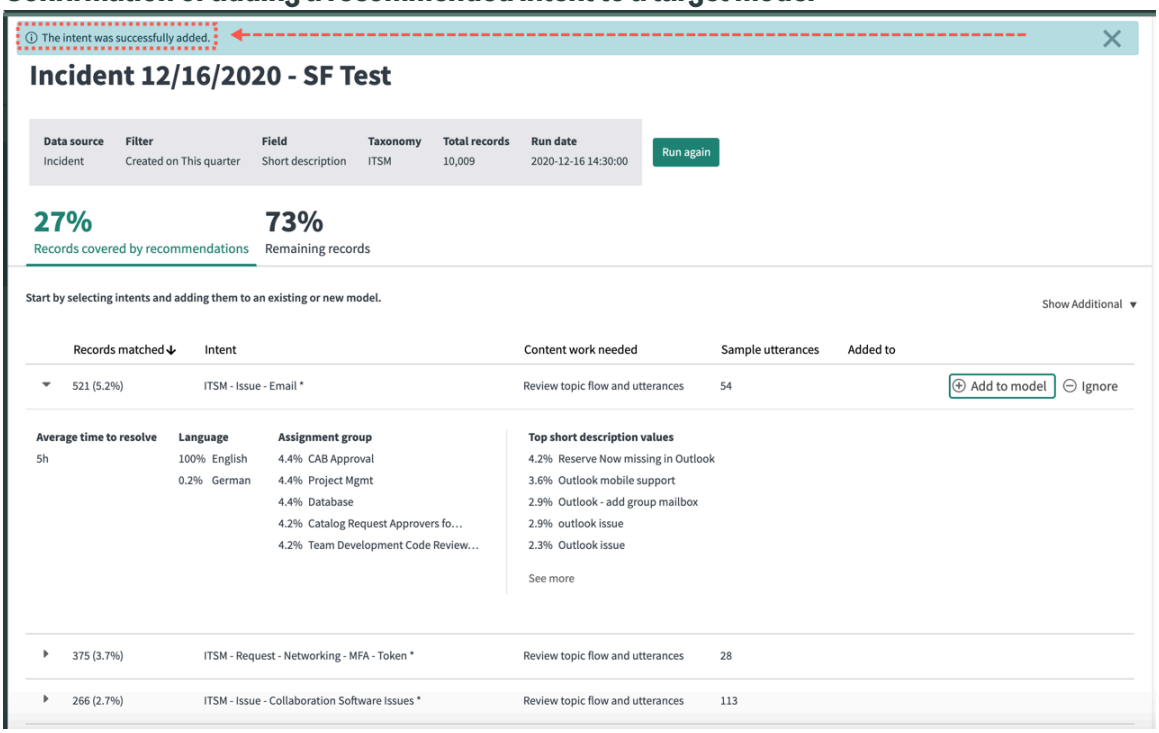
4. Select **Save**.

Saving a recommended intent to a model



Result: A banner appears on the screen, confirming the intent is added to the target model.

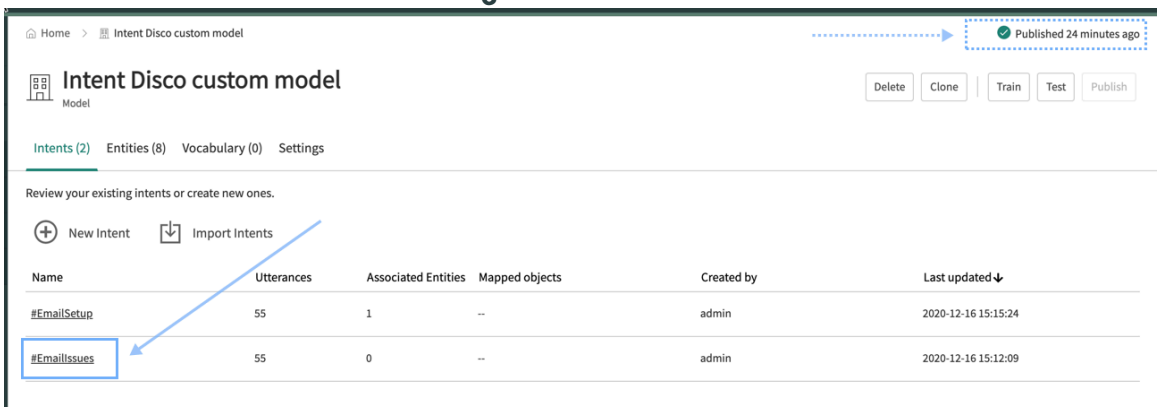
Confirmation of adding a recommended intent to a target model



The recommended intent also appears on the Model screen of the target model, as shown in the image below.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View a recommended intent in the target model



Adding clustered utterances to an intent and its model

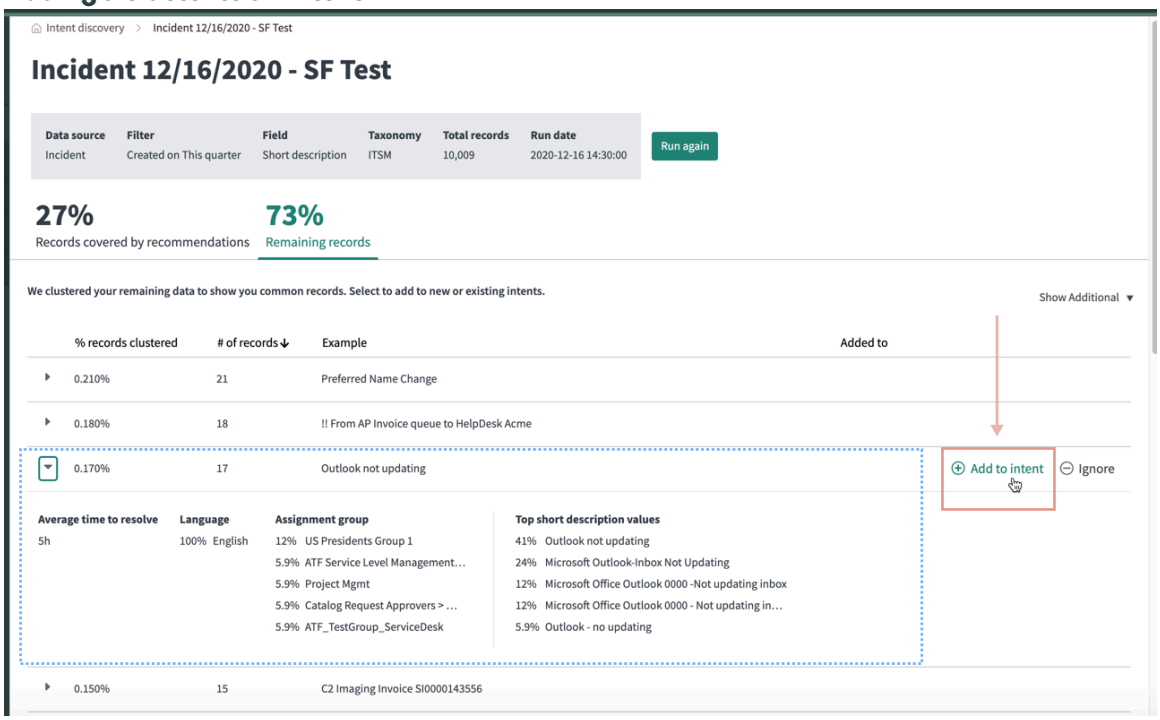
1. On the Remaining records section of the intent discovery records screen, select and open a cluster of utterance and short description data that you want to add to an intent and its associated model.

As you continue to build out new intents from these clusters, you can click the **Ignore** icon to remove any unwanted intents from the report.

There's also a **Show Additional** filter you can use to show or hide the added intents, and the ignored intents as well.

2. Select **Add to intent**.

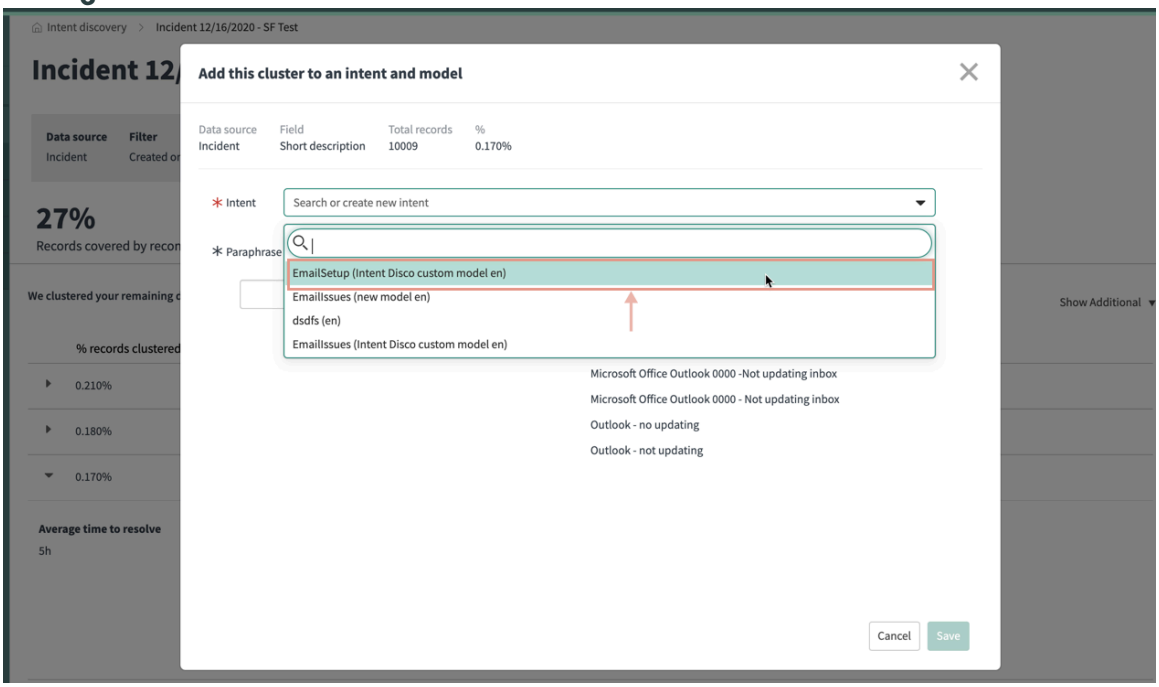
Adding a cluster to an intent



3. In the Add this cluster to an intent and model screen, select an intent and model pair you want to associate to this cluster.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

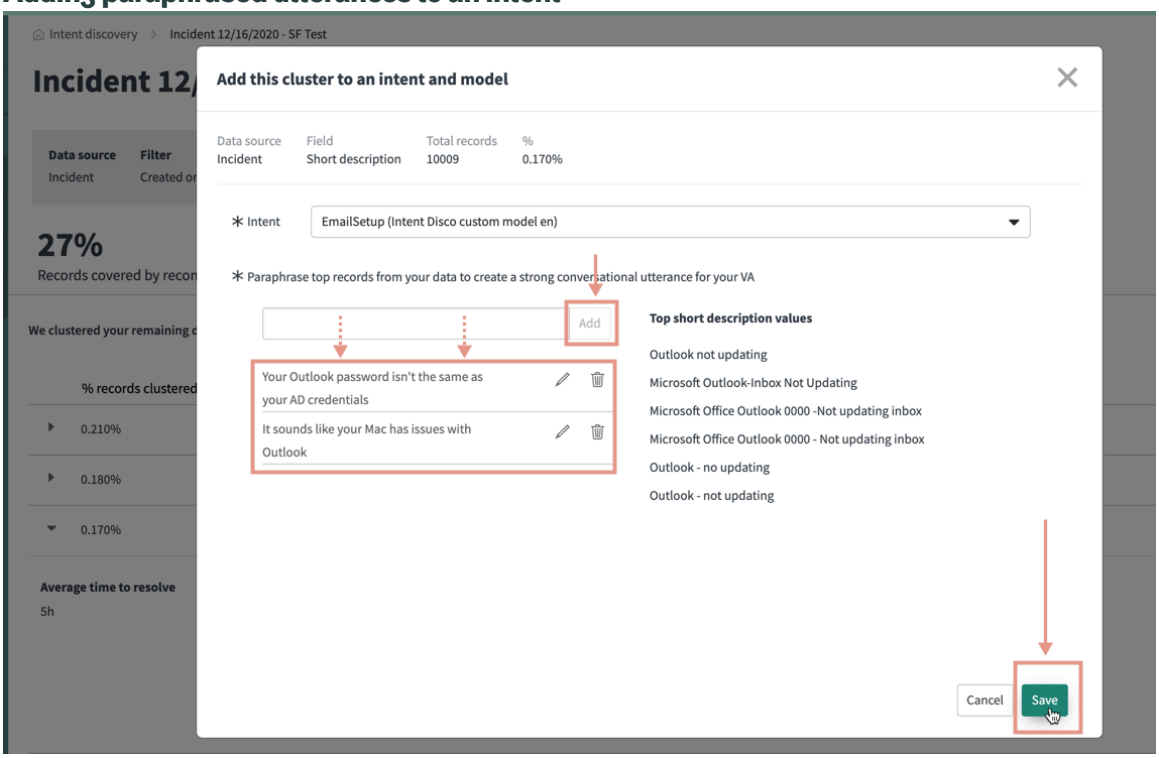
Adding a cluster to an intent and model



4. Enter a few utterance examples into the open text field. Select **Add** each time you complete your entry to save it in the system. Use the pencil icon or the trash can icon respectively to edit or delete your entry.

5. Select **Save**.

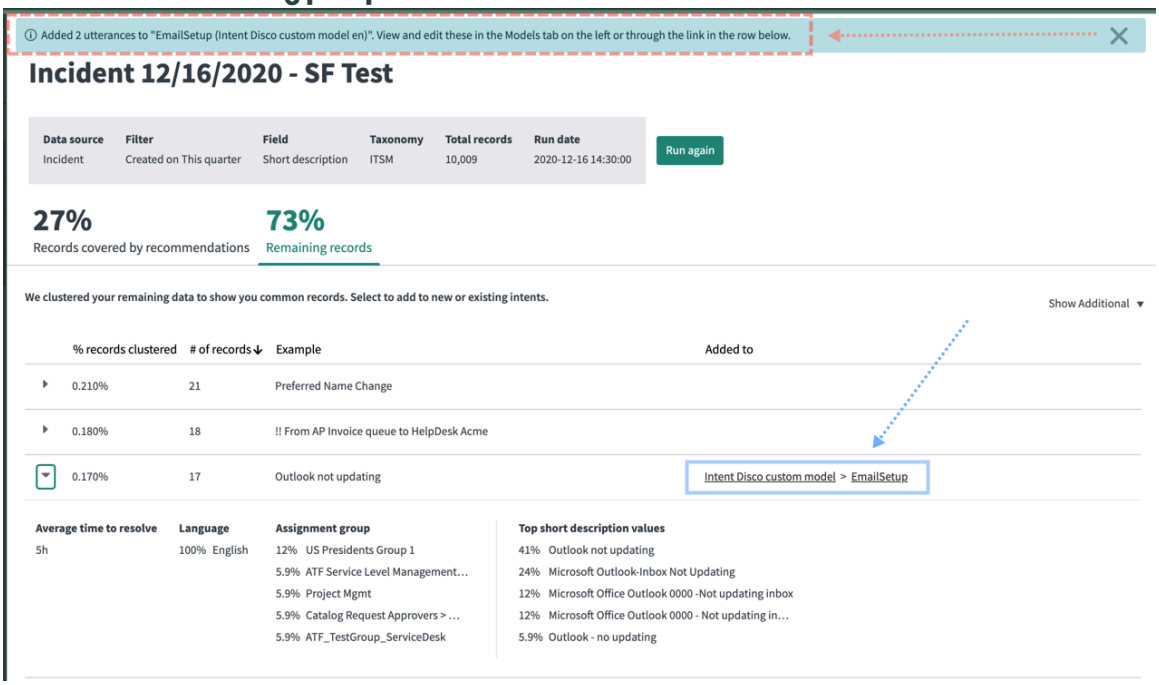
Adding paraphrased utterances to an intent



Result: The records screen appears, showing a banner confirming you added two new utterances to the target intent and its associated model. The model and intent pair appears in the **Added To** column, as shown in the image below.

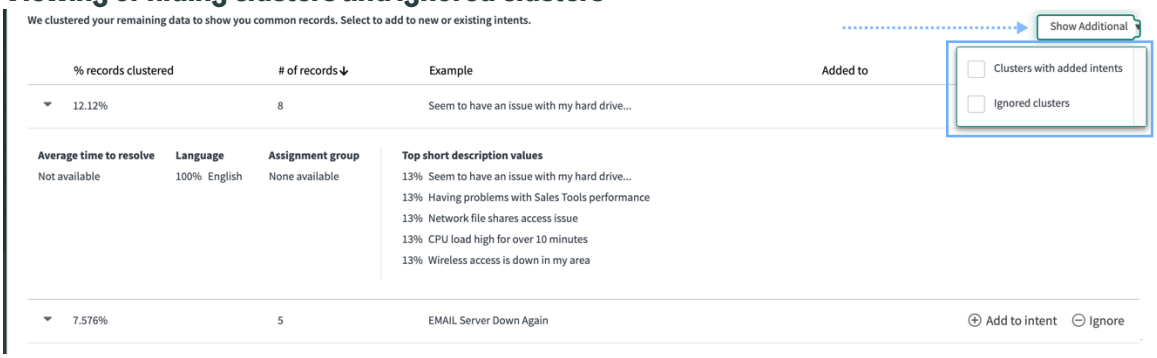
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Confirmation of adding paraphrased utterances to an intent



Use the **Show Additional** filter if you want to show or hide the clusters that have added intents, and the clusters that are ignored.

Viewing or hiding clusters and ignored clusters

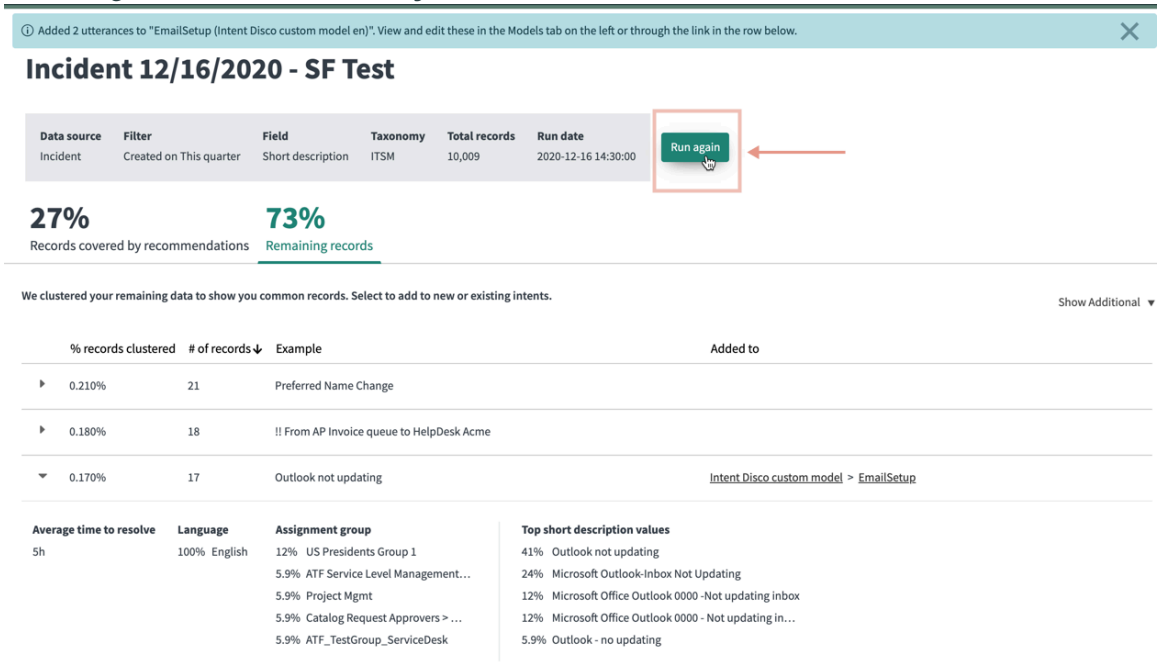


Running another analysis on your Intent Discovery report

1. Select **Run Again**.

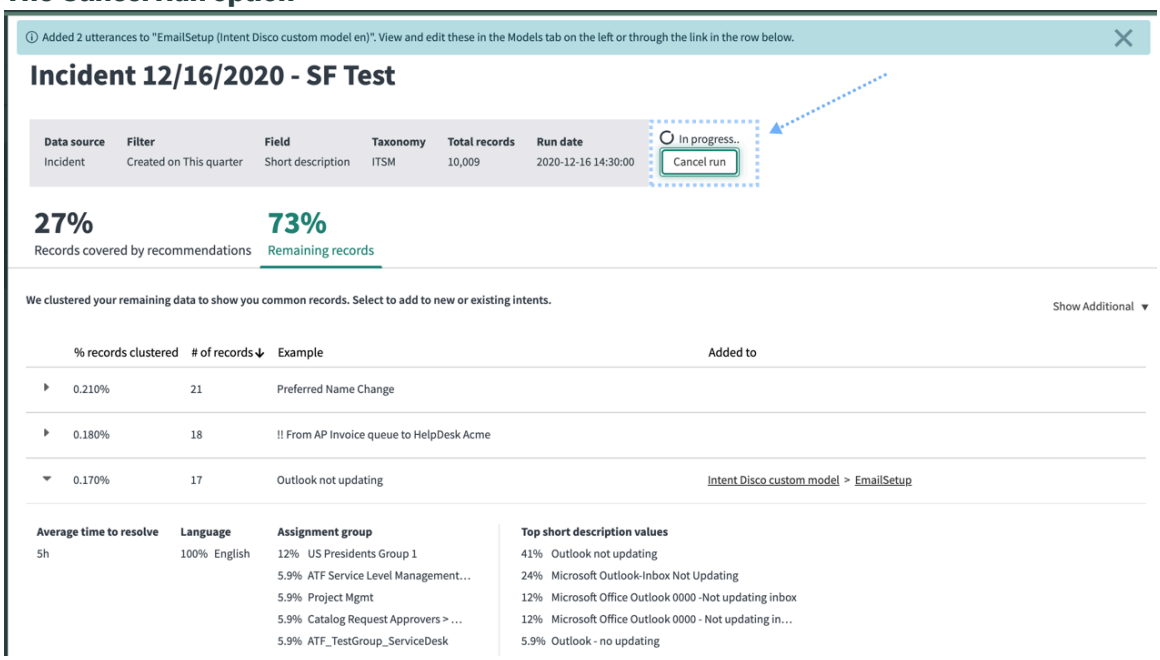
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Selecting the version of the analysis to run



Result: The new run begins. When it's in progress, the option to cancel the run appears, as shown in the image below.

The Cancel Run option

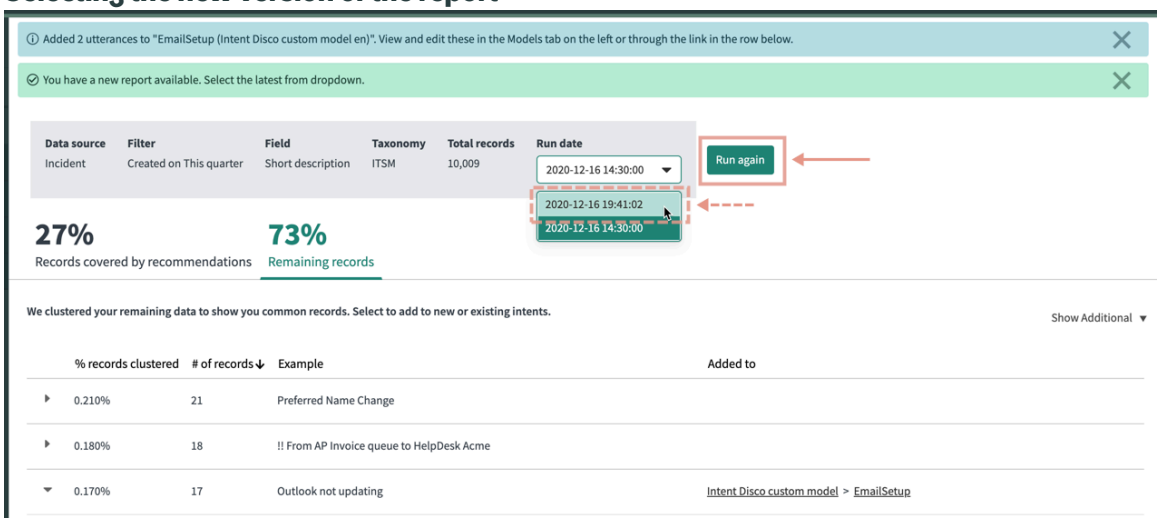


When the run is complete, a new banner appears that states you have a new version of the report.

2. Select the new version, then select **Run Again**.

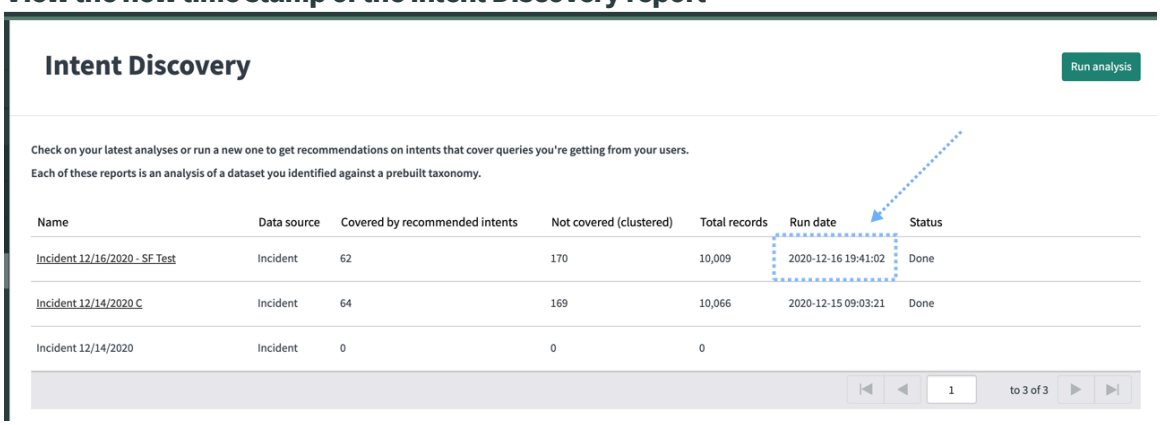
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Selecting the new version of the report



Result: The time stamp you selected for the most recent run appears in the **Run date** column of the Intent Discovery screen.

View the new time stamp of the Intent Discovery report



Install Intent Discovery

You can install the Intent Discovery application (sn_nlu_discovery) if you have the admin role.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the Intent Discovery application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.
- Intent Discovery requires the following plugins. Ensure that these plugins are activated before you install Intent Discovery.

Required ServiceNow plugins

Predictive Intelligence (com.glide.platform_ml)

Enables various Predictive Intelligence and Machine Learning capabilities for training models. See Install Predictive Intelligence.

NLU Workbench - Core (com.glide.nlu)

Adds NLU Model capabilities. See Activate the NLU Workbench.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: admin

About this task

Tables are installed with Intent Discovery:

For more information, see Components installed with Intent Discovery.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Intent Discovery application (sn_nlu_discovery) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you cannot find the application, you might have to request it from the ServiceNow Store.

   Visit the ServiceNow Store 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. In the Application installation dialog box, review the application dependencies.

   Dependent plugins and applications are listed if they will be installed, are currently installed, or need to be installed. If any plugins or applications need to be installed, you must install them before you can install Intent Discovery.
4. Select **Install**.

Components installed with Intent Discovery

Several types of components are installed with activation of the sn_nlu_discovery plugin, including tables.

Note: The Application Files table lists the components that are installed with this application. For instructions on how to access this table, see Find components installed with an application.

Tables installed

Table	Description
Discovery Report [sn_nlu_discovery_report]	Reference table for discovery clusters and intents.
Discovery Message [sn_nlu_discovery_processed_message]	Contains flexible references to the source record and field that was used for analysis.
Discovery Report Trace [sn_nlu_discovery_report_trace]	Contains the information on the utterances added to intents.
Discovery Intent [sn_nlu_discovery_intent]	Contains the intents for a report.
Discovery Cluster	Contains the clusters for a report.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
[sn_nlu_discovery_cluster]	
Discovery Report Definition [sn_nlu_discovery_report_definition]	Includes the necessary information for generating reports.

Natural Language Query

Natural Language Query (NLQ) enables you to query the data in your instance by entering plain language requests into the user interface.

Get started

Choose one of these tiles to get started.

Explore



Learn about NLQ concepts and features.

Use



Use NLQ to query your data with requests in natural language.

Configure



Set up NLQ for your environment.

References



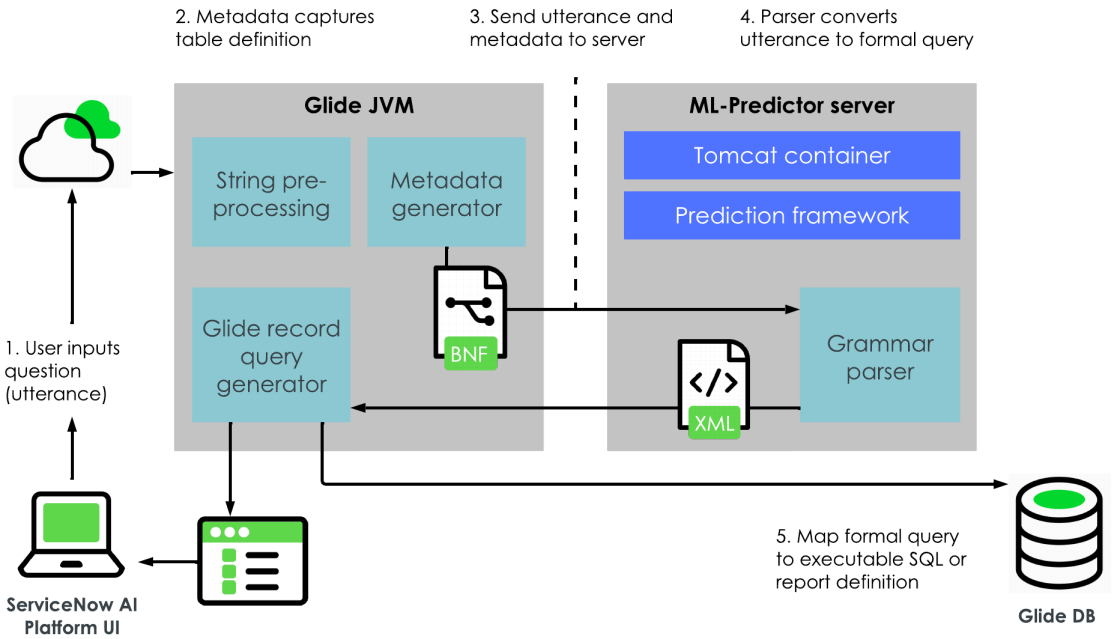
Get details about properties and roles installed with NLQ.

Exploring Natural Language Query

NLQ is a ServiceNow AI Platform feature that is active by default. Use NLQ to query the data in your instance by entering plain language requests into the user interface.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

NLQ data flow for user input and processing



ServiceNow® NLQ translates natural language user input into glide record queries. The queries are rendered into an executable structured format, such as a JavaScript Object Notation (JSON) file or a visual definition. The output, in whichever format, is the response to the user's request.

Note: When a user enters a request directly into the user interface (UI), the text of the request is called an utterance in tables such as NLQ logs.

NLQ is a ServiceNow AI Platform feature that is active by default and supports the following data operations:

- Driving table suggestion
- Filtering
- Grouping and aggregations
- Sorting
- Data visualization (single score, list, bar chart, pie chart, time series)
- Business calendar
- Single number
- Multi-table

For more information, see Using Natural Language Query.

NLQ doesn't support domain separation. It also doesn't support on-premise instances.

Language support

NLQ supports American English by default. For all applications and features except CMDB, NLQ also supports queries in Spanish, French, French Canadian, German, and Japanese.

Note that you must first activate the languages on your instance for NLQ to parse queries in those languages. For more information, see Activate a language.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Using NLQ in other applications and features

Other ServiceNow® applications and features can consume NLQ functionality. See the following resources for more information.

Features that can consume NLQ

Application or feature	Information
Analytics Q & A	Create a report with Analytics Q&A
Configuration Management Database (CMDB)	Query your CMDB data without needing to know table relationships or data structures. • Querying the CMDB • Intelligent Search for CMDB English is the only supported language for CMDB.
NLQ with AI Search in global search	NLQ Genius Results in AI Search (supports English only)
List V2	Filter through any platform list by entering in a plain-language request.

Using Natural Language Query

With Natural Language Query (NLQ), you can query data in your tables by entering requests in natural, everyday language.

NLQ overview

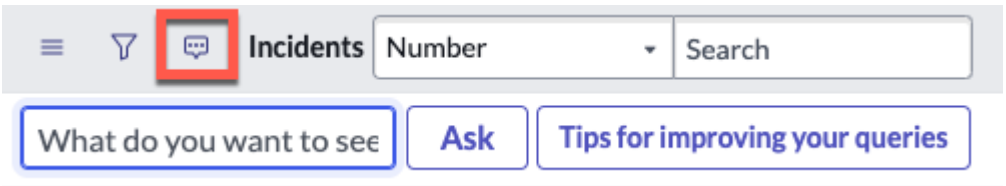
NLQ turns your plain-language requests into structured queries of your data. You don't need to know how to use the condition builder, because NLQ constructs and displays the conditions for you.

If you have a role such as itil that can view and interact with tables, you can use NLQ by selecting the natural language filter icon.

NLQ works on any list on the platform. It returns results only from the table or list you query on.

Find and use the natural language filter

Selecting the natural language filter icon brings up the NLQ interface.

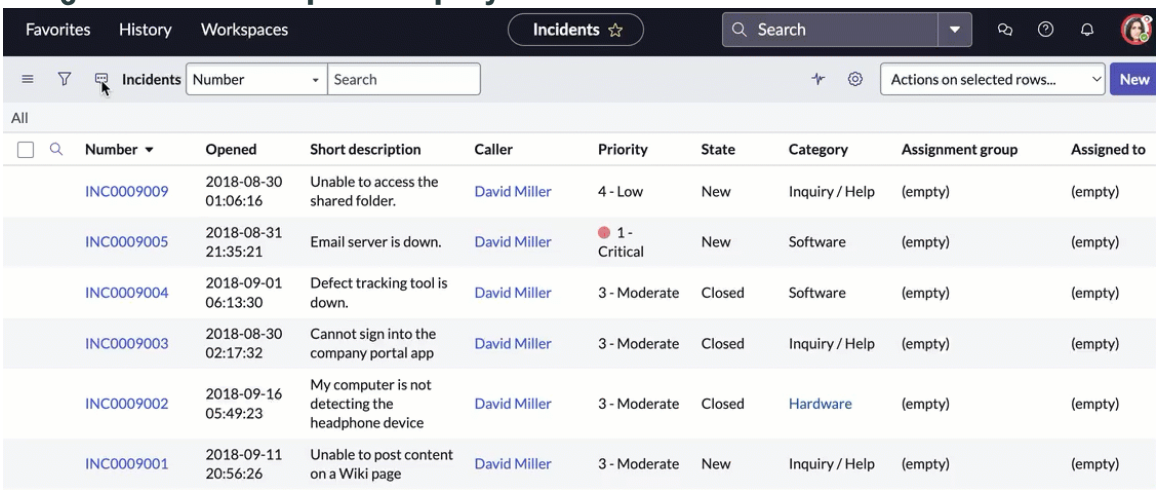


Enter your request in the What do you want to see field, then select **Ask**. NLQ parses your request, then displays the query in the condition builder. The results of the query are displayed in the list.

With the default configuration of NLQ, you can continue to refine your query by entering another request into the What do you want to see field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Using NLQ to build on a previous query



Number	Opened	Short description	Caller	Priority	State	Category	Assignment group	Assigned to
INC0009009	2018-08-30 01:06:16	Unable to access the shared folder.	David Miller	4 - Low	New	Inquiry / Help	(empty)	(empty)
INC0009005	2018-08-31 21:35:21	Email server is down.	David Miller	1 - Critical	New	Software	(empty)	(empty)
INC0009004	2018-09-01 06:13:30	Defect tracking tool is down.	David Miller	3 - Moderate	Closed	Software	(empty)	(empty)
INC0009003	2018-08-30 02:17:32	Cannot sign into the company portal app	David Miller	3 - Moderate	Closed	Inquiry / Help	(empty)	(empty)
INC0009002	2018-09-16 05:49:23	My computer is not detecting the headphone device	David Miller	3 - Moderate	Closed	Hardware	(empty)	(empty)
INC0009001	2018-09-11 20:56:26	Unable to post content on a Wiki page	David Miller	3 - Moderate	New	Inquiry / Help	(empty)	(empty)

This example image and procedure illustrates how to build a query using NLQ:

1. On the Incident table in the natural language filter, enter `show me active hardware tickets` and select the **Ask** button.
2. The condition builder displays **All** > **Active = true** > **Category = Hardware** as the query. The filtered results are displayed in the list.
3. To narrow down the list of results further, enter `without assignment group`. Notice that as you type, NLQ displays possible matches for columns and fields. Select assignment group from the list of suggestions, and then **Ask**.
4. In the condition builder, NLQ adds **>Assignment group is empty** to the query. The list refreshes to display only the matching rows.
5. To reset and start a new query, delete everything in the condition builder so that only **All** remains.

Useful information

Keep the following information in mind when using NLQ.

- Your requests can contain periods and apostrophes, but not wildcard characters such as asterisks or regex.
- To group by a field or column, that column must be visible in the list view. Use the personalize list icon ( ) to hide or display columns.
- For information about querying CMDB tables, see Querying the CMDB.

Tips for improving your queries

When you select the **Tips for improving your queries** button, a modal window appears in the user interface. This window offers the following information about terms you can use in your NLQ questions and requests:

- Sorting or grouping: grouped by; sorted by; A-Z; z-a
- Dates: today; yesterday; last; this; next day(s); week(s); quarter(s); year(s)
- Filtering: starts with; ends with; more than; less than; empty; not empty; and; or
- Other information: my; my team; created by; unassigned

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Query	Example
Sorting or grouping	incidents grouped by category
Dates	created last month
Filtering	short description starts with computer
Other information	unassigned tickets
Single number	INC0777

Configuring NLQ

Enhance your users' query experience by supplementing NLQ with words and terms used in your environment. Review your users' actual requests in the NLQ logs to find possible synonyms and shortcuts to add.

Natural Language Query functionality is included in the base system. Admins can optionally expand NLQ's range of understanding by adding terms that commonly occur in users' requests. Review the following topics to learn more.

For information about system properties related to NLQ, see Natural Language Query References.

Create an NLQ synonym

Add synonyms to improve the ability of NLQ to recognize the various ways your users request data. With synonyms, you can map commonly used words or terms to table columns.

Before you begin

Role required: admin, nlq_admin, or pa_analyst

About this task

NLQ synonyms enable you to map common words in your users' requests to the columns in your instance tables. When these words are detected in natural-language queries, NLQ replaces them with actual column and table names, then submits the formal query.

Several synonyms are provided in the base system, and you can add more for your use cases and business requirements. Review NLQ logs of actual user utterances to find possible synonyms to add. For more information, see View NLQ logs.

You can create a new synonym, or modify an existing synonym, as follows.

- If any synonym already exists for your target table and column, you must add your new synonym to the current record. Separate multiple synonyms with a comma.
- Synonyms can point to reference fields, using dot-walking. For more information, see Dot-walking examples.
- Synonyms are case-insensitive in queries.
- Synonyms can contain apostrophes and periods, but not commas.
- Synonym records are associated to one table. You can use the same synonym term for more than one table, but you must create a record for each table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Ensure that you are in the application scope you want for your synonym, then navigate to **All** > **NLQ** > **Synonyms**.
2. Select **New**.
   If you are updating an existing record, select its **Column name** in the list.
3. On the form, fill in the fields.

Field	Description
Synonym Type	Type of synonym. ○ **Table or Column**: Synonym for a value on a table or a specific column in that table. ○ **Record**: Synonym for a value on the CMDB tables [cmdb_rel_type_table]. For more information, see Querying the CMDB. **Note:** Supports cmdb_rel_type,cmdb_group, cmdb_ci_service_technical, cmdb_ci_service_discovered, cmdb_ci_query_based_service tables.
Application	[Read only] Application scope that can use the synonym in a query. The default is **Global**. When creating a synonym, ensure that you are in the scope you want for the synonym.
Table	Source table. You can use the same term for more than one table, but you must create a synonym record for each table. The synonym's mapping is unique to the table.
Column name	Specific column on the source table.
Language	Language of the synonym. Should match the language of the source value.
Synonyms	Words or short phrases that the system should map to column names or tables, when converting the utterance to a formal query. Separate multiple synonyms with a comma.
Extend to All Child Tables	**Table or Column** type only. Select this option to make any child tables of the source table inherit the synonyms.

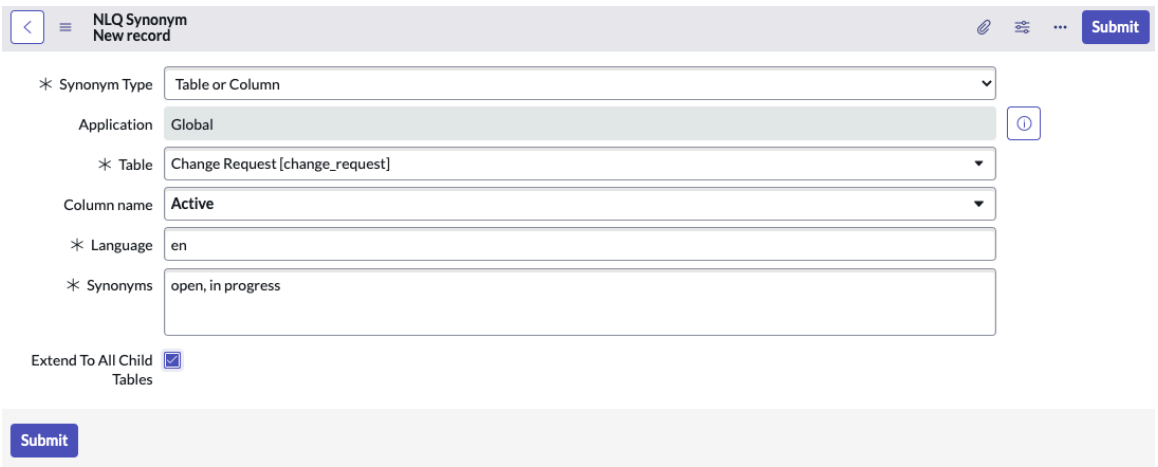
4. Select **Submit** if new, or **Update** if you are modifying an existing record.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
The new synonym is available to your users as soon as they refresh the browser window of the list.

Example: NLQ Synonym for the Active field

The following image shows an example of an NLQ synonym record for the **Active** field on a change request. The synonyms open and in progress are replaced by the term active when the system submits a query.



With this synonym record, a user can type `show open change requests` or `change requests in progress` and the system displays active change requests.

Create an NLQ shortcut

Create a semantic shortcut to help improve the ability of NLQ to recognize the various ways your users request data. Semantic shortcuts operate similarly to NLQ synonyms by mapping common words to columns, but for a selected table when certain conditions are met.

Before you begin
Role required: admin, nlq_admin, or pa_analyst

About this task
Like NLQ synonyms, semantic shortcuts enable you to map common words in your users' requests to the columns in your instance tables. When these words are detected in natural-language queries, NLQ replaces them with actual column and table names, then submits the formal query.

Semantic shortcuts provide a condition builder so that you can filter which records are covered by the terms you provide.

Some semantic shortcuts are provided in the base system, and you can add more for your use cases and business requirements. Review NLQ logs of actual user utterances to find possible terms to add. For more information, see View NLQ logs.

You can create a new shortcut, or modify an existing shortcut, as follows.

- If any shortcut already exists for your target table and filter conditions, you must add your new shortcut to the current record. Separate multiple shortcuts with a comma.
- Shortcuts can point to reference fields, using dot-walking. For more information, see Dot-walking examples.
- Synonyms are case-insensitive in queries.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Synonyms can contain apostrophes and periods, but not commas.
- Synonym and shortcut records are associated to one table. You can associate the same synonym term to more than one table, but you must create a record for each table.

Procedure

1. Ensure that you are in the application scope that you want for your shortcut, then navigate to **All** > **NLQ** > **Semantic Shortcuts**.
2. Select **New**.
   If you are updating an existing record, select its row in the list.
3. On the form, fill in the fields.

Field	Description
Table	Source table. You can use the same term for more than one table, but you must create a record for each table. The synonym's mapping is unique to the table.
Conditions	Conditions on the source table that must be met for the synonym to work.
Synonyms	Words or phrases that the system should map to column names or tables, when converting the utterance to a formal query. Separate multiple synonyms with a comma.
Application	[Read only] Application scope that can use the synonym in a query. Default is **Global**. Ensure that you are in the target application scope for your shortcut before creating it.

4. Select **Submit** if new, or **Update** if you are modifying an existing record.

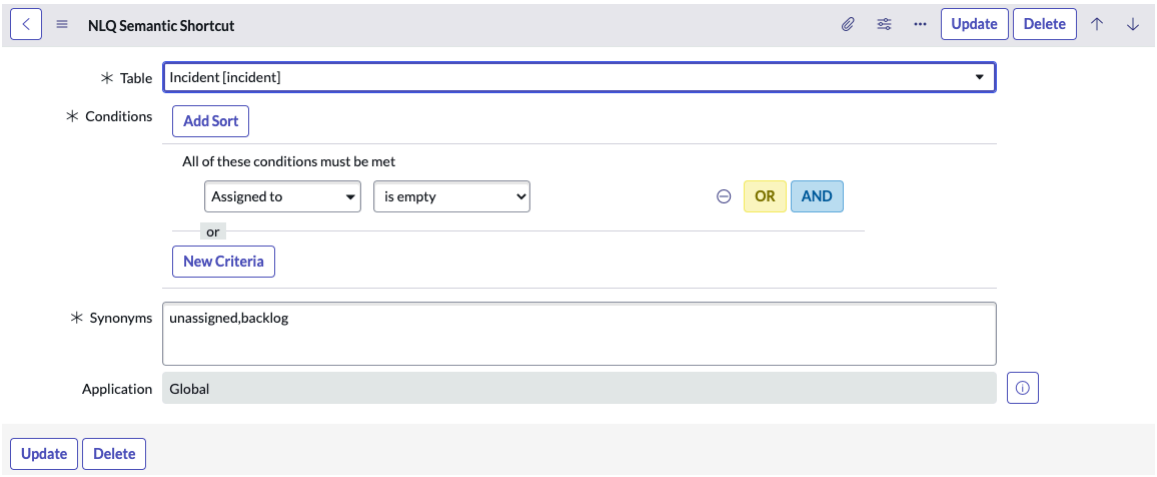
Result

The new shortcut is available to your users as soon as they refresh the browser window of the list.

Example: Semantic shortcut for Incident backlog

The following image shows an example of a semantic shortcut used on incident records. Incidents contain the Assigned to field. Using the condition builder, you can provide synonyms for when the field is empty. When NLQ detects unassigned and backlog in your users' input, it replaces them with the SQL clause `assigned_toISEMPTY`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



A user can enter `show me incident backlog` or `unassigned incidents` and the system displays incidents with an empty Assigned to field.

View NLQ logs

Review NLQ logs to see how the system has handled your users' plain-language requests. Use log records from attempted requests to expand NLQ synonyms or shortcuts.

Before you begin

Role required: nlq_admin, pa_analyst, or admin

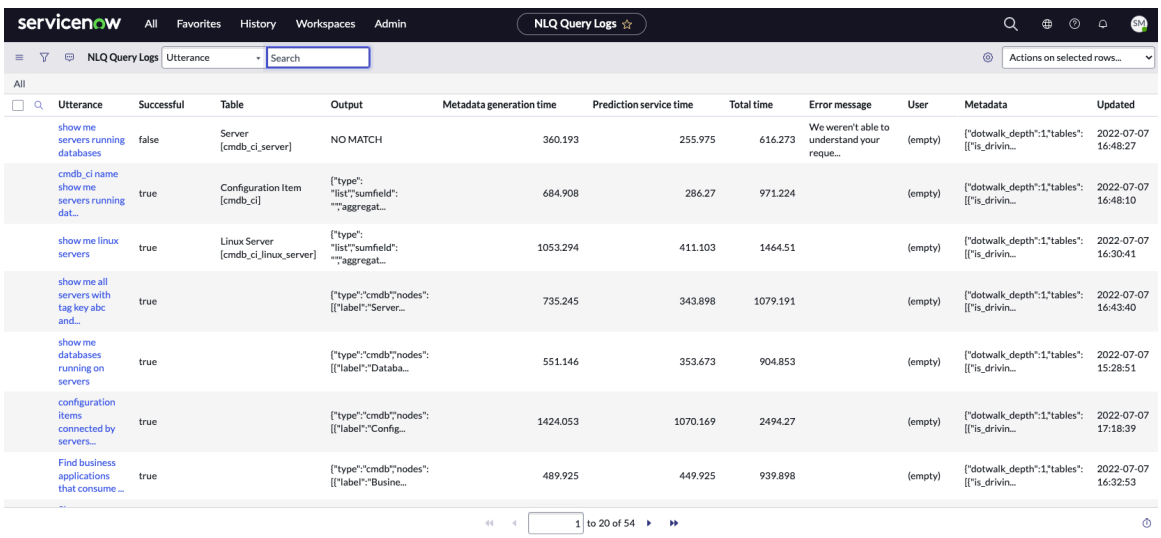
About this task

Every natural language query is logged in the *NLQ* **Query Logs** table [nlq_query_log]. Each log entry provides details such as the table that was queried, whether the query succeeded, and how the results were generated. Other available fields include the following.

- Output Source: how the results were generated. The value **BNF** indicates a rules-based method. The value **GAI** indicates the Now LLM Service (fallback) method.
- Source: the location from which the query was initiated. The value **AC** indicates Analytics Overview. The value **CMDB_WS** indicates CMDB Workspace.
- Utterance: the original natural-language request which triggered NLQ.

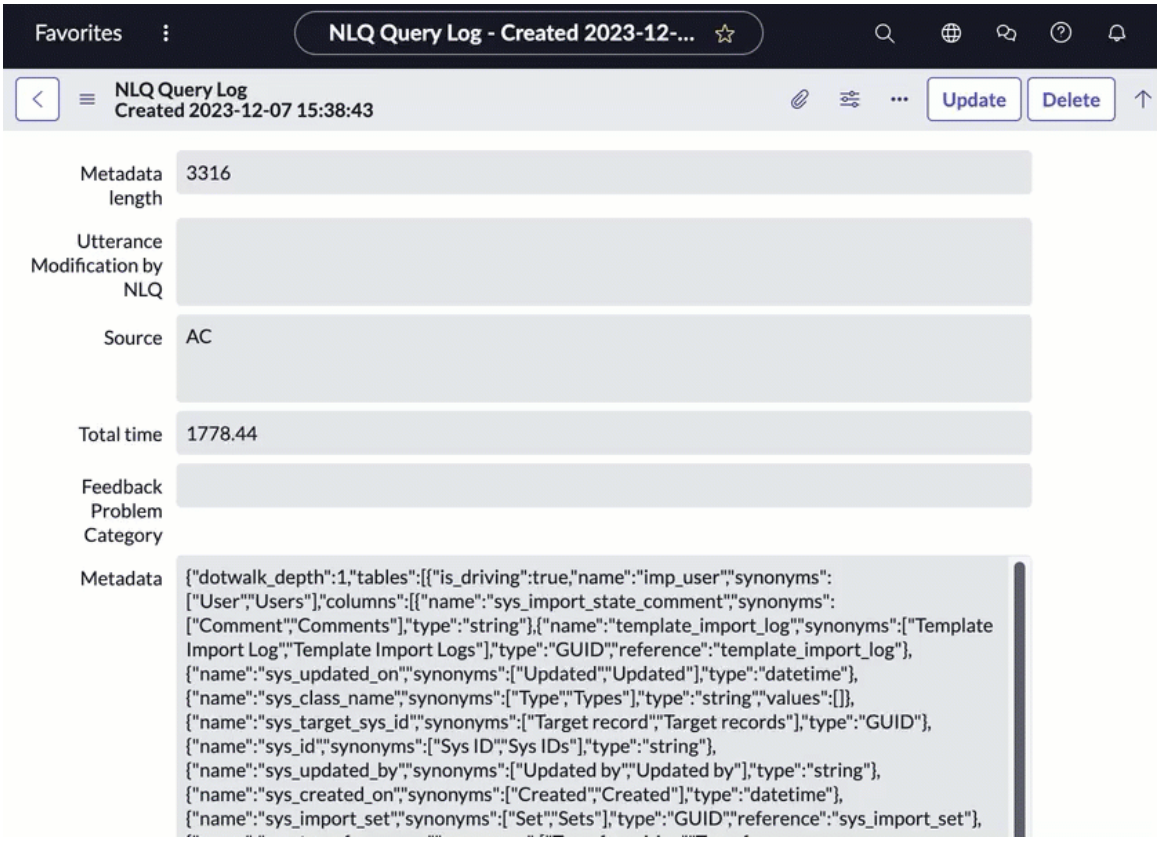
Procedure

1. Navigate to **All** > **NLQ** > **Logs**.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Select the value in the **Utterance** column to open the full log entry.



What to do next

Based on your users' attempted queries, consider adding more synonyms or shortcuts.

View NLQ Table Guesser logs

Use the Table Guesser logs to review the CMDB tables that were picked by NLQ in response to plain-language queries.

Before you begin

This module is read-only.

Role required: nlq_admin.

About this task

When a plain-language request for CMDB data is parsed by NLQ, the system tries to determine the specific tables the user intended. The system's guesses, including confidence levels, are recorded in the NLQ Table Guesser logs module [nlq_table_guesser_log].

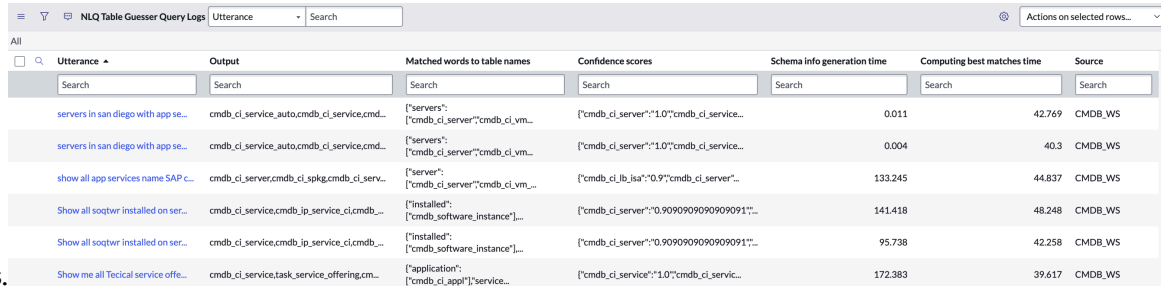
Review these logs to troubleshoot which CMDB tables were inferred by NLQ.

Note: This list is read-only for NLQ admins. There are no buttons or actionable functions in the UI.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **NLQ** > **Table Guesser Logs**.
2. Select the personalize list icon ( ) to display the columns **Matched words to table names** and **Confidence scores**.



   The **Utterance** column shows the user's natural language query.
3. Use the information in the **Matched words to table names** column for troubleshooting if needed.
   In this column's value field, the first word was found in the utterance. Next is a list of the CMDB tables that were matched to that word.

What to do next

For more information on CMDB queries, see Exploring CMDB Query Builder.

Natural Language Query References

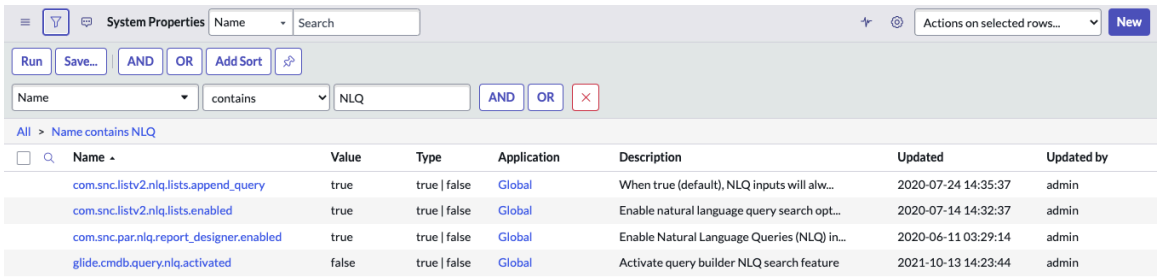
The following components are installed with Natural Language Query.

NLQ properties

The Natural Language Query (NLQ) properties control how and where NLQ operates.

Admins can edit properties of NLQ by navigating to **All** > **System Properties** > **All Properties**. Filter for the NLQ properties.

Note: Editing these system properties requires the admin role. The nlq_admin role does not have permission to edit records in this table.



Property	Description
com.snc.listv2.nlq.lists.append_query	• True: NLQ inputs add onto existing queries via an "and" operator • False: New NLQ input replaces any existing queries

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	Example: You run two queries. • Query 1: Incidents with critical priority • Query 2: assigned to John Smith If the property is set to true, the results show incidents with critical priority that are assigned to John Smith. If the property is set to false, the results show only items assigned to John Smith.
com.snc.listv2.nlq.lists.enabled	• True: Enables NLQ search option for List v2 • False: Removes NLQ search option for List v2
com.snc.nlq.gai_enabled	• True: The Now LLM Service fallback is available • False: The Now LLM Service fallback is not available Initially, queries are interpreted using a rules-based method. If that method fails, queries are passed to the Now LLM Service as a fallback. Queries that fail both of these methods are marked as unsuccessful in the NLQ log.
com.snc.par.nlq.report_designer.enabled	• True: Enables NLQ in Report Designer • False: Removes NLQ in Report Designer
glide.cmdb.query.nlq.activated	• True: NLQ search feature is active in CMDB Query Builder • False: NLQ search feature is inactive in CMDB Query Builder
glide.service_portal.ais_nlq_enabled	• True: Enables NLQ in global search • False: NLQ is not available in global search

Natural Language Query roles

Natural Language Query (NLQ) is installed with these roles.

To learn more about managing per-user subscriptions, see Managing per-user subscriptions in Subscription Management and contact your account representative.

NLQ Admin [nlq_admin]

The administrator for Natural Language Query.

Module Access

Has full access to the following modules:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- NLQ Cmdb Implicit Relationships. For more information see Intelligent Search for CMDB.
- NLQ Query Logs
- NLQ Semantic Shortcuts.
- NLQ Synonyms.

Has read-only access to the following module:

NLQ Table Guesser Query Logs.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

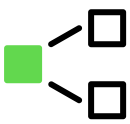
To avoid granting an admin role, use the pa_analyst role for NLQ Cmdb Implicit Relationships, NLQ Semantic Shortcuts, and NLQ Synonyms.

Agentic Desktop

ServiceNow® Agentic Desktop enables you to design, configure, and manage desktop actions that automate repetitive tasks in your desktop environment. AI agents can autonomously and semi-autonomously process instructions, generate execution plans, and run desktop actions across legacy systems, thick client applications, and business applications without APIs.

Agentic Desktop is a client application that is installed on the Windows operating system. Agentic Desktop offers two workspaces, the Design workspace for design experience and the Execution workspace for the runtime experience. Design workspace enables you to automate multi-step processes with intelligence and adaptability. Execution workspace enables AI agents to execute desktop actions in an isolated desktop session.

Get started

Explore	Configure	Design
		
Learn about Agentic Desktop concepts and features	Set up Agentic Desktop to automate desktop tasks	Design automations for legacy applications that do not have APIs

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create an AI agent	**Use**	Reference
		
Create your own custom AI agents with advanced multi-agent reasoning frameworks	Use the Agentic Desktop application to execute automations using AI agents.	Get details about the Agentic Desktop properties and components

Important:

- 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.

Helpful resources

- Workflow Automation product on the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

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. In addition, this application will collect information about scripts (and associated script records) in which Now Assist for code generation is called. Customers can opt out of future data collection at any time, as described in the Now Assist Opt-Out page.

For more information, see the Now Assist documentation.

Exploring Agentic Desktop

Design, configure, and manage desktop actions with Agentic Desktop to automate repetitive tasks on your desktop using AI agents and agentic workflows.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agentic Desktop overview

Agentic Desktop is a no-code solution that helps you automate repetitive tasks in legacy desktop applications lacking APIs or backend integrations. Agentic Desktop leverages AI agents created in the ServiceNow AI Platform to interact with desktop applications, perform UI-based tasks, and automate end-to-end workflows.

Agentic Desktop is a client application installed on the Windows operating system. Agentic Desktop provides two workspaces: the Design workspace, where you create and configure desktop automations, and the Execution workspace, where those automations run.

The Design workspace lets you build multi-step desktop actions with intelligence and adaptability. The Execution workspace runs desktop actions in an isolated desktop session and is launched automatically when you test a desktop action or trigger an automation from the Now Assist panel. You do not open the Execution workspace manually. For more information, see Design workspace and Execution workspace.

How it fits into ServiceNow workflows

Agentic Desktop integrates with AI Agent Studio, enabling you to publish, manage, and incorporate desktop actions into your broader ServiceNow workflows. This integration lets you automate both cloud and desktop applications, giving your AI agents broader capabilities within ServiceNow.

Agentic Desktop users

Users

User	Description
Administrators	Set up and deploy the Agentic Desktop application.
Developers	Design, manage, and deploy automation solutions within their organization.
Business users	Automate repetitive desktop tasks without extensive technical knowledge.
Process owners	Optimize business processes by automating manual desktop interactions.

Key interfaces related to Agentic Desktop

Design workspace

Create and test desktop actions that define how automations interact with desktop applications.

Execution workspace

Automatically runs desktop actions in an isolated desktop session during testing or execution. You do not open this workspace directly.

AI Agent Studio

Create and manage AI agents that run desktop actions.

Now Assist panel

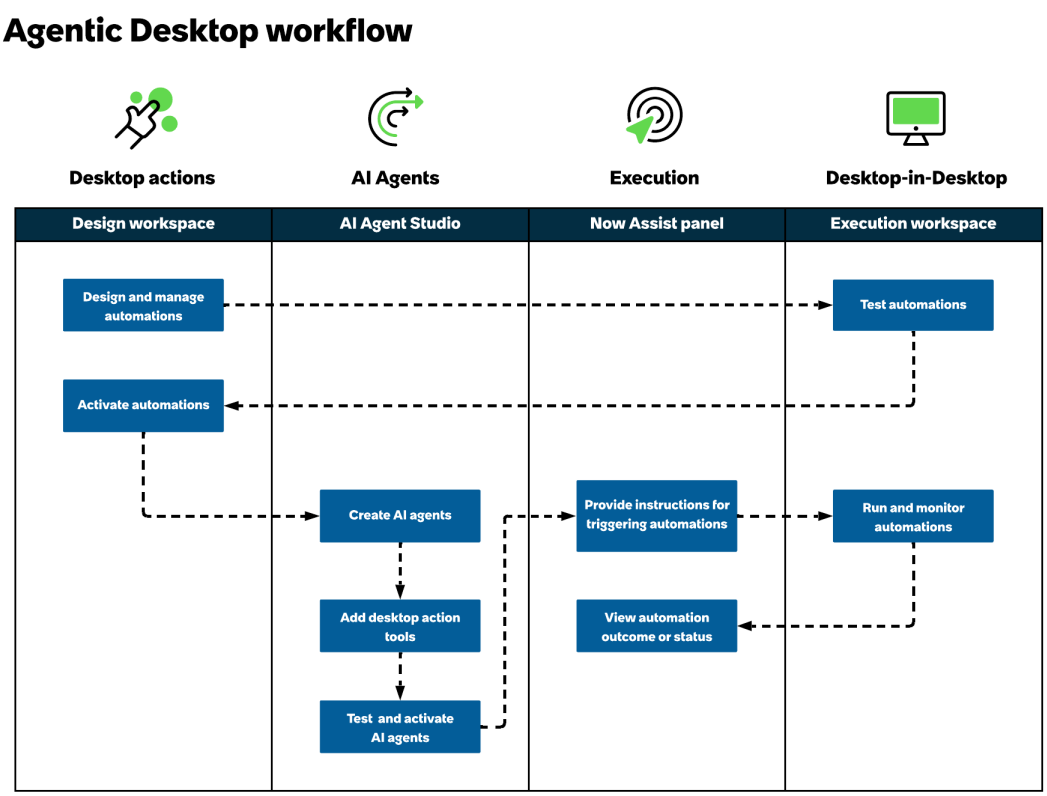
Trigger automations from within ServiceNow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

These components work together to separate design, execution, and monitoring, ensuring secure and reliable desktop automation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

How Agentic Desktop works end-to-end



This workflow shows how Agentic Desktop works end to end from designing desktop actions to running them during execution. As you move through the workflow, note which steps you perform manually and which steps are handled automatically by the system.

Important: Access to the Design workspace and Execution workspace depends on the user's role.

- When users with the AI Agent Admin (sn_aia.admin) role sign in from their desktop, they can access the Home page and Design workspace to create desktop actions. When they test a desktop action, the Execution workspace launches automatically.
- When users with the Now Assist panel user (now_assist_panel_user) role trigger an automation from the Now Assist panel, the Execution workspace launches automatically to run the desktop action.

You don't open the Execution workspace directly. It launches automatically when you test or run a desktop action.

- With the AI Agent Admin role, use the no-code Design workspace to create reusable desktop actions that AI agents can execute. You can visually create and simulate application interactions, such as clicks, text entry, and extraction that mimic human behavior using screen and UI element recognition.
- When you test a desktop action, the system automatically launches the Execution workspace in an isolated desktop session. You do not open or manage the Execution workspace directly.
- Activate your desktop actions from the Design workspace so that AI agents can use them as tools in AI Agent Studio.
- Create AI agents by providing simple text instructions in AI Agent Studio.
- Add your designed desktop actions as tools to your AI agents.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc. in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see Desktop actions in Agentic Desktop and Creating AI agents for Agentic Desktop.

Agentic Desktop capabilities

Design desktop actions

- Create, configure, and manage reusable desktop actions with application metadata.
- With the Recorder feature, automatically record and capture user interactions and contextual information while you perform actions on desktop applications.
- Additionally, manually capture screens, define UI interactions, and structure steps in a no-code Design workspace.

Run in the background

Enable silent and non-intrusive execution of agent tasks, supporting uninterrupted user workflows.

Support for diverse applications on Windows

Automate tasks across a wide variety of environments and UI types, including legacy systems, thick client applications, and business applications on Windows operating system.

Core desktop capabilities

Support common desktop actions such as form filling, application clicks, and OS file handling.

Performance

Supports automations with native integration to AI Agent Studio.

Seamless integration

Publish configured desktop actions directly to AI Agent Studio for easy access and deployment.

Impersonating users

You can trigger AI agents from the Now Assist panel while impersonating another user, provided the impersonated user has the required roles. The sn_aia.admin role is required to use AI Agent Studio, and the now_assist_panel_user role is required to trigger AI agents that execute desktop actions in the Execution workspace. For more information, see Impersonating users.

What to explore next

To learn more about configuring and using Agentic Desktop, see:

- Configure Agentic Desktop
- Desktop actions in Agentic Desktop
- Creating AI agents for Agentic Desktop
- Examples of creating desktop actions
- Examples of executing desktop actions using AI agents
- Components installed with Agentic Desktop
- System requirements and limitations in Agentic Desktop

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agentic Desktop Design workspace

The Design workspace is an interactive environment within Agentic Desktop that enables you to create desktop actions by recording and configuring user interactions with desktop applications. The workspace provides a visual canvas where you can design multi-screen automation workflows that capture business processes across different applications.

Important: Access to the Design workspace and Execution workspace depends on the user's role.

- When users with the AI Agent Admin (sn_aia.admin) role sign in from their desktop, they can access the Home page and Design workspace to create desktop actions. When they test a desktop action, the Execution workspace launches automatically.
- When users with the Now Assist panel user (now_assist_panel_user) role trigger an automation from the Now Assist panel, the Execution workspace launches automatically to run the desktop action.

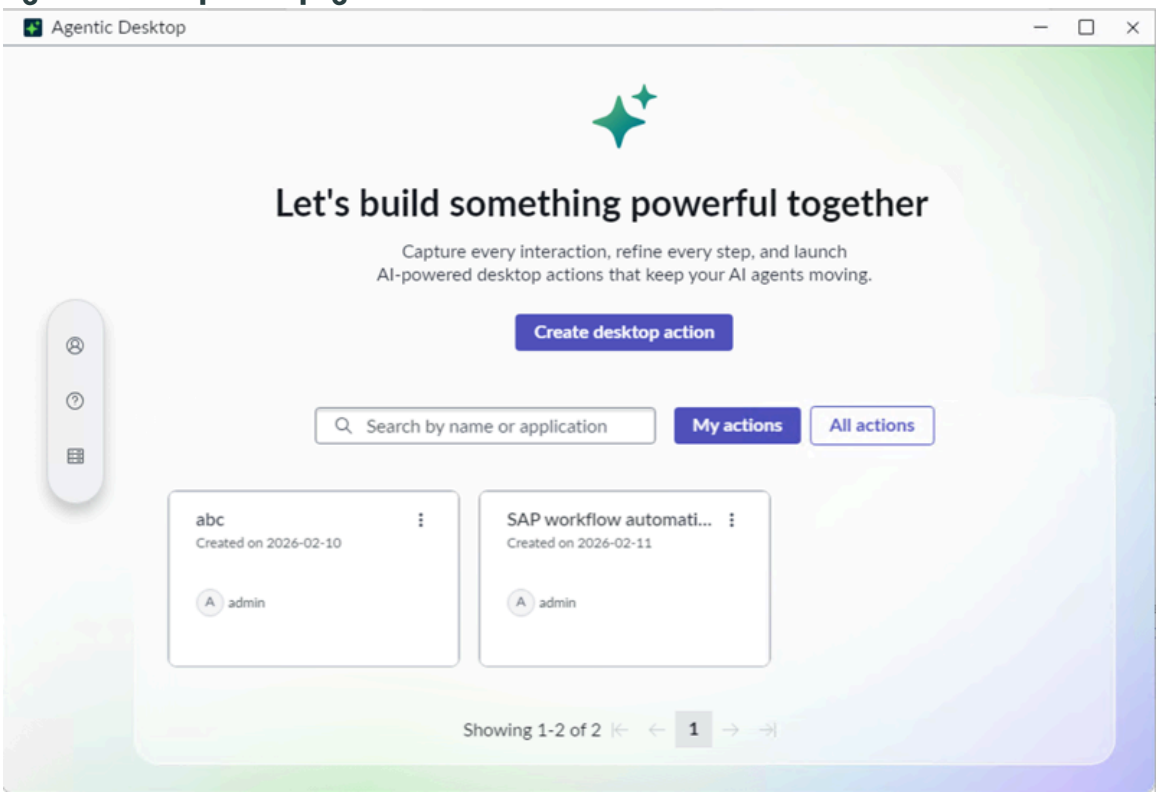
You don't open the Execution workspace directly. It launches automatically when you test or run a desktop action.

Agentic Desktop home page

The home page provides an intuitive interface to help you quickly create, track, and manage your desktop actions. You can easily find and manage your existing desktop actions and monitor their status.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agentic Desktop home page



When you log in to Agentic Desktop, the home page is the first screen you see. The home page lets you interact with various functionalities.

Create desktop action

You can create desktop actions by manually capturing steps or auto-capture steps through recording.

Search and manage actions

You can search for specific actions and manage them directly from the home page.

My Actions and All Actions

My Actions tab shows only the actions published by the logged in user. **All Actions** tab shows list of all actions available on the instance the logged in user is connected to.

Desktop action cards

Configured desktop actions with details, such as created date, user who created, applications added, and activation status.

Edit desktop actions

You can modify a desktop action by selecting the **Edit** option for each card. Any changes you make are saved and reflected on the home page.

Access to profile and documentation

A dedicated section enables you to view your profile information, access product documentation, and check the product version.

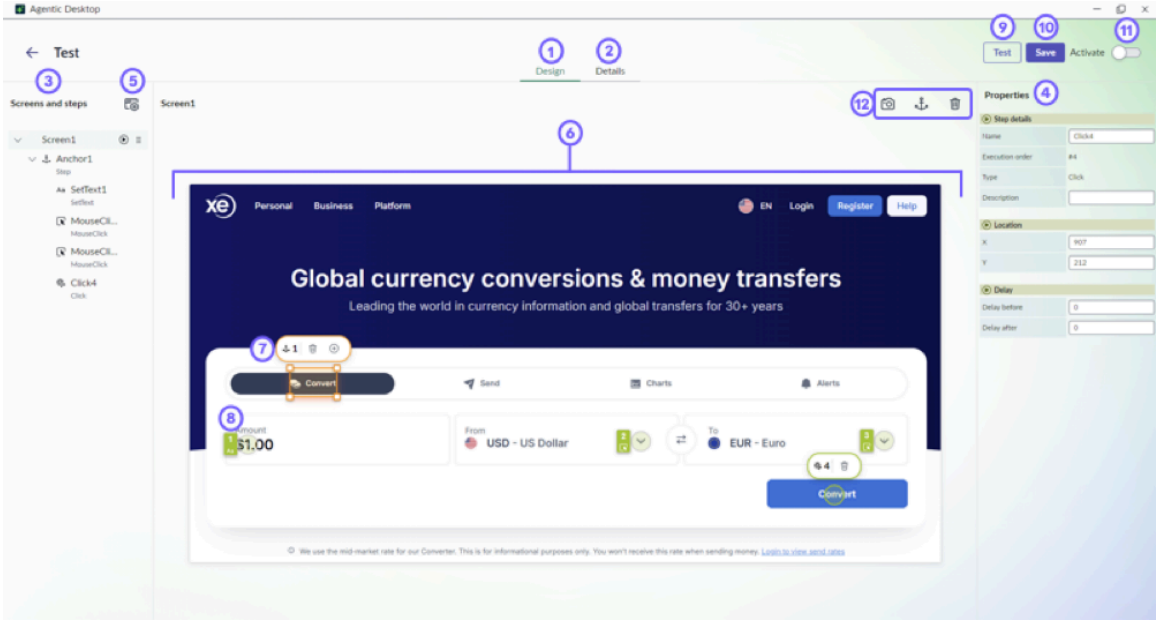
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Agentic Desktop Design tab

The Design workspace provides a no-code environment for creating, managing, and testing your desktop actions. You can either auto-capture or manually record a series of steps you perform on your computer, such as clicking buttons, entering text, selecting from drop-down, or interacting with different applications. You can then save this sequence as a reusable desktop action.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Design tab



No.	Element	Description
1	**Design** tab	Enables you to design your desktop actions by adding screens, anchors, and steps.
2	**Details** tab	Enables you to add details to your desktop actions, such as name, description, and applications that this desktop action interacts with. You can also review the inputs and outputs of this desktop action.
3	**Screens and Steps** panel	Shows all the captured screens and added anchors and steps in a sequence the desktop action must perform them. You can drag and move these items to change the order. You can also run a screen-level test by selecting the Run screen test icon.
4	**Properties** panel	Enables you to configure the properties for screens, anchors, and steps.
5	**Capture options**	Provides two options for capturing interactions with applications that you want to automate. • **Auto-capture steps**: Automatically records the steps that you perform for the task that you want to automate. • **Manual capture screens**: Enables you to capture the screens manually, insert anchors, and add steps that you want to automate.
6	Captured screen	The application screen that you captured. You can delete any screen by selecting the **Delete** icon from the Screens and steps panel.
7	**Anchor**	An anchor is a reference point on the screen that helps the automation identify and interact

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc. in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- Desktop actions in Agentic Desktop
- Automate repetitive tasks by auto-capturing steps in Agentic Desktop
- Automate repetitive tasks by manually capturing steps in Agentic Desktop
- Example: Automate badge request management using Agentic Desktop
- Example: Automate shipping management tasks using Agentic Desktop
- Screen, anchor, and step properties in Agentic Desktop

Action recorder in Agentic Desktop

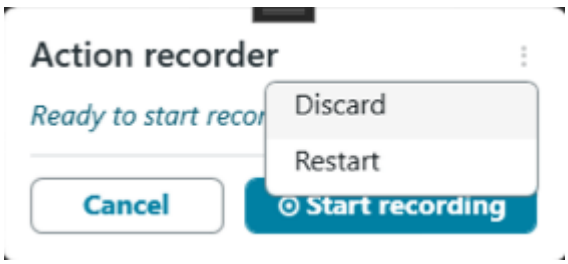
With Action recorder, you can auto-capture steps to automate repetitive tasks in Agentic Desktop. You can save the steps that you perform on the application elements as a reusable desktop action of type UI block.

Action recorder helps you record your interactions with desktop applications to create automated workflows. By recording steps, you can automate tasks that replicate your interactions. The recorded actions are displayed as screenshots in the Design workspace with anchors and steps automatically added.

The recorder captures the following items:

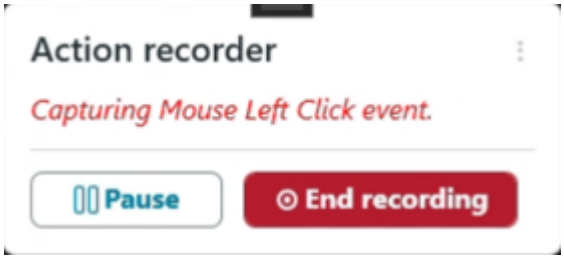
- Visual snapshots of your screen at key interaction points.
- Steps in the form of the buttons, fields, and other interface components you interact with.

You can select any of the following options when needed from the **More options** menu on the Action recorder panel:



- **Start recording**: Start recording the steps
- **Pause**: Skip recording steps
- **Restart**: Restart recording the steps
- **Discard**: Discard the recording if it doesn't meet your needs
- **Stop recording**: Stop recording steps

Each step that you perform is captured sequentially and the type of UI action is displayed for each step. For example, Capturing Mouse Left Click event.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

During recording, the steps that you perform are captured as screens. Related steps are captured in the same screen, for example, the steps related to filling in the text fields in the same window. A new screen is captured in the following conditions:

1. Window is changed
2. Button is clicked or option is selected that might change the layout
3. Page is scrolled
4. Special keyboard keys are used, such as Enter or Tab

This helps maintain the correct screen context and reliable automation replay.

You can capture maximum of 50 steps using the recorder in a recording session. While auto-capturing steps, a counter displays the remaining number of steps you can capture using the recorder (for example, “35 of 50 max”). Recording stops automatically after you capture 50 steps. If you need to capture additional steps, start a new recording session. The new recording adds screens and steps to those captured in the previous recording.

Related topics

Agentic Desktop Design workspace

Automate repetitive tasks by auto-capturing steps in Agentic Desktop

Create badge desktop action in Agentic Desktop

Agentic Desktop Execution workspace

Execution workspace enables you to test, run, and monitor your desktop actions. It enables you to observe how your automations interact with desktop applications, including handling situations where human input is needed.

Important: Access to the Design workspace and Execution workspace depends on the user’s role.

- When users with the AI Agent Admin (sn_aia.admin) role sign in from their desktop, they can access the Home page and Design workspace to create desktop actions. When they test a desktop action, the Execution workspace launches automatically.
- When users with the Now Assist panel user (now_assist_panel_user) role trigger an automation from the Now Assist panel, the Execution workspace launches automatically to run the desktop action.

You don’t open the Execution workspace directly. It launches automatically when you test or run a desktop action.

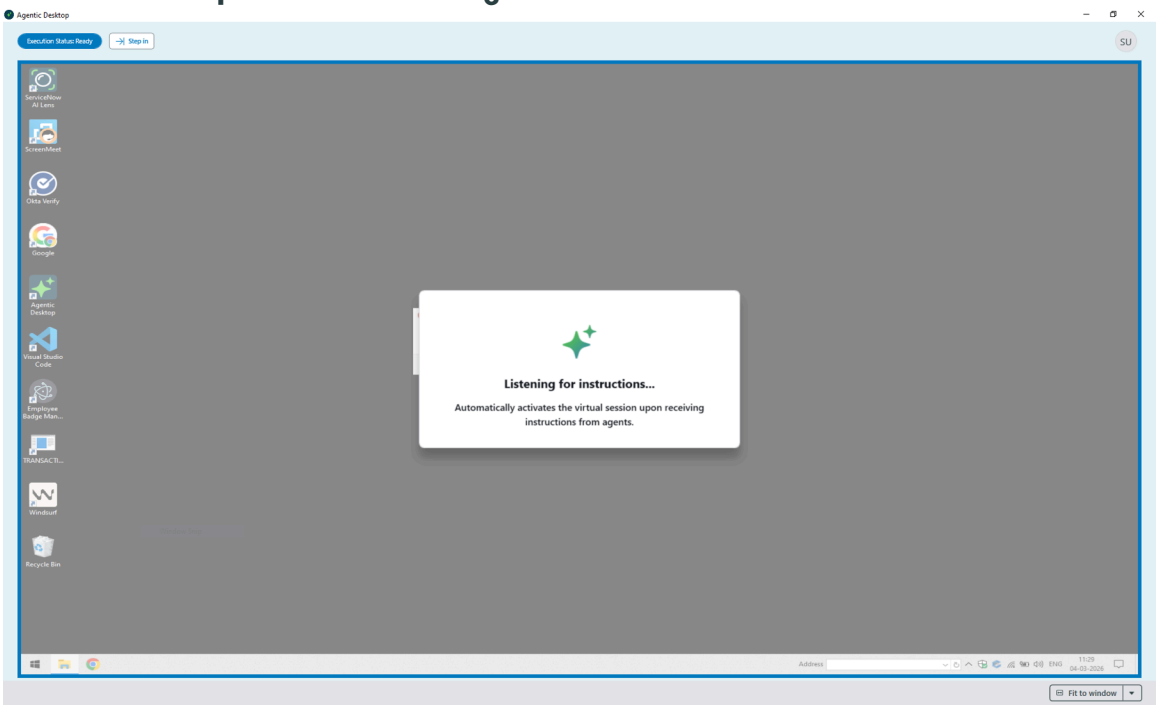
This session acts as a virtual environment where the automations run in isolation. While desktop actions run in the Execution workspace, you can continue working on other desktop tasks.

Note:

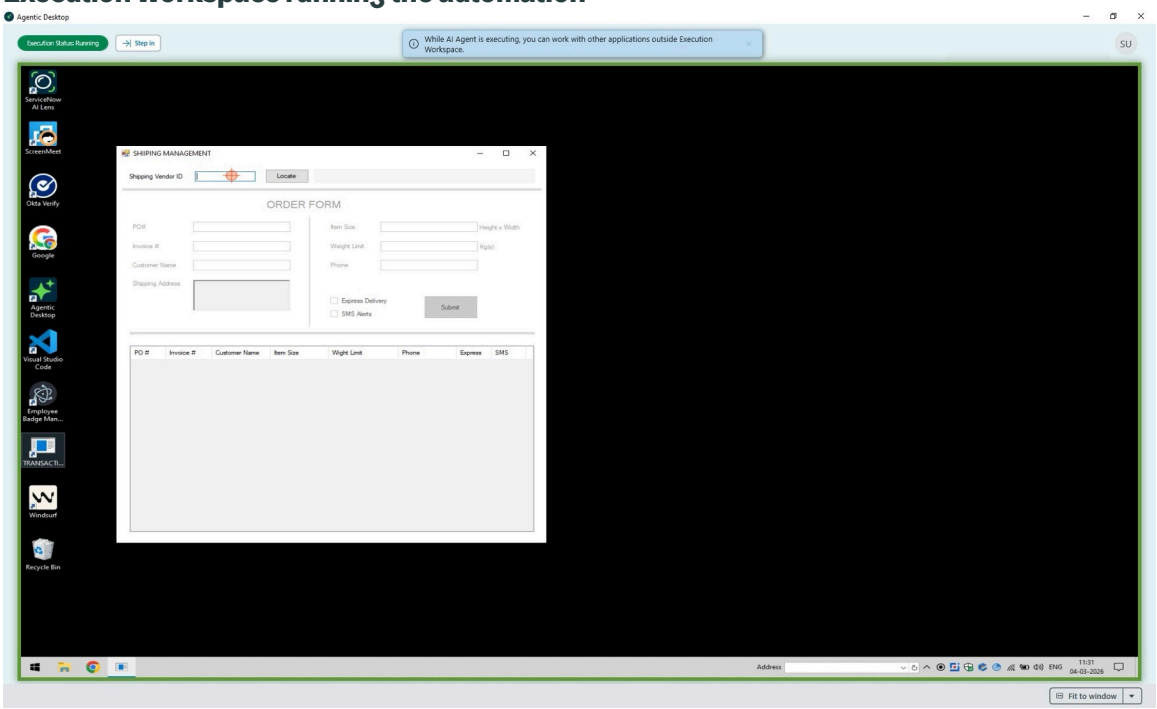
To avoid conflicts, do not run the Agentic Desktop Execution workspace and RPA Attended Desktop mode at the same time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Execution workspace before executing automation



Execution workspace running the automation



Execution workspace features

Execution workspace offers the following features:

Desktop-in-Desktop (DiD) mode

A virtual environment where your automation executes. You can monitor the execution of your desktop actions and how they interact with desktop applications.

Step in/Step out

You can manually take control to provide an input by selecting **Step in**. Once done, you can leave the control by selecting **Step out**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If your automation requires manual inputs, such as entering an OTP or CAPTCHA, you must provide instructions the AI Agent to wait for user input during execution. Otherwise, the automation cannot proceed.

Single execution

The system supports one automation execution at a time within the workspace.

Smart sizing

Smart Sizing enables your desktop sessions automatically adapt to your display. This ensures each captured desktop session is fully visible, readable, and usable, regardless of the device resolution and scaling.

Fit to window: Automatically scales the desktop session to fit within the Execution workspace while keeping it fully visible and readable.

Original resolution: Displays the desktop session at its original resolution. Scroll bars appear if the desktop session is larger than the Execution workspace.

Execution statuses

Execution workspace shows the following statuses throughout the execution process.

Execution status	Description
Ready	Automation execution is ready to be initiated. The workspace is waiting for instructions from AI Agent Studio.
Initiated	Automation execution is initiated. The workspace has received the execution plan from AI Agent Studio. **Note:** When you launch Execution workspace for the first time, it takes longer to move to the Running state because it is setting things up for you.
Running	Automation execution is running. The workspace executes the steps defined in the desktop actions.
Success	Automation execution is completed successfully. The workspace executed the steps defined in the desktop actions and the outcome is shown on the Now Assist panel, such as the automation execution completed successfully and incident created successfully.
Failed	Automation execution is failed. The workspace couldn't execute the automation due to an error.
Canceled	Automation execution is canceled. The user canceled the execution manually.

Related topics

Example: Use AI agents to process badge-related requests automatically

Example: Use AI agents to automatically enter data into the shipping management app

Supporting information for Agentic Desktop

Get a quick overview of the important information that is related to the Agentic Desktop application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Supported versions

Agentic Desktop is supported from Zurich Patch 4.

Supported user interfaces

To design and run AI-powered desktop actions using Now Assist AI agents, install the Agentic Desktop application. For more information, see Configure Agentic Desktop.

Licensing requirements

The Agentic Desktop application requires a ServiceNow Pro Plus or Enterprise Plus license.

Application information

This app has the following dependencies:

- Agentic Desktop Core (sn_desktop_core)
- Now Assist AI agents

Configure Agentic Desktop

You can enable the Agentic Desktop application (sn_desktop_agents) if you have the admin role.

Before you begin
To get started with Agentic Desktop, you must have:

- A ServiceNow Pro Plus or Enterprise Plus license.
- An instance on Yokohama Patch 11+ or Zurich Patch 4+.
- A .NET 9.0 runtime v9.0.10 and .NET 9 Desktop Runtime v9.0.10 is installed.

Role required: administrator

About this task
Agentic Desktop isn't a standalone application that you can install directly. To enable Agentic Desktop on your instance, you must install other Now Assist applications, such as Now Assist for IT Service Management (ITSM) or Now Assist for Customer Service Management (CSM).

- Review the Agentic Desktop application listing in ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.
- Perform these steps in your ServiceNow instance.

By enabling the Agentic Desktop application (sn_desktop_agents), you can view the Agentic Desktop and download modules in your ServiceNow instance.

The following items are installed with Agentic Desktop:

- Roles
- Tables

For more information, see Components installed with Agentic Desktop.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Definition** > **Plugins**.
2. Search for and select a Now Assist application, such as Now Assist for IT Service Management (ITSM) or Now Assist for Platform.
3. Select **Install**.

What to do next

Download and install the Agentic Desktop installer on your system to automate repetitive tasks in your desktop environment. For more information, see Download Agentic Desktop installer.

Download Agentic Desktop installer

Download the Agentic Desktop installer so that you can install Agentic Desktop on your Windows machine for designing and running desktop actions.

Before you begin

- A .NET 9.0 runtime v9.0.10 and .NET 9 Desktop Runtime v9.0.10 is installed.
- Ensure that you install the Agentic Desktop installer using Windows administrator privileges.
- Ensure that the Desktop-in-Desktop settings are configured correctly.
- Add the end users who interact with the Execution workspace of Agentic Desktop to the Remote Desktop Users group on the target machine and provide Remote Desktop access permissions for seamless automation execution.

  If your organization uses Group Policy, add the end users to a Microsoft Active Directory group that is permitted to use Remote Desktop through Group Policy on each target machine where desktop actions run.

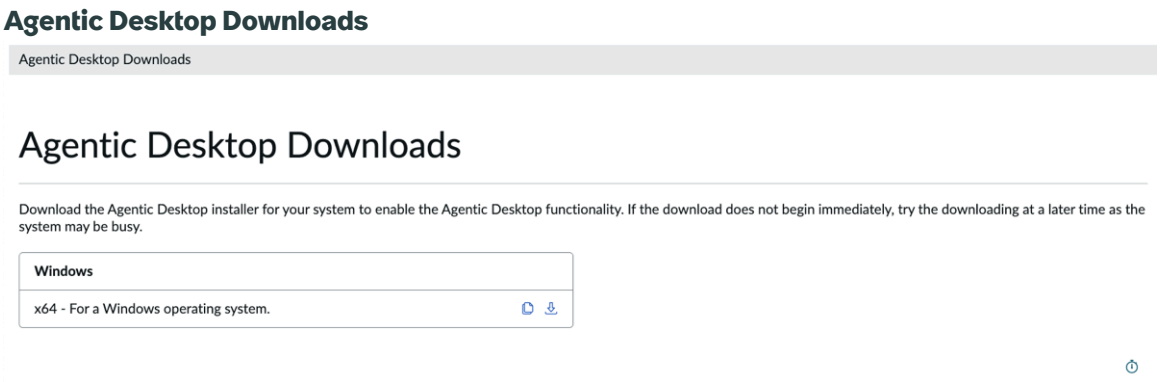
  - Local changes to the Remote Desktop Users group are temporary unless they align with Microsoft Active Directory entitlements.
  - If the user is not entitled, Group Policy refresh automatically removes them from the group.

Role required: Windows administrator

Procedure

1. Navigate to **All** > **Agentic Desktop** > **Downloads**.
2. In the Agentic Desktop Downloads page, do any of the following actions to download the required application for Windows:
   - Select the download icon ( ).
   - Select the copy link ( ). In a browser, right-click and select the **paste and go** option.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

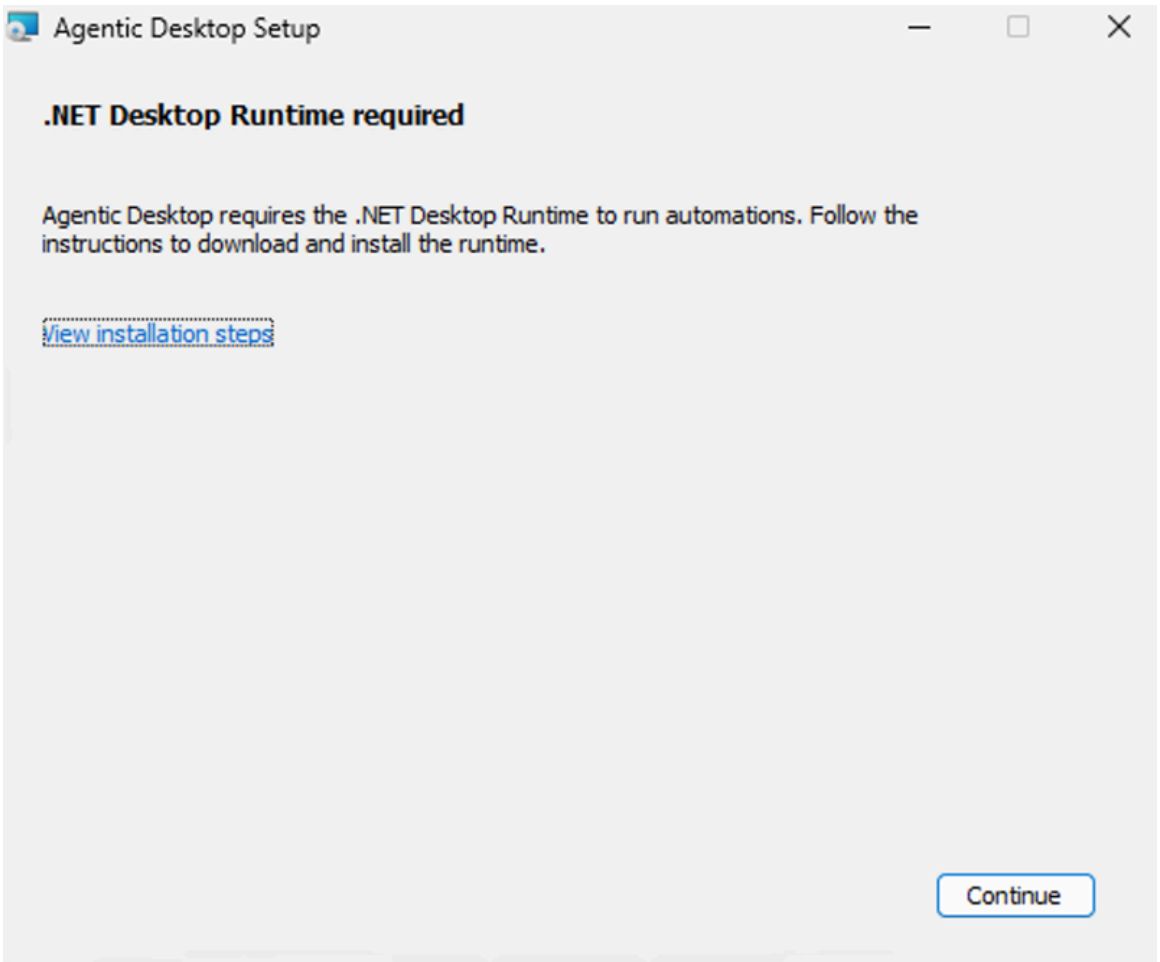


The system might prompt you to save or open the file.

Note: Depending on your browser setting, the browser might automatically save the file to your Downloads folder.

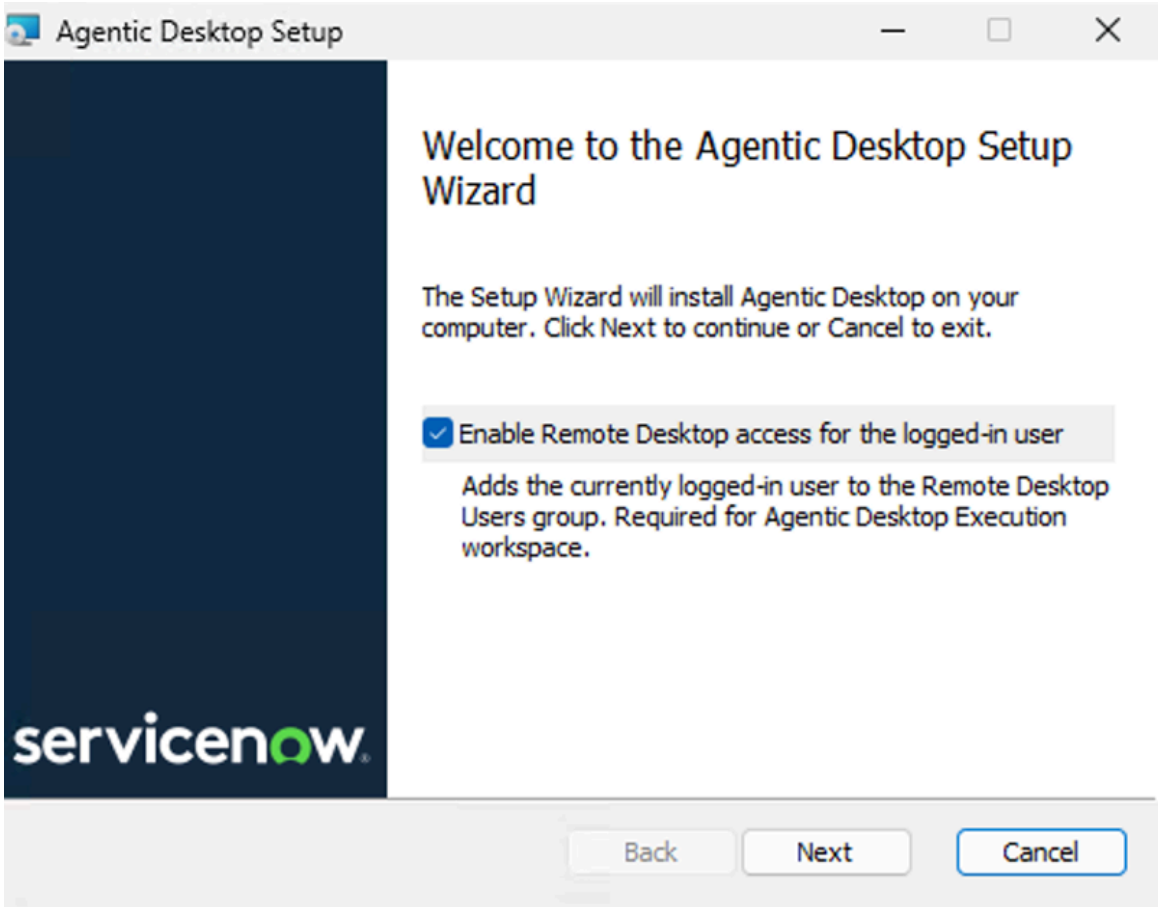
3. Execute the application installation file (.msi file).
4. Follow the on-screen instructions to install the application.
5. If not already installed, download and install the .Net Desktop Runtime framework.

   The installer validates if required .NET Desktop Runtime is installed or not on the user machine. If not installed, the installer provides a link to instructions for downloading and installing the same. For more information, see Download and install .Net Desktop Runtime for Agentic Desktop.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Keep the option **Enable Remote Desktop access for the logged-in user** selected to automatically add the logged-in user to the Remote Desktop Users group.



Trouble?

If Agentic Desktop is not installed using Windows administrator privileges, you may encounter issues with the desktop session in the Execution workspace. To fix the issue, perform the following steps.

1. Navigate to the **SessionRegistry** folder of the Agentic Desktop installation directory.
2. Right-click **SN.AGENTICDESKTOP.SESSION.REGISTER.exe** and select **Run as administrator**.
3. Restart your system.

What to do next

Use Agentic Desktop to design desktop actions. For more information, see Desktop actions in Agentic Desktop.

Download and install .Net Desktop Runtime for Agentic Desktop

Reduce setup time and prevent installation errors by downloading and installing .Net Desktop Runtime following the instructions.

Before you begin

Ensure that you have installed the Agentic Desktop MSI. For more information, see Download Agentic Desktop installer.

Role required: admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

The Agentic Desktop installer guides you through installation of .NET Desktop Runtime during the setup.

Procedure

1. Copy the following script to your clipboard.

```
Write-Host "--- Initializing Agentic Environment Setup ---"
 -ForegroundColor Cyan
$ErrorActionPreference = "Stop"

$downloadUrl = "https://dot.net/v1/dotnet-install.ps1"
$workingDir = $PSScriptRoot
$scriptPath = Join-Path $workingDir "dotnet-install.ps1"

try {
    Write-Host "[1/3] Downloading official Microsoft
 installer..." -ForegroundColor Yellow
    Invoke-WebRequest -Uri $downloadUrl -OutFile $scriptPath

    Write-Host "[2/3] Installing .NET 9.0.10..."
 -ForegroundColor Yellow
    Set-Location $workingDir

    & .\dotnet-install.ps1 `
        -Architecture x64 `
        -InstallDir "C:\Program Files\dotnet\" `
        -Runtime windowsdesktop `
        -Version 9.0.10

    & .\dotnet-install.ps1 `
        -Architecture x64 `
        -InstallDir "C:\Program Files\dotnet\" `
        -Runtime dotnet `
        -Version 9.0.10

    Write-Host ""
    Write-Host "--- SUCCESS: System is ready for Agentic
 Desktop ---" -ForegroundColor Green
}
catch {
    Write-Host ""
    Write-Host "--- ERROR: Installation failed ---"
 -ForegroundColor Red
    Write-Host $_.Exception.Message -ForegroundColor Red
}

Write-Host ""
Write-Host "Press any key to exit..."
$null = $Host.UI.RawUI.ReadKey("NoEcho,IncludeKeyDown")
```

2. Paste the script in any text editor.
3. Save the file with a `.ps1` extension.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Select **File** > **Save As**.

b. In the **File name** field, enter a name and save the file with a `.ps1` extension.
Example: `myscript.ps1`

c. Set **Save as type** to **All Files (.)**.

d. Choose a location such as Documents or Desktop.

e. Select **Save**.

4. Run the script in Windows PowerShell.

a. Select **Start**.

b. Search for **Windows PowerShell**.

c. Select **Run as administrator**.

d. Use the `cd` command to move to the folder where you saved the script.
Example:

```
cd C:\Users\YourName\Documents
```

e. Run the script by entering the following command.
Example:

```
.\myscript.ps1
```

Result
.NET Desktop Runtime is downloaded and installed on your system.

What to do next
Use Agentic Desktop to design desktop actions. For more information, see Desktop actions in Agentic Desktop.

Enable AI agents to securely access SSH parameters in Agentic Desktop

Enable AI agents to securely access credentials, such as usernames and passwords, through Desktop Action Parameter records. This approach protects sensitive values during desktop action execution without exposing them in agent instructions.

Note: Currently, Parameter and Parameter Value records are only supported for SSH connector, non-UI block desktop actions.

Only the users with the sn_aia.admin role can create Parameter records for various parameters that the AI agent securely accesses during desktop action execution. To store both a username and a password, the AIA admin must create two separate Parameter records, one for the username "un_username_group" and one for the password "un_password_group".

Then, users with the sn_aia.admin or now_assist_panel_user role can create Parameter Value records under the Parameter record to store the values. For example, under the "un_username_group" Parameter record, users can create a Parameter Value record to store usernames and under the "un_password_group" Parameter record, users can create a Parameter

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Value record to store passwords. Only one Parameter Value record can be created per user for each Parameter record. If multiple users need to trigger the same AI agent, each user must create a Parameter Value record for their own credentials.

You can create Parameter records only for username and password. For any other SSH parameters, provide the values either in the agent instructions or in your instructions in the Now Assist panel.

Example: AI Agent instructions during execution

The following example shows how an AI agent instruction can reference stored parameter names:

```
Connect to SSH server and retrieve my session info. Here are my details:
```

- `IP address: 172.27.50.123`
- `Port: 22`
- `Retrieve the user name stored in "un_username_group" and the password stored in "un_password_group" parameter records.`

Note: When triggering an AI agent, explicitly specify in your instructions whether the credentials are provided directly or stored in Parameter records. If Parameter records are used, clarify that the record names are for reference only and that the agent must retrieve the username and password from those records.

Ensure that you use the exact names of the Parameter records. Parameter record names are case sensitive. For example, "UserName" and "username" are treated as different values.

Create Parameter records for SSH credentials

Create a Parameter record with a name that you can reference in your instructions so the AI agent can securely sign in to the SSH server and execute commands.

Before you begin

Perform this task in the ServiceNow instance.

Ensure that you have an active SSH server.

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > **Agentic Desktop** > **Desktop Action Parameters**.
2. Select **New** to create a Parameter record.
3. Enter the name and description.

   Example
   For example, un_username_group or un_password_group
4. Select **Submit**.

Store username and password values in the Parameter Value records

Create Parameter Value records for storing the actual username and password values so that AI agents can securely connect to the SSH server.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Perform this task in the ServiceNow instance.

Role required: sn_aia.admin or now_assist_panel_user

Procedure

1. Navigate to **All** > **Agentic Desktop** > **Desktop Action Parameters**.
2. Select a Parameter record from the list.
3. Create a parameter value record.

   a. In the Desktop action parameter values related list, select **New**.

   b. Fill in the following fields.

Field	Description
Name	Unique name for this Parameter Value record that stores user name or password.
User	Read-only. Name of the user who creates this record. **Note:** Only user with the sn_aia.admin role can edit this field and assign this record to any other user.
Is sensitive	Option to encrypt the value
Value	Value such as user name or password for connecting to the SSH server.

   c. Select **Submit**.

Desktop actions in Agentic Desktop

Desktop actions enable you to automate repetitive tasks on your desktop and web applications. This capability helps you streamline repetitive tasks, improve efficiency, and integrate desktop application workflows into your ServiceNow processes.

Types of desktop actions

AI agents use these desktop actions to interact with desktop applications, perform UI-based tasks, and automate end-to-end workflows. There are two types of desktop actions that can be added as tools to AI agents:

UI blocks

The UI block actions help you simulate humans interacting with UI elements on your thick client applications, legacy systems, or SaaS applications without APIs. These actions include clicking buttons, typing into text boxes, selecting from dropdown menus, and more. These desktop actions make it easier to automate repetitive tasks. They encapsulate repeatable UI interactions, such as screens, anchors, and steps. You can create, manage, and test your desktop actions using the no-code Design workspace.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example: Filling in fields on a payroll application, updating inventory in a point-of-sale system, or submitting a request through a desktop insurance claims app.

You can access and view the existing desktop actions in the ServiceNow instance by navigating to **All** > **Agentic Desktop** > **Desktop Actions**. You can see the details of related desktop action applications, AI agents, and desktop action execution records for UI block, and AI agents and desktop action execution records for non-UI block desktop actions.

Non-UI block

The non-UI block actions include pre-built connectors that enable your AI agents to interact with various applications and system components in the background. These connectors streamline automation by offering actions for common tasks, reducing the need for complex scripting. Each connector focuses on a specific application or system area, providing a collection of related methods. You can't create non-UI block actions.

Example: Reading data from Microsoft Excel, extracting emails from Microsoft Outlook, generating a PDF, compressing files into a ZIP, fetching records from a database, or sending a system notification.

Non-UI block actions are supported for the following applications:

Note: For the AI agents to perform non-UI block automations, the following applications must already be installed on the end-user's system.

Application	Description
Microsoft Excel	Enables AI agents to perform standard actions on the Microsoft Excel documents. For example, data manipulation, content modification, and information retrieval from spreadsheets. The following methods are supported: • ReadData • WriteData • FindAndReplace **Note:** CSV and password-protected files aren't supported.
Microsoft Outlook	Enables AI agents to perform standard actions on the Microsoft Outlook application. For example, you can automate sending or replying to emails. The following methods are supported: • DeleteMail • ForwardMail • GetMailsFromFolder • MoveMail • Reply • SaveAllAttachments

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	Description
	• SendMail • GetMail • MarkAsRead • MarkAsUnread **Note:** Only the classic view is supported. Shared inboxes aren't supported.
Microsoft Word	Enables AI agents to perform standard actions on the Microsoft Word documents. For example, you can replace text in a document. The following methods are supported: • GetText • InsertText • ReplaceText **Note:** Password-protected files aren't supported.
PDF	Enables AI agents to perform standard actions on the PDF documents. For example, extracting text, converting to Excel or Images, and merging and spiting files. The following methods are supported: • GetText • ConvertToExcel • ConvertToImages • Merge • Split **Note:** Password-protected files aren't supported.
PowerShell	Enables AI agents to execute PowerShell commands and scripts and returns the results. The following methods are supported: • InvokeCommand • InvokeScript **Note:** Each PowerShell step runs in a new session. Output from one step isn't carried over to the next step.
SQL	Enables AI agents to execute custom SQL queries against supported databases (MSSQL, ORACLE, MYSQL, OLEDB, ODBC) and returns results as JSON.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Application	Description
	The following method is supported: ExecuteQuery
Secure Shell (SSH)	Enables AI agents to securely connect to remote servers to run SSH commands and scripts; and returns the execution results. This connector supports non-interactive commands. The following method is supported: RunCommand **Note:** Each SSH step runs in a new session. Output from one step isn't carried over to the next step.
System Actions	Enables AI agents to perform standard Windows system operations. For example, starting an app, creating a ZIP file, or deleting any file or folder. The following methods are supported: • StartApp • Terminate • DateTimeNow • SetEnvironmentVariable • GetEnvironmentVariable • CopyFileOrFolder • DeleteFileOrFolder • WriteToFile • ReadFromFile • GetFilesFromFolder • ExtractFile • CreateZip

Ways to create desktop actions

Note: If your automation requires manual inputs, such as entering an OTP or CAPTCHA, you must provide instructions to the AI Agent to wait for user input during execution. Otherwise, the automation can't proceed.

You can create desktop actions of type UI block in the following ways:

Auto-capture screens and steps

With Action recorder, you can record your interactions with desktop applications to create automated workflows. It records every step you take, including clicks, keystrokes, and data entry, along with visual and contextual information. By recording steps, you can automate tasks that replicate your interactions. You can

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

save the recorded screens and steps as a reusable desktop action. For more information, see Action recorder in Agentic Desktop.

Manually capture screens and steps

You can manually take a screen capture, add anchors and steps to automate a series of steps you perform on your computer, such as clicking buttons, typing text, or interacting with different applications, and then save this sequence as a reusable desktop action.

Benefits of desktop actions

- Eliminate manual effort by automating repetitive tasks for frequent desktop activities.
- Execute complex sequences of steps quickly and consistently.
- Connect desktop application processes with your broader ServiceNow automations.
- Standardize task execution to minimize human error.
- You can record steps directly on your desktop applications.
- Automate tasks that span across multiple desktop applications.
- Save recorded sequences as desktop actions to use them in various automations.
- Review, rearrange, edit, or delete individual steps within a recorded action.

How the Design workspace works

Record steps

Capture your interactions with desktop applications, including clicks, text input, and screen navigation.

Define steps

Specify individual steps like Click, Set Text, or Get Text to interact with elements on your screen.

Application details

Select and store the desktop applications involved in your automations in the **Details** tab.

Organize screens

Structure your recorded steps by application screens, making it easy to review and edit.

Save and publish

Store your desktop actions in your ServiceNow instance for reuse and deployment.

What you can automate

Desktop actions support a wide range of desktop application interactions, including:

- Entering data into forms.
- Extracting information from applications.
- Navigating through application menus and screens.
- Performing steps like saving files or generating reports.
- Transferring data between different desktop applications.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- Agentic Desktop Design workspace
- Action recorder in Agentic Desktop
- Automate repetitive tasks by auto-capturing steps in Agentic Desktop
- Automate repetitive tasks by manually capturing steps in Agentic Desktop
- Add details to desktop actions in Agentic Desktop
- Test and activate a desktop action in Agentic Desktop
- Add a desktop action to an AI agent

Automate repetitive tasks by auto-capturing steps in Agentic Desktop

Create desktop actions by auto-capturing steps to automate repetitive tasks in Agentic Desktop. You can save the steps that you perform on the application elements as a reusable desktop action of type UI block.

Before you begin

To access the Agentic Desktop functionality, perform the following steps:

- Enable Agentic Desktop on your ServiceNow instance. For more information, see Configure Agentic Desktop.
- Download the Agentic Desktop installer to automate repetitive tasks across applications and systems. For more information, see Download Agentic Desktop installer.

Confirm that the following system requirements are met:

- Windows 11 operating system is used.
- A .NET 9.0 runtime v9.0.10 and .NET 9 Desktop Runtime v9.0.10 is installed.
- No extended monitors are connected.
- Theme must match between the systems used for recording and execution.

Familiarize yourself with the Design workspace and Action recorder. For more information, see Agentic Desktop Design workspace and Action recorder in Agentic Desktop.

Role required: sn_aia.admin

About this task

You can record steps that you want to automate for desktop and web applications and save them as desktop actions of type UI block that AI agents can simulate while executing automations.

With Action recorder, you can record every step you perform on desktop applications, including clicks, keystrokes, and data entry, along with visual and contextual information. This detailed recording of steps allows automated tasks to reliably replicate your interactions.

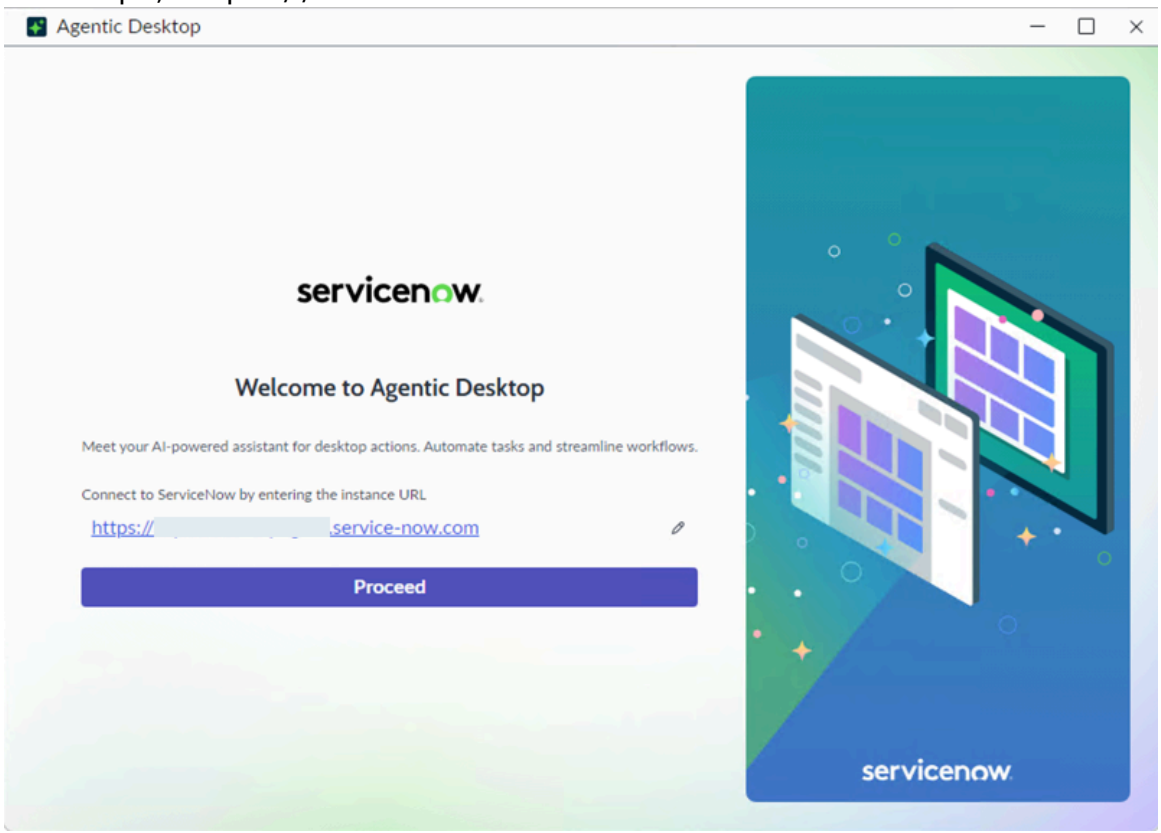
Note: If your automation requires manual inputs, such as entering an OTP or CAPTCHA, you must provide instructions to the AI Agent to wait for the user input during execution. Otherwise, the automation can't proceed.

Procedure

1. From your Windows system, launch the Agentic Desktop application.
2. On the login page, in the **Add ServiceNow URL** field, enter the ServiceNow instance URL.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example
For example, `https://<instance name>.service-now.com`.

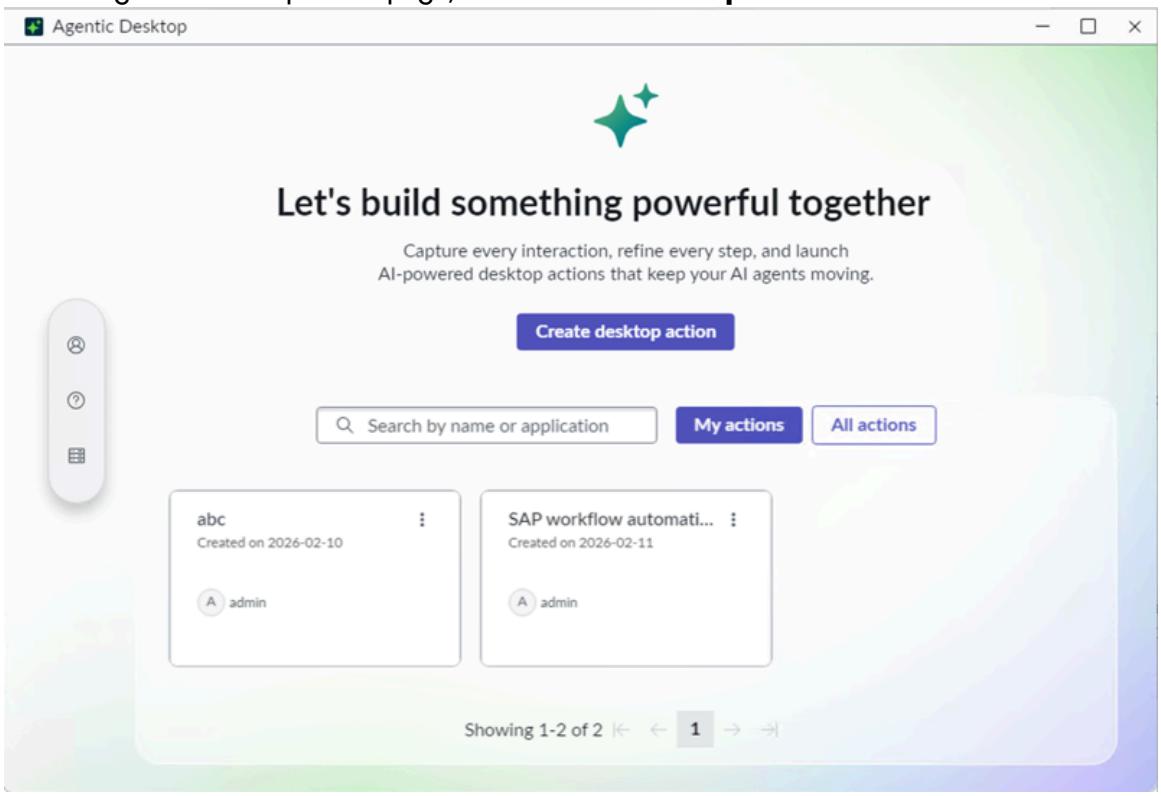


3. Select **Proceed**.
4. Log in to your ServiceNow account by entering your user name and password. Your account must have the sn_aia.admin role.

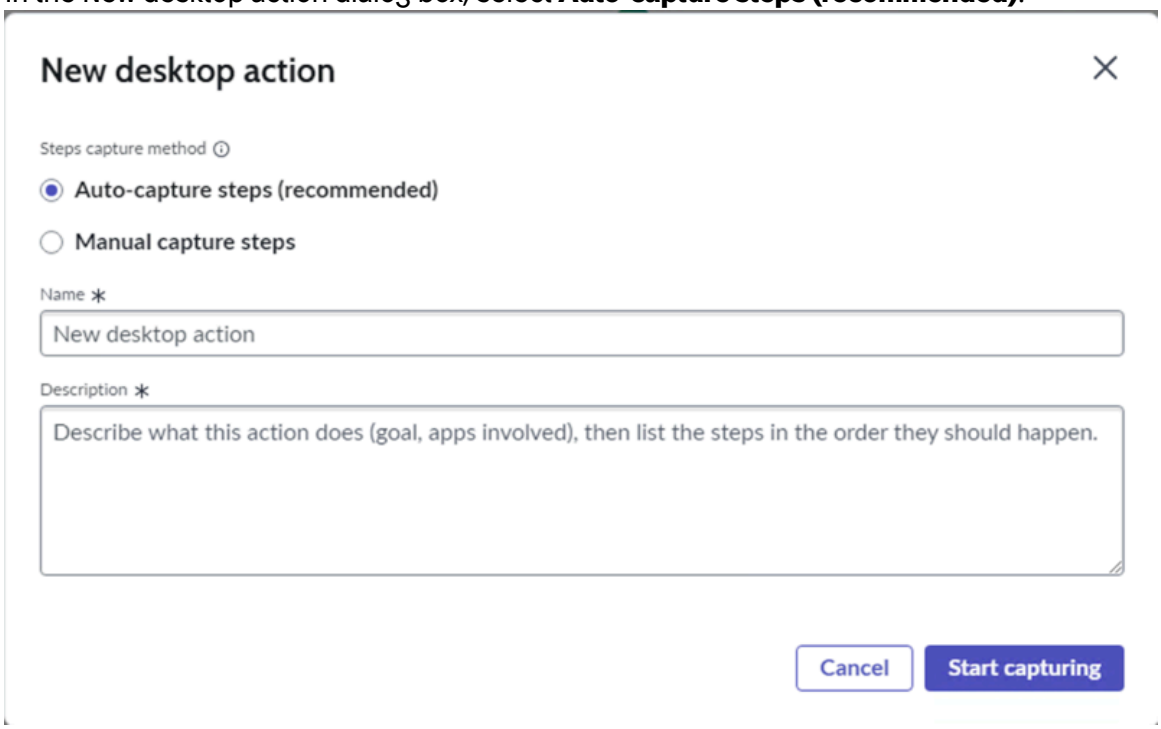


© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. On the Agentic Desktop home page, select **Create desktop action**.



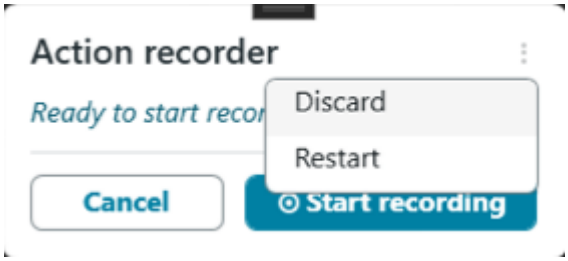
6. In the New desktop action dialog box, select **Auto-capture steps (recommended)**.



7. Enter a name and description for the desktop action.
8. Select **Start capturing**.

   The Agentic Desktop window is minimized and the Action recorder panel is launched. You can freely drag and reposition the Action recorder panel anywhere on your desktop screen.

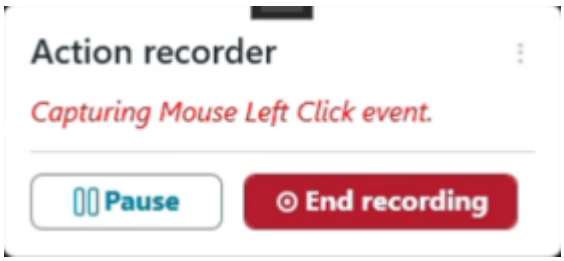
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



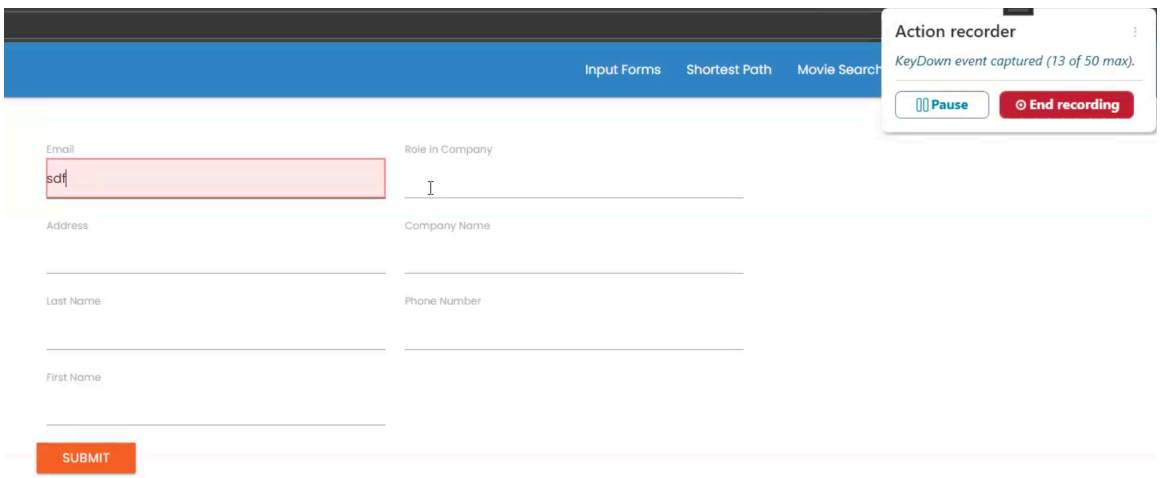
9. Open the applications that you want to record steps for.
10. From the Action recorder panel, select **Start recording**.
    You'll see a "Recording started" message on the Action recorder panel. You can select any of the following options when needed from the **More options** menu:
    - **Pause**: Skip recording steps
    - **Restart**: Restart recording the steps
    - **Discard**: Discard the recording if it doesn't meet your needs
11. Perform the steps that you want to automate.

    Each step that you perform is captured sequentially and the type of UI action is displayed for each step. For example, Capturing Mouse Left Click event.

    Avoid performing steps too quickly while recording because the recorder may miss capturing steps.
    - When the message in the Action recorder panel appears in red, wait for the recorder to finish processing the current step before continuing.



    - Proceed with capturing the next step when the red highlight box appears around the element you interact with and the message in the Action recorder panel changes to blue.



12. After you're done with all the steps, select **End recording** on the Action recorder panel.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can capture maximum of 50 steps using the recorder in a recording session. While auto-capturing steps, a counter displays the remaining number of steps you can capture using the recorder (for example, “35 of 50 max”). Recording stops automatically after you capture 50 steps. If you need to capture additional steps, start a new recording session. The new recording adds screens and steps to those captured in the previous recording.

You’ll see the "Drafting recorded action(s)" message on the Action recorder panel.

The recorded steps are displayed as screenshots in the Design workspace with anchors and steps automatically added.

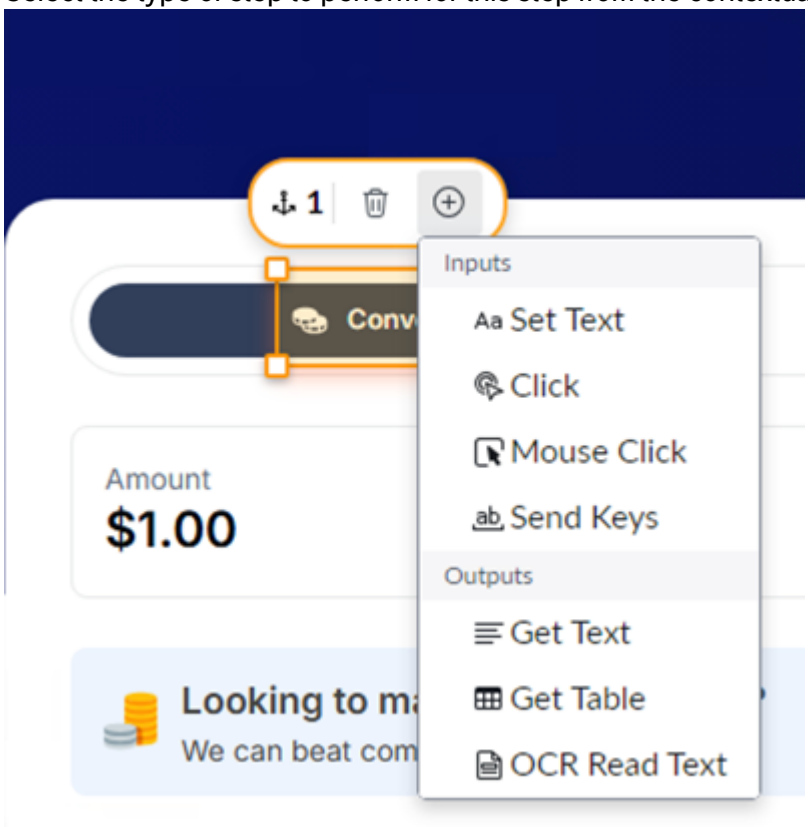
13. Review the sequence of captured screens and anchors.
    You can adjust the sequence by dragging the screens and anchors in the Screens and steps panel.
14. Review the location of anchors and adjust as necessary.

    Note: The auto-anchoring accuracy depends on the application's accessibility, performance, and UI complexity.

15. Review and adjust steps.
    You can modify the sequence by adding, removing, or dragging steps.

    a. From the Anchor control menu, select the **Add step** icon ⊕.

    b. Select the type of step to perform for this step from the contextual menu.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Description of the actions

Goal	Step	Type	Example
Enter text in a field	Set Text	Input	Enter any text data such as a user name, an address, a survey response, or in any situation where text entry is accepted. **Note:** If you set a static value for this field, the automation uses it during execution and doesn't prompt you for input from the Now Assist panel.
Simulate a mouse click	Click	Input	Click a button, open a menu, or perform any step typically performed by a mouse click.
Simulate an alternative mouse action (for example, right-click, drag, scroll, or paste)	Mouse Click	Input	Perform various mouse device actions, such as right-click and select an object or scroll on a web page.
Simulate a key press or a key function	Send Keys	Input	Perform keyboard shortcuts, such as copying text by entering `Ctrl + C` on fields and elements. **Note:** If you set a static value for this field, the automation uses it during execution and doesn't prompt you for input from the Now Assist panel.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Goal	Step	Type	Example
Capture text from a window or web page	Get Text	Output	Receive text from the source area.
Capture a table	Get Table	Output	Receive table from the source area when the text is in the table format. **Note:** For the step to capture table data successfully, the data must already be in the table form. The step can't convert ordinary text to table data.
Read text from an image	OCR Read Text	Output	Recognize text from images and return it in the standard text format.

You can add multiple steps representing your automation steps.

16. Configure the properties for added screens, anchors, and steps.
For more information, see Screen, anchor, and step properties in Agentic Desktop.
17. Provide names for all added screens, anchors, and steps.

    When you create these elements, you can edit the auto-generated name, but follow these naming guidelines.

    - Name fields must not be empty.
    - Name fields must contain only alphanumeric characters. Spaces and special characters are not permitted.
    - Each name must be unique at its respective parent level.
      - Each screen must have a unique name at the desktop-action level.
      - Each anchor must have a unique name at the screen level.
      - Each step must have a unique name at the anchor level.

What to do next

1. Configure the details of your desktop action. For more information, see Add details to desktop actions in Agentic Desktop.
2. Test and activate the desktop action so that it can be added as a tool to AI agents. For more information, see Test and activate a desktop action in Agentic Desktop.
3. Add the desktop action as a tool to AI agents in AI Agent Studio. For more information, see Add a desktop action to an AI agent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- Agentic Desktop Design workspace
- Action recorder in Agentic Desktop
- Screen, anchor, and step properties in Agentic Desktop
- Automate repetitive tasks by manually capturing steps in Agentic Desktop
- Add a desktop action to an AI agent
- Examples of creating desktop actions
- Examples of executing desktop actions using AI agents

Automate repetitive tasks by manually capturing steps in Agentic Desktop

Create desktop actions by manually capturing steps to automate repetitive tasks in Agentic Desktop. The steps that you define on one or more application screens are saved as a reusable desktop action of type UI block.

Before you begin

To access the Agentic Desktop functionality, perform the following steps:

- Enable Agentic Desktop on your ServiceNow instance. For more information, see Configure Agentic Desktop.
- Download the Agentic Desktop installer to automate repetitive tasks across applications and systems. For more information, see Download Agentic Desktop installer.

Confirm that the following system requirements are met:

- Windows 11 operating system is used.
- A .NET 9.0 runtime v9.0.10 and .NET 9 Desktop Runtime v9.0.10 is installed.
- No extended monitors are connected.
- Theme must match between the systems used for recording and execution.
- Remote Desktop must be enabled on your machine and your account must be granted Remote Desktop access permissions before you start using the Agentic Desktop Execution workspace.

Familiarize yourself with the Design workspace and Action recorder. For more information, see Agentic Desktop Design workspace and Action recorder in Agentic Desktop.

Role required: sn_aia.admin

About this task

You can simulate a user interaction in an automation by manually capturing screens and defining steps.

Using the controls in the window, you can capture an area of an application window. You can then set one or more anchors, and define steps that represent the user interactions in that window.

- **Anchor**: Anchors help specify the target area for the interaction by defining a static area from which steps can be defined at a relative distance.
- **Step**: Steps are your sequential interactions of type click, select, type, scroll, and more.

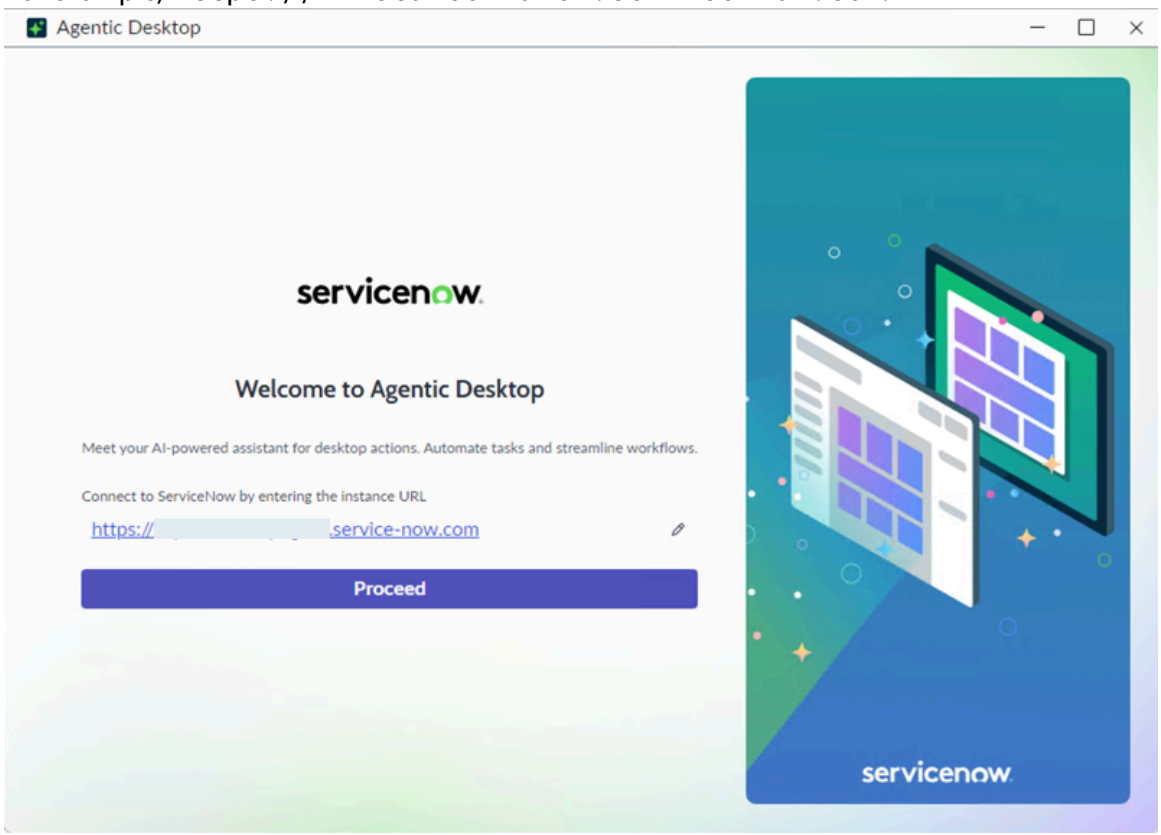
Note: If your automation requires manual inputs, such as entering an OTP or CAPTCHA, you must provide instructions to the AI Agent to wait for the user input during execution. Otherwise, the automation can't proceed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. From your Windows system, launch the Agentic Desktop application.
2. On the login page, in the **Add ServiceNow URL** field, enter the ServiceNow instance URL.

   Example
   For example, `https://<instance name>.service-now.com`.

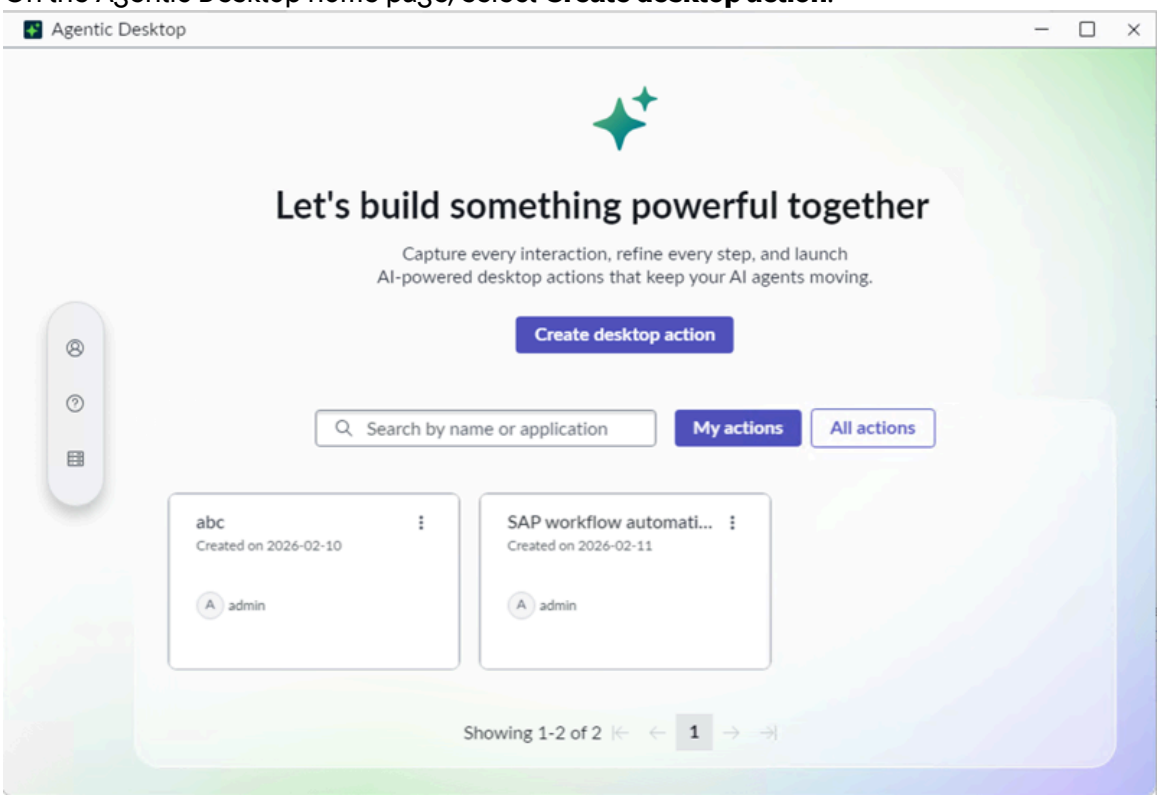


3. Select **Proceed**.
4. Log in to your ServiceNow account by entering your user name and password.
   Your account must have the sn_aia.admin role.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

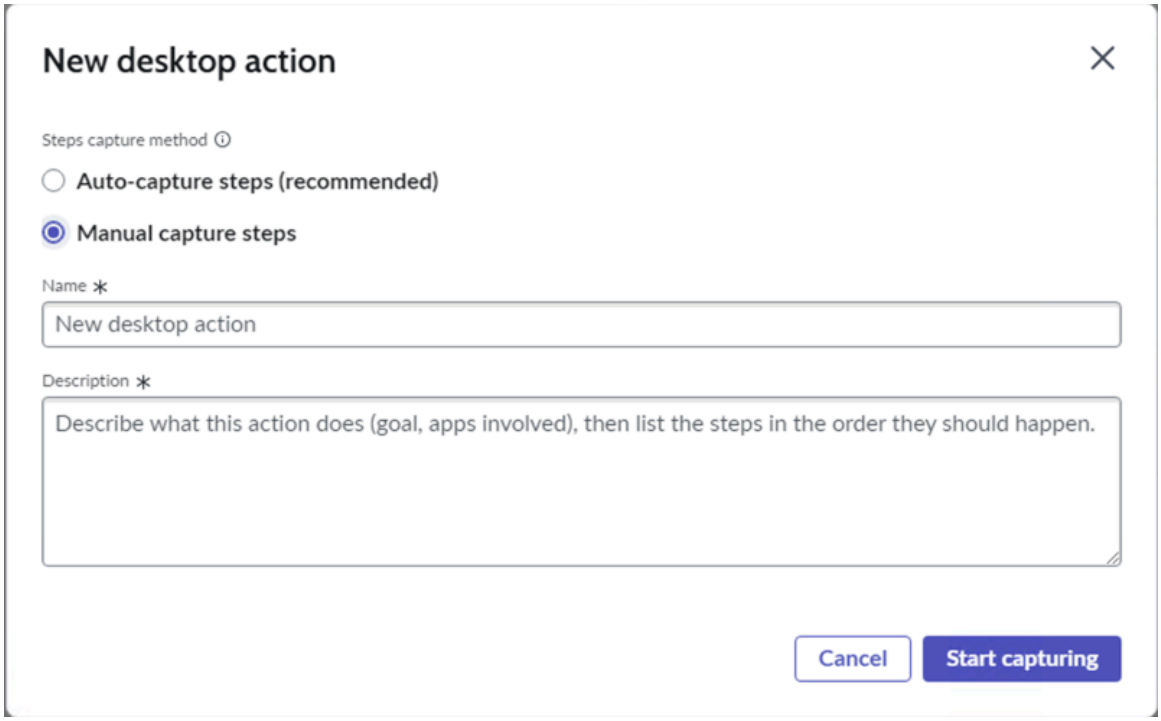


5. On the Agentic Desktop home page, select **Create desktop action**.



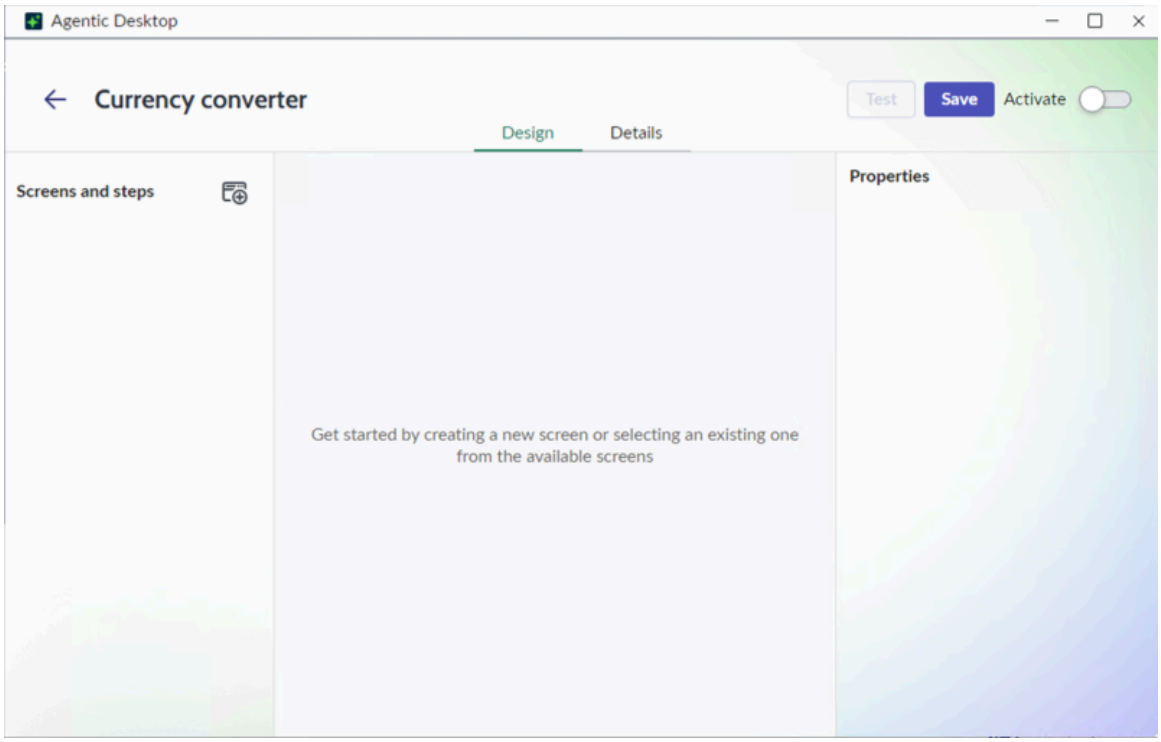
6. In the New desktop action dialog box, select **Manual capture steps**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



7. Enter a name and description for the desktop action.
8. Select **Start capturing**.

   The Design workspace is displayed.

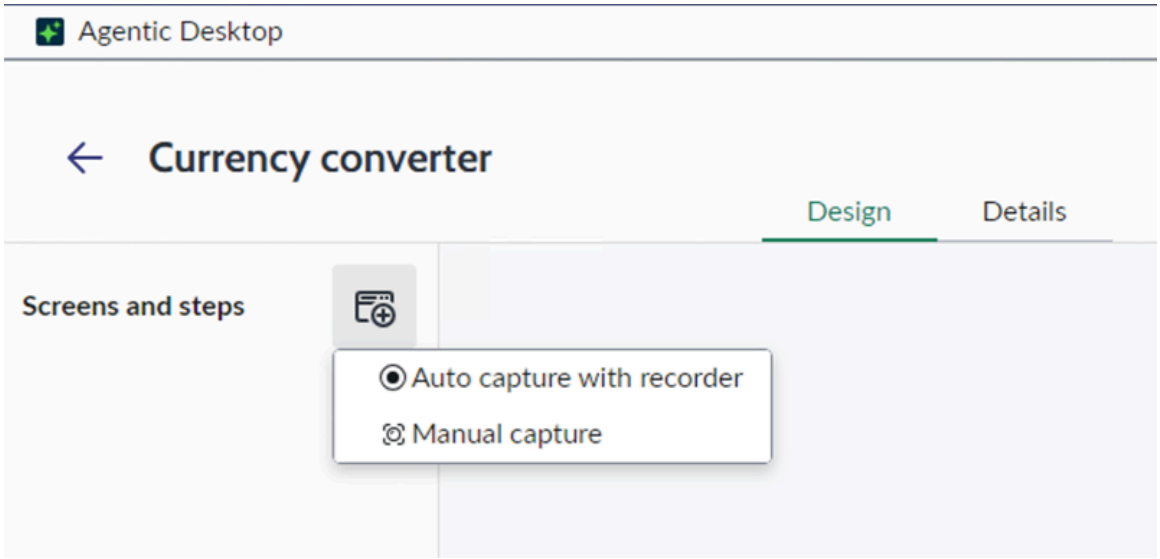


9. Capture screens.

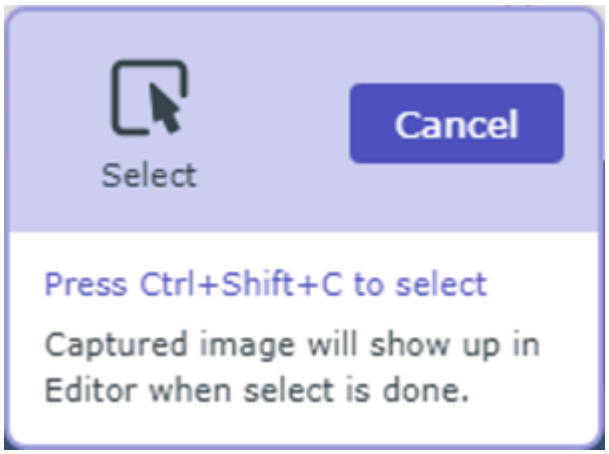
   a. In the Design tab, select the **Capture Options** icon .

   b. Select **Manual capture screens**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



The Agentic Desktop window is minimized and the Capture panel is launched.



c. Open the applications that you want to automate steps for.

d. Capture the area of the external application's window by selecting the **Select** icon on the Capture panel or pressing the `Ctrl + Shift + C` on the keyboard.

> **Note:** If the `Ctrl + Shift + C` shortcut conflicts with another application on the your machine, such as Zoom, then you must use the **Select** button to initiate manual screen capture.

For example, you can capture the area surrounding a button or a text field in a web browser.

The cursor icon changes to the + icon.

e. Drag the + icon and select the required screen area.

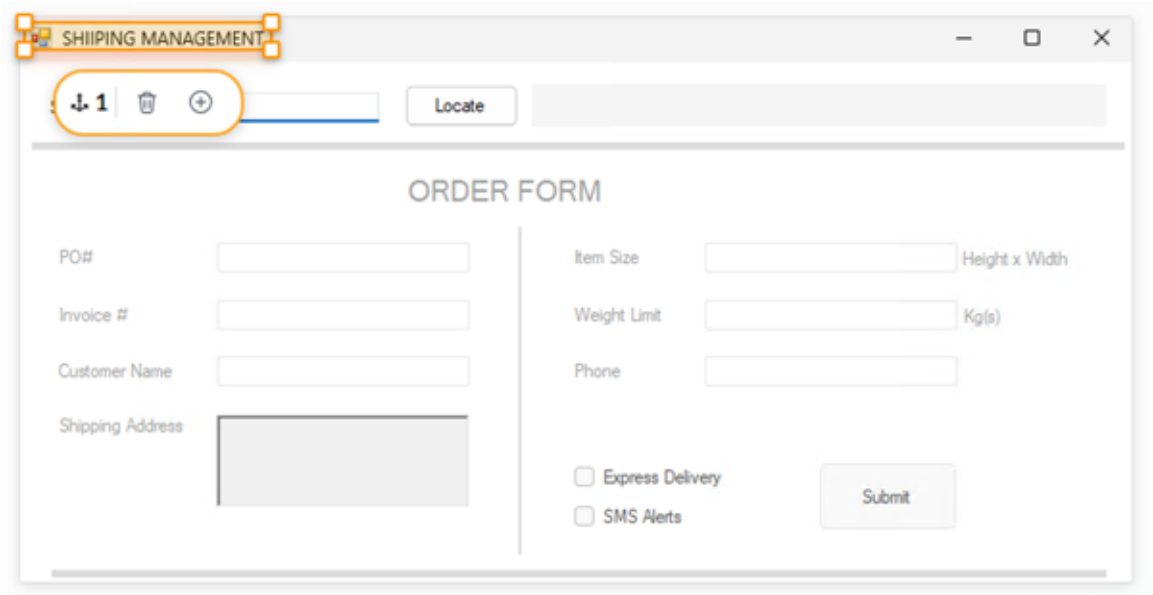
When you leave the + icon, the selected area is captured as a screenshot in the Design workspace.

If you are not satisfied with the captured screen, you can recapture the screen area by selecting the **Capture image** icon .

10. Insert anchors.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Insert an anchor on the captured screen by selecting the **Add anchor** icon .



An anchor is a reference point on the screen that helps the automation identify and interact with a nearby UI elements. During execution, the system uses computer vision to locate the anchor and then identifies the UI elements at a related distance from the anchor. Anchors improve the stability and accuracy of steps when the target element's location may shift or when the UI layout varies across sessions.

Note: Do not use dynamically changing UI elements as anchors. If an element changes its color, text, or state after an action (for example, after a click), select a different anchor that remains static on the screen.

b. Move the anchor to a part of the captured image that won't change.
For example, move the anchor to a title or field label.

If the area under the anchor doesn't exactly match the corresponding area of the captured image, the anchor isn't recognized, and the steps are not performed as intended. Choose a static area of the image for setting your anchor.

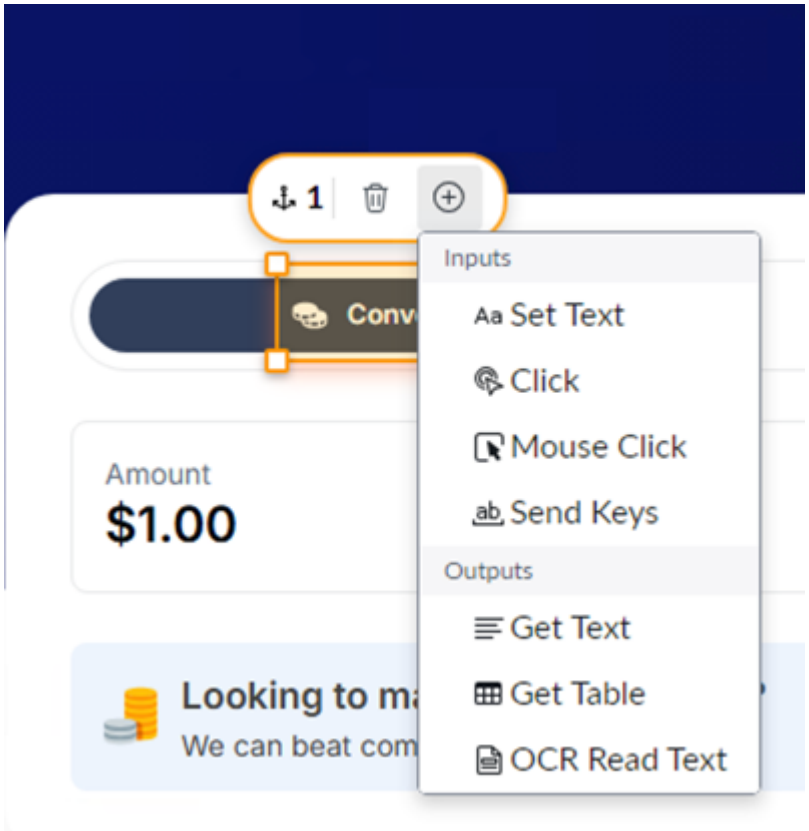
You can add multiple anchors on each screen. Multiple anchors let you define the geographical relationship between anchor and target with greater accuracy when targeting different locations in the image.

11. Configure the steps.

a. From the Anchor control menu, select the **Add step** icon ⊕.

b. Select the type of step to perform for this step from the contextual menu.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Description of the actions

Goal	Step	Type	Example
Enter text in a field	Set Text	Input	Enter any text data such as a user name, an address, a survey response, or in any situation where text entry is accepted. **Note:** If you set a static value for this field, the automation uses it during execution and doesn't prompt you for input from the Now Assist panel.
Simulate a mouse click	Click	Input	Click a button, open a menu, or perform any step typically performed by a mouse click.
Simulate an alternative mouse	Mouse Click	Input	Perform various mouse device

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Goal	Step	Type	Example
action (for example, right-click, drag, scroll, or paste)			actions, such as right-click and select an object or scroll on a web page.
Simulate a key press or a key function	Send Keys	Input	Perform keyboard shortcuts, such as copying text by entering `Ctrl + C` on fields and elements. **Note:** If you set a static value for this field, the automation uses it during execution and doesn't prompt you for input from the Now Assist panel.
Capture text from a window or web page	Get Text	Output	Receive text from the source area.
Capture a table	Get Table	Output	Receive table from the source area when the text is in the table format. **Note:** For the step to capture table data successfully, the data must already be in the table form. The step can't convert ordinary text to table data.
Read text from an image	OCR Read Text	Output	Recognize text from images and return it in the standard text format.

You can add multiple steps representing your automation steps.

12. Configure the properties for added screens, anchors, and steps.
    For more information, see Screen, anchor, and step properties in Agentic Desktop.
13. Provide names for all added screens, anchors, and steps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When you create these elements, you can edit the auto-generated name, but follow these naming guidelines.

- Name fields must not be empty.
- Name fields must contain only alphanumeric characters. Spaces and special characters are not permitted.
- Each name must be unique at its respective parent level.
  - Each screen must have a unique name at the desktop-action level.
  - Each anchor must have a unique name at the screen level.
  - Each step must have a unique name at the anchor level.

What to do next

1. Configure the details of your desktop action. For more information, see Add details to desktop actions in Agentic Desktop.
2. Test and activate the desktop action so that it can be added as a tool to AI agents. For more information, see Test and activate a desktop action in Agentic Desktop.
3. Add the desktop action as a tool to AI agents in AI Agent Studio. For more information, see Add a desktop action to an AI agent.

Related topics

Agentic Desktop Design workspace

Screen, anchor, and step properties in Agentic Desktop

Automate repetitive tasks by auto-capturing steps in Agentic Desktop

Add a desktop action to an AI agent

Examples of creating desktop actions

Examples of executing desktop actions using AI agents

Add details to desktop actions in Agentic Desktop

Add desktop action details, such as name, description, and associated applications, and review inputs and outputs.

Before you begin

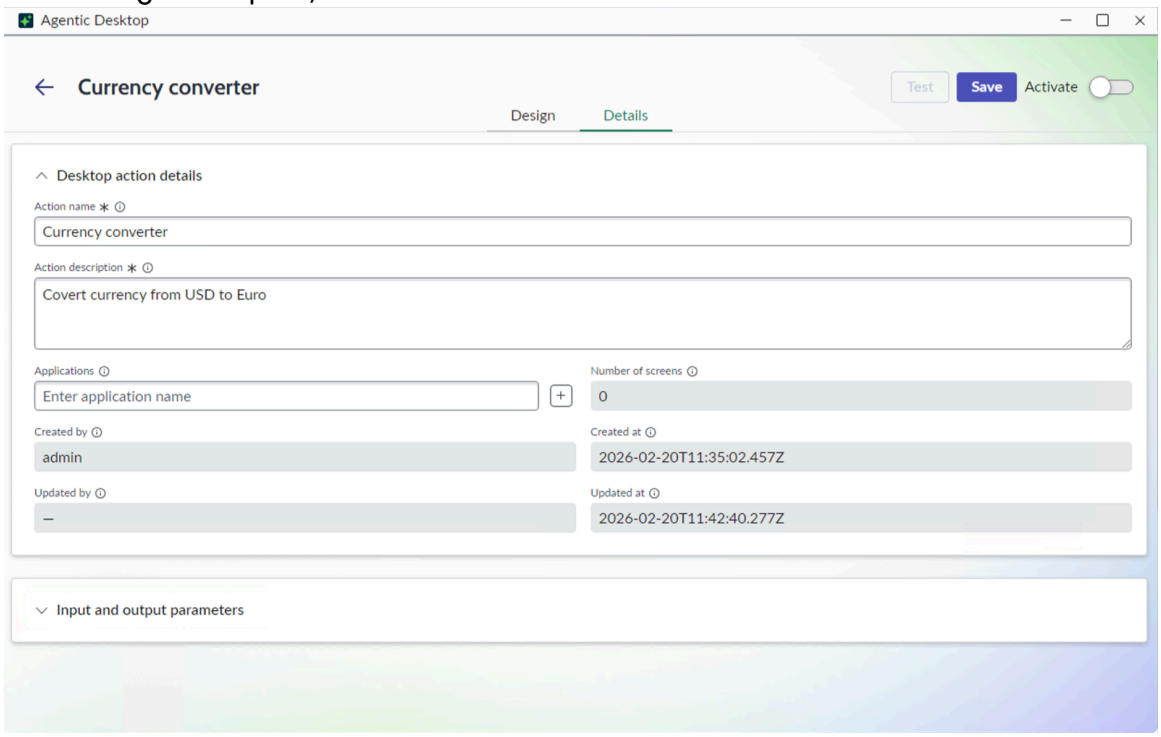
- Capture your automation steps. For more information, see Automate repetitive tasks by auto-capturing steps in Agentic Desktop or Automate repetitive tasks by manually capturing steps in Agentic Desktop.
- Configure the properties for screens, anchors, and steps. For more information, see Screen, anchor, and step properties in Agentic Desktop.

Role required: sn_aia.admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the Design workspace, select the Details tab.



2. **Optional:** In the Desktop action details section, edit the name of the desktop action.
3. **Optional:** In the **Application** list, select the applications you want to automate.
   The list displays the applications available on the ServiceNow instance. If an application isn't listed, you can add it by selecting the **Add** icon ⊞.
4. In the Input and output parameters section, review the parameters related to steps for each screen.
5. Select **Save**.

   You can save the desktop action after all the required fields are configured.

What to do next

Test and activate the desktop action so that it's available as a tool in the AI Agent Studio. You can add this tool to the AI agents that execute desktop actions in your desktop environment. For more information, see Test and activate a desktop action in Agentic Desktop.

Test and activate a desktop action in Agentic Desktop

Test the desktop action and then activate it so that it's available as a tool in the AI Agent Studio. You can add this tool to the AI agents that execute desktop actions in your desktop environment.

Before you begin

- Capture your automation steps. For more information, see Automate repetitive tasks by auto-capturing steps in Agentic Desktop or Automate repetitive tasks by manually capturing steps in Agentic Desktop.
- Configure the properties for screens, anchors, and steps. For more information, see Screen, anchor, and step properties in Agentic Desktop.
- Add desktop action details, such as name, description, and associated applications, and review inputs and outputs. For more information, see Add details to desktop actions in Agentic Desktop.

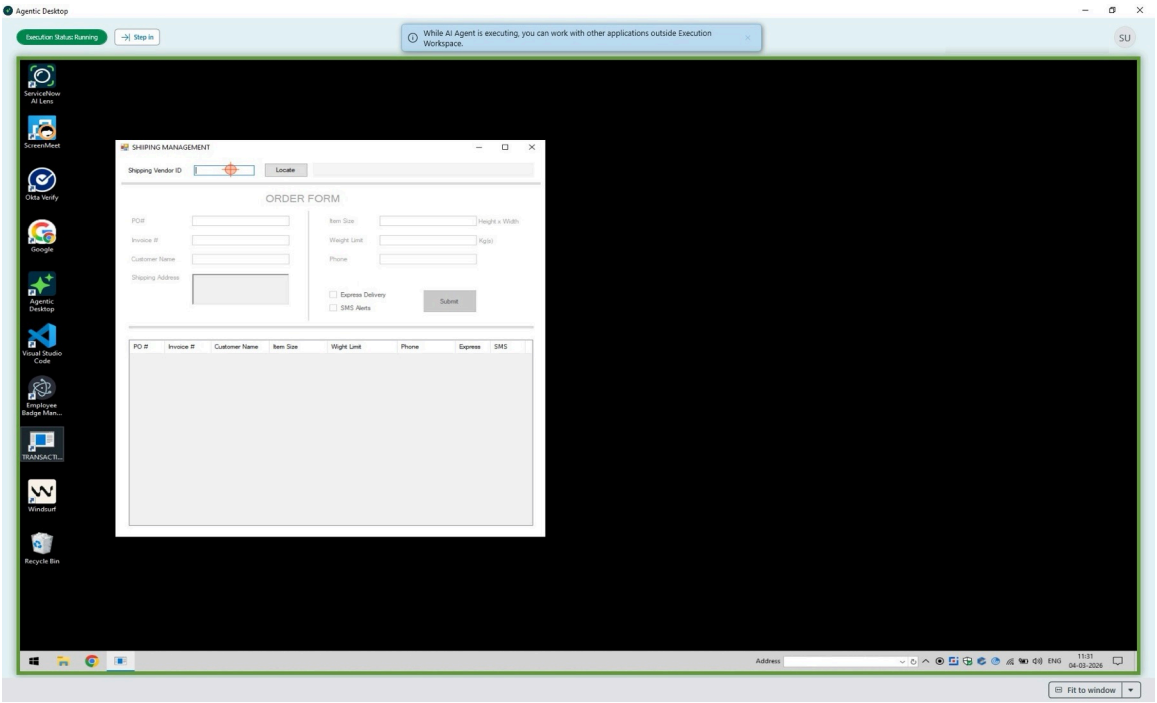
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: sn_aia.admin

About this task

You can test the desktop action before you activate it. You can also test individual screens to quickly identify and troubleshoot issues.

Execution workspace is launched to run and test the selected screen or the desktop action.



After you run the test by entering the required inputs manually, test results are shown either for a screen or for entire desktop action in the Execution workspace. While running tests, if you do not see the Execution workspace coming forth, switch to the Execution workspace manually.

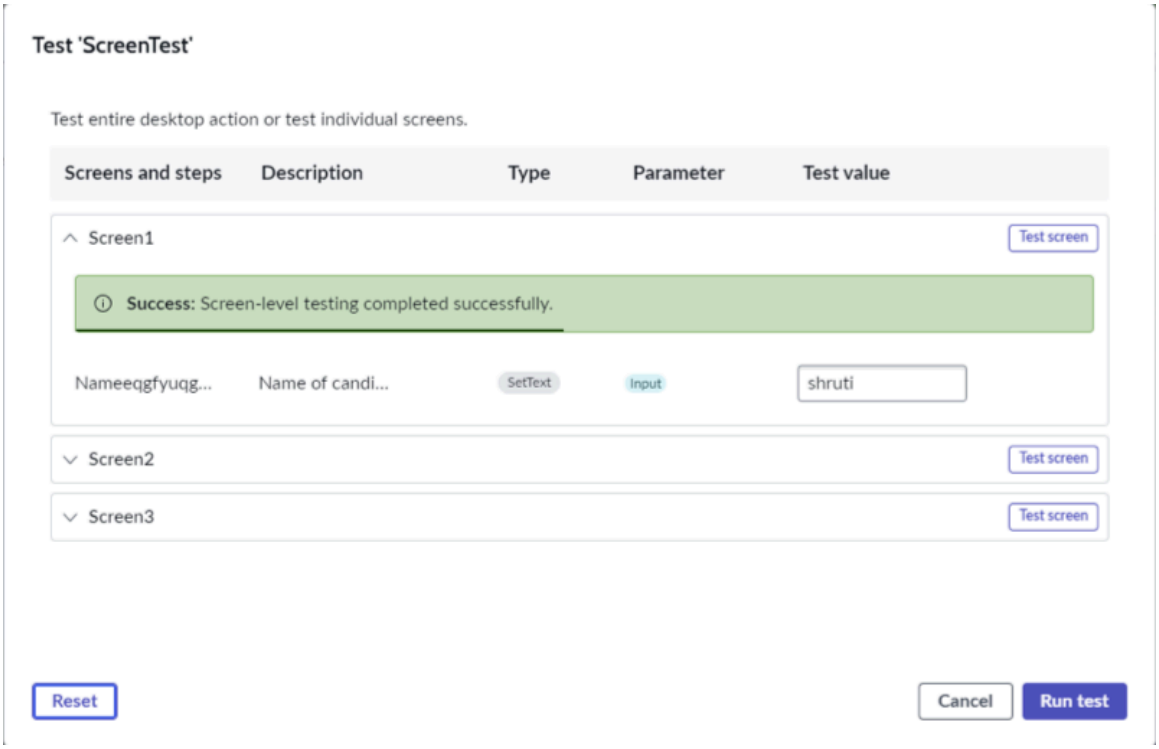
While working on a desktop action, the test values you enter are retained across test runs until you close the desktop action or reset the test values.

Procedure

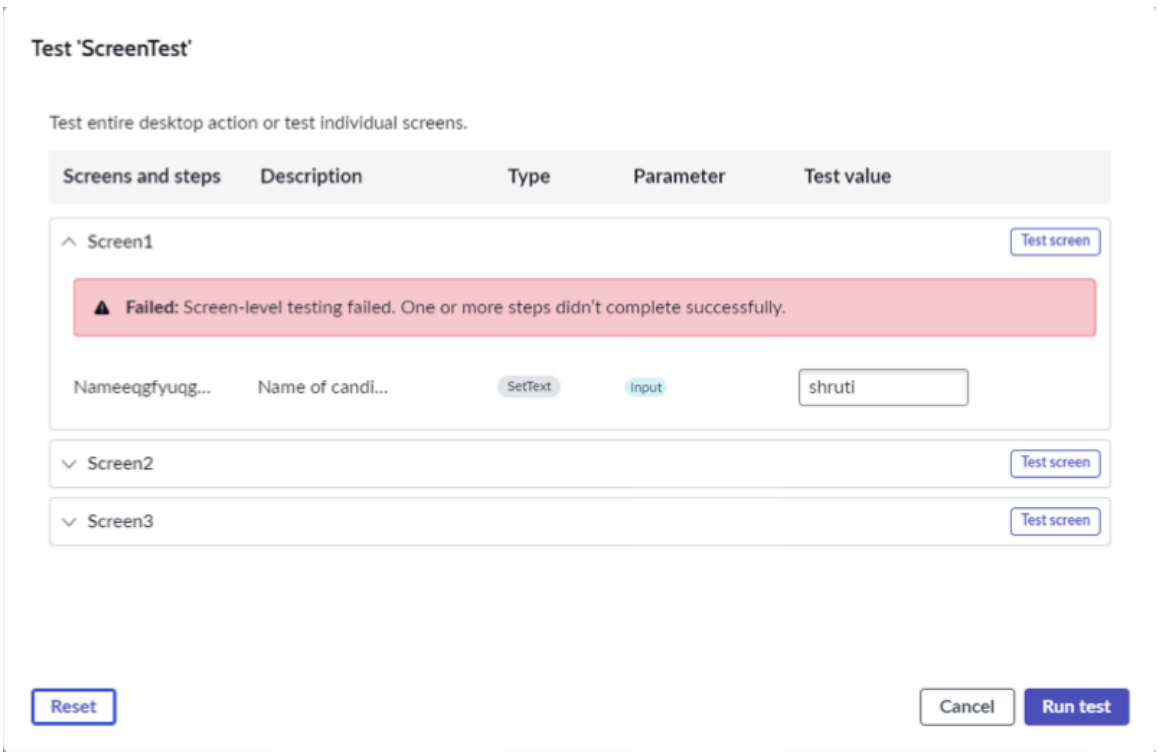
1. From the desktop action header, select **Test**.
2. Option 1: Run the screen-level test.
   a. Enter required inputs for the screen that you want to test.
   b. Select **Test screen**.
   c. After the test run is complete, check the status.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If test is successful, proceed with testing the next screen.



- If test isn't successful, fix the issue highlighted in the error message and run the screen-level test again.



d. Optional: Reset test data by selecting **Reset**.

e. Optional: Cancel running test by selecting **Cancel**.

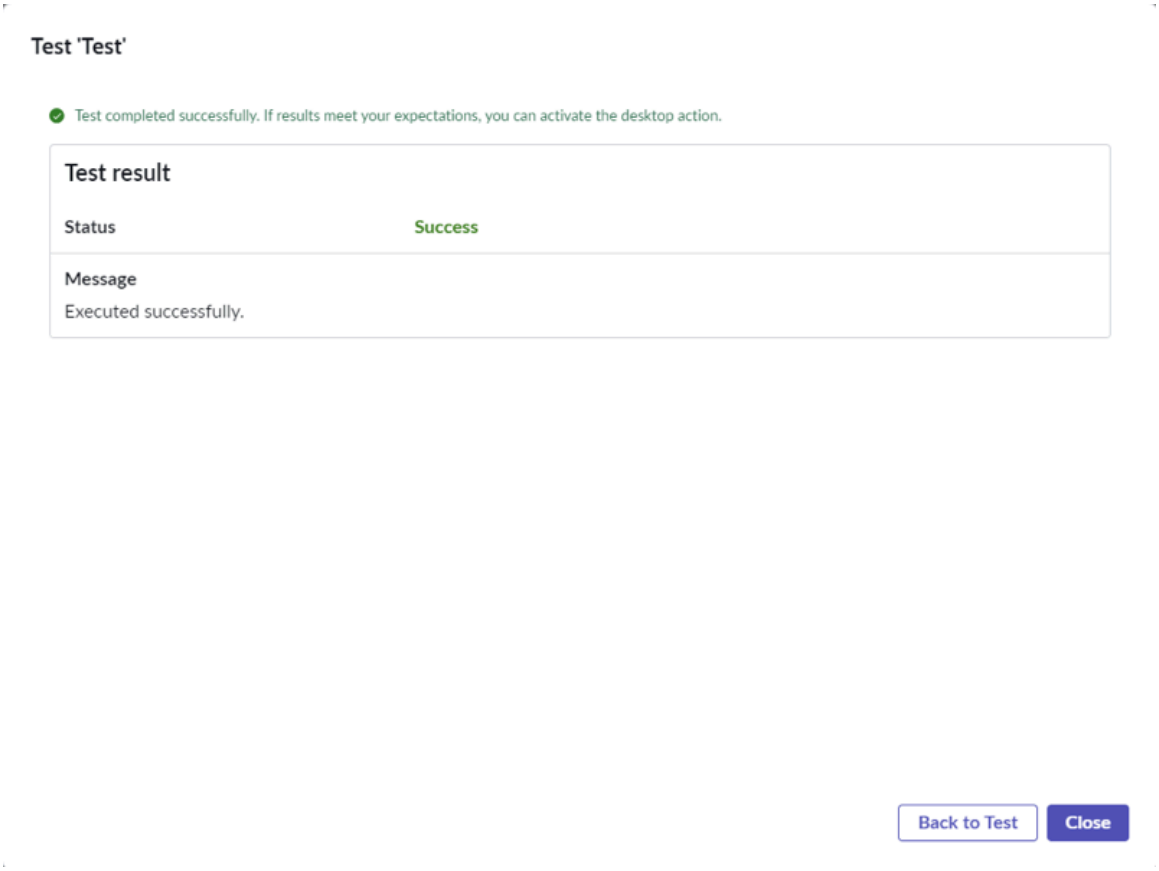
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Option 2: Run the Desktop action-level test.

   a. Enter required inputs for all screens in a desktop action.

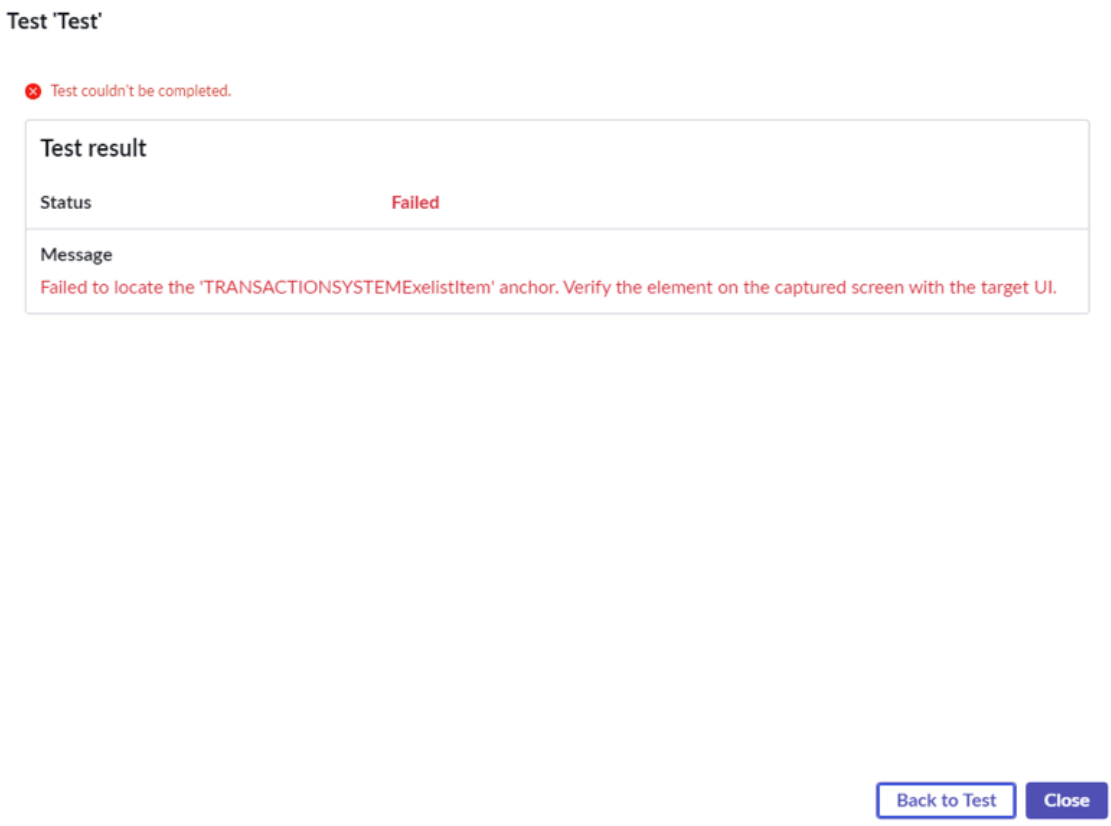
   b. Select **Run test**.

   c. After the test run is complete, check the status.
      - If test is successful, proceed with activating the desktop action.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If test isn't successful, select **Back to Test** to check the configuration and run the test again.



4. **Optional:** Cancel running test by selecting **Cancel**.
5. After a successful completion of the desktop action test and when you are satisfied with the result, select **Activate** to publish the desktop action in AI Agent Studio.
   The **Activate** button sets the desktop action as active or inactive. Only the activated desktop actions can be selected from the drop-down of UI block actions in AI Agent Studio.

   Note: After activation, save the desktop action to make it available as a tool in AI Agent Studio.

Trouble?
If you get en error that the file size is beyond 10 MB, you can increase the file size in a payload by creating system properties in the ServiceNow instance. For more information, see Increase payload limit through system properties.

What to do next

Create an AI agent. For more information, see Creating AI agents for Agentic Desktop.

Enable AI agents to perform desktop actions by adding desktop actions as tools. For more information, see Add a desktop action to an AI agent.

Screen, anchor, and step properties in Agentic Desktop

Learn about the properties of screens, anchors, and steps. There are multiple types of steps and each step type has distinct properties. You can update the properties to modify the behavior of the steps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Screen

Screen properties

Property	Description
Name	Unique name for the captured screen. The updated name reflects in the Screens and steps panel.
Execution order	Order in which the screen is executed when multiple screens are captured. You can drag the screen at the desired order in the Screens and steps panel.

Anchor

Anchor properties

Property	Description
Name	Unique name of the added anchor. The updated name reflects in the Screens and steps panel.
Execution order	Order in which the anchor is executed when multiple anchors are added.
X	Location of the top-left corner of the anchor on the screen along the X-axis.
Y	Location of the top-left corner of the anchor on the screen along the Y-axis.
Width	Width of the area highlighted as anchor.
Height	Height of the area highlighted as anchor.
Convert to grayscale	Option to convert the captured screen as a grayscale image. The default value is False.
Search in Scale	Option to specify whether the component scales the target image before trying to match the source image. The scaling is done in an increment of 10% each between 100% and 150%. The default value is False.
Confidence threshold	Specifies the accuracy in matching the image captured before the component performs a step. The value of 1 defines a 100% match while 0.5 defines 50% match. The default value is 0.95 or 95% match.
Retry with scroll	Option to scroll the target image before trying to match the source image.
Wait for image	Option to specify if the automation must wait for the image to appear on the screen. The default value is True.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Anchor properties (continued)

Property	Description
Max wait time	Number of seconds to wait for the image or application to appear on the screen.
Delay before	Delay in number of seconds before the anchor is executed. The automation waits until the specified time elapses before executing the anchor.
Delay after	Delay in number of seconds after the anchor is executed. The automation will wait until the specified time elapses before executing the next anchor.

Step

Common properties for steps

The following properties apply to all steps - Set Text, Click, Mouse Click, Send Keys, Get Text, Get Table, and OCR Read Text.

Common properties

Property	Description
Name	Unique name of the step. The updated name reflects in the Screens and steps panel.
Execution order	Order in which this step is executed when multiple steps are set up. This order is assigned when you add a step from the anchor. You can drag the step under its anchor at the desired order in the Screens and Steps panel.
Type	Type of the step, such as Set Text, Get Text, Mouse Click.
Description	Helpful description of this step.
X	Location of the green plus icon  along the X-axis relative to the anchor.
Y	Location of the green plus icon  along the Y-axis relative to the anchor.
Delay before	Delay in seconds before the action is executed.
Delay after	Delay in seconds after the action is executed.

Set Text

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set Text properties

Property	Description
Value	New text to set in the field.

Mouse Click

Mouse Click properties

Property	Description
Mouse click type	Type of mouse click action to set for this step: • Left click • Left double click • Middle click • Right click • Right double click • Move • Drag • Drop • Scroll down • Scroll up • Paste

Send Keys

Send Keys properties

Property	Description
Clear existing value	Option to specify if the step clears the existing value on a field before setting the text.
Value	New text to set in the field.
Focus with mouse click	Option to use the mouse click to focus on the area in the application where the mouse inputs are passed.

OCR Read Text

OCR Read Text properties

Property	Description
Width	Width of the area in image to recognize and get the text.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

OCR Read Text properties (continued)

Property	Description
Height	Height of the area in image to recognize and get the text.

Increase payload limit through system properties in Agentic Desktop

By default, maximum 10 MB of file size is allowed in a scripted REST API request payload. Increase the payload limit to 15 MB by creating system properties in the global scope.

Before you begin

You must do this procedure in the ServiceNow instance.

Role required: system admin

About this task

If you already have these properties configured for your instance, you can update the value to 15 MB to increase the limit.

Procedure

1. Navigate to **All**.
2. In the filter navigator, enter `sys_properties.list`.
3. In the System Properties table, select **New** to add the glide.rest.max_content_length property.
   a. On the form, fill in the fields.

Field	Description	Value
Name	Name of this system property	glide.rest.max_content_length
Description	Explanation for this property	For example: `Revised payload for a scripted REST request body.`
Type	Data type of this system property	integer
Value	The maximum size, in megabytes, for a scripted REST request body. The default value is 10 MB.	15

   b. Select **Submit**.
4. In the System Properties table, select **New** to add the glide.rest.scripted.max_inbound_content_length_mb property.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. On the form, fill in the fields.

Field	Description	Value
Name	Name of this system property	glide.rest.scripted.max_inbound_content_leng
Description	Explanation for this property	For example: `Revised payload for a scripted REST request body.`
Type	Data type of this system property	integer
Value	The maximum size, in megabytes, for a scripted REST request body. The default value is 10 MB. **Note:** Even if glide.rest.scripted.max_inbound_content_length_mb is set, the request body is limited to the value of glide.rest.max_content_length.	15

b. Select **Submit**.

Creating AI agents for Agentic Desktop

Create an AI agent in AI Agent Studio to mimic human-like intelligence while executing desktop actions for repetitive tasks.

Overview of AI agents for Agentic Desktop

In the ServiceNow agentic ecosystem, an AI agent is a set of large language model (LLM) instructions and tools that can perform specific tasks. The Now Assist AI agents can perform specific tasks and functions, often using natural language instead of traditional code. For more information, see Create an AI agent.

AI agents process instructions, generate execution plans, and run desktop actions autonomously and semi-autonomously across legacy systems, thick client applications, and business applications lacking APIs or backend integrations. AI agents can interpret your goal and map them to one or more desktop actions via metadata (capabilities, inputs, and outputs).

Use AI agents to do the following tasks for your organization:

- Dynamically figure out a plan to complete a desktop task.
- Collaborate with other AI agents for subtasks as needed.
- Take feedback from humans as needed.
- Plan to solve an issue with human collaboration.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Define the specialty of an AI agent

   Write a clear name and description of the AI agent, define role, and list of steps this AI agent must complete, define supported LLMs, enable third-party access, and manage long-term memory. The LLM analyzes the specific wording that you use to understand the specialty of this AI agent.

2. Add a desktop action to an AI agent

   Add a desktop action as a tool to the AI agent to enable desktop automations. Tools provide your AI agents with the capabilities necessary to complete their tasks. Providing your AI agents with the appropriate tools help with the robustness and quality of their performance. An AI agent selects a tool based on the tool's name and description, which must be clearly written.

3. Define security controls for an AI agent

   Define security controls for who can access the AI agent and what data the AI agent has access to. The Define security controls step is divided into two parts:

   - **Define user access**: Creates an ACL that determines which users can discover or invoke the AI agent.
   - **Define data access**: Defines the data that the AI agent has access to once it's invoked.

   Trigger conditions are not supported for AI agents that execute desktop actions. You must manually trigger these agents from the system where the Agentic Desktop application is installed.

4. Select channels and status for an AI agent

   Activate the AI agent to use in an assistant in Now Assist for Virtual Agent and set the processing message. This AI agent can engage with users who initiate an interaction when it's available for use in channels. Select channels where you want this AI agent to be available to engage with users.

Related topics

Agentic Desktop Design workspace

Desktop actions in Agentic Desktop

Add a desktop action to an AI agent

Add a desktop action as a tool to an AI agent in AI Agent Studio so that AI agents can perform desktop actions for repetitive tasks.

Before you begin

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > *AI Agent Studio* > **Create and manage** > **AI agents**.
2. Open the AI agent that you want to add a desktop action to.
   For creating an AI agent, see Create an AI agent.
3. Navigate to the Add tools and information step.
4. In the Add tool drop-down list, select **Desktop action**.
5. On the form, fill in the fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Add a desktop action form

Fields	Description
Select a type of desktop action	○ **UI block**: These actions include UI interactions, such as clicking buttons, typing into text boxes, selecting from drop-down menus, and more. Example: Filling in fields on a payroll application, updating inventory in a point-of-sale system, or submitting a request through a desktop insurance claims app. ○ **Non UI blocks**: These actions are prebuilt components that don't interact with the UI, but run in the background. Non UI block actions enable your AI agents to interact with various applications and system components and offer prebuilt actions for common tasks, reducing the need for complex scripting. Example: Reading data from Microsoft Excel, extracting emails from Microsoft Outlook, generating a PDF, compressing files into a ZIP, fetching records from a database, or sending a system notification.
Select a desktop action	Desktop action that you want to add to the AI agent from the list. **Note:** This field appears when **UI block** is selected as the type of desktop action.
Select an application	Select an application from the list to use the associated actions. The following applications are supported for non-UI block actions: ○ Microsoft Excel ○ Microsoft Outlook ○ Microsoft Word ○ PDF ○ PowerShell Connector ○ SQL ○ SSH ○ SystemAction **Note:** This field appears when **Non UI block** is selected as the type of desktop action.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fields	Description
Desktop action description	Read-only. Description of the desktop action or application that you selected.
Applications	Read-only. Applications that this desktop action interacts with. **Note:** This field appears when **UI block** is selected as the type of desktop action and a desktop action is selected from the list.
Created by	Read-only. User who created the desktop action. **Note:** This field appears when **UI block** is selected as the type of desktop action and a desktop action is selected from the list.
Last updated	Read-only. Date when this desktop action was last updated. **Note:** This field appears when **UI block** is selected as the type of desktop action and a desktop action is selected from the list.
Inputs	Read-only. Input parameters associated with this desktop action. **Note:** This field appears when a desktop action or application is selected from the list.
Outputs	Read-only. Output parameters associated with this desktop action. **Note:** This field appears when a desktop action or application is selected from the list.
Name	Name that you want to specify for your selected desktop action.
Tool description	Description of the desktop action tool and what it's going to do to assist your AI agent. **Note:** This description is sent to the large language model (LLM).
Execution mode	Mode of execution for your selected desktop action:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fields	Description
	○ **Supervised**: Inputs from your live agent are required during the execution of this desktop action while the AI agent runs. ○ **Autonomous**: Doesn't require any input from your live agent during the execution of this desktop action while the AI agent runs.
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: ○ **Yes** ○ **No** If you want the AI agent to work in Off Glide architecture with AI-native experience, you must turn-on the **Display output** toggle. When the toggle is turned-on, you can add widgets that can be used in assistants built with AI-native experiences. The widget configuration includes: ○ **Widgets**: Defines the display output to render the content in a better user experience. You can select the widget from the drop-down. ○ **Require widget transformation**: An additional LLM call is required to transform the raw tool. If you choose to skip this transformation step, the tool output will be directly mapped to the widget. ▪ **Yes** ▪ **No** ○ **Display refined widget message**: Refines the widget message when configured. ▪ **Yes** ▪ **No** **Note:** The display output as a toggle is exclusively available for the AI-native experience when the Off Glide Conversation Server plugin (com.glide.cs.offglide) is installed. If the plugin is not installed, you will continue to access the standard display output options.
Advanced settings	
Select an output transformation format	Style for the LLM to present the results as it passes information between tools and to other agents. Out transformation formats:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Fields	Description
	○ None ○ Concise ○ Paragraph ○ Summary ○ Custom
Write processing messages for users	Message to display to users during tool execution. ○ In-progress message: Write an in-progress message to be displayed to end-users while the tool is running. ○ Completion message: Write a completion message to be displayed to end-users once the tool finishes running.

6. Select **Add**.
   The desktop action is added in the Desktop actions list on the Add tools and information page.

Examples of creating desktop actions

Learn the key concepts and workflow for creating end-to-end desktop actions and building automations for your desktop environments.

The tasks that you can automate are done manually on desktop-native applications and repetitive in nature. Here are examples:

Badge management automation

Shows how an HR team automates the end-to-end workflow for issuing, replacing, and disabling badges for their employees.

Shipping order processing

Demonstrates how to extract order data from Excel, interact with a shipping management tool, and complete the data processing life cycle.

Example: Automate badge request management using Agentic Desktop

Automate various tasks related to badge requests through desktop actions using Agentic Desktop and AI agents.

Your HR representatives manage repetitive badge-related tasks. For example, issuing new badges, distributing temporary badges, replacing lost badges, and disabling badges during offboarding. To streamline and automate this work, you can create a desktop action for each task and assign these actions to an AI Agent in AI Agent Studio.

When new requests come in, HR representatives can trigger the AI agent from the Now Assist panel. The AI Agent automatically selects and runs the appropriate desktop action. This automation reduces manual effort and enables them to focus on higher-value work.

Create badge desktop action in Agentic Desktop

Automate various badge-related tasks through desktop actions in Agentic Desktop.

Before you begin

To access the Agentic Desktop functionality, perform the following steps:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Enable Agentic Desktop on your ServiceNow instance. For more information, see Configure Agentic Desktop.
- Download the Agentic Desktop installer to automate repetitive tasks across applications and systems. For more information, see Download Agentic Desktop installer.

Confirm that the following system requirements are met:

- Windows 11 operating system is used.
- A .NET 9.0 runtime v9.0.10 and .NET 9 Desktop Runtime v9.0.10 is installed.
- No extended monitors are connected.
- Theme must match between the systems used for recording and execution.

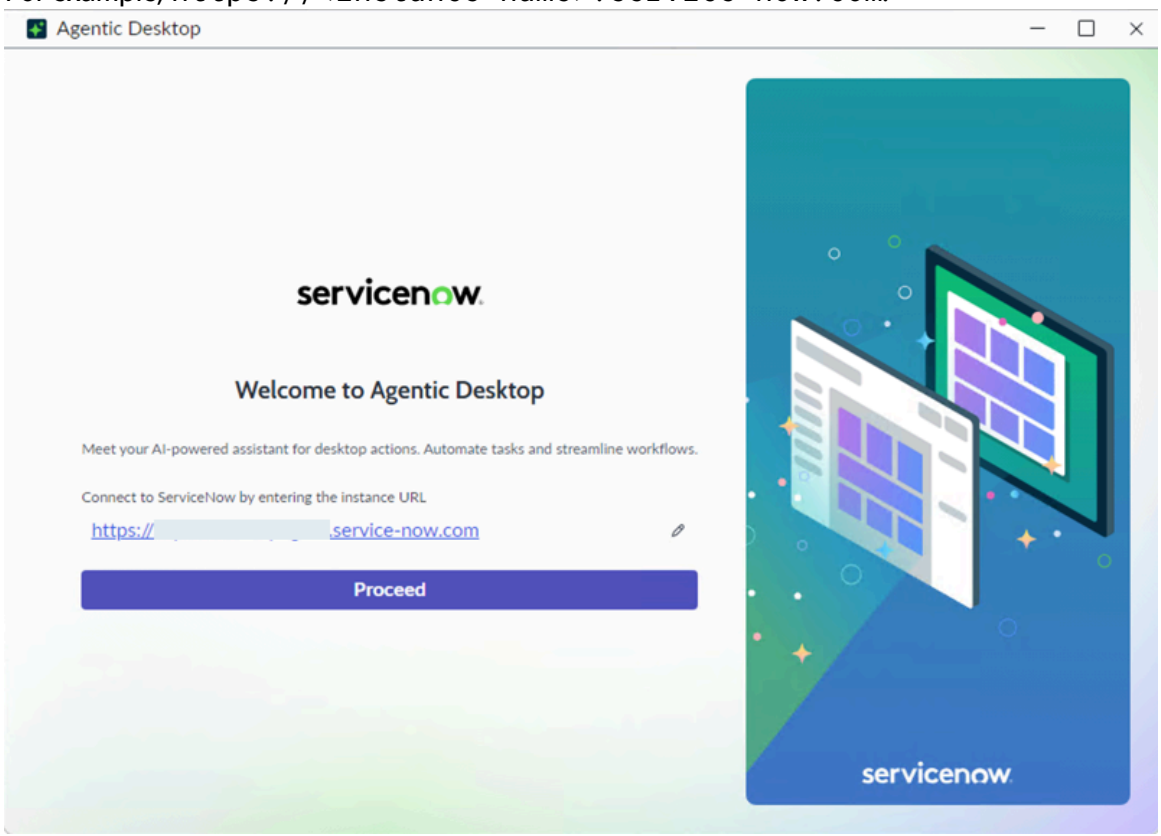
Familiarize yourself with the Design workspace and Action recorder. For more information, see Agentic Desktop Design workspace and Action recorder in Agentic Desktop.

Role required: sn_aia.admin

Procedure

1. From your Windows system, launch the Agentic Desktop application.
2. On the login page, in the **Add ServiceNow URL** field, enter the ServiceNow instance URL.

   Example
   For example, `https://<instance name>.service-now.com`.

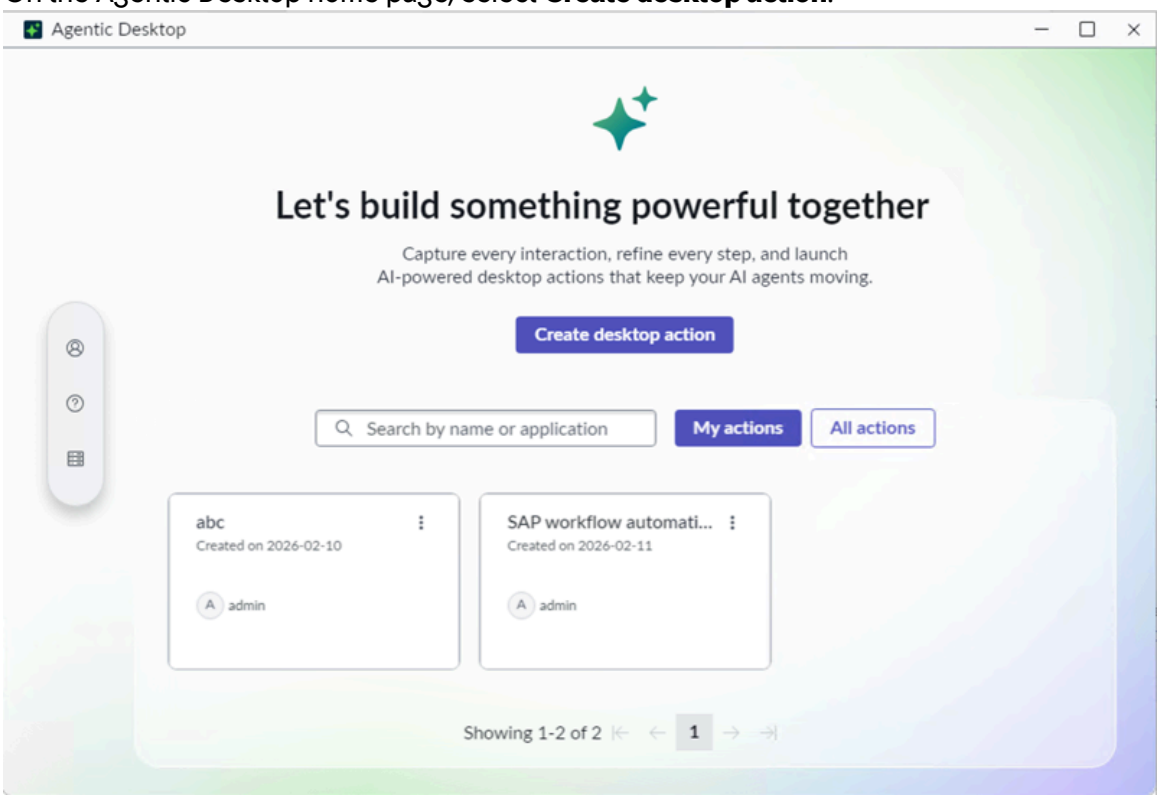


3. Select **Proceed**.
4. Log in to your ServiceNow account by entering your user name and password.
   Your account must have the sn_aia.admin role.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

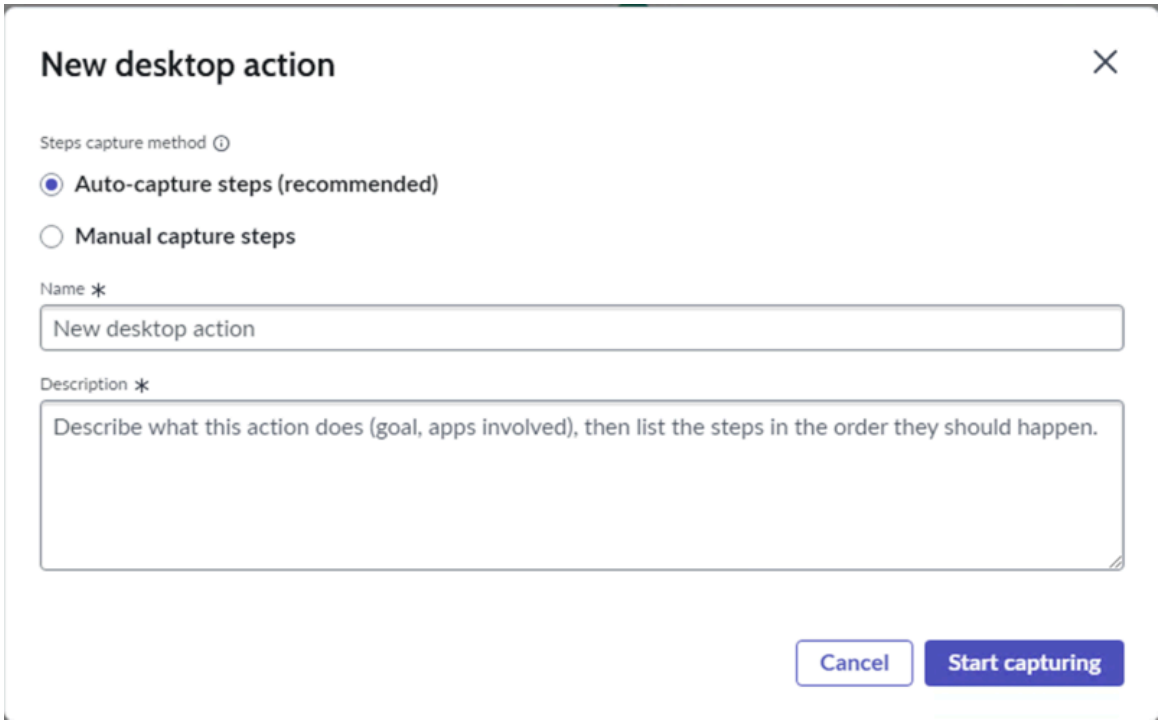


5. On the Agentic Desktop home page, select **Create desktop action**.



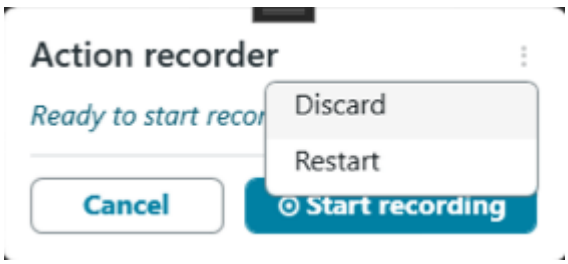
6. In the New desktop action dialog box, select **Auto-capture steps (recommended)**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



7. In the **Name** field, enter `Create new badge`.
8. In the **Description** field, enter `Desktop action for issuing a badge to a new employee.`
9. Select **Start capturing**.

   The Agentic Desktop window is minimized and the Action recorder panel is launched. You can freely drag and reposition the Action recorder panel anywhere on your desktop screen.



10. Open the **Employee Badge Management** application on your desktop.
11. From the Action recorder panel, select **Start**.
    You see that a "Recording started" message on the Action recorder panel.
    - Pause recording: You can skip recording steps by selecting **Pause**.
    - Discard recording: You can discard the recording if it doesn't meet your needs by selecting **Discard**.
12. On the **Employee Badge Management** application, perform the following steps for creating a badge.

    The recorder feature records the steps that you perform for creating a badge.

    a. Select **CREATE BADGE** on the **New Badge** tile.

    b. In the New Badge Creation page, search for the employee using the employee ID.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

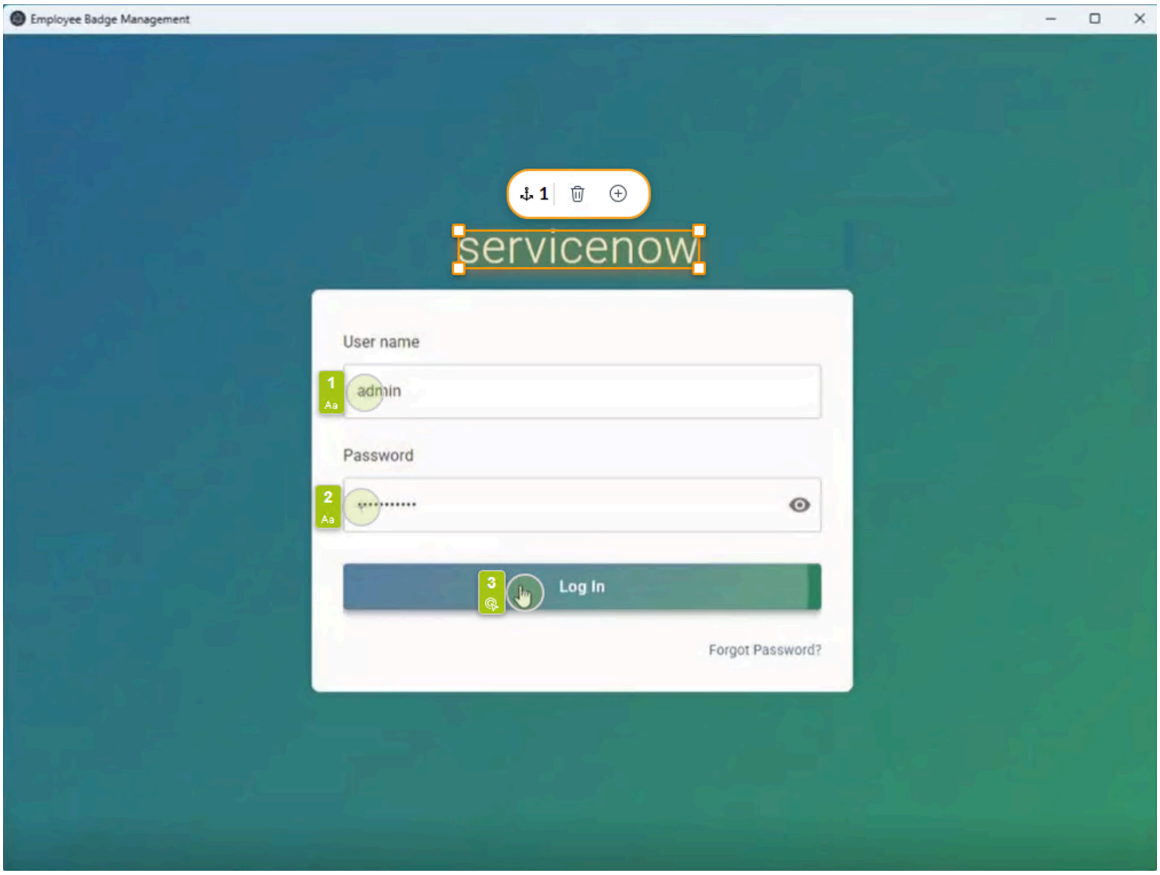
c. After the employee record is found, select **CREATE BADGE** for creating badge.

You see the success message that the new badge is created.

13. After you're done with all the steps, select **Stop** on the Action recorder panel.
You see a "Drafting recorded action(s)" message on the Action recorder panel.

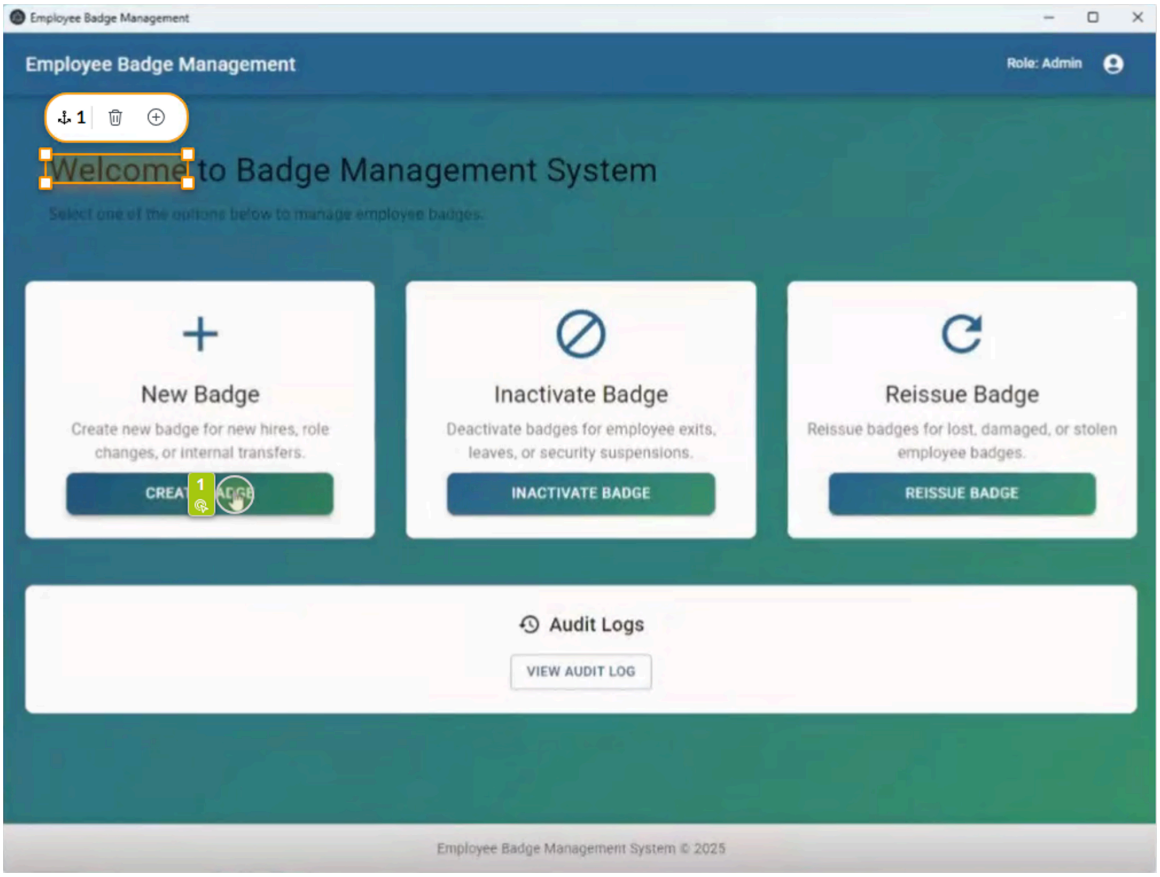
The recorded steps are displayed as screenshots in the Design workspace with anchors and steps automatically assigned.

Screen1

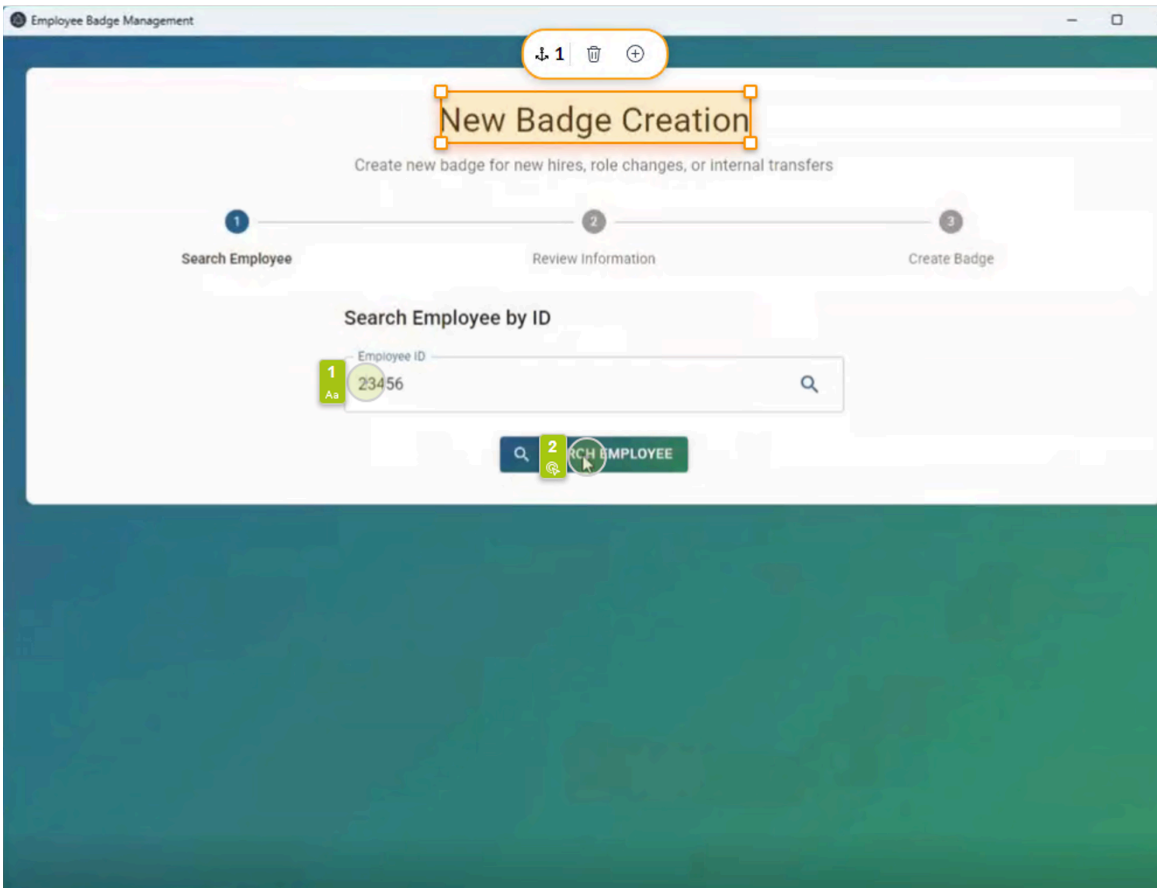


© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Screen2

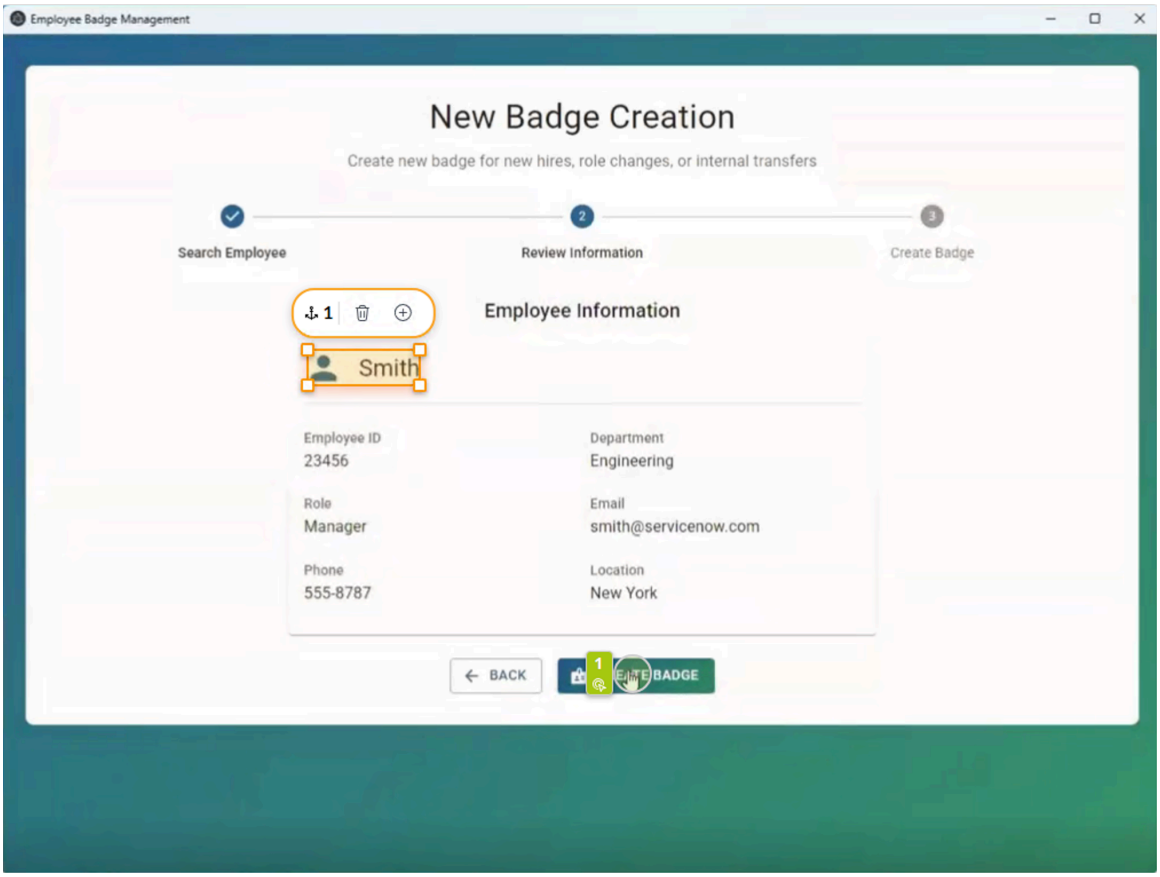


Screen3

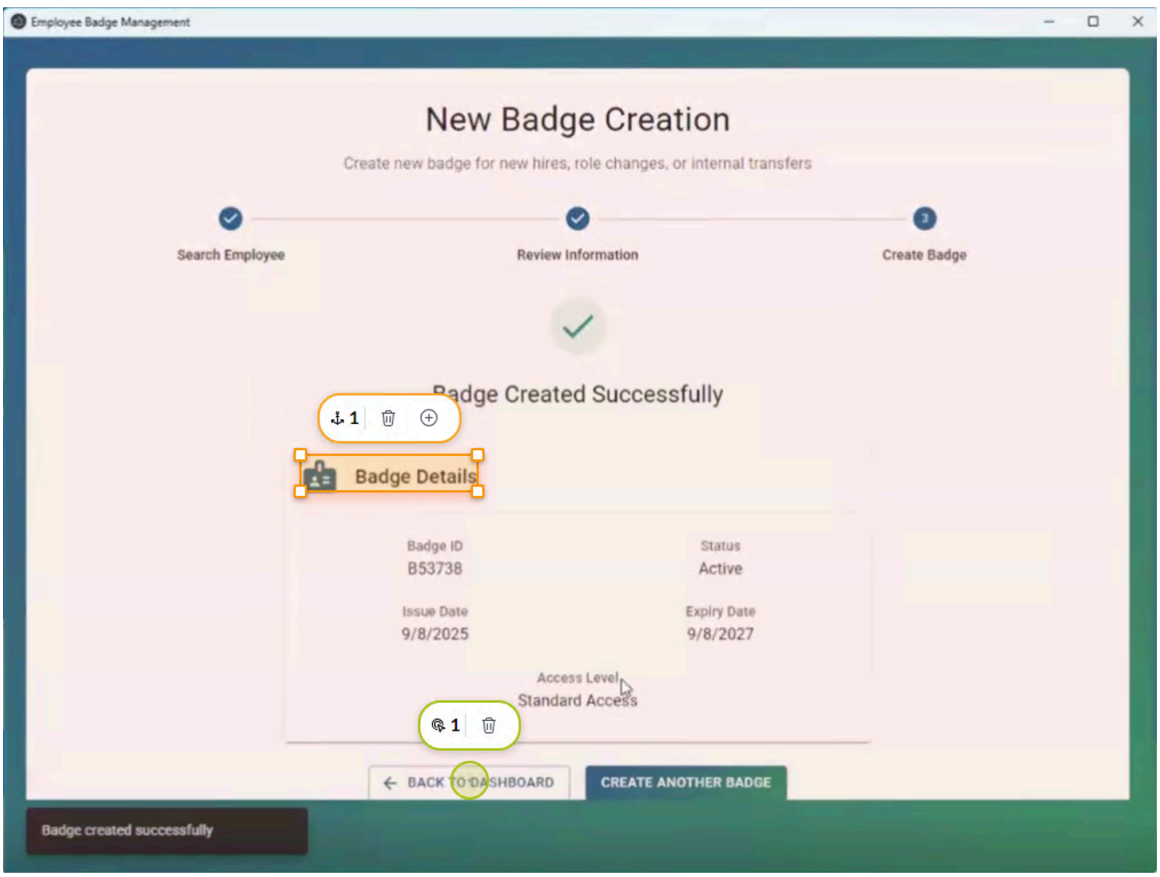


© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Screen4



Screen5



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

14. Configure the following properties for the captured steps.

Screen > Step	Property	Value
Screen1 > SetText1	Value	admin
Screen1 > SetText2	Value	Enter your password
Screen1 > Click1	Delay after	5
Screen2 > Click2	Delay after	5
Screen3 > Click2	Delay after	10

For more information, see Screen, anchor, and step properties in Agentic Desktop.

15. Provide names for all added screens, anchors, and steps.

When you create these elements, you can edit the auto-generated name, but follow these naming guidelines.

- Name fields must not be empty.
- Name fields must contain only alphanumeric characters. Spaces and special characters are not permitted.
- Each name must be unique at its respective parent level.
  - Each screen must have a unique name at the desktop-action level.
  - Each anchor must have a unique name at the screen level.
  - Each step must have a unique name at the anchor level.

16. Select the Details tab.
17. In the Applications list, add Badge Management Application.
18. Select **Save**.
19. Test and activate the desktop action.
For more information, see Test and activate a desktop action in Agentic Desktop.
20. Similarly, create and activate the following desktop actions.
    - Badge application login
    - Badge application logout
    - Reissue badge
    - Deactivate badge
    - Issue a temporary badge
    - Read Request details

Create AI agents and add tools for badge management

Create an AI agent in AI Agent Studio and add desktop action tools for automating badge-related requests.

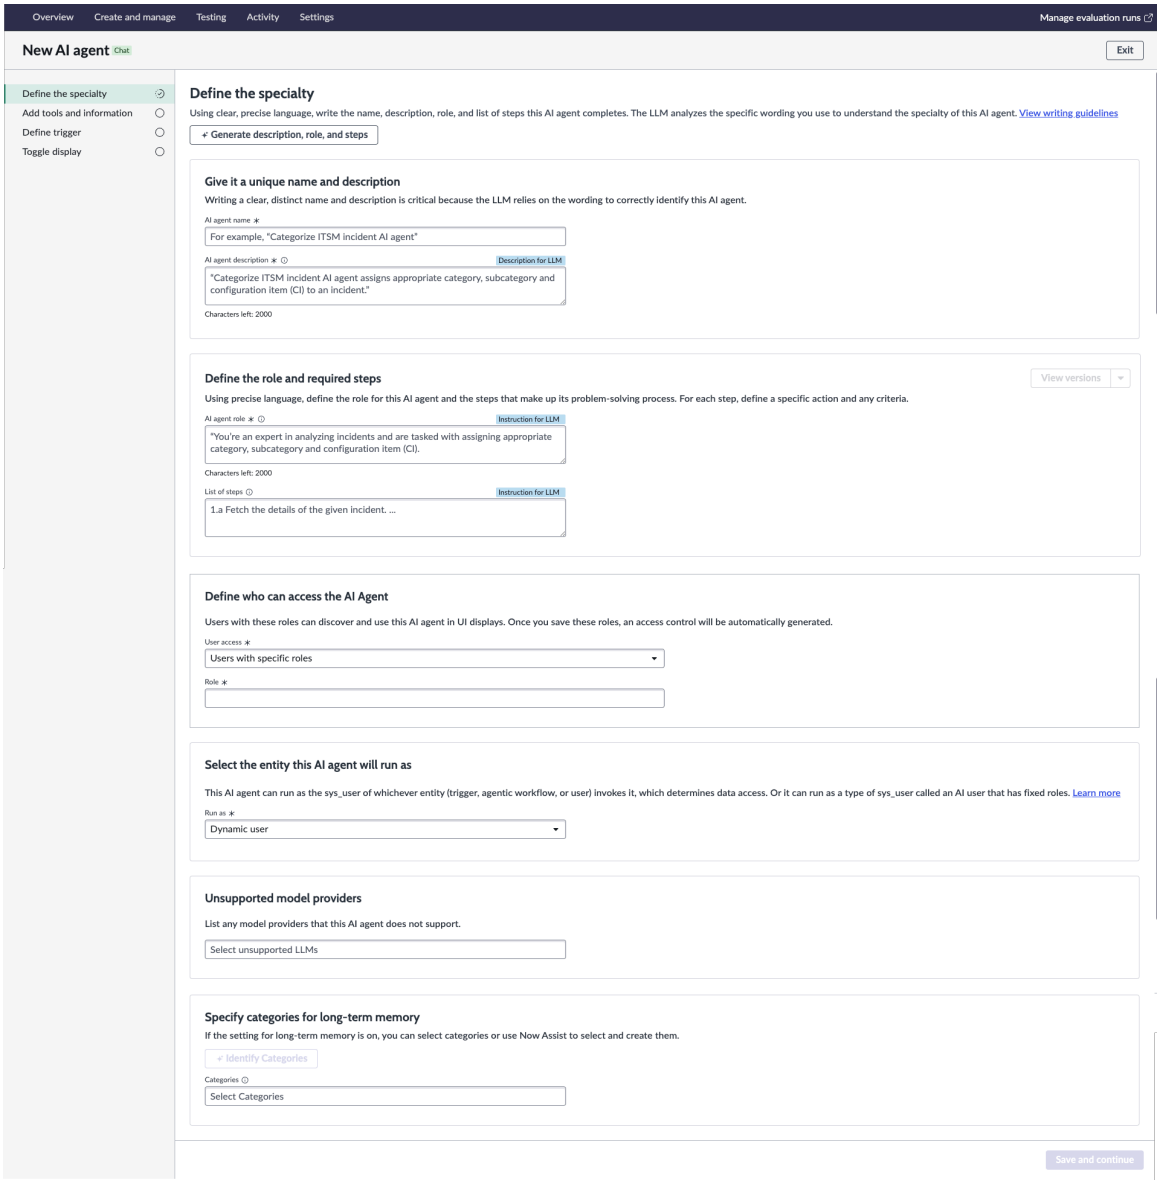
Before you begin

Role required: sn_aia.admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **AI agents**.
2. From the **Add** drop-down, select **Chat**.
3. On the New AI Agent page, in the Define the specialty step, define your AI agent and provide the specialties that this agent contains so that the LLM can analyze the wording you use to understand the purpose of the AI agent.



Note: The more details that you provide, the more accurately your AI agent can perform.

a. Describe your AI Agent by giving it a unique name and description.

Give it a unique name and description

Field	Description
AI agent name	Badge Management Agent
AI agent Description	This AI agent takes Requests details and launches the Employee Badge Management

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	application and perform various tasks—create badges, deactivate badges, replace lost badges, and Reissue the badges. The agent executes the appropriate desktop actions automatically. This reduces manual effort, ensures process consistency, and speeds up the overall employee lifecycle experience.

b. Define the role and necessary steps so that the AI agent can carry out its tasks.

Note: The AI agent uses this information as guidance to tailor its responses and actions.

Define the role and required steps

Field	Description
AI agent role	Automates the intake of badge-related requests and performs accurate, end-to-end execution of badge actions—such as issuance, replacement, temporary provisioning, and deactivation—across the HR and security systems.
List of steps	**Note:** If your automation requires manual inputs, such as entering an OTP or CAPTCHA, you must provide instructions to the AI Agent to wait for the user input during execution. Otherwise, the automation can't proceed. Read the comma-separated Request numbers from the input by invoking the **Read Request details** desktop action. Launch and log in to the Employee Badge Management application by invoking the **Badge application login** desktop action. Based on the details of each Request, perform one of the following actions that is applicable. ▪ Create a badge i. Create a badge in Employee Badge Management application by invoking **Create new badge** desktop action. ii. If badge creation failed, create an incident with the issue details. Otherwise, continue with other steps. ▪ Deactivate a badge

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	i. Deactivate a badge in Employee Badge Management application by invoking **Deactivate badge** desktop action. ii. If badge deactivation failed, create an incident with the issue details. Otherwise, continue with other steps. ▪ Reissue Badge i. Reissue a badge in Employee Badge Management application by invoking **Reissue badge** desktop action. ii. If badge deactivation failed, create an incident with the issue details. Otherwise, continue with other steps. ▪ Issue a temporary badge i. Issue a temporary badge in Employee Badge Management application by invoking **Issue temporary badge** desktop action. ii. If badge reissue failed, create an incident with the issue details. Otherwise, continue with other steps. Log out of the Employee Badge Management application by invoking the **Badge application logout** desktop action.

c. Select **Save and continue**.
You're directed to the Add tools and information page.

4. On the Add tools and information page, add Desktop actions tools for AI agents to automate your desktop tasks.

a. In the Add tool drop-down list, select **Desktop action**.

b. In the **Select a type of desktop action** field, select **UI block**

c. In the **Select a desktop action** drop-down list, select **Create new badge** desktop action.
This desktop action enables AI agents to create a badge in the Employee Badge Management application.

d. Provide a name and tool description for this desktop action configuration.
Tool description of the desktop action helps with what it's going to do to assist your AI agent.

Note: This description is sent to the large language model (LLM).

e. In the **Execution mode** field, select **Autonomous**.

f. Select **Add**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The desktop action is added in the Desktop actions list on the Add tools and information page.

g. Similarly, add the following desktop actions related to the badge management to this AI agent.
   - Badge application login
   - Badge application logout
   - Reissue badge
   - Deactivate badge
   - Issue a temporary badge
   - Read Request details

5. Select **Save and continue**.
6. Complete the remaining steps that are necessary.
   For more information, see Create an AI agent.
7. Select **Save and test** to complete the configuration steps or review a previous step by selecting **Back**.

   Selecting Save and test leads you to the AI agent testing page, where you can test the AI agent that you created. For more information, see Manually test the execution of an AI agent.

   To test the AI agent, you must have the sn_aia.admin role and any roles the ACLs configured for the AI agent and its tools require, if applicable.

What to do next

Enable your HR representatives to trigger AI agents from the Now Assist panel to address badge-related requests.

For more information, see Example: Use AI agents to process badge-related requests automatically.

Example: Automate shipping management tasks using Agentic Desktop

Automate various tasks related to shipping management through desktop actions using Agentic Desktop and AI agents.

Your shipping coordinators frequently handle repetitive data-entry tasks. For example, entering PO numbers, invoice details, customer information, shipping addresses, item dimensions, and weights from Excel into the Shipping Management application on their desktops. When there are dozens of rows to process, it becomes time-consuming, error-prone, and significantly slows down their daily workflow.

To streamline this process, you can create a dedicated desktop action that reads each row from the Excel file and automatically inputs the data into the Shipping Management application. You can then assign this action to an AI Agent. When new shipping batches arrive, shipping coordinators simply trigger the AI Agent from the Now Assist panel. The agent then automatically processes each line item by selecting and running the appropriate desktop action.

Create Process shipping orders desktop action in Agentic Desktop

Automate various tasks related to shipping management through desktop actions in Agentic Desktop.

Before you begin

To access the Agentic Desktop functionality, perform the following steps:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Enable Agentic Desktop on your ServiceNow instance. For more information, see Configure Agentic Desktop.
- Download the Agentic Desktop installer to automate repetitive tasks across applications and systems. For more information, see Download Agentic Desktop installer.

Confirm that the following system requirements are met:

- Windows 11 operating system is used.
- A .NET 9.0 runtime v9.0.10 and .NET 9 Desktop Runtime v9.0.10 is installed.
- No extended monitors are connected.
- Theme must match between the systems used for recording and execution.

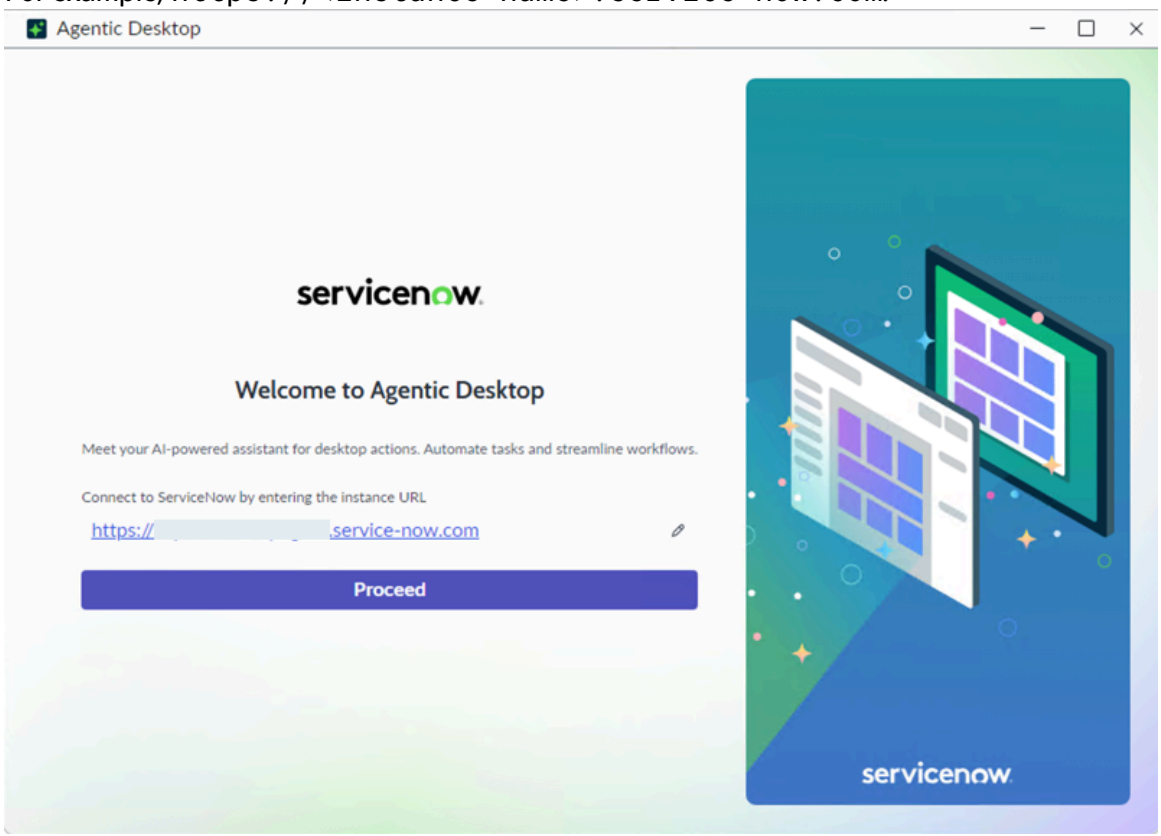
Familiarize yourself with the Design workspace and Action recorder. For more information, see Agentic Desktop Design workspace and Action recorder in Agentic Desktop.

Role required: sn_aia.admin

Procedure

1. From your Windows system, launch the Agentic Desktop application.
2. On the login page, in the **Add ServiceNow URL** field, enter the ServiceNow instance URL.

   Example
   For example, `https://<instance name>.service-now.com`.

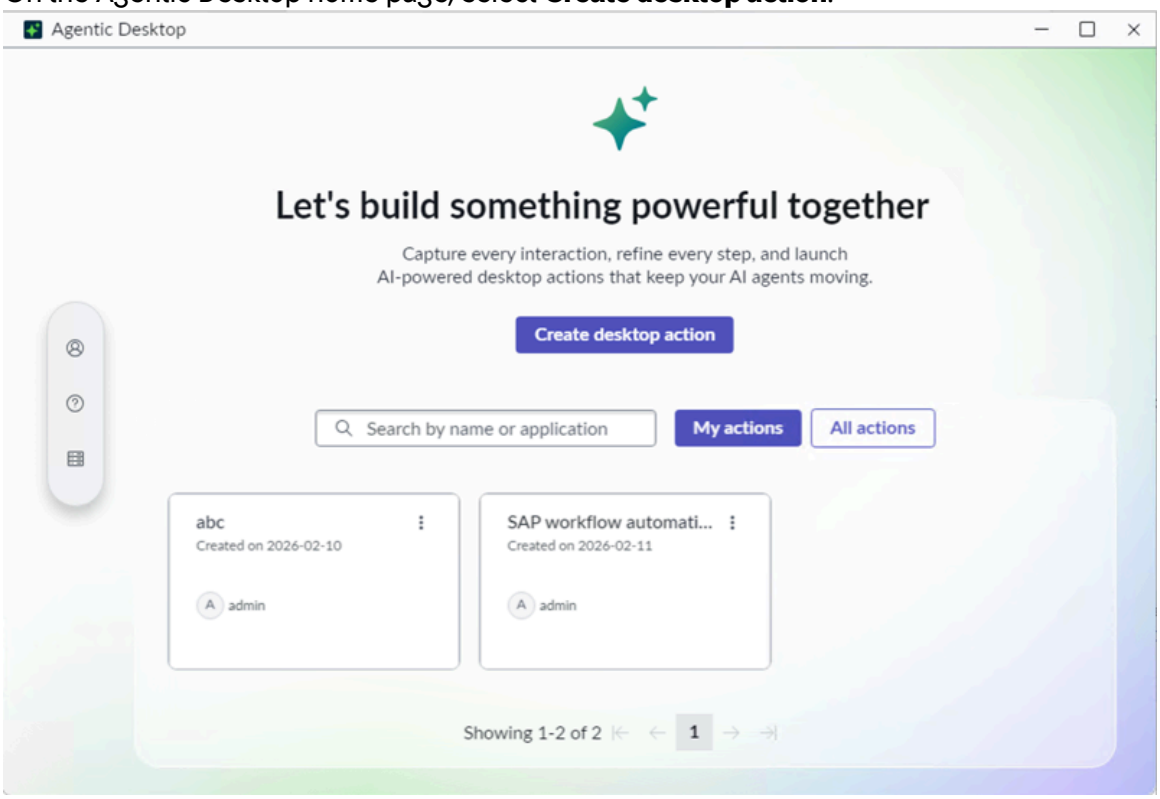


3. Select **Proceed**.
4. Log in to your ServiceNow account by entering your user name and password.
   Your account must have the sn_aia.admin role.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

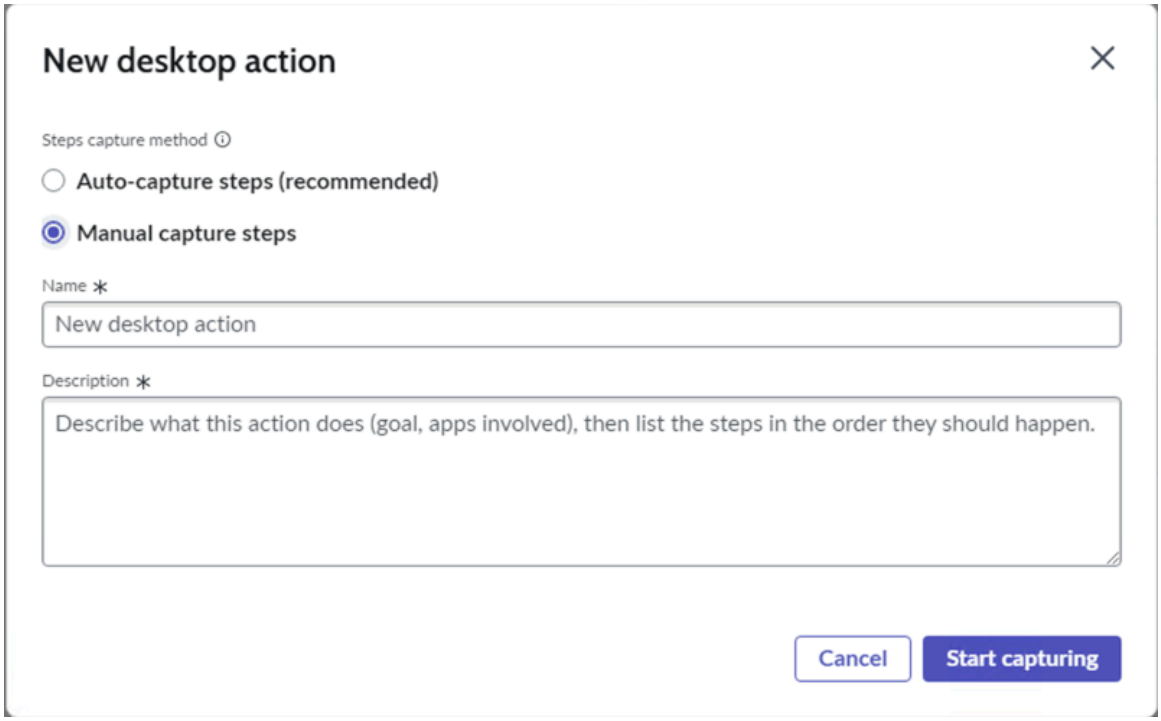


5. On the Agentic Desktop home page, select **Create desktop action**.



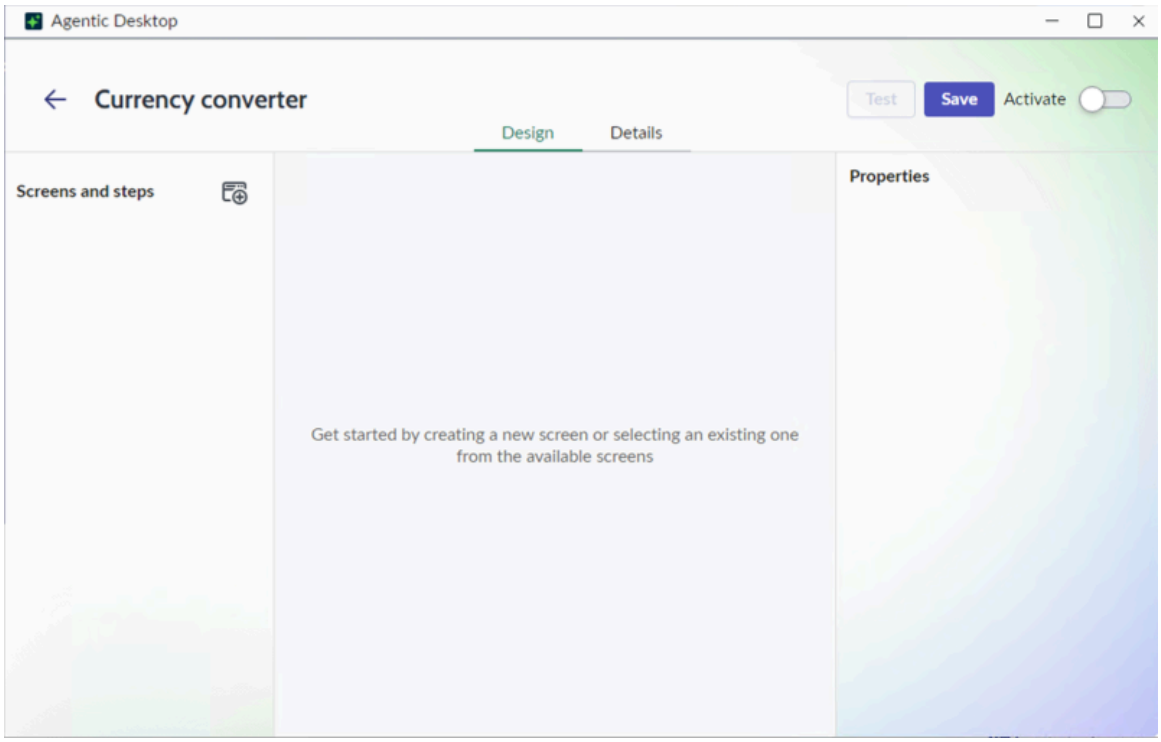
6. In the New desktop action dialog box, select **Manual capture steps**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



7. In the **Name** field, enter `Process shipping orders`.
8. In the **Description** field, enter `Enter Shipping order details from Excel to Shipping Management tool.`
9. Select **Start capturing**.

   The Design workspace is displayed.

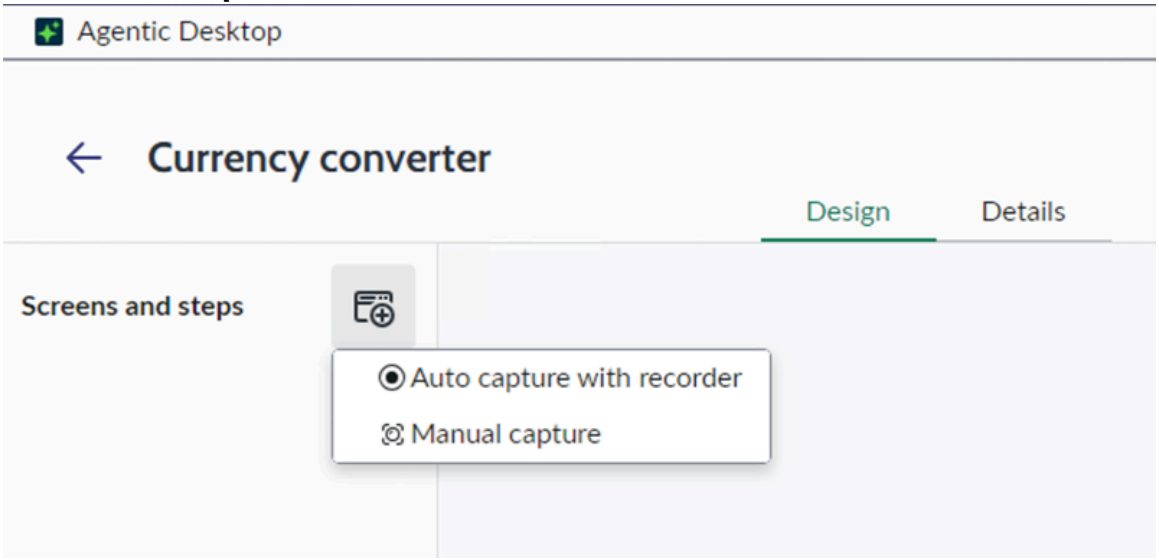


10. Capture screens.

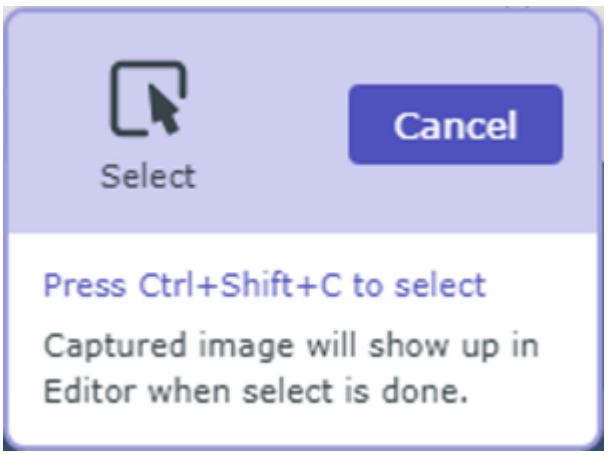
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the Design tab, select the **Capture Options** icon .

b. Select **Manual capture screens**.

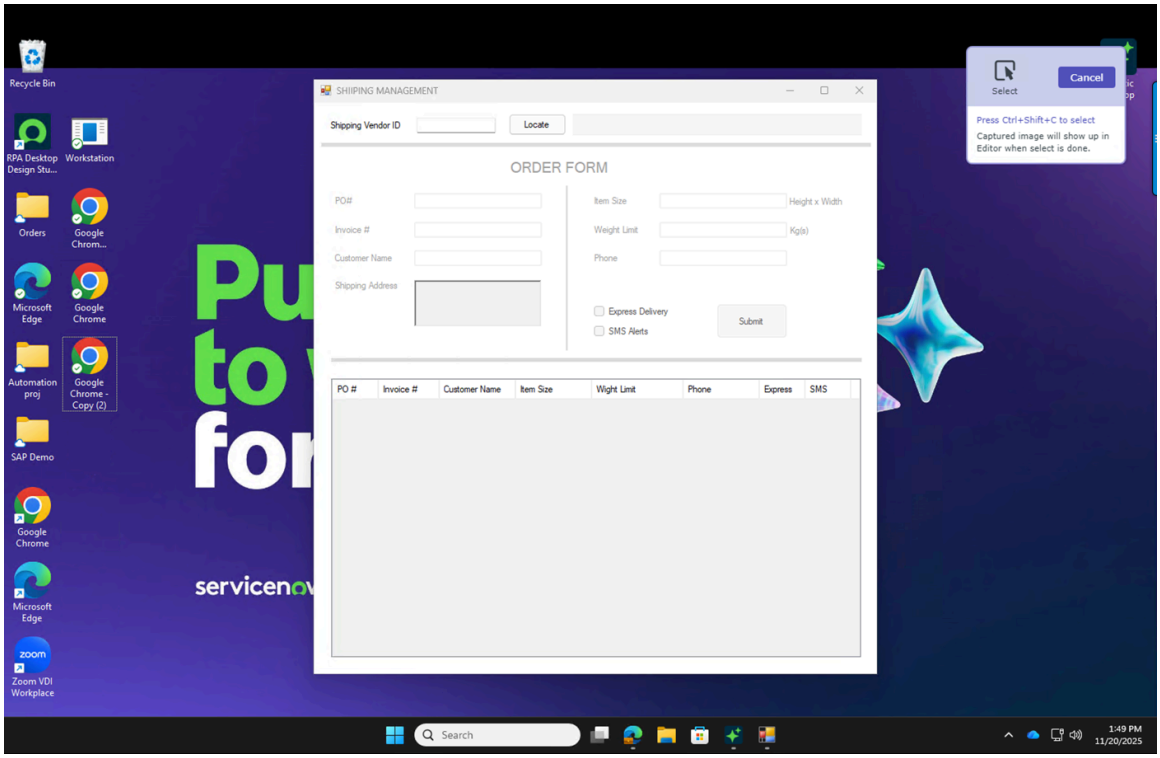


The Agentic Desktop window is minimized and the Capture panel is launched.



c. Open the Shipping Management application from your desktop.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



d. Capture the area of the external application's window by selecting the **Select** icon  on the Capture panel or pressing the `Ctrl + Shift + C` on the keyboard.

> **Note:** If the `Ctrl + Shift + C` shortcut conflicts with another application on the your machine, such as Zoom, then you must use the **Select** button to initiate manual screen capture.

For example, you can capture the area surrounding a button or a text field in a web browser.

The cursor icon changes to the + icon.

e. Drag the + icon and select the required screen area.

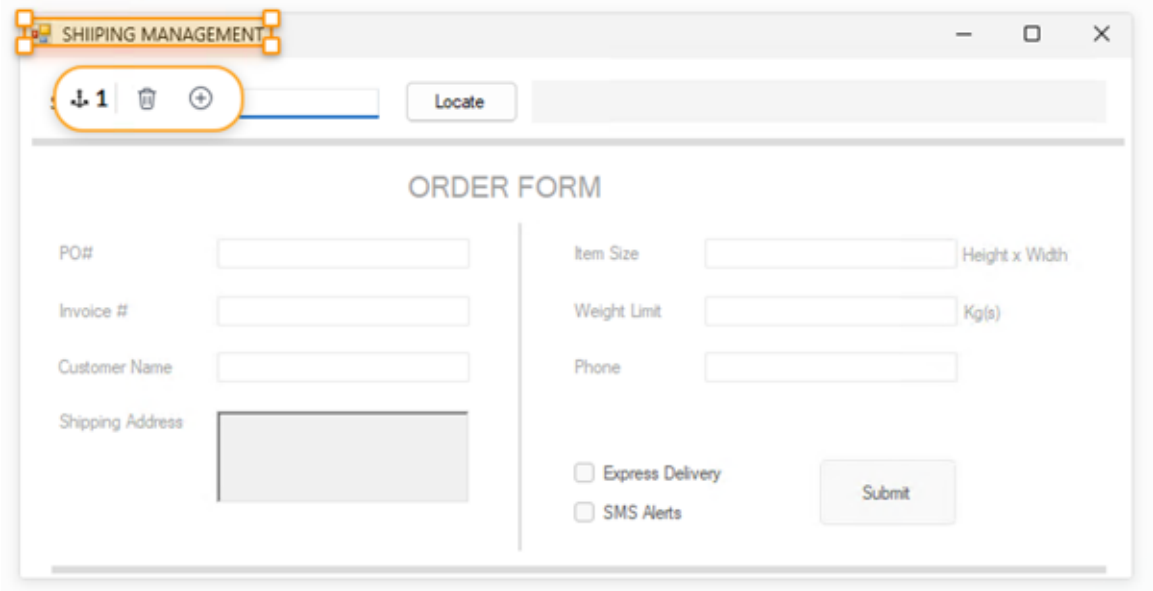
When you leave the + icon, the selected area is captured as a screenshot in the Design workspace.

If you are not satisfied with the captured screen, you can recapture the screen area by selecting the **Capture image** icon .

11. Insert anchors.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Insert an anchor on the captured screen by selecting the **Add anchor** icon .



An anchor is a reference point on the screen that helps the automation identify and interact with a nearby UI elements. During execution, the system uses computer vision to locate the anchor and then identifies the UI elements at a related distance from the anchor. Anchors improve the stability and accuracy of steps when the target element's location may shift or when the UI layout varies across sessions.

Note: Do not use dynamically changing UI elements as anchors. If an element changes its color, text, or state after an action (for example, after a click), select a different anchor that remains static on the screen.

b. Move the anchor to a part of the captured image that won't change.
For example, move the anchor to a title or field label.

If the area under the anchor doesn't exactly match the corresponding area of the captured image, the anchor isn't recognized, and the steps are not performed as intended. Choose a static area of the image for setting your anchor.

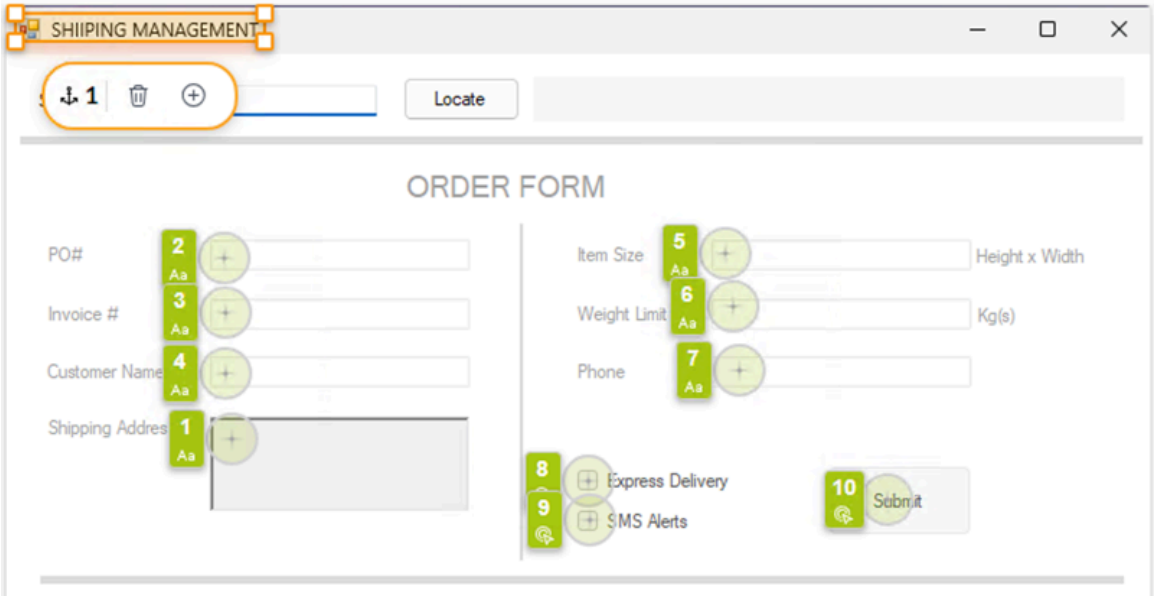
You can add multiple anchors on each screen. Multiple anchors let you define the geographical relationship between anchor and target with greater accuracy when targeting different locations in the image.

12. Configure the steps.

a. From the Step control menu, select the **Add step** icon .

b. Select the type of step for each field on the ORDER FORM from the contextual menu.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Number	Field	Step
1	PO#	Set Text
2	Invoice#	Set Text
3	Customer Name	Set Text
4	Shipping Address	Set Text
5	Item size	Set Text
6	Weight limit	Set Text
7	Phone	Set Text
8	Express Delivery	Click
9	SMS Alerts	Click
10	Submit	Click

13. Configure the properties for added screens, anchors, and steps.
For more information, see Screen, anchor, and step properties in Agentic Desktop.
14. Provide names for all added screens, anchors, and steps.

    When you create these elements, you can edit the auto-generated name, but follow these naming guidelines.

    - Name fields must not be empty.
    - Name fields must contain only alphanumeric characters. Spaces and special characters are not permitted.
    - Each name must be unique at its respective parent level.
        - Each screen must have a unique name at the desktop-action level.
        - Each anchor must have a unique name at the screen level.
        - Each step must have a unique name at the anchor level.
15. Select the Details tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

16. In the Applications list, add Shipping Management.
17. Select **Save**.
18. Test and activate the desktop action.
    For more information, see Test and activate a desktop action in Agentic Desktop.
19. Similarly, create and activate the **Shipping Management login** desktop action.

Create AI agents and add tools for shipping orders

Create an AI agent in AI Agent Studio and add desktop action tools for Shipping Management task automation.

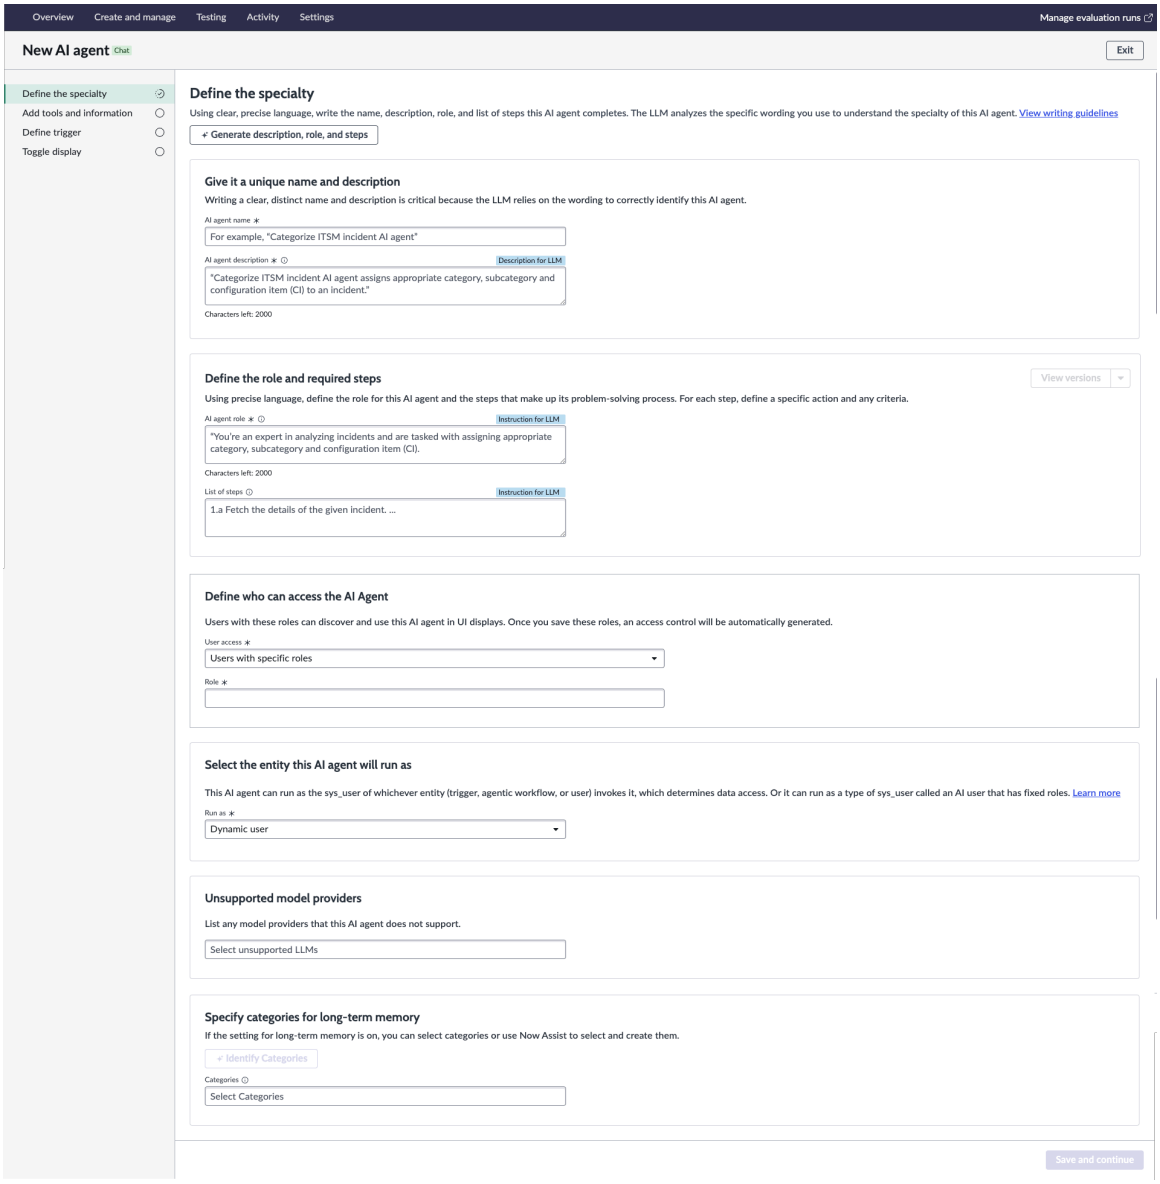
Before you begin

Role required: sn_aia.admin

Procedure

1. Navigate to **All** > **AI Agent Studio** > **Create and manage** > **AI agents**.
2. From the **Add** drop-down, select **Chat**.
3. On the New AI Agent page, in the Define the specialty step, define your AI agent and provide the specialties that this agent contains so that the LLM can analyze the wording you use to understand the purpose of the AI agent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Note: The more details that you provide, the more accurately your AI agent can perform.

a. Describe your AI Agent by giving it a unique name and description.

Give it a unique name and description

Field	Description
AI agent name	Shipping Management Agent
AI agent Description	The Shipping Management Agent automates end-to-end handling of shipping orders across multiple platforms. It retrieves order data from an Excel sheet and seamlessly uploads the information into the shipping management tool. This reduces manual effort, ensures data consistency, and accelerates the overall invoice life cycle.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. Define the role and necessary steps so that the AI agent can carry out its tasks.

Note: The AI agent uses this information as guidance to tailor its responses and actions.

Define the role and required steps

Field	Description
AI agent role	Automates the retrieval of shipping orders from Excel and performs seamless data entry into the shipping management application.
List of steps	**Note:** If your automation requires manual inputs, such as entering an OTP or CAPTCHA, you must provide instructions to the AI Agent to wait for the user input during execution. Otherwise, the automation can't proceed. **i.** Launch and log in to the Shipping Management application. **ii.** Open the Excel file from the user-specified folder. **iii.** Navigate to the designated worksheet within the Excel file. **iv.** Read the required columns for each order from the Excel sheet. **v.** Enter the extracted data into the corresponding fields in the shipping management application. **vi.** Select **Submit** to register the order. **vii.** Repeat steps 4–6 until all pending orders in the Excel sheet have been processed.

c. Select **Save and continue**.
You're directed to the Add tools and information page.

4. On the Add tools and information page, add Desktop actions tools for AI agents to automate your desktop tasks.

a. In the Add tool drop-down list, select **Desktop action**.

b. In the **Select a type of desktop action** field, select **UI block**

c. In the **Select a desktop action** drop-down list, select **Shipping Management login** desktop action.
This desktop action enables AI agents to log in to the Shipping Management application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. Provide a name and tool description for this desktop action configuration.
Tool description of the desktop action helps with what it's going to do to assist your AI agent.

> **Note:** This description is sent to the large language model (LLM).

e. In the **Execution mode** field, select **Autonomous**.

f. Select **Add**.
The desktop action is added in the Desktop actions list on the Add tools and information page.

g. In the Add tool drop-down list, select **Desktop action**.

h. In the **Select a type of desktop action** field, select **UI block**

i. In the **Select a desktop action** drop-down list, select **Process shipping orders** desktop action.
This desktop action enables AI agents to enter data extracted from the spreadsheet into the Shipping Management application.

j. Provide a name and tool description for this desktop action configuration.
Tool description of the desktop action helps with what it's going to do to assist your AI agent.

> **Note:** This description is sent to the large language model (LLM).

k. In the **Execution mode** field, select **Autonomous**.

l. Select **Add**.
The desktop action is added in the Desktop actions list on the Add tools and information page.

m. In the Add tool drop-down list, select **Desktop action**.

n. In the **Select a type of desktop action** field, select **Non UI block**

o. In the **Select an application** drop-down list, select Microsoft Excel.
The Microsoft Excel tool, non-UI block desktop action, enables AI agents to contextually extract information from the spreadsheet.

p. Provide a name and tool description for this desktop action configuration.
Tool description of the desktop action helps with what it's going to do to assist your AI agent.

> **Note:** This description is sent to the large language model (LLM).

q. In the **Execution mode** field, select **Autonomous**.

r. Select **Add**.
The desktop action is added in the Desktop actions list on the Add tools and information page.

5. Select **Save and continue**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. Complete the remaining steps that are necessary.
   For more information, see Create an AI agent.
7. Select **Save and test** to complete the configuration steps or review a previous step by selecting **Back**.

   Selecting Save and test leads you to the AI agent testing page, where you can test the AI agent that you created. For more information, see Manually test the execution of an AI agent.

   To test the AI agent, you must have the sn_aia.admin role and any roles the ACLs configured for the AI agent and its tools require, if applicable.

What to do next

Enable shipping coordinators to trigger AI agents from the Now Assist panel to enter data automatically into the Shipping Management application.

For more information, see Example: Use AI agents to automatically enter data into the shipping management app.

Examples of executing desktop actions using AI agents

These examples explain how to trigger AI Agents that execute automations you designed using desktop actions. The following topics walk you through the end-to-end execution flow so you understand how AI-driven execution works in your desktop environment.

Example: Use AI agents to process badge-related requests automatically

As an HR representative, automatically process various badge requests by triggering AI agents that use desktop action tools from the Now Assist panel.

Before you begin

To access the Agentic Desktop functionality, perform the following steps:

- Enable Agentic Desktop on your ServiceNow instance. For more information, see Configure Agentic Desktop.
- Download the Agentic Desktop installer to automate repetitive tasks across applications and systems. For more information, see Download Agentic Desktop installer.

Confirm that the following system requirements are met:

- Windows 11 operating system is used.
- A .NET 9.0 runtime v9.0.10 and .NET 9 Desktop Runtime v9.0.10 is installed.
- No extended monitors are connected.
- Remote Desktop must be enabled on your machine and your account must be granted Remote Desktop access permissions before you start using the Agentic Desktop Execution workspace.
- Theme must match between the systems used for recording and execution.

Note: Screen resolution and scaling must be the same between the systems used for recording and execution of desktop actions that are created prior to Agentic Desktop v1.0.1.

Familiarize yourself with the Agentic Desktop Execution workspace. For more information, see Agentic Desktop Execution workspace.

Role required: now_assist_panel_user

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

AI agents use desktop actions that are designed in the Agentic Desktop Design workspace as tools. When an AI agent is triggered from the Now Assist panel, it determines which desktop actions it can use to perform the automation. Once triggered, the automation is executed in the desktop-in-desktop mode within the Agentic Desktop Execution workspace.

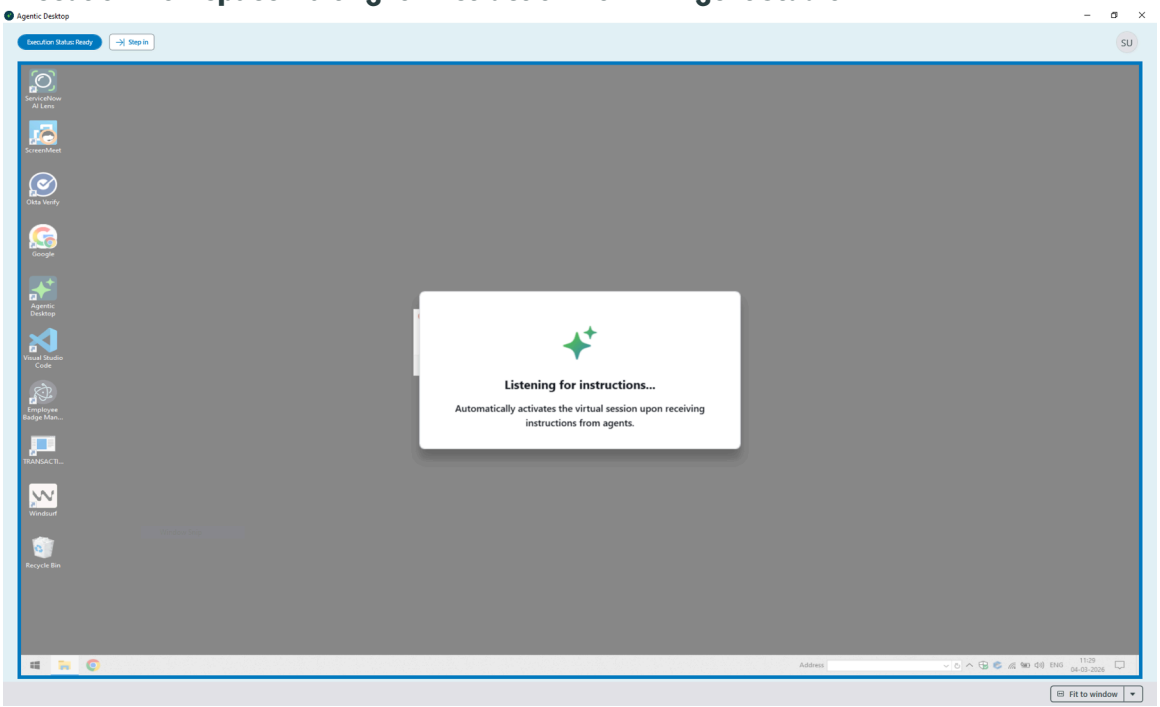
Note:

To avoid conflicts, do not run the Agentic Desktop Execution workspace and RPA Attended Desktop mode at the same time.

Procedure

1. Navigate to **All** > **Requests** > **Assigned to you** and select all the pending requests that you would like to resolve.
2. Open the Now Assist panel by using the Now Assist icon.
3. On the Now Assist panel, enter `Use the badge management tool and process the selected requests`.
4. Monitor the automation execution.

Execution workspace waiting for instruction from AI Agent Studio



Automation execution process on Now Assist panel and Agentic Desktop

On Now Assist panel	On Agentic Desktop Execution workspace
An AI agent is triggered from the Now Assist panel to complete your request in the Execution workspace.	The Execution workspace is launched.
The AI agent prepares a plan and shows which desktop actions it uses for the execution.	AI agent logs in to the Agentic Desktop Execution workspace. If it's the first time you launch the Execution workspace using AI agents, you must enter

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

On Now Assist panel	On Agentic Desktop Execution workspace
	the Windows Security credentials before proceeding. The Execution workspace waits for instructions from AI Agent Studio.
The AI agent shows the steps as it executes them in the panel. The outcome of execution is shown in the Now Assist panel.	Once the plan is received, the AI agent starts executing the automation in the Execution workspace. The AI agent performs the tasks same as a human would. The Execution workspace shows the status of the execution. For more information, see Execution statuses.

5. **Optional:** Interact with the automation when your inputs are required.
   - **Step in**: take control whenever human inputs are required
   - **Step out**: give the control back to the AI agent.

   Note: If your automation requires manual inputs, such as entering an OTP or CAPTCHA, you must provide instructions to the AI Agent to wait for the user input during execution. Otherwise, the automation can't proceed.

6. **Optional:** Use the smart sizing options to enable your desktop executions automatically adapt to your display.

Option	Description
Fit to window	Scales the execution screen to fit within the display area of the Execution workspace. The entire screen is visible without scrolling. (Optional) Shortcut: `ctrl+shift+w`
Original resolution	Displays the execution screen at its original resolution. Scroll bars appear if the screen is larger than the display area of the Execution workspace. (Optional) Shortcut: `ctrl+shift+d`

Trouble?

If the desktop session isn't sized correctly and mouse actions aren't working as expected after you enter the session, use the following keyboard shortcuts to resize the session:

`Ctrl + Shift + W`: Resize to window view.

`Ctrl + Shift + D`: Resize to actual desktop view.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example: Use AI agents to automatically enter data into the shipping management app

As a shipping coordinator, enter shipping details automatically from Excel to the Shipping Management app by triggering AI agents that use desktop action tools from the Now Assist panel.

Before you begin

To access the Agentic Desktop functionality, perform the following steps:

- Enable Agentic Desktop on your ServiceNow instance. For more information, see Configure Agentic Desktop.
- Download the Agentic Desktop installer to automate repetitive tasks across applications and systems. For more information, see Download Agentic Desktop installer.

Confirm that the following system requirements are met:

- Windows 11 operating system is used.
- A .NET 9.0 runtime v9.0.10 and .NET 9 Desktop Runtime v9.0.10 is installed.
- No extended monitors are connected.
- Remote Desktop must be enabled on your machine and your account must be granted Remote Desktop access permissions before you start using the Agentic Desktop Execution workspace.
- Theme must match between the systems used for recording and execution.

Note: Screen resolution and scaling must be the same between the systems used for recording and execution of desktop actions that are created prior to Agentic Desktop v1.0.1.

Familiarize yourself with the Agentic Desktop Execution workspace. For more information, see Agentic Desktop Execution workspace.

Role required: now_assist_panel_user

About this task

AI agents use desktop actions that are designed in the Agentic Desktop Design workspace as tools. When an AI agent is triggered from the Now Assist panel, it determines which desktop actions it can use to perform the automation. Once triggered, the automation is executed in the desktop-in-desktop mode within the Agentic Desktop Execution workspace.

Note:

To avoid conflicts, do not run the Agentic Desktop Execution workspace and RPA Attended Desktop mode at the same time.

Procedure

1. Navigate to **All** > **Incidents** > **Assigned to you** and open an incident that you would like to resolve.
   When you open an incident, the Now Assist application checks if a plan is available by using AI agents and displays the `Now Assist has a plan for solving INCXXXXXXX. Open Now Assist Panel to view the plan.` message in a banner.

   Note: You can select the banner and directly go to the conversation on the Now Assist panel to complete the task.

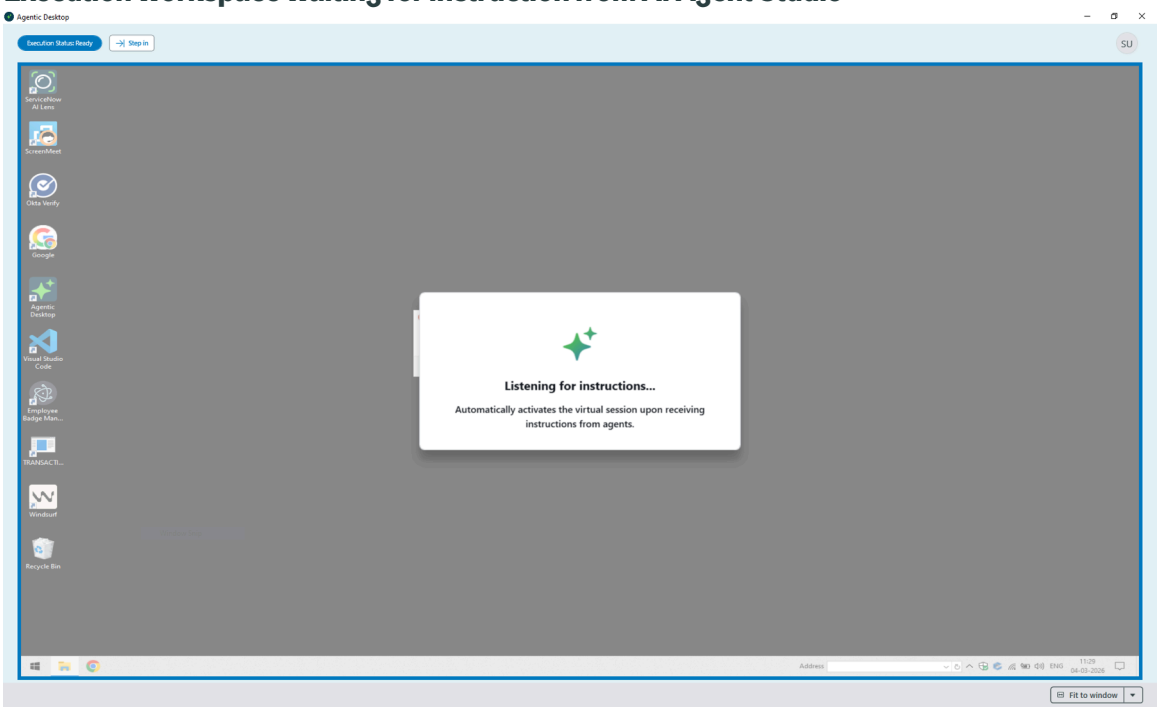
2. Open the Now Assist panel by using the Now Assist icon.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist provides the resolution steps for the incident.

3. On the Now Assist panel, enter `Read the shipping order details from the Excel spreadsheet located at: "C:\temp\Shipping orders.xlsx\". First, launch the Shipping Management application and login. Then, iterate through each order row in the spreadsheet and enter the data into the Shipping Management application.` .
4. Monitor the automation execution.

Execution workspace waiting for instruction from AI Agent Studio



Automation execution process on Now Assist panel and Agentic Desktop

On Now Assist panel	On Agentic Desktop Execution workspace
An AI agent is triggered from the Now Assist panel to complete your request in the Execution workspace.	The Execution workspace is launched.
The AI agent prepares a plan and shows which desktop actions it uses for the execution.	AI agent logs in to the Agentic Desktop Execution workspace. If it's the first time you launch the Execution workspace using AI agents, you must enter the Windows Security credentials before proceeding. The Execution workspace waits for instructions from AI Agent Studio.
The AI agent shows the steps as it executes them in the panel. The outcome of execution is shown in the Now Assist panel.	Once the plan is received, the AI agent starts executing the automation in the Execution workspace. The AI agent performs the tasks same as a human would.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

On Now Assist panel	On Agentic Desktop Execution workspace
	The Execution workspace shows the status of the execution. For more information, see Execution statuses.

5. **Optional:** Interact with the automation when your inputs are required.
   - **Step in**: take control whenever human inputs are required
   - **Step out**: give the control back to the AI agent.

   Note: If your automation requires manual inputs, such as entering an OTP or CAPTCHA, you must provide instructions to the AI Agent to wait for the user input during execution. Otherwise, the automation can't proceed.

6. **Optional:** Use the smart sizing options to enable your desktop executions automatically adapt to your display.

Option	Description
Fit to window	Scales the execution screen to fit within the display area of the Execution workspace. The entire screen is visible without scrolling. (Optional) Shortcut: `ctrl+shift+w`
Original resolution	Displays the execution screen at its original resolution. Scroll bars appear if the screen is larger than the display area of the Execution workspace. (Optional) Shortcut: `ctrl+shift+d`

Trouble?

If the desktop session isn't sized correctly and mouse actions aren't working as expected after you enter the session, use the following keyboard shortcuts to resize the session:

`Ctrl + Shift + W`: Resize to window view.

`Ctrl + Shift + D`: Resize to actual desktop view.

Agentic Desktop reference

Reference topics provide additional information about the roles and tables that are installed with the Agentic Desktop application.

Components installed with Agentic Desktop

Several types of components are installed with activation of the sn_desktop_agents plugin, including user roles and tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Roles installed

Role title [name]	Description	Contains roles
AI Agent Admin [sn_aia.admin]	Enables you to create, manage, and test desktop actions in Agentic Desktop and deploy them as tools in AI Agent Studio.	• agent_role_config_admin • sn_skill_builder.admin • connection_admin • flow_designer • sn_aia.viewer • sn_mcp_client.admin • sn_ace.ace_user
Now Assist panel user [now_assist_panel_user]	Enables you to trigger desktop actions from Now Assist panel and execute desktop actions in the desktop environment using Execution workspace.	sn_nowassist_admin.user

Tables installed

Table	Description
Desktop action [sn_desktop_core_action]	Stores details, such as input and output parameters and number of screens for desktop actions of type UI block.
Desktop action application [sn_desktop_core_action_application]	Stores the desktop application associated with the desktop action.
Desktop action execution [sn_desktop_core_action_execution]	Stores the execution details of a desktop action, such as request, response, tool execution, and execution plan.
Desktop action parameter [sn_desktop_core_action_parameter]	Stores the desktop action parameter records as parent records.
Desktop action parameter value [sn_desktop_core_action_parameter_value]	Stores the desktop action parameter value records as child records.

Plugins installed

Plugin	Description
Agentic Desktop [com.sn_desktop_agents]	Contains configuration of desktop action tool used by AI Agent Studio to discover available desktop actions while creating AI agents and invoke desktop actions via AI agents during execution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugin	Description
Agentic Desktop Core [com.sn_desktop_core]	Packages tables, custom functionality specific to Agentic Desktop application, and Agentic Desktop download interface within a scoped application.

System requirements and limitations in Agentic Desktop

Be aware of system requirements and a few limitations when you're using the Agentic Desktop application.

Hardware requirements

Hardware requirements for Agentic Desktop

Resource	Minimum	Recommended
CPU	Intel Processor (Core i5)	Intel Processor (Core i7)
RAM	4-GB RAM	8-GB RAM
Disk space	Minimum 20-GB free disk space after installing the OS, patches, and base software	Minimum 50-GB free disk space after installing the OS, patches, and base software

Software requirements

- Windows 11 operating system is used.
- A .NET 9.0 runtime v9.0.10 and .NET 9 Desktop Runtime v9.0.10 is installed.
- No extended monitors are connected.
- Add the end users who interact with the Execution workspace of Agentic Desktop to the Remote Desktop Users group on the target machine and provide Remote Desktop access permissions for seamless automation execution.

  If your organization uses Group Policy, add the end users to a Microsoft Active Directory group that is permitted to use Remote Desktop through Group Policy on each target machine where desktop actions run.

  - Local changes to the Remote Desktop Users group are temporary unless they align with Microsoft Active Directory entitlements.
  - If the user is not entitled, Group Policy refresh automatically removes them from the group.
- Theme must match between the systems used for recording and execution.

Note: Screen resolution and scaling must be the same between the systems used for recording and execution of desktop actions that are created prior to Agentic Desktop v1.0.1.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Limitations

Limitations with Agentic Desktop

Limit	Description
RPA Attended robot and Agentic Desktop parallel execution	To avoid conflicts, do not run the Agentic Desktop Execution workspace and RPA Attended Desktop mode at the same time.
Pause for human input	If your automation requires manual inputs, such as entering an OTP or CAPTCHA, you must provide instructions to the AI Agent to wait for user input during execution. Otherwise, the automation can't proceed.
Single execution	The system supports one automation execution at a time within the Execution workspace.
Screen elements	During recording, auto-assignment accuracy of screen elements depends heavily on the target application's technology and interface characteristics. Applications with strong accessibility support, such as well-defined UI elements and controls, stable rendering performance, and simple, consistent layouts enable anchors to be detected more reliably. In contrast, UI complexity, such as dynamic elements, frequent layout shifts, animations, or poor accessibility metadata can reduce accuracy and may require manual anchor configuration or additional tuning.
Getting table information	The **Get Table** step doesn't convert ordinary text to table data. For the step to capture table data successfully, the data must already be in the table form.
Browsers parallel execution	The Chrome or Edge browser must not be used in parallel in the main session and the desktop session on your system. You must close the session in the main session to execute automations that use the Chrome or Edge browser in the Execution workspace.
Sensitive information	The sensitive information isn't masked during recording. Confirm that you don't record any sensitive information.

Agentic Desktop glossary

Learn about the terms and concepts that are unique to Agentic Desktop.

A

Glossary terms are grouped alphabetically.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Action recorder

Recording feature in the Design workspace that auto-captures user interactions with desktop applications to create automated workflows. The recorder captures clicks, keystrokes, data entry, and visual and contextual information as screens and steps.

AI agent

Set of large language model (LLM) instructions and tools in AI Agent Studio that can process natural language inputs, generate execution plans, and run desktop actions autonomously or semi-autonomously.

AI Agent Studio

ServiceNow platform interface for creating, managing, testing, and activating AI agents. Desktop actions are published to AI Agent Studio as tools that AI agents can use during execution.

Anchor

Reference point on a captured screen that helps the automation identify and interact with nearby UI elements. During execution, the system uses computer vision to locate the anchor and then identifies UI elements at a relative distance from it. Anchors improve the stability and accuracy of steps when the target element location may shift or when the UI layout varies across sessions.

Auto-capture

Method of creating desktop actions by recording user interactions with desktop applications using the Action recorder. The recorder captures clicks, keystrokes, and data entry along with visual and contextual information.

C

Glossary terms are grouped alphabetically.

Confidence threshold

Accuracy level required for matching a captured image before the system performs a step. A value of 1 defines a 100% match while 0.5 defines a 50% match. The default value is 0.95 (95% match).

D

Glossary terms are grouped alphabetically.

Design workspace

Interactive no-code environment within Agentic Desktop for creating, configuring, managing, and testing desktop actions. The workspace provides a visual canvas where you can design multi-screen automation workflows that capture business processes across different applications. Accessible to users with the AI Agent Admin (sn_aia.admin) role.

Desktop action application

Record that stores the association between a desktop action and the desktop application it interacts with. Stored in the Desktop action application [sn_desktop_core_action_application] table.

Desktop action

Reusable automation that defines how AI agents interact with desktop applications. Desktop actions consist of screens, anchors, and steps. There are two types: UI block and non-UI block.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Desktop-in-Desktop (DiD) mode

Virtual environment within the Execution workspace where automations run in isolation from the main desktop session. You can monitor the execution of desktop actions and how they interact with desktop applications while continuing to work on other tasks.

Desktop session

Isolated Windows session within the Execution workspace where desktop actions run. The desktop session launches automatically when you test a desktop action or trigger an automation from the Now Assist panel.

E

Glossary terms are grouped alphabetically.

Execution status

Current state of an automation in the Execution workspace. Statuses include Ready, Initiated, Running, Success, Failed, and Canceled.

Execution workspace

Isolated desktop session where desktop actions run during testing or AI agent execution. This workspace launches automatically when you test a desktop action from the Design workspace or trigger an automation from the Now Assist panel. You do not open the Execution workspace directly.

M

Glossary terms are grouped alphabetically.

Manual capture

Method of creating desktop actions by taking screen captures, inserting anchors, and defining steps individually rather than recording interactions automatically.

N

Glossary terms are grouped alphabetically.

Non-UI block

Type of desktop action that uses pre-built connectors to interact with applications and system components in the background without UI interaction. Supported applications include Microsoft Excel, Microsoft Outlook, Microsoft Word, PDF, PowerShell, SQL, SSH, and System Actions. Non-UI block actions can't be created by users.

Now Assist panel

ServiceNow interface from which users trigger AI agent automations. When you provide instructions through the Now Assist panel, the AI agent selects and runs the appropriate desktop actions in the Execution workspace. Accessible to users with the Now Assist panel user (now_assist_panel_user) role.

P

Glossary terms are grouped alphabetically.

S

Glossary terms are grouped alphabetically.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Screen

Captured image of an application window within a desktop action. Screens represent the application states that the automation moves through during execution. Each screen contains one or more anchors and associated steps.

Step

Individual interaction within a desktop action, such as clicking a button, entering text, or extracting data. Steps are positioned relative to an anchor.

- Input step types include Set Text, Click, Mouse Click, and Send Keys.
- Output step types include Get Text, Get Table, and OCR Read Text.

Step in / Step out

Manual control options in the Execution workspace. Step in transfers control to the user for manual input such as entering an OTP or CAPTCHA. Step out returns control to the AI agent to continue the automation.

T

Glossary terms are grouped alphabetically.

Tool

Desktop action that has been activated and added to an AI agent in AI Agent Studio. Tools provide AI agents with the capabilities to complete specific tasks during execution. An AI agent selects a tool based on its name and description.

U

Glossary terms are grouped alphabetically.

UI block

Type of desktop action that simulates human interactions with UI elements on thick client applications, legacy systems, or SaaS applications without APIs. Interactions include clicking buttons, entering text, and selecting from menus. UI block actions are created and managed in the Design workspace.

Predictive Intelligence

Predictive Intelligence is a powerful interface to train machine learning models. With Predictive Intelligence, you can improve performance, efficiency, and flexibility to your systems across multiple business units.

Get started with Predictive Intelligence

Administrators can harness the power of machine learning to improve productivity and efficiency for their agents and fulfillers. Predictive Intelligence uses artificial intelligence to improve processes across the platform. Predictive Intelligence enables you to do things such as the following:

- Automatically populate fields during case creation.
- Categorize and route work based on how records have been handled in the past.
- Recommend resolutions to cases that are similar to previous ones.
- Identify major outages based on incoming incidents.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

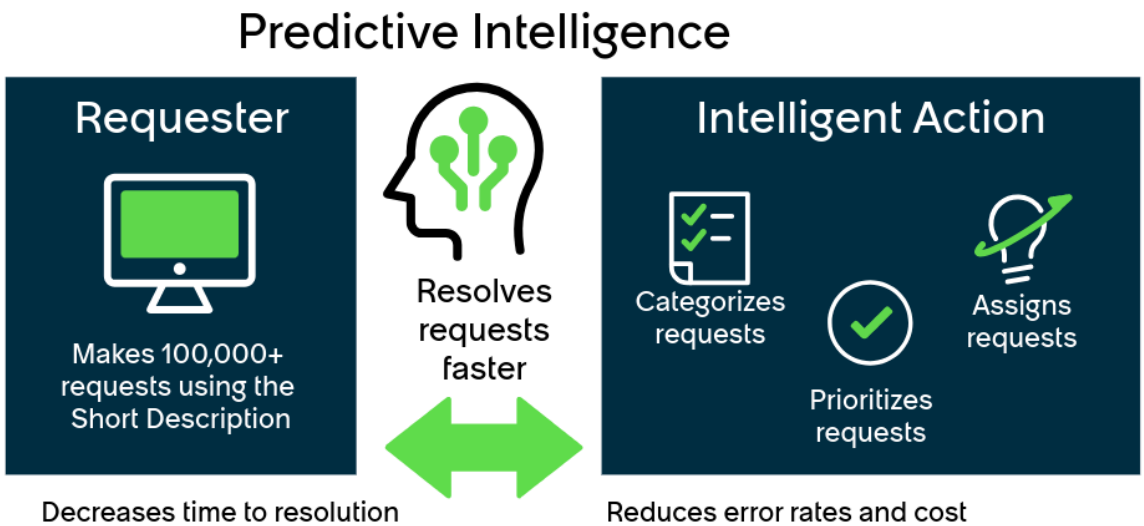
Explore Learn about Predictive Intelligence and machine learning	**Install** Install Predictive Intelligence and its associated apps	**Configure** Configure Predictive Intelligence on the platform
Train Create and train machine learning solutions.	**Use** Use Predictive Intelligence for machine learning.	**Reference** Get details about Predictive Intelligence components such as properties and language support.

Explore Predictive Intelligence

ServiceNow® Predictive Intelligence is a platform function that provides a layer of artificial intelligence that empowers features and capabilities across ServiceNow® applications to provide better work experiences.

Overview of Predictive Intelligence

Predictive Intelligence is a powerful set of tools applying artificial intelligence and machine learning to make predictions. You can create and train models in three different frameworks: classification, clustering, and similarity. A trained solution can be invoked by any ServiceNow application through an API.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To learn more about ways to use existing models, see Using Predictive Intelligence.

Predictive Intelligence for on-premise customers

Predictive Intelligence is also available for on-premise customers. If you're interested in deploying this product on-premise, contact your account manager. For on-premise installation and configuration instructions, see the complete instructions for Machine Learning Engine installation and configuration for self-hosted customers [KB0782052] in the Now Support Self-Hosted Knowledge Base.

Note: Only on-premise accounts can access the Now Support Self-Hosted Knowledge Base.

Terminology

Artificial intelligence

Systems designed to do work that needs a level of human intelligence to accomplish.

Machine learning

Ability for models to improve over time with more experience.

Models

Collections of algorithms, math, and statistics that make predictions and decisions based on input-output data.

Training

Adding or changing data that the model is based on to affect future predictions.

Supervised Training

Providing input-out pairs so that the model can generate rules that connect the two.

Unsupervised Training

Providing raw data so that the model can identify structures in the data set.

Training frequency

How often models are retrained to incorporate new data into an existing model.

Word corpus

Vocabulary that a model can use to look for textual similarity.

Predictive model components

A predictive model includes these components, some of which you must provide.

Solution definition

A data record you create and configure that specifies these values for training a predictive model.

- The records used to train the model. For example, only train on incidents that are resolved or closed within the last six months.
- The input fields that the model uses to make predictions. For example, use the incident short description to make a prediction.
- The output field whose value the model predicts. For example, set the incident category based on the short description.
- The frequency to retrain the model. For example, retrain the model every 30 days.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Solution

The solution is the result of a solution definition that you've trained in a ServiceNow datacenter. Predictive Intelligence uses the solution to predict a target field value given one or more input field values. All solutions specify these values.

- The solution precision is the aggregate percentage of correct predictions. For example, a precision of 50 means that out of 100 predictions, half of them should have the correct value.
- The solution coverage is the aggregate percentage of records that receive a prediction. For example, a coverage of 50 means half of all eligible records actually receive a prediction.
- The solution classes are the output field values for which the model can make predictions. Each class is an output field value with a list of possible precision, coverage, and distribution metrics to choose from. For example, the Incident Categorization solution has a class for each category such as software, inquiry, and database.
- The class distribution is the percentage of records from the entire table that have this particular output field value. For example, a distribution of 50 for the inquiry class means that half of incidents have the inquiry category.

Predictive Intelligence frameworks

Predictive Intelligence provides three different model frameworks in the Australia release: classification, similarity, and clustering. Each framework specializes in different types of predictions.

Predictive Intelligence classification framework

The Predictive Intelligence classification framework enables you to use machine-learning algorithms to set categorical field values during record creation. For example, you can use the model to set the incident category based on the short description. You can train predictive models so they act as an agent to categorize and route work automatically based on your past record-handling experience.

Enable Predictive Intelligence to handle volumes of incoming requests at lower costs. Automate the categorization and assignment of requests to reduce:

- Task resolution times.
- The number of interactions required to resolve tasks.
- The error rates of categorizing and assigning work.

For more information, see Create and train a classification solution.

Predictive Intelligence similarity framework

The Predictive Intelligence similarity framework identifies existing records that have similar values to a new record. For example, you can train a subset of your incident records to recommend a resolution based on the information of a similar incident record. By borrowing from similar closed incidents that have a proven resolution, you can help agents and fulfillers quickly provide the best resolution for an incoming incident.

The similarity framework doesn't need an exact match of keywords for its text comparisons because its algorithms identify similar words and synonyms based on similar contexts. For example, the phrases printer not working and printer broken are both recognized as similar. The framework also collects, learns, and applies your industry-specific context. For example, the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

phrase unable to join network has a different context in a computer networking company than it does in a healthcare insurance company.

The similarity framework uses a workflow similarity solution. For more information, see Create and train a similarity solution.

Predictive Intelligence clustering framework

Clustering divides data into groups that can then be used to identify patterns. You can then address records collectively or find gaps in existing data. For example, you can group similar new incidents to identify a major outage.

The clustering framework uses a workflow clustering solution. For more information, see Create and train a clustering solution.

Deprecated in the Washington DC release: Predictive Intelligence regression framework

Important: Support for creating new regression solutions was removed in the Washington DC release. You can train and edit existing solutions, but you can't create new ones. This information is provided for legacy context.

Regression is a machine-learning framework that uses historic data to predict numeric outputs, such as a temperature or a stock price.

For more information, see Create and train a regression solution.

ServiceNow® apps and features that use Predictive Intelligence

Learn about ServiceNow applications and features that leverage Predictive Intelligence. Solutions that you can adapt are available for various business units and industries.

ServiceNow teams work together to offer models that apply the artificial intelligence and machine learning capabilities provided by Predictive Intelligence. For example, Customer Service Management (CSM) partners with Predictive Intelligence to deliver BU-specific solutions to ServiceNow CSM customers.

Following is a listing of some ServiceNow products that use Predictive Intelligence functionality.

- Machine learning solutions for Customer Service Management
- Machine learning solutions for Event Management
- Issue assignment using the Governance, Risk, and Compliance Predictive Intelligence plugin
- Machine learning solutions for HR Service Delivery
- Predictive Intelligence for Innovation Management
- Machine learning solutions for IT Service Management
- Machine learning solutions for Knowledge Management
- Machine learning solutions for Search administration
- Machine Learning solutions for Vulnerability Response
- Work order insights powered by Predictive Intelligence
- Machine Learning solutions for Flow Designer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Install Predictive Intelligence

Activate Predictive Intelligence on your instance and get started with basic configuration.

Before you begin
Role required: admin

About this task
The Predictive Intelligence plugin (com.glide.platform_ml) is included in the base system, but if necessary you can use the following procedure to activate it.

Other related plugins and store applications are available. For example, Predictive Intelligence for Field Service Management (com.snc.fsm_ml) provides solutions relevant to FSM. Some of these plugins may require a separate license.

When you activate the Predictive Intelligence plugin for the first time on your instance, the system launches a Homepage. The Homepage provides an overview of your solution definitions in the classification, similarity, clustering, and regression frameworks. You can create, train, and test solutions directly on the Homepage. A summary of the latest trained solution is also available.

Procedure

1. Navigate to **System Definition** > **Plugins**.
2. Use the search bar to locate the Predictive Intelligence (com.glide.platform_ml) plugin.
3. Select **Install** and then in the Activate Plugin dialog box, select **Activate**.
   When you activate a plugin, any dependent plugins are activated automatically.
4. Confirm that the activation has successfully created a sharedservice.worker user.
   When training your solutions, Predictive Intelligence operates as this user.

   Note: The sharedservice.worker user includes the following roles:
   - platform_ml_read
   - platform_ml_write
   - platform_ml_create

   These roles are required to create, train, and view solutions. They are internal roles and not meant to be edited or assigned to other users.

Implement Predictive Intelligence

Implement initial setup and configuration steps for Predictive Intelligence to train a machine-learning (ML) algorithm to make predictions based on your past record data.

Before you begin
Role required: admin or ml_admin

About this task
The training process requires sending record data to a training service in the nearest datacenter. Since every datacenter has its own dedicated training server and the data doesn't leave the datacenter, this service is also available to customers who have data sovereignty requirements. For more information on this process, see Explore Predictive Intelligence.

For frequently asked questions regarding initial configuration and setup, see KB0781894.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Activate Predictive Intelligence on a non-production instance.
2. From your production instance, export the records that you want your Predictive Intelligence solutions to process.
   For example, export 12 months of incident records to a non-production instance.
3. On the non-production instance, import the records you exported.
4. On the non-production instance, review the default solution definition records to determine if the filter, input fields, and output field are sufficient to predict your incident or task records.
   If necessary, create a solution definition for each record set you want to predict.
5. On the non-production instance, train the solution definition records.
6. Test the solution predictions on the non-production instance by either creating test records or importing more records from production.
7. For classification solutions, review the prediction reports to determine the accuracy and coverage of your solution and individual classes.
8. For similarity solutions, review the similarity examples to update the similarity score threshold if needed.
9. If necessary, update the solution definition filter to include more or different training records.
10. Retrain and retest any updated solution definition records.
11. When you are satisfied with your solutions, activate Predictive Intelligence on your production instance.
12. Recreate any custom solution definition records and train the solution, or import the solution from your non-production instance to your production instance.

Related topics

Create and train a classification solution

Create and train a similarity solution

Configure Predictive Intelligence

Configure solution definitions for Predictive Intelligence to train, test, and monitor solution definitions

Configuration

Use the following links to guide you as you set up your solution definition.

Database View support for Predictive Intelligence

Use database views to join two or more tables as input for your Machine Learning (ML) solution.

Using database views in an ML solution

Database views help expand the amount of fields your solution can use for training. By using more than one table in your ML solution definition form, you can access more input data to help enrich the solution outcome.

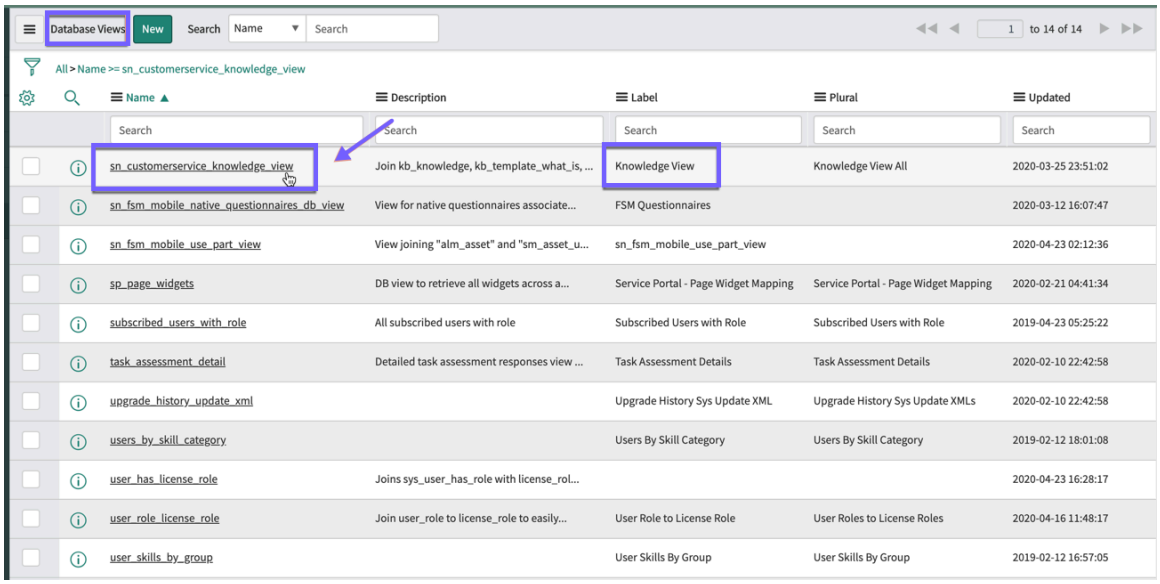
Database views enable you to join two or more tables into one consolidated view. For this to work in an ML solution you must first create the database view. See Database views.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Database views are supported in all four of the Predictive Intelligence capability frameworks: classification, similarity, clustering, and regression.

A database view example for Predictive Intelligence

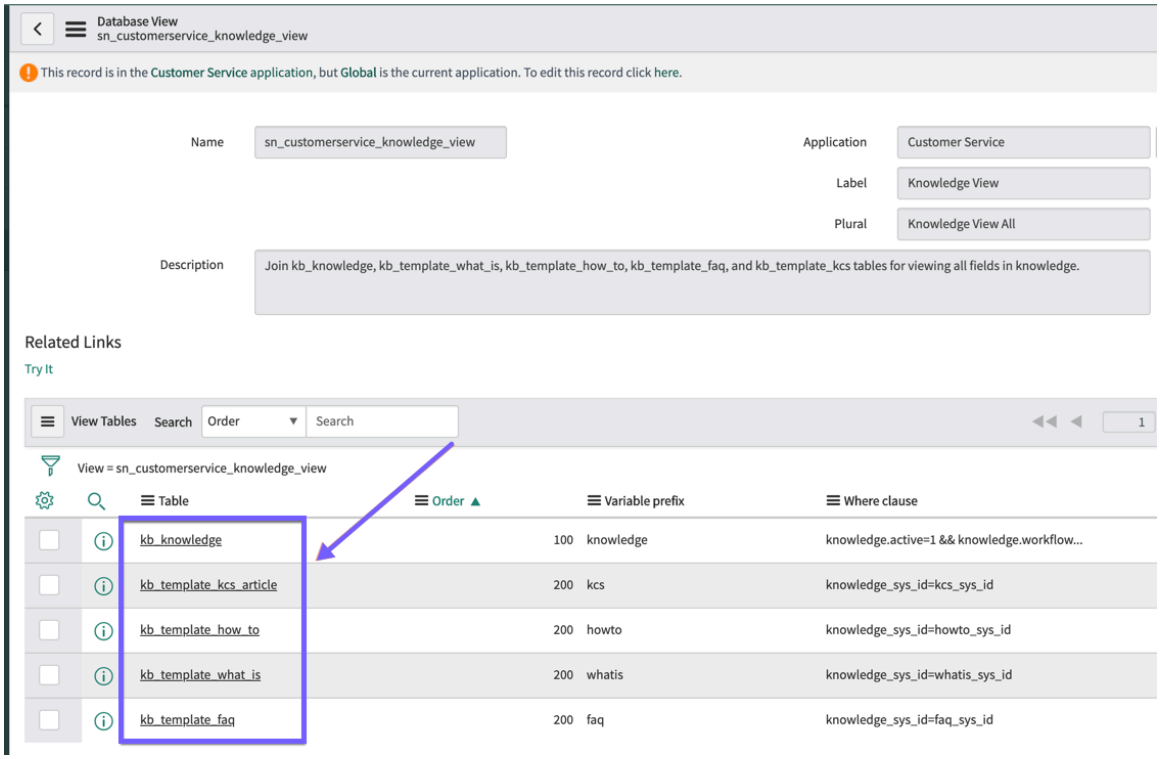
In the following example scenario, you've created a database view for use in an ML similarity solution. The image below shows the database view record you've created, including its **Name** and **Label**.



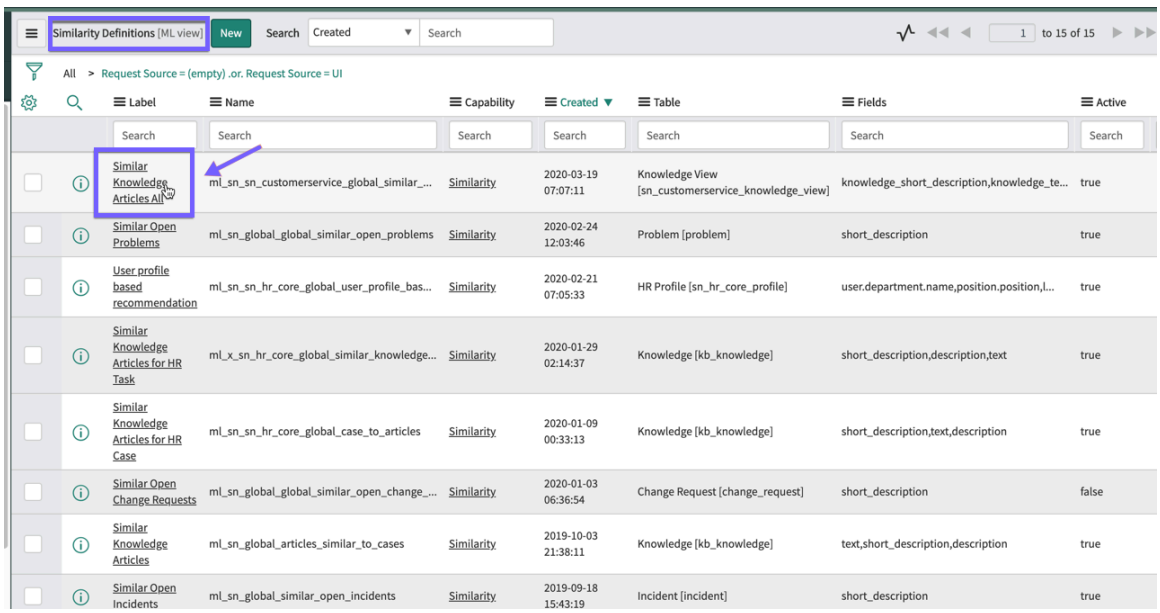
When you click the record **Name**, its content appears, as shown in the image below. Within the database view content, the five Knowledge Base tables you've joined to the view are listed. Most of these tables contain different Knowledge article template types, such as an FAQ or a How To article.

When a database view is used as input to a similarity solution, each of the tables that constitute the view must have at least the required number of records set in the configuration of your ServiceNow instance. The default minimum number required is 10,000. For example, the **Knowledge View** database view has five tables and each table must have 10,000 records. If a table doesn't have 10,000 records, you may not see the results from that table. If you must change that value, contact Customer Support.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



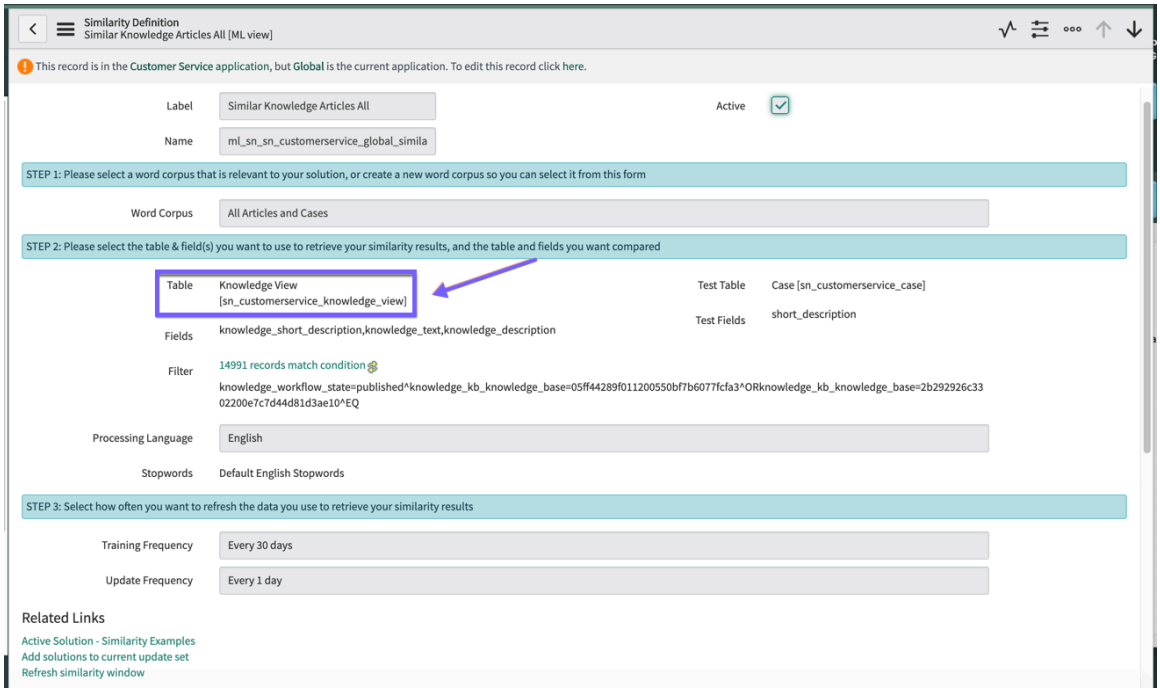
In the image below you can see the similarity solution definition record you've already created, which you plan to associate to your database view. When you click the **Label** for your similarity record, its Similarity Definition form appears.



Similarity Definition forms compare your existing table records based on their similarity by using a table in the **Table** field and another table in the **Test Table** field.

To use a database view in your similarity solution, instead of selecting a table in the **Table** field, you select the database view you created, which in this example scenario is the **Knowledge View** database view. This configuration increases the number of records your solution uses in training because the system compares and processes five tables of data instead of one.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Configuration tips for Predictive Intelligence

If you encounter issues during your solution training and solution prediction, follow these suggested resolutions.

Input Data

It is recommended to have at least 30,000 records to train your models with, but the accuracy of the model is determined by the input data.

There are three primary factors that determine the quality of the input data used to train solutions:

- Cleanliness: Sanitized data reduces noise, making the model more accurate.
- Quality: The input and output should be valid and correct to train the model to make accurate predictions.
- Distribution: Data that represents the entire dataset as a whole will result in a model that can make more generalized predictions.

Most raw data sets contain dirty and unusable data. Reviewing your input sets before training is essential to keeping accurate predictive models.

It is recommended to use approximately 80% of your input data to train your model and about 20% of the data to evaluate whether the model is accurate. You can compare the model's predicted results against the real values for the 20% of remaining data.

Solution training

Issue	Resolution or suggested action
The solution training remains in Waiting for Training status for too long, as the scheduler job is using an incorrect Glide callback instance URL.	Ensure the *glide.servlet.uri* property in the Glide instance is set to the correct instance URL. This issue can occur when:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Issue	Resolution or suggested action
	• An instance is cloned from production, yet it still refers to the production URL for the *glide.servlet.uri* property. • The Glide instance is provisioned and runs the training for the first time.
New categories have been added and aren't yet having an impact on training.	This is expected behavior, as the new categories may not yet have sufficient data until the solution is retrained.
The solution training fails.	When the training fails, click the **Show Training Progress** related link on the solution screen to determine where the potential problem resides.
The solution training fails due to user authentication.	Navigate to **System Security> Users** and ensure the sharedservice.worker user is set to Active.
The model training returns saying the model cannot be created. The training fails and shows the “Error while training solution” message. The training progress window shows this message: “Solution training failed as either the data used isn't sufficient or the input field isn't predictive of the output field."	This issue can occur when the data quantity or the distribution of field values isn't sufficient for a model to build successfully. Follow these steps to troubleshoot: **1.** Ensure the distribution of the output field isn't skewed. **2.** Retrain the model by changing the date filters to use a larger amount of data. **3.** If the input fields are not fully populated, add a filter to remove null records.
The solution has data in multiple languages but the coverage and precision results are poor.	Use the following options to help improve your metrics. Option 1: Update the processing language of the solution to the most prominent non-English language. **Note:** English is applied by default for all datasets. Option 2: If there's sufficient data for each language/region: **1.** Add a filter criteria for a specific language/region where the primary language can be identified (Dutch, English, French, German, Japanese, or Spanish). **2.** Generate a solution for each language/region and apply the proper processing language to each solution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Solution prediction

Issue	Resolution or suggested action
The prediction fails and returns a Java exception where the cause is unknown.	1. Search for the exception in the Predictive Intelligence Glide logs. 2. Submit an Incident record for Predictive Intelligence including all relevant details, such as the exception, the impacted instance, the solution name, and the input string.
There is no prediction applied to the incident/case record but the prediction returns a value when tested in the Rest API Explorer.	This can occur when the confidence of the prediction is less than the threshold required to make a prediction. After your solution is trained, use the following steps to confirm if your solution settings need adjusting. 1. Navigate to **System Web Services > REST > REST API Explorer** to find the confidence level for the prediction. See Test a classification solution prediction. 2. On your ML Solution Definition record, check the threshold set for your outcome class that was returned in the prediction by clicking on the name of the class. The **Class** page appears. 3. Check the **Estimated Precision** and **Estimated Coverage** values. If the corresponding threshold is more than the prediction confidence of the outcome, this is the root cause for why you did not see any prediction. 4. Adjust the class precision and coverage values to increase coverage or precision. See Tune a trained classification solution.

Instance cloning

Issue	Resolution or suggested action
After an instance is cloned, predictions for your existing solutions fail.	The ML solution artifacts in the [ml_artifacts] table are stored in the [sys_attachment table]. If the [ml_artifacts] table isn't included in the clone when you run it, the predictions fail. Ensure your clone includes the machine-learning artifacts, as these are critical components of your Predictive Intelligence solution.
After an instance is cloned, the solution training fails.	As the cloning run proceeds, it is possible that the sharedservice.worker user has either been

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Issue	Resolution or suggested action
	inactivated, locked out, or the user ID isn't set. Resolve these problems so that the solution training succeeds.

Create a custom stopwords list

Exclude common words you want the system to ignore during training and prediction.

Before you begin

Role required: admin or ml_admin

About this task

Stopwords lists enable the system to exclude extraneous words that can impede search and the overall natural language processing of your data.

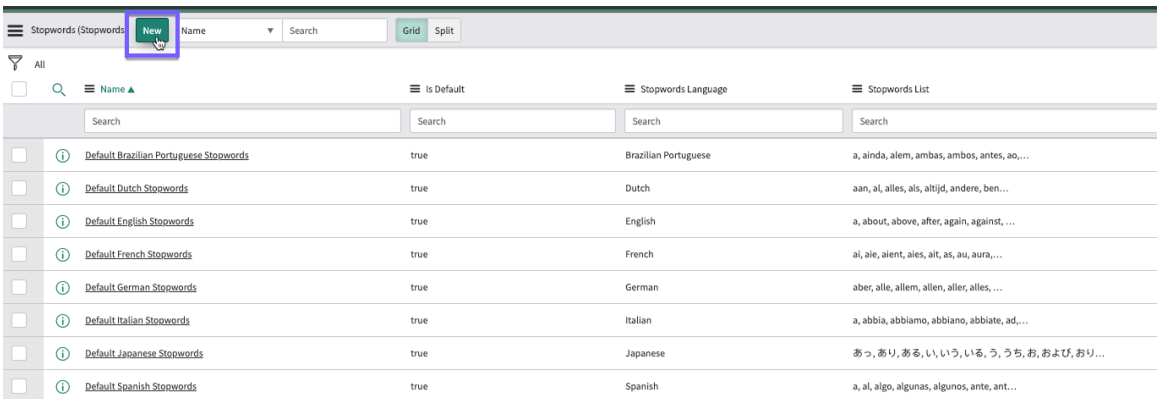
Predictive Intelligence provides you with default stopwords lists for each language the system supports. Examples of stopwords include words such as in, the, and the names of people and companies. You can also define your own stopwords list that's comprised of words specific to your organization and industry.

The custom list you provide works alongside those that the system already uses by default. For example, if incident records are used in a classification solution, and a company name is used in those records, consider adding that name to your list, as it's unlikely to provide any relevant information for the solution you're building.

In this example scenario, you create a custom stopwords list for the Brazilian Portugese language.

Procedure

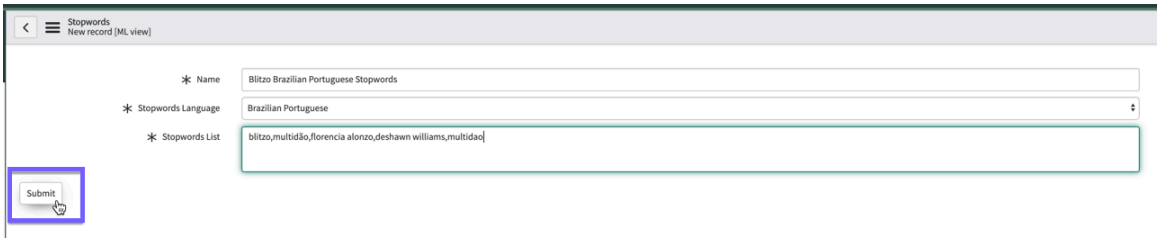
1. Navigate to **All** > **Predictive Intelligence** > **Stopwords**.
2. On the Stopwords list, click **New**.



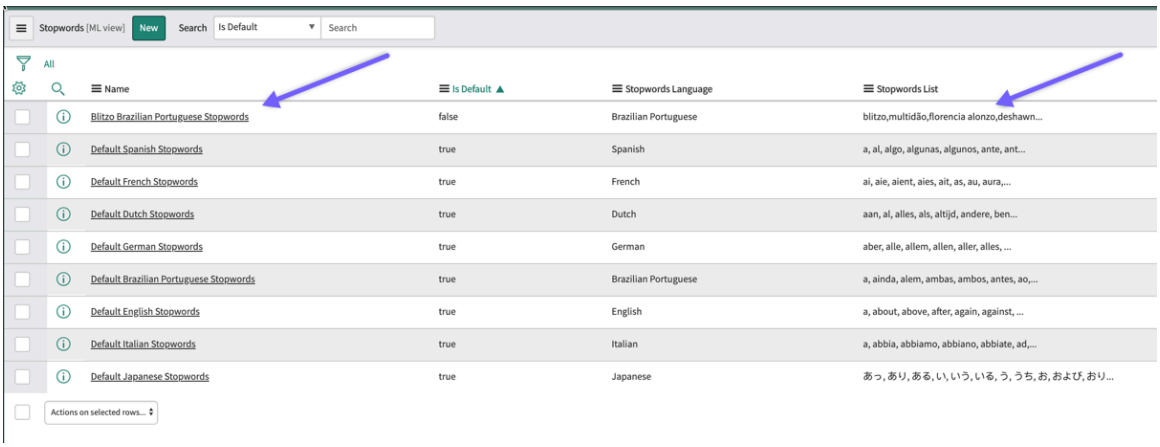
3. In the Stopwords form, configure these fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

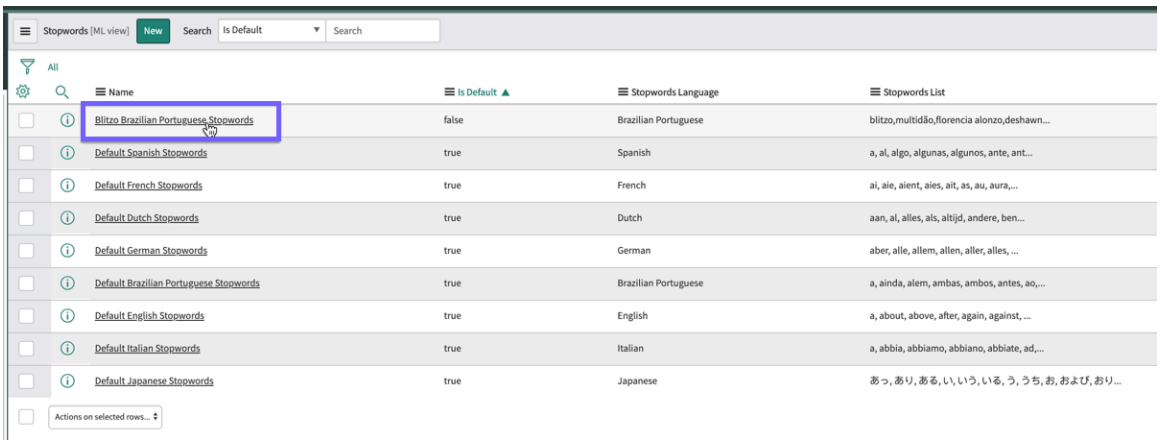
Field	Value
Name	Enter a unique name for the list, such as the name of your company and the processing language. For example, `Blitzo Brazilian Portuguese Stopwords`.
Stopwords Language	Select **Brazilian Portuguese**
Stopwords List	Manually enter the stopwords using a comma-separated format. For more examples of stopwords, see the image in Step 2 of this procedure.



4. Click **Submit**.
   Your custom stopwords list appears in the Stopwords list view.



5. **Optional:** If you need to update your stopwords list, just click its Name, add or remove words from the list, and click **Update**.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next
Assign a custom or default stopwords list to a classification, similarity, clustering, or regression solution definition.

Create a word corpus

Build a collection of words and phrases that functions as the vocabulary the system uses to compare your instance records based on their textual similarity. You can think of the word corpus as a dictionary that you want your machine-learning system to understand.

Before you begin
Role required: admin or ml_admin

Important: In the Australia release, models in the classification, clustering, and similarity frameworks use Workflow solutions. These are pre-trained, so a word corpus isn't needed for your new solutions.

After upgrading, your existing solutions with a word corpus become Workflow solutions the next time they are re-trained. Also the Word Corpus field is removed from the form.

The following information is provided for legacy context.

About this task
The primary purpose for a word corpus is to infer textual data for training your NLU model. If using a word corpus in a solution, you must specify it for training in the solution definition phase of a solution. A trained word corpus can be reused across solutions and capabilities.

You can use a word corpus to help compare similar record text in a table or across multiple tables. A word corpus can also be helpful in other scenarios, such as clustering, where you group similar records together for data analysis, reuse, or review. The items you add to your corpus should be specific to your company and your industry so you can reuse it in other similarity or clustering solutions and apply it to various use cases.

In this example procedure, you're working on incident records and you want to locate relevant knowledge base (KB) articles that could provide resolutions to those incident cases. Your goal here is to create a word corpus that you can apply to a new similarity solution that compares active incidents to published KB articles.

Procedure
1. Navigate to **All** > **Predictive Intelligence** > **Word Corpus**.
2. In the Word Corpus form, click **New**.
3. Configure these fields according to the following guidance.

Field	Description
Name	A unique title that references the contents of your corpus. For example, in this use case you could enter a name such as `Active Incidents and Published KBs`, as the name indicates the tables that your corpus will mine to help create your solution.
Active	Select this check box if you're creating more than one word corpus at a time and you plan to configure their detail components later. Otherwise, leave it blank because you can select it in a later step.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Select **Submit**.
5. In the Word Corpus list view, locate your new word corpus and click its **Name** value to open the record.
6. In the Word Corpus Contents section, Click **New**.
7. In the Word Corpus Content form, configure these fields per the following guidance to define a content component for your word corpus.

Field	Description
Name	Enter a title that references the data you want to add to your corpus, such as `Incidents closed in last 6 months`.
Table	Select the table that contains the data you want to include in your word corpus. For this use case, select **Incident [incident]**. **Note:** The number of records per table for Word Corpus creation used in Similarity and Clustering solutions is limited to 300,000.
Filter	Select the following filter condition values: **[Closed] [is not empty] and [Created in last 6 Months]**.
Field List	For this use case, select **Short description**, **Description**, and **Resolution notes**.
Domain	The system automatically displays the user group for your corpus. For example, in this use case it shows the global user group. You can select other user groups as well.

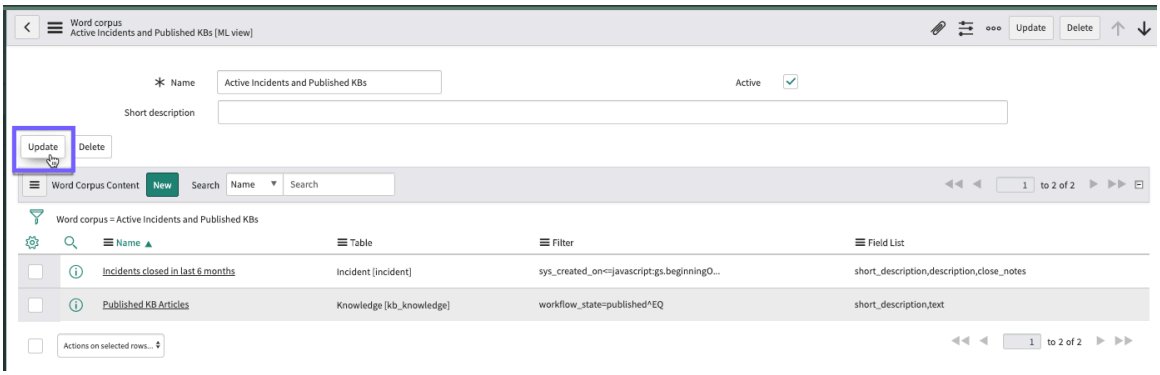
8. Select **Submit**.
9. In the Word Corpus Details section, select **New**.
10. Configure these fields according to the following guidance to define a second content component for your word corpus.

Field	Description
Name	Enter a title that references the data you want to compare to your first content component, such as `Published KB Articles`.
Table	Select the table that contains the data you want to compare to your first content component. For this use case, select **Knowledge [kb_knowledge]**. **Note:** The number of records per table for word corpus creation used in Similarity and Clustering solutions is limited to 300,000 records per table.
Filter	Select the following Filter Condition values: **[Workflow] [is] [Published]**.
Field List	Select **Short description** and **Article body**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

11. Select **Submit**.

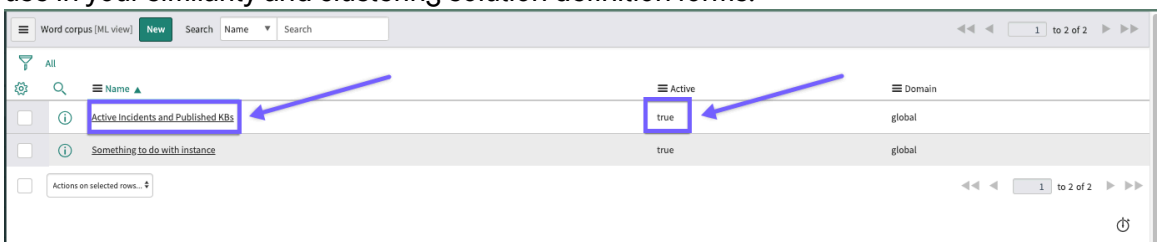
Your two word corpus content components appear on the word corpus form.



12. Select **Update**.

Result

The completed word corpus you created appears on the word corpus form and is available for use in your similarity and clustering solution definition forms.



What to do next

Create a solution in the appropriate framework. For more information, see the links in the Related Content panel on this page.

Related topics

- Create and train a classification solution
- Create and train a similarity solution
- Create and train a clustering solution

Quick start tests for Predictive Intelligence

Validate that Predictive Intelligence still works after you make any configuration change such as applying an upgrade or developing an application. Copy and customize these quick start tests to pass when using your instance-specific data.

Predictive Intelligence quick start tests require activating the Predictive Intelligence [com.glide.platform_ml] plugin. In order to execute critical upgrade tests on existing machine learning solutions, you need to create a basic authorization profile named ml_atf in the Basic Auth Configurations table (sys_auth_profile_basic.list). To run the tests successfully, the user attached to the ml_atf authorization profile must have the ml_admin role.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Predictive Intelligence: ATF Test Suite

Test	Description
PI: Presence of ML model artifacts persisted in glide	Verify all the trained ML model artifacts are persisted in glide (sys_attachments table) after model training/instance cloning so that ML prediction calls are successful.
PI: Valid setup of ML user (sharedservice.worker) in glide	Validate if the ML user in glide (sharedservice.worker) is active and not logged out so that model training is successful.
PI: Glide upgrade test for Classification solution	Validate that the classification model prediction on existing active models is producing the same class membership and confidence value results after a glide upgrade.
PI: Glide upgrade test for Similarity solution	Validate that the similarity model prediction API calls on active models are successful after a glide upgrade.

Related topics

Quick start tests

Activate a solution version

Predictive Intelligence activates the most recent version of the solution when it completes training a solution. However, you can activate any previously trained solution version. Only one solution version can be active at a time, and only the active version is used when making predictions.

Before you begin

- Manually train a solution multiple times or set a training schedule.
- Role required: admin or ml_admin

About this task

The system creates a solution version each time you train a solution definition. Typically, you manually create a solution version only when you change the solution definition filter and want to test it. Otherwise, most solution versions are created during scheduled solution training.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Classification** > **Solutions** or **Predictive Intelligence** > **Similarity** > **Solutions** or **Predictive Intelligence** > **Clustering** > **Solutions** or **Predictive Intelligence** > **Regression** > **Solutions** .
2. In the list view of the ML Solutions table, find the row that you want to activate, then select the value in the first column to open the individual record.

   You can also open the individual record by selecting the row's reference lookup icon  , then selecting **Open Record**.
3. In the solution record, select **Activate**.
   The system activates this solution version and deactivates any other solution version.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next
For classification solutions, review the trained solution precision and coverage statistics. For similarity solutions, review the similarity examples.

Export trained solutions to production

Refine and test your ML solutions iteratively on a non-production instance, and then use update sets to export the changes to your production instance. This practice mitigates the risk of retraining solutions on your live production instance.

Before you begin
Prior to testing on a test instance, ensure that the instance hosts recently-cloned data from your production instance so that the solutions you train on the test instance remain valid when you export them to production.

Role required: admin or ml_admin

About this task
Plan your changes carefully, and remember that update sets match records based on the system ID [sys-id] and not the version number. Note that moving solutions to instances can sometimes deliver unpredictable results. If you run into such an issue, retrain the solution again, which takes 5 minutes per solution. For more information on update sets, see System update sets .

Procedure
1. Navigate to **All** > **Predictive Intelligence** > **Classification** > **Solution Definitions** or **Predictive Intelligence** > **Similarity** > **Solution Definitions**.
2. Click the name of your trained ML Solution Definition record to open it.
3. In the Related Links section, click **Add solutions to the current update set**.
4. Click **Update**.

Result
Your trained ML solution artifacts, such as solution definitions, template records, and predictive model statistics, are added to the current update set.

What to do next
Depending on where you are in your solution testing, schedule your update set for export to another non-production instance for further testing, or on to production.

Note: After you export a similarity solution, click **Refresh similarity window (Required after Solution Import)** in the Related Links section of the corresponding ML Solution Definition form.

Configuring advanced settings for your ML solutions

Learn about advanced settings for your Predictive Intelligence machine learning (ML) solutions. Apply optional technology and algorithms for classification, clustering, similarity, and regression capabilities.

Using advanced settings in your ML solutions is optional. If you choose to implement any of these settings, make sure you're well informed regarding the technology you're enabling in the solution, and that your use case benefits from what the technology offers. For more information about several of these parameters, see the Dive deeper with Clustering Advanced Parameters article on ServiceNow Community.

Most of the advanced settings are limited to only one capability because their functionality is targeted to specific use cases.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Classification, clustering, similarity, and regression advanced settings

Using the admin or ml_admin role, you apply these technologies by configuring a parameter in the **Advanced Solution Settings** tab on your ML solution definition form.

To update or remove an existing setting: on the solution definition form, select the **Advanced Solution Settings** tab, then select the name of the solution parameter. For details, see the following procedures.

- Configure class recall for a classification solution
- Configure TF-IDF for solutions
- Configure DBSCAN for a clustering solution
- Configure HDBSCAN for a clustering solution
- Configure XGBoost for classification or regression solutions
- Configure Connect Component algorithm and Levenshtein Distance method for a clustering solution
- Apply purity on a clustering solution
- Analyze a cluster with Cluster Insight

Advanced Solution Settings for Classification

From the Zurich release, classification solutions offer the following advanced solution settings.

- Config parameters for model config in classification
- Configure include only top N labels
- Minimum records needed for label to include it
- Remove others label
- Use LightGBM algo for classification model training

Configure class recall for a classification solution

Create and apply a class recall parameter to an ML solution prior to training its data. For example, you set and apply this solution parameter to 90% recall for all records in the Email class.

Before you begin

Note: Using advanced settings in your ML solutions is optional. If you choose to implement any of these settings, make sure you're well informed regarding the technology you're enabling in the solution, and that you have a use case that benefits from what the technology offers.

- Create and save a classification solution definition or use an existing one.
- Role required: admin or ml_admin

About this task

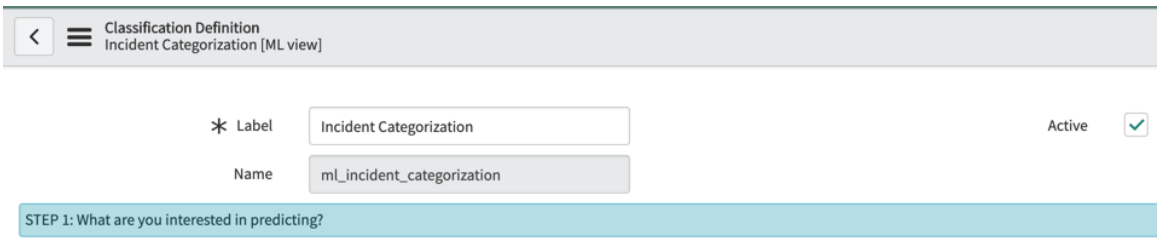
The class recall solution parameter enables you to steer a solution's training to bias a specific class. For example, classifying an incoming email as a Phish or not can be an important use case in a security-related machine learning solution. In this situation, it's very important to identify every Phish, and it may be acceptable to report a non-Phish as a Phish occasionally. However, no actual Phish should be classified as a non-Phish. In such situations, the recall metric must have a high value to minimize false negatives. High recall values might lead to lower percentages for precision and coverage.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

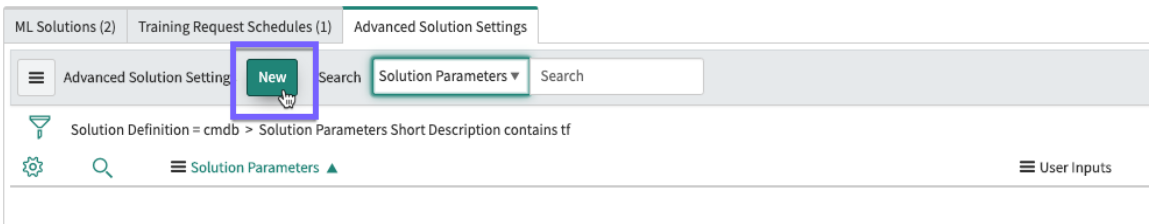
You can apply a recall value to one class per solution definition.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Classification** > **Solution Definitions**.
2. Open a saved classification solution definition form.
   In this example scenario, you use an Incident Categorization solution definition form that you haven't trained yet.



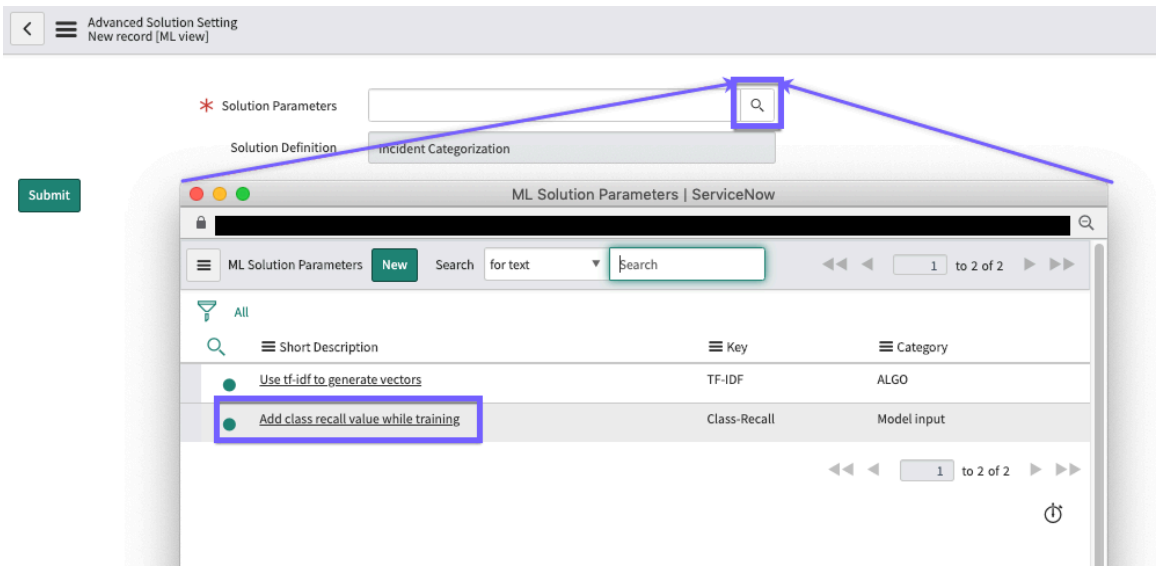
3. On the Advocated Solution Settings tab in the Related Links section of the form, select **New**.



4. Create a parameter record.

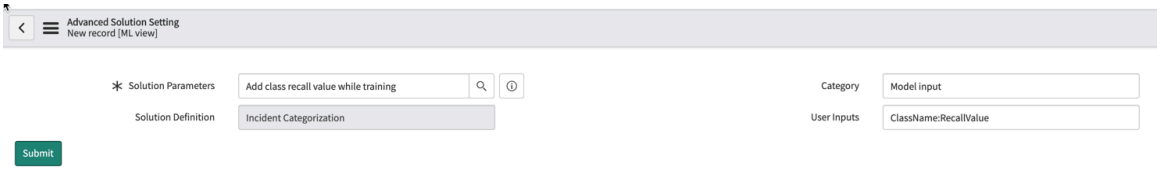
   a. In the **Solution Parameters** field, select the search icon.

   b. In the ML Solution Parameters screen, select **Add class recall value while training**.



The Advanced Solution Setting [ml_advanced_solution_settings] record appears.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

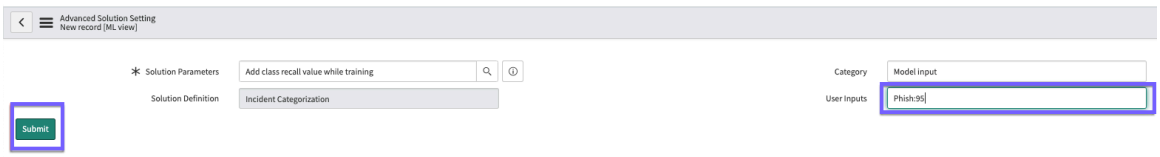


5. Add your values to the **User inputs** field, in the format `ClassName:RecallValue`.

   a. Enter the name of the class for which you want to specify a recall value.
   In this scenario, you enter `Phish` for the **ClassName**.

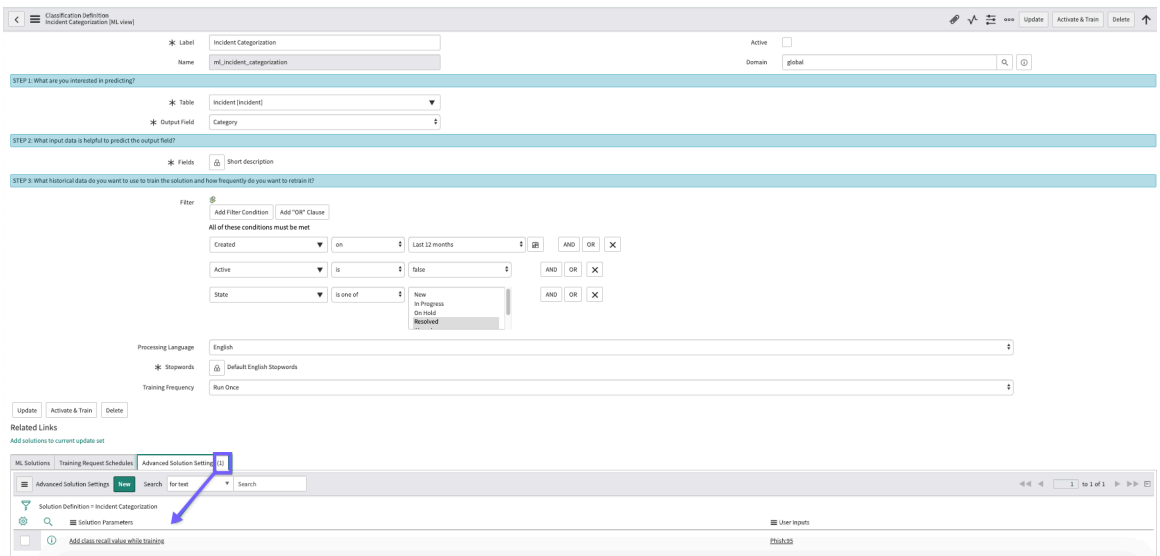
   b. Enter a colon character (:), then the Recall value.
   Enter `95` for the **RecallValue** in the example scenario.

   In other words you're specifying Phish as the target class, and 95 is the Recall percentage you're requesting the system to deliver during solution training (`Phish:95`).



6. Select **Submit**.

   Result: Class recall is configured for your classification solution. Its solution parameter appears on the Advanced Solution Settings tab of your classification solution definition form.



What to do next

Train your saved classification solution so that your updates are incorporated.

Related topics

- Create and train a classification solution
- Configuring target metrics for a trained classification solution
- Configure TF-IDF for solutions

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure TF-IDF for solutions

Apply Term Frequency–Inverse Document Frequency (TF-IDF) encoding to classification, clustering, or similarity solutions for Predictive Intelligence.

Before you begin

Note: Configuring advanced settings on your ML solutions is optional. If you choose to configure any of these settings, make sure you're well informed regarding the technology you're enabling in the solution, and that you have a use case that benefits from what the technology offers. For more information, see https://www.servicenow.com/community/intelligence-ml-articles/dive-deeper-with-clustering-advanced-parameters/ta-p/2695847.

- Create a classification, clustering, or similarity solution definition or use an existing one.
- Role required: admin or ml_admin

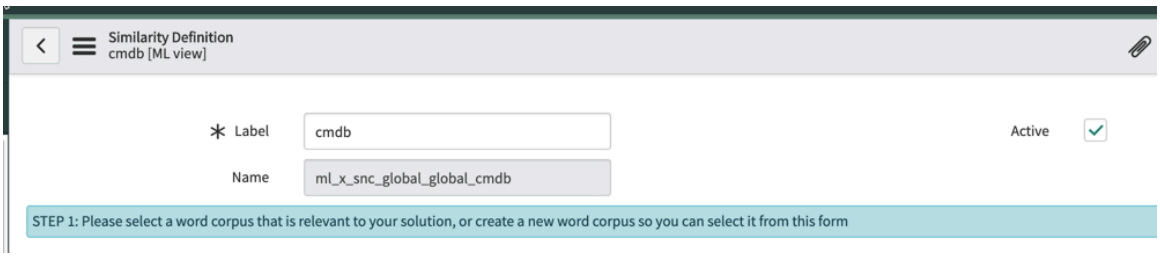
About this task

Predictive Intelligence uses paragraph vector word embedding by default in its classification and similarity frameworks, which is highly effective for processing data comprised of primarily human-readable content. However, TF-IDF might return better prediction results for records that have machine-generated content, such as alerts and error messages for log files. Choose advanced settings that are appropriate for the kind of data your solution is processing.

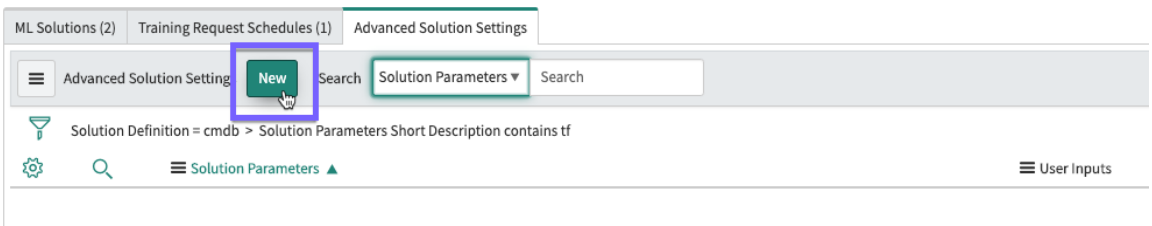
Note: The steps for configuring TF-IDF are the same for all model frameworks, but TF-IDF support for clustering solution definitions is applicable only if you have a Professional subscription.

Procedure

1. Navigate to a **Solution Definition**, such as **All** > **Predictive Intelligence** > **Similarity** > **Solution Definitions**.
2. Open a solution definition form.
   In this example scenario, you use a CMDB similarity definition form.

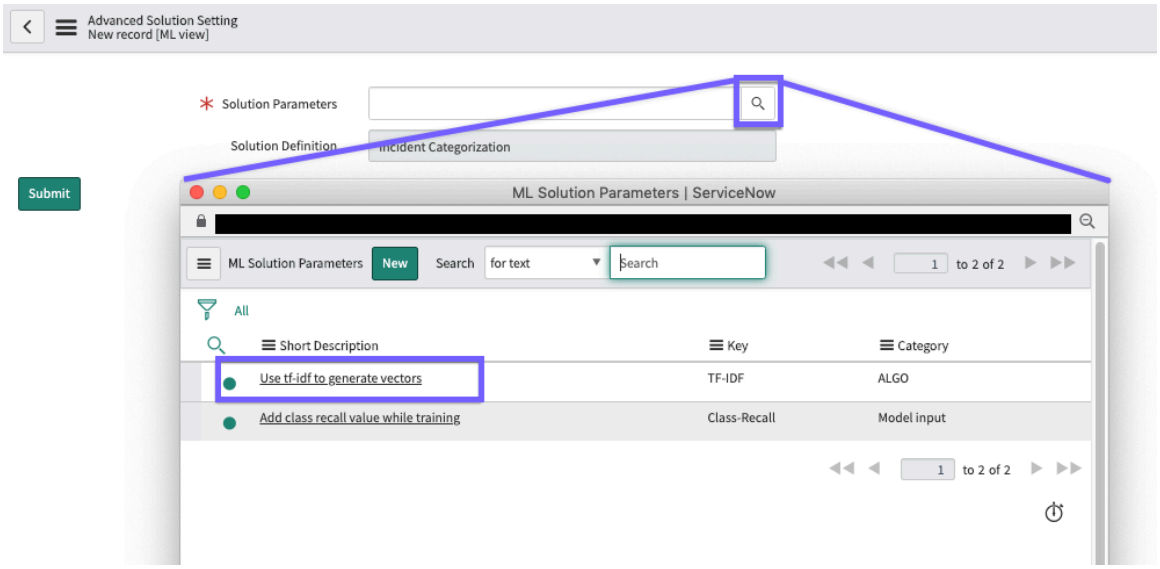


3. On the Advanced Solution Settings tab in the Related Links section of the form, click **New**.



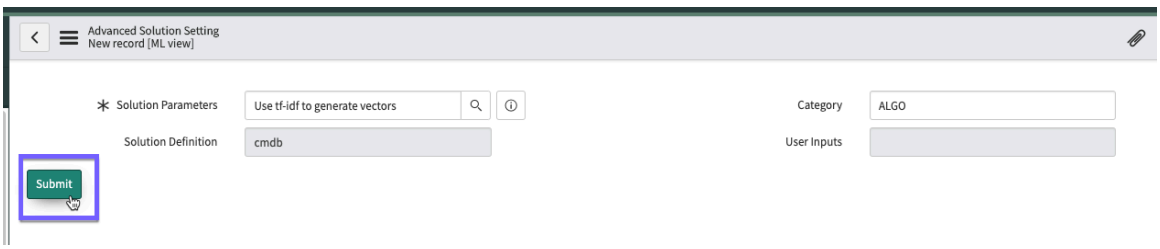
4. Create a parameter record.
   a. In the **Solution Parameters** field, click the search icon.
   b. In the ML Solution Parameters screen, select **Use tf-idf to generate vectors**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



5. Click **Submit**.

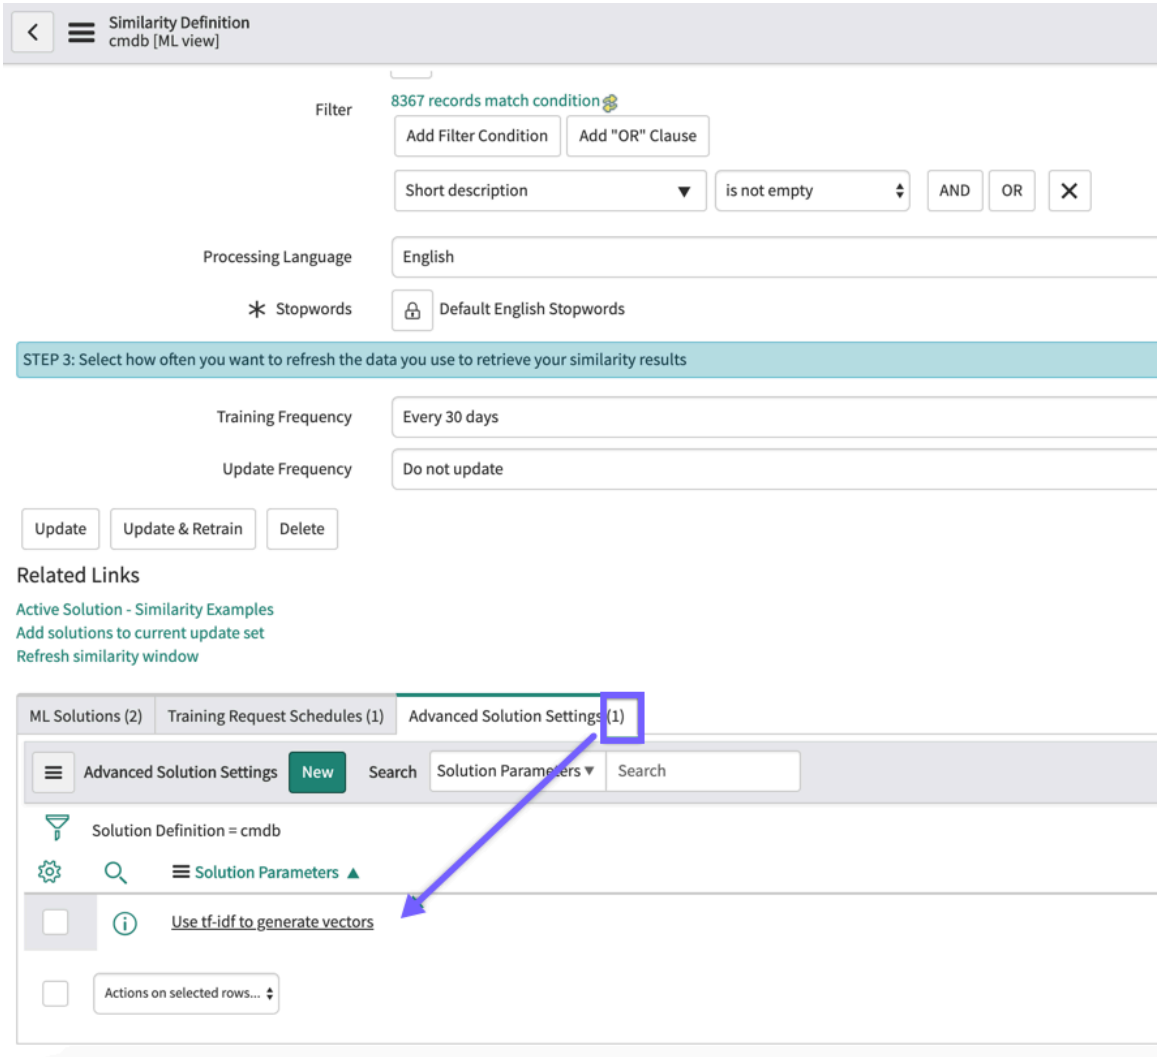
   The Advanced Solution Setting record screen refreshes.



6. Click **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result: TF-IDF is configured for your similarity solution. Its solution parameter appears on the Advanced Solution Settings tab of your similarity definition form.



Related topics

- Create and train a classification solution
- Create and train a similarity solution
- Create and train a regression solution
- Create and train a clustering solution

Configure XGBoost for classification or regression solutions

Apply XGBoost encoding to optimize the training for your classification or regression solutions.

Before you begin

- Create a classification solution definition or use an existing one.
- Create a regression solution definition or use an existing one.
- Role required: admin or ml_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

XGBoost is an optional gradient boosting framework that uses multiple decision trees and supports both Paragraph Vector-based text and TF-IDF distance-based text. LogR is the default distance-based model algorithm.

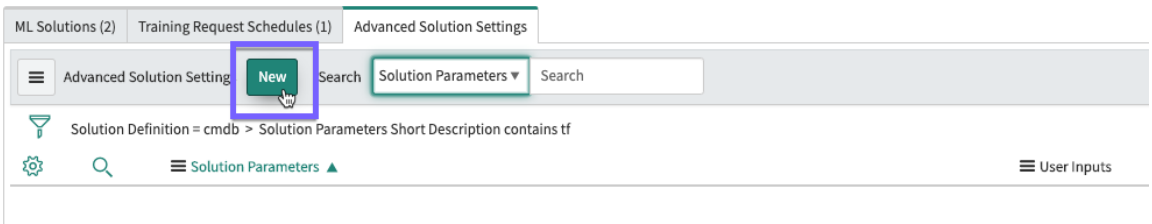
Configuring advanced settings on your ML solutions is optional. If you choose to configure any of these settings, make sure you're well informed regarding the technology you're enabling in the solution, and that you have a use case that benefits from what the technology offers.

In this example scenario, you apply XGBoost to both a classification solution and a regression solution.

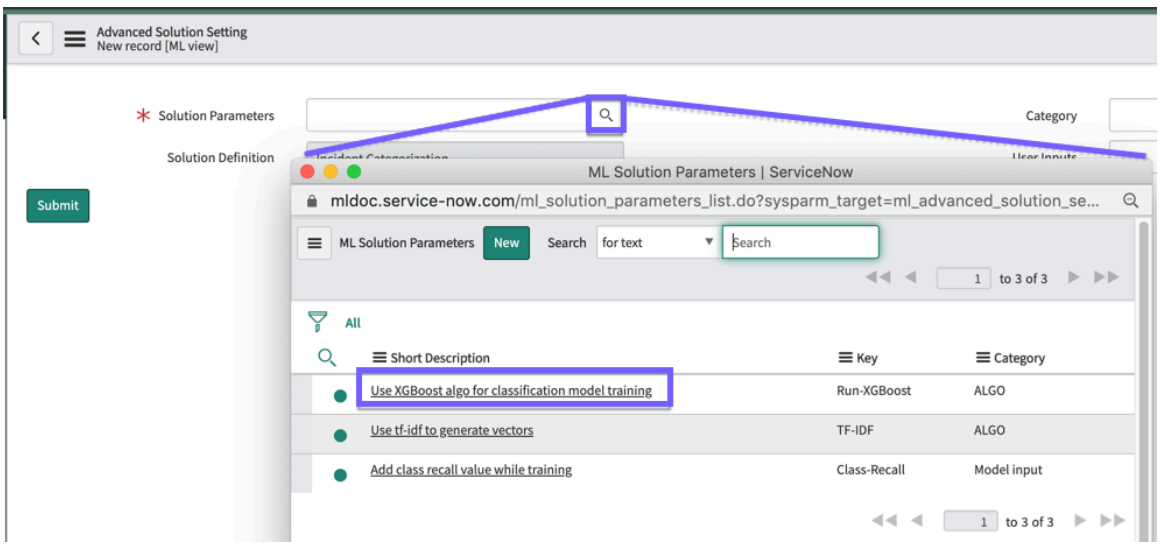
Note: The regression framework is deprecated in the Australia release. You can continue to use existing regression solutions but you can't create new ones.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Classification** > **Solution Definitions**.
2. Open a classification solution definition form.
3. On the Advanced Solution Settings tab in the Related Links section of the form, select **New**.

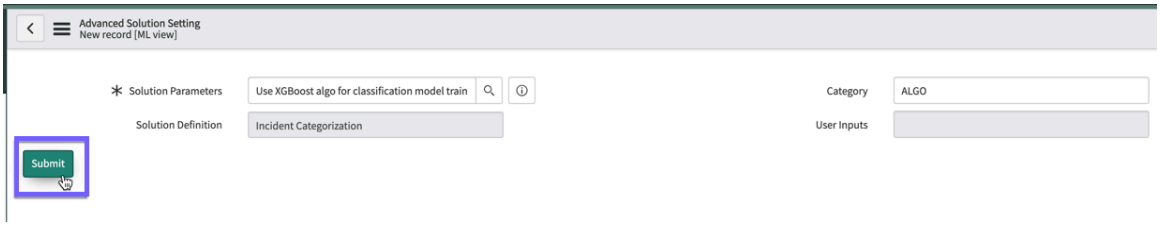


4. Create a parameter record.
   a. In the **Solution Parameters** field, select the search icon.
   b. In the ML Solution Parameters screen, select **Use XGBoost algo for classification model training**.



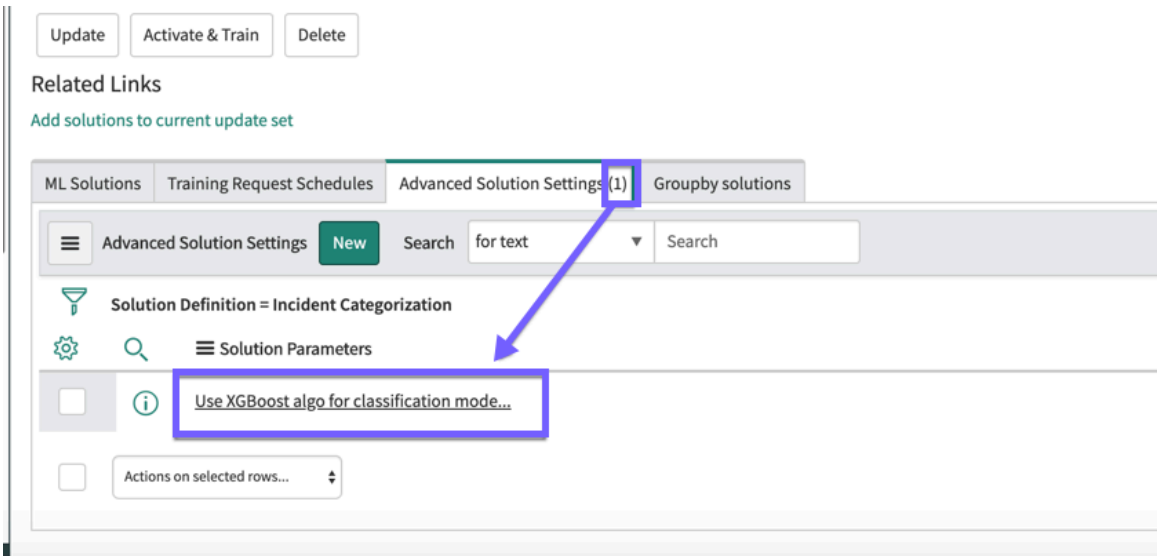
5. Select **Submit**.
   The Advanced Solution Setting record screen refreshes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



6. Select **Submit**.

   Result: XGBoost is configured for your classification solution. Its solution parameter appears on the Advanced Solution Settings tab of your classification definition form.



Note: Follow the steps below if you want to configure XGBoost on a regression solution.

7. Navigate to **Predictive Intelligence** > **Regression** > **Solution Definitions** (available only if you had created a regression solution before the regression framework was deprecated).
8. In this second scenario, you open a regression solution definition form.
9. Repeat the steps from the previous classification solution example, except this time you're using a regression solution.
10. Select **Submit**.

    Result:

    XGBoost is configured for your regression solution. Its solution parameter appears on the Advanced Solution Settings tab of your regression solution definition form.

Related topics

Create and train a classification solution

Create and train a regression solution

Configure DBSCAN for a clustering solution

Consider applying the Density Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm to your clustering solution. DBSCAN is available as an alternative to the default clustering algorithm, k-means.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Note: Configuring advanced settings on your ML solutions is optional. If you choose to configure any of these settings, make sure you're well informed regarding the technology you're enabling in the solution, and that your use case benefits from what the technology offers. For more information, see the Dive deeper with Clustering Advanced Parameters article on ServiceNow Community.

- Create a clustering solution definition or use an existing one.
- Role required: admin or ml_admin

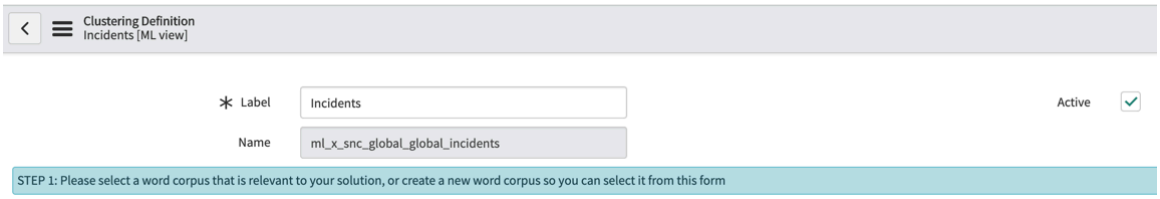
About this task

Predictive Intelligence uses the k-means algorithm by default in its clustering framework. DBSCAN is another clustering algorithm that's also used in data mining and machine learning. Some users prefer DBSCAN as it doesn't require you to specify the number of clusters in the data before clustering. For a summary of the pros and cons for each algorithm, see this conversation and this article.

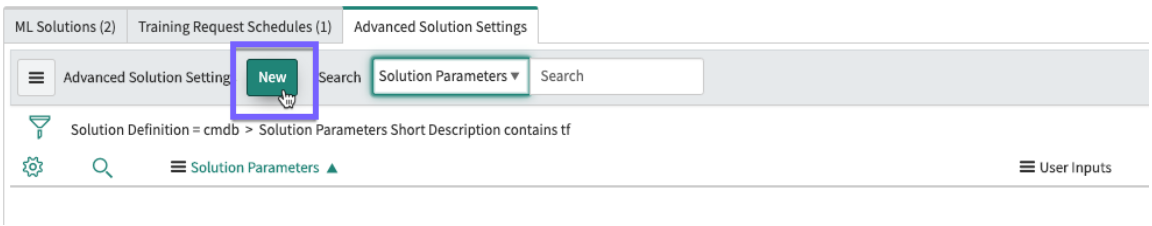
In this example scenario, you apply DBSCAN to a clustering solution.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solution Definitions**.
2. Open a clustering solution definition form.
   In this example scenario, you use an Incidents clustering solution definition form.



3. On the Advanced Solution Settings tab in the Related Links section of the form, select **Solution Parameters** from the picker, then click **New**.

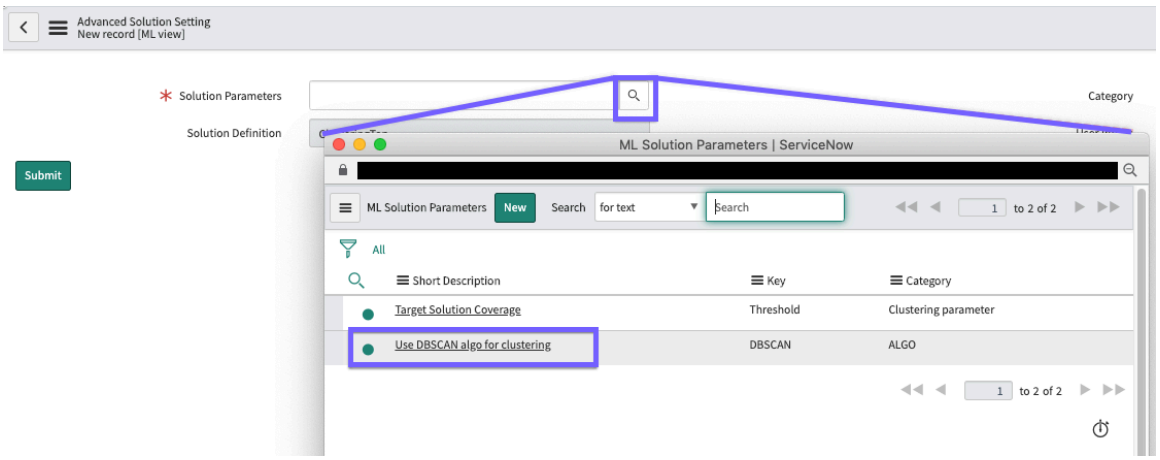


4. Create a parameter record.

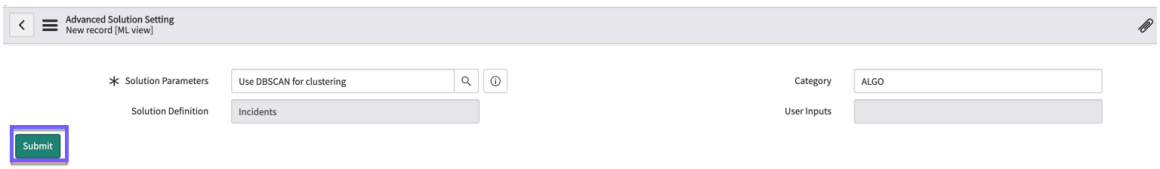
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the **Solution Parameters** field, select the search icon.

b. In the ML Solution Parameters screen, select **Use DBSCAN algo for clustering**.



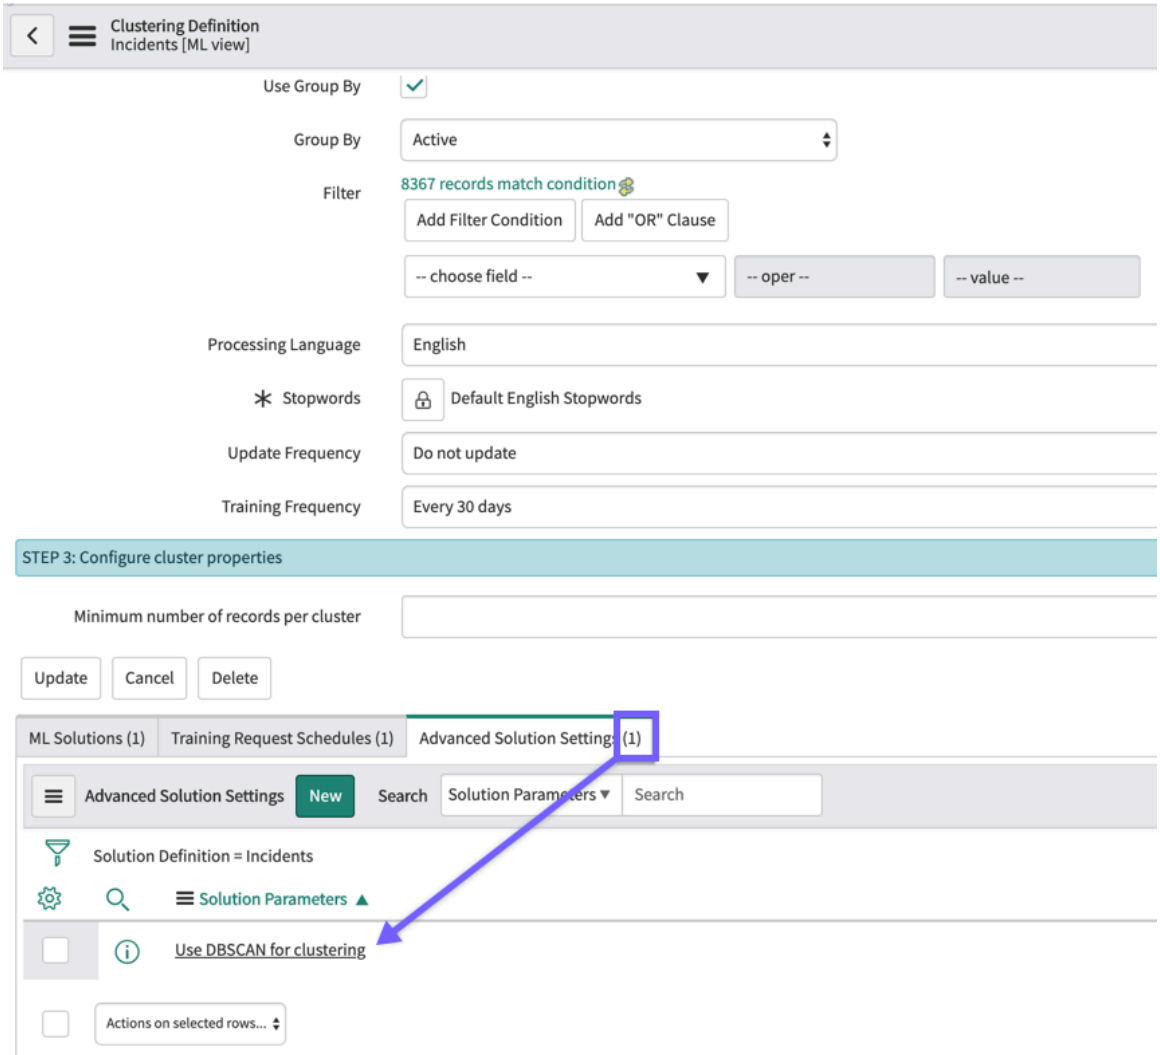
5. Select **Submit**.
The Advanced Solution Setting record appears. The field **User Inputs** is grayed out because it does not apply to this algorithm.



6. Select **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result: DBSCAN is configured for your clustering solution. Its solution parameter appears on the Advanced Solution Settings tab of your clustering solution definition form.



Related topics

Create and train a clustering solution

Configure HDBSCAN for a clustering solution

Consider applying the Hierarchical Density Based Spatial Clustering of Applications with Noise (HDBSCAN) algorithm to your clustering solution. HDBSCAN is available as an alternative to the default clustering algorithm, k-means.

Before you begin

Note: Configuring advanced settings on your ML solutions is optional. If you choose to configure any of these settings, make sure you're well informed regarding the technology you're enabling in the solution, and that your use case benefits from what the technology offers. For more information, see the Dive deeper with Clustering Advanced Parameters article on ServiceNow Community.

- Create a clustering solution definition or use an existing one.
- Role required: admin or ml_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

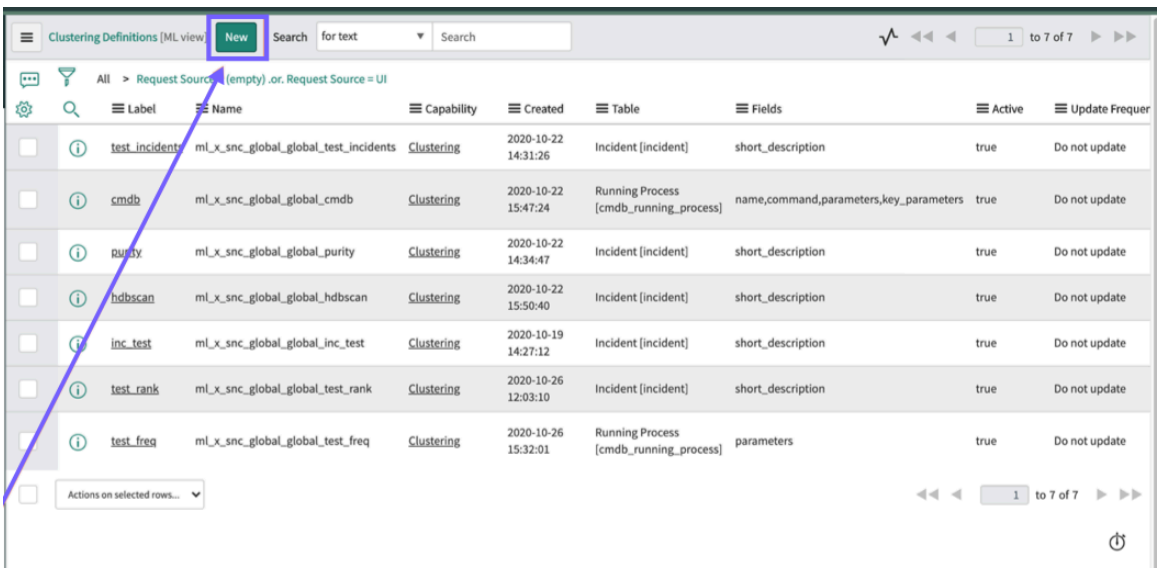
You can apply the HDBSCAN algorithm to help the system identify data samples that aren't assigned to any cluster. For example, you can apply HDBSCAN to support Topic Discovery.

Predictive Intelligence implements the k-means algorithm by default in its clustering framework. HDBSCAN is similar to the DBSCAN clustering algorithm except that it works with minimum-sized clusters and can help deliver more stable and persistent clusters. For a summary of how HDBSCAN works, see this article. For a comparison between DBSCAN and HDBSCAN, see this article and this article.

Note: Clustering solutions trained with HDBSCAN do not support cluster updates. Updates on these solutions fail and the solutions are not logged in the ml_cluster_detail_table. Use DBSCAN or k-means training methods if you want to enable cluster updates.

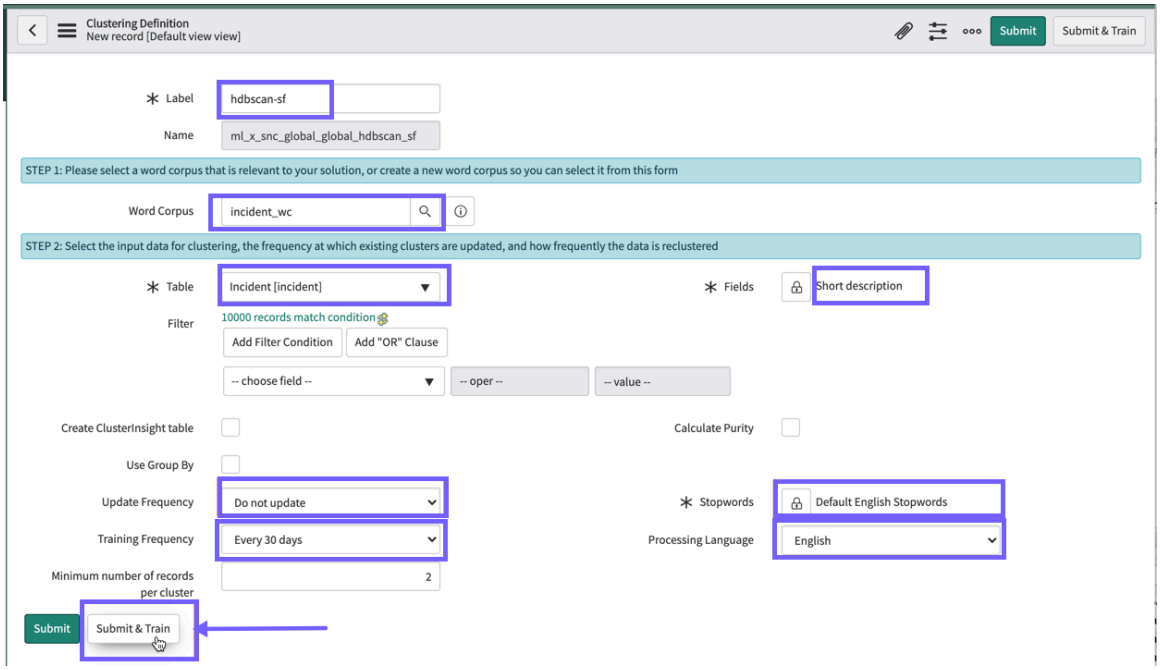
Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solution Definitions**.

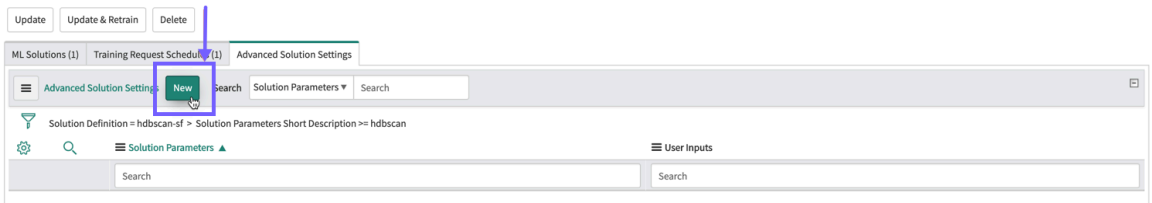


2. Select **New**.
3. Create a new clustering solution definition form or use an existing one.
   In this example scenario, you create the hdbscan-sf clustering definition form as in the image below. Configure the fields as follows:
   - **Label**: `hdbscan-sf`
   - **Word Corpus**: `incident_wc`, or any other word corpus that has incident record data (from the Washington DC release, a word corpus is not needed, so this field does not appear).
   - **Table**: Incident [incident]
   - **Fields**: Short description
   - **Update Frequency**: Do not update
   - **Stopwords**: Default English Stopwords
   - **Training Frequency**: Every 30 days
   - **Processing Language**: English

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. Select **Submit & Train**.
5. On the Advanced Solution Settings tab in the Related Links section of the trained form, select **Solution Parameters** from the picker, then select **New**.

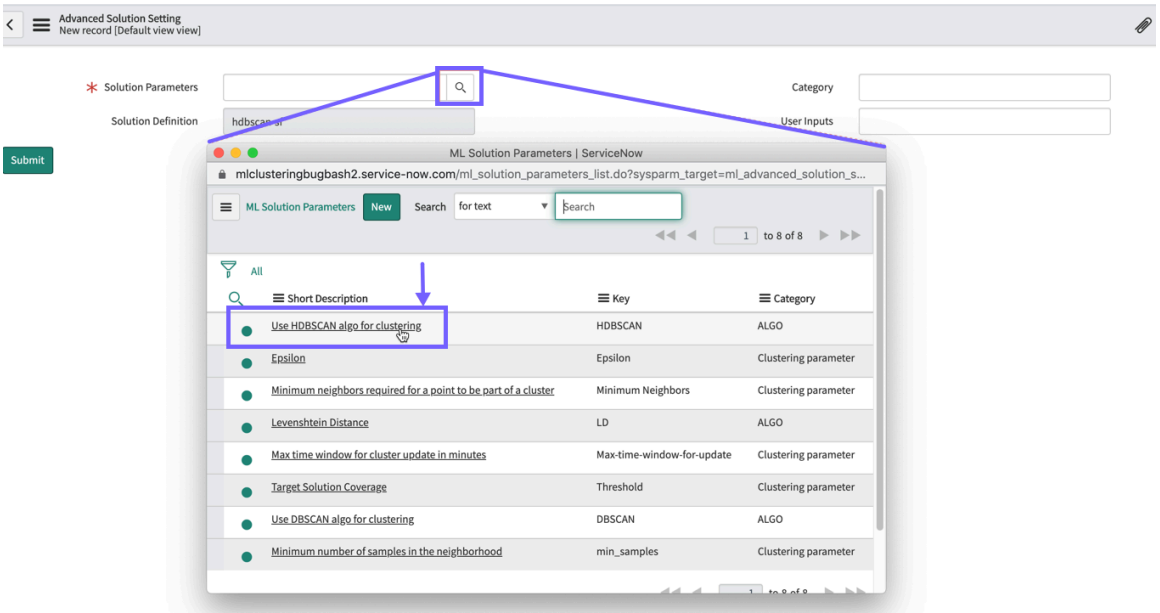


6. Create a parameter record.

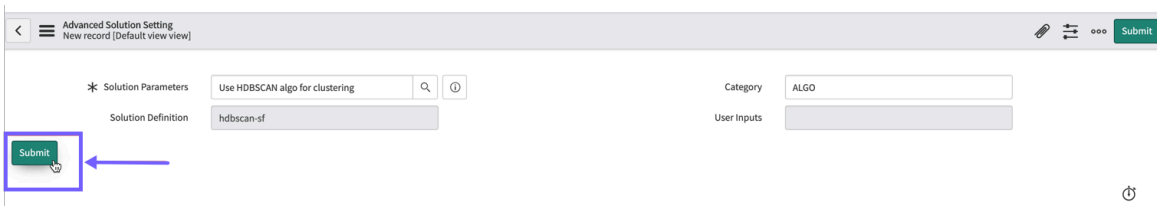
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the **Solution Parameters** field, click the search icon.

b. In the ML Solution Parameters screen, select **Use HDBSCAN algo for clustering**.



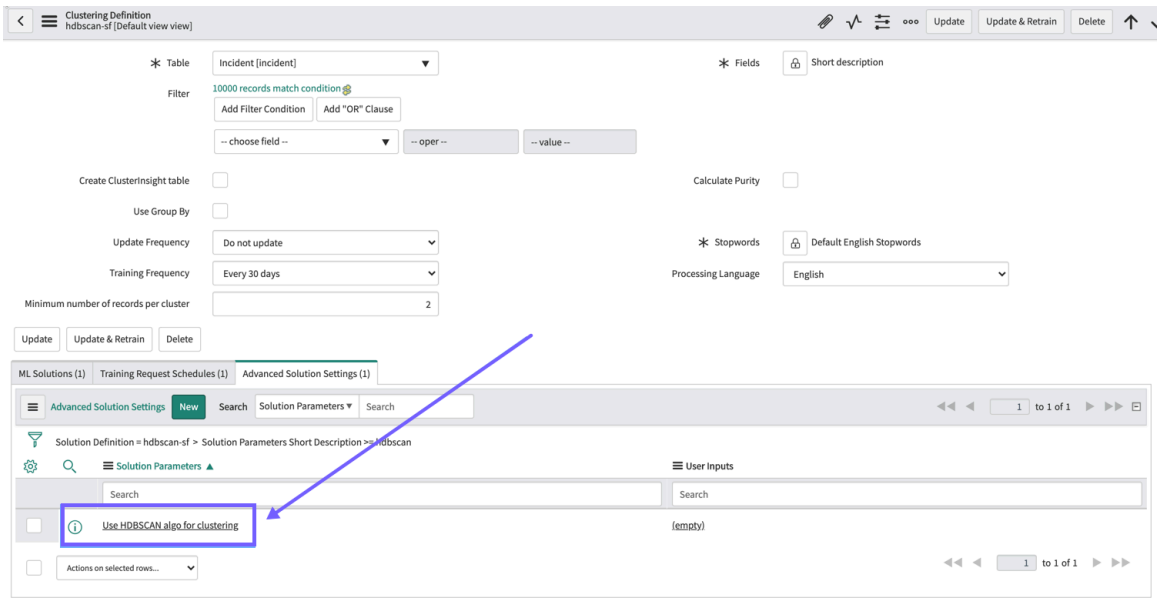
7. Select **Submit**.
The Advanced Solution Setting record appears with the HDBSCAN algorithm applied to the record. The field **User Inputs** is grayed out because it does not apply to this algorithm.



8. Select **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result: HDBSCAN is configured for your clustering solution. Its solution parameter appears on the Advanced Solution Settings tab of your clustering solution definition form.



Related topics

Create and train a clustering solution

Configure Connect Component algorithm and Levenshtein Distance method for a clustering solution

Apply Configure Connect Component and Levenshtein Distance method encoding to optimize the training for your clustering solutions.

Before you begin

Role required: admin or ml_admin

Note: Configuring advanced settings on your ML solutions is optional. If you choose to configure any of these settings, make sure you're well informed regarding the technology you're enabling in the solution, and that your use case benefits from what the technology offers. For more information, see the Dive deeper with Clustering Advanced Parameters article on ServiceNow Community.

- Create and train a clustering solution definition or use an existing one.
- Role required: admin or ml_admin

About this task

When training clustering solutions, you have the following three options.

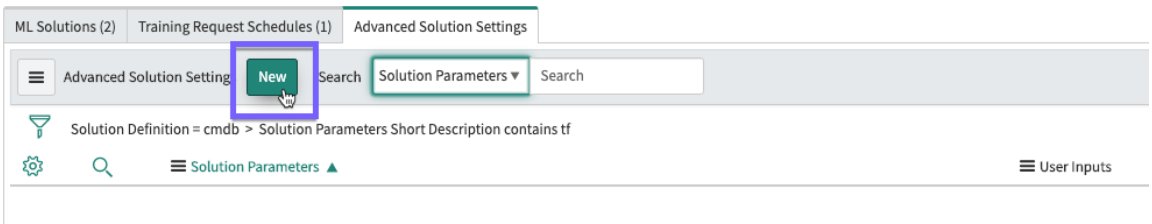
- Use the default k-means algorithm.
- Use the optional DBSCAN solution parameter with the Euclidean distance method as a metric.
- Use the optional DBSCAN, Minimum Neighbors, and Levenshtein Distance solution parameters. Connect Component is enabled by DBSCAN and Minimum Neighbors, and supports both Paragraph Vector-based text and Levenshtein Distance-based text. If you train your solution using the Levenshtein Distance method, you don't need to use a word corpus in your clustering solution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

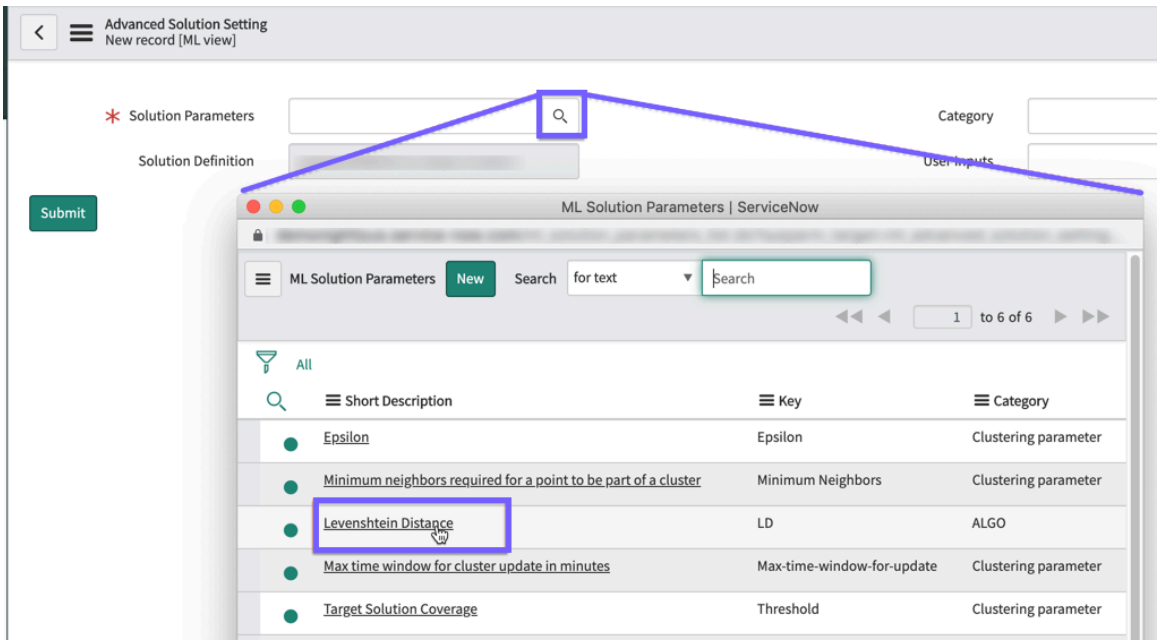
In this example scenario, you train your solution definition by using the third option referenced above.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solution Definitions**.
2. Open a trained clustering solution definition form.
3. On the Advanced Solution Settings tab in the Related Links section of the form, select **New**.

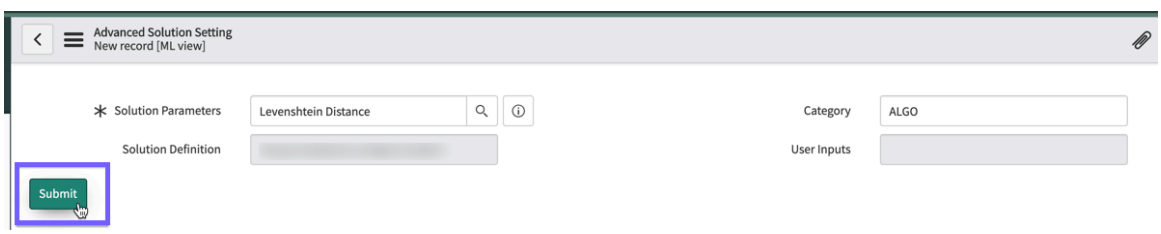


4. Create a parameter record.
   a. In the **Solution Parameters** field, select the search icon.
   b. In the ML Solution Parameters screen, select **Levenshtein Distance**.



5. Select **Submit**.

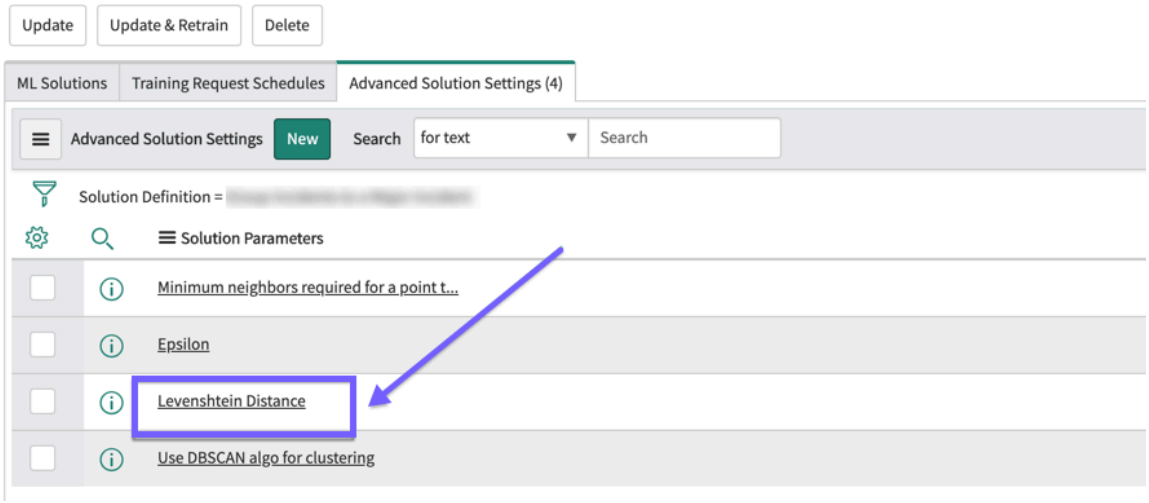
   The Advanced Solution Setting record screen refreshes.



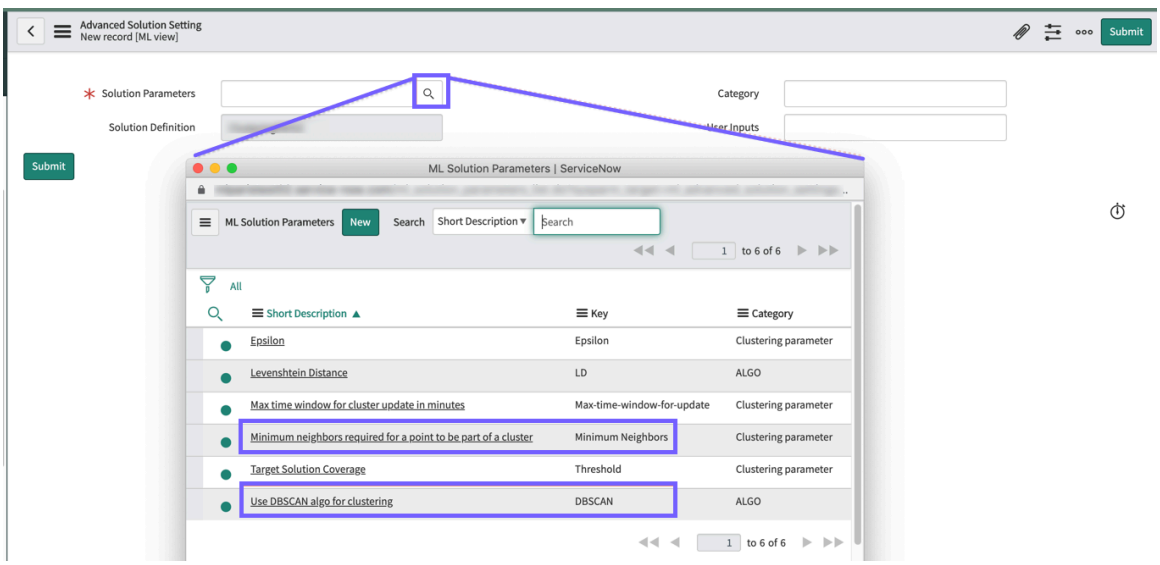
6. Select **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

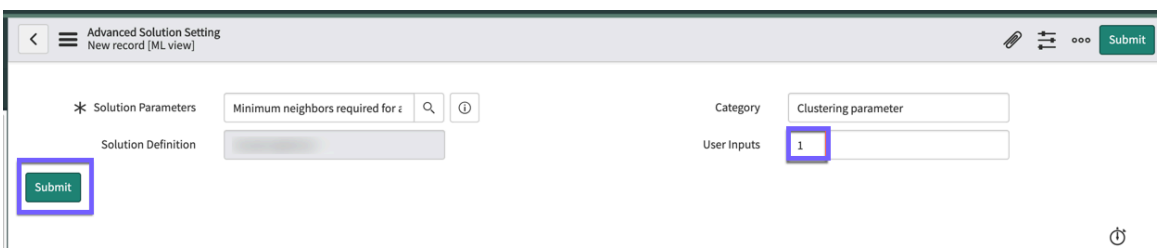
Result: Levenshtein Distance is configured for your clustering solution. Its solution parameter appears on the Advanced Solution Settings tab of your clustering definition form.



7. Repeat steps 1-6 from the previous Levenshtein Distance example, except this time you're creating the **Minimum Neighbors** and **DBSCAN** solution parameters, which together enable the Connect Component feature.



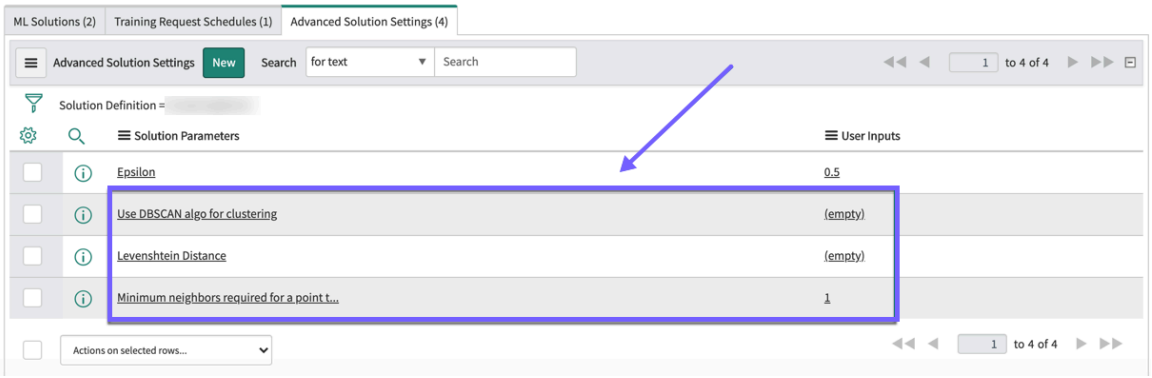
When you select, configure, and submit the **Minimum Neighbors** solution parameter, be sure to set the **User Inputs** field with a value of 1. Only some parameters have a User Inputs field.



Result:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connect Component is configured for your clustering solution. Its two solution parameters appear on the Advanced Solution Settings tab of your clustering definition form, alongside the Levenshtein Distance parameter you configured in steps 1-6 of this procedure.



Related topics

Create and train a clustering solution

Config parameters for model config in classification

Customize the training behavior of your classification solution by passing in a dictionary of parameters in JSON format.

Before you begin

Role required: ml_admin or admin

About this task

You can fine-tune model performance by passing in a dictionary of parameters to adjust training behavior.

Understand the machine learning hyperparameters relevant to your model objective. This option is advanced and should be modified with care.

You can add these parameters in the solution definition form by using the following procedure.

Alternatively, you can use a script to add these parameters. For the format to use when scripting, see the **What to do next** section of this page.

Procedure

1. Navigate to **All** > *Predictive Intelligence* > **Classification** > **Solution Definitions**.
2. Open an existing classification solution or create a new one.
3. Select the **Advanced Solution Settings** tab in the **Related Links** section of the form.
4. Select **New** to open the **Advanced Solution Setting** (ml_advanced_solution_settings) form.
5. In the **Solution Parameters** field, search for `Config parameters for model config in classification`.
6. In the **User Inputs** field, enter the following key-value pairs, after editing the values per your requirements.

> **Note:** When pasting the parameters into the solution definition form, don't include JSON's enclosing label, braces, or parentheses. Enter the key-value pairs in the format shown here, after editing the values according to your requirements.

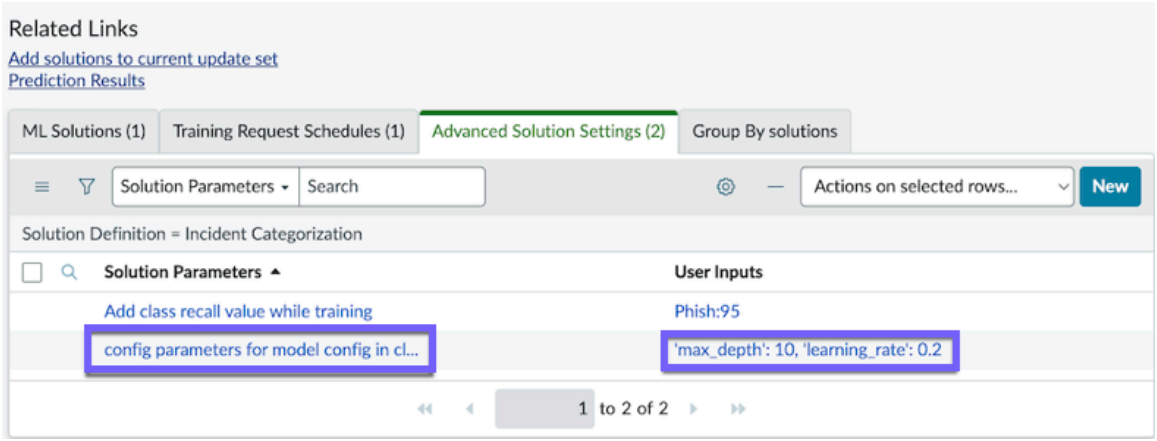
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
'max_depth': 10,'learning_rate': 0.1,'objective':
 'multi:softprob','num_class': 11,'random_state':
 10,'eval_metric': 'mlogloss','booster':
 'gbtree','n_estimators': 150,'verbosity':
 1,'use_label_encoder': false,'tree_method':
 'hist','num_parallel_tree': 1,'gamma': 0.25,'subsample':
 0.8,'max_delta_step': 5,'reg_alpha': 0,'reg_lambda':
 1,'colsample_bytree': 0.8,'colsample_bylevel':
 1,'colsample_bynode': 1,'min_child_weight': 1.25,'n_jobs':
 11,'validate_parameters': true
```

7. Select **Submit** to update the solution definition.

Result

The setting appears as a row on the **Advanced Solution Setting** tab on your solution's form.



What to do next

Retrain your solution.

You can also pass in this dictionary with a script. Be sure to use the enclosing label, braces, and parentheses when scripting.

Script example:

```
config.setModelParams({
                    'max_depth': 10,
                    'learning_rate': 0.1,
                    'objective': 'multi:softprob',
                    'num_class': 11,
                    'random_state': 10,
                    'eval_metric': 'mlogloss',
                    'booster': 'gbtree',
                    'n_estimators': 150,
                    'verbosity': 1,
                    'use_label_encoder': false,
                    'tree_method': 'hist',
                    'num_parallel_tree': 1,
                    'gamma': 0.25,
                    'subsample': 0.8,
                    'max_delta_step': 5,
                    'reg_alpha': 0,
                    'reg_lambda': 1,
                    'colsample_bytree': 0.8,
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
                    'colsample_bylevel': 1,
                    'colsample_bynode': 1,
                    'min_child_weight': 1.25,
                    'n_jobs': 11,
                    'validate_parameters': true
                    });
```

JSON example:

```
            {
              "classification_model_params": {
                'max_depth': 10,
                'learning_rate': 0.1,
                'objective': 'multi:softprob',
                'num_class': 11,
                'random_state': 10,
                'eval_metric': 'mlogloss',
                'booster': 'gbtree',
                'n_estimators': 150,
                'verbosity': 1,
                'use_label_encoder': false,
                'tree_method': 'hist',
                'num_parallel_tree': 1,
                'gamma': 0.25,
                'subsample': 0.8,
                'max_delta_step': 5,
                'reg_alpha': 0,
                'reg_lambda': 1,
                'colsample_bytree': 0.8,
                'colsample_bylevel': 1,
                'colsample_bynode': 1,
                'min_child_weight': 1.25,
                'n_jobs': 11,
                'validate_parameters': true
                }
              }
```

Related topics

- Using Machine Learning APIs
- Create and train a classification solution

Use LightGBM algo for classification model training

Enable the LightGBM (Light Gradient-Boosting Machine) algorithm as an advanced option for training classification models.

Before you begin

Role required: ml_admin or admin

About this task

The LightGBM algorithm offers performance and efficiency for large datasets. For detailed information about this algorithm see https://lightgbm.readthedocs.io.

Use the following procedure to add this option in the solution definition form. Alternatively, you can use a script to add this algorithm.

To add this option using a script:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

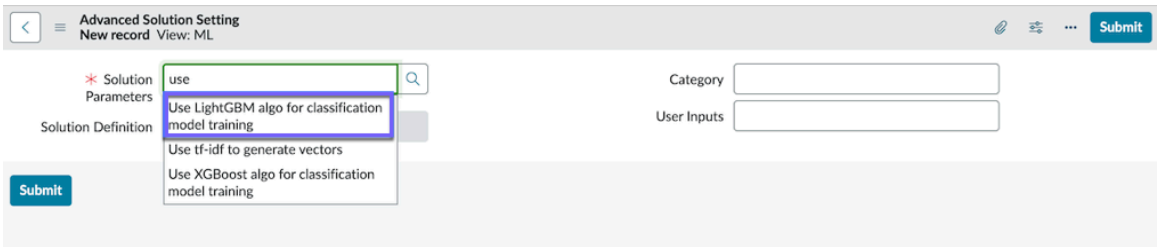
```
var config = new MLPredictorConfig('ml_solution_ID');
                    config.useLightGBM();
                    config.updateConfig();
```

In JSON format:

```
{
    "use_lightgbm": true
    }
```

Procedure

1. Navigate to **All** > *Predictive Intelligence* > **Classification** > **Solution Definitions**.
2. Open an existing classification solution or create a new one.
3. Select the **Advanced Solution Settings** tab in the **Related Links** section of the form.
4. Select **New** to open the **Advanced Solution Setting** (ml_advanced_solution_settings) form.
5. In the **Solution Parameters** field, search for `Use LightGBM algo for classification model training`.
6. Leave the **User Inputs** field blank, and select **Submit**.



Result

The setting appears as a row on the **Advanced Solution Setting** tab on your solution's form.

What to do next

Retrain your solution.

Related topics

Using Machine Learning APIs

Create and train a classification solution

Minimum records needed for label to include it

Set a threshold for the minimum number of records a label must have in your dataset to be included in model training.

Before you begin

Verify that your dataset includes label values suitable for classification.

Role required: ml_admin or admin

About this task

This advanced option helps to eliminate labels that occur infrequently. Setting this minimum threshold can reduce overfitting and unreliable predictions.

To configure, you provide a number for a minimum count of records. Any class value with fewer records than the minimum is dropped from the training set. Class values that are dropped before training can't appear in predictions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To add this option using a script:

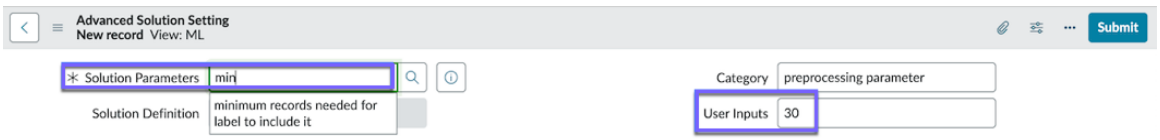
```
config.setMinLabelCount(50);
```

In JSON format:

```
{
    "min_label_count": 50
    }
```

Procedure

1. To add this option in the solution definition form, navigate to **All** > *Predictive Intelligence* > **Classification** > **Solution Definitions**.
2. Open an existing classification solution or create a new one.
3. Select the **Advanced Solution Settings** tab in the **Related Links** section of the form.
4. Select **New** to open the **Advanced Solution Setting** (ml_advanced_solution_settings) form.
5. In the **Solution Parameters** field, search for `minimum records needed for label to include it`.
6. In the **User Inputs** field, enter the number of records you want as the minimum, then select **Submit**.



Result
The setting appears as a row on the **Advanced Solution Setting** tab on your solution's form.

What to do next
Retrain your solution.

Related topics

Using Machine Learning APIs

Create and train a classification solution

Configure include only top N labels

Limit your classification model to use only the most common classes when training. You can specify the number of classes to use.

Before you begin
Role required: ml_admin or admin

About this task
Restricting the model to a fixed number of the most frequent labels improves focus and performance, especially with large datasets.

You can add this option in the solution definition form by using the following procedure. Alternatively, you can use a script to add this option.

To add this option using a script:

```
config.setTopNLabels(100);
```

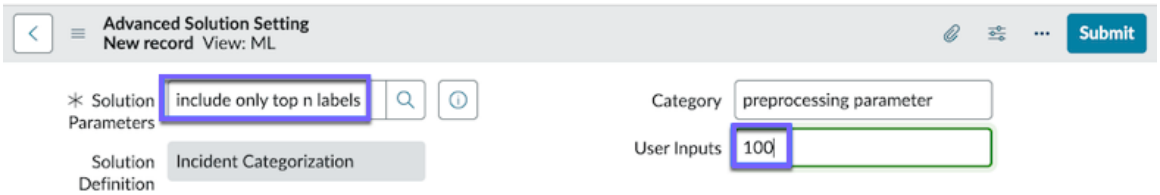
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In JSON format:

```
{
    "include_top_n_labels": 100
    }
```

Procedure

1. To add this option in the solution definition form, navigate to **All** > *Predictive Intelligence* > **Classification** > **Solution Definitions**.
2. Open an existing classification solution or create a new one.
3. Select the **Advanced Solution Settings** tab in the **Related Links** section of the form.
4. Select **New** to open the **Advanced Solution Setting** (ml_advanced_solution_settings) form.
5. In the **Solution Parameters** field, search for `include only top n labels`.
6. In the **User Inputs** field, enter the number of classes you want as the limit, then select **Submit**.



Result
The setting appears as a row on the **Advanced Solution Setting** tab on your solution's form.

What to do next
Retrain your solution.

Related topics

Using Machine Learning APIs

Create and train a classification solution

Remove others label

Reduce noise in your classification model and enhance predictive accuracy by removing records with the label "others" from training data. These are records with a distribution frequency of under one percent.

Before you begin
Role required: ml_admin or admin

About this task
When this option is set to true, records under 1% of total data distribution are excluded from training. These records are labeled "others" in the interface. The 1% threshold can't be modified.

You can add this option in the solution definition form by using the following procedure. Alternatively, you can use a script to add this option.

To add this option using a script:

```
config.removeOthersLabel(true);
```

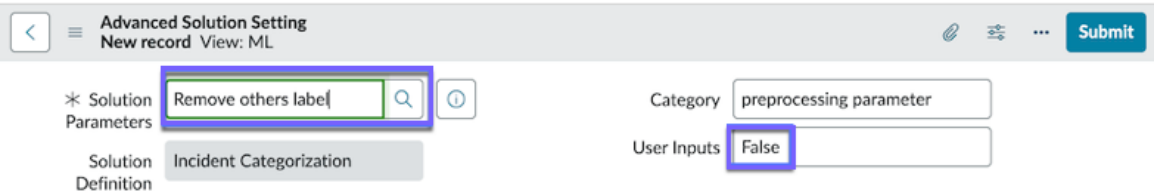
In JSON format:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
{
    "remove_others_label": true
    }
```

Procedure

1. To add this option in the solution definition form, navigate to **All** > *Predictive Intelligence* > **Classification** > **Solution Definitions**.
2. Open an existing classification solution or create a new one.
3. Select the **Advanced Solution Settings** tab in the **Related Links** section of the form.
4. Select **New** to open the **Advanced Solution Setting** (ml_advanced_solution_settings) form.
5. In the **Solution Parameters** field, search for `Remove others label`, then select **Submit** to save your update.



Result

The setting appears as a row on the **Advanced Solution Setting** tab on your solution's form.

What to do next

Retrain your solution.

Related topics

Using Machine Learning APIs

Create and train a classification solution

Creating and training solutions

Use one of the Predictive Intelligence (PI) frameworks to create and train machine-learning solutions. Each framework delivers a different solution type for training the system to predict, recommend, and organize data outcomes.

Types of solutions

The three PI frameworks provide different solutions that can be invoked by any application through a prediction API to make a prediction. Create and train your own solutions using your previous data. Navigate to **All** > **Predictive Intelligence** > **Homepage** to view and create solutions.

Select the best framework for your desired prediction:

- Classification solutions:

  Sets field values during record creation to automatically categorize and route work based on past records. See Create and train a classification solution.

- Similarity solutions:

  Identifies similarities between new and existing records to recommend resolutions. See Create and train a similarity solution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Clustering solutions:

  Groups similar records into clusters to identify patterns and major incidents. See Create and train a clustering solution.

- Regression solutions:

  Note: From the Washington DC release, support for creating new regression solutions was removed. You can still edit and train existing regression solutions, but you won't be able to initiate new ones.

  Uses historic data to predict numeric outputs, such as estimating the time it takes to resolve an incident or case. See Create and train a regression solution.

Selecting data records for training your solution

A solution is only as good as the record data you use to train it. In general, a good training dataset has these characteristics.

- The solution definition input fields are available to users when creating records. To make predictions at record creation, the solution must have the input field values at record creation.
- The solution definition output field is a choice field. To make more accurate predictions, limit the output field to a finite set of possible values.
- The training records only contain correct values for the output field. To make more accurate predictions, filter out any records that have unreliable output field values. For example, if recently closed incidents are subject to review and change for a month, filter out any recently closed incidents.
- The training records contain multiple examples of each output field value that you want the solution to predict. To provide more record coverage, include multiple examples of each output field value.
- The training records include common variations of the input fields. To provide more record coverage, include multiple examples of input field values.

Exporting your solution for training

To train a solution, you export its solution definition and associated records to a centralized training server within the same datacenter. When the training completes, the training server exports the solution back to your instance and deletes all of your training data from the server. Each datacenter has its own dedicated training server and the data doesn't leave the datacenter. Confirm that your configuration aligns with your compliance requirements.

Predictions occur on a centralized prediction server within the same datacenter as the instance. The trained model artifacts are sent from the instance server to the prediction server when the prediction is invoked for the first time. After that, the trained model artifacts are cached on the prediction server for subsequent predictions.

Note: All communication between the instance and the training service occurs within the same datacenter firewall. Even so, all communications occur over HTTPS.

Solution training troubleshooting

For troubleshooting common training issues, see the Predictive Intelligence Common issues [KB781893] article in the Now Support Knowledge Base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create and train a classification solution

Specify the records used to train a classification solution, what fields trigger a prediction, and how often you want to retrain your solution.

Before you begin

- Create a custom stopwords list if needed.
- Role required: admin or ml_admin

Important: In the Australia release, models in the classification, clustering, and similarity frameworks use Workflow solutions. These are pre-trained, so a word corpus isn't needed for your new solutions. When your existing solutions with a word corpus are re-trained after upgrading, they become Workflow solutions, and the Word Corpus field is removed from the form.

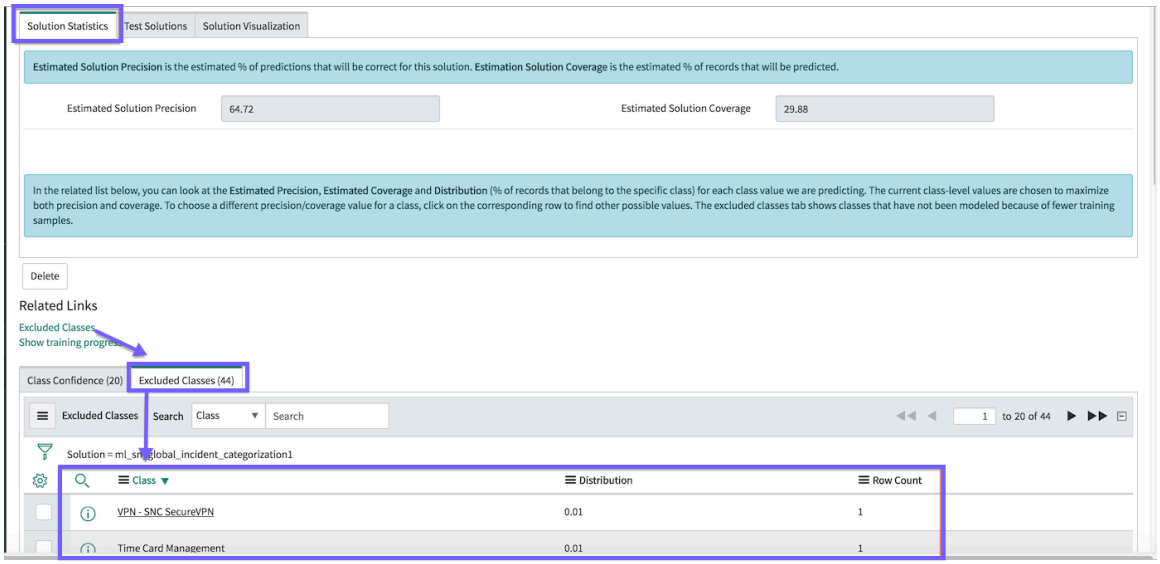
About this task

A predictive model is only as good as the data that you use to train it. To select appropriate records for training, examine the table's database dictionary as well as the current quality of the record values that you want to use.

For information on using encrypted training data, see Data Encryption in Predictive Intelligence.

For information about the minimum and maximum number of records you can use for training, see Predictive Intelligence properties.

Note: Classes that have fewer than 30 records in your training dataset are excluded from solution training. When your solution is trained and complete, any excluded classes are listed in the Solution Statistics section of your ML Solution form.



You must create a separate solution definition for each predictive model you want to support. The following procedure explains how to create a new classification solution, but you can also copy an existing solution definition and its configuration into a new record by selecting **Copy Solution Definition** from the context menu. Edit the field values on the new record as needed.

From the Yokohama release, you can also create a Workflow Classification solution using a script if you want to include an analysis of the key features influencing the model's predictions. For more information see Model Explainability.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Ensure that you are in the application scope you want for your solution definition, then navigate to **All** > **Predictive Intelligence** > **Classification** > **Solution Definitions**.
2. On the Classification Definitions list, select **New**.
3. On the empty Classification Definition form, configure the fields according to the following guidance.

Field	Value
Label	Enter a unique name for the solution record.
Name	The system generates the value of this read-only field based on the Label value that you entered.
Word Corpus	Select a word corpus that's relevant to your solution. For more information, see Create a word corpus. **Note:** Word Corpus is not a required field for customers implementing Predictive Intelligence for the first time starting in Utah. A pre-trained model is used instead. The Word Corpus field is removed for pre-trained models.
Table	Select the table containing the target records that you want the system to predict.
Output Field	Select the field whose value you want the predictive model to set. In general, a good output field has these characteristics. ○ It's a choice field or a string field with a finite set of possible values. ○ It has some causal connection to the input fields. For example, on the default **Incident Categorization** solution definition, the output field is set to **Category**.
Fields	Select the input fields that you want the solution to use to generate a prediction. Input fields are fields within a record that may contain the classification information your prediction solution needs to succeed. For example, if you're predicting the correct class for triaging an incident record, the prediction should gather records containing text that references the class. Most records have contextual text in their **Short description** field, so it's a great input field to use in general. You could also use **Resolution notes** as an input field, as it too might reference the incident class in the detailed notes for the incident. In general, good input fields have these characteristics. ○ The fields are available to users when creating records. ○ The field data type can be string, reference, choice, or HTML. The more information that a field provides, the more often a solution can make a prediction, and the more often predictions are accurate. ○ The field has a default value and should not be blank.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	All default solution definitions use the **Short description** field.
Filter	Click **Add Filter Condition** to apply conditions to the records you're training. For example, the **Incident Categorization** solution definition uses a filter with these conditions: **[Created][on][Last 12 months] AND [Active][is][false] AND [State][is one of][Resolved | Closed]** To train a solution, the filter must return at least one record. If your filter returns no records, update it until it returns records for training. **Note:** The recommended number of records for training a good solution is from 30,000 through 300,000. If you submit more than 300,000 records, the most recent 300,000 records are used to train the solution. Use only authentic records from the database. In general, a good filter has these characteristics. ○ The training records are inactive and their states indicate work completed within your standard process, such as resolved or closed. ○ The target fields contain only correct values. Filter out records with unreliable target field values. For example, if you're predicting the assignment group/ category and your historic incident data contains assignment groups/ categories that are no longer used, add a filter to remove such records from the training. ○ The training records contain multiple examples of each target field value you want the solution to predict. ○ The training records include common variations of the input fields. Use relative date filters such as last 3 months or last 12 months. Don't use hard-coded dates because these filters aren't updated when the solutions are retrained, unless you update them manually.
Processing Language	Select the dominant language of the dataset that you're training on the solution definition. If the dataset language is Italian, choose `Italian`. Also, English processing is applied to all datasets by default. For example, if you select Italian, the system processes the data in both English and Italian. **Note:** The term processing indicates some of the language-specific steps used as part of training a solution. For example, tokenizing words, removing stop words, and stemming.
Stopwords	When you select your processing language, the system automatically adds a Stopwords list for that language. For example, if your processing language is Italian, the **Default Italian Stopwords** list appears. The **Default English Stopwords** list is also included. If you create a custom stopwords list, you can select it from the Stopwords field to add to your solution.
Training Frequency	Select how often the system regenerates the solution. The available options range from **Run once** up to **Every 180 days**. **Note:** The minimum number of records required for classification solution training is set at 10,000.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	By default, the system runs training once. This provides you time to review and update the solution definition until it provides acceptable coverage and precision values. When your solution definition is fairly stable, consider scheduled trainings, as data can age over time, thus degrading the accuracy of your prediction model. **Note:** The ML scheduler limits the number of trainings an instance can commit to 50 new ML training requests per instance within a 24-hour window. This limit excludes scheduled retraining requests, clustering updates, and similarity updates, even if the new training requests exceed 50 within a 24-hour window.

4. Click the appropriate context menu option or button for your solution definition.

Option	Description
Save or Save & Train	Save your solution definition record so you can return to it later, or save and submit it for training.
Submit or Submit & Train	Create your solution definition record and submit it, or submit and train it.

5. If you submitted the solution for training, click **OK** on the **Training Activation** window to confirm.

   - The system schedules the solution for training with the nearest training service. The system sends you a notification when the training completes, including any errors that may have occurred in the training. Other users can subscribe to the Predictive Intelligence Notifications category. When training completes, the system uploads the solution as an Attachment record.
   - A bubble chart populates the Solution Visualization tab of your solution form, showing the estimated precision and coverage for each of the classes covered by the solution. The

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

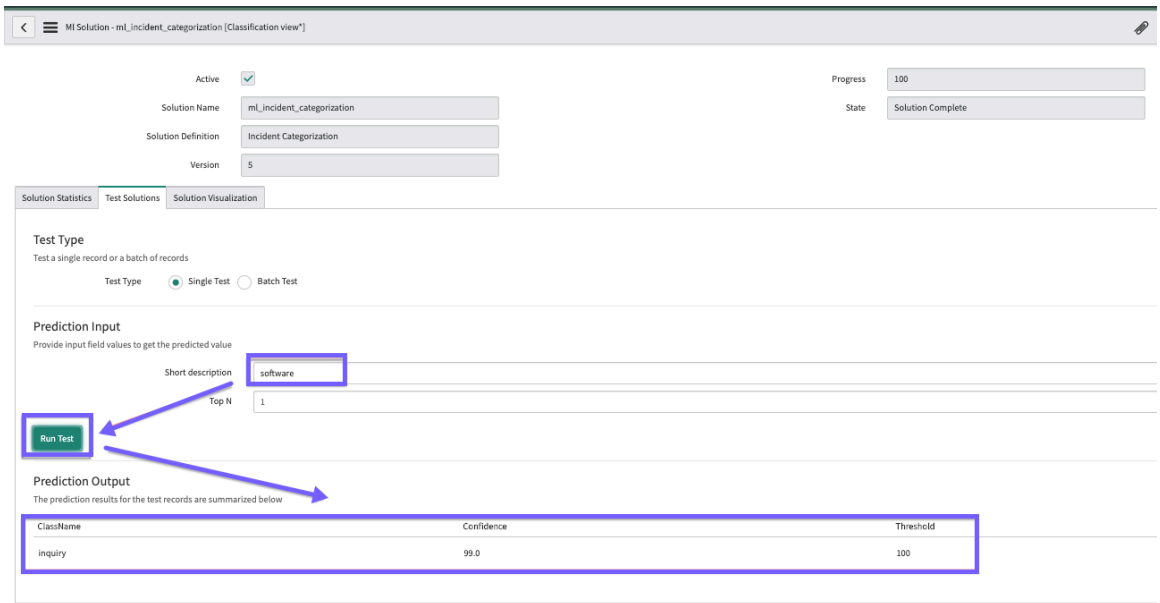
size of the bubble indicates the percent of records (distribution) that belong to the class. When you point to a bubble, you can see its estimated coverage, estimated precision, and distribution.



What to do next

In the Class Confidence section of the Solution Statistics tab in your solution, review the trained solution precision and coverage statistics.

In the Test Solutions tab in your solution, you can test the prediction output by entering values from the input fields, such as the Short Description.



Exclude a class from prediction

Exclude a class from prediction if its precision or coverage aren't satisfactory. Excluding a class prevents the classification model from predicting a particular output field value.

Before you begin

- Train the solution definition whose output field values you want to exclude.
- Role required: admin or ml_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task
As part of the testing and refinement of your classification solution, you can try excluding an output class to check the effect on the model's results.

Excluding a class from prediction is temporary. The class is restored the next time you train your solution, as long as the solution definition remains the same. If your precision or coverage targets are still not met, consider deactivating the solution or changing the solution definition.

Typically, you exclude a class from prediction if you want a person to manually set the excluded class value. For example, exclude a particular incident category when the solution doesn't offer sufficient precision or coverage, or because the class triggers other business logic that requires review or approval.

Procedure
1. Navigate to **All** > **Predictive Intelligence** > **Classification** > **Solutions**.
2. In the ML Solutions list, select the solution whose classes you want to exclude.
   This solution must have a **State** of **Solution Complete**.
3. In the **Class Confidence** related list, select the class you want to exclude.
4. In the Class Confidence record, review the precision and coverage combinations available from the **Precision Coverage Lookups** embedded list.
5. Select the check box for the 100 precision and 0 coverage combination.

   You can select only one check box.

6. From the Actions on selected rows control, select **Apply Values**.
   The system shows a **Precision / Coverage Setting** confirmation window.
7. Click **OK** to confirm the change or **Cancel** to discard it.

Result
The solution excludes the class from all predictions until the next training cycle.

What to do next
If you conclude that this class does not contribute to meaningful predictions, consider deactivating the solution or changing the solution definition.

Exclude a class from solution training

Exclude a class from solution training to prevent the model from ever making predictions for a particular output field class. For example, you can exclude a particular incident category from training if you plan to retire or change the category.

Before you begin
Role required: admin or ml_admin

About this task
Excluding a class from training doesn't prevent the solution from making predictions for records that use the excluded class. Solution training still uses the input and output field values as data and attempts to match input field values to a new output field class. This attempt may cause undesirable prediction results unless you have another suitable class to replace the excluded class value.

Typically, you only exclude a class from training if you change the list of valid output field values. For example, if you replaced one incident category with another incident category, you may exclude the old category from training so that the solution only uses the new category for predictions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: If you specify a target recall for a class, don't exclude the class from training even if the number of records are less than 30 for that class.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Classification** > **Solution Definitions**.
   The system shows the current list of solution definitions.
2. Select the solution definition you want to edit.

   Example
   For example, select **Incident Categorization** to exclude an incident category from training.
3. Edit the filter to exclude the class.
   You can use the **[is one of]** or **[is not one of]** operators to exclude a particular class.

   Example

   For example, if you want to exclude the Hardware class, add this condition: **[Category] [is not one of] [Hardware]**.
4. Click **Update & Train**.
   The system schedules the solution for training with the nearest training service. When training is complete, the system uploads the solution as an Attachment record.

Result
The solution excludes the class from all predictions.

What to do next
Review the trained solution precision and coverage statistics.

Tune a trained classification solution

Tune the performance of a trained classification solution by configuring class level precision and coverage values.

Before you begin

- Train the solution definition whose output field values you want to configure.
- Role required: admin or ml_admin

About this task
The system creates a class record for each output field value that it can predict. Each class record includes a list of possible precision and coverage combinations to choose from. By default, solutions use the highest combination of precision and coverage available. You can select another combination to refine predictions based on acceptable precision and coverage values.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Classification** > **Solutions**.
   The system shows the list of available solutions.
2. Select the solution whose classes you want to configure.
   This solution must have a **State** of **Solution Complete**.
   The system shows the Solution record.
3. From the **Class Confidence** related list, select the class you want to configure.
   The solution only lists output field values for which it can make predictions. If the output field value is missing from this list, update the solution definition filter to provide more data for this output field value, and retrain the solution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The system shows the Class Confidence record.

4. Review the precision and coverage combinations available from the **Precision Coverage Lookups** embedded list.
5. Select the check box for the precision and coverage combination you want to use to make predictions for this class.

   You can only select one check box. Some combinations produce special prediction results.

Special prediction combinations

Prediction result	Precision	Coverage
Never include class in predictions	100	0
Always include class in predictions	0	100

6. From the Actions on selected rows control, select **Apply Values**.
   The system shows a **Precision / Coverage Setting** confirmation window.
7. Click **OK** to confirm the change or **Cancel** to discard it.

What to do next

Test predictions for this class to verify that the system produces acceptable results.

Configuring target metrics for a trained classification solution

Set values for precision, coverage, and recall statistics for a trained machine learning solution.

Setting classification metric values at the class or solution level

Predictive Intelligence provides three classification metric types: precision, coverage, and recall. You configure these metrics on the Solution Statistics tab of a trained classification solution form. While you can manually set values to these metrics at the class level, doing so can be challenging if you have a large number of classes to cover. In many cases, you may not know the best value to set until your solution is trained. This topic focuses on setting the metric values at just the solution level.

Configuring solution metrics

When you apply a value to one metric, it changes the values of the other two. This behavior enables you to modify your metrics iteratively in real time to see which value combinations render particular results. When you apply a new value to a metric, the system recomputes it by considering its new targets.

Applying a value to a metric asks the system to train its predictions to favor the metric you set based on the highest percentage value, and at a cost to the other metrics. The system tries to meet these values but may not set them exactly as you request due to how the data you're training is distributed.

When you apply metric values at the solution level, the system automatically sets the appropriate values at the class level.

Here are the basic steps for configuring a target metric for your solution.

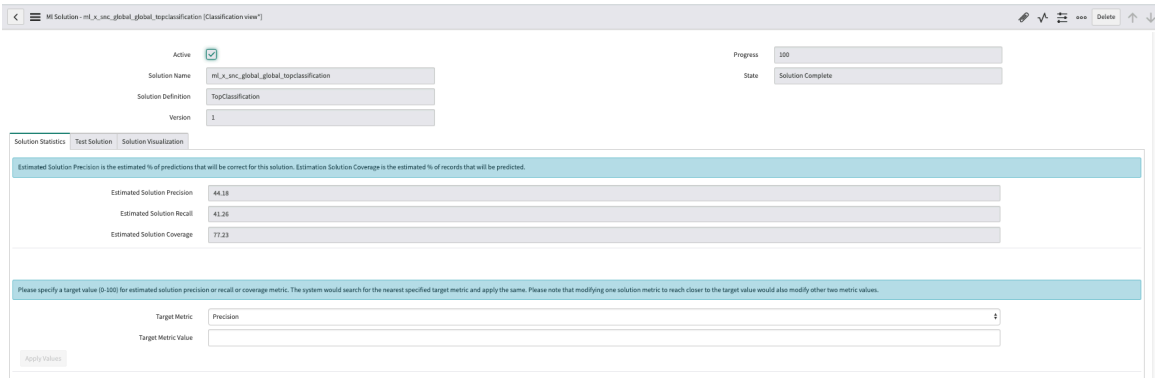
1. Navigate to the Solution Statistics tab of a trained ML solution.
2. Review the messages on the green banners of the screen which define each of the metrics so you can better understand the values you want to assign to the solution. The first two message

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

banners address estimated solution-level metrics. The third banner addresses class-level results based on the solution values you applied.

3. In the **Target Metric** choice list, select the metric you want to configure.
4. In the **Target Metric Value** field, enter a numeric percentile value between 0-100.
5. Click **Apply Values**.
6. Result: On the Solutions Statistics tab, you can review the change in values to the **Estimated Solution Precision**, **Estimated Solution Recall**, and **Estimated Solution Coverage**. The system calculates these values based on the **Target Metric** you select and the **Target Metric Value** you enter for the solution.

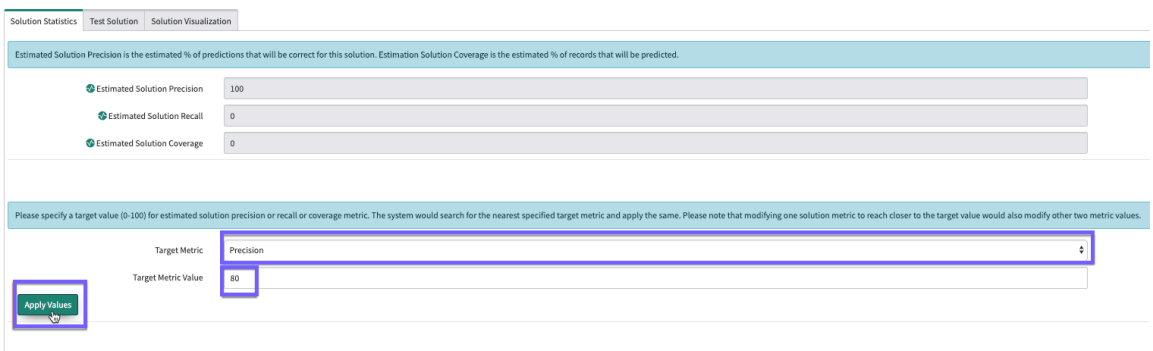
Here's a sample landing page for a recently trained classification solution. As you can see, the precision metric is 44.18, recall is 41.26, and coverage is 77.23.



If you need to adjust these default values for a use case, refer to the sample configurations below. For example, based on the classification solution you're implementing, you might want to change the target metric value for precision, recall, or coverage. Keep in mind that when you change the target metric value for one metric, such as precision, it impacts the values of the recall and coverage metrics as well.

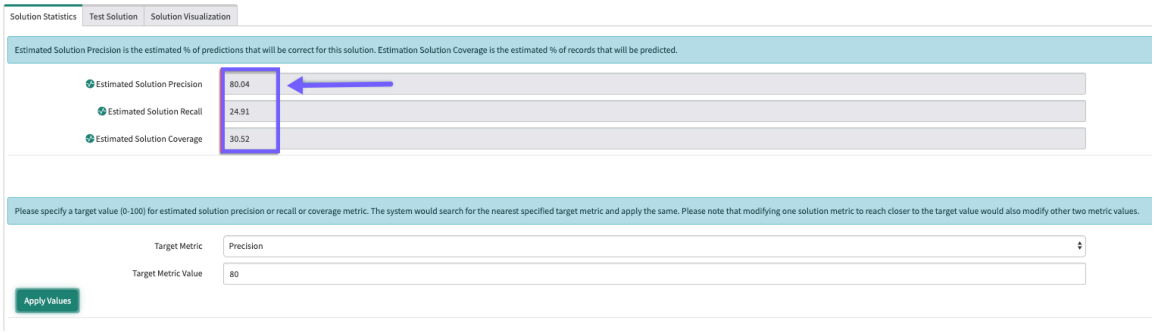
Precision configuration example

In this example scenario, you're replacing a manual triage process for routing incident records with an ML classification solution that automatically assigns the records to the correct assignment group. For this scenario, you have a target value in mind and the system must predict correctly at least 80% of the time. So you set the precision metric value to 80 and click **Apply Values**.



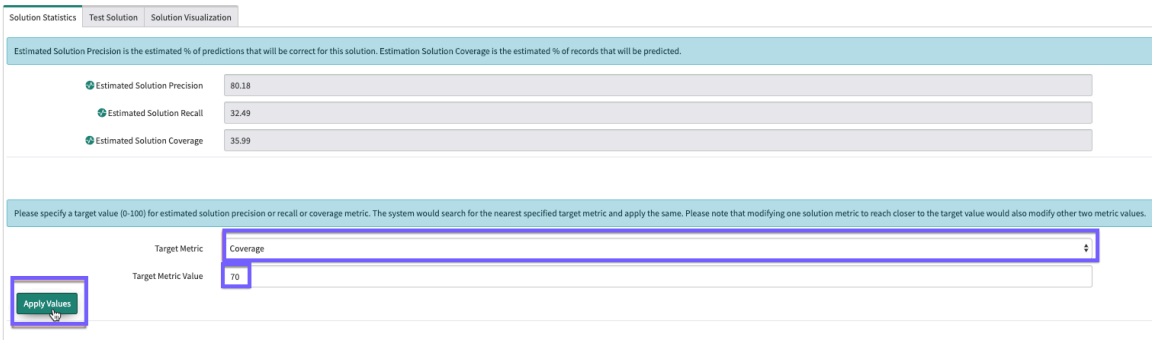
Here are the metric values the system applied to the solution. In this scenario, the precision value of 80.04 slightly exceeded your request for 80%, so you're likely satisfied with that value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Coverage configuration example

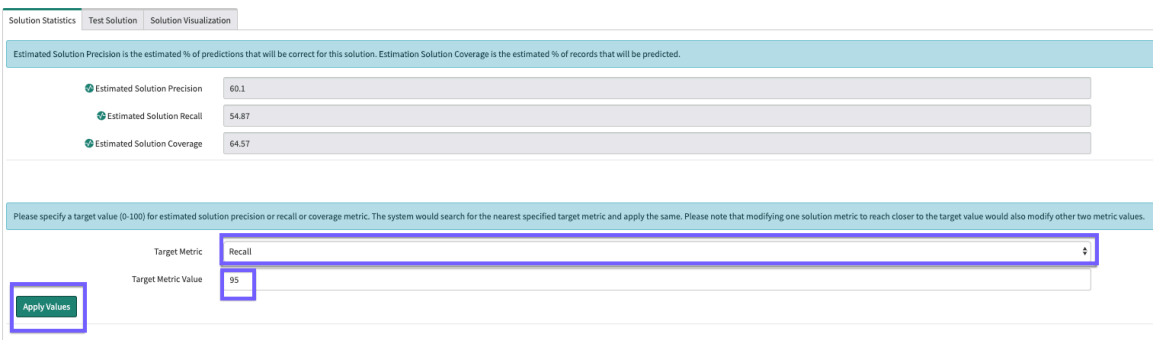
In another example scenario where you're replacing a manual triage process for routing incident records, your minimum goal is to predict at least 70% of incoming incidents in the first quarter of the year. So you set the coverage metric value to `70` and click **Apply Values**.



The metric values the system applied to the solution are shown in the following image. The coverage metric value increased from 35.99 to 55.98. However, the precision metric decreased from 80.18 to 64.97. This could be because you set the coverage metric to a relatively high value of 70, or perhaps because of how the data you're training is distributed.

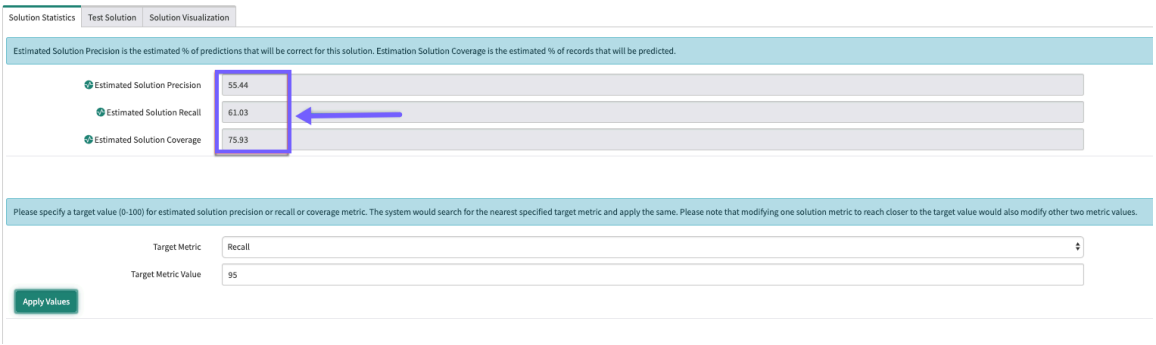
Recall configuration example

In another scenario, classifying if an incoming email is a Phish or not can be an important use case in a security-related machine learning solution. In this situation, it's very important to identify every Phish, and it may be okay to report a non-Phish as a Phish occasionally. However, no real Phish should be classified as a non-Phish. In such situations, the recall metric must have a high value, which might lead to lower percentages for precision and coverage. So here you can set the recall metric to `95` and click **Apply Values**.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

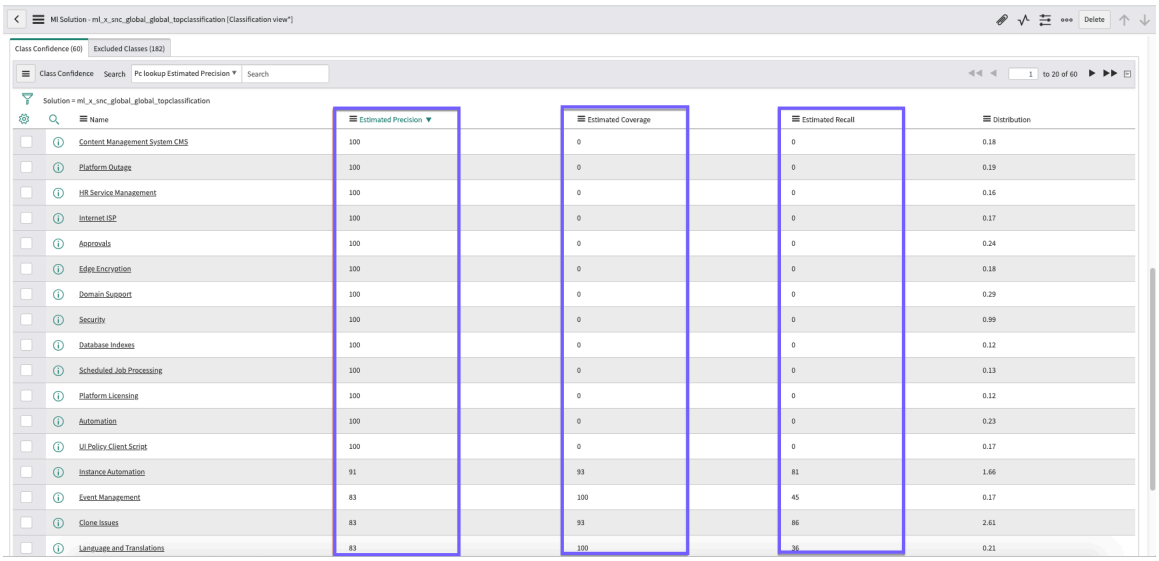
Here are the metric values the system applied to the solution. The recall metric value increased from 54.87 to 61.03. However, the precision metric decreased from 60.1 to 55.44. This is likely because you set the recall metric to the high value of 95.



Class-level results for the solution metric values you apply to your solution

The following image shows an example of the class-level results the system applied to a solution's precision, coverage, and recall statistics for 37 classes. You can keep modifying the metric values until you're fully satisfied with the results.

By Sorting (z to a) on the Estimated Precision column you can see which classes have the highest precision for the solution.



Using Group By for classification

Use APIs to simultaneously submit multiple classification solutions for training based on the Group By field.

You can use the optional Group By capability to train and maintain one classification solution that covers more than one data area, such as geographical location or domain.

To train a solution using Group By, you must add the groupby parameter while creating a classification solution definition using APIs. The groupby parameter accepts only categorical columns as inputs, where individual models are created on the subset of data belonging to each of the groupby values. Only those child solutions that pass the minimum records criteria set for the capability are created. Here, the prediction calls are routed to the corresponding Group By model based on the Group By value present in the prediction input. Batch predictions are not supported.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

A Group By scenario for geographical locations

Let's say your global company uses classification routing for incoming records, with one support center in the US and one in Europe. Here, you want to create a single classification solution that has one model for your United States incidents and another model for your European incidents.

In this scenario, you could use one of these two approaches:

- Create and train two separate ML classification solution definitions, where one is filtered by US incidents only, and one by European incidents only.
- Use the groupby parameter to create Groupby for the country location so that all US definitions create a US model and all European definitions create a European model. Then, based on the incident, the system identifies which model it uses to predict the correct classification category.

The second approach has benefits in that the models you use can even be in different domains, such as healthcare or finance. This approach is especially beneficial if you have several country locations or domains to maintain.

Example usage for training and prediction using Group By via API

```
var myIncidentData = new sn_ml.DatasetDefinition({
    'tableName' : 'incident',
    'fieldNames' :
 ['category','short_description','assignment_group','description
','priority'],
    'encodedQuery' : 'activeANYTHING'
});

var mySolution = new sn_ml.ClassificationSolution({
    'label': 'solution label',
    'dataset' : myIncidentData,
    'groupByFieldName' : 'assignment_group',
    'predictedFieldName': 'category',
    'inputFieldNames':
 ['short_description','description','priority']
});
//Add solution definition
var solution_gr =
 sn_ml.ClassificationSolutionStore.add(mySolution)
//Get existing solution
var my_unique_name =
 sn_ml.ClassificationSolutionStore.get('solution name');
// submit training job
var solutionVersion = my_unique_name.submitTrainingJob();


// Run prediction
var input = new GlideRecord("incident");
input.get("sys_id");
// configure optional parameters
var options = {};
options.apply_threshold = false;
var mlSolution = sn_ml.ClassificationSolutionStore.get('solution
 name');
//Prediction using glide record
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
var results = mlSolution.getActiveVersion().predict(input,
 options);
//Prediction using map
var results =
 mlSolution.getActiveVersion().predict([{ 'short_description':
 input.short_description,

 'assignment_group': input.assignment_group }], options);
```

For more context regarding this example and the general usage of Machine Learning APIs, see the links in the Related Content section on this page.

Related topics

- DatasetDefinition - Global
- ClassificationSolution - Global
- ClassificationSolutionStore - Global
- ClassificationSolutionVersion - Global

Model Explainability

Analyze the importance of each input field to your model's predictions using model explainability. Create a Workflow Classification model that includes a graphical analysis of feature importance by executing the provided script.

Before you begin

- This method uses the Workflow Classification Solution API, instead of the Solution Definition form, to create and train a model with explainability added. For information about the components of Workflow Classification models, see Create and train a classification solution.
- Role required: ml_admin or admin

About this task

Model explainability helps identify the key features that influence the model's predictions during training.

> **Note:** Explainability can't be added to an existing model. This method uses a script to create and train a new Workflow Classification model. For more information about scripting the creation of Classification solutions, see ClassificationSolution - Global.

The script provided in the procedure creates and trains a model with explainability set to true. On the new model's solution form, an additional tab labeled **Feature Importance** appears. This tab offers a graph of the relative contribution of each input to the prediction.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System Definition** > **Scripts - Background**.
2. Edit the query filter and table, field, and variable values in the following script according to your planned model, then execute the script.

```
// Define a dataset
                    var myIncidentData = new
sn_ml.DatasetDefinition({
                    'tableName': 'incident',
                    'fieldNames': ['category',
'short_description', 'sys_updated_by', 'assignment_group',
'description', 'priority'],
                    'encodedQuery': 'activeANYTHING'
                    });

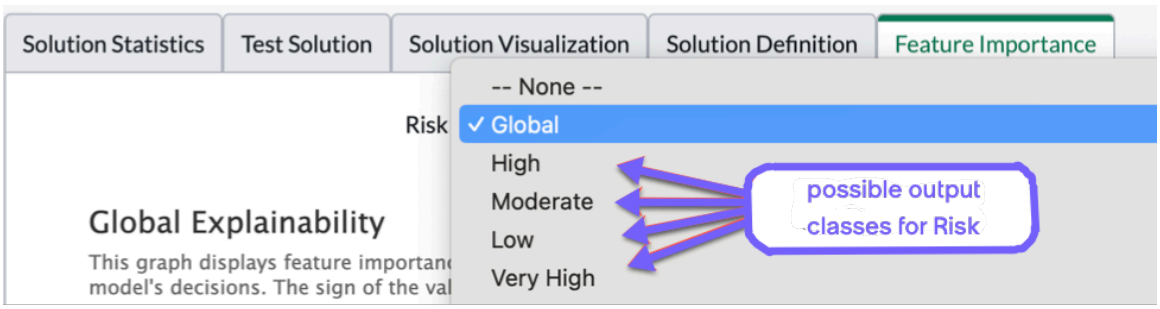
                    // Define a classification solution
definition with explainability field
                    var mySolution = new
sn_ml.ClassificationSolution({
                    'label': 'model explainability',
                    'dataset': myIncidentData,
                    'predictedFieldName': 'category',
                    'inputFieldNames': ['short_description',
'priority'],
                    //setting the explainability field to true.
                    'explainability': true,
                    });

                    // Add solution to
ClassificationSolutionStore
                    var my_unique_name =
sn_ml.ClassificationSolutionStore.add(mySolution)

                    // Submit training job
                    var solutionVersion =
mySolution.submitTrainingJob();
```

Note: Substitute the query filter and table, field, and variable names in this script with your own values.

3. Navigate to the ML Solutions [ml_solution] table and open your new solution by selecting the value of its **Active** field.
4. On the solution form, locate and open the **Feature Importance** tab.
Feature Importance displays a drop-down list.



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- The label for this drop-down list is the name of your output (predicted) field, so the label is specific to each model.
- The values in the drop-down list are the possible output classes for your output field, plus the **Global** option.

List option	Description
Global	Provides an overview of how the model be haves across all predictions, showing the overall impact of each input feature. Select **Global** to open a graph of the impor tance of your input fields to predictions for all output classes as a whole.
`Your output class value`	Focuses on the model's behavior for the cho sen class only, showing how input features contribute to predictions on a per-class basis. Select one of the possible output classes to open a graph of the importance of your input fields to predictions for that class.

5. Launch the graphical analysis by selecting a value from the drop-down list.
   The y axis shows your input fields and the x axis shows the numerical importance for each input. The label of the graph reflects the class that you chose in the drop-down list.



   You can hover over a bar to display the numerical percentage for each input's importance.

Result

A positive importance value means that the input field increases the model's prediction score. A negative value means that the input field decreases the prediction score.

What to do next

Consider dropping input fields with low importance scores. Retrain your model after modification.

Related topics

Using Machine Learning APIs

ClassificationSolution - Global

Create and train a classification solution

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create and train a similarity solution

Create and train a machine learning solution to collect and compare your existing records to new similar records. For example, you can compare the text in an open Incident record to a resolved Incident record to reuse its resolution.

Before you begin

- Role required: ml_admin or admin

Important: From the Washington DC release, models in the classification, clustering, and similarity frameworks use Workflow solutions. These are pre-trained, so a word corpus isn't needed for your new solutions. When your existing solutions with a word corpus are re-trained after upgrading, they become Workflow solutions, and the Word Corpus field is removed from the form.

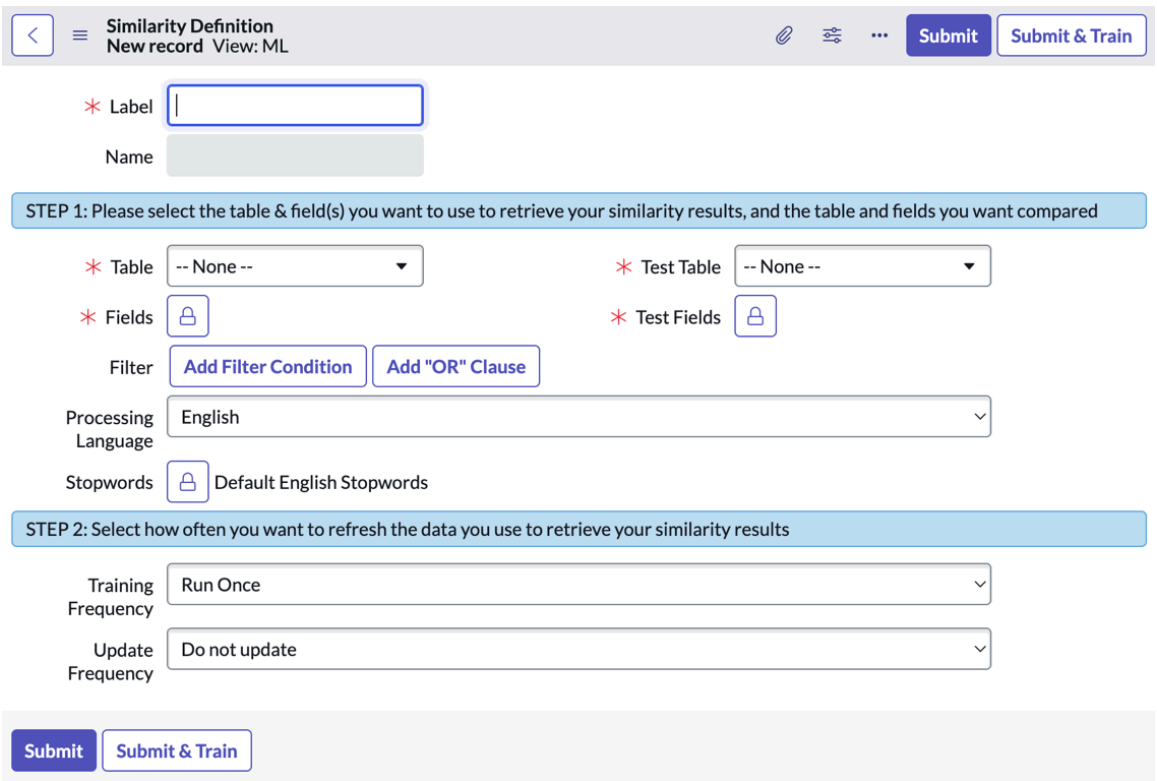
About this task

After comparing your existing records based on similarity, the system recommends examples that you can review and reuse in your solution.

When applied in your forms and flows, similarity solutions are domain-aware, so records from other domains on the instance are not displayed to users. For more information, see KB article Similarity prediction behavior in domain separated environment on Now Support.

For information on using encrypted training data, see Data Encryption in Predictive Intelligence.

In this example procedure, you're working on Incident records and you want to locate relevant Knowledge Base articles that could provide resolutions to those incidents.



Procedure

1. Ensure that you are in the application scope that you want for your solution definition, then navigate to **All** > **Predictive Intelligence** > **Similarity** > **Solution Definitions**.
2. In the Similarity Definitions list, select **New**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. On the Similarity Definition form, fill in the fields.

Field	Value
Label	Enter a unique name for your similarity solution. For example, in this use case you could enter `Match Knowledge Articles to Incidents`.
Name	As you enter a Label value, this field automatically populates with a system-assigned, read-only name based on your Label value.
Word Corpus	If you have a legacy similarity solution, you can select a relevant word corpus from the **Word Corpus** field in the definition form. **Note:** Starting from the Washington DC release, a word corpus is not required because a pre-trained model is used instead. The **Word Corpus** field is not visible in the definition form for pre-trained models. For more information, see Create a word corpus.
Table	In the Table field, select the table that contains records you want to use as an information source. In this use case you select the **Knowledge [kb_knowledge]** table, because its KB Article records might provide information relevant to the Incidents that you're trying to resolve. After you assign a Table, the number of records matching your filter conditions is displayed as a link. Select this link to view the list of records.
Test Table	In the Test Table field, select the table containing the records that you want to target. In this use case, you select the **Incident [incident]** table, as it contains the Incident records that you're trying to resolve. **Note:** You can select the same table for Table and Test Table. For example: using filter conditions, you could collect information from recent Incidents to help with target Incidents.
Fields	For the Table that you selected, enter fields that are likely to contain words and phrases relevant to the Incidents you're trying to resolve. In this example, you choose **Short description** and **Article body**. Including the article body increases your chances of capturing informative details regarding the subject. **Note:** **Journal Type** is not a supported data type.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
Test Fields	For the Test Table that you selected, enter fields that contain text that you want to compare to other similar records. In this example, you choose the **Short description** of the Incident records you're trying to resolve.
Filter	Select **Add Filter Condition** to apply conditions to the Fields records you're using as an information source. For example, in this use case you could set a **workflow_state=published** condition to retrieve published KB articles only. **Note:** Script includes can't be referenced from the Filter. Use database views as an alternative.
Processing Language	Select the dominant language of the dataset you're training on. Also, English processing is applied to all datasets by default. For example, if you select Italian, the system processes the data in both Italian and English. **Note:** The term processing indicates some of the language-specific steps used as part of training a solution, such as tokenizing words, removing stop words, and stemming.
Stopwords	When you select your processing language, the system automatically adds a Stopwords list for that language. For example, if your processing language is Italian, the **Default Italian Stopwords** list appears. The **Default English Stopwords** list is also included. To use a custom stopwords list, select the lock icon( ) and then search in the **Select target record** field.
Training Frequency	Select a retraining frequency. The available options range from **Run Once** up to **Every 180 days**.
Update Frequency	Select how often you want to refresh the data you use to retrieve your similarity results. For example, for open incident records, you could select an update frequency of **Every 15 minutes**, as new incidents typically occur frequently throughout the day. This frequency may increase the likelihood that newly opened records are included in the refresh. However, for KB Knowledge article records, which are typically not created often, you could choose a less frequent update frequency such as **Every 1 day**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	Note: The ML scheduler limits the number of trainings an instance can commit to 50 new ML training requests per instance within a 24 hour window. This excludes scheduled re-training requests. In addition, clustering and similarity updates are also excluded from this limit, even if the new training requests exceed 50 within a 24 hour window.

4. Select the appropriate button for the solution definition.

Option	Description
Save	Save your solution definition record so you can return to it later.
Submit & Train	Create your solution definition record and train it.

5. If you submitted the solution for training, select **OK** on the **Training Activation** window to confirm.

Result

- The system schedules the solution definition for processing with the nearest training service and sends you a notification when the training completes. The notification includes any errors that may have occurred during the training. Other users can subscribe to the Predictive Intelligence Notifications category.
- The trained solution delivers paired example records ranked by their degree of similarity. The list of examples is available from **Similarity Examples** under Related Links on your solution's form.
- When training completes, the system uploads the solution as an Attachment record.

What to do next

Review the similarity examples on the Related Links section of your Solution form. See Review solution similarity examples.

Update your similarity score threshold

After you review the similarity examples provided by the system, update your solution similarity score threshold if you want the results returned by the solution to be more or less similar.

Before you begin

- Review your similarity examples and their scores. For more information see Review solution similarity examples.
- Role required: ml_admin or admin

About this task

A deployed solution returns records only when the similarity score is higher than the **Similarity Score Threshold** value:

- The higher the threshold value, the higher the precision and the lower the coverage.
- The lower the threshold value, the lower the precision and the higher the coverage.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Similarity** > **Solutions**.
2. In the ML Solutions list, locate your solution by filtering for the Solution Name, and open its active version.
3. On the solution form, locate and open the **Solution Statistics** tab.
4. In the **Similarity Score Threshold** field, enter a new numerical value that represents a percentage.
   For example, imagine that the current threshold is 80. In your similarity example review you decided to increase the accuracy of your similarity recommendations at the cost of lowering the prediction coverage. So you update the field by entering the higher value of 90.
5. In the form's context menu (Additional actions), select **Save**.
   Your solution uses the new threshold value that you assigned to it and returns similar results that have a score higher than 90. If you set the threshold to 90, the similarity score must be 91% or above to return records.

Create and train a clustering solution

Group similar records into clusters so you can address them collectively or identify patterns.

Before you begin

Role required: ml_admin or admin

Important: In the Australia release, models in the classification, clustering, and similarity frameworks use Workflow solutions. These are pre-trained, so a word corpus isn't needed for your new solutions. When your existing solutions with a word corpus are re-trained after upgrading, they become Workflow solutions, and the Word Corpus field is removed from the form.

About this task

In this example procedure, you're creating a solution to identify a major incident by grouping similar incidents that have occurred recently.

For information on using encrypted training data, see Data Encryption in Predictive Intelligence.

Procedure

1. Ensure that you are in the application scope that you want for your solution definition, then navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solution Definitions**.
2. On the Clustering Definitions list, select **New**.
3. On the Clustering Definition form, configure the fields according to the following guidance.

Field	Value
Label	Enter a unique name for your clustering solution. For example, in this use case you could enter `Group Incidents to a Major Incident.`
Name	As you enter your solution Label, this field automatically populates with a system-assigned name based on your Label value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
Word Corpus	If you have a legacy clustering solution, you can select a relevant word corpus from the **Word Corpus** field in the definition form. **Note:** With the Australia release, a word corpus is not required, because a pre-trained model is used instead. The **Word Corpus** field is not visible in the definition form for pre-trained models. For more information, see Create a word corpus.
Table	Select the table that contains record types that you want to group into one or more clusters. For example, in this use case you select the **Incident [incident]** table as it contains incident records you want to group together for a major incident analysis. When you assign a table value, a link appears in the form that shows the number of records that match your current conditions.
Fields	Select one or more input fields types that help the system identify the records you want to include in your cluster. In this use case, use **Short description**. **Note:** When selecting a reference type field, you must dot-walk to the field's property name. For example, instead of `short_description`, enter `short_description.name`.
Use Group By	Select this check box only if you want to group input records by a field before creating clusters. **Note:** Selecting this check box activates the **Group By** list. If you don't select the check box, all the table records are grouped into clusters.
Group By	Selecting a value from this list is optional. If you do so, the system groups records into one or more clusters based on your selection.
Purity Fields	Choose fields from your table that can help the system identify the class that is most frequent in the cluster. In this example scenario, select **Category** and **Assignment group.Name**.
Filter	Add filter conditions to apply to the input field records that you want to include in your clusters. ○ The maximum number of records for clustering is limited to 300,000. ○ For best results, aim for at least 2000 records as a minimum.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	Note: Script includes can't be referenced from the Filter. Use database views as an alternative.
Processing Language	Select the dominant language of the dataset you're training on the solution definition. If the dataset language is Italian, choose **Italian**. Also, English processing is applied to all datasets by default. For example, if you select Italian, the system processes the data in both English and Italian. **Note:** The term processing indicates some of the language-specific steps used as part of training a solution. For example, tokenizing words, removing stop words, and stemming.
Stopwords	When you select your processing language, the system automatically adds a Stopwords list in that language. For example, if your processing language is Italian, the **Default Italian Stopwords** list appears. The **Default English Stopwords** list is also included. If you create a custom stopwords list, you can select it from the Stopwords field to add to your solution.
Update Frequency	Select how often you want the system to update your clusters with new and updated records. **Note:** The system pulls records based on the Group By filter conditions that you set on your clustering solution, if any. For example, if you select **Every 15 minutes**, the system identifies which records have arrived within that time frame. The system tries to assign them to the existing clusters, or creates a new cluster if possible. In this example, 20 new records arrive. If 16 of those records make it into an existing cluster and 4 don't, the system forms a new cluster for the four unassigned records. You can also choose not to update your clusters at all.
Training Frequency	Select how often you want the system to discard all previous cluster results and recreate clusters from the beginning. Your options range from daily, every third day, every seven days, or monthly. You can also choose to train your cluster once.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	Note: The ML scheduler limits the number of trainings an instance can commit to 50 new ML training requests per instance within a 24-hour window. The limit excludes scheduled retraining requests. In addition, clustering and similarity updates are also excluded from this limit, even if the new training requests exceed 50 within a 24-hour window.
Minimum number of records per cluster	Enter the minimum number of records you want a cluster to contain. The value you enter must be 2 or higher.

4. Select the appropriate context menu option or button for your solution definition.

Option	Description
Save or Save & Train	Save your solution definition record so you can return to it later, or save and submit it for training.
Submit or Submit & Train	Create your solution definition record and submit it, or submit and train it.

5. If you submitted the solution for training, select **OK** on the Training Activation window to confirm.

Result

The system trains the solution and notifies you in real time when the training completes.

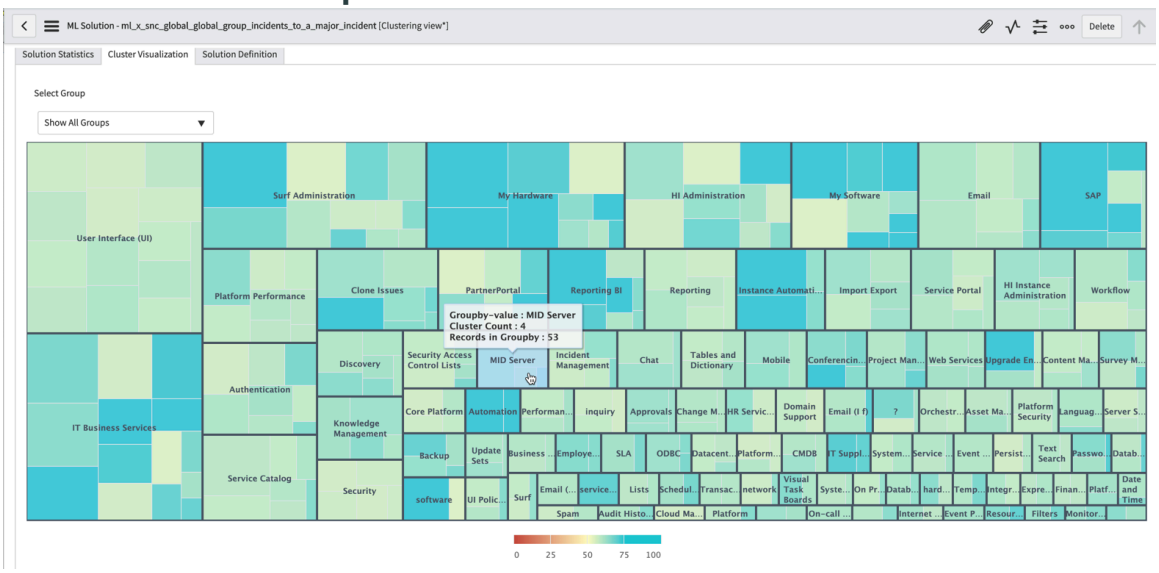
A treemap plot appears on the Cluster Visualization tab of your Clustering Solution Definition form. The plot shows the clusters the system formed for your solution in descending order from the top-left corner to the bottom-right corner. The treemap node labels are the Cluster Concept, which is created by the top words from the cluster, and helps you see the most prominent content found in each cluster.

Note: The Cluster Concept displays the top words from the processed input data, in the data's language. Depending on the language, the Cluster Concept may display words in their root form and so appear truncated.

Each node is colored from red to green depending on the cluster quality for that node. The Select Group filter appears only when you select the **Use Group By** and **Group By** fields on your Clustering Definition form. When you point to a cluster, you can see its Groupby value, Cluster Count, and Records in Groupby.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

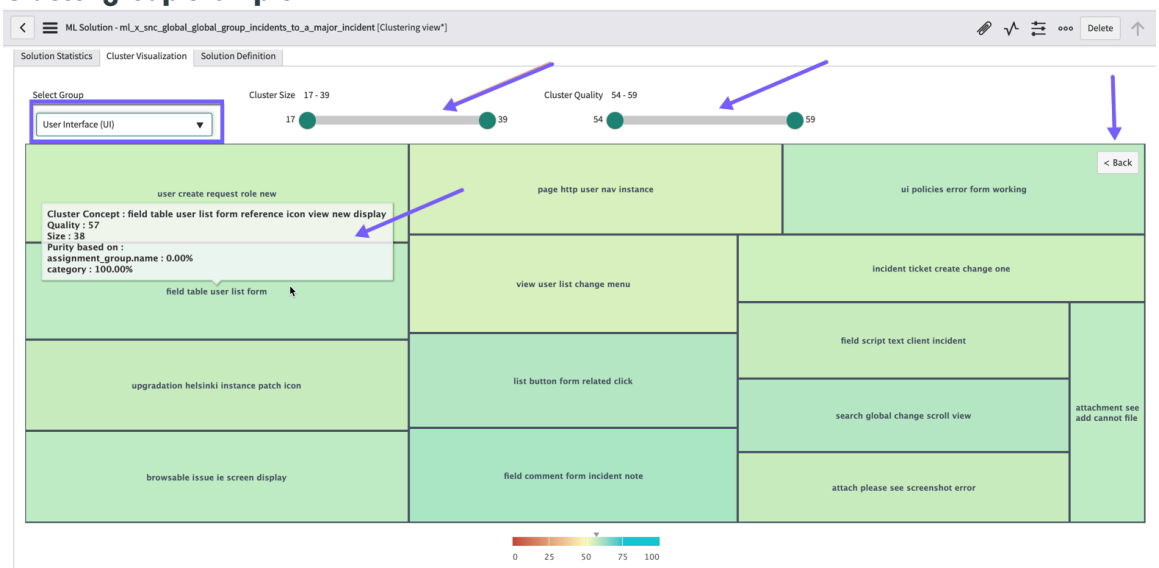
Cluster visualization example



To open a cluster, you can click it, or select it from the Show All Groups filter.

Inside the cluster grouping, you can filter the results further by using the two slide bars for cluster size and cluster quality, respectively. You can also navigate backward by clicking the Back button, which only appears when a clustering hierarchy is present. When you point to a cluster at this level, the Purity field percentile values appear along with the Cluster Concept, Quality, and Size values.

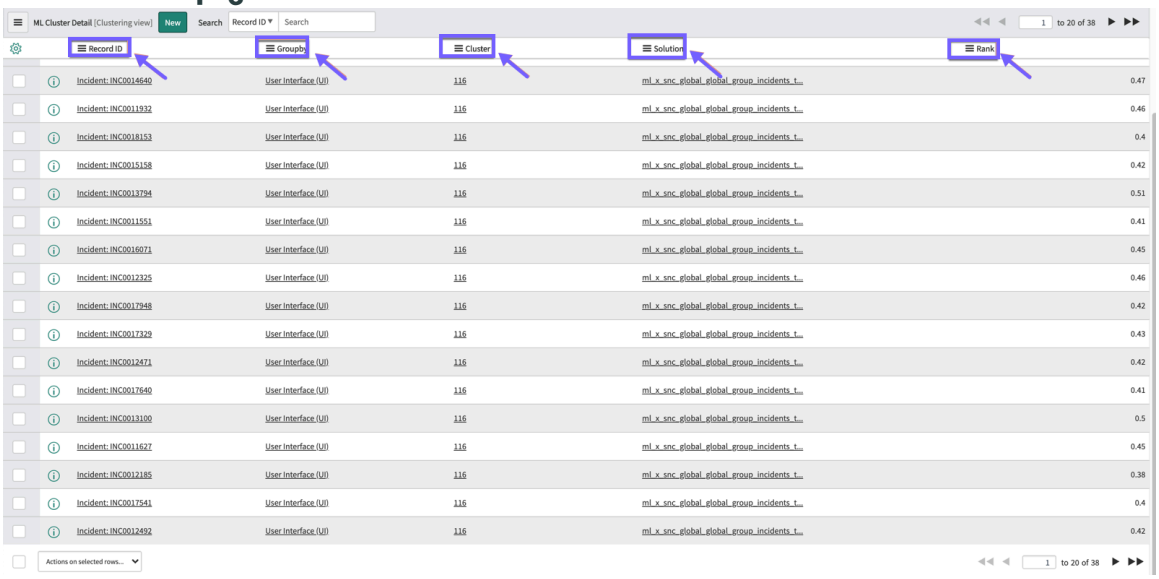
Cluster group example



When you click a cluster node, its ML cluster details appear in a list view format.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cluster details page

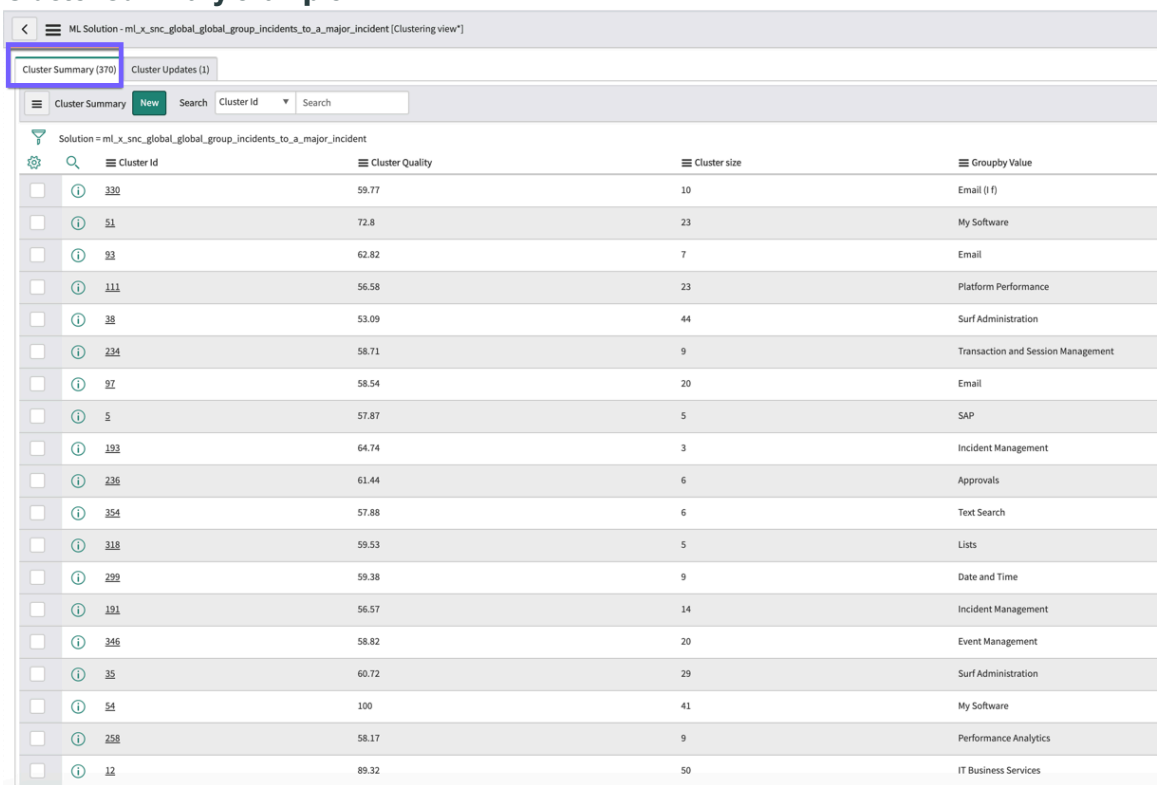


What to do next

- Review the solution output on the Solution Statistics tab of your solution. If you aren't satisfied with your clustering solution results, reconfigure the values you've set to your solution and retrain it until the results are to your satisfaction.

- Review the Cluster Summary tab for a list view of the cluster IDs, quality size, and Groupby values.

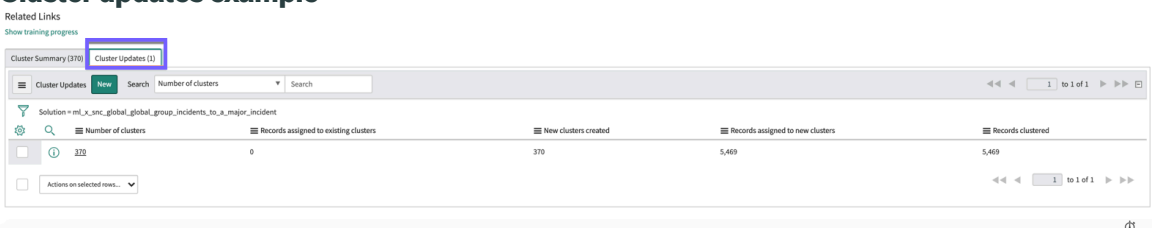
Cluster Summary example



© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- On the Cluster Updates tab, review the summary of changes to clusters for each cluster update interval you configured in the solution definition.

Cluster updates example



Assign a name to a cluster

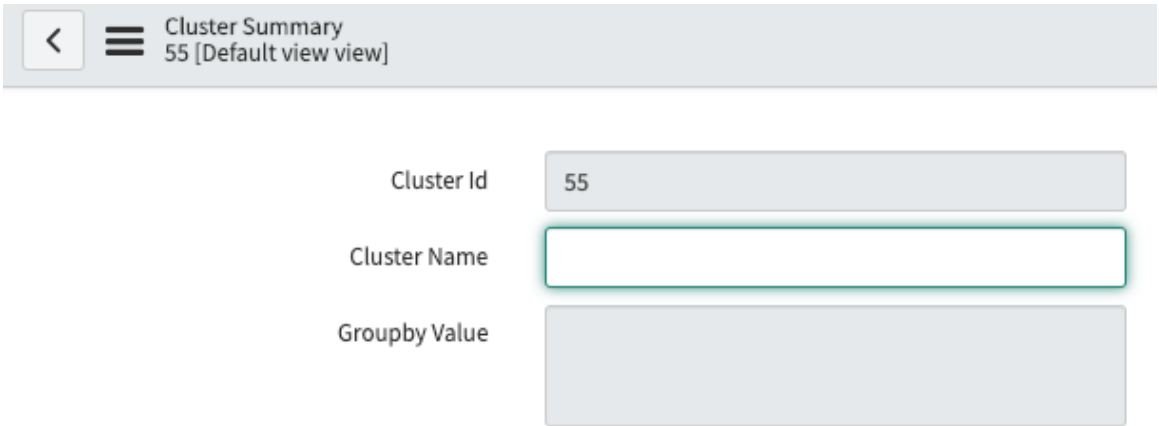
Name your clusters to help identify and organize them.

Before you begin

- Create a clustering solution definition or use an existing one.
- Role required: admin or ml_admin

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solutions**.
2. Select a solution.
3. Under the **Cluster Summary** related list, select a cluster ID.
4. Enter a name in the **Cluster Name** field.



5. Select **Update**.

Generate a representative sample of a cluster

View the top 25 most representative records of a cluster.

Before you begin

- Create a clustering solution definition or use an existing one.
- Role required: admin or ml_admin

About this task

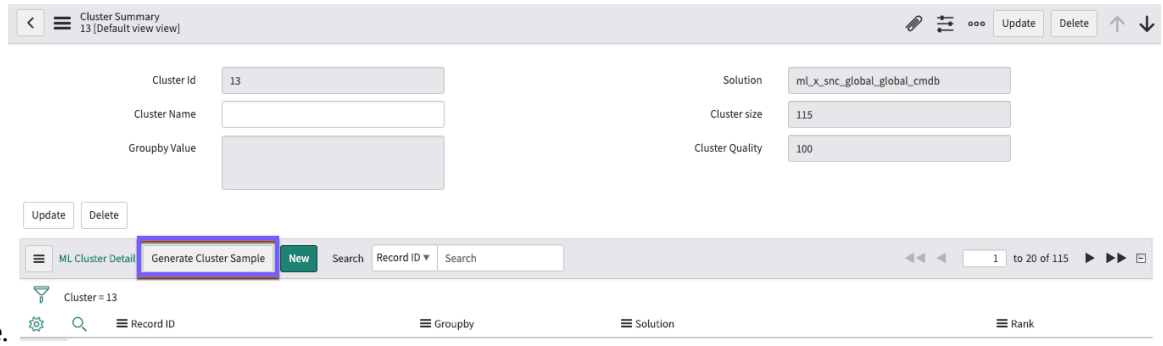
Generating a sample of a cluster filters the records for that cluster. Generating a sample works on clusters with over 25 records.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solutions**.
2. Select a clustering solution.
3. Under the **Cluster Summary** related list, select a cluster ID.
4. Select **Generate Cluster**

Sample.



Result

The **ML Cluster Detail** list shows the top 25 records in the cluster. The cluster sample also applies to the cluster visualization.

Apply purity on a clustering solution

Apply purity to learn details about the composition of each cluster. For example, see what percentage of incidents in a cluster have the same value for Assignment Group. You can specify which fields to focus on, or you can let auto-purity display fields by default.

Before you begin

- Create a clustering definition solution or use an existing one.
- Role required: admin or ml_admin

About this task

With purity, you can choose field distribution metrics for your clusters to display. The table assigned to your solution determines which fields are available. For more information about purity, see the following articles on ServiceNow Community:

- https://www.servicenow.com/community/intelligence-ml-articles/using-purity-fields-to-better-understand-your-clusters/ta-p/2302441
- https://www.servicenow.com/community/intelligence-ml-articles/predictive-intelligence-using-the-cluster-insight-table-to/ta-p/2301006

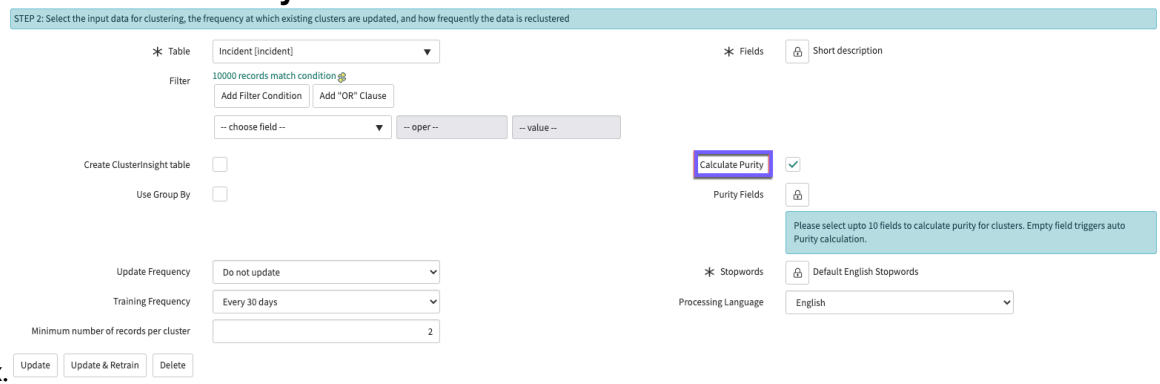
If you do not select any fields, auto-purity automatically determines which insights to show based on distribution significance. You can change the default auto-purity selections by editing the ML Auto Purity Configuration [ml_autopurity_config] table.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solution Definitions**.
2. Select a solution definition or create a new one.

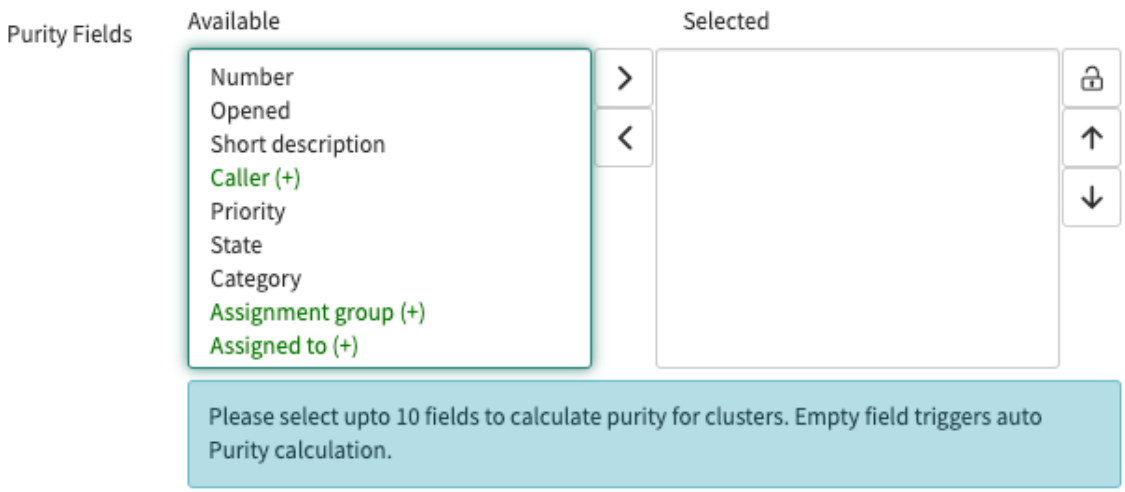
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select the **Calculate Purity** check box.



4. **Optional:** Select the **Purity Fields** lock icon and choose the fields that you want displayed in cluster results.

> **Note:** If you don't select any fields, auto-purity displays the most significant fields based on distribution.



5. Select **Update** or **Update & Retrain**.

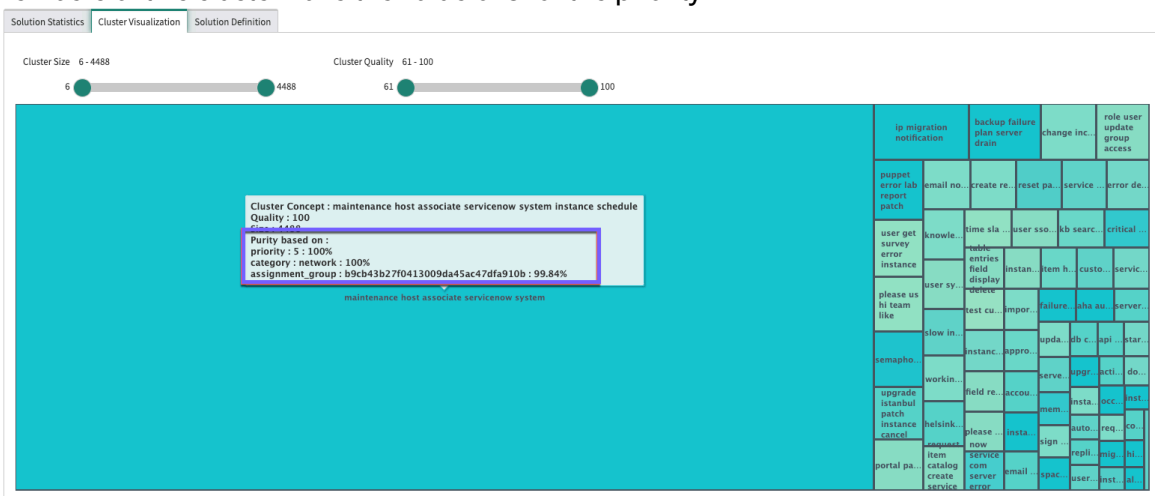
Result

View your insights by opening the **Cluster Visualization** tab of your retrained solution, then pointing to a cluster. Under **Purity based on**, the format displays as:

field name : value : percent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, the row priority : 5 : 100% means that 100% of the members of this cluster have the value of 5 for the priority field.



Analyze a cluster with Cluster Insight

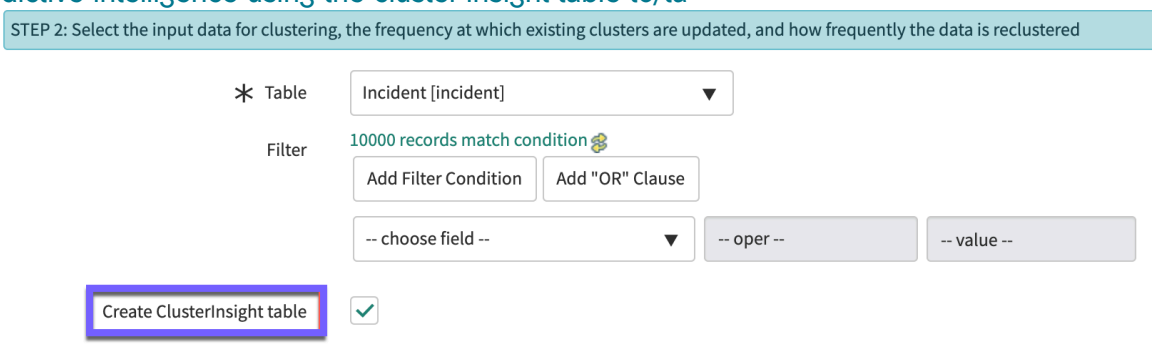
Analyze a cluster by a field available on the source table. With the Cluster Insight check box, you can add a filter condition on your input field when you review the list of results.

Before you begin

- Create a clustering solution definition or use an existing one.
- Role required: admin or ml_admin

About this task

Drill down into a cluster and filter its records with Cluster Insight. You can add cluster insight analysis when creating a solution or editing an existing solution. For more information, see https://www.servicenow.com/community/intelligence-ml-articles/predictive-intelligence-using-the-cluster-insight-table-to/ta-p/2301006.



Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solution Definitions**.
2. Select a solution definition or create a new one.
3. On the definition form, select the **Create Cluster Insight table** check box, then update and retrain the model.

Result

In the **Solutions** [ml_solution] table, open the trained solution, then open the **Cluster Visualization** tab. When you select a cluster, all records included in the cluster are displayed as a list. Using the filter icon ( ), open the list's filter conditions. Search for the input field from your

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

source table. For example, if your source table is Incident and the input field is Short Description, you could filter for records containing the word "help" in the Short Description.

To make other dimensions from your source table available for filtering, add them as input fields on the solution definition form.

Create and train a regression solution

Regression solutions have been deprecated. They could be used to predict numeric outputs, such as a temperature or a stock price.

Before you begin

Note: Support for new regression solutions is deprecated in the Australia release. You can still edit and train any existing solutions, but you can't initiate new ones. The following information is provided for legacy context.

Role required: ml_admin or admin

About this task

Regression solutions enable you to predict a point estimate and prediction interval. The resulting model delivers the following statistics:

- Mean Absolute Error (MAE), which measures the mean deviation of a predicted value from the actual value. This metric is useful as it's easy to understand as its scale is the same as that of its target. However, MAE is unbounded, making it difficult to compare across models.
- Symmetric Mean Absolute Percentage Error (SMAPE) is a percentage value of the deviation from the predicted to the actual. SMAPE is a bounded version of MAE except that it has a value range between 0 and 100. The lower the SMAPE value, the better the model accuracy.
- Range Accuracy is the percentage of actual values between a predicted range. In other words, it's the range between the upper and lower bounds of the prediction. For example, if four out of five actual values lie within the predicted range, the range accuracy is 80%.
- Average Interval Width is the difference between the upper and lower bounds of the prediction. This metric explains how informative the interval is. The smaller the average width, the better the model

When making predictions, regression also enables you to specify a confidence level for the prediction interval (range).

In this example procedure, you create and train a regression solution definition to predict the amount of time it takes to restore a cloud database.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Regression** > **Solution Definitions**.
2. On the Regression Definitions list, select **New**.
3. On the Regression Definition form, configure these fields per the following guidance.

Field	Value
Label	Enter a unique name for your regression solution. For example, in this use case you could enter `Regression Test for DB Restore`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
Name	As you enter your solution Label value, this field automatically populates with a system-assigned name that's similar to your label value.
Word Corpus	Select an existing word corpus that's relevant to your solution. For example, in this use case you select a word corpus that has a title such as **Incidents in the last 3 months**. If you don't have a relevant word corpus, follow the steps to create a word corpus first. When the word corpus is complete, you can select it from the Word Corpus field in your Regression Definition form. However, the word corpus selection is optional. If your input data has text columns and you don't choose a word corpus, your regression solution trains a new word corpus model by using the text columns in your input data. The resulting word corpus can be reused in any other regression solution or other ML solution type. **Note:** A pre-trained model is used instead of the Word Corpus for users who activated Predictive Intelligence starting in the Utah release.
Table	Select the database table on which you're applying regression. The table should contain historical records the system can use to predict the length of time for its database restore.
Output Field	Select the field whose value you want the predictive model to set. In general, a good output field is a numeric, integer, or floating point field. In this example scenario, you use the **Duration** field to measure a length of time. The output field should generate a numeric value.
Fields	Select one or more field types that help the system identify the records you want to train using regression. In this example scenario, you use **Short description**, **Source datacenter**, **Target datacenter**, and **Database size**. (short_description, Sourcedc, Targetdc, and Dbsize.) Input field types can be string, nominal, or numeric.
Filter	(Optional) Add filter conditions to the output field records that you want to train using regression.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	Note: ○ The minimum number of records for regression training is 10,000 records. ○ The maximum number of records for regression training is limited to 300,000.
Processing Language	Select the primary language of the dataset that you're training on the solution definition. If the dataset language is Italian, choose **Italian**. Also, English processing is applied to all datasets by default. For example, if you select Italian, the system processes the data in both English and Italian. **Note:** The term processing indicates some of the language-specific steps used as part of training a solution. These steps include tokenizing words, removing stop words, and stemming.
Stopwords	When you select your processing language, the system automatically adds a Stopwords list that uses the same language. For example, if your processing language is Italian, the **Default Italian Stopwords** list is displayed. The **Default English Stopwords** list also displays in your selection. If you create a custom stopwords list, you can select it from the Stopwords field to add it to your solution. In this scenario, you use the **Default English Stopwords** list.
Training Frequency	Select how often the system regenerates the solution based on records matching the **Filter**. Your options include: ○ Run Once ○ Every 30 days ○ Every 60 days ○ Every 90 days ○ Every 120 days ○ Every 180 days In this scenario, you select Every 30 days By default, the system runs training once. This practice provides you time to review and update the solution definition as needed until it provides acceptable coverage and precision values.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	Note: ○ The minimum number of records required for regression solution training is set at 10,000. ○ The ML scheduler limits the number of trainings an instance can commit to 50 new ML training requests per instance within a 24-hour window. This limit excludes scheduled requests for retraining. In addition, clustering and similarity updates are also excluded from this limit, even if the new training requests exceed 50 within a 24-hour window.

4. Select the appropriate context menu option or button for your solution definition.

Option	Description
Save or Save & Train	Save your solution definition record so you can return to it later, or save and submit it for training.
Submit or Submit & Train	Create your solution definition record and submit it, or submit and train it.

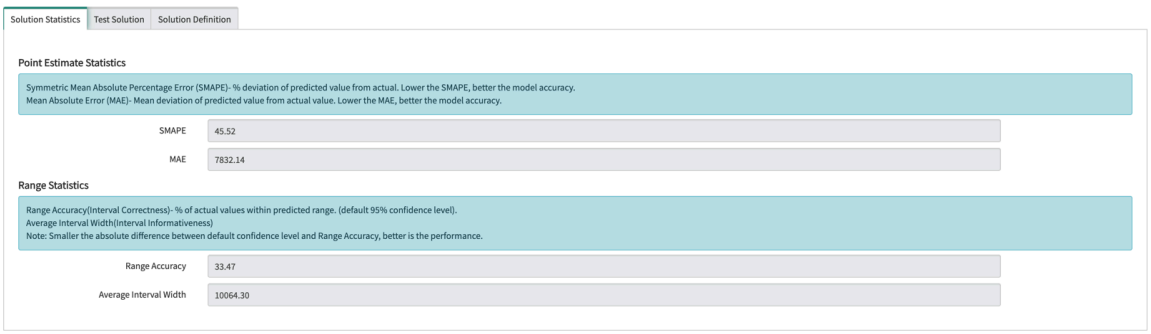
5. If you submitted the solution for training, select **OK** on the **Training Activation** window to confirm.

   The system schedules the solution for training with the nearest training service. The system sends you a notification when the training completes, including any errors that may have occurred in the training. Any other users can subscribe to the Predictive Intelligence Notifications category. When training completes, the system uploads the solution as an Attachment record.

What to do next

In this example scenario, you created an ML solution from your solution definition. The Solution Statistics, Test Solution, and Solution Definition tabs appear in the Related Links section of your ML solution.

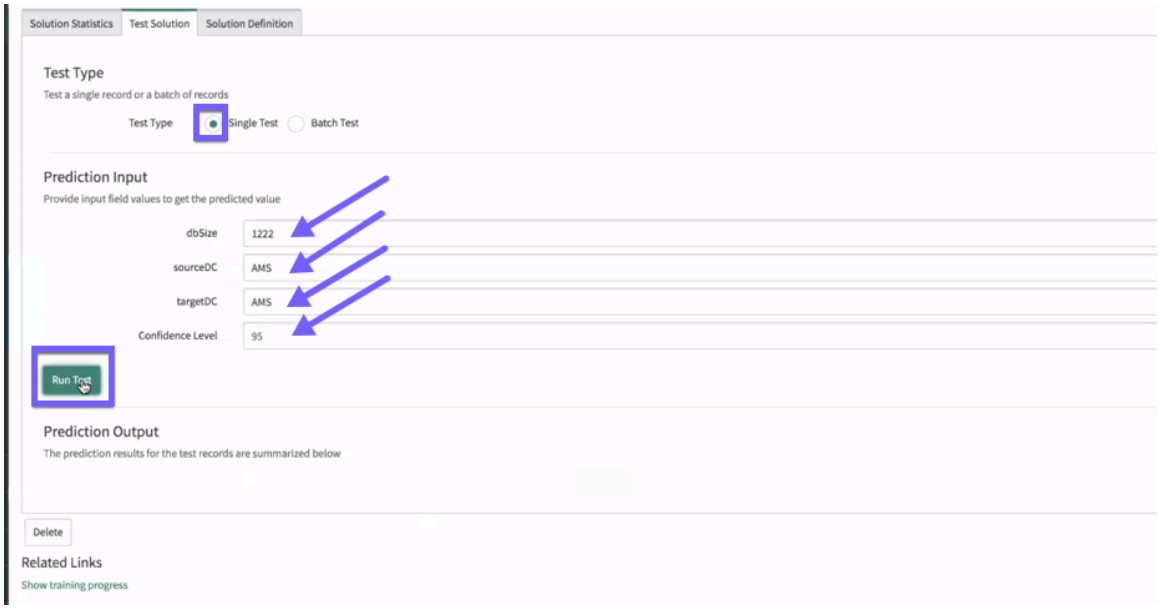
On the Solution Statistics tab, review the Point Estimate and Range (prediction interval) statistics generated by your solution.



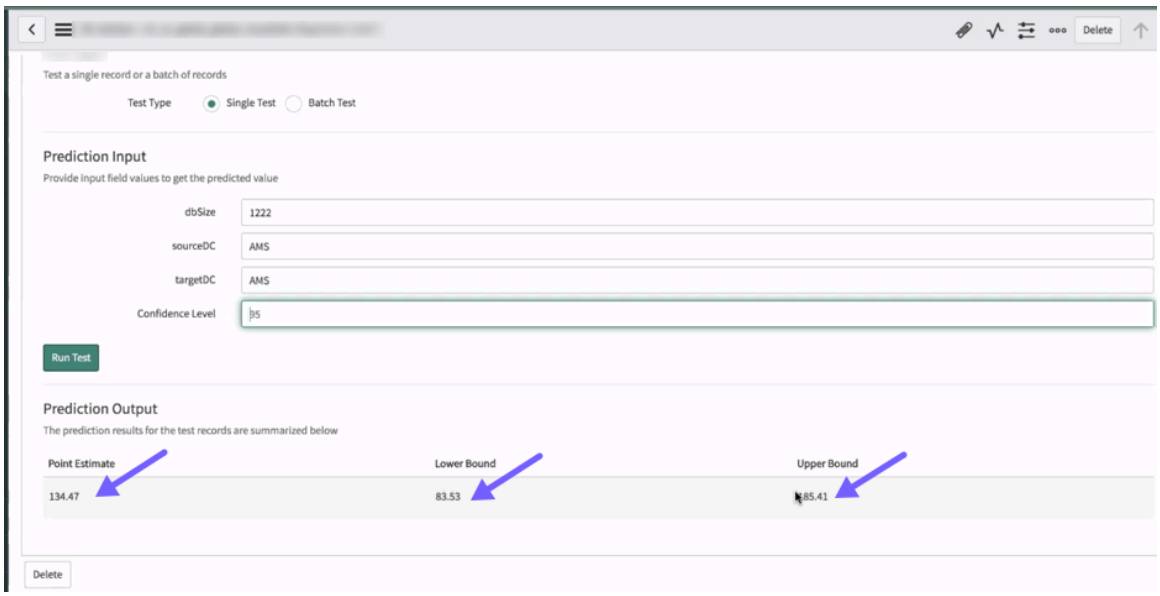
On the Test Solutions tab of your solution, you can test the prediction output for the records you used as input to the prediction by entering values for the input fields, such as the **Source**

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

datacenter, **Target datacenter**, and **Database size**. You can also use the default prediction confidence level of 95, or enter a different level between 0 and 100. Using 95 as the value means that the system is 95% confident that the actual prediction falls within the prediction interval. Select the **Run Test** button to find the prediction output.



After you run the test, the prediction output statistics appear. The Point Estimate on the screen is a single value at one point in time. For example, the database restore takes 134.47 seconds to complete. The Lower and Upper bounds on the screen signify a range accuracy value. For example, the database restore takes from 84.53 to 185.41 seconds to complete.



View solution training progress

View your solution training progress or statistics to determine if a solution is available, or how long the next training cycle might take to complete.

Before you begin

Role required: admin or ml_admin

About this task

Solution training involves these steps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

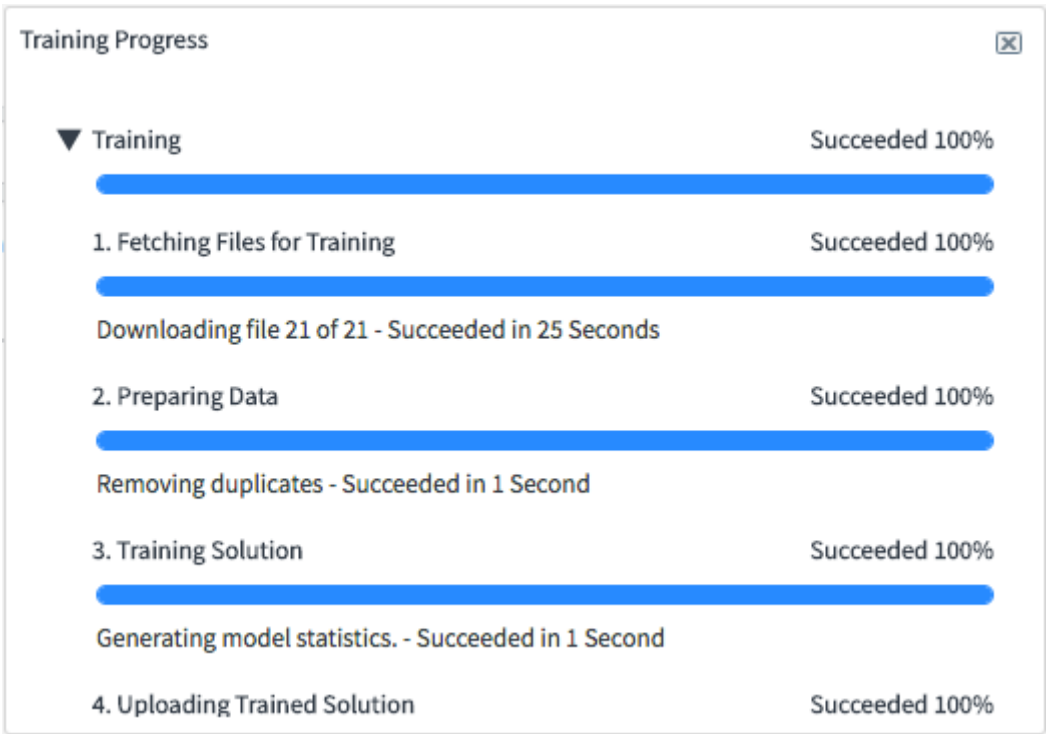
1. Fetching files for training. The system downloads the training records and sends them to the nearest training service.
2. Preparing the data. The system removes duplicate records from the training set.
3. Training the solution. The training service trains the solution.
4. Uploading the trained solution. The training service uploads the solution as attachment records.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Classification** > **Solutions** or **Predictive Intelligence** > **Similarity** > **Solutions**.
2. In the ML Solutions list, select the solution whose progress or statistics you want to view.

   Example
   For example, select **Incident Categorization** to see the training history.
3. In the **Related Links** section, click **Show training progress**.
   Training times vary based on the number of records and classes within the training set. The more records and classes you use, the longer the training can take. For example, a data set containing 100,000 records and several hundred classes can take around five hours to complete.

   The system shows a Training Progress pop-up window.



What to do next

For classification solutions, see Review classification solution statistics.

For similarity solutions, see Review similarity solution examples and scores.

Review classification solution statistics

The Solution Statistics dashboard in Predictive Intelligence has been deprecated in the Xanadu release. It provided precision and coverage statistics for each class in a classification solution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Role required: admin, ml_admin, or ml_report_user

About this task

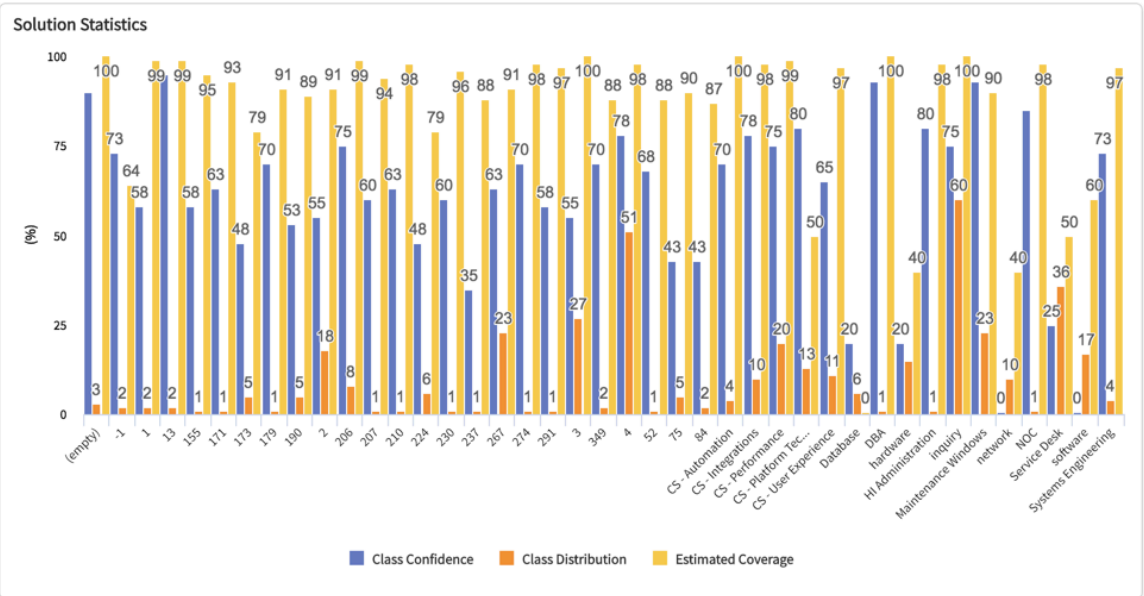
Important: With the Australia release, the Solution Statistics dashboard is deprecated. Upgrading customers can continue to use their existing Solutions Statistics dashboards from the application menu. For new customers onboarding with the Australia release, the Solutions Statistics dashboard is not available. The following information is provided for legacy context.

The Solution Statistics dashboard lists the precision, coverage, and distribution for each class of active solutions. The system uses the classes with the highest number of records when it builds a solution. The number of classes predicted may be less than 50, and may skip a class if there is not enough historical data to build a solution that can predict the class confidently.

The Solution Statistics dashboard is different from the Solution Statistics tab in an ML Solution record. For more information, see Create and train a classification solution.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Classification** > **Solution Statistics**.
2. From **Filter by solution**, select the solution whose statistics you want to review.
3. From **Filter by version**, select the solution version whose statistics you want to review.
4. Click **Apply.**
   The system updates the dashboard based on the filters selected.



5. Identify classes with unwanted combinations of precision, coverage, and distribution values.

   Example
   For example, identify classes that have low precision or coverage but a high distribution.

6. Identify any missing classes you want the model to include.

   Example
   For example, identify any missing incident categories from the Incident classification solution.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

If you're satisfied with the solution you've reviewed, the latest version is already active and ready to use. If you're not satisfied, you can choose a different version of the solution and make it active. You can also tune and retune the solution by configuring the class precision and coverage to use acceptable values.

Review solution similarity examples

Review the similarity examples generated during solution training to determine whether the similarity score threshold meets your business requirements.

Before you begin

- Train a similarity solution
- Role required: ml_admin or admin

About this task

Solution training generates a set of paired record examples. Each pair includes a percentile score that estimates the degree of similarity between the two records. The higher the score, the higher the similarity, as follows:

- A score of 100 indicates identical records.
- A score of 0 indicates dissimilar records.

A deployed solution returns records only when the score is higher than the **Similarity Score Threshold** value:

- The higher the threshold value, the higher the precision and the lower the coverage.
- The lower the threshold value, the lower the precision and the higher the coverage.

On domain-separated instances, the following procedure displays records from all domains to ML admins. However, when solutions are applied in your forms and flows, similarity predictions are domain-aware. Examples from other domains on the instance are not displayed to users.

Note: The similarity filters specified in the solution definition aren't applied for similarity examples and are applied only during prediction.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Similarity** > **Solutions**.
2. In the ML Solutions list, locate your solution by filtering for the Solution Name, and open the active version.
3. In the Related Links section of the form, select **Similarity Examples**.
   A list of paired records opens on the Similarity Examples (ml_similarity_example) table. By default the rows are sorted by the **Similarity Score** column in descending order.
4. Review the similarity examples for accuracy and coverage.
   Based on your review, determine whether to increase or decrease the Similarity Score Threshold value of your solution.

What to do next

If you decide to adjust the Similarity Score Threshold for your similarity solution, see update its similarity solution threshold.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Using Predictive Intelligence

Train and use Predictive Intelligence solutions to accomplish various tasks and that integrate with other ServiceNow products, such as Document Intelligence and Task Intelligence.

Overview of Predictive Intelligence

Predictive Intelligence is the interface by which you can train models on the ServiceNow AI Platform. These models enable you to predict, estimate, and identify patterns that can be used to route work, populate form fields, estimate wait times, and more.

- Show suggestions for relevant articles.
- Assign, categorize, and prioritize tasks.
- Detect major incidents.
- Recommend case resolutions.
- Prevent duplicate articles and ideas.
- Detect phishing attempts.

For more information about the different types of solutions available, see Explore Predictive Intelligence.

Training your ML solutions

Predictive Intelligence enables you to train predictive models and machine-learning solutions that you can apply using data on your instances. The solutions you create use the frameworks to predict, recommend, and organize data. To get started, see Creating and training solutions.

You can also extend Predictive Intelligence to other processes and applications, such as:

- Incident categorization: Predicts the incident category based on the short description. See Predictive Intelligence for Incident Management.
- CSM case assignment: Predicts the case record assignment group based on the short description. See Predictive Intelligence for case management.

For more information, see ServiceNow apps and features that use Predictive Intelligence.

Testing and monitoring predictions

After creating and training your solutions, call on the Predictive Intelligence API to make a solution prediction. Use the results to gauge the performance of the solution and make changes as needed.

You can track the coverage and precision of deployed predictive models using the Solution Statistics dashboard, which provides reporting on these prediction areas by default.

Report	Description
Average Prediction Coverage (last 30 days)	The percentage of predictions that yielded an outcome out of the total number of predictions attempted. Click the coverage score to see a breakdown by class.
Daily Prediction Coverage	The percentage of records created on a given day in which the solution was able to predict an outcome.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Report	Description
Average Prediction Precision (last 30 days)	The percentage of predictions in which the predicted value was the same as the final value of the field when the record closed. Click the precision score to see a breakdown by class.
Daily Prediction Precision	The percentage of records closed on a given day in which the predicted field value was the same as the final value.

For more information, see Testing and monitoring predictions.

Preparing your instance

For you to get the most out of Predictive Intelligence, you will want to prepare. You don't need to write code or do calculations, but deciding what you hope to do with the solution definitions will make implementation easier.

- Identify the problems that you want to solve with Predictive Intelligence.
- Have 30,000–300,000 high-quality records from which Predictive Intelligence can learn.
- Set your expectations.

Note: Inconsistencies or gaps in training data can cause incorrect or unreliable predictions.

Implementation Process

Predictive Intelligence takes approximately 14 days to implement on a production instance.

- Day 1: Clone production instance over to a non-production instance.
- Days 2–10: Create a solution definition, train it on historical records, and validate that the solution works as desired on the non-production instance.
- Days 11–13: Create import and update sets to move the solution to production, train and validate on the new instance, and set the retraining frequency.
- Day 14 and on: Monitor the solution.

In general, non-production environments are where workflows can be tested and formatted before being moved over to the production instance to further train models and test predictions.

For more information about getting started with Predictive Intelligence, see our guide on how to get started with Predictive Intelligence.

Using Machine Learning APIs

Use ServiceNow Machine Learning (ML) APIs to train Machine Learning models and run inferences.

ML APIs enable training solutions and managing solution versions. You can get and set active versions, monitor training status, and more. The ML API also provides encoders, which enable using term frequency–inverse document frequency (TF-IDF) as a word corpus. Predictability estimates enable assessing the predictive value of table columns.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: Predictive Intelligence APIs run with full privileges before the Vancouver Patch 7 Hotfix 2b and Washington DC Patch 7 releases. With later releases, grant access using ACLs. For more information see Query ACLs.

ML API class overview

This section briefly describes classes for training ML solutions and running inferences with trained solutions.

Datasets

A dataset is a set of records including a table name, columns, and row selection criteria to use as input for ML training algorithms. Datasets don't contain the actual data.

For more information, see DatasetDefinition.

ML objects – Solutions, Encoders, and Estimates

ML objects define a specific training configuration to apply on a dataset. Some operations are common across ML objects. Solution objects include classification, clustering, regression, and similarity.

Encoders are text processing objects that are either pre-trained or trained based on the language datasets you provide. You can train encoders that determine how the system interprets and processes text fields. For ML solutions that include text, you can train an encoder to specify how to process text and use the trained encoder in a solution.

PredictabilityEstimate objects estimate which fields in a dataset are predictable and the features on which this predictability is based.

For more information, see:

- ClassificationSolution
- ClusteringSolution
- Encoder
- PredictabilityEstimate
- RegressionSolution
- SimilaritySolution

Stores

ML objects are maintained in a specific store for each object type. Each store class includes methods for add, get, update, and delete operations.

For more information, see:

- ClassificationSolutionStore
- ClusteringSolutionStore
- EncoderStore
- PredictabilityEstimateStore
- RegressionSolutionStore
- SimilaritySolutionStore

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Versions

Each trained object results in a new version that you can run tasks on. Use the version API to get any solution version and run tasks on it.

For more information, see:

- ClassificationSolutionVersion
- ClusteringSolutionVersion
- EncoderVersion
- PredictabilityEstimateVersion
- RegressionSolutionVersion
- SimilaritySolutionVersion

Putting it together: ML API flows

You can use the following flow to configure and train solutions, encoders, and predictability estimates:

1. Define a dataset (DatasetDefinition)
2. Create an ML object (Solution/Encoder/PredictabilityEstimate)
3. Add to store (Store)
4. Train (Solution/Encoder/PredictabilityEstimate)

Note: The encoder definitions support multiple dataset definitions, but have the same training flow.

To train a solution with an encoder, create the encoder first, then include the encoder in the solution configuration.

1. Create an encoder (Encoder)
2. Define a dataset (DatasetDefinition)
3. Create solution specifying encoder (Solution)
4. Add to store (SolutionStore)
5. Train (Solution)

ML object encoder requirements:

- Required in similarity API solutions.
- Required in clustering API solutions, unless using the Levenshtein distance algorithm, in which case encoders are optional.
- Optional for classification and regression solutions.
- Unavailable for predictability estimates.

Getting started with ML API solution training

Follow this example procedure to learn how to configure and train a solution.

Configure and train a solution

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Define a dataset using the *DatasetDefinition* API.

```
var myData = new sn_ml.DatasetDefinition({

  'tableName' : 'incident',
  'fieldNames' : ['assignment_group',
 'short_description', 'description'],
  'encodedQuery' : 'activeANYTHING'

});
```

2. Use the constructor to define the solution, including the dataset in the configuration.

```
var mySolution = new sn_ml.ClassificationSolution({

  'label': "my solution definition",
  'dataset' : myData,
  'predictedFieldName' : 'assignment_group',
  'inputFieldNames':['short_description']

});
```

   - *ClassificationSolution()*
   - *ClusteringSolution()*
   - *Encoder()*
   - *PredictabilityEstimate()*
   - *RegressionSolution()*
   - *SimilaritySolution()*

3. Add the solution definition to the store using the *add()* method.

```
var my_unique_name =
 sn_ml.ClassificationSolutionStore.add(mySolution);
```

   - *ClassificationSolutionStore - add()*
   - *ClusteringSolutionStore – add()*
   - *EncoderStore - add()*
   - *PredictabilityEstimateStore - add()*
   - *RegressionSolutionStore - add()*
   - *SimilaritySolutionStore - add()*

4. Train the solution using the *submitTrainingJob()* method. After training is complete, you can manage the trained solution using a solution version API. A solution can be retrained multiple times. Each training results in a new solution "version" on which you can run inferences.

```
var myClassifierVersion =
 mySolution.submitTrainingJob();
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- *ClassificationSolution - submitTrainingJob()*
- *ClusteringSolutionVersion - submitTrainingJob()*
- *Encoder - submitTrainingJob()*
- *PredictabilityEstimate - submitTrainingJob()*
- *RegressionSolution - submitTrainingJob()*
- *SimilaritySolution - submitTrainingJob()*

View all classification solutions in a store

You can use the store *getAllNames()* method to see a list of all solutions that have been added to the store.

```
gs.print(JSON.stringify(JSON.parse(sn_ml.Classification
SolutionStore.getAllNames()), null, 2));
```

In the output, the system has named the solution `ml_x_snc_global_global_my_solution_definition`. Use this name in subsequent examples to get version information.

```
*** Script: [
  "ml_incident_assignment",
  "ml_x_snc_global_global_my_solution_definition",
  "ml_incident_categorization"
]
```

- *ClassificationSolutionStore - getAllNames()*
- *ClusteringSolutionStore - getAllNames()*
- *EncoderStore - getAllNames()*
- *PredictabilityEstimateStore - getAllNames()*
- *RegressionSolutionStore - getAllNames()*
- *SimilaritySolutionStore - getAllNames()*

Getting started with ML API solution versions

Follow these example breakdowns to learn how to manage trained solution versions.

Check training status

Get the classification solution from the store, choose a version, and check its training status. The methods used for checking training status are applicable to all ML object types.

1. Get the solution from the classification solution store using the *get()* method.

```
// Get the solution created in the previous example
 from the classification solution store
var mlSolution =
 sn_ml.ClassificationSolutionStore.get('ml_x_snc_globa
l_global_my_solution_definition');
```

- *ClassificationSolutionStore - get()*
- *ClusteringSolutionStore - get()*

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- *EncoderStore - get()*
- *PredictabilityEstimateStore - get()*
- *RegressionSolutionStore - get()*
- *SimilaritySolutionStore - get()*

2. Access the most recent solution version using the *getLatestVersion()* solution method and get its training status using the *getStatus()* version method.

```
// Access the latest version of the solution and print
 its training status
gs.print(JSON.stringify(JSON.parse(mlSolution.getLates
tVersion().getStatus(), null, 2)));
```

Output when training is complete:

```
*** Script:
 {"state":"solution_complete","percentComplete":"100",
"hasJobEnded":"true"}
```

getLatestVersion()	getStatus()
ClassificationSolution - getLatestVersion()	*ClassificationSolutionVersion - getStatus()*
ClusteringSolution – getLatestVersion()	*ClusteringSolutionVersion – getStatus()*
Encoder - getLatestVersion()	*EncoderVersion - getStatus()*
PredictabilityEstimate - getLatestVersion()	*PredictabilityEstimateVersion - getStatus()*
RegressionSolution - getLatestVersion()	*RegressionSolutionVersion - getStatus()*
SimilaritySolution - getLatestVersion()	*SimilaritySolutionVersion - getStatus()*

Get predictions using a solution version

After the solution has been trained, get the trained version and run a prediction on it. Get the solution you created from the store. Next, choose the trained version and predict the trained version.

Note: Predictions cannot be made on encoders and predictability estimates.

1. Get the solution from the classification solution store using the *get()* method.

```
// Get the solution created in the first example from
 the classification solution store
var mlSolution =
 sn_ml.ClassificationSolutionStore.get('ml_x_snc_globa
l_global_my_solution_definition');
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Use the GlideRecord API *get()* method to provide a record from the Incident [incident] table.

```
// single GlideRecord input
var input = new GlideRecord("incident");
input.get("<sys_id>");
```

3. Optional. Configure the *ClassificationSolutionVersion - predict()* method *options* parameter to return the top three results and return all results.

```
// configure optional parameters
var options = {};
options.top_n = 3;
options.apply_threshold = false;
```

4. Declare a variable called `results` and assign it to the prediction job.
To run the prediction job, get the most recent solution version using the *ClassificationSolution - getLatestVersion()* method and call the *ClassificationSolutionVersion - predict()* method on it.

```
var results =
 mlSolution.getLatestVersion().predict(input,
 options);
```

   - *ClassificationSolutionVersion - predict()*
   - *ClusteringSolutionVersion - predict()*
   - *RegressionSolutionVersion - predict()*
   - *SimilaritySolutionVersion - predict()*

5. Print the predicted results output.

```
gs.print(JSON.stringify(JSON.parse(results), null,
 2));
```

Predicted results example output:

```
*** Script: {
  "<sys_id>": [
    {
      "confidence": 99,
      "threshold": 24.75,
      "predictedValue": "Email",
      "predictedSysId": ""
    },
    {
      "confidence": 5.88210244009169,
      "threshold": 100,
      "predictedValue": "Email (I/f)",
      "predictedSysId": ""
    },
    {
      "confidence": 2.3461203499840932,
      "threshold": 14.81,
      "predictedValue": "Authentication",
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
      "predictedSysId": ""
    }
  ]
}
```

Reviewing your ML solution training jobs

Use the ML Solutions (ML Training Jobs view) module to monitor the training status and progress for Predictive Intelligence solutions. The module displays training jobs for both user interface and API solutions.

Background and usage

When you submit an ML solution or ML solution definition for training, it goes to a ServiceNow data center for processing to predict and deliver a data outcome. Depending on the solution size, the training can take hours or sometimes days to complete. The ML Training Jobs view helps you stay on top of all in-progress and completed ML solution training jobs in your instance.

To access this view, you use the admin or ml_admin role and the following navigation path: **Predictive Intelligence** > **Training Jobs**.

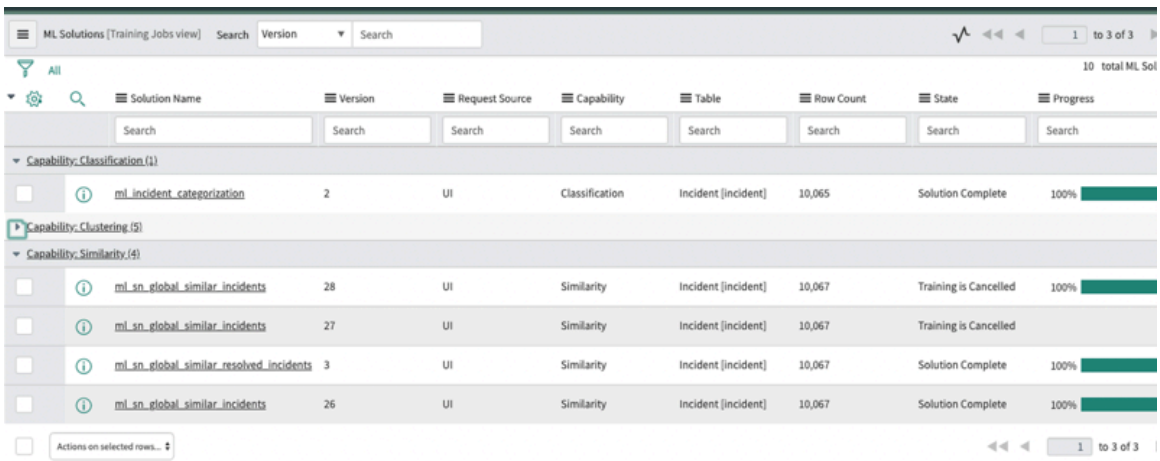
Note: The ML scheduler limits the number of trainings an instance can commit to 50 new ML training requests per instance within a 24 hour window. This excludes scheduled re-training requests. In addition, clustering and similarity updates are also excluded from this limit, even if the new training requests exceed 50 within a 24 hour window.

ML Training Jobs view summary

The view shows all ML training jobs grouped by the four Predictive Intelligence capability frameworks: classification, similarity, clustering, and regression.

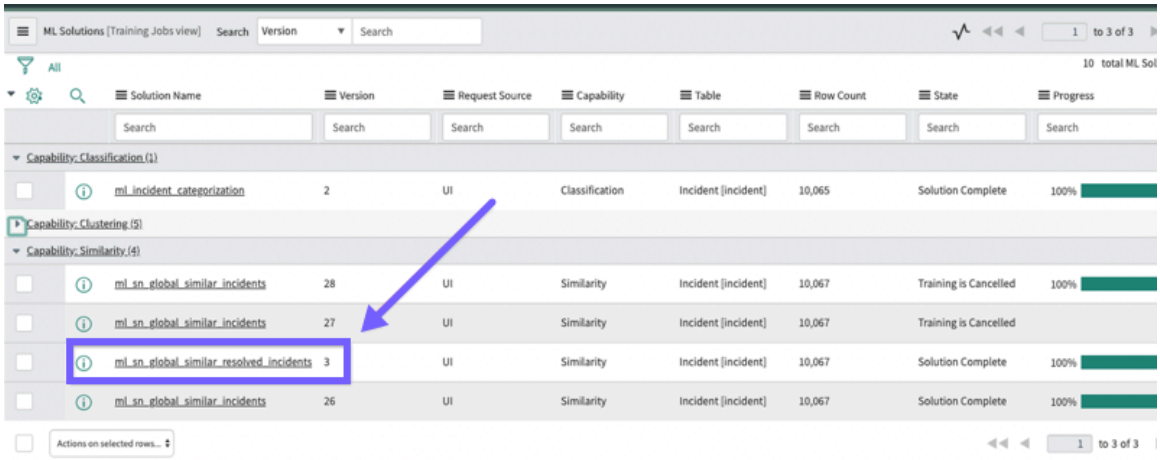
Each record displays values such as solution name, version, training state, and training completion percentage.

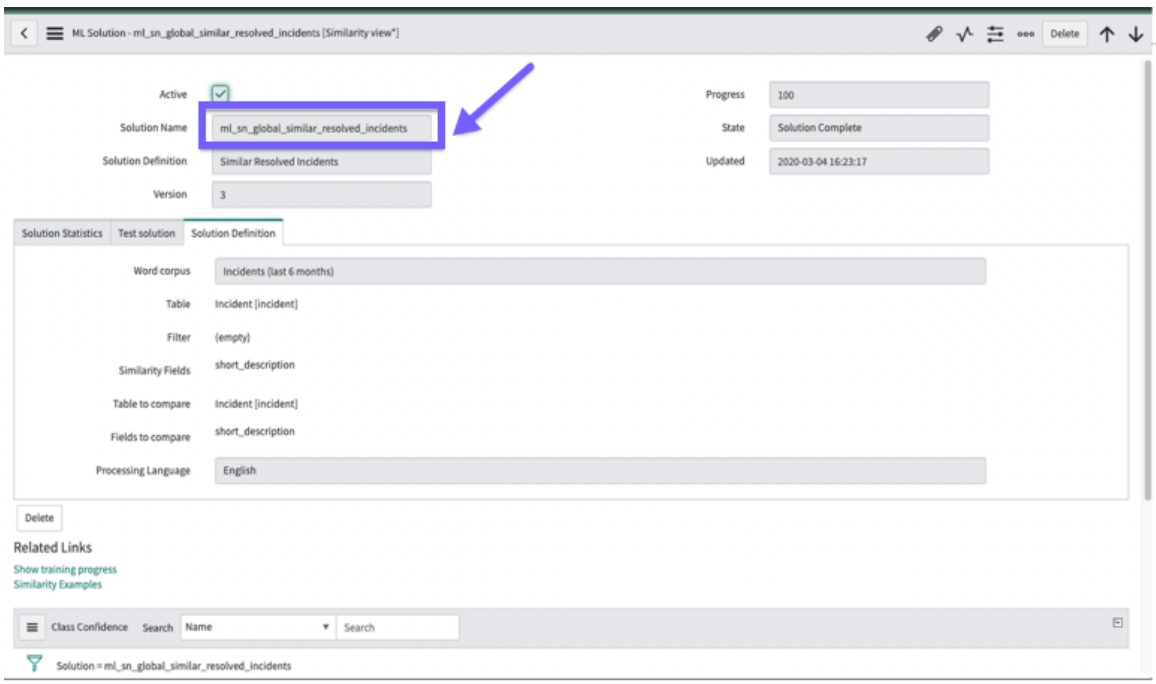
If you don't have any training jobs for a particular capability, the list doesn't display a group for that capability. For example, in this scenario, there is no group for regression because you don't have any regression solutions that you've submitted for training yet.



Select the Solution Name to see the details of the ML solution, as demonstrated in the images below.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.





Use Predictive Intelligence in Workflow Studio with ML actions

Use Predictive Intelligence actions in Workflow Studio to create flows that incorporate your model predictions.

Before you begin

- Make sure the following plugins are activated in your instance: Predictive Intelligence (com.glide.platform_ml) and Predictive Intelligence for Workflow Studio (com.snc.ml_flowdesigner).
- Create or use an existing trained Predictive Intelligence solution.
- Roles required: ml_admin or admin, and flow_designer or delegated_developer.

About this task

Workflow Studio enables you to automate complex processes. The first thing to identify is what process you want to automate. In this example scenario, you're automating the assignment of a Category to an incident record. When you complete the flow, the next incident record created

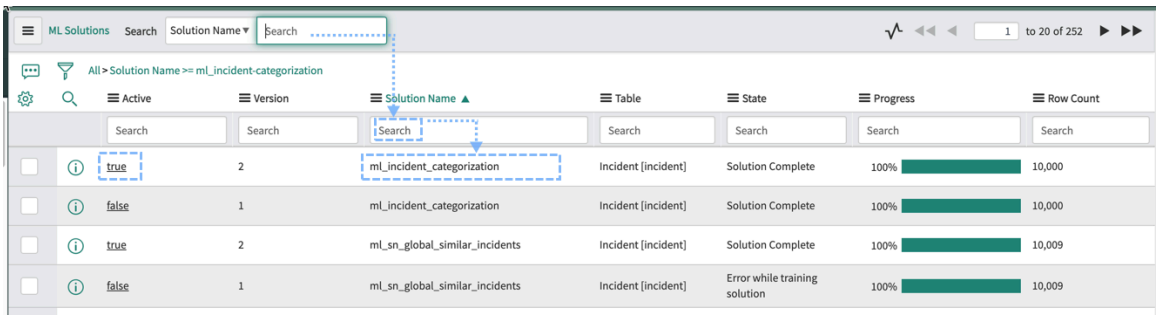
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

in your instance updates the Category field in the record based on the text entered in the **Short description** field.

You can deploy any active and trained classification, similarity, or regression ML solution in your flow, as appropriate for your use case.

Note: The regression framework is deprecated in the Australia release. You can continue to use existing regression solutions but you can't create new ones.

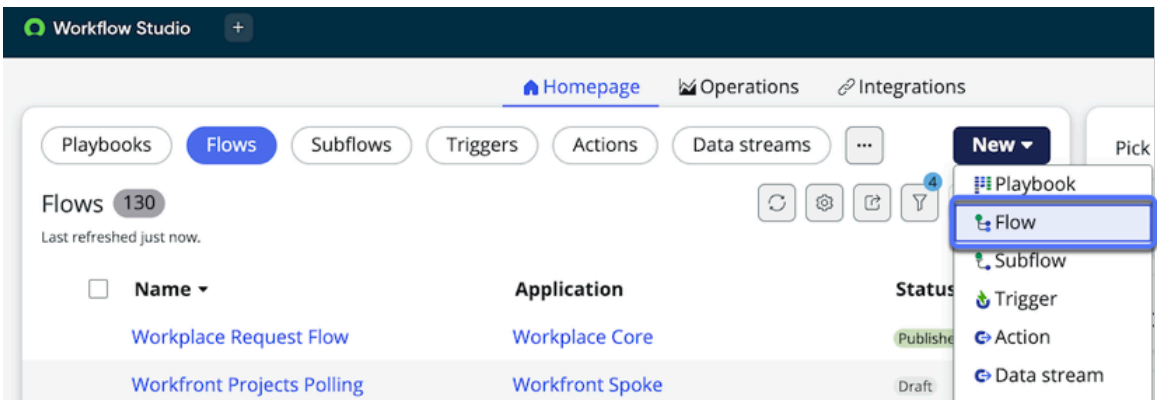
In this example procedure, you create a flow that implements the ml_incident_categorization solution in a Workflow Studio action. You can find this solution by searching on the ML Solutions [ml_solution] table, as shown in the image below. Confirm that the solution you use has been trained and its **Active** value is set to **true**.



For more information on how to use Flow Designer in Workflow Studio, see Exploring flows. For information about the Actions included in Predictive Intelligence for Workflow Studio, see the Spoke actions table in Machine Learning solutions for Flow Designer.

Procedure

1. Navigate to **All** > **Process Automation** > *`Workflow Studio`*.
2. Select **New** > **Flow**.



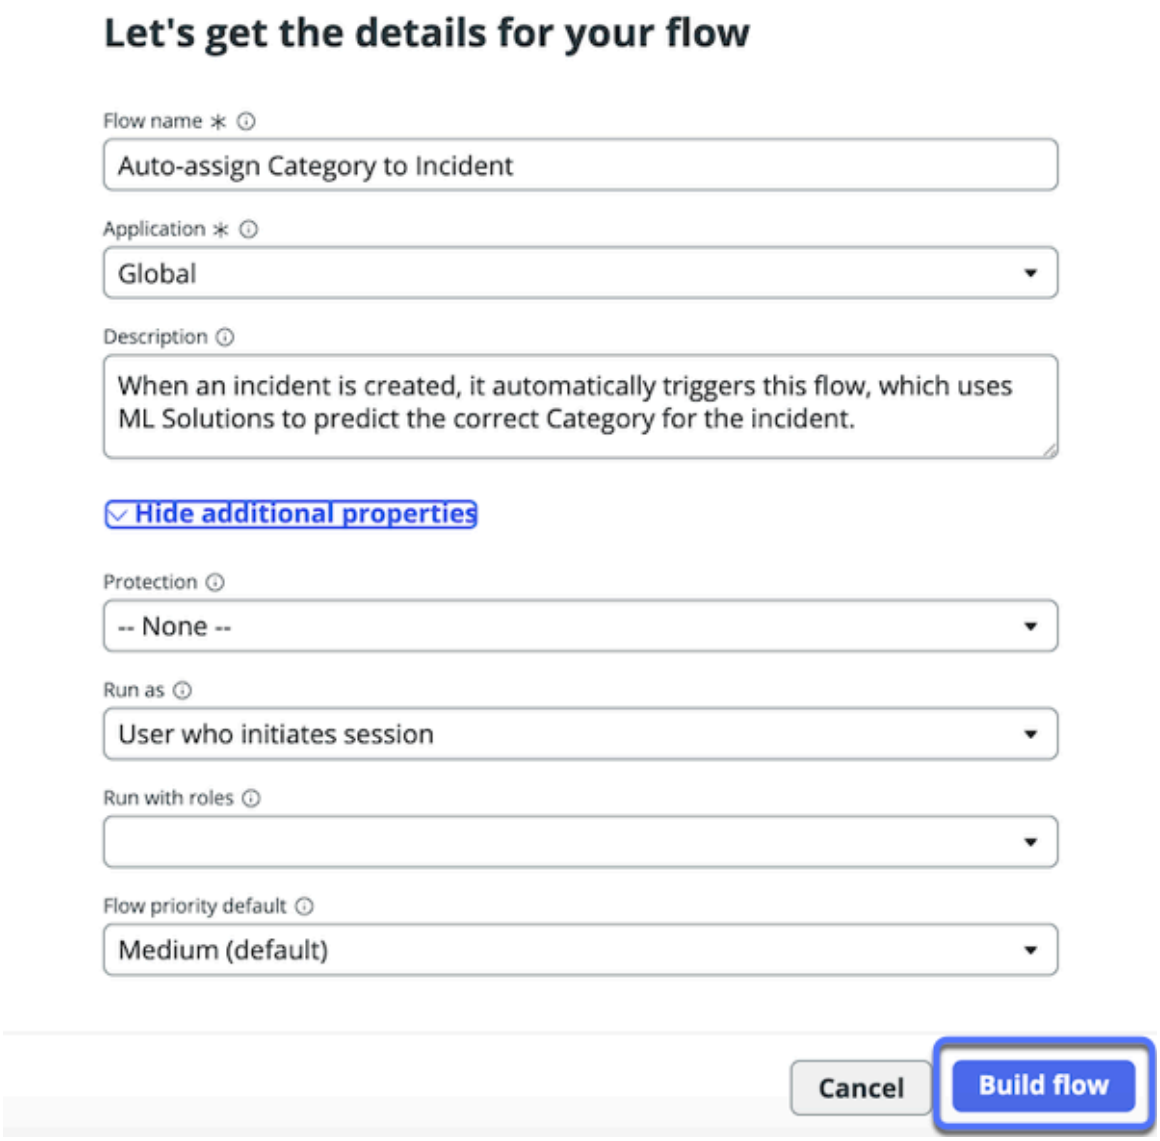
3. On the **Let's get the details for your flow** screen, configure the following fields. Expand **Show additional properties** to view all fields.

Field	Description
Flow Name	Provide a name for the flow. In this scenario, you enter `Auto-assign Category to Incident`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Description	Enter a brief summary description of what the flow delivers. For example, you enter the following: `When an incident is created, it automatically triggers this flow, which uses ML Solutions to predict the correct Category for the incident.`
Application	Select **Global**.
Protection	Select **--None--** or **Read-only**. In this scenario, you select **--None--**.
Run As	Select **User who initiates session**.
Run with role(s)	Leave blank.
Flow priority default	Medium (default).

4. Select **Build flow**.

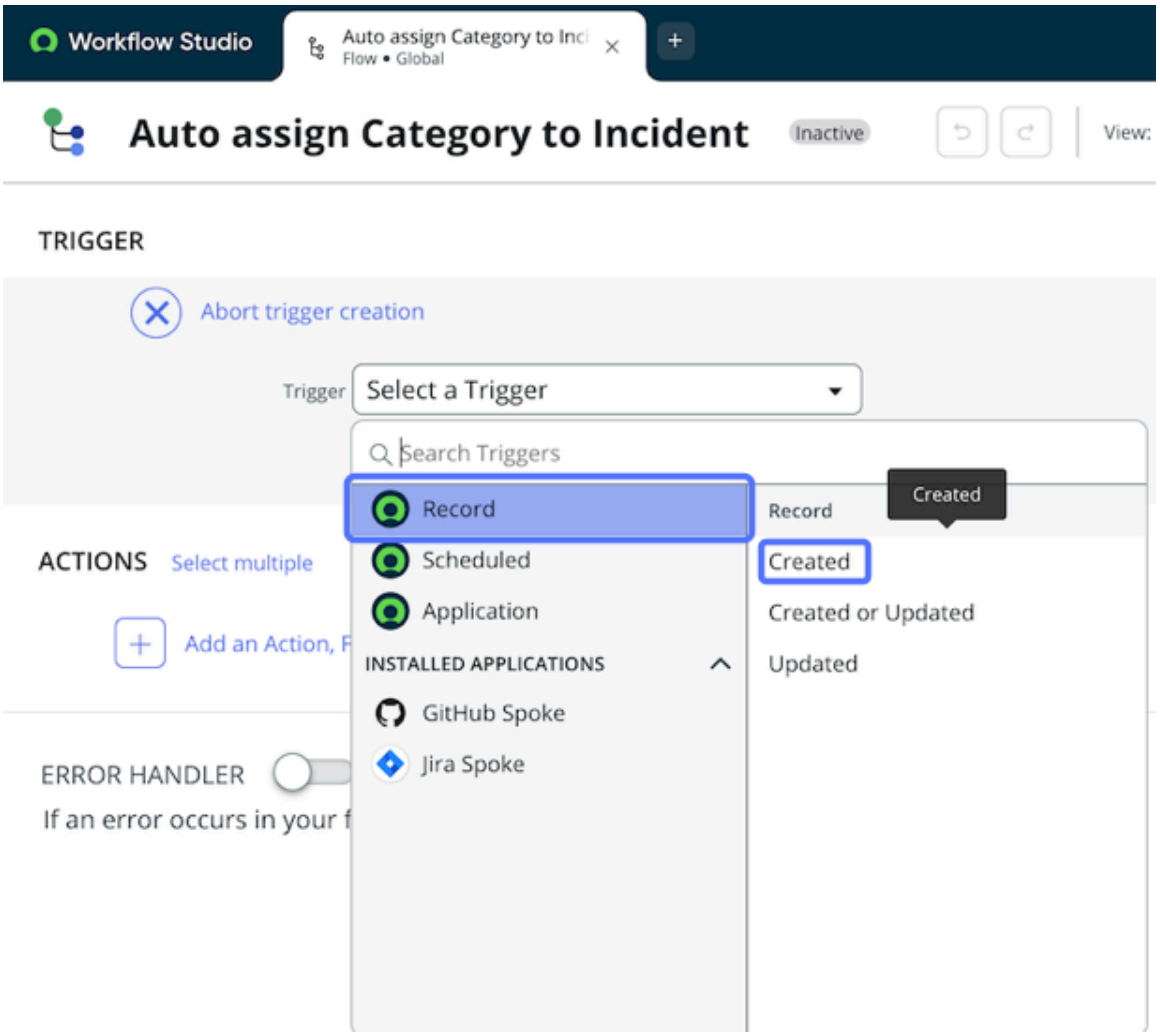


© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Flow screen appears, showing the Auto-assign Category to Incident name you assigned to the flow. If a Getting started screen appears, select **Skip tour**.

5. In the TRIGGER section of the Flow screen, configure the following fields to create a trigger for the flow.

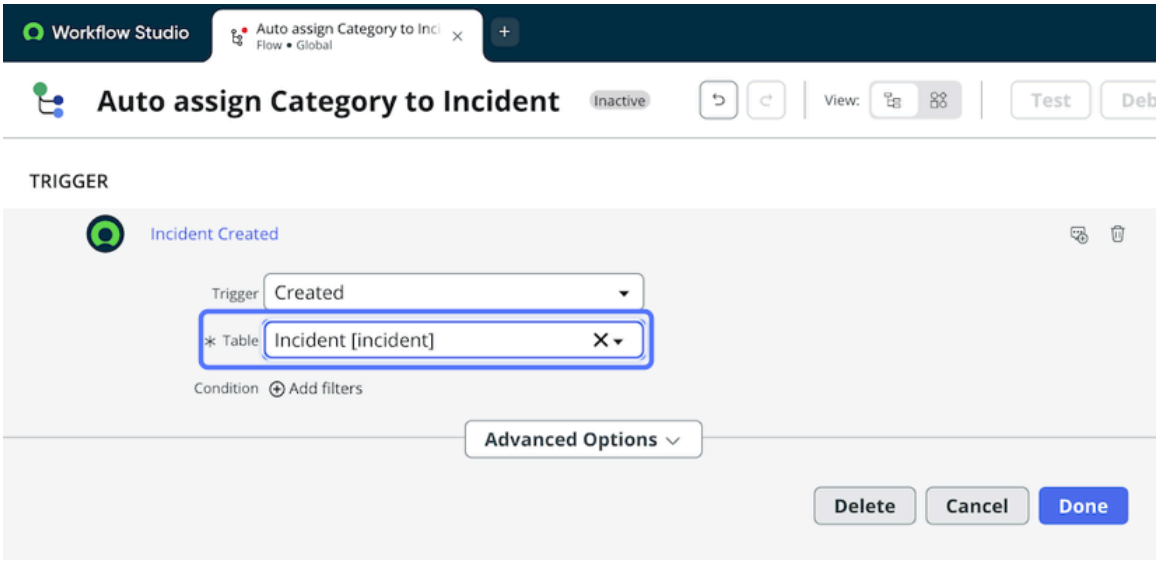
   a. **Select a Trigger**: Select **Record**, then from the list of possible options for Record, select **Created**.



   b. **Table**: Select **Incident [incident]**.

> **Note:** After you configure both the **Trigger** and **Table** fields, record data pills appear in the **Data** section of the screen so you can use them in your flow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



c. **Condition**: Select **Add filters** if you want to add any conditions to the flow.

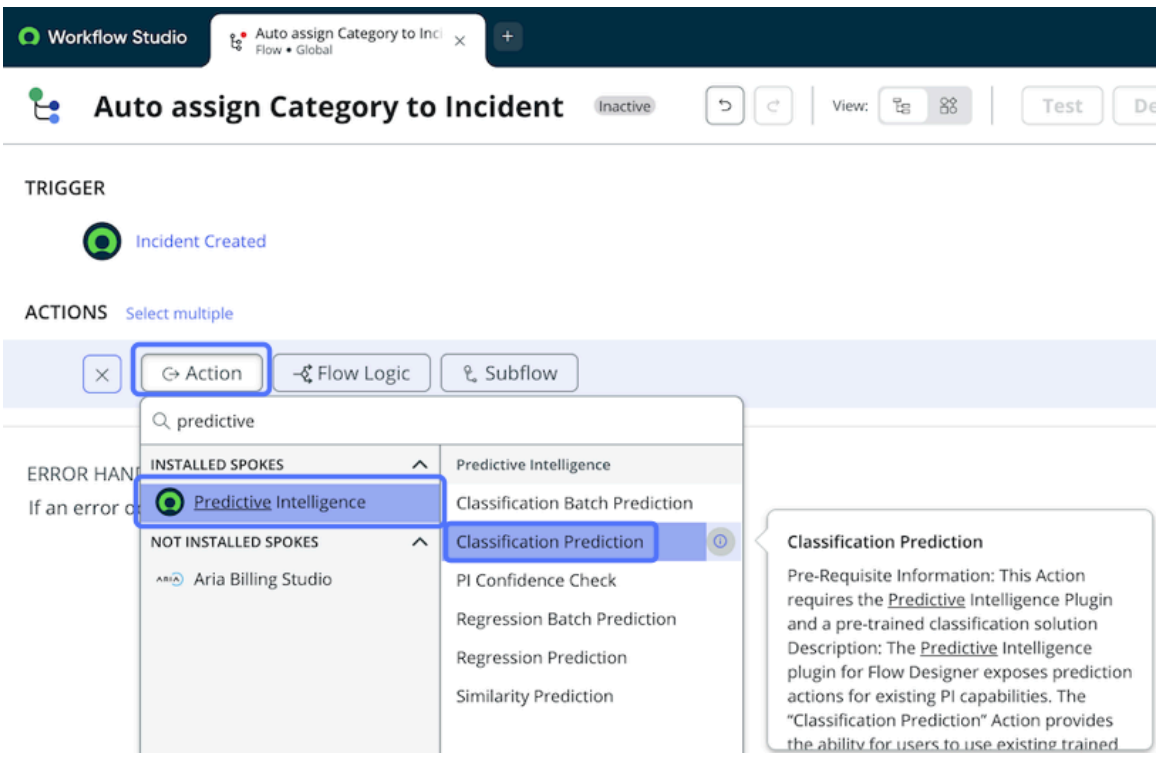
d. **Optional:** Open the **Advanced Options** panel to view additional conditions you can apply to the flow.

e. **Optional:** To close the panel, select **Advanced Options**.

f. Select **Done**.

6. In the ACTIONS section of the screen, configure the following fields to create a Classification Prediction action.

a. **Action** tab: Select **Action** > **Predictive Intelligence** > **Classification Prediction**.



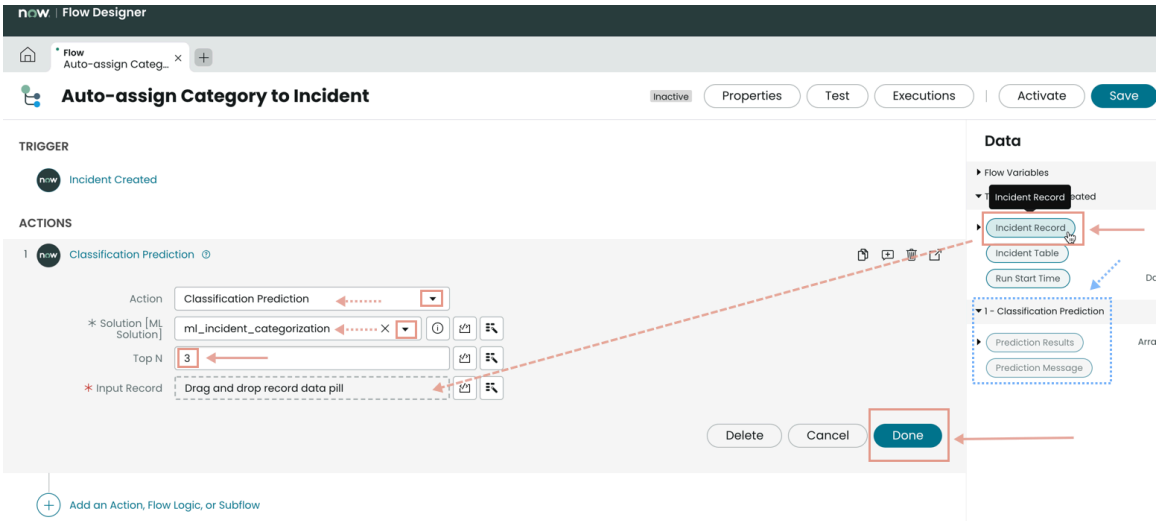
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Select the information icon ( ) to see the description of a Classification Prediction.

b. Solution [ML Solution]: Select **ml_incident_categorization**.

c. Top N: Enter 3 for the example scenario.
When you enter a number, such as 3, the system uses the top three ML predictions that have the highest prediction confidence score. If you don't enter anything, the system sets the default value to 1.

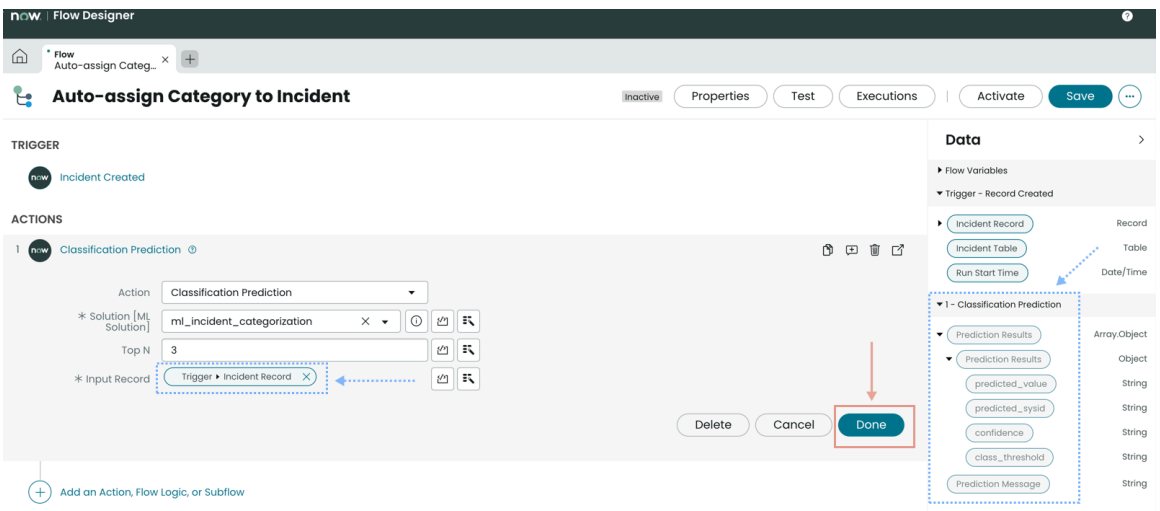
d. Input Record: Drag and drop your **Trigger → Incident Record** data pill into the Input Record field.



The Action, Solution Name, Top N, and Input Record values provide a base for the Category prediction.

Note: The data pill you drop into this record must also be a record. For example, don't try to drop a table pill or a date/time pill into the Input Record field.

e. Select **Done**.
Result: The Classification Prediction action is completed in the flow and its data pills appear in the Data section of the screen.

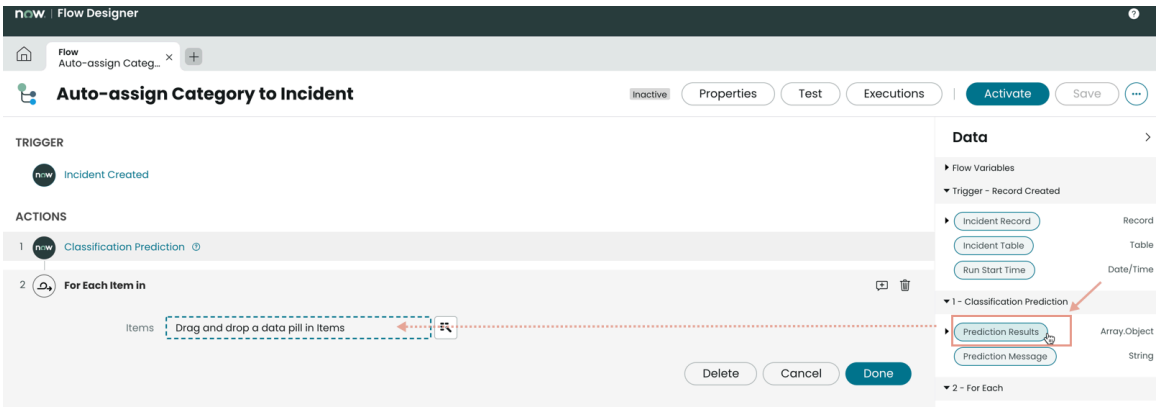


© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. In the ACTIONS section of the screen, use the following steps to create actions and flow logic for the incident's Prediction Results.

   Note: Although you can use a loop to iterate through every prediction result, the scenario shown in this documentation uses a relatively small number of actions. For more advanced flow configurations, see the Flow Designer.

   a. **For each item in list of items**: Drag and drop the **Prediction Results** data pill into the **Items** field.



   Note:

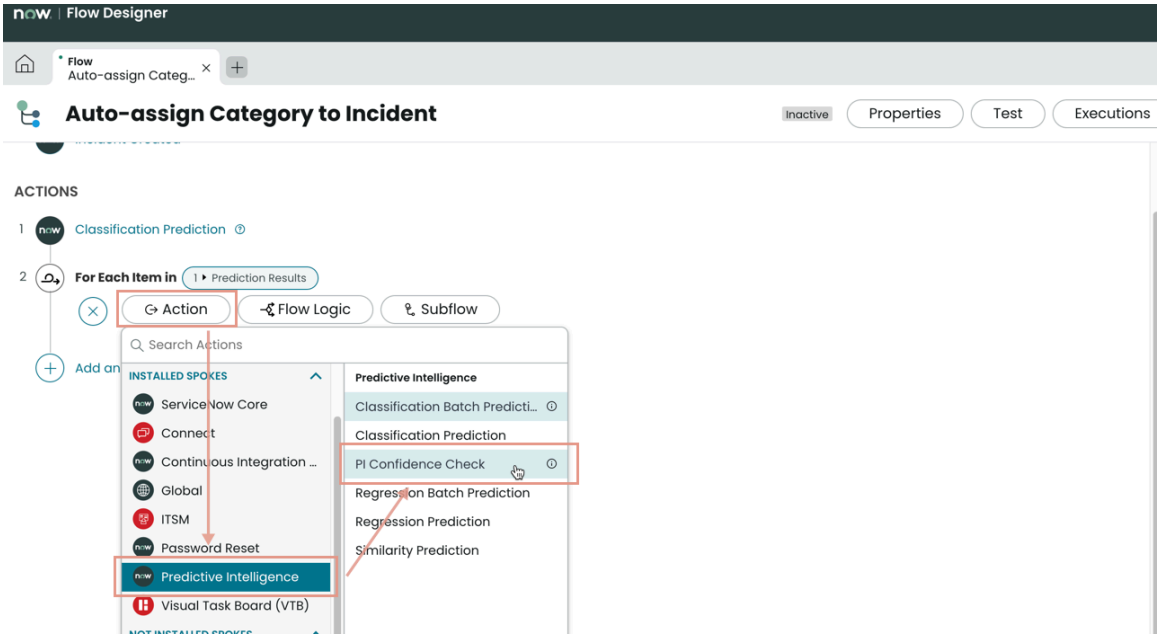
   In order to access the list of items in the Regression Prediction action, you don't need the **For Each Item in** flow logic.

   (The regression framework is deprecated in the Australia release. You can continue to use existing regression solutions but you can't create new ones.)

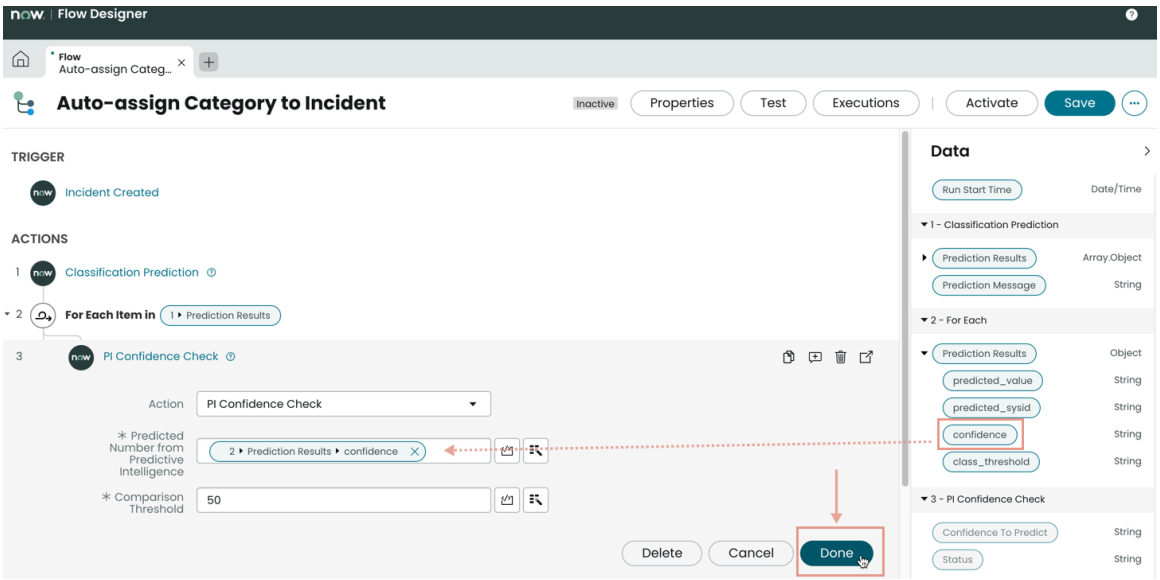
   b. Select **Done**.
   Result: The Prediction Results action is started in the flow and its data pills appear in the Data section of the screen.

8. In the ACTIONS section of the screen, select **Action** > **Predictive Intelligence** > **PI Confidence Check**.
   The PI Confidence Check is a tool you can use to compare values in a flow. In this use case, it compares prediction result values, and the output from the check is either True or False.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

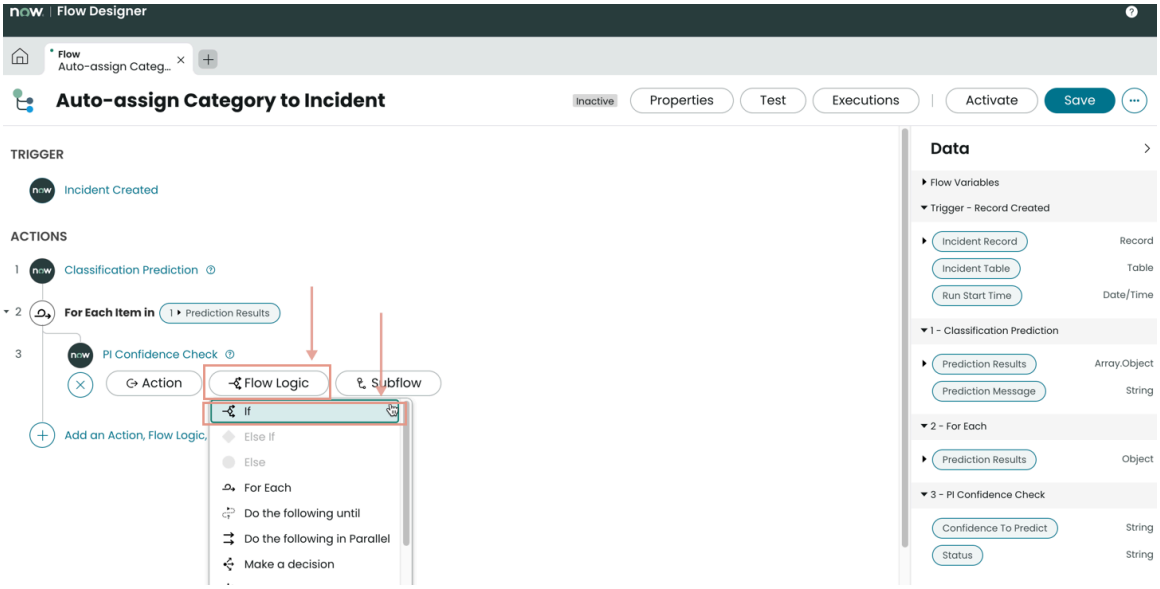


9. Drag and drop the **confidence** data pill into the **Predicted Number from Predictive Intelligence** field.
10. Enter 50 in the **Comparison Threshold** field.
    In this example scenario, you enter the number 50, which tells the system to use predictions that have a confidence score above 50%.
11. Select **Done**.

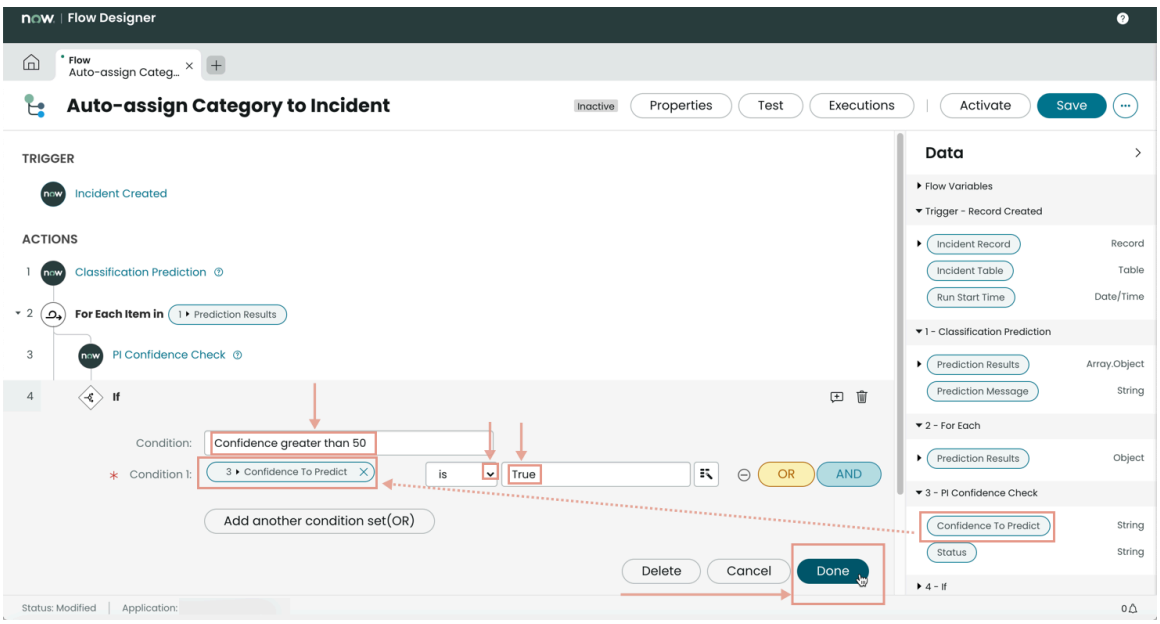


12. Select **Flow Logic** > **If** to add a condition to the flow.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

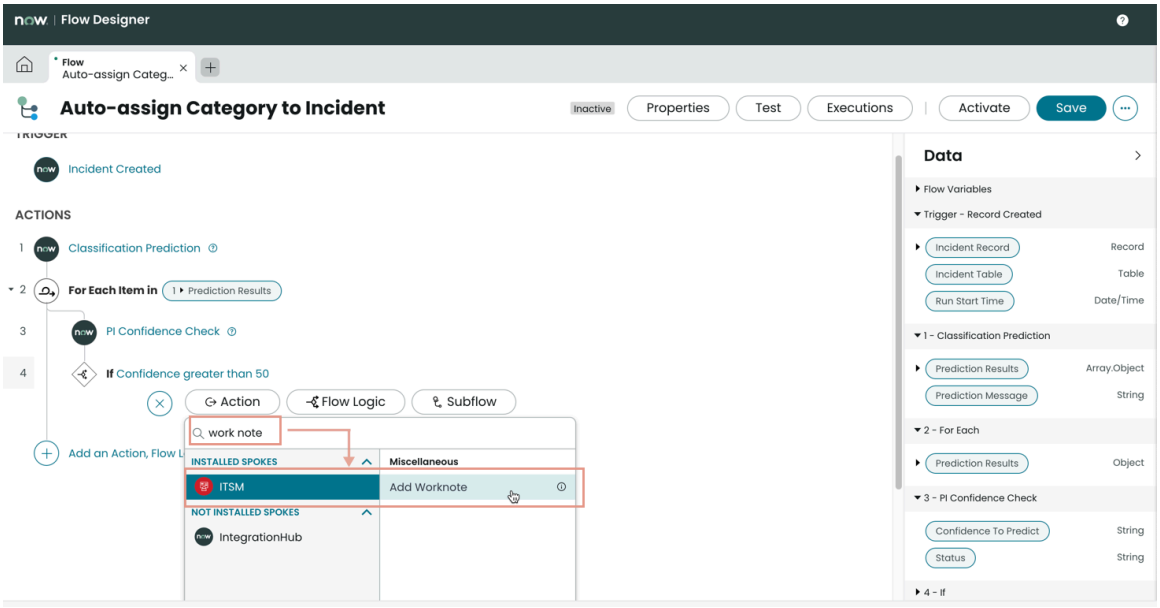


13. Configure the following fields to define the first part of the condition flow logic.
    - **Condition**: Enter a name for the condition that defines what it does. In this example scenario, you enter `Confidence greater than 50`.
    - **Condition 1**: Drag and drop the **Confidence To Predict** data pill into the field. Select **is**, and enter the value `True`. This step completes the first part (the antecedent) of the condition flow logic.
    - Select **Done**.

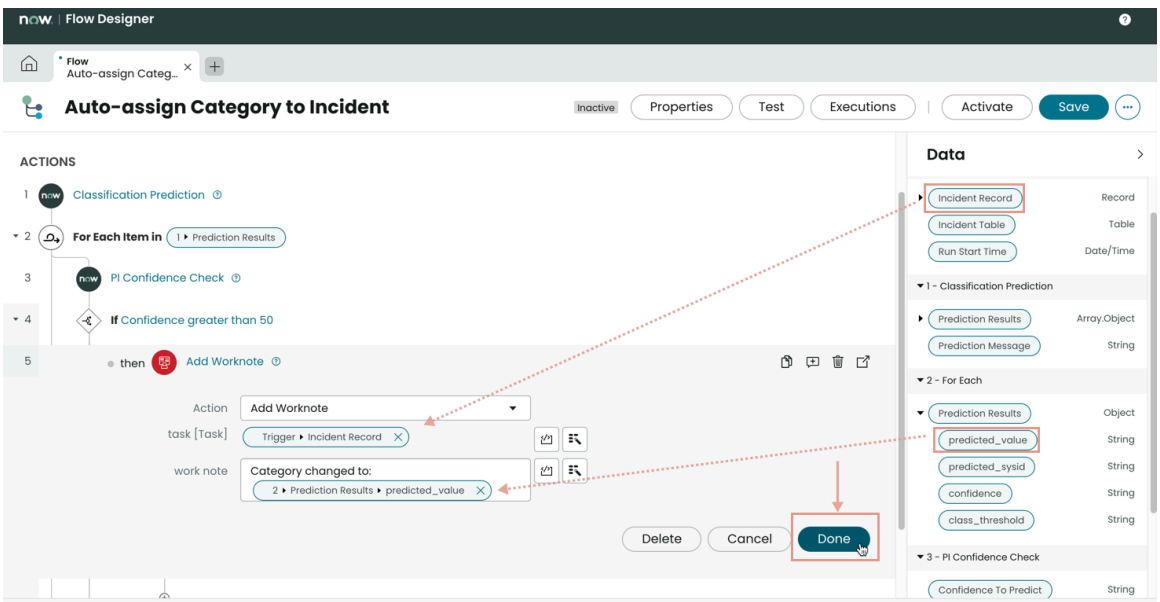


14. Select **Action** and enter `worknote` into the search field.
15. Select **ITSM** > **Add Worknote** to add a work note as the second part (the conclusion) of the condition.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

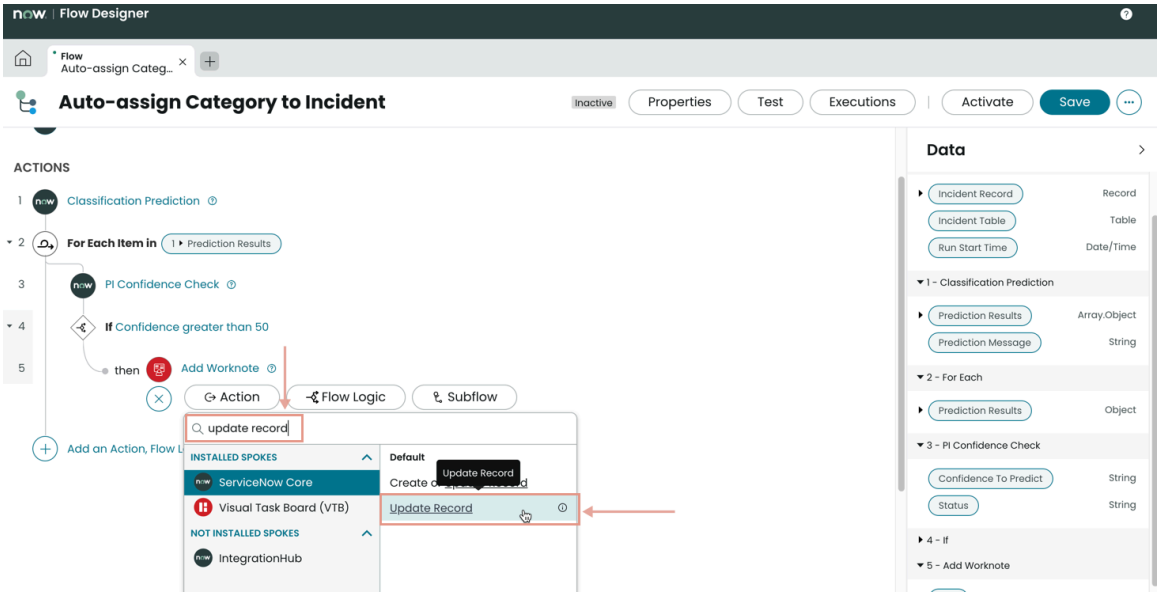


16. Configure the following fields to define the second and final part of the condition flow logic.
    - **Action**: As a result of Step 14 above, **Add Worknote** appears automatically in this field.
    - **task [task]**: Drag and drop the **Incident Record** data pill into the field.
    - **work note**: Drag and drop the **predicted_value** data pill into the field. This step completes the condition flow logic conclusion.
    - Select **Done**.



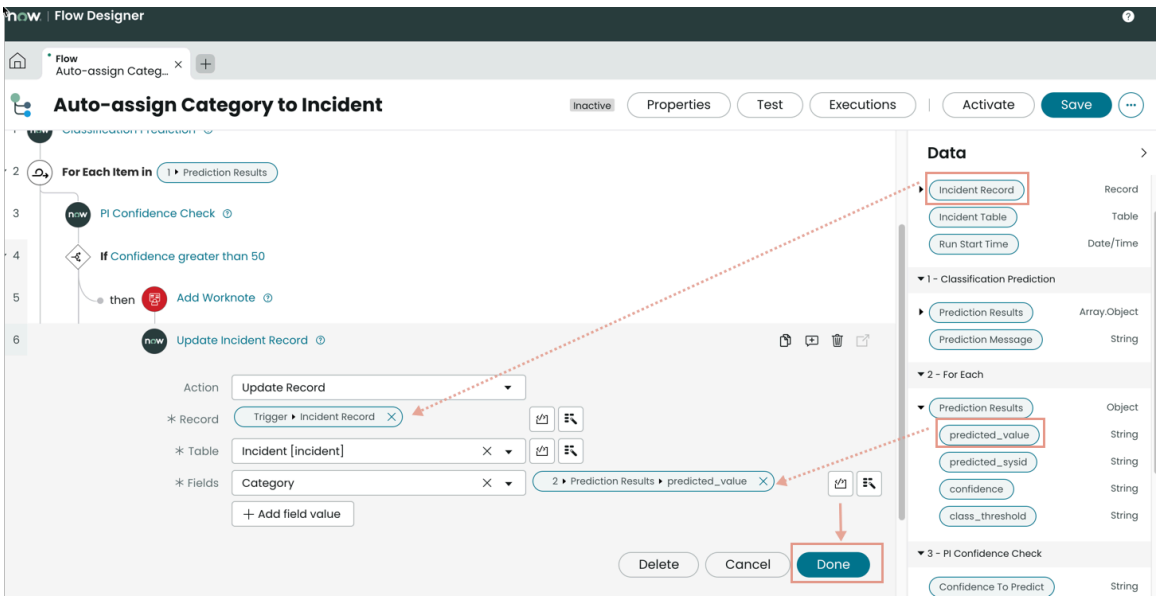
17. Select **Action** and enter `update record` into the search field.
18. Select **Update Record**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



19. Configure the following fields to update the Incident Record.

- **Action**: As a result of Step 16 above, **Update Record** appears automatically in this field.
- **Record**: Drag and drop the **Incident Record** data pill into this field.
- **Table**: Select **Incident [incident]**.
- **Fields**: Select **Category**. Then drag and drop the **predicted_value** data pill into this field, next to the **Category** value.
- Select **Done**.



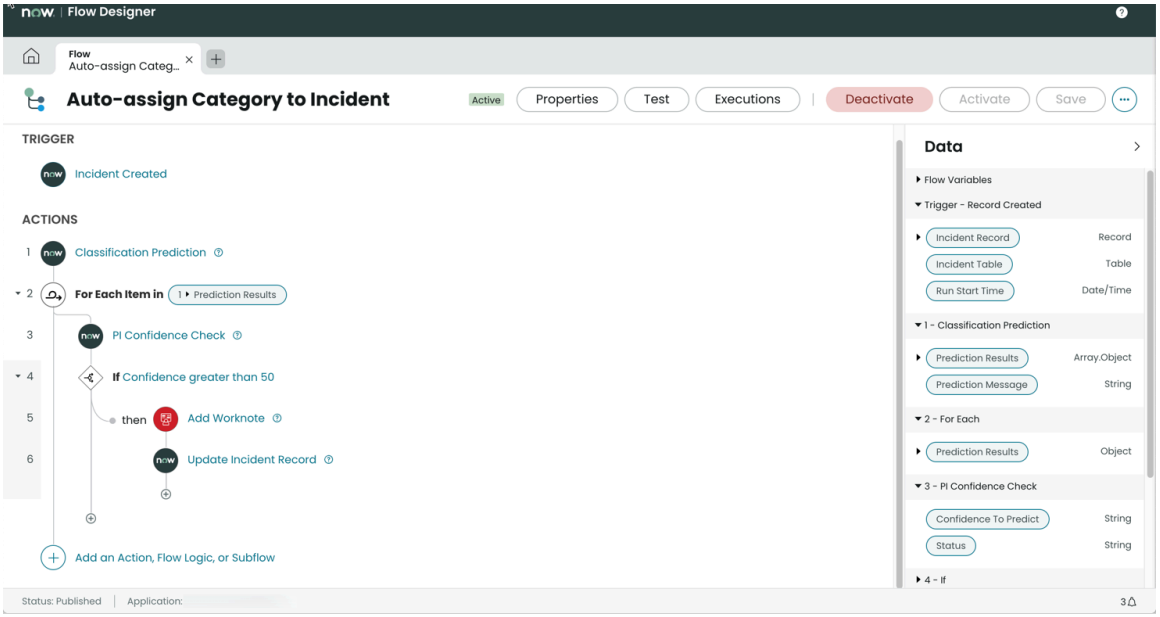
20. Select **Save**.

21. Select **Activate**.

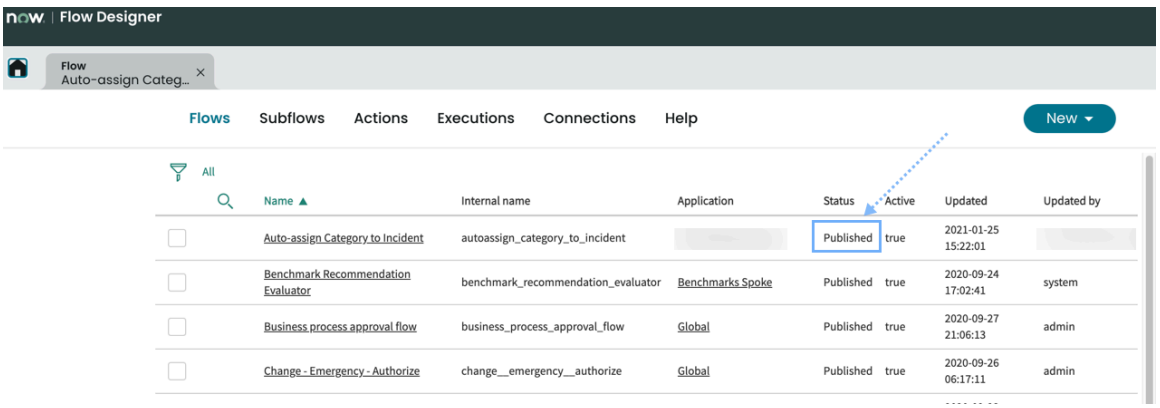
Result

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Your Auto-assign Category to Incident flow is activated and complete.

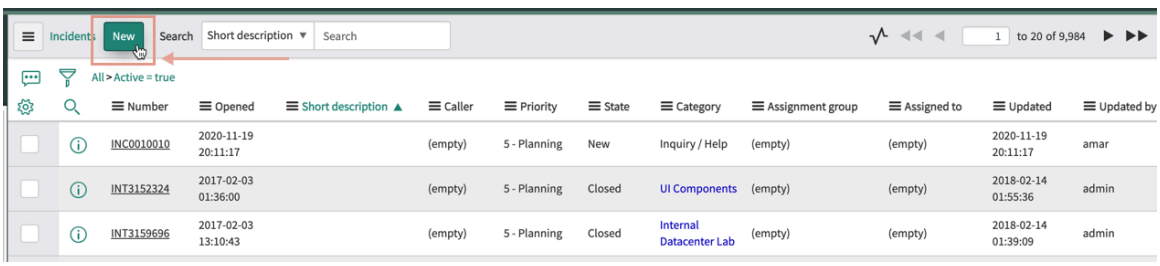


- It also appears as published in the Flows column on the Workflow Studio home screen.



22. Navigate to **Incidents**.

23. Select **New** to create a test incident record in the Incidents table.
In this example scenario, you create record INC0010011.

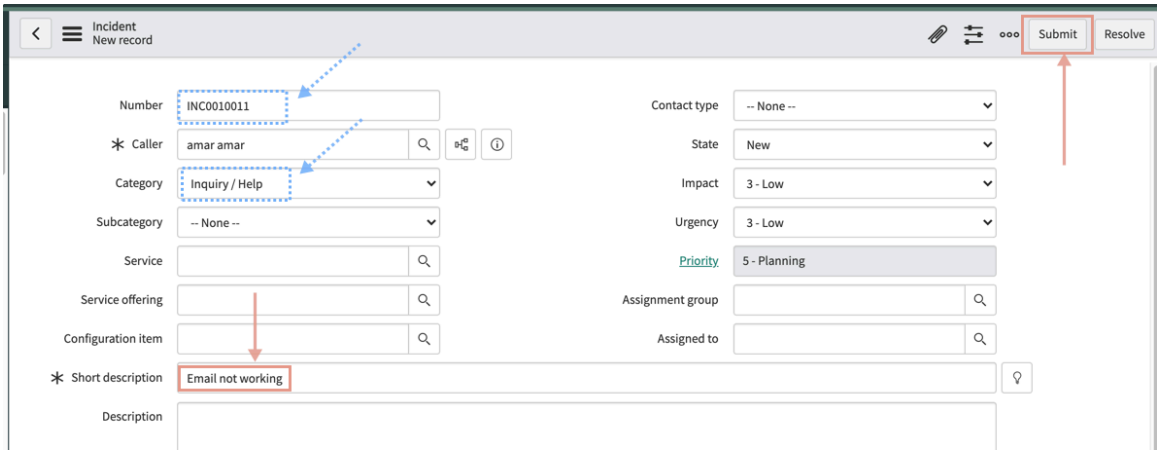


24. In the record you created, note in the following image that the Category value is set to **Inquiry / Help**.

25. In the **Short description** field, enter `Email not working`.

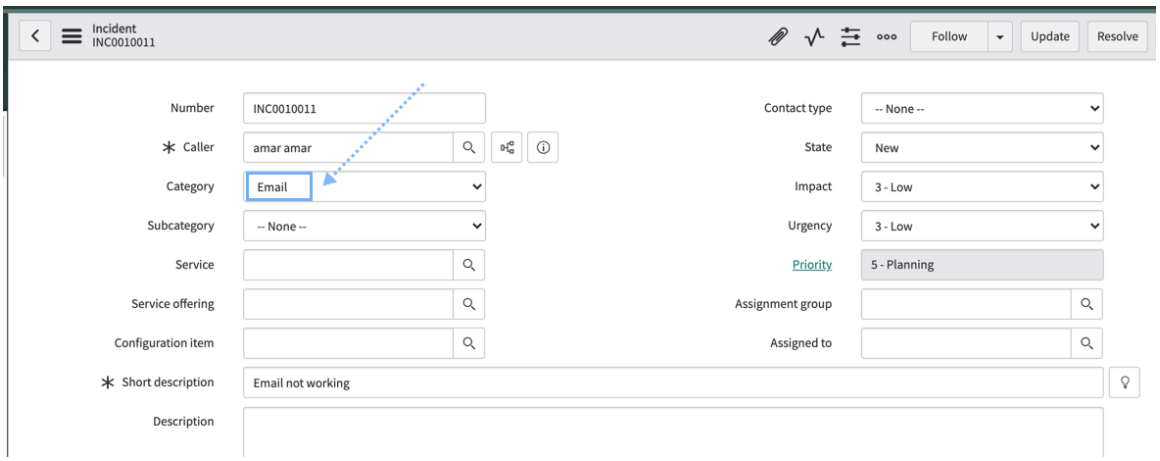
26. Select **Submit**.

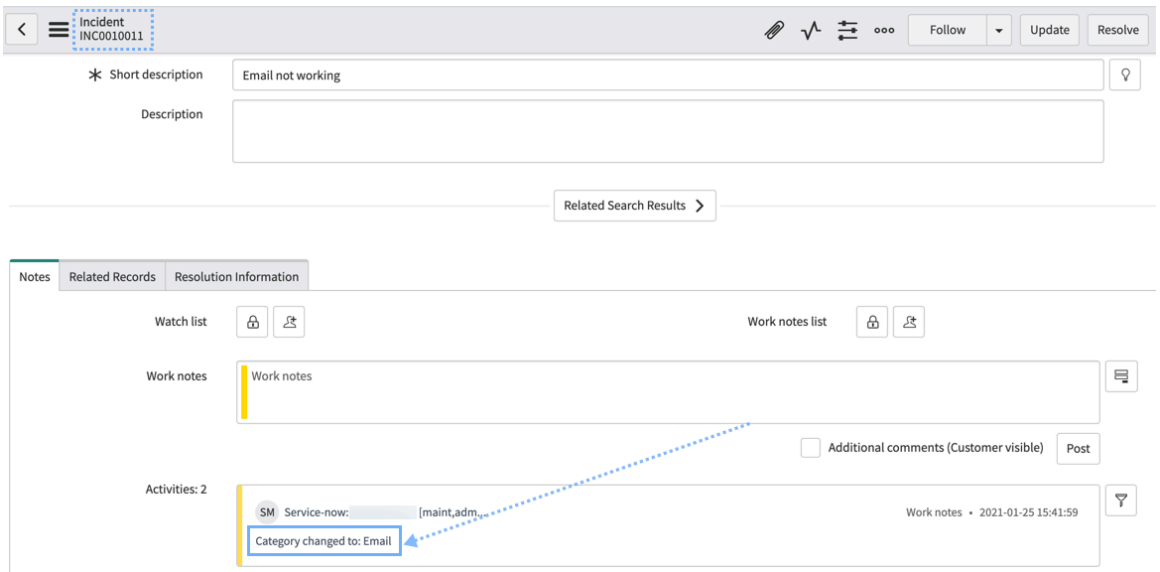
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Result

The system updates the incident record to show that its Category value has changed from **Inquiry / Help** to **Email**.

- 

- 

Related topics

Machine Learning solutions for Flow Designer

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Using MLSolutionFactory scriptable objects

MLSolutionFactory scriptable objects enable defining ML functionality. You can use the APIs to compose data-driven functionality, such as subclustering large clusters or clusters with multiple PRBs attached.

You can use scriptable objects in a scripted extension point to modify the Predictive Intelligence REST API method to address unique business use cases.

ServiceNow applications on the NOW platform can call scriptable objects and script includes. External applications can call scripted REST APIs. By default, the Predictive Intelligence REST API - Prediction for multiple solutions (GET) method uses the MLSolutionUtil scripted extension point to take a list of active solution names, run predictions on them given the input, and return results. The MLSolutionUtil extension point enables creating custom implementations for specific usage scenarios, for example, running a second solution prediction only after the condition of a first solution prediction is satisfied.

Here's the high-level process for creating a custom usage scenario.

1. Developers customize a MLSolutionUtil scripted extension point implementation using the MLSolutionFactory scriptable object.
   - Listed as global.MLSolutionUtil in the Extension Point [sys_extension_point] table
   - See Register a custom script include
2. The MLSolutionUtil scripted extension point implementation uses `MLSolutionFactory` to get the scriptable object, and invokes prediction methods on that object.
3. The Predictive Intelligence REST API - Prediction for multiple solutions (GET) method invokes the `MLSolutionUtil extension point` implementation, depending on the scope of the request.
4. Applications call the Predictive Intelligence REST API - Prediction for multiple solutions (GET) end point from a script include form.

For more information, see

- Predictive Intelligence REST API
- MLSolutionFactory - Global

Preserve ML solutions during a system clone

Save your trained machine-learning (ML) solution data during a system clone.

Before you begin

Role required: clone_admin or admin

About this task

The system stores trained ML solutions as Attachment records. These records include your solution artifacts, such as solution definitions, template records, and predictive model statistics, all of which are required components of the Predictive Intelligence prediction functionality. To preserve these records, follow the high-level steps below each time you run a system clone.

Note:

For troubleshooting common issues with cloning solutions, see the Predictive Intelligence Common issues [KB0781893] article in the Now Support Knowledge Base.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Enter `sys_properties.list` in the application navigator to access the System Properties list.
2. Ensure the *`glide.platform_ml.clone_artifacts`* system property is set to **True**.
3. If you want to preserve only the ML solution records and not the numerous other records in the sys_attachments table, exclude the sys_attachments table from your clone run.
4. Request a system clone.
   The system preserves your ML solution records during the system clone.

Testing and monitoring predictions

Evaluate the coverage and precision of your machine-learning (ML) solutions by testing them. Once deployed, track their performance over time. Improve predictions by using performance information to refine your solutions.

Testing solutions

After your ML solutions are trained, you can call on the Predictive Intelligence API to make a solution prediction. Use the API to test the following solutions:

- Test a classification solution prediction
- Test a similarity solution prediction

Monitoring predictions

For classification predictions, use the Prediction Results dashboard to track the coverage and precision of predictions over time.

To learn how, see Track classification prediction results over time.

Prediction troubleshooting

For troubleshooting common issues with solution predictions, see the Predictive Intelligence Common issues [KB781893] article in the Now Support Knowledge Base.

Test a classification solution prediction

Once your machine-learning (ML) solutions are trained, you can call on the Predictive Intelligence API to make a solution prediction. In this example procedure, we use the REST API Explorer to test a classification solution prediction for incident categorization.

Before you begin

Train your ML solution prior to testing a prediction.

Role required: web_service_admin, rest_api_explorer, or admin or ml_admin

About this task

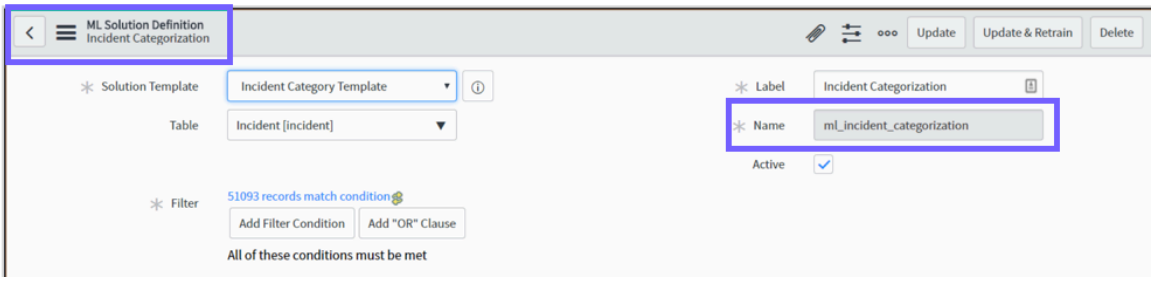
This procedure uses sample data to illustrate what you can do in your instance, and may not represent data or records that are actually in your instance.

This scenario illustrates a classification solution prediction for a hypothetical ML solution that you have previously created and trained. You can also use the REST API Explorer to test a similarity solution prediction.

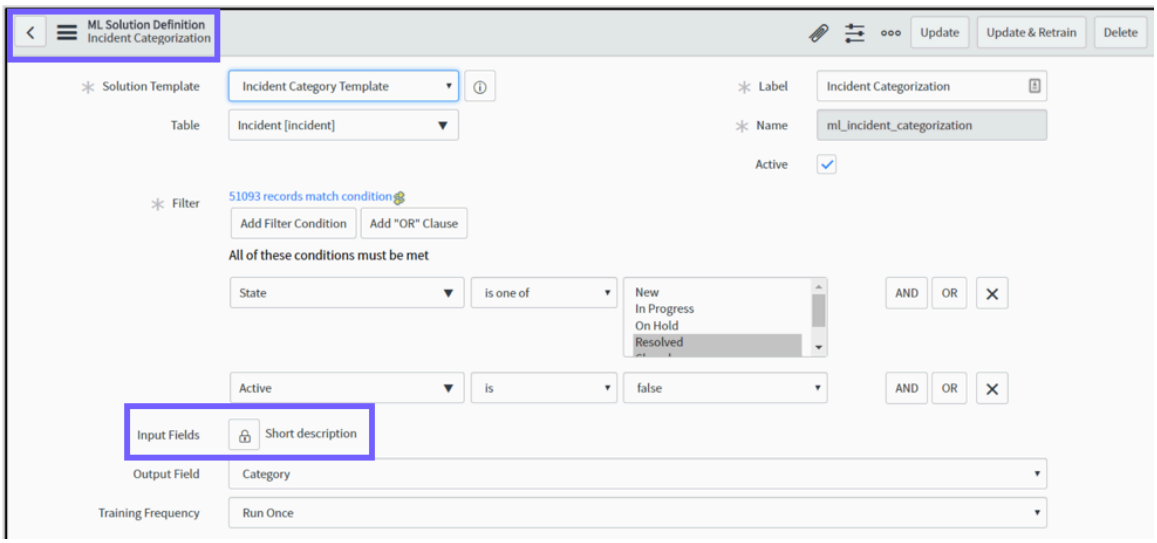
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Classification** > **Solution Definitions**.
2. Locate the ML solution definition whose prediction you want to test, and copy its **Name** value to your clipboard or a Notepad file.
   In this case, use the **Name** field value in your ML Solution Definition Incident Categorization record, as illustrated in the following example.



3. Write down and save the **Input Fields** used in your ML Solution Definition record that you want the REST API Explorer to use in its call to the Predictive Intelligence API.
   In this case, we use the **short _description** field, as the prediction model has been trained to use this field to learn its category definition.



4. Navigate to **System Web Services** > **REST** > **REST API Explorer**.
5. Set these choice fields as follows.

Field	Value
Namespace	now (leave as default)
API Name	Predictive Intelligence
API Version	latest (leave as default)

   The Predictive Intelligence form appears. You use this form to prepare your call request to the Predictive Intelligence API.
6. In the solution-name **Value** field, type `ml_incident_categorization`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: This is the Name value you captured in Step 1 of this procedure.

7. Click **Add query parameter**.
   The Predictive Intelligence form refreshes to show the **Query parameters** section.
8. Type `short_description` in the first field.

   Note: This is the input field you captured in Step 2 of this procedure.

9. Type a short description of an incident in the second field.
   For instance, type `Unable to connect`.
10. Click the **Send** button.
    The REST API Explorer sends your request to the Predictive Intelligence API.
    The system predicts the output value in the Response Body section of the API output. You can use other short descriptions to test what the solution is predicting.
11. **Optional:** Send a different request to the Predictive Intelligence API so that you can test the prediction model again.

    a. Return to the **Query parameters** section of the Predictive Intelligence form.

    b. Type a short description that references a different kind of incident in the second field.
    For example, type `Unable to connect to MSSQL`.

    c. Click the **Send** button.

    The Response Body section may refresh to show a different outcome than what you saw in Step 9, depending on which incident categories you configured in your solution definition setup. In other words, changing the short description text can recategorize the incident as a different kind of issue.

Example:
You can also test the Predictive Intelligence prediction model when you create a new incident record using the incident form.

1. Navigate to **Incident** > **Create New**.
2. In the new Incident form that loads, set the fields as follows.
   - **User:** Enter the Caller name.
   - **Category:** Leave as default.
   - **Short description:** Enter a short description that you want to test.
3. Submit the incident form.

Result: When the form refreshes, an information message appears with the incident category automatically set to a specific value.

Note: For some short descriptions, the prediction might not process because the solution does not have enough confidence in predicting the value for this input.

Related topics

Predictive Intelligence API

MLPredictor - Global

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Test a similarity solution prediction

Once your machine-learning (ML) solutions are trained, you can call on the Predictive Intelligence API to make a solution prediction. In this example procedure, we use the REST API Explorer application to test a similarity solution prediction for resolved incident recommendations.

Before you begin

Train your ML solution prior to testing a prediction.

Role required: web_service_admin, rest_api_explorer, or admin or ml_admin

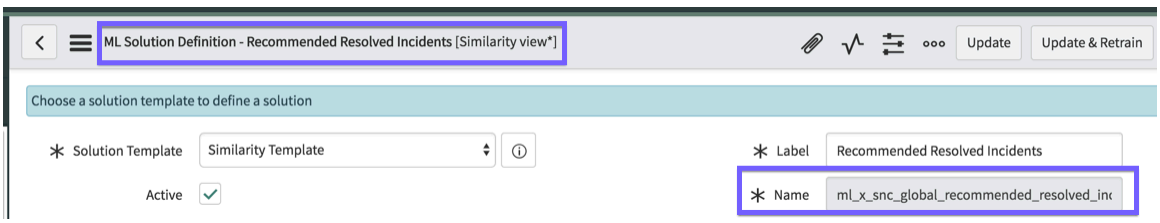
About this task

This procedure uses sample data to illustrate what you can do in your instance, and may not represent data or records that are actually in your instance.

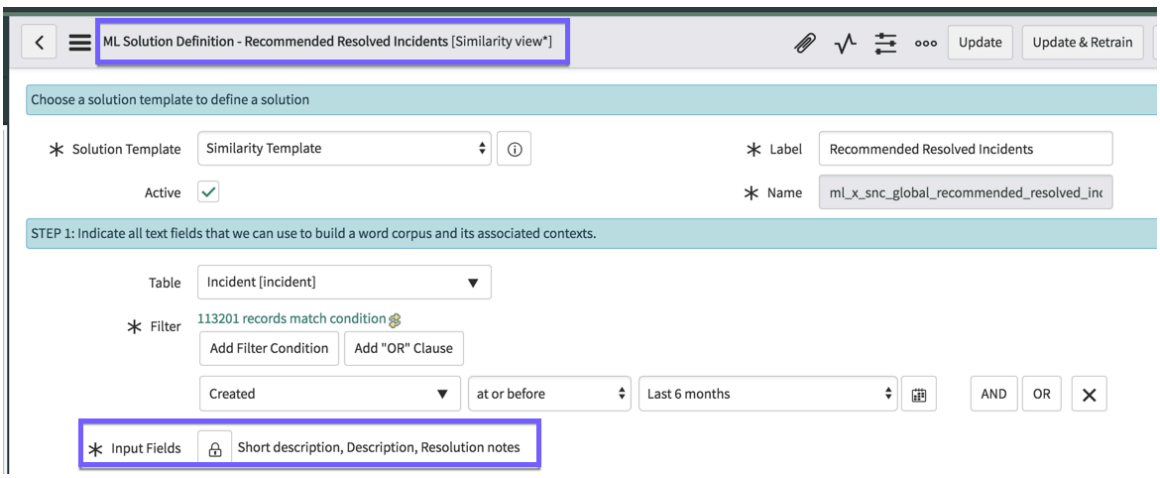
This scenario illustrates a similarity solution prediction for a hypothetical ML solution that you have previously created and trained. You can also use the REST API Explorer to test a classification solution prediction.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Similarity** > **Solution Definitions**.
2. Locate the ML solution definition whose prediction you want to test, and copy its **Name** value to your clipboard or a Notepad file.
   In this case, we use the **Name** field value in your ML Solution Definition Recommended Resolved Incidents record, as illustrated in the following example.



3. Copy the **Input Fields** value(s) used in your ML Solution Definition record that you want the REST API Explorer to use in its call to the Predictive Intelligence API.
   In this case, we use the **Short description** field type, as the prediction model has been trained to use this field to learn, pair, and recommend similar records for your review.



4. Right-click the browser tab you're using to view your instance, and select **Duplicate**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. In the duplicate browser tab, navigate to **System Web Services** > **REST** > **REST API Explorer**.
6. Click **Explore**.
7. Set these choice fields as follows.

Field	Value
Namespace	now (leave as default)
API Name	Predictive Intelligence
API Version	latest (leave as default)

The Predictive Intelligence form appears. You use this form to prepare your call request to the Predictive Intelligence API.

8. In the solution-name **Value** field, enter `ml_x_snc_global_recommended_resolved_incidents`.

Note: This is the Name value you captured in Step 2 of this procedure.

9. Click **Add query parameter**.
10. In the Query parameters section, enter the value of one of the Input Fields from the solution you're testing.

a. In the first field, paste `short_description`.

Note: This is one of the input fields you captured in Step 2 of this procedure.

You can use other field types, such as **Description** or **Resolution notes** to test what the solution is predicting.

b. In the second field, enter some text that you might find in an incident record.
For instance, enter `Discovery errors`.

11. Click the [+] button to create a second query condition that defines the number of results you want to query.

a. In the first field, enter `top_n`.

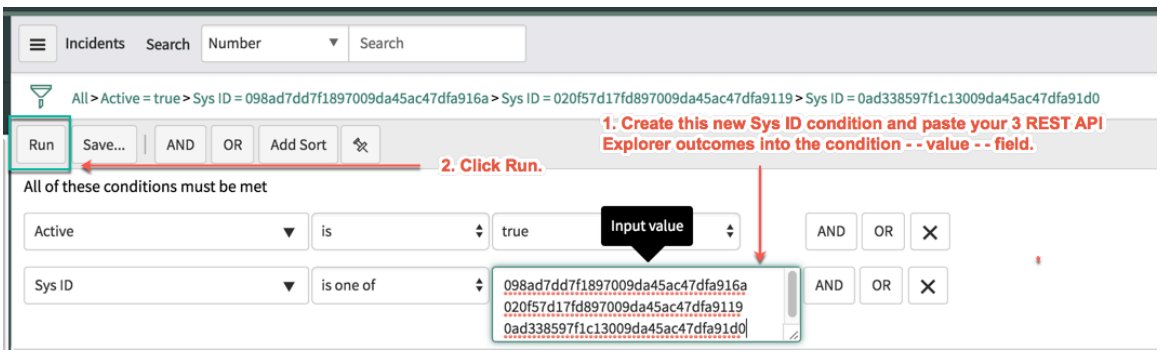
b. In the second field, enter `3`.
These conditions set the query to retrieve the top three most similar incident records.

12. Click **Send**.
The REST API Explorer sends your request to the Predictive Intelligence API.
13. In the Response body section, copy the three outcome values that your API call returned, as illustrated in the image below.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



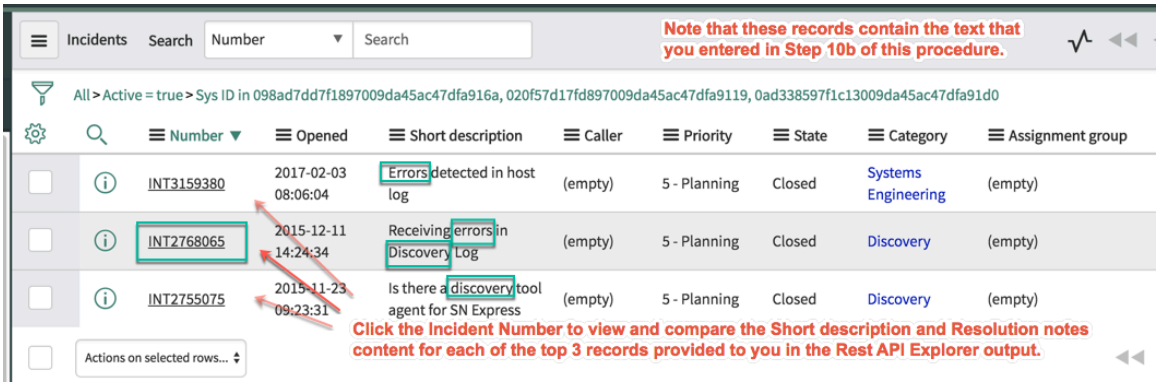
14. In your original browser tab, navigate to **Servicedesk** > **Incidents**.
15. As shown in the image below, set filter conditions for the three REST API outcomes to the Incidents table list view.
    a. Add the Active and Sys ID conditions below to the Incidents list view Filter icon.
    b. Paste the three REST API outcomes into the **Input value** field of the Sys ID condition that you created.
    c. Click **Run**.



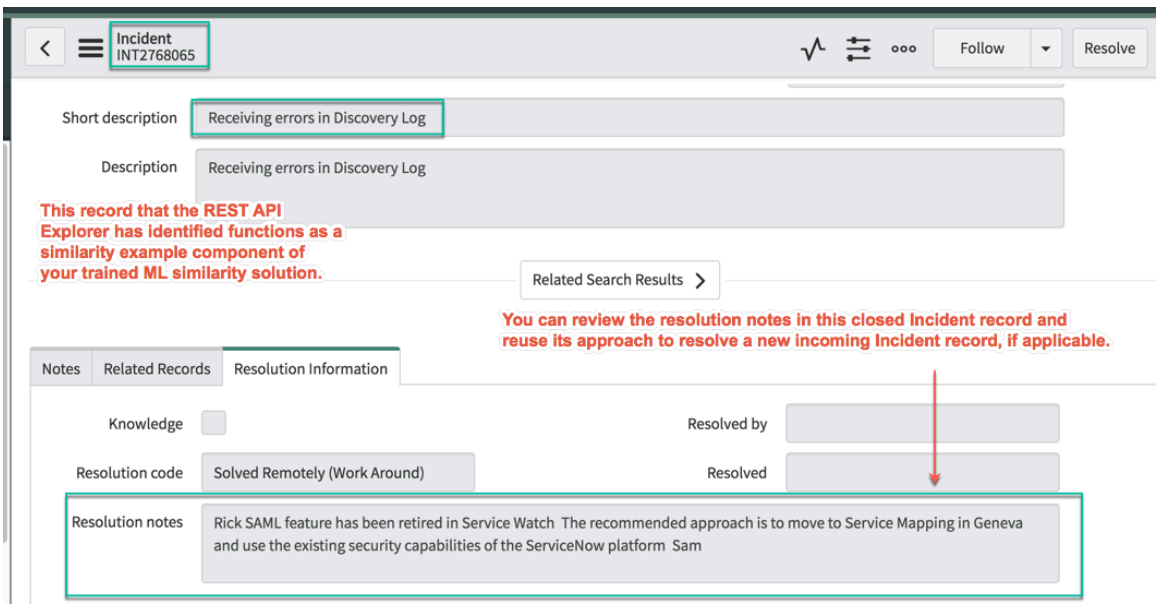
16. Per the image below, compare the returned list of incidents with the input for the prediction output in the REST API Explorer.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Click the Incident Number to open the Incident record.



b. Per the image below, review the Resolution notes text in the Incident record.



Related topics

Predictive Intelligence API

MLPredictor - Global

Track classification prediction results over time

Use the Prediction Results dashboard to determine if classification solution predictions are improving over time. Identify solutions that need refining or retraining.

Before you begin

- Role required: admin, ml_admin, or ml_report_user

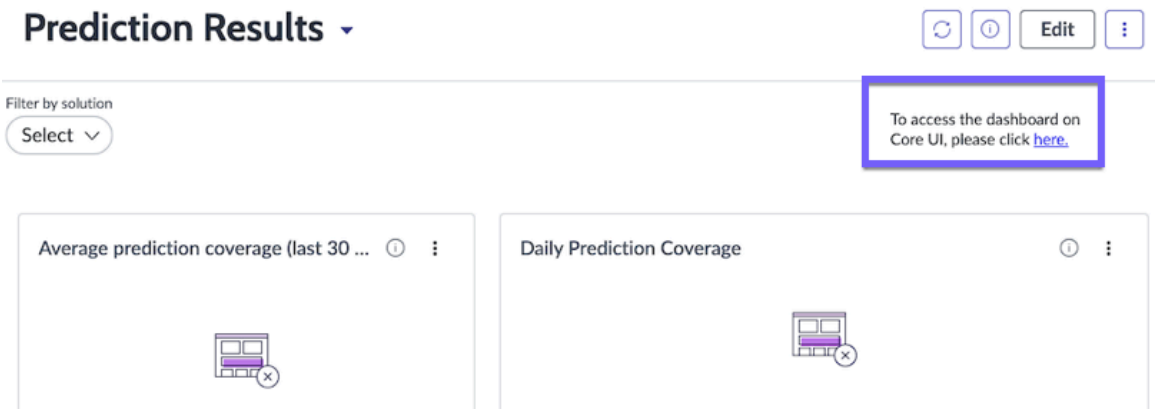
About this task

The Prediction Results dashboard reports on coverage, precision, and recall over time for classification solutions.

With the Xanadu release, this dashboard has been migrated to the Next Experience UI. Customers upgrading from previous releases

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

can access the Core UI version from the current dashboard.



On the Prediction Results dashboard, statistics are provided in two timeframes: the average for the past 30 days, and daily. The indicators coverage, precision, and recall are defined as follows.

Prediction Results indicators

Report type	Definition
Coverage	The percentage of predictions that yielded an outcome out of the total number of predictions that were attempted.
Precision	The percentage of predictions where the predicted value was the same as the final value of the field when the report closed.
Recall	The percentage of correct predictions that yielded an outcome out of the total number of predictions that were attempted.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Classification** > **Prediction Results Report**.
2. In the Prediction Results dashboard **Filter by solution** prompt, select the solution statistics you want to review.
   The system updates the dashboard based on the solution you selected.
3. Identify classes with anomalous coverage, precision, or recall values.

   Example
   For example, identify solutions where coverage, precision, or recall is declining over time.

What to do next

Refine the solution definition filter by including or excluding classes as needed. After updating, retrain the solution.

Reviewing prediction errors with the Observability Dashboard

The Observability Dashboard offers a unified view and actionable insights for errors detected in Predictive Intelligence. Use this dashboard to visualize logged errors and gain information on prediction reliability and potential problem areas.

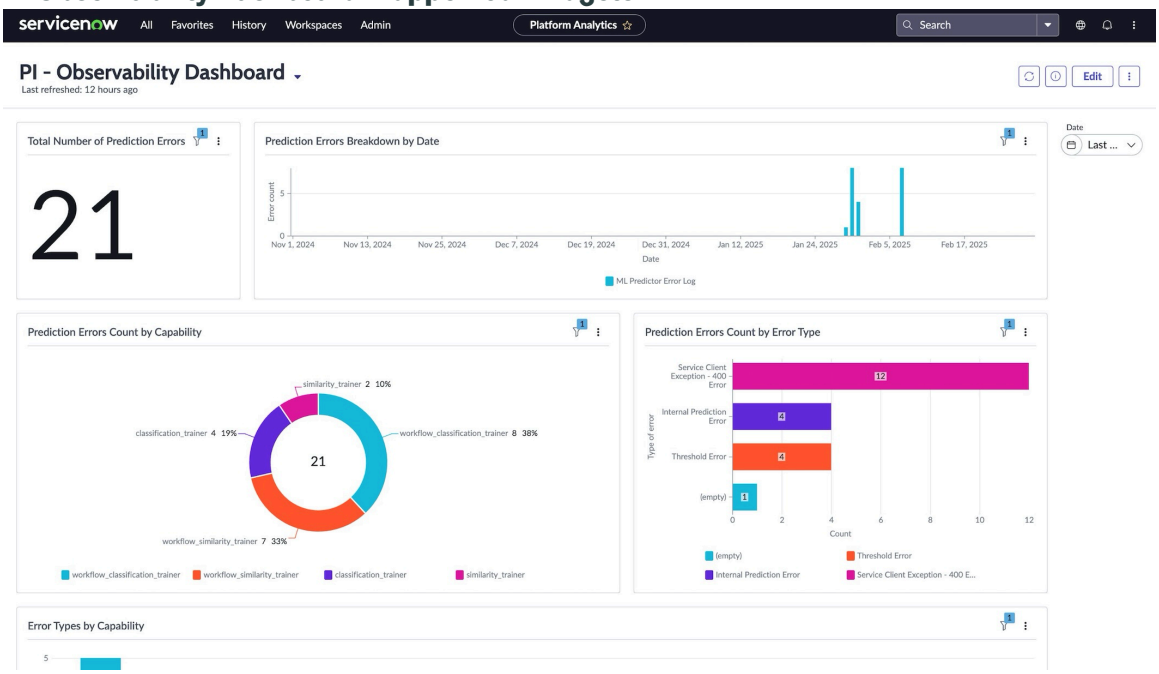
View the PI - Observability Dashboard by navigating to **Predictive Intelligence** > **Dashboards** > **Observability**. The dashboard contains the following widgets.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Total Number of Prediction Errors
- Prediction Errors Breakdown by Date
- Prediction Errors Count by Capability
- Prediction Error Count by Error Type
- Error Types by Capability
- Successful and Unsuccessful Predictions Breakdown by Date

You can drill down to the underlying records from the widget graphics. You can also change the date range of all widgets by selecting **Date** to open the selector.

PI Observability Dashboard — upper four widgets



The PI - Observability Dashboard draws from a table dedicated to logging prediction errors: **ML Predictor Error Logs** [ml_predictor_error_logs].

- Fields in the table include Error Type, Error Message, Status code, Solution, and Capability.
- View this table's records directly by navigating to **Predictive Intelligence** > **Observability** > **Prediction Error Logs**.
- Roles required to access the table: ml_report_user or ml_admin.

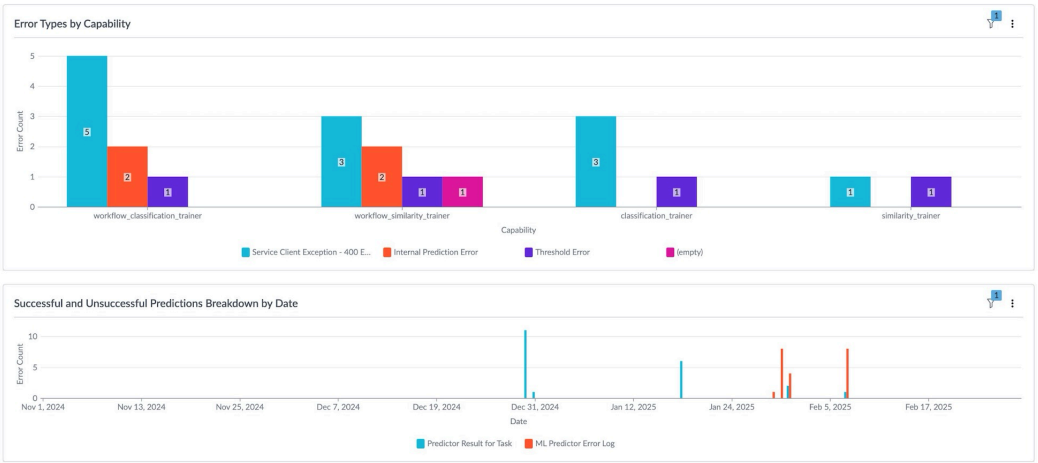
The table logs the following types of granular errors.

- Client-side issues (400 Series) — captures request errors such as timeouts and invalid inputs.
- Server-side issues (500 Series) — tracks internal server errors encountered during prediction.
- Internal prediction failures — identifies instances when the model was unable to generate a prediction.
- Low confidence predictions — records log results falling below a defined confidence threshold.

Note: Errors in training aren't included in this table. For a dashboard reporting on training errors, see Predictive Intelligence Usage Analytics dashboard.

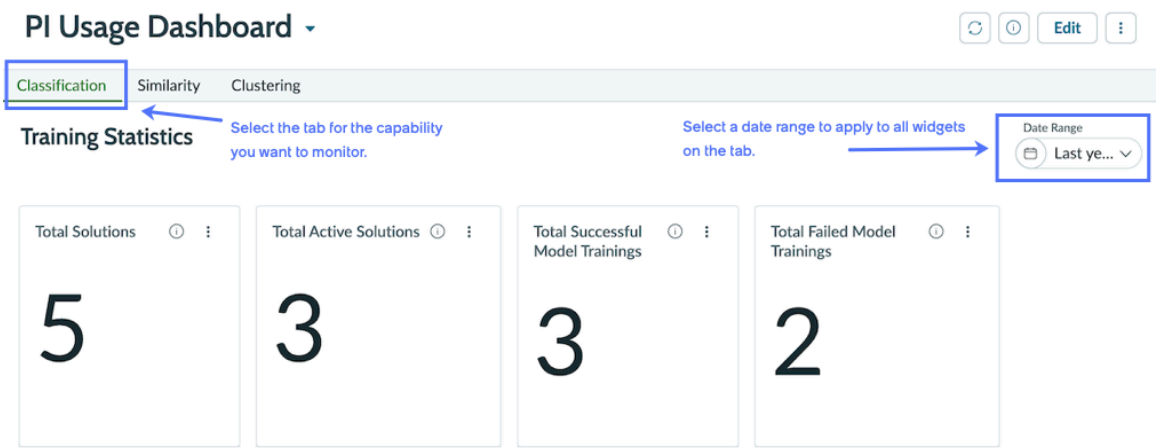
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

PI Observability Dashboard — lower two widgets



Predictive Intelligence Usage Analytics dashboard

The Predictive Intelligence Usage Analytics dashboard is a central location to understand the effectiveness and overall value of all your Predictive Intelligence solutions. View metrics for model training successes and failures. Monitor prediction statistics including breakdowns by individual model.



For your models in the Classification, Similarity, and Clustering capabilities, the Predictive Intelligence Usage Analytics dashboard gives you visibility into training and predictions. You can specify the time periods to report on, as well as select specific solutions for a deep dive.

With the ml_report_user or ml_admin role, view the Predictive Intelligence Usage Analytics dashboard by navigating to **Predictive Intelligence** > **Dashboards** > **Monitoring**.

The dashboard has separate tabs for each capability. The Classification tab opens by default, or you can select Similarity or Clustering. Each tab contains a set of reports in the form of widgets.

You can drill down into the underlying records by selecting a report's chart or graph, which opens a query-filtered list. Many reports in the Training Statistics section of PI Usage Dashboard are drawn from the ml_solution table. Reports in the Prediction Statistics section are drawn from the ml_predictor_results table.

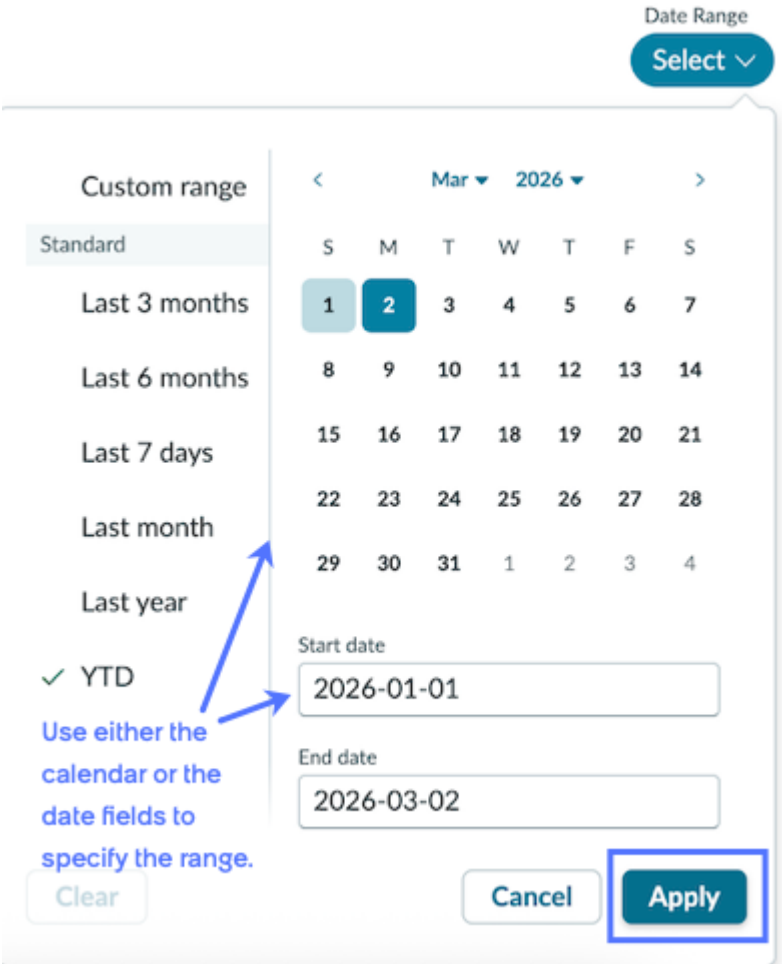
Date Range

The **Date Range** selector opens a calendar picker where you can select the time frame for the reports. Your selected date range applies to all widgets on the current tab.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The calendar picker offers a variety of preset dates under its **Standard** subhead. Choose from **Last 3 months**, **Today**, **YTD** (year to date, from January 1 of the current calendar year), and so forth.

You can also choose **Custom range** to select your own dates. Enter dates directly into the **Start date** and **End date** fields, or select a date range using the calendar. Then select **Apply** to update the reports in all widgets on the page.



Available metrics for Classification and Similarity

The following tables describe the widgets available on both the Classification and Similarity tabs, unless otherwise noted. The Clustering capability has different widgets, so refer to the **Available metrics for Clustering** section on this page.

Widgets available in Training Statistics for Classification and Similarity solutions

Widget	Description
Total Solutions	Total number of Classification (or Similarity) solutions created.
Total Active Solutions	Total number of Classification (or Similarity) solutions that are active.
Total Successful Model Trainings	The number of Classification (or Similarity) jobs that have been successfully completed in the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

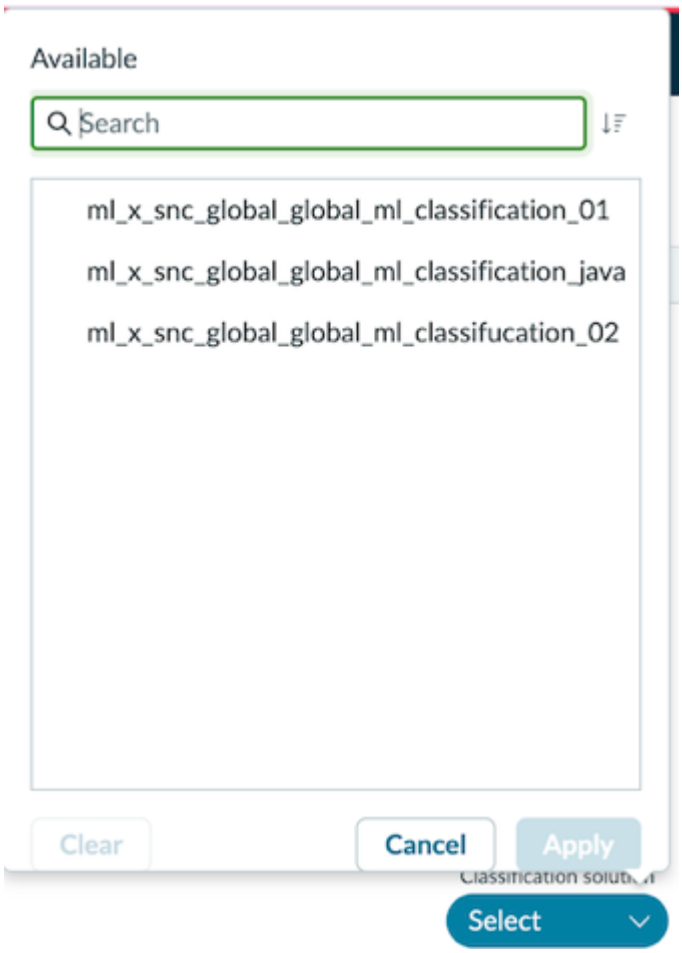
Widgets available in Training Statistics for Classification and Similarity solutions (continued)

Widget	Description
Total Failed Model Trainings	The number of Classification (or Similarity) jobs that have failed in the instance.
Trained Solutions - By Month	Total trained Classification (or Similarity) solutions aggregated monthly.
Active Solutions - By Month	Total active Classification (or Similarity) solutions aggregated monthly.
Classification (or Similarity) Training Status	The distribution of Classification (or Similarity) model training states.
Active Status Distribution	The distribution of active versus inactive records. The "true" portion indicates active items, while "false" represents inactive ones.
Training Frequency Distribution	The training frequency distribution across all Classification (or Similarity) solutions. The available options range from **Run once** up to **Every 180 days**.
Update Frequency Distribution (Similarity only)	The update frequency distribution across all similarity solutions. The available options range from **Run once** up to **Every 180 days**.

Prediction Statistics

Widget	Description
Total Successful Predictions - By Month	Count of total Classification (or Similarity) predictions calculated per month.
Total Successful Predictions	Total number of successful predictions made on all Classification (or Similarity) solutions.
Total Failed predictions	Total number of failed predictions made on all Classification (or Similarity) solutions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



For the following widgets, use the selector labeled **Classification solution** or **Similarity solution** to choose an individual model.

Selected Active Classification Solution Details

Widget	Description
Successful Predictions per Solution	The number of successful predictions for the selected Classification (or Similarity) solution.
Failed Predictions per Solution	The number of failed predictions for the selected Classification (or Similarity) solution.

Available metrics for Clustering

The following table describes metrics available on the Clustering tab of the Predictive Intelligence Usage Analytics dashboard.

Training statistics for Clustering

Widget	Description
Total Solutions	The total number of clustering solutions created in the instance.
Total Active Solutions	The number of active clustering solutions in the instance.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Training statistics for Clustering (continued)

Widget	Description
Total Successful Model Trainings	The number of clustering jobs that have been successfully completed in the instance.
Total Failed Model Trainings	The number of clustering jobs that have failed in the instance.
Trained Solutions - By Month	The number of clustering trainings triggered per month.
Active Solutions - By Month	The number of active clustering solutions trained per month.
Clustering Training Status	A pie chart showing the training status distribution across all clustering solutions in the instance.
Active Status Distribution	A pie chart showing the distribution of active and inactive clustering solutions.
Training Frequency Distribution	The percentage of training frequencies set across the clustering solutions.
Update Frequency Distribution	The percentage of update frequencies set across the clustering solutions.

For the following widgets, use the selector labeled **Clustering solution** on the Clustering tab to pick an individual model.

Selected Cluster Solution Breakdown

Widget	Description
Cluster details	The details of each cluster created in the specified solution.
Number of clusters	The number of clusters created in the specified clustering solution.
Total Records Clustered	The total number of records clustered in the specified clustering solution.
Number of updates completed (in last 7 days)	The number of update jobs completed for a specified clustering solution.
Details of Last Update (in last 7 days)	When the clustering solution was last updated, number of new clusters created, records assigned to existing clusters, and new clusters in the last update.

Predictive Intelligence references

Reference pages for components of Predictive Intelligence.

Domain separation and Predictive Intelligence

Domain separation is supported in the Predictive Intelligence application. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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 of domain separation and Predictive Intelligence

With Predictive Intelligence, you can create machine learning solutions using historic datasets. A machine learning solution is created and trained in the global domain, and that solution's predictions can be applied in any domain on the instance.

Similarity solutions are domain-aware when applied in your forms and flows, so records from other domains on the instance are not displayed to users. For more information, see KB article Similarity prediction behavior in domain separated environment on Now Support.

How domain separation works in Predictive Intelligence

An instance owner can train a machine learning (ML) solution for each domain by creating a solution definition for each domain and training those solutions. In this way each solution uses data specific to the corresponding domain.

- Data can be domain separated
- 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

Related topics

Domain separation for service providers

Data Encryption in Predictive Intelligence

Learn which types of encryption are supported for training Predictive Intelligence solutions.

Predictive Intelligence support for encrypted training data

Encryption type	Supported	Notes
Field Encryption	Yes	Ensure the sharedservice.worker user has the same encryption module role that's been used for encryption.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Encryption type	Supported	Notes
Edge Encryption	No	None.
FDE (Full Disc Encryption)	Yes	None.

Predictive Intelligence language support

Predictive Intelligence provides international language support. Learn which languages are available to Predictive Intelligence solutions.

When you create a Predictive Intelligence solution, you can choose the language you want the system to use for processing your training records. English is the default language.

For an example of how to assign a language to a solution, see the Processing Language section in Step 3 of the Create and train a classification solution documentation.

You can also create custom stopwords lists for a language. For more information, see Create a custom stopwords list.

Language support coverage

The current available languages for Predictive Intelligence solutions and stopwords lists are as follows: Brazilian Portuguese, Chinese (simplified), Danish, Dutch, English, Finnish, French, French Canadian, German, Italian, Japanese, Korean, Norwegian, Polish, Portuguese, Spanish, and Swedish.

Predictive Intelligence properties

The properties for Predictive Intelligence control certain parameters of its machine-learning solutions, solution training process, and caching.

Using the maint role, navigate to **Predictive Intelligence** > **Configuration** to review or edit.

Note: Most of the properties related to training require the maint role to review or edit. Some properties may require other roles.

Predictive Intelligence properties

Property	Property Name	Description
Override ml_solution_definition record's readonly state during training	glide.platform_ml.override_training_lock	• True: Overrides the **readonly** state during training, enabling you to edit ml_solution_definition. • False: ml_solution_definition remains in the **readonly** state.
Minimum number of records for Regression	glide.platform_ml.api.min_regression_records	Sets the minimum number of records required for Regression solution training.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Predictive Intelligence properties (continued)

Property	Property Name	Description
		• Type: Integer • Default value: 10000 **Note:** Support for new regression solutions is deprecated as of the Washington DC release. You can edit and train any existing solutions, but you can't create new ones.
Maximum number of records for Regression	glide.platform_ml.api.max_regression_records	Sets the maximum number of records that can be used in Regression solution training. • Type: Integer • Default value: 300000 **Note:** Support for new regression solutions is deprecated as of the Washington DC release. You can edit and train any existing solutions, but you can't create new ones.
The time (in ms) an in-training solution will wait without updates before timing out	glide.platform_ml.training_timeout	Sets the time-out time in milliseconds for an in-training solution. • Type: Integer • Default value: 21600000 (ms)
Maximum model size	glide.platform_ml.api.model_size	Sets the maximum number of records you can use to train a model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Predictive Intelligence properties (continued)

Property	Property Name	Description
		• Type: Integer • Default value: 524288000
Maximum number of records used in training	glide.platform_ml.api.csv_max_line	Sets the maximum number of records that can be used in Classification solution training. • Type: Integer • Default value: 300000
Minimum number of records used in training	glide.platform_ml.api.csv_min_line	Sets the minimum number of records required for Classification solution training. • Type: Integer • Default value: 10000
Maximum number of days worth of records on request can retrieve	glide.platform_ml.api.csv_split_days	Sets the maximum number of days a single request can retrieve from your records. • Type: Integer • Default value: 30
Maximum number of records per table for word corpus	glide.platform_ml.api.max_wordcorpus_records	Sets the maximum number of records per table for Word Corpus creation for Similarity and Clustering solutions. • Type: Integer • Default value: 300000

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Predictive Intelligence properties (continued)

Property	Property Name	Description
		Note: With the Washington DC release, classification, clustering and similarity solutions use Workflow solutions. These are pre-trained so a word corpus isn't needed. The Word Corpus field is removed from Workflow solution forms.
Maximum number of records for similarity window (to lookup results)	glide.platform_ml.api.max_similarity_window_records	Sets the maximum number of records that the Similarity Window can retrieve to look up results. • Type: Integer • Default value: 100000
Maximum number of records for Clustering	glide.platform_ml.api.max_clustering_records	Sets the maximum number of records you can include in a cluster. • Type: Integer • Default value: 100000

Shared Service Scheduler

Property	Property Name
Shared service scheduler url	glide.shared_service_scheduler.url

Machine Learning Artifacts

Property	Property Name
Maximum number of artifacts cached (in MB)	glide.cache.size.ml_object_cache
Artifact cache compression scheme	glide.platform_ml.artifact.cache_compression_scheme

Predictive Intelligence roles

Predictive Intelligence is installed with these roles.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To learn more about managing per-user subscriptions, see Managing per-user subscriptions in Subscription Management and contact your account representative.

ML Admin [ml_admin]

Can create, read, write or delete the ml_predictor_results table and the ml_predictor_results_task table.

Contains Roles

List of roles contained within the role.

- *pa_data_collector*
- *pa_viewer*

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.

ML Labeler [ml_labeler]

Can create, read, write or delete the ml_label_candidate table.

Contains Roles

List of roles contained within the role.

- *sn_ace.ace_user*
- *nlu_user*

Groups

List of groups this role is assigned to by default.

None.

Special considerations

None.

ML Report User [ml_report_user]

Can read the ml_predictor_results table and the ml_predictor_results_task table.

Contains Roles

List of roles contained within the role.

pa_viewer

Groups

List of groups this role is assigned to by default.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

None.

Special considerations

None.

Now Assist in Document Intelligence

With ServiceNow® Now Assist in Document Intelligence, you can use generative AI to get key information from digital documents into your automation workflows.

https://player.vimeo.com/video/1164703930?h=61afcd70f4&badge=0&autopause=0&player_id=0&app_id=58479

Get started

Explore



Learn about the generative AI skills that are available in Now Assist in Document Intelligence

Configure



Activate Now Assist in Document Intelligence and configure the generative AI skills

Use



Use the generative AI capabilities that are offered by Now Assist in Document Intelligence

Reference



Get details about forms, tables, and more

Note: Depending on your license, you will have access to certain application features, generative AI skills, agentic workflows, and AI agents. For more information, see ServiceNow product tiers.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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

- Additional resources for AI products and solutions.
- ServiceNow Community on AI and Intelligence.
- Access real-time courses, self-paced training, and career resources at ServiceNow University.
- Search the Known Error Portal for known error articles.
- Contact Customer Service and Support.

None of the Advanced AI and Data Products provide legal or professional advice. The outputs provided by the Advanced AI and Data Products are for informational purposes only and are not a substitute for advice from a qualified professional. Customer assumes all responsibility and obligations with respect to any decisions, advice, conclusions, legal opinions, recommendations made or given as a result of the use of the services, including without limitation, any decision made, or action taken by Customer in reliance upon the Advanced AI and Data Products.

The Advance AI and Data Products and services do not and are not intended to constitute legal advice and do not create an attorney-client relationship.

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. In addition, this application will collect information extracted from documents. Customers can opt out of future data collection at any time, as described in the Now Assist Opt-Out page.

For more information, see the Now Assist documentation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Exploring Now Assist in Document Intelligence

With Now Assist in Document Intelligence, you can use generative AI to extract key information from your documents and images for use in your workflows.

Now Assist in Document Intelligence overview

Now Assist in Document Intelligence makes the following generative AI capabilities available:

- Extract information from documents. Extract information in the form of fields, tables, and answers to predefined questions from documents and image files. The information can be reviewed by agents in the Document Intelligence workspace, stored in mapped fields, and used as defined in the workflow.
- Document Q&A: Agents can save time when reviewing documents by using the predictive capabilities of generative AI to provide answers to predefined questions.
- Document and visual insights AI agent. The AI agent gathers context from user input and document or image attachments, generates the requested information based on the content, and provides the information along with any relevant task details. For more information on AI agents, see Document and visual insights AI agent.
- Contract metadata extraction and contract analysis. The skills for contract metadata extraction and contract analysis are only available with the Now Assist in Contract Management application. For more information, see Now Assist in Contract Management.
- Attachment summarization. Agents can view a summary of attachment content along with the record summary in ITSM. Attachment summarization is available in Now Assist for ITSM. For more information, see Customize a Now Assist for IT Service Management (ITSM) skill.
- Document chat. Agents can receive chat responses based on document content. Document chat is available in Now Assist for Virtual Agent. For more information, see Upload documents in a chat.
- Invoice data extraction. Extract invoice details from PDFs submitted through email or supplier portal and automatically process the digitized invoice within Accounts Payable Operations (APO). For more information, see Now Assist for Accounts Payable Operations (APO).
- ESG information extraction. Extract ESG-related information from invoices to support Operational Sustainability Management (ESG) initiatives and regulatory compliance. For more information, see Filter citations documents for Operational Sustainability Management (formerly ESG Management).

Now Assist in Document Intelligence skills

Now Assist in Document Intelligence skills are enabled in the Platform workflow on the Now Assist Admin Console. For more information on activating the skills, see Activate a Now Assist in Document Intelligence skill.

Extract information from documents

The extract information from documents skill allows you to use Now Assist predictions to extract fields and tables and find answers to predefined questions from a document or image. Agents can review and edit predictions in the Document Intelligence workspace. For more information, see Review extracted information in the Document Intelligence workspace.

Multimodal chat

The multimodal chat skill is used to enable chat responses about the content of uploaded documents and images. It is only used on the server side by the document and visual insights AI agent and by the question answering capability

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

in Now Assist for Virtual Agent. This skill doesn't require configuration in the Now Assist Admin console.

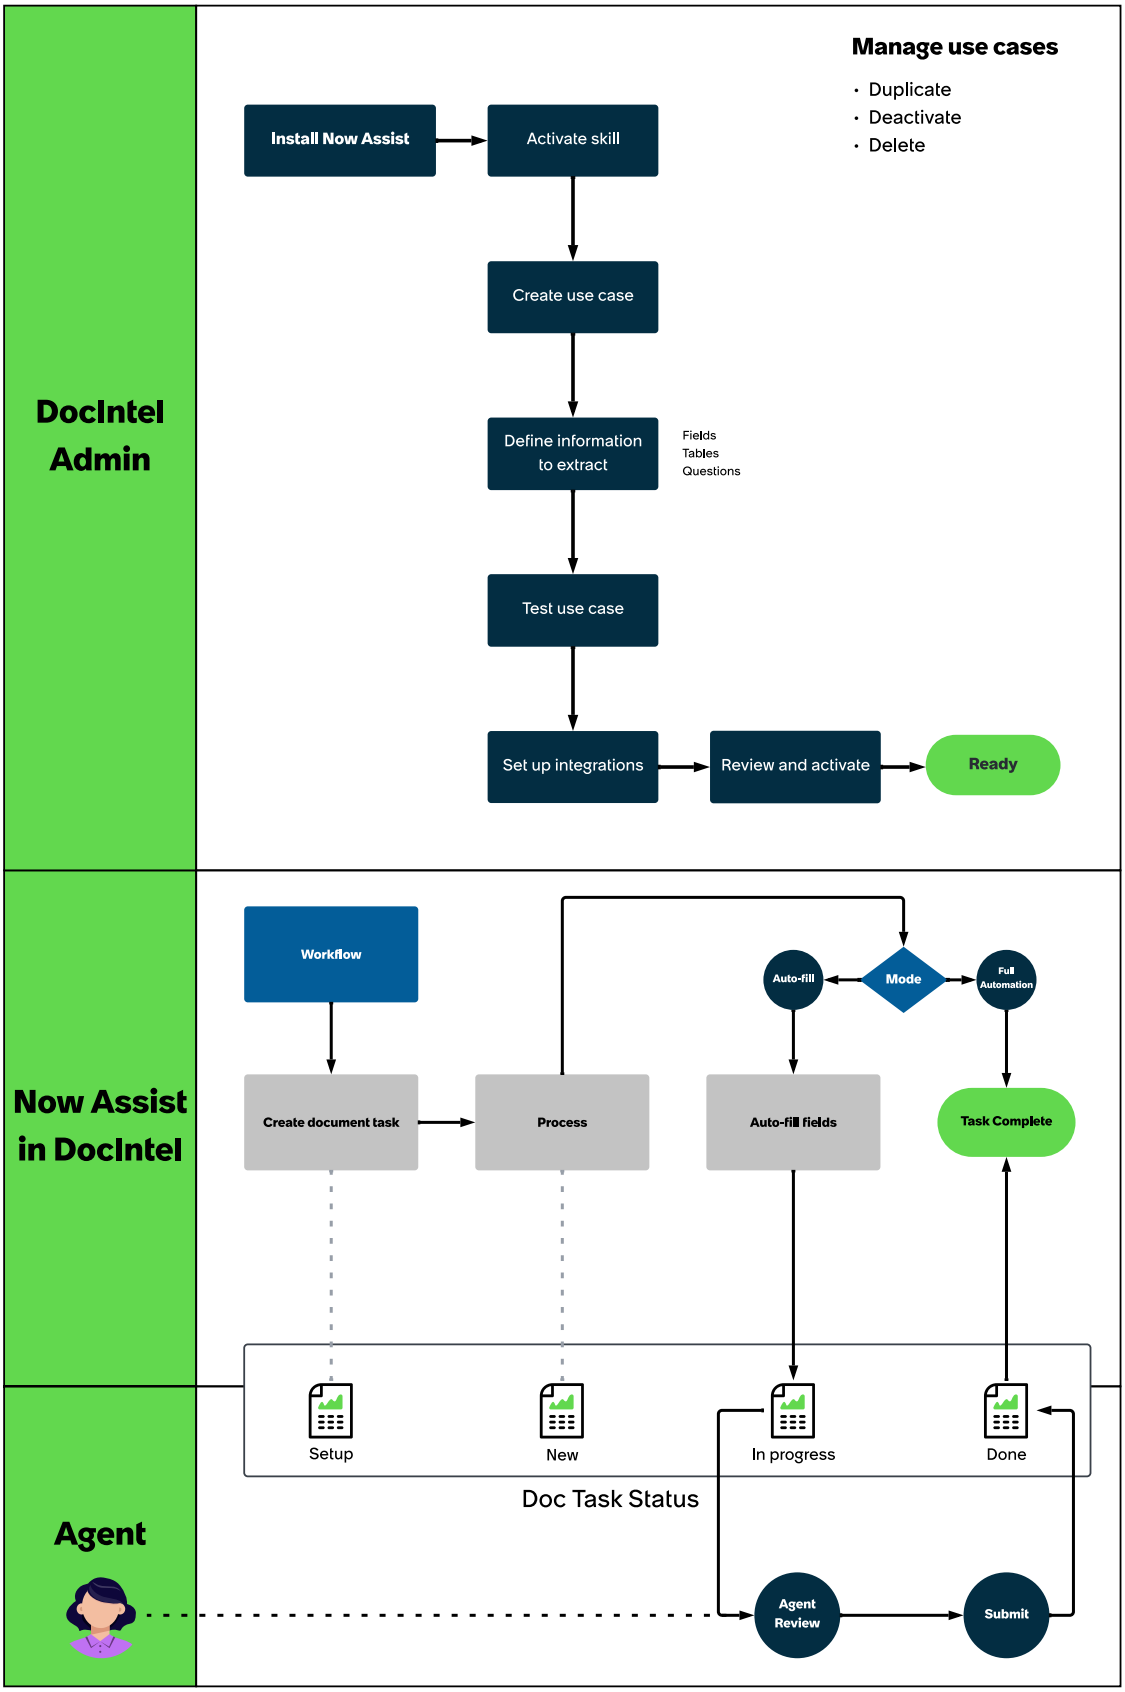
For more information on these capabilities, see Document and visual insights AI agent and Upload documents in a chat.

Now Assist in Document Intelligence workflow

The following diagram shows how the Now Assist in Document Intelligence skills are set up and used to process documents.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist in Document Intelligence flow



In this workflow:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- An admin activates a skill and sets up a use case for it.
- A workflow integration creates a document task as part of its flow.
- A document is uploaded for processing in a document task.
- Now Assist processes the document and makes predictions based on the fields defined in the use case.
- If the use case is not set to full automation, the task is sent to a live agent for review.
- The agent provides input to validate or correct the values predicted by Now Assist.
- The task is completed and the integrated workflow proceeds as defined.

Now Assist in Document Intelligence benefits

Now Assist in Document Intelligence provides the following benefits.

Benefit	Feature	User
Start fast with a guided set up of your use cases to identify the information you want to get from your documents.	Set up document intelligence use cases	DocIntel Admin [sn_docintel.admin] DocIntel Manager [sn_docintel.manager]
Accelerate the extraction of information from documents in the form of fields, tables, and answers to predefined questions, and turn it into structured data in the platform.	Review extracted information in the Document Intelligence workspace	DocIntel Extraction Agent [sn_docintel.extraction_agent]

Use cases in Now Assist in Document Intelligence

In Now Assist in Document Intelligence, a use case is used to define the information you want generative AI to detect from a document.

Use cases for Now Assist in Document Intelligence skills

Now Assist in Document Intelligence supports use cases for the following skills:

Extract information from documents

In an extract information use case, define the information you want Now Assist to extract from a document in the form of fields, tables, and questions.

You can define a specific field to extract, a table structure with multiple columns, or a question you want generative AI to answer based on the text in a document.

For example, if you need to process information from purchase order forms, you may want a use case with fields to extract the form date and number, a table to extract the itemized list with prices, and questions such as “Does the order include any items with ‘gift card’ in the description".

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use case structure

A use case is made up of the use case record and its related fields, tables, and questions as well as anyintegrations, and flows.

Now Assist in Document Intelligence uses fields to identify the specific information in documents to consider when making predictions. Tables are used to extract information in the cells that appear under table columns. Questions are used to produce an answer based on analysis of the text in a document.

Use case setup

Use the Now Assist Features interface to set up use cases for Now Assist in Document Intelligence skills. Use cases for the Document Intelligence application have a separate setup process. For more information, see Configuring Document Intelligence.

Use case setup involves defining the use case name and details, defining the fields, tables, and questions, setting up integrations and flows, and testing the use case.

For more information, see Set up a use case for Now Assist in Document Intelligence.

Tip: To save time when you need to create a new use case that shares a similar structure to another, make a copy of the existing use case and edit the details of the copy. For more information, see Make a copy of a use case in Now Assist in Document Intelligence.

Once you completed the setup of a use case, agents can begin processing documents for it.

Document Intelligence workspace with Now Assist

Use the Document Intelligence workspace to review the generative AI predictions made by Now Assist for a document or image file. You can also identify missing information in the file.

Overview of the Document Intelligence workspace with Now Assist

The Document Intelligence workspace provides document management features that enable you to quickly review and process documents and image files.

With the Document Intelligence workspace, you can:

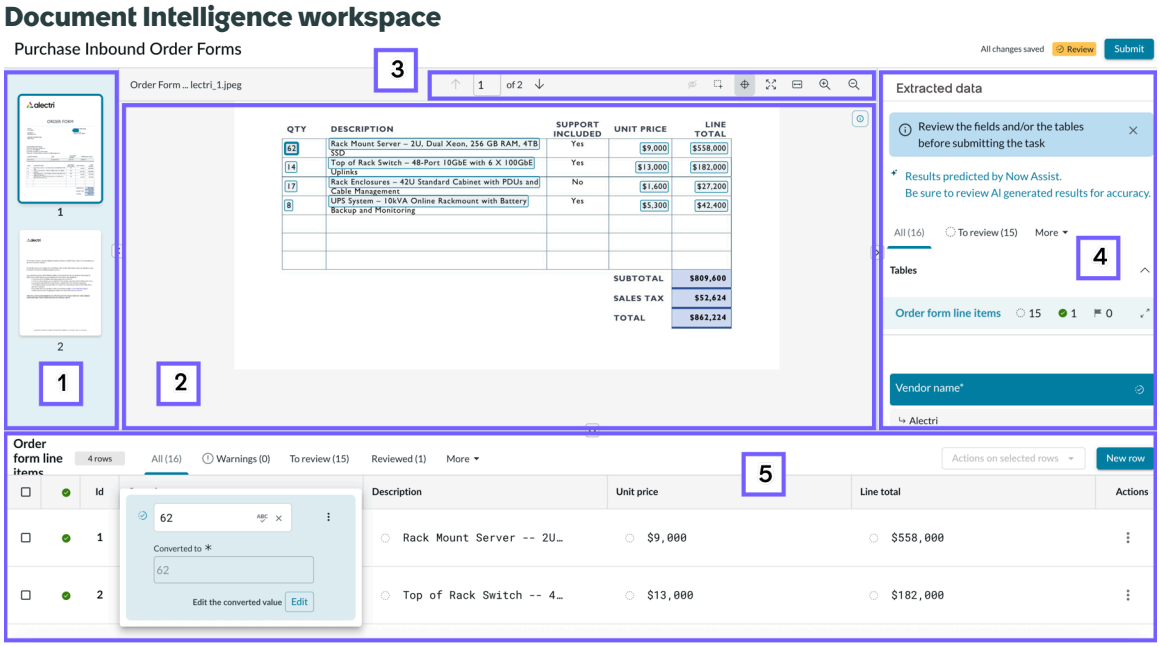
- Efficiently review the generative AI recommendations and confirm your file's extracted text.
- Review generative AI's answers to predefined questions to derive key information from your file.
- Identify missing information and review pending items.

Layout of the Document Intelligence workspace

The following illustration shows the Document Intelligence workspace. The workspace includes the following areas:

- 1 - Thumbnail panel
- 2 - Document image panel
- 3 - Document controls
- 4 - Extraction review panel
- 5 - Table panel

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



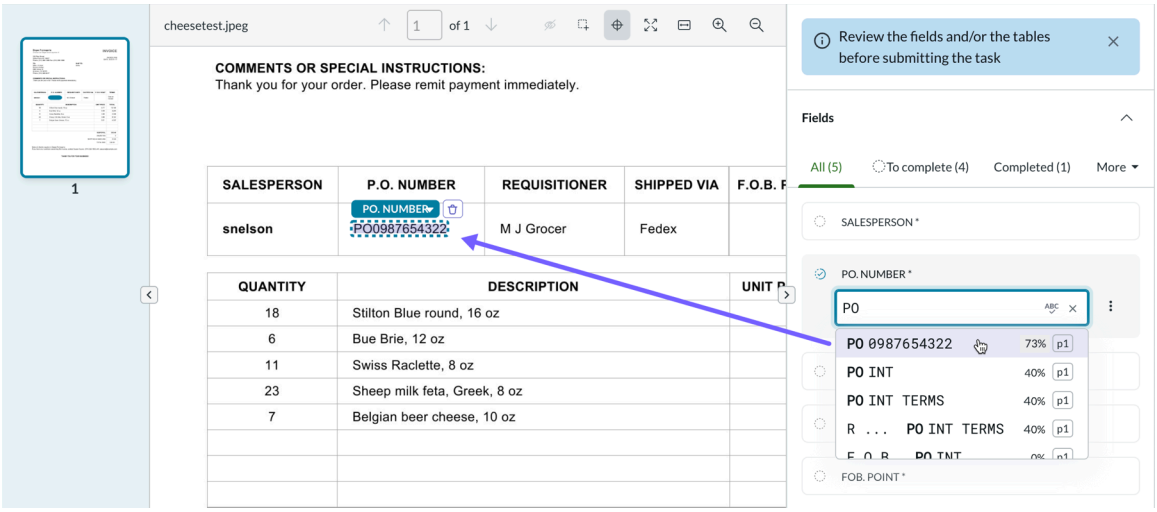
Thumbnail panel

In the thumbnail panel of the workspace, you can select one page from a multiple-page document. The selected page is displayed in the document image panel. Selecting a page does not affect what is displayed in the document fields panel.

Document image panel

The document image panel displays the page selected in the thumbnail panel.

If Now Assist can detect the source of its prediction in the document, it will highlight it in the document image panel when you select a field to review.



Document controls

When you are reviewing a document for extraction, you can use various controls to maximize the viewing area, zoom, or focus on the areas that you need.

You can also extract information directly from the document image using the draw tool. For more details, see Extract fields using the draw tool.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Extraction review panel

The extraction review panel enables you to open items for review, including viewing the generative AI predictions. You can also mark fields as missing in the document.

Note: An asterisk indicates a required field.

Table panel

The table panel enables you to open table rows for review, including viewing the generative AI recommendations. You can also flag fields or mark fields as missing in the document. Additional table controls allow you to insert rows and resize columns as needed.

Note: An asterisk in the column heading indicates a required field.

Supporting information for Now Assist in Document Intelligence

Get a quick overview of the important information that is related to the Now Assist in Document Intelligence application.

Supported versions

Now Assist in Document Intelligence is supported starting with Yokohama (Patch 1).

- Now Assist Admin Console v5.0
- Now Assist for Platform 7.0
- Document Intelligence v6.0
- Now Assist in Document Intelligence v2.0

Supported user interfaces

The Now Assist in Document Intelligence application is supported for the Document Intelligence workspace.

For more information, see Document Intelligence workspace with Now Assist.

Supported languages

Now Assist in Document Intelligence supports text in multiple languages. For information on supported languages, see Languages supported by Now Assist in Document Intelligence.

Application information

Now Assist in Document Intelligence skills are available in the Platform workflow.

Activate the Now Assist in Document Intelligence store app (sn_docintel_gen_ai) to configure the document extraction and Q&A skills.

This store app has the following dependencies:

- Now Assist Admin Console (sn_nowassist_admin)
- Now Assist for Platform (sn_genai_platform)
- Document Intelligence (sn_docintel)

For more information, see Configuring Now Assist in Document Intelligence.

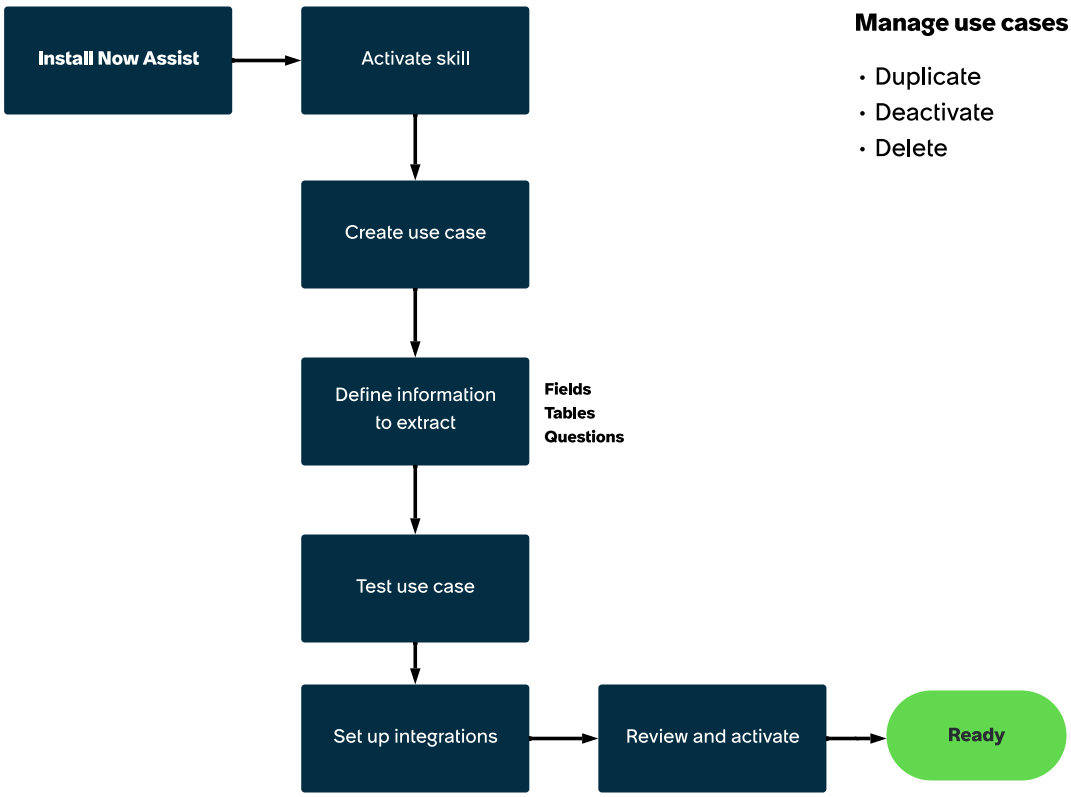
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configuring Now Assist in Document Intelligence

If you have the admin role, you can configure the Now Assist in Document Intelligence feature, enabling agents to use its generative AI skills within their application workspace.

The following diagram shows the process for setting up the extract information platform skill in Now Assist.

Now Assist in Document Intelligence configuration steps



In this workflow:

- Install Now Assist.
- Activate the extract information skill in the Platform workflow.
- Create one or more use cases for the skill.
- Define fields, tables, or questions for the use case.
- Test use case performance with a sample document.
- Set up integrations with your workflows.
- Review the use case and activate it.

Activate a Now Assist in Document Intelligence skill

Activate the Now Assist in Document Intelligence skills that agents can use to help analyze and extract information from documents with generative AI.

Before you begin

Now Assist in Document Intelligence skills are turned on by default. The skills will be automatically available to appropriate role users for the application. The new default behavior works as follows:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

New customers

When you install a Now Assist product, designated skills are turned on automatically.

Existing customers who are upgrading (starting with Australia Patch 1)

There is no change to skills that are currently enabled and customized. A skill is turned on if:

- The Now Assist plugin is installed, but the skill was never turned on.
- An admin has never adjusted roles for the skill.

A skill is not turned on if:

- The skill was previously turned on, and then turned off again.
- An admin has adjusted roles for the skill.

For more information, see Now Assist skills, agents, and agentic workflows on by default.

Before activating a document intelligence skill in Now Assist, the following applications and their respective plugins must be installed.

- Now Assist

  For more information, see Install Now Assist plugins.

- Document Intelligence

  For more information, see Install Document Intelligence.

- Now Assist in Document Intelligence

  For more information, see Install Now Assist plugins.

Role required: admin

About this task

Now Assist in Document Intelligence skills are now turned on by default. If a skill was previously turned on and then turned off, or if an administrator has adjusted roles for the skill, perform the steps in this task to reactivate the skill.

This task doesn't apply to the Multimodal chat skill. The multimodal chat skill is used to enable chat responses about the content of uploaded documents and images. It is only used on the server side by the document and visual insights AI agent and by the question answering capability in Now Assist in Virtual Agent. This skill doesn't require configuration in the Now Assist Admin console.

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Skills**.
2. In the workflow list, select **Platform** > **Other**.
3. In the list of Platform skills, select **Activate skill** for the document intelligence skill that you want to activate.

   For more information on document intelligence skills, see Exploring Now Assist in Document Intelligence.

   The guided use case setup opens.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. Create a use case.
   a. To create a new use case, follow the steps in Set up a use case for Now Assist in Document Intelligence.
   b. Select **Save and continue**.
5. Define access.
   a. **Optional:** Define a different access control list (ACL) for the skill.

      (Optional) An ACL enables you to restrict who is able to access and execute the skill to only users with the correct role.

      i. Select the **Edit** icon ( ) for **Roles**.
      ii. Select an option.

Option	Description
Any authenticated user	(Optional) As long as a user is logged in, they can access and execute the skill.
Select roles	(Optional) Select the roles that a user must have to execute the skill. **Note:** If you select multiple roles, a user must only have one of the roles to execute the skill.

      iii. Select **Apply**.
   b. Select **Save and continue**.
6. Review and activate.
   a. Review the configuration selections you've made for the skill.
   b. **Optional:** Select **Back** to return to a previous step and make a change.
   c. Select **Activate**.

Result
The skill is active and available to the selected user roles.

Set up a use case for Now Assist in Document Intelligence

Create a use case record to define a document you want to process with Now Assist in Document Intelligence.

Before you begin

- If the Now Assist in Document Intelligence skill has been deactivated, activate the skill. For more information, see Activate a Now Assist in Document Intelligence skill.
- Role required: Admin, DocIntel Admin, DocIntel Manager

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

In a use case, you define the information you want Now Assist to get from a document by specifying the type of document to process, the fields and tables to detect or questions you want to ask, and the location where document processing results are stored.

Once you have defined a use case, users can begin processing documents for it in the related workflows.

For more information on use cases, see Use cases in Now Assist in Document Intelligence.

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Skills**.
2. In the workflow list, select **Platform**.
3. In the Platform skills list, find the applicable document intelligence skill and select **Edit** in the options menu ( ).
4. Select **New use case**.

   The guided use case setup opens.
5. Define the use case.

   a. Enter a name for the use case.

   b. Select a target table to store the document processing results for this use case.

   c. Select the language of the files processed for the use case.

      If the files contain multiple languages, select the primary language.

      For more information, see Languages supported by Now Assist in Document Intelligence.

   d. Select the large language model (LLM) that will make predictions for the documents processed with this use case.

      For more information, see Large language models used by Now Assist in Document Intelligence.

   e. **Optional:** Turn on image mode to process images more efficiently.

      (Optional) Image mode sends pages to the LLM as images to leverage the visual capability of the multimodal LLM as well as any of the languages supported by it.

      Note: Selecting image mode reduces the page count limit to 10 pages per file.

      The image mode option is available when a multimodal LLM is selected.

      Note: Now LLM Service is a text-only model and doesn't support image mode.

   f. Select **Save and continue**.
6. Define information or questions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Select **Add a field**.

If you have already defined one or more fields, tables, or questions and you want to add another, select **New field**.

b. Select the type of information you want to get from the document.

You can choose one of the following:

Field

Fields are used to extract a single piece of information in the document. For example, a document number or a customer name.

Table

Tables are used to extract lists or tables of information. A table can have multiple columns. The number of list items or table rows doesn't have to be known in advance.

Question

Define the question you want to ask about the document.

A form displays based on the information type you selected.

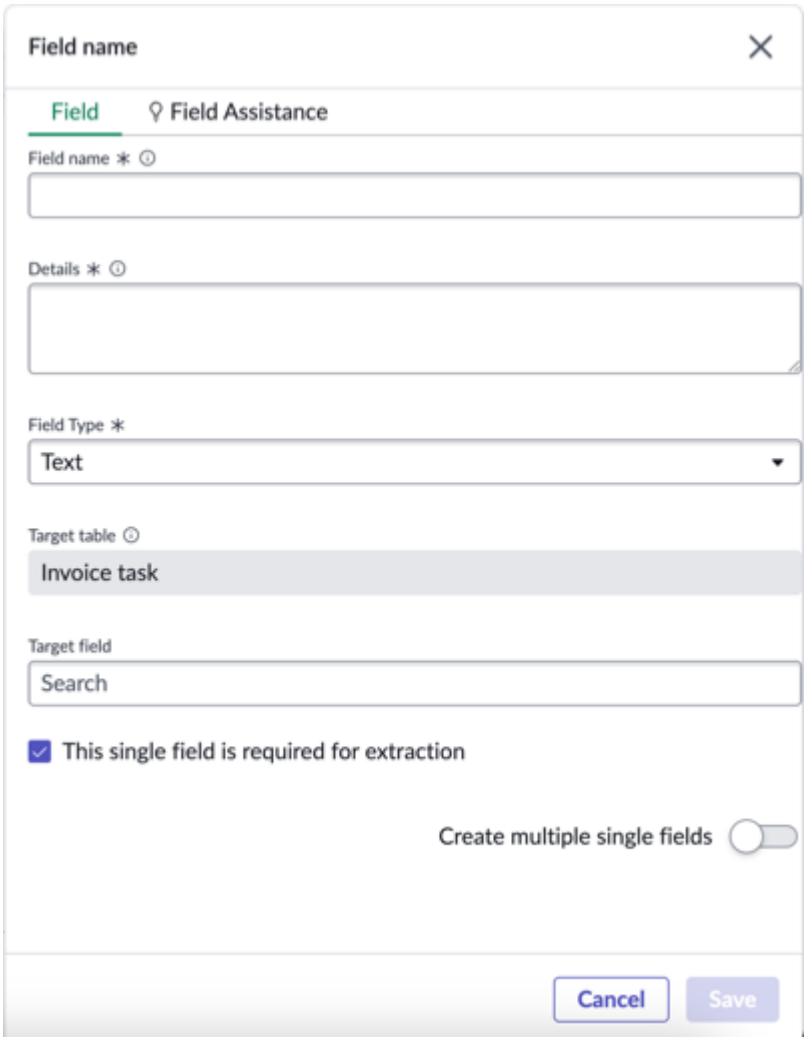
c. On the form, fill in the fields.

The type of form depends on the type of field.

- Question form for use case setup
- Field form for use case setup
- Table form for use case setup

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field form for document extraction



d. Select **Save**.
The system adds the new fields, tables, or questions to the Information list associated with the use case.

e. Optional: To edit or delete a field, table, or question, select the appropriate option in the options menu (⋮) of the row.

f. Select **Save and Continue**.

7. Test the use case performance with a sample document.
Select **Test a new document**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Select a document.

Option	Description
Upload from record	i. Select **Upload from record**. ii. Enter search criteria in the search field. iii. Select a record from the list. This option is available when a target table is selected for the use case.
Upload from this device	i. Select **Upload from this device**. ii. Select **Add file**. iii. Select a file and select **Open**. iv. Select **Upload**.

b. Select **Continue**.

Tip: Select the **Open in a new tab** button ( ) to view the document in a larger workspace on a separate browser tab.

The Document Intelligence workspace appears in a frame on the Test output screen.

c. Review the performance of the skill for the test document.

d. Select **Save and continue**.

8. Add integrations.

Integrate the document intelligence use case with a workflow.

This option is available when a target table is selected for the use case.

For more information on Document Intelligence integrations, see Integrate with a custom application or workflow.

a. Select **Add integration**.

If you have already defined one or more integrations and you want to add another, select **New integration**.

b. Enter a name for the integration.

c. Select the type of integration you want to use.

The `Process task` type creates an integration point to automatically create and process document tasks based on specific triggers happening in the target table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The `Extract values` type creates an integration point to automatically propagate the extracted values to the target table when extraction has been completed in Now Assist in Document Intelligence.

d. Use the conditions to select certain fields as specific triggers for the integration.

Conditions are available if you selected `Process task` in the previous step. For more information on conditions, see OR conditions.

e. Optional: Select the **Create Flow** option to create a flow for this integration in Workflow Studio.

Tip: This option should be selected, unless you are planning to write your own custom script to set up the integration. Be sure the integration is activated on Workflow Studio. For more information, see Building flows.

f. Select **Save**.

g. Select **Save and continue**.

9. Review and activate.

a. Review the configuration selections you've made for the use case.

b. Optional: Select **Back** to return to a previous step and make a change

c. Select **Complete setup**.

Turn on Full automation mode for a use case

Turn on Full automation mode to automatically complete and submit document tasks without an agent review. Full automation mode is turned off by default in use cases.

Before you begin

- Set up a use case for a Now Assist in Document Intelligence skill. For more information, see Set up a use case for Now Assist in Document Intelligence.
- Role required: sn_docintel.manager

About this task

The extraction mode determines how Now Assist in Document Intelligence processes document tasks for a use case. For more information, see Data extraction modes in Now Assist in Document Intelligence.

Turn on Full automation mode if you want Now Assist to bypass the agent review used to check the accuracy of the predicted values. Now Assist auto-fills the values for all required fields or marks them as missing in the document. Document tasks created for the use case are automatically completed and submitted by Now Assist.

Warning: Now Assist may not always produce accurate, complete, or appropriate information. By choosing to bypass the agent review, there is no way to check the accuracy of the predicted values before using the data in your workflows.

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Skills**.

2. In the workflow list, select **Platform**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. In the Platform skills list, find the applicable document intelligence skill and select **Edit** in the options menu ( ).
4. Select the use case you would like to configure.
5. Click the settings icon ( ).
6. On the Extraction mode screen, select the **Full automation mode (no agent review required)** option.
7. **Optional:** Select the **Any required field is missing in the document** option to turn off the automation and require an agent review when any of the required fields are missing in the document.
8. Close the Settings box.

Change the language models for a use case

Choose the language models for a Now Assist in Document Intelligence use case.

Before you begin

- Set up a use case for the extract information from documents skill. For more information, see Set up a use case for Now Assist in Document Intelligence.
- Role required: sn_docintel.manager

About this task

Language models are used to detect information in documents and make predictions for information extraction.

Third-party large language model (LLM) providers are available for Now Assist skills and AI agents in addition to Now LLM Service. For more information on LLMs in Now Assist, see Manage AI models.

For each Now Assist in Document Intelligence use case, only one LLM can be enabled at a time. Now Assist employs the selected LLM when processing documents for the use case.

For image files that need optical character recognition (OCR) to detect the text in them, OCR models are used to support different language groups. The language selected during use case setup helps the OCR model to detect text in the images.

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Skills**.
2. In the workflow list, select **Platform**.
3. In the Platform skills list, find the applicable document intelligence skill and select **Edit** in the options menu ( ).
4. Select the use case you would like to configure.
5. Select the settings icon ( ).
6. Select the Manage LLMs screen.
7. Select the LLM that will make predictions for the documents processed with this use case.

   For more information, see Large language models used by Now Assist in Document Intelligence.

8. **Optional:** Turn on image mode to process images more efficiently.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

(Optional) Image mode sends pages to the LLM as images to leverage the visual capability of the multimodal LLM as well as any of the languages supported by it.

The image mode option is available when a multimodal LLM is selected.

Note: Selecting image mode reduces the page count limit to 10 pages per file.

9. Select the **Language of the files** processed for the use case.

   If the files contain multiple languages, select the primary language.

   For more information, see Languages supported by Now Assist in Document Intelligence.

10. Select **Save**.

Result
The selected languages are enabled for the use case.

Edit a use case in Now Assist in Document Intelligence

Edit a use case to change the name, fields, tables, questions, and integrations.

Before you begin
Role required: Admin, DocIntel Admin, DocIntel Manager

About this task
Follow these steps to edit a use case along with its fields, tables, questions, integrations, and flows.

Use cases marked as read-only can't be edited. However, you can make a copy of a read-only use case that is editable. For more information, see Make a copy of a use case in Now Assist in Document Intelligence.

Procedure
1. Navigate to **All** > **Now Assist Admin** > **Skills**.
2. In the workflow list, select **Platform**.
3. In the Platform skills list, find the applicable document intelligence skill and select **Edit** in the options menu ( ).
4. In the row of the use case you want to edit, select **Edit** in the options menu ( ).
5. Make changes as needed.

   For more information on use case setup, see Set up a use case for Now Assist in Document Intelligence

Make a copy of a use case in Now Assist in Document Intelligence

Make a copy of a use case to save time when you need to create a new use case with a similar structure.

Before you begin
Role required: Admin, DocIntel Admin, DocIntel Manager

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task
Follow these steps to create a copy of a use case along with its fields, tables, questions, integrations, and flows.

Procedure
1. Navigate to **All** > **Now Assist Admin** > **Skills**.
2. In the workflow list, select **Platform**.
3. In the Platform skills list, find the applicable document intelligence skill and select **Edit** in the options menu ( ).
4. In the row of the use case you want to copy, select **Make a copy** in the options menu ( ).
5. In the confirmation box, select **Make a copy**.
6. Select **Continue**.
7. Enter a name for the use case.
8. Select **Make a copy**.

Result
The duplicated use case appears in the use cases list.

What to do next
Edit the new use case to make any necessary changes and test it to make sure it functions properly.

For more information, see Edit a use case in Now Assist in Document Intelligence.

Deactivate a use case in Now Assist in Document Intelligence

Deactivate a use case that you don't want to use for your documents.

Before you begin
Role required: Admin, DocIntel Admin, DocIntel Manager

Procedure
1. Navigate to **All** > **Now Assist Admin** > **Skills**.
2. In the workflow list, select **Platform**.
3. In the Platform skills list, find the applicable document intelligence skill and select **Edit** in the options menu ( ).
4. In the row of the use case you want to deactivate, select **Deactivate** in the options menu ( ).
5. In the confirmation box, select **Yes**.

Result
The use case is deactivated. The deactivated use case will not be used for document processing.

Delete a use case in Now Assist in Document Intelligence

Delete a use case when it is no longer needed for your documents.

Before you begin
Role required: Admin, DocIntel Admin, DocIntel Manager

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Skills**.
2. In the workflow list, select **Platform**.
3. In the Platform skills list, find the applicable document intelligence skill and select **Edit** in the options menu ( ).
4. In the row of the use case you want to delete, select **Delete** in the options menu ( ).
5. In the confirmation box, select **Delete**.

Result

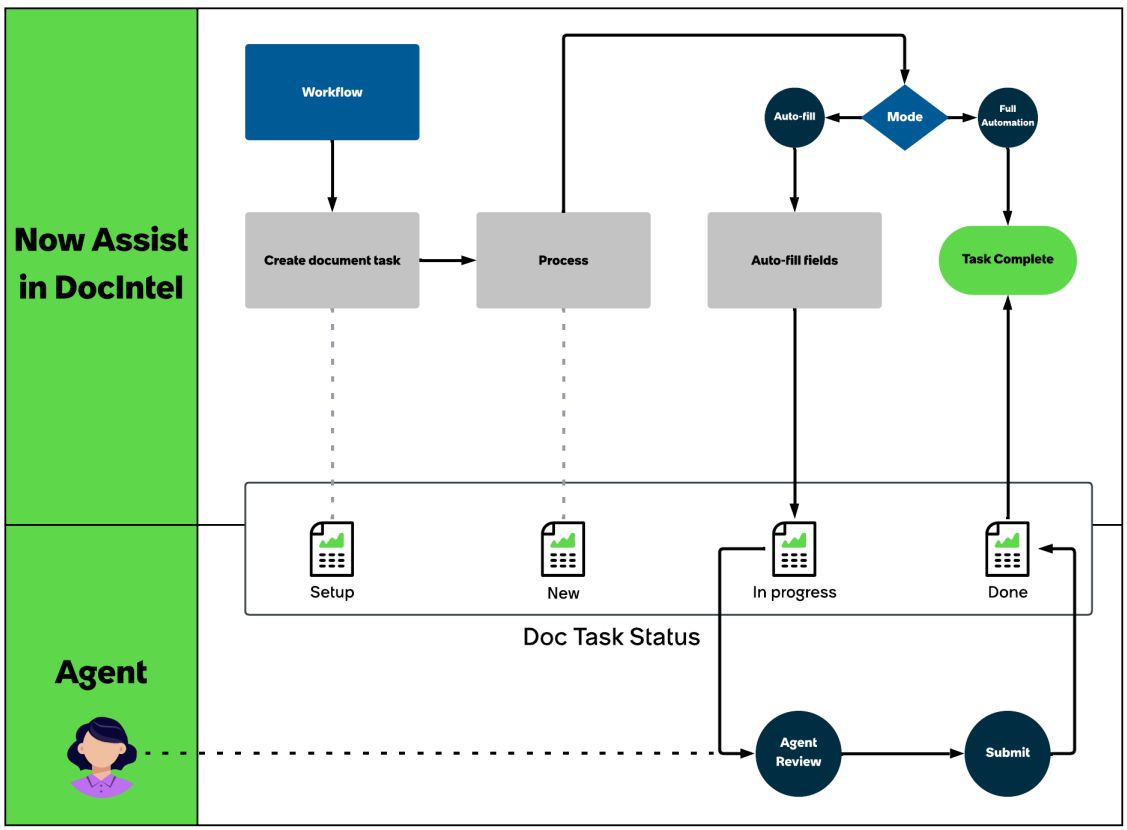
The use case along with the related fields, tables, questions, integrations, and flows are deleted.

Using Now Assist in Document Intelligence

If you have an agent role, you can use the Now Assist in Document Intelligence workspace to analyze and extract information from your documents with generative AI.

The following diagram shows how the Now Assist in Document Intelligence skills are set up and used to process documents.

Now Assist in Document Intelligence document processing workflow



In this workflow:

- A workflow integration creates a document task as part of its flow.
- A document is uploaded for processing in a document task.
- Now Assist processes the document and makes predictions based on the fields defined in the use case.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If the use case is not set to full automation, the task is sent to a live agent for review.
- The agent provides input to validate or correct the values predicted by Now Assist.
- The task is completed and the integrated workflow proceeds as defined.

Review extracted information in the Document Intelligence workspace

Use the Document Intelligence workspace to review the information that was extracted from a document or image file by Now Assist in Document Intelligence.

Before you begin

Role required: sn_docintel.admin, sn_docintel.creation_agent, or sn_docintel.extraction_agent

About this task

Now Assist uses generative AI capabilities to make predictions on the information to extract from a file for the fields, tables, and questions defined in the use case. The Document Intelligence workspace provides document management features that enable you to quickly review and process the extracted information.

With the Document Intelligence workspace, you can:

- Review the Now Assist predictions and confirm or edit the extracted text.
- Review the answers generated by Now Assist and edit them for accuracy.
- Identify fields or answers as missing in the document.

For more information, see Document Intelligence workspace with Now Assist.

Procedure

1. Open the Document Intelligence workspace to view the document and the extracted information.
2. In the extraction review panel, select a field or question.

   The item's review status is set to In progress.

   If Now Assist can detect the source of its prediction in the document, it will highlight it in the document image panel.

   Collapse a side panel using the Collapse icon ( ), if needed.
3. Review the field or question for accuracy.

   Note: The field value or answered question is a generative AI prediction based on the document text or image. Review it carefully to ensure that it is accurate.
4. **Optional:** Take the appropriate action on the field value or answered question.

   (Optional) For information on using the draw tool, see Extract fields using the draw tool.

Option	Description
Edit the field value or answer	To make a change: **a.** Select the **Edit** icon ( ) in the field or question. **b.** Select the **Edit** icon again.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Option	Description
	c. Edit or type a new value in the text box. **d.** Select **Save**.
Mark the field value or answer as missing in the document	If you can't find an appropriate match, mark it as missing in the document. **a.** Select the **Edit** icon ( ) in the field or question. **b.** Select the **Edit** icon again. **c.** Select **Answer is missing in the document**. **d.** Select **Save**.
Mark the field value or answer as reviewed	Confirm that you've reviewed all fields and questions for accuracy before you submit. **a.** Select the **Edit** icon ( ) in the field or question. **b.** Select **Mark as reviewed**. **c.** [Optional] If needed, select the **Reviewed** icon ( ) to mark the question as unre viewed.

5. Select the next field or question and repeat steps 3 to 4.
6. Select **Submit**.

> **Note:** Generative AI may not always produce accurate, complete, or appropriate information. It is important to review all field values and answers for accuracy before submitting.

Result
The extracted values are saved. Empty fields and answers are submitted with no values.

Now Assist in Document Intelligence reference

The following topics provide additional information about the features and properties installed with Now Assist in Document Intelligence.

Components installed with Now Assist in Document Intelligence

Components are installed with the activation of Now Assist in Document Intelligence.

Roles installed

There are no roles installed with Now Assist in Document Intelligence. It uses the same roles included with the Document Intelligence application.

For more information, see Components installed with Document Intelligence.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tables installed

Now Assist in Document Intelligence uses the same tables included with the Document Intelligence application.

For more information, see Components installed with Document Intelligence.

Data extraction modes in Now Assist in Document Intelligence

The extraction mode determines how Now Assist in Document Intelligence processes a document task.

Now Assist in Document Intelligence uses the following data extraction modes.

Extraction mode	Description
Default	Now Assist auto-fills each field in the Document Intelligence workspace if it recognizes a value for it in the document. Now Assist may not always produce accurate, complete, or appropriate information. All fields should be reviewed for accuracy. A reviewer can accept the predicted value or manually enter a value.
Full automation	Now Assist automatically completes and submits the document task without an agent review. The values for all required fields are auto-filled by Now Assist or are identified as missing in the document. By turning on Full automation mode, the agent review used to check the accuracy of the predicted values is bypassed. There is a configuration option to suspend the automation and require an agent review when any of the required fields are missing in the document.

Document and visual insights AI agent

The document and visual insights AI agent gathers context from user input and document or image attachments, generates the requested information based on the content, and provides the information along with any relevant task details.

Document and visual insights AI agent overview

The document and visual insights AI agent performs tasks to process documents and images.

Data extraction

Extract specific information.

Summarization

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Summarize the key topics of documents and images.

Question answering (Q&A)

Answer questions about the document or image content.

The document and visual insights AI agent is not typically used in standalone mode and any use case can access it. For more information on AI agents, see Now Assist AI agents.

AI agents use role masking to determine which users can access them. Ones installed with Now Assist applications have specific roles that come included with the application. If you select *`Users with specific roles`* for user access, you must configure the security controls to include these roles. For the instructions to change the security controls, see Define security controls for an AI agent.

AI agent actions

When used, the AI agent may attempt the following actions:

- Determine the type of task to perform based on information provided by the user.
  - Data extraction
  - Summarization
  - Question answering (Q&A)
- Retrieve the relevant record details.
- Initiate the process based on the type of task.
- Provide the task results as directed.
- Prompt the user to upload a file for processing.
- Display the results of the AI agent actions in a document view screen.
- Provide citation sources.
- Notify the user of task completion and next steps.

Document Intelligence tool for Now Assist Skill Kit

Use the Document Intelligence tool to leverage extraction, question answering, and summarization capabilities for a skill created with Now Assist Skill Kit.

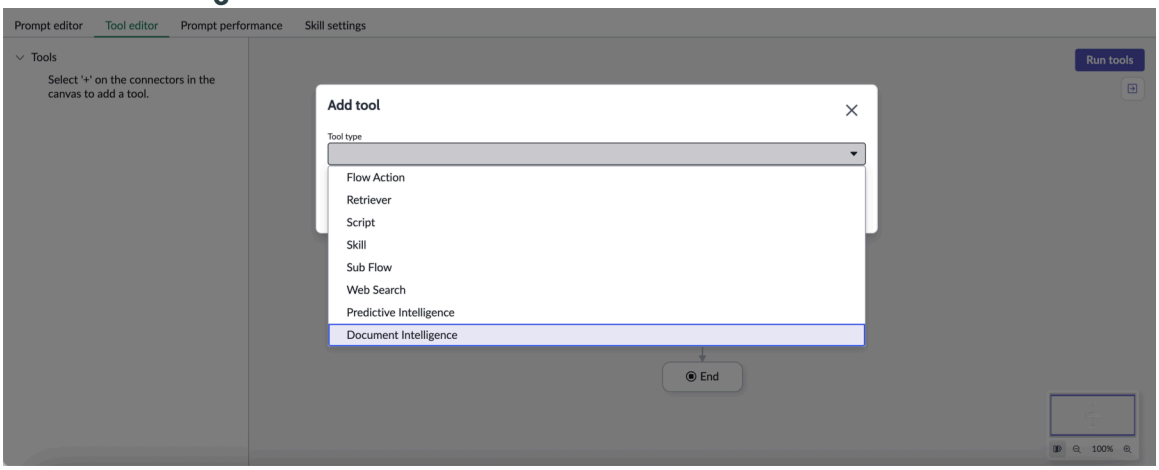
Document Intelligence tool overview

A tool is a utility in Now Assist Skill Kit that is configured to convert skill inputs into skill outputs. The Document Intelligence tool uses generative AI capabilities to extract information and insights from document and image files. The extracted information is turned into structured data that can be used in the platform.

You can use the tool editor to configure the Document Intelligence tool.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence tool in the Tool Editor



For more information on adding the Document Intelligence tool to a custom skill in Now Assist Skill Kit, see Add a tool.

Document Intelligence tool actions

The Document Intelligence tool provides the following capabilities for a custom skill to process documents and images.

Ask a question

Define a question to ask the large language model (LLM) about the content of a file attachment.

Extract information

Pull structured data and insights from file content.

Extract values from the files based on the fields, tables, and questions defined in a Now Assist in Document Intelligence use case. Use an existing use case or create a new one that meets the needs of your custom skill.

Summarize files

Generate a concise summary the key topics of the file content.

For more information on creating custom Now Assist prompts and skills, see Now Assist Skill Kit.

Field types in Now Assist in Document Intelligence

The field type specifies the information that is retrieved from a document with Now Assist in Document Intelligence.

The following field types are available to administrators when configuring fields for use cases.

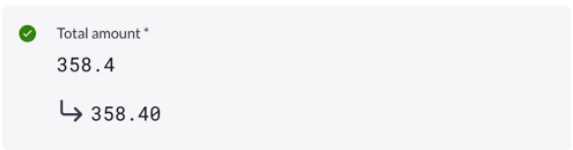
Note: Some field types convert the extracted value into a standard format. For more information, see Data normalization.

Now Assist in Document Intelligence field types

Field type	Description
Boolean (True/False)	Holds a true/false value. In document Q&A questions, the value is displayed as a yes or no answer. The answer can be followed by

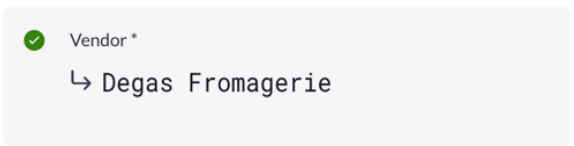
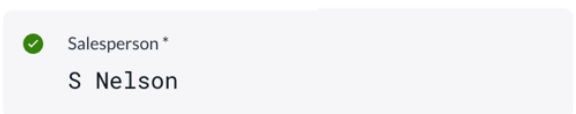
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist in Document Intelligence field types (continued)

Field type	Description
	an explanation generated by Now Assist, if enabled by the admin. This field type is available for questions defined in the use case setup. 
Date	Date, displayed in the date format extracted from the document. This field type is available for fields and tables defined in the use case setup. 
Decimal	A number with up to two decimal places (for example, 12.5 or 12.55). This field type is available for fields and tables defined in the use case setup. 
Float (floating point number)	A number with up to seven decimal places (for example, 12.0 to 12.0000000). This field type is available for fields and tables defined in the use case setup.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist in Document Intelligence field types (continued)

Field type	Description
	
Integer	Whole number (for example, 12). This field type is available for fields and tables defined in the use case setup. 
Reference	A reference field stores a reference to a field on another table. For example, the Caller field on the Incident table is a reference to the User [sys_user] table. For more information, see Reference field type. This field type is available for fields and tables defined in the use case setup. 
Text	Text field. This field type is available for fields, tables, and questions defined in the use case setup. 

Now Assist in Document Intelligence forms

Use forms to view and update Now Assist in Document Intelligence information.

Question form for use case setup

The Question form enables you to define a question you want to ask about the document.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Question form includes the following fields.

Single field form

Field	Description
Question	The question about the document. This question is sent to the generative AI for a prediction. Write a clear and concise question about the information you want to find in the document. For example, "Is this document subject to expiration?". Below the question, add details about how generative AI can best provide the correct answer. For example, “Review the sections where validity, duration or time frames are discussed and look for any details about expiration dates, renewal requirements, or explicit statements declaring the document as permanent. Answer ‘Yes’ if the document is subject to an expiration date. Answer ‘No’ if there is no mention of expiration or contains an explicit statement that it is permanent.” **Tip:** View the **Question assistance** tab for additional guidance on forming an effective question. This question is displayed to the agent when reviewing the answers in the Document Intelligence workspace.
Field Type	The type of field (for example, a text or Boolean field). For more information, see Field types in Now Assist in Document Intelligence.
Provide an explanation for the answer	Option to have the generative AI provide an explanation based on the document text that supports a yes or no answer. This field is available when the **Field Type** field is set to `Boolean (True/False)`.
Target table	The table that stores the document processing results for this use case.
Target field	Field on the target table you want to align this field with.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Single field form (continued)

Field	Description
	The use case must have a target table selected.
This single field is required for extraction	Option to make a question required. Required questions can't be left empty or unreviewed.
Create multiple single fields	Option to keep the form displayed on the screen. Enable this if you are adding more than one question to the use case.

Field form for use case setup

The Field form enables you to define a single field for extraction.

The Field form includes the following fields.

Single field form

Field	Description
Field name	The name for the field as it appears in the Document Intelligence workspace. Clearly state the information you want to extract. Be specific, clear and concise. For example: "Order number"
Details	A description of the information you want to extract. Along with the field name, the description is used to help the LLM predict the text to extract from the document. Include any relevant context or additional details to identify the right information. This might include keywords. For example, "The contract number or the number of the reference contract". Provide examples of what information you want to extract. For example, "AGR-2023-0042", "CON2039739", or "BV-22122KEY". Example of a description: Field name: "Currency" Details: "The currency in which the contract is denominated. Only valid for Order Forms. Otherwise, leave it empty ('). If the currency

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Single field form (continued)

Field	Description
	symbol is '$', answer 'USD'. Examples: 'USD', 'EUROS', 'GBP', etc.” **Tip:** View the **Field assistance** tab for additional guidance on forming an effective field.
Field Type	The type of field. For example, a text or date field. For more information, see Field types in Now Assist in Document Intelligence. Some field types convert the extracted value into a standard format. For more information, see Data normalization.
Target table	The table that stores the document processing results for this use case.
Target field	Field on the target table you want to align this field with. The use case must have a target table selected.
This single field is required for extraction	Option to make a field required. Required fields can’t be left empty or unreviewed.
Create multiple single fields	Option to keep the form displayed on the screen. Enable this if you are adding more than one single field to the use case.

Table form for use case setup

The Table form enables you to define a table for extraction.

The Table form includes the following fields.

Table form

Field	Description
Table	
Table name	The name for the table as it appears in the Document Intelligence workspace.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table form (continued)

Field	Description
Additional Details	A description of the table information you want to extract. Along with the table name, the description is used to help the LLM predict the table fields to extract from the document. Include any relevant context or additional details to identify the right information.
Target table	The table that stores the document processing results for these table fields.
Parent mapping to field	Field on the target table you want to align this table with. You must first select a target table.
This table is required for extraction	Option to make the table fields required. Required table fields can't be left empty or unreviewed.
Columns	
Column title	Name of the column header in the table.
Column type	The type of field in the table column. For example, a text or date field. For more information, see Field types in Now Assist in Document Intelligence. Some field types convert the extracted value into a standard format. For more information, see Data normalization.
Select target field	Field on the target table you want to align this field with. The use case must have a target table selected.
New column	Option to add a column to the table.
Create multiple tables	Option to keep the form displayed on the screen. Enable this option if you're adding more than one table to the use case.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Limitations in Now Assist in Document Intelligence

There are several important limitations to be aware of when you're using Now Assist in Document Intelligence.

Now Assist in Document Intelligence skill limits

The following table lists the important limitations for the skills in Now Assist in Document Intelligence.

Skill limits for Now Assist in Document Intelligence

Limit	Description
File formats	The supported file formats are JPEG, PNG, PDF, and DOCX. **Note:** Encrypted files are not supported.
File size limit	The file size limit is 20 MB.
Page count limit	If image mode is turned on for the use case, the page count limit is 10 pages per file. If image mode is turned off, the page count limit is: • 200 pages per file if no tables are defined for the use case. • 20 pages per file if a table is defined for the use case. Image mode is selected during use case setup. For more information, see Set up a use case for Now Assist in Document Intelligence.
Maximum number of fields per use case	The maximum number of fields per use case is 50.
Supported languages	For image files that need optical character recognition (OCR) to detect the text in them, there are OCR models to support different language groups. For more information, see Languages supported by Now Assist in Document Intelligence. For text-based files, the skill recognizes any language supported by the selected or default model, as described in the model card for the LLM. For more information on LLMs, see Large language models used by Now Assist in Document Intelligence.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Limits for the document and visual insights AI agent

The following table lists the important limitations for the document and visual insights AI agent.

Document and visual insights AI agent limits

Limit	Description
File formats	The supported file formats are JPEG, PNG, PDF, DOCX, XLSX, CSV, TXT, and PPTX.
File size limit	The file size limit is 20 MB.
Page count limit	For each document task, the page count limit is 200 pages per file. If the AI agent performs the extraction based on a use case that includes a table, the limit decreases to 20 pages per file.
Supported languages	For image and text-based files, the AI agent recognizes any language supported by the selected or default model, as described in the model card for the LLM. For more information on LLMs, see Large language models used by Now Assist in Document Intelligence.

Large language models used by Now Assist in Document Intelligence

Now Assist in Document Intelligence uses large language models (LLMs) to perform generative AI and agentic AI capabilities.

You can select the LLM when creating a use case. The LLM can also be selected for the AI agent to use when processing documents and images.

Multimodal LLMs can support image mode, as they are capable of processing multiple types of inputs like text, sound, images, and more, to generate a text response. Processing may take longer for image inputs.

Note: The document and visual insights AI agent only uses multimodal LLMs. The text-only Now LLM Service is not a supported model for the AI agent.

The following table lists the available LLMs for Now Assist in Document Intelligence.

LLM options for Now Assist in Document Intelligence

LLM	Highlights
Now Assist default	Use the LLM selected as the default for all Now Assist skills. For more information, see Manage AI models.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

LLM options for Now Assist in Document Intelligence (continued)

LLM	Highlights
Now LLM Service - Large	Text-only model used for AI-driven solutions to support natural language understanding, automation, and decision support.
Now LLM Service - Small	Text-only model used for enterprise AI applications by enhancing text-based automation and content generation within ServiceNow workflows.
Google Cloud - Gemini Large	Multimodal model with advanced reasoning and problem-solving capabilities.
Google Cloud - Gemini Small	Multimodal model with strong performance in summarization, rewriting, and content transformation.
Azure OpenAI - GPT Large	Multimodal model with a strong performance on tasks involving analysis, summarization, and multi-step problem solving.
Azure OpenAI - GPT Small	Multimodal model that is lightweight, high-efficiency, and effective for routine workflows with summarization, paraphrasing, and Q&A.
Amazon Bedrock - Claude Large	Multimodal model with strong context management for long documents and dialogs.
Amazon Bedrock - Claude Small	Multimodal model with lower latency and higher efficiency for real-time applications.

For more information, see Large language models on the ServiceNow AI Platform.

Languages supported by Now Assist in Document Intelligence

Now Assist in Document Intelligence supports text in multiple languages.

For text-based files, the Now Assist recognizes any language supported by the selected or default model, as described in the model card for the LLM. For more information on LLMs, see Large language models used by Now Assist in Document Intelligence.

For image files that need optical character recognition (OCR) to detect the text in them, OCR models are used to support different language groups. The language selected during use case setup helps the OCR model to detect text in the images. For more information, see Set up a use case for Now Assist in Document Intelligence.

The following table lists the available OCR models.

OCR models for Now Assist in Document Intelligence

OCR model	Languages supported
Latin model (Default)	English, Dutch, French, German, Italian, Portuguese, and Spanish
CJ model	Japanese and Chinese (simplified)

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence

Document Intelligence (DocIntel) is an AI solution that enables any organization to automate and accelerate the process of extracting data from documents. That data can easily be integrated into larger automation workflows to save time and resources.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Request apps on the Store

Visit the ServiceNow Store 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.

Get started

Explore  Learn about DocIntel concepts and features	**Configure**  Set up DocIntel to process different types of documents	**Integrate**  Integrate your DocIntel application with other applications
Use  Use the DocIntel application to process documents by detecting and analyzing text in documents	**Monitor**  Track DocIntel usage and effectiveness	**Reference**  Get details about the properties that are installed with DocIntel

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

	Now Assist in Document Intelligence  Use generative AI to analyze and extract information from documents	

Troubleshoot and get help

- Learn more about what's new and changed in the Document Intelligence release notes.
- Additional resources for AI products and solutions.
- ServiceNow Community on AI and Intelligence.
- Access real-time courses, self-paced training, and career resources at ServiceNow University.
- Search the Known Error Portal for known error articles.
- Contact Customer Service and Support.

Exploring Document Intelligence

Document Intelligence helps you to quickly and accurately classify and extract information from documents using artificial intelligence (AI).

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Document Intelligence overview

Document Intelligence uses optical character recognition (OCR) together with AI models to detect and extract information from documents with various structures and formatting. This enables you to accurately extract information and automate document processing.

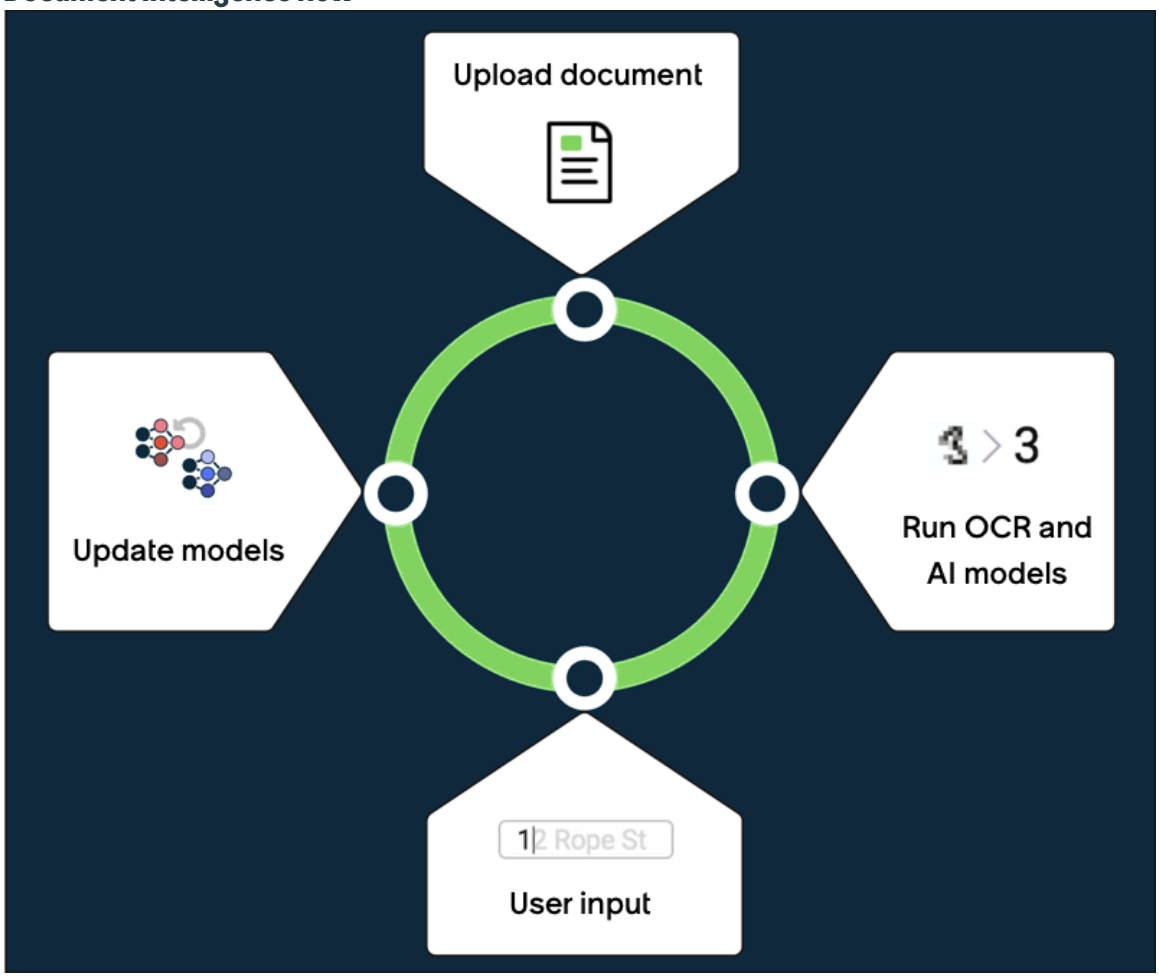
Document Intelligence workflow

With Document Intelligence (DocIntel) you can process single or multi-page documents in JPEG, PNG, or PDF formats. You can process documents that contain typed text such as forms, invoices, identity documents, and more.

The following diagram shows how document extraction works in Document Intelligence.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence flow



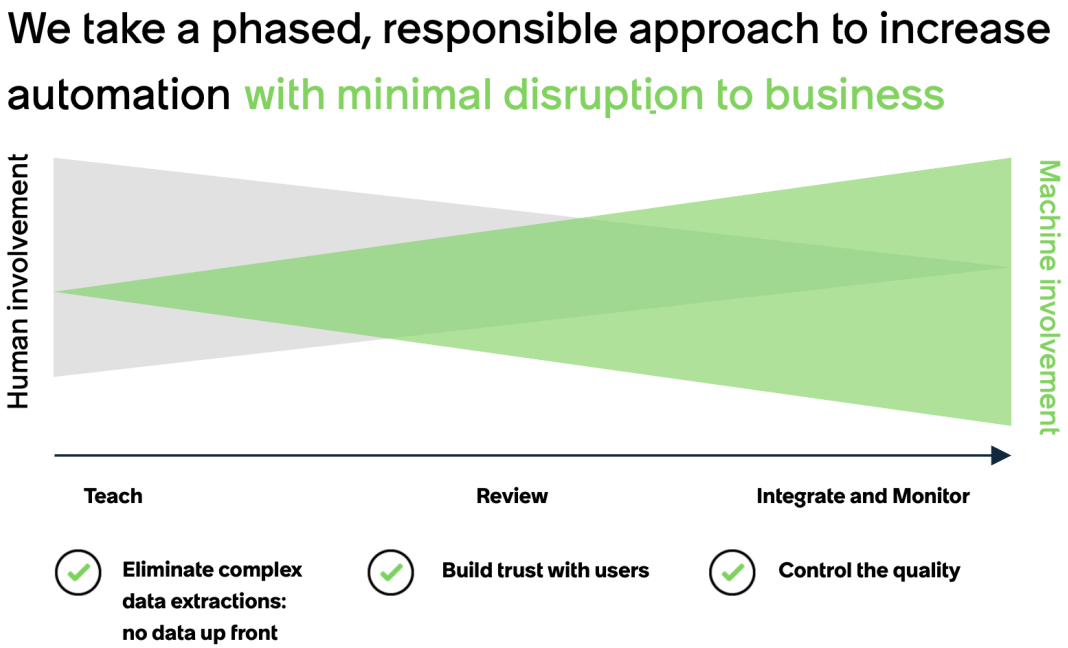
In this workflow:

1. A document is uploaded for processing in a document task.
2. DocIntel extracts the data from the document using OCR and AI models.
3. The user provides input to validate or correct the DocIntel recommendations.
4. The models are updated and trained to provide more accurate results.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence benefits

Benefits of Document Intelligence



Document Intelligence provides the following benefits.

Benefit	Feature	User
Start fast with a no-code set-up that enables data extraction from many document types including PDF and scanned paper documents.	Set up document extraction use cases	DocIntel Admin [sn_docintel.admin] DocIntel Manager [sn_docintel.manager]
Enable categorization for any type of document you define.	Set up document classification use cases	DocIntel Admin [sn_docintel.admin] DocIntel Manager [sn_docintel.manager]
Automate intelligently with responsible, feedback-driven AI for continual learning.	Configure data extraction modes	DocIntel Admin [sn_docintel.admin] DocIntel Manager [sn_docintel.manager]
Seamlessly integrate document processing steps into workflows.	Integrating Document Intelligence with other applications	DocIntel Admin [sn_docintel.admin] DocIntel Manager [sn_docintel.manager]

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Benefit	Feature	User
Accelerate extraction of structured and semi-structured documents such as forms, invoices, IDs, and more.	Extract fields using the Document Intelligence workspace	DocIntel Creation Agent [sn_docintel.creation_agent] DocIntel Extraction Agent [sn_docintel.extraction_agent]
Accelerate classification of single and multi-page documents.	Classify documents using the Document Intelligence workspace	DocIntel Creation Agent [sn_docintel.creation_agent] DocIntel Extraction Agent [sn_docintel.extraction_agent]

Configuring Document Intelligence

Activate Document Intelligence on your instance and get started with basic configuration.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

1. Install and configure Document Intelligence.
   - Set up Document Intelligence.

     Review important information before you start setting up Document Intelligence.
   - Install Document Intelligence.

     For the best configuration experience for document extraction use cases, install the Document Intelligence Admin application.
   - Upgrade to Document Intelligence 3.0 or later from version 2.4 or earlier.

     Document Intelligence 3.0 or later includes an updated database schema to support its transition from a scoped application to a ServiceNow AI Platform plugin.
   - Configure Document Intelligence settings.

     Use general settings to control how Document Intelligence displays information to users.
2. Configure Document Intelligence use cases.

   Set up Document Intelligence to process different types of documents by telling it what data you want to extract and what classes or categories you want to apply.
   - Set up document extraction use cases.

     Create extraction use cases, add fields and field groups, and configure data extraction modes to begin extracting data from your documents.
   - Manage document extraction use cases.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Manage your Document Intelligence use cases to support your document extraction requirements.

- Set up document classification use cases.

  Create classification use cases, add fields (classes), and train your use cases to begin classifying documents.

- Manage document classification use cases

  Manage your Document Intelligence use cases to support your document classification requirements.

Set up Document Intelligence

Review the following information before you start setting up Document Intelligence.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Checklist

Setup task	Description
Verify that the ServiceNow core applications or plugins that are required to support Document Intelligence are installed and activated.	Verify that the following applications or plugins are installed and activated from the ServiceNow Store. When you activate the first plugin, its dependent plugins are activated automatically. If not installed, install and activate one application at a time in the following order to ensure a smooth implementation. • Predictive Intelligence (com.glide.platform_ml) • Platform Document Intelligence (com.glide.platform_ml_di) • Document Intelligence UIB Component (sn_docintel_iframe) For more information, see Install Document Intelligence.
(Optional) Install the Document Intelligence Admin application for an improved user experience for process owners to set up, configure, and monitor their document extraction solutions.	Install the Document Intelligence Admin application from the ServiceNow® Store. Review the listing for information on dependencies, licensing or subscription requirements, and release compatibility. Key features of the application include:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Checklist (continued)

Setup task	Description
	• Easily configure your document extraction solutions by defining what you want to extract from your documents. • Monitor the performance of your document extraction solutions and adjust the level of automation.
Verify that you've assigned the required ServiceNow AI Platform roles.	The following roles are used across the Document Intelligence features: • The administrator (admin) installs the application from the ServiceNow Store and assigns the roles. • sn_docintel.admin • sn_docintel.creation_agent • sn_docintel.extraction_agent • sn_docintel.manager • sn_docintel.viewer
Domain separation	Review the domain separation topic information if you intend to separate data, processes, and administrative tasks. See Domain separation and Document Intelligence.
Verify that file extensions required to support Document Intelligence are included in any glide.attachment.extensions system property customizations.	The glide.attachment.extensions system property is empty by default. When customizing it, ensure that the following file extensions are included: • jpeg • jpg • png • pdf For more information, see Restrict attachment file extensions.
Copy any use cases needed to support your Document Intelligence implementation.	Follow the steps in Duplicate a document extraction use case to copy a use case along with its fields, field groups, integrations, flows, and all the related machine learning (ML) models.
Import any use cases needed to support your Document Intelligence implementation.	Follow the steps in Import a document extraction use case to import a use case along

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Checklist (continued)

Setup task	Description
	with its fields, field groups, integrations, flows, and all the related machine learning (ML) models.

Install Document Intelligence

You can install Document Intelligence (sn_docintel) and Document Intelligence Admin (sn_docintel_admin) if you have the admin role. The sn_docintel_admin application installs related ServiceNow® Store dependencies if they aren't already installed.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the Document Intelligence application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.
- Review the Document Intelligence Admin application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.
- Role required: admin

About this task

The following items are installed with Document Intelligence:

- Plugins
- Roles
- Tables

For more information, see Components installed with Document Intelligence.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Document Intelligence Admin application (sn_docintel_admin) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you can't find the application, you might have to request it from the ServiceNow Store.

   In the list next to the **Install** button, the versions that are available to you're displayed.
3. Select a version from the list and select **Install**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In the Review Installation Details dialog box, any dependencies installed with your application are listed.

4. If you're prompted, follow the links to the ServiceNow Store to get any additional entitlements for dependencies.
5. Select **Install**.

Upgrade to Document Intelligence 3.0 or later from version 2.4 or earlier

Document Intelligence 3.0 or later includes an updated database schema to support its transition from a scoped application to a ServiceNow AI Platform plugin.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Upgrade tasks

Review the information in this topic before you upgrade to Document Intelligence (DocIntel) 3.0 or later from version 2.4 or earlier.

For more information on installation, see Install Document Intelligence.

Before the upgrade

Document Intelligence 3.0 or later pre-upgrade checklist

Pre-upgrade task	Description
Choose a time to schedule the upgrade	Avoid performing the upgrade during the run time of the nightly task definition (use case) upgrade job. Upgrading the application at that time would prevent use cases from upgrading until the run time of the next day.
Update any custom code that points directly to DocIntel database tables	After upgrading to Document Intelligence 3.0 or later from version 2.4 or earlier, DocIntel will use new flow actions. DocIntel will also use ServiceNow AI Platform database tables (sys_di_) in place of the scoped application tables (di_). All DocIntel data will be migrated automatically to these platform tables. **Note:** Document extraction use cases may not be available for use until the migration is complete. If you have custom code in your instance that points directly to the scoped application tables (di_), it should be updated, preferably to the platform component using DocIntel APIs.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence 3.0 or later pre-upgrade checklist (continued)

Pre-upgrade task	Description
	If you have custom code using DocIntel 2.4 or earlier flow actions, it should be updated to use 3.0 or later flow actions. For the list of tables, see Components installed with Document Intelligence.
Add cross-scope records for integrations	For any integrations with Document Intelligence, add cross-scope records for the new DocIntel database tables. **Note:** Ensure that the existing cross-scope records for the old tables are not removed. See the Cross-scope records section for the list of records to be added. For more detail on cross-scope records, see Cross-scope privilege record.

During the upgrade

Document Intelligence 3.0 or later upgrade checklist

Upgrade task	Description
Avoid using the Document Intelligence workspace to extract fields	Wait until a document task is migrated to the sys_di_task table before completing it using the Document Intelligence workspace. Completing a task during the upgrade may lead to data loss.

After the upgrade

Document Intelligence 3.0 or later post-upgrade checklist

Post-upgrade task	Description
Test DocIntel integrations and custom workflows	Test your use case integrations to ensure they function with the new database schema.
Check for data loss	Some cases where you can lose data include: • Any action resulting in the attachment deletion in an old di_task record would result in a missing attachment for the new sys_di_task record.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence 3.0 or later post-upgrade checklist (continued)

Post-upgrade task	Description
	• Deletion of any new tasks created during the data migration (that is, during or right after the 3.0 or later upgrade) would result in permanent deletion as those records might not have entries in di_task tables. • Deletion of an old "process_task" flow (or any action resulting in that flow deletion) would result in an empty flow reference in the new integration setup records.
Post-upgrade migration of legacy use cases imported through update sets	At any time after the completion of the Document Intelligence 3.0 or later upgrade and data migration, you can rerun the migration batch script that ran during the upgrade. You may want to run this script if you must migrate any imported use cases requiring migration to the platform tables (sys_di_). 1. Navigate to **All** > **System Definition** > **Scheduled jobs**. 2. Open the *`DocIntel migrate remaining data`* scheduled job. 3. Select the **Active** check box. 4. Select **Execute Now**. ⚠ **Warning:** Never use the *`DocIntel migrate remaining data`* batch job to rerun migration after deleting all records from sys_di tables. You should only use it to migrate missing data from di_ tables to sys_di tables.

Cross-scope records

Cross-scope records for Document Intelligence 3.0 or later

Source scope	Target scope	Target name	Operation	Target type	Status
Scope of integrating BU	global	sys_di_task	Read	Table	Allowed
Scope of integrating BU	global	sys_di_key	Read	Table	Allowed
Scope of integrating BU	global	sys_di_key_group	Read	Table	Allowed
Scope of integrating BU	global	sys_di_ocr_input	Read	Table	Allowed
Scope of integrating BU	global	sys_di_pdf_input	Read	Table	Allowed

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cross-scope records for Document Intelligence 3.0 or later (continued)

Source scope	Target scope	Target name	Operation	Target type	Status
Scope of integrating BU	global	sys_di_prediction_input	Read	Table	Allowed
Scope of integrating BU	global	sys_di_training_input	Read	Table	Allowed
Scope of integrating BU	global	sys_di_key_hint	Read	Table	Allowed

Configure Document Intelligence settings

Use general settings to control how Document Intelligence displays information to users.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin

Role required: sn_docintel.admin, sn_docintel.manager, or admin

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Settings**.
2. Adjust the settings as desired.

Note: The following settings only apply to document extraction features.

DocIntel settings

Setting	Description
Exact match option	Select **Auto-sync** to enable the exact match option. This lets the user select an icon to view only the recommendations that match exactly what they type.
Candidate score	Select **Auto-sync** to make confidence scores visible next to each recommendation. For more information, see Confidence scores.
Side panels	Select the location of the navigation panel and document fields panel.
Default width of the extraction panel	Enter the default width of the document fields (extraction) panel in pixels.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Setting	Description
Default width of the thumbnail panel	Enter the default width of the navigation (thumbnail) panel in pixels.
Default document fit for the image panel	Select how to display the document in the document (image) panel.

Set up document extraction use cases

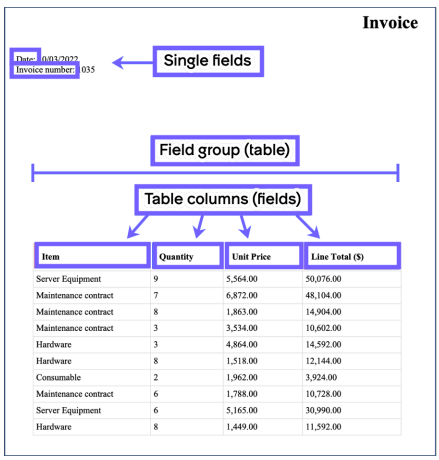
In Document Intelligence, a use case is used to define the structure of a type of document you want to process. It's made up of the use case record and its related fields, field groups, integrations, flows, and all the related machine learning (ML) models.

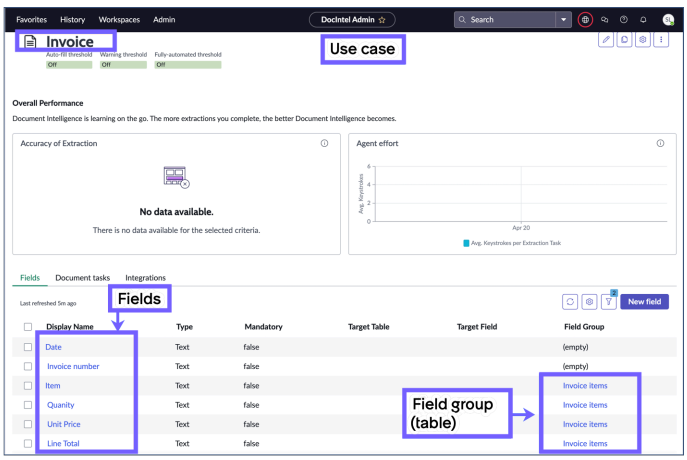
Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Overview of document extraction use cases

In a document extraction use case, you define the information that you want the AI to detect in a document. Do this by specifying the type of document to process, the fields to detect, and the location where document processing results are to be stored.

For example, if you want to process invoice documents, you may want an “Invoice” use case. This use case could have fields for date, invoice number, item, and so on, to define which information needs to be extracted from the document.





After you've defined a document extraction use case, agents can begin processing documents for it in document tasks.

Note: Use cases for Now Assist in Document Intelligence skills have a separate setup process. For more information, see Configuring Now Assist in Document Intelligence.

Workflow

Set up a document extraction use case in the following steps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. Create a use case.

   Define the name, target table, and language for the use case.

2. Create a field for data extraction.

   Define the fields that the AI will learn to detect and extract values from.

   Define any groupings of fields to help extract and organize data gathered from tables or information patterns, like check box lists.

3. Configure data extraction modes.

   Define how fields should be extracted from documents in a document task.

4. Set up integrations.

   Configure an integration to trigger document task processing or value extraction for workflows with other applications.

As agents work on document tasks to extract field values from individual documents, the AI will learn from the feedback and continue to improve.

Create a document extraction use case

Create a use case record to define a document you want to process in Document Intelligence. For example, invoices or driving licenses.

> **Important:** Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin
Role required: sn_docintel.manager

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Use Cases**.
2. Select **New use case**.
3. Enter a name for the use case.
4. Select a target table to store the document processing results for this use case.
5. **Optional:** Change the language model used to support document extraction.
   For more information, see Languages supported by Document Intelligence.
6. Select **Save**.

What to do next
After creating a use case, finish setting it up by adding fields and configuring the data extraction modes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a field for data extraction

Set up fields as part of your use case. Document Intelligence uses fields to identify and extract data from documents. Fields can be grouped to help DocIntel extract data from documents with tables, check box lists, and other logical groupings of fields.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin

Role required: sn_docintel.manager

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Use Cases**.
2. Select a use case in the list.
3. Go to the **Fields** tab and select **Define your fields**.
   If you have already defined one or more fields and you want to add another field, select **New field**.
4. Select the type of data that you want to extract from the document.

   You can choose one of the following types of data to extract:

   Single field

   Single fields are used to extract a single piece of information in the document. For example, a document number or a customer name.

   Check box list

   Check box lists are used to extract a check box or a group of check boxes. Each check box can be checked or cleared.

   Table

   Tables are used to extract lists or tables of information. A table can have multiple columns. The number of list items or table rows doesn't have to be known in advance.

   Single field group

   Single field groups are used to extract values that are grouped in the document. For example, a location with an address, city, and country. Only one item can be extracted for a single field group, as opposed to the multiple rows extracted for a table.

   A form displays based on the data type you selected.
5. On the form, fill in the fields.
   The type of form depends on the type of field that you selected in the previous step.
   - Single field form
   - Single field group form
   - Check box list form
   - Table form
6. Select **Save**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
The system added the new fields to the Fields list associated with the use case.

Deactivate a data extraction field

Deactivate fields that you don't want to use as part of your use case.

> **Important:** Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin
Role required: sn_docintel.manager

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Use Cases**.
2. Select a use case in the list.
3. Go to the **Fields** tab and select the field you want to deactivate.
4. On the form, select **Deactivate Field**.
5. In the confirmation box, select **Deactivate**.

> **Warning:** There is no way to reactivate the field after it has been deactivated.

Result
The deactivated field will not be used for data extraction or training.

Configure data extraction modes

Configure the extraction modes for use cases to define how Document Intelligence extracts fields from documents.

> **Important:** Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin
Role required: sn_docintel.manager

About this task
Extraction modes determine how the data is extracted in the document task and how the task is processed. The mode changes the behavior of the fields in the Document Intelligence workspace.

DocIntel uses the following extraction modes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Extraction mode	Description
Recommendation	DocIntel provides recommendations for the fields in the Document Intelligence workspace. Choose the recommendation or manually enter the value. All fields must be reviewed. Recommendations are ordered based on how confident the AI is in the prediction. As DocIntel continues processing your documents, recommendations improve over time.
Auto-fill	DocIntel auto-fills the fields in the Document Intelligence workspace. All fields must be reviewed. Auto-fill works only if the AI has enough confidence to make the prediction. You can change the confidence threshold by updating the **Auto-fill threshold** field in the use case.
Fully automated (Straight-through processing)	DocIntel automatically extracts the data for all fields and processes the document task if the confidence scores for all required fields are above the defined confidence threshold. Fields don't need to be reviewed. **Note:** For Fully automated mode to process tasks successfully, the use case must have at least one field designated as required. DocIntel becomes more confident over time, as it processes more and more documents. Choose Fully automated mode for frequently processed documents or if you're confident in the system.

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Use Cases**.
2. Select the use case that you would like to configure.
3. Select the settings icon ().
4. Select the extraction mode for the use case.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If the DocIntel AI model needs further training to make recommendations with higher confidence scores, leave the default recommendation mode in place without selecting any other mode.
- If DocIntel provides recommendations with confidence scores above the specified threshold, select the **Auto-fill mode** option.
- If DocIntel auto-fills the required fields with very high confidence scores, select the **Fully Automated mode** option.

5. **Optional:** Adjust the auto-fill threshold and warning threshold for Auto-fill mode.

   - (Optional) **Auto-fill threshold**: DocIntel only auto-fills the fields if the confidence score of the top recommendation is at or above the percentage you define.

     Fields with a confidence score lower than the threshold are left empty in the Document Intelligence workspace, and the recommendation mode is available to extract these fields.

   - (Optional) **Warning threshold**: DocIntel shows a warning for empty fields and auto-filled fields with a confidence score at or below the percentage that you define.

6. **Optional:** Adjust the confidence threshold for required fields in Fully automated mode.

   - (Optional) DocIntel automatically completes and submits the document task if all required fields have a confidence score at or above the percentage you define in the **Fully-automated threshold** field.

   - (Optional) If any required field is below the threshold, the document task isn't submitted automatically and requires agent review. Fields that aren't required may be left empty or unreviewed.

   - (Optional) If there are no fields defined as required for the document task, then DocIntel automatically completes and submits the document task.

7. Select **Save**.

Manage document extraction use cases

Manage your Document Intelligence use cases to efficiently support your document extraction requirements.

> **Important:** Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

The following topics describe features that you can use to manage your document extraction use cases.

- Duplicate a use case and modify it as needed to quickly set up a new use case with similar fields, field groups, integrations, flows, and all the related machine learning (ML) models.
- Export and import trained use cases to share them across your ServiceNow® instances.
- Delete a use case if you no longer need to process documents with it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Duplicate a document extraction use case

Make a copy of a use case to save time when you want to create a new use case that shares a similar structure to another.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin

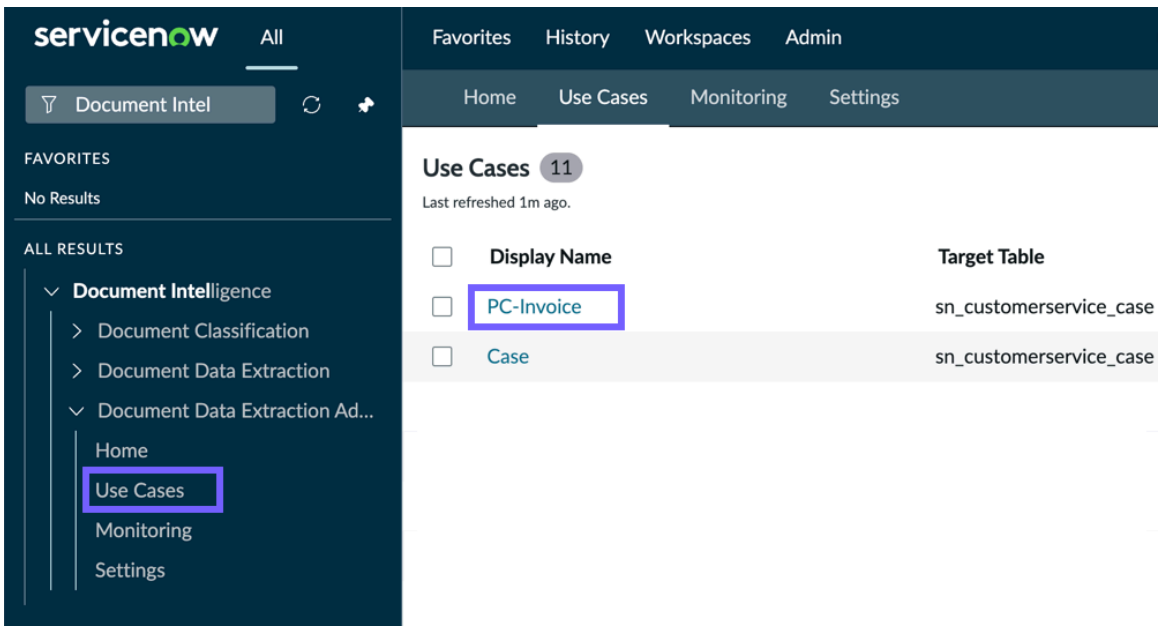
Role required: sn_docintel.manager

About this task

Follow these steps to create a copy of a use case along with its fields, field groups, integrations, flows, and all related machine learning (ML) models.

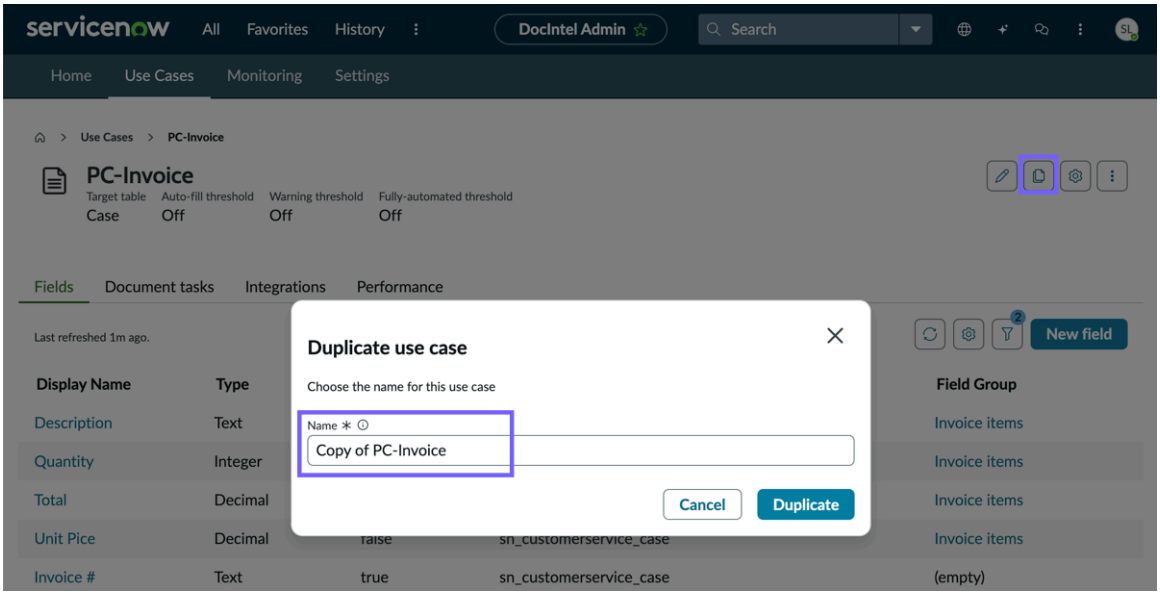
Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Use Cases**.
2. In the list, select the display name of the use case that you want to copy.



3. On the use case screen, select the Duplicate this use case icon ( ).
4. In the Duplicate use case box, type a name for the new use case.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



5. Select **Duplicate**.

Result
The duplicated use case is displayed in the use cases list.

Export a document extraction use case

Export a document extraction use case for use in another ServiceNow instance by adding it to an update set.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin

- Ensure both instances have the same family release and the same version of Document Intelligence installed when exporting and importing use cases.
- Role required: sn_docintel.manager

About this task
Follow these steps to add a document extraction use case to an update set along with its fields, field groups, integrations, flows, and all related machine learning (ML) models. The update set can then be exported for use in another instance.

The update set(s) are automatically created and set to "Completed" as a background process. This process takes several minutes. The resulting update set(s) should not be manually set to "Completed" or exported before the background job finishes.

For more information, see System update sets.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Use Cases**.
2. In the list, select the display name of the use case you want to export.
3. On the use case screen, select the options icon () and select **Add to update set**.
4. Select **Add to update set**.
   The use case is added to a system update set.
5. Navigate to **System Update Sets** > **Local Update Sets**.
6. Select the update set you added.
7. On the update set form, select **Export to XML** under Related Links.

Import a document extraction use case

Import a document extraction use case for use in your ServiceNow instance.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin

- Update sets for a document extraction use case are downloaded according to the steps provided in Export a document extraction use case.
- Ensure both instances have the same family release and the same version of Document Intelligence installed when exporting and importing use cases.
- Role required: sn_docintel.manager

About this task

Follow these steps to import a document extraction use case along with its fields, field groups, integrations, flows, and all related machine learning (ML) models.

For more information, see System update sets.

Procedure

1. Navigate to **All** > **System Update Sets** > **Retrieved Update Sets**.
2. Under Related Links, select **Import Update Set from XML**.
3. Select the exported XML file and upload it.
4. Open the update set record.
5. Select **Preview Update Set**.

   There may be a few errors for sys_di_extracted_values, stating `Could not find a record in sys_di_image for column image referenced in this update`.

   This error is because those extracted values are empty and don't have a di_image.

6. Select **Accept remote update**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. If there are no other errors, select **Commit Update Set**.
8. Repeat steps 3 through 7 for each update set.

Result
The use case is imported into the instance and appears in the use cases list.

Delete a document extraction use case

Delete a use case when it's no longer needed for your documents.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin
Role required: sn_docintel.manager

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Use Cases**.
2. In the list, select the display name of the use case that you want to delete.
3. On the use case screen, select the options icon ( ) and select **Delete**.
4. In the confirmation box, select **Delete**.

Result
The use case is deleted, along with the related fields, field groups, integrations, flows, and machine learning (ML) models.

Set up document classification use cases

A document classification use case is a set of categories used to classify your documents and their individual pages. It's made up of the use case record and its related fields (classes), and all related machine learning (ML) models.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Overview of document classification use cases

In a document classification use case, you define the classes or categories that you want the AI to detect and apply to a document. Do this by specifying the type of document to process, the classes to apply, and the location where document processing results are stored.

For example, if you want to process identification documents, you may want an "Identity Documents" use case. Then, add classes for passports, driver's licenses, military IDs, and so on, to label which type of documents are being processed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

After you've defined a document classification use case, agents can begin processing documents for it in document tasks.

Workflow

Set up a document classification use case in the following steps.

1. Create a use case.

   Define the name and properties for the use case.

2. Create document classes using fields.

   Define the classes or categories that the AI will learn to detect and apply to documents.

3. Train the use case.

   Initiate a training job to provide user inputs from completed document tasks to the AI for continuous improvement.

As agents work on document tasks to classify documents and their individual pages, the AI learns from the feedback and continues to improve.

Create a document classification use case

Create a use case record to begin defining the classes or categories that you want to apply to a type of document or pages within the document.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin
Role required: sn_docintel.manager

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Classification** > **Use Cases**.
2. Select **New**.
3. On the form, fill in the fields.

Use case form

Field	Description
Display Name	The name for the use case as it appears in the Document Intelligence workspace.
Document Type	Type of document to be processed for the use case.
Document Config	Configuration of the document to be processed for the use case.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Autofill Threshold	DocIntel only auto-fills the classes (fields) if the confidence score of the top recommendation is at or above the percentage you define. Fields with a confidence score lower than the threshold are left empty in the Document Intelligence workspace, and the recommendation mode is available to extract these fields. This field is available only if the Auto-fill mode is enabled.
Fully Automated Threshold	Confidence score threshold for document classifications, which enables a document task to be fully automated in the Fully Automated (Straight Through Processing) mode.
Warning Threshold	DocIntel shows a warning for empty fields and auto-filled fields with a confidence score at or below the percentage that you define. This field is available only if the Auto-fill mode is enabled.

4. Select **Submit**.

What to do next

After creating a use case, finish setting it up by adding fields to create document classes. Then, train the use case.

Create a document class

Create fields as part of your document classification use case. Document Intelligence uses fields to define the classes or categories to apply to documents.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin

Role required: sn_docintel.manager

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Classification** > **Use Cases**.
2. Select a use case in the list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Go to the **Fields** tab and select **New**.
4. On the form, fill in the fields.

Define a new field form

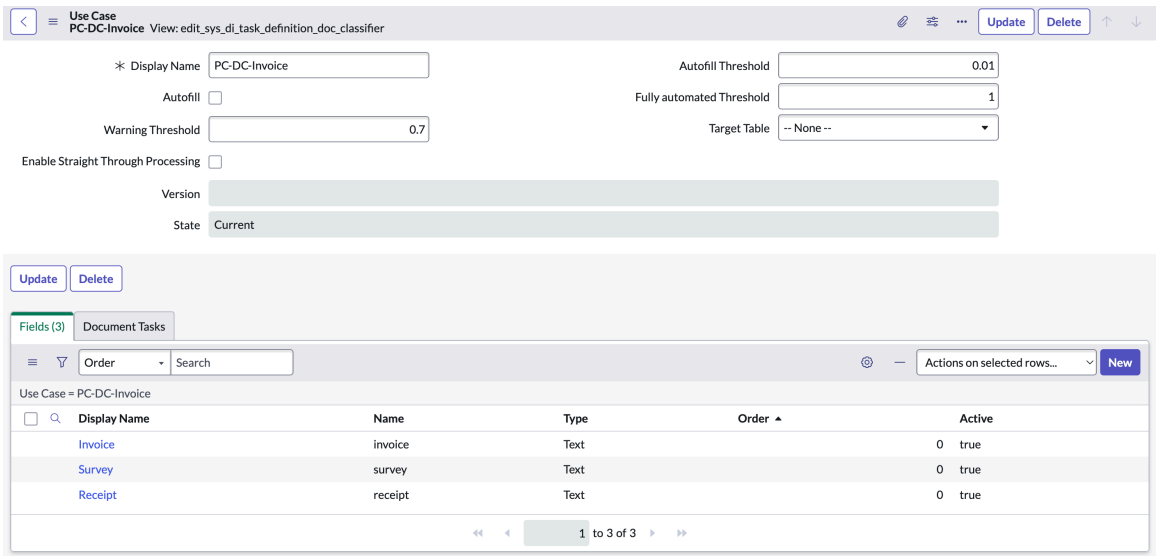
Field	Description
Display Name	The name for the class as it appears in the Document Intelligence workspace.
Use Case	The use case associated with this field (class) record.
Type	The type of field (for example, a text field or a check box option). Select text.
Active	Option to indicate whether the class is being used.

5. Select **Submit**.

Result

The system added the new class field to the Fields list associated with the use case.

Document classification use case with classes



Train a use case

Train the document classification use case with user input from completed document tasks to improve Document Intelligence recommendations over time.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- There must be at least one reviewed document task associated with the use case to train it. Across all reviewed tasks, there must be at least two attachments for each type you defined in your use case.

  Begin by creating a document task and completing it using the Document Intelligence workspace.

  Note: You won't be able to process the task until the use case is trained, but you can complete it by completing all the fields and submitting it.

- Role required: sn_docintel.manager

About this task

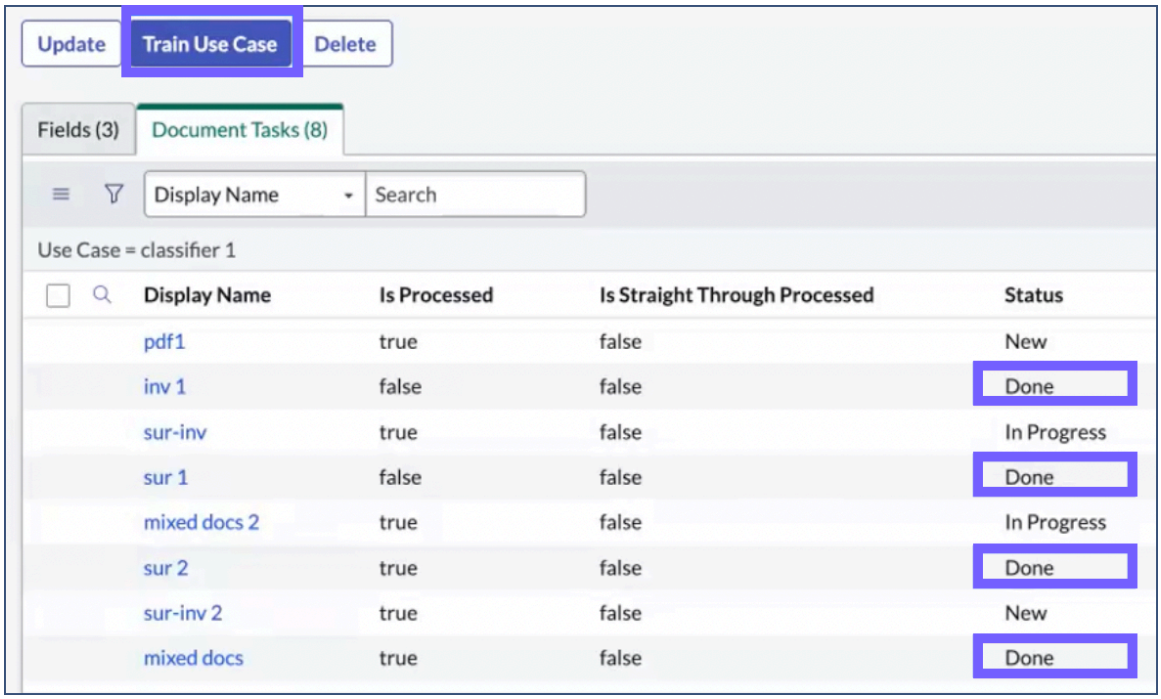
Document classification use cases don't begin with pre-trained AI models, so it's important to train the models with user input from completed document tasks.

Note: To reduce server load and minimize performance issues, the default limit for training a use case is once every 30 days.

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Classification** > **Use Cases**.
2. Select a use case in the list.
3. Select **Train Use Case**.

   DocIntel uses the extracted values from the document tasks in a Done status to train the model.



Result

The train use case job begins. This job may take several hours to complete.

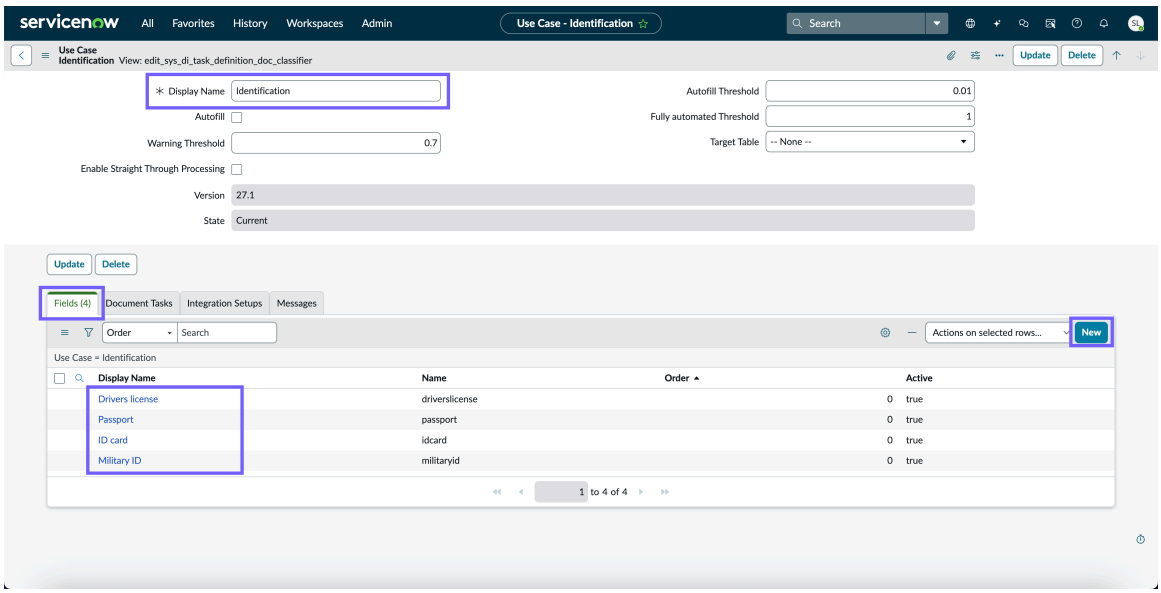
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Example: Test a use case

The following example shows how to train a document classification use case to process identification documents.

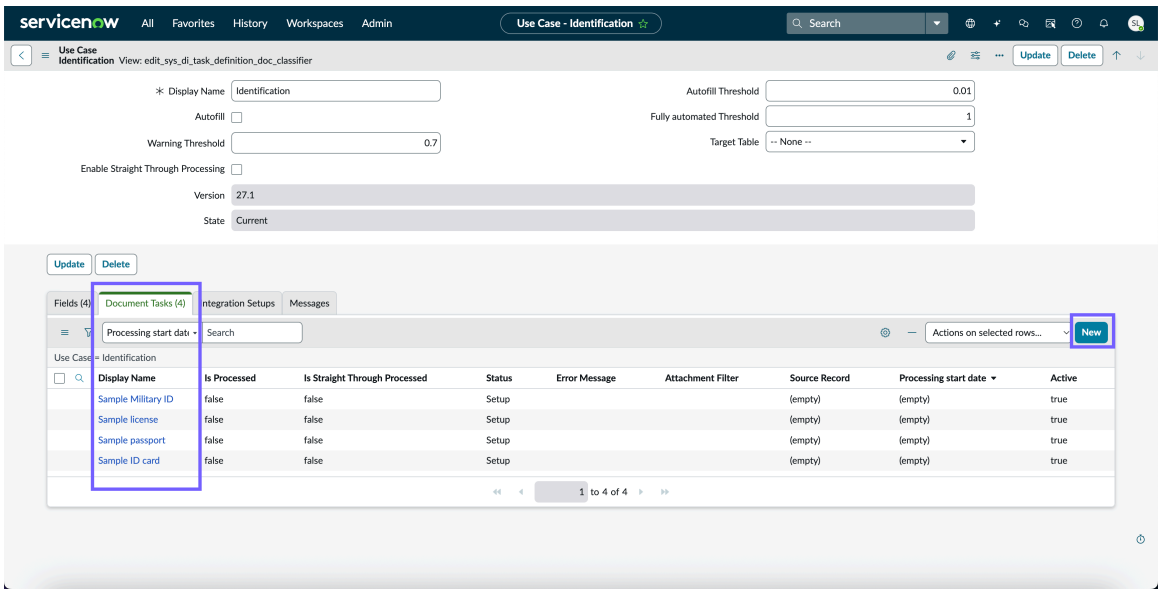
1. You start by creating a document classification use case called Identification. You create a field for each type of category you want to apply to your identification documents.

   In this case, there are fields for a **Drivers license**, **Passport**, **Military ID**, and **ID card**.



2. Then you create document tasks containing attached examples of the categories you defined in the fields.

   Tip: You should submit at least 50 documents to adequately train the AI model.

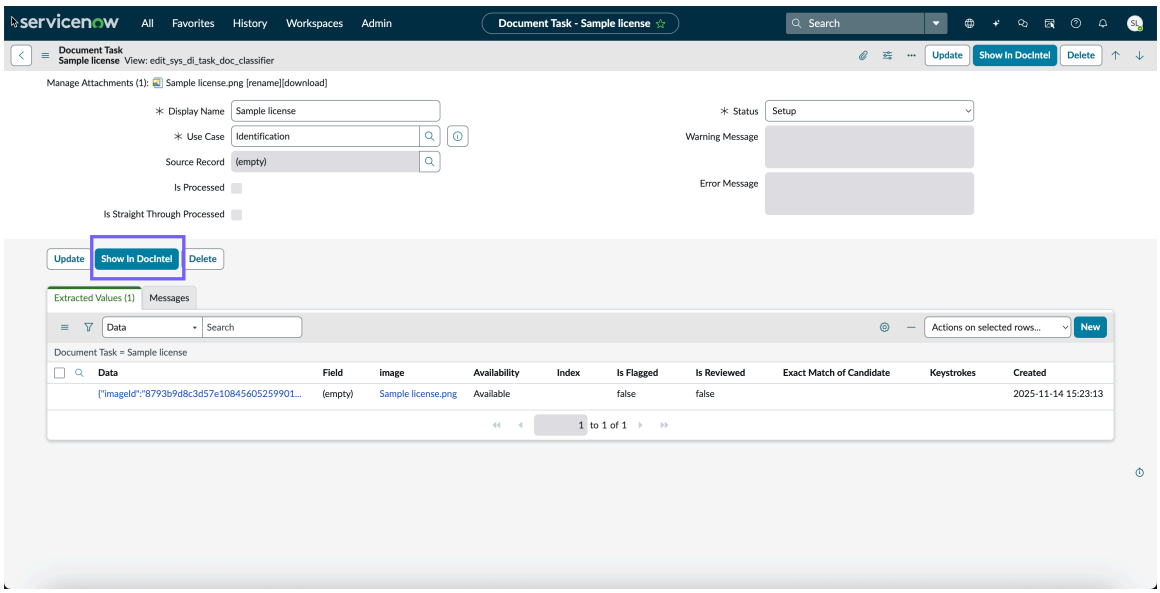


3. You complete the tasks by opening the documents in Document Intelligence workspace, completing all the fields, and submitting them.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Select each document task to open it.

b. Select **Show in DocIntel**.



c. Fill in the category.

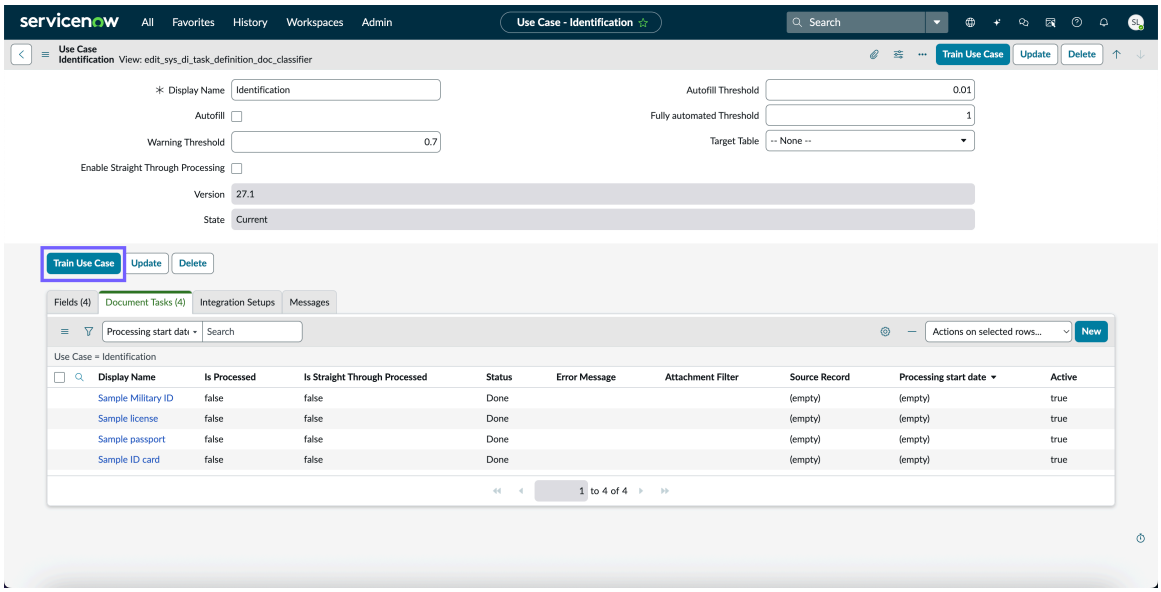
d. Select **Submit**.



The document tasks will be in a Done state and so become available to train the model.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. You go back to the Identification use case screen and select **Train Use Case**.



Manage document classification use cases

Manage your Document Intelligence use cases to efficiently support your document classification requirements.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

The following topics describe features you can use to manage your document classification use cases.

Delete a document classification use case

Delete a use case when it's no longer needed for your documents.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin

Role required: sn_docintel.manager

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Classification** > **Use Cases**.
2. In the list, select the display name of the use case you want to delete.
3. On the use case screen, select **Delete**.
4. In the confirmation box, select **Delete**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
The use case is deleted along with the related fields, field groups, integrations, flows, and machine learning (ML) models. The deletion of the related records may trigger their own related deletions.

Manage field values

View the field values gathered from your processed document tasks. Review the values and add any additional information.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin
Role required: sn_docintel.admin or sn_docintel.manager

About this task
When your document task completes, the extracted data or classifications appear in the Field Values list. The field value form includes more information about the value.

Procedure
1. For document data extraction fields, navigate to **All** > **Document Intelligence** > **Document Data Extraction** > **Field Values**.

   For document classification fields or classes, navigate to **All** > **Document Intelligence** > **Document Classification** > **Field Values**.
2. Select the name of the value (data).
3. On the form, review the fields.

Field	Description
Data	The field value extracted from the document.
Is Reviewed	Indicates whether this field value has been reviewed by a user.
Candidate ID (Recommendation ID)	Internal system ID for the selected recommendation.
Index	For fields that are part of a field group: the order of the field value in reference to other field values for the same field. For regular fields, the index is always 0.
Exact Match of Candidate (Recommendation)	Indicates whether the field value exactly matches the AI's top recommendation.
Field	Field from the use case that the value belongs to.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Candidate (Recommendation) Rank	Rank that the AI assigned to the selected recommendation.
Is Flagged	Whether a user has flagged this field value in the Document Intelligence workspace.
Document task	Document task with the attached document from which the data value was extracted.
Availability	Indicates whether this field value was available or missing in the document.
Keystrokes	Number of keystrokes that were needed to extract this field value.
Target Record	The record where the field value is used.
Metadata	The metadata associated with the field value.
Target Table	The table that stores the field values.
Domain	Domain linked to this field value. See Domain separation and Document Intelligence.

4. Select **Update**.

Integrating Document Intelligence with other applications

Extend the capabilities of Document Intelligence to other ServiceNow applications. Other applications can take advantage of document classification and extraction using Document Intelligence.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Integrate with a custom application or workflow

Configure an integration to trigger document task processing or value extraction. Integrations can be used to quickly set up flows with other applications.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin

- You must first add a target table to your use case before creating an integration.
- The target table must be readable and writable. Ensure the **Can read** and **Can update** check boxes are selected in the **Application Access** tab of the target table record. For more information, see Create a table.
- Role required: sn_docintel.admin or sn_docintel.manager

About this task

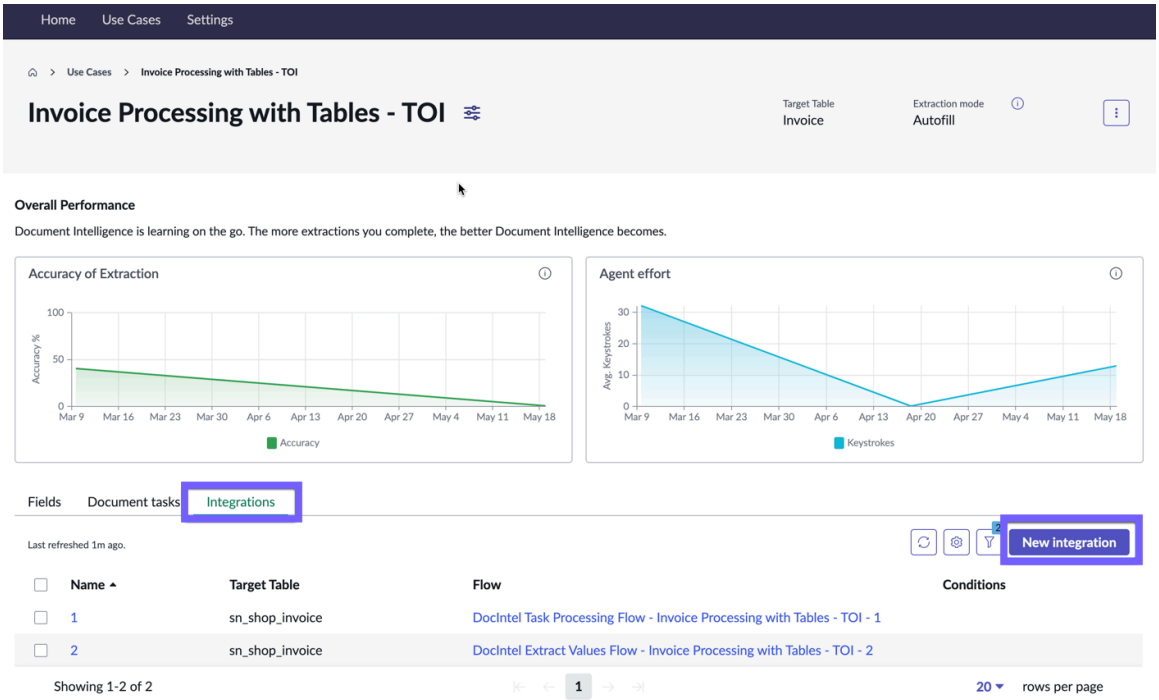
Define integration points for your Document Intelligence solution. Two integration points are provided for data extraction use cases: one to automatically create and process Document Intelligence document tasks, and one to automatically propagate the field values to another application when extraction has been completed in Document Intelligence. For document classification use cases, there's also an integration point to map to a data extraction use case.

Procedure

1. Do one of the following:
   a. To integrate a data extraction use case, navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Use Cases**.
   b. To integrate a document classification use case, navigate to **All** > **Document Intelligence** > **Document Classification** > **Use Cases**.
2. Select the use case for which you want to set up integration points.
3. Do one of the following:
   a. For a data extraction use case, go to the **Integrations** tab and select **Set up your first integration**.

      If you have already defined one or more integrations and you want to add another, select **New integration**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



b. For a document classification use case, go to the **Integration Setups** tab and select **New**.

4. On the form, fill in the fields.

Field for data extraction	Field for document classification	Description
Name your Integration	Name	Name for the integration.
Use case	Use case	Use case to use for the integration task.
Where do you want to take the documents from and store the extracted data?	Target table	Table to receive data from or send data to. **Note:** The target table is taken from the use case.
What type of integrations you want to set?	Type	For data extraction and document classification use cases, the options include **Process Task** or **Extract Values**. Document classification use cases also have a **Map To Document Data Extraction Use Case** option. The **Process Task** type creates an integration point to automatically create and process DocIntel document tasks based on specific triggers happening in the target table. The **Extract Values** type creates an integration point to automatically propagate the extracted values to the target

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field for data extraction	Field for document classification	Description
		table when extraction has been completed in DocIntel (that is, when the document task status changes to Done). The **Map To Document Data Extraction Use Case** type creates an integration point that allows a processed document classification value to automatically create a new data extraction task.
Conditions		Filters used to select certain fields to use as specific triggers for the integration. Process Task only.
Create Flow		Select this option to create a flow for this integration in Workflow Studio. This option should be selected, unless you're planning to write your own custom script to set up the integration.

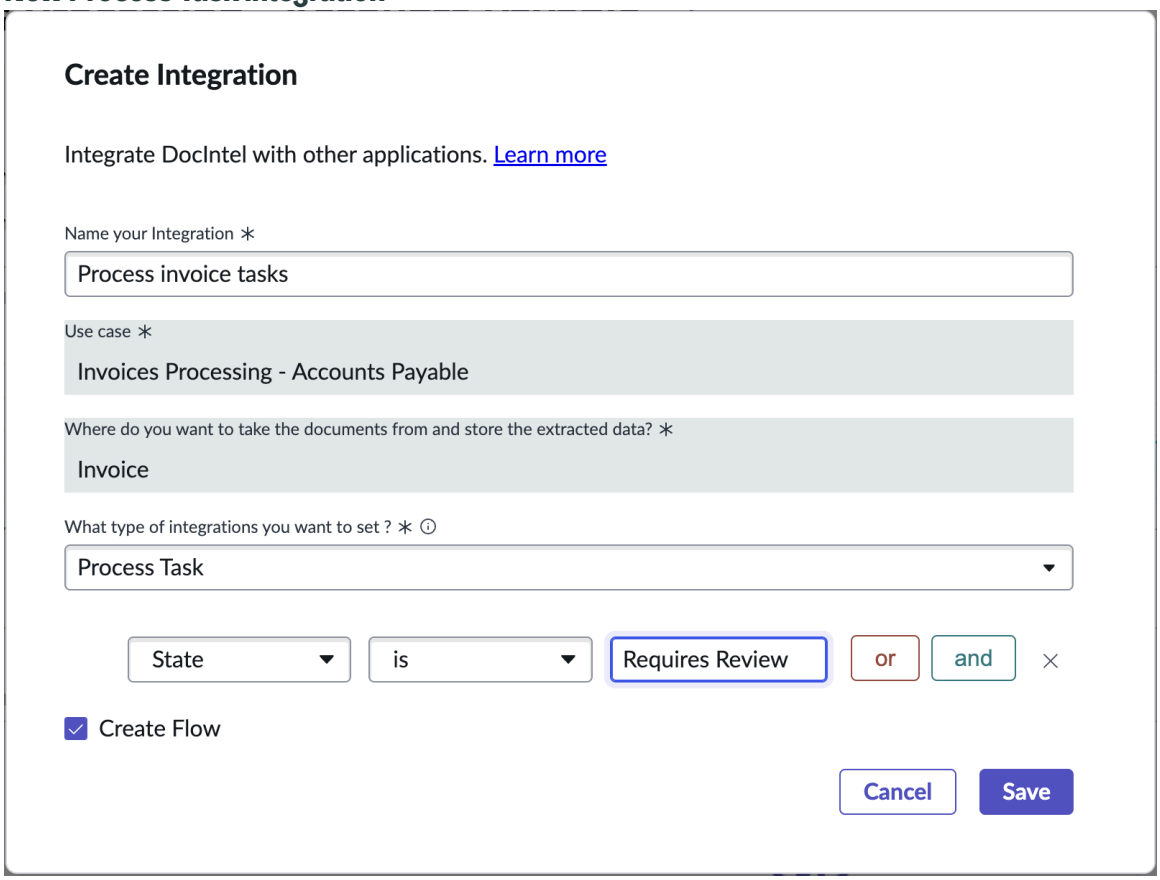
5. Do one of the following:
   a. For data extraction, select **Save**.
   b. For document classification, select **Submit**.

Example: Integration

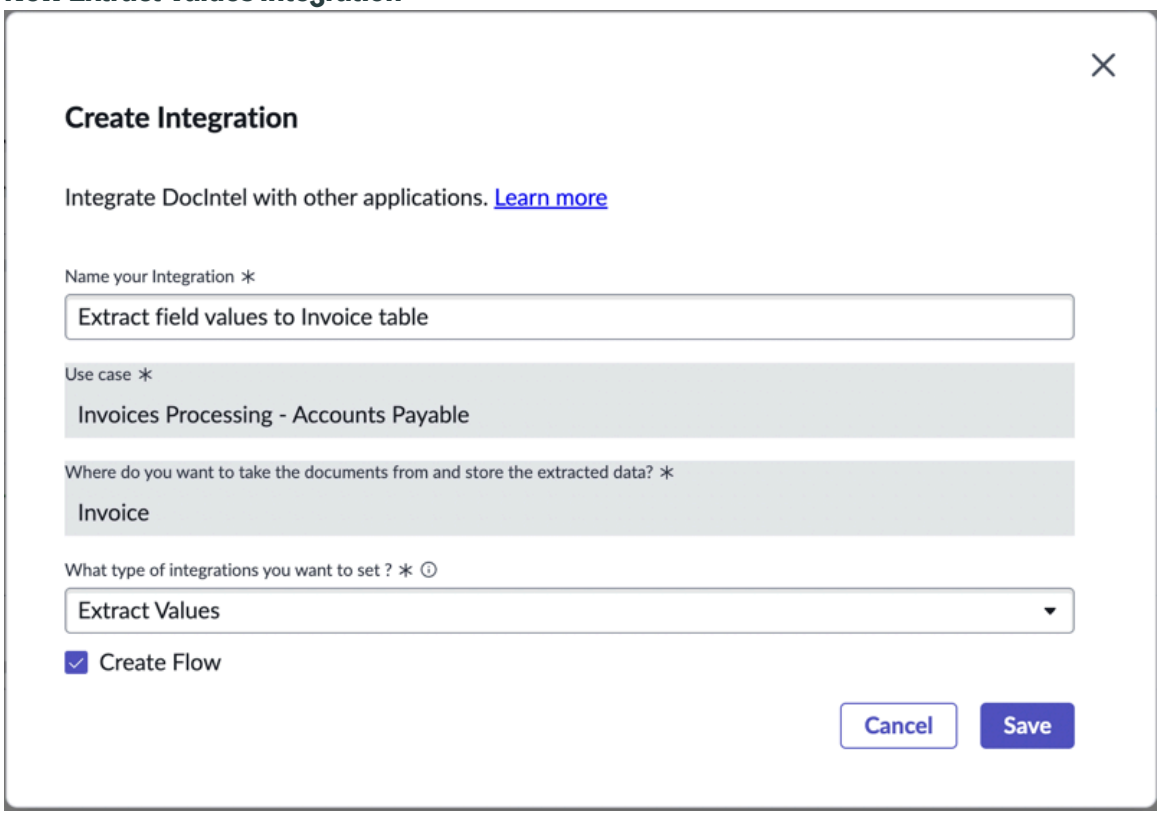
The following images show two example integrations. The first image is a Process Task integration that triggers when a record needs review. The second image is an Extract Values integration that can automatically send extracted fields to the invoice table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

New Process Task integration



New Extract Values Integration



What to do next

If you selected to create a flow, finish the activation in Workflow Studio.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see Building flows.

Integrate with Customer Service Management

Document Intelligence provides document extraction capabilities to Customer Service Management (CSM). Extract relevant information from email and case attachments, such as credit card numbers or customer addresses, and add that information to cases.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Agents can review values for extracted fields and make corrections as needed by accessing the Document Intelligence interface directly from the case. From this interface, agents can confirm correct values, fix incorrect values, and continue to train the model. This Human In the Loop (HITL) interaction of verifying the recommended values enables agents to refine the model and continually improve performance.

Create use cases that identify the information to extract from attachments, such as invoices, and automatically add that information to case fields, depending on the configuration. Labels identify the extracted fields on the case form.

For more information, see Document Intelligence for CSM.

How Document Intelligence works with CSM

When a case is created, the Document Intelligence for CSM feature checks to see:

- If the case has one or more attachments
- If the attachment types are specified in the sn_csm_ml_task.case.docintel.parsing_supported_types system property

If the case meets those requirements, the feature:

- Identifies the right use case to use based on the case or case type and the use case filters.
- Creates a task using the use case, the sys_id of each attachment, and the case reference.
- Sends the task, attachment sys_ids, and case reference as inputs to the prediction model.
- Uses optical character recognition (OCR) solutions to extract data from the documents.
- Sets the status of the task to Done after the solution has completed.
- If the extraction mode in the use case is set to Fully automated (straight through processing), the extracted field values are added to the case.
- If the extraction mode is set to Auto-fill or Recommendation, the agent can validate the field values in the Document Intelligence interface.

Enable Document Intelligence for CSM

Set the system properties and activate the required flows to enable Document Intelligence for CSM.

See Document Intelligence for Customer Service.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Integrate with Financial Services Operations

Document Intelligence provides document extraction capabilities to Financial Services Operations (FSO).

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Document Intelligence integration with FSO enables machine learning (ML) to assist in quickly automating document processing and accurately extracting information from documents. Admins can further integrate Document Intelligence with the FSO Document Processor to enable users to extract and store fields from a document.

For more information, see Document Intelligence for FSO.

How Document Intelligence works with FSO

Document Intelligence works within the Document Processor as follows:

- Creating a Document Type (sn_doc_processor_type) creates a record on DocIntel Use Case (sys_di_task_definition)
- Creating a Document Attribute (sn_doc_processor_attribute) creates a DocIntel field (sys_di_key) record
- Updating the Document Attribute record updates the field record

Creating a Document Verification Task (sn_doc_processor_verification_task) creates a DocIntel Document Task (sys_di_task) if it meets the following requirements:

- Document Verification task has "ocr_processing_needed" checked
- Document Verification task has a document attached
- Document Verification task is in state "Submitted"

After the DocIntel document task is complete, the values should be extracted to the Extracted Fields (sn_doc_processor_extracted_value) table.

Enable Document Intelligence for FSO

Ensure all the necessary applications and plugins are installed and activated to enable Document Intelligence in the FSO Document Processor.

See Enable Document Intelligence for FSO.

Integrate with Accounts Payable Operations

Document Intelligence provides document extraction capabilities to Accounts Payable Operations (com.sn_ap_ic).

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Accounts Payable Operations uses the capabilities of Document Intelligence to extract information quickly and accurately from invoice documents that are received as email attachments and then create invoice records in the Accounts Payable Operations application.

For more information, see Accounts Payable Operations integration with Document Intelligence.

How Document Intelligence works with Accounts Payable Operations

DocIntel for Accounts Payable Operations extracts the required information from invoices received as email attachments, creates invoice records in Accounts Payable Management, and adds the extracted information into the invoice records.

For more information, see How Accounts Payable Operations integration with Document Intelligence works.

Enable Document Intelligence for Accounts Payable Operations

Ensure all the necessary applications and plugins are installed and activated to enable Document Intelligence in the Accounts Payable Operations application plugin.

See Components installed with Accounts Payable Operations integration with Document Intelligence.

Integrate with Automation Center

Use Automation Center to discover opportunities to automate document processing with Document Intelligence.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

In Automation Center, you can view Document Intelligence automation metrics and discover automation opportunities. For more details on Automation Center, see Automation Center.

Using Document Intelligence

Use document tasks to process documents for classification and data extraction in Document Intelligence.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

In a document task, you upload the single or multi-page documents you want to process. Then DocIntel processes them by detecting and analyzing text in those documents and auto-filling or providing recommendations to populate the fields defined in the use case.

For document tasks that aren't yet fully automated, training the DocIntel AI is an important part of achieving full automation.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Process document tasks in the following steps:

1. Create a document task.
   a. Create the task record.
   b. Process the document task.
2. Complete a document task.
   a. Use the Document Intelligence workspace to train the AI.
   b. Submit the document task.

Create a document task

Create a document task and upload single or multi-page documents that are in JPEG, PNG, or PDF formats to start extracting text or classifying documents.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin

Role required: sn_docintel.creation_agent

About this task

These are the manual steps to create a document task. If you have integrations set up with other workflows, this task may be automated. See Integrating Document Intelligence with other applications.

Procedure

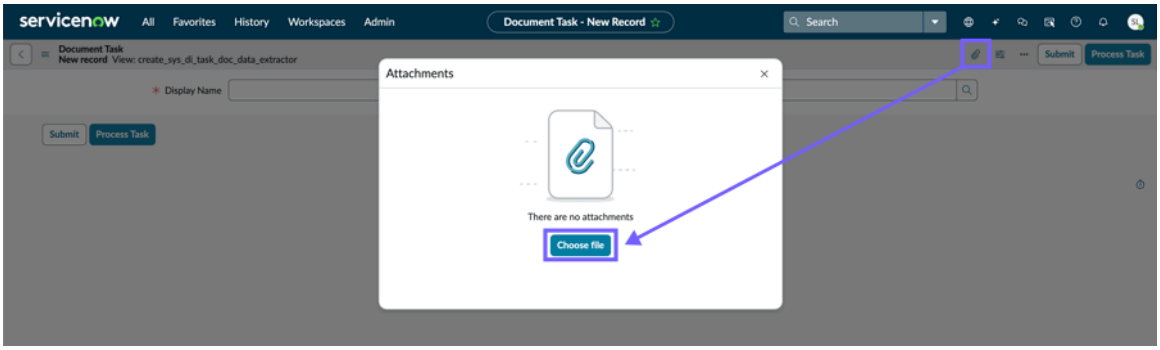
1. Do one of the following:
   - To create a document task for document data extraction, navigate to **All** > **Document Intelligence** > **Document Data Extraction** > **Create Document Task**.
   - To create a document task for document classification, navigate to **All** > **Document Intelligence** > **Document Classification** > **Create Document Task**.
2. On the form, fill in the fields.

New document task record

Field	Description
Display Name	The name associated with this use case. For example, PO-1.
Use Case	The use case created for this document type.

3. Select the **Manage Attachments** icon and select the document you want to upload for processing.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. Select **Process Task**.

   Document Intelligence analyzes the document and extracts data or classes for the fields defined in the use case.

What to do next

Find the document task in the document tasks list. After the **Is Processed** field changes to True, Document Intelligence has completed the analysis of the document. You can proceed to complete the document task, which helps train the AI through your input or review of the extracted fields.

Complete a document task

After the document task processing is finished, complete the task by providing input or review to train the AI.

> **Important:** Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin

Role required: sn_docintel.extraction_agent

Procedure

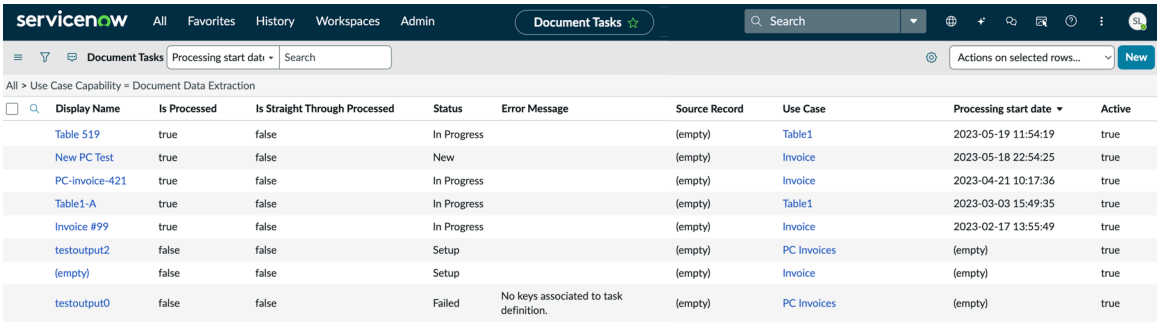
1. Do one of the following:
   - To complete a document task for document data extraction, navigate to **All** > **Document Intelligence** > **Document Data Extraction** > **Document Tasks**.
   - To complete a document task for document classification, navigate to **All** > **Document Intelligence** > **Document Classification** > **Document Tasks**.
2. Review the document task list and ensure that the document task's Is Processed field is set to true.

   The following image shows the list of document tasks. If the **Is Processed** field for a task is true, that indicates that the task has been processed and is available for user input.

   If the status appears as Setup, go to the document task record and select **Process Task**.

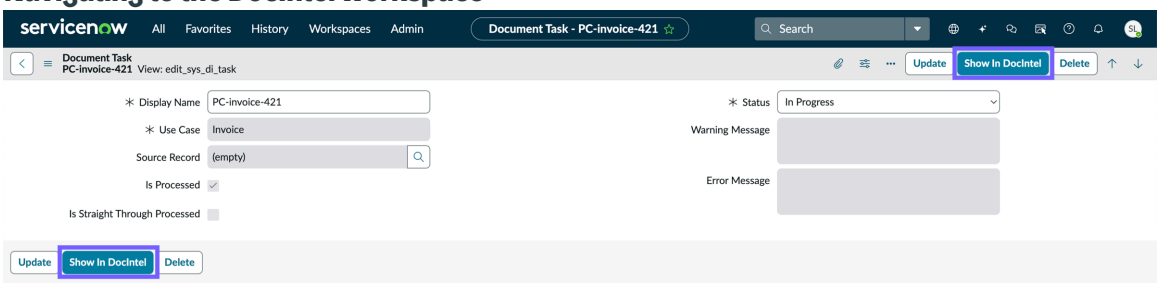
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document tasks list



3. Select the processed document task.
4. Select **Show in DocIntel**.
   The following image shows how to navigate to the Document Intelligence workspace for data extraction.

Navigating to the DocIntel workspace



   The Document Intelligence workspace opens in a new tab.
5. Use the Document Intelligence workspace to review the document fields and select the appropriate recommendations.

   Learn how to use the Document Intelligence workspace to extract fields.

   Learn how to use the Document Intelligence workspace to classify documents.
6. Select **Submit** to complete the document task.

Extract fields using the Document Intelligence workspace

Use the Document Intelligence workspace for field extraction, searching for recommendations, flagging fields, and identifying missing fields to complete document tasks.

> **Important:** Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Overview of the Document Intelligence workspace

The Document Intelligence (DocIntel) workspace provides document management features that enable you to quickly review and process text extraction.

With the Document Intelligence workspace, you can:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Efficiently review the AI's recommendations, and confirm your document's extracted text.
- Flag fields, identify missing fields, and review pending fields.

To get started with the Document Intelligence workspace:

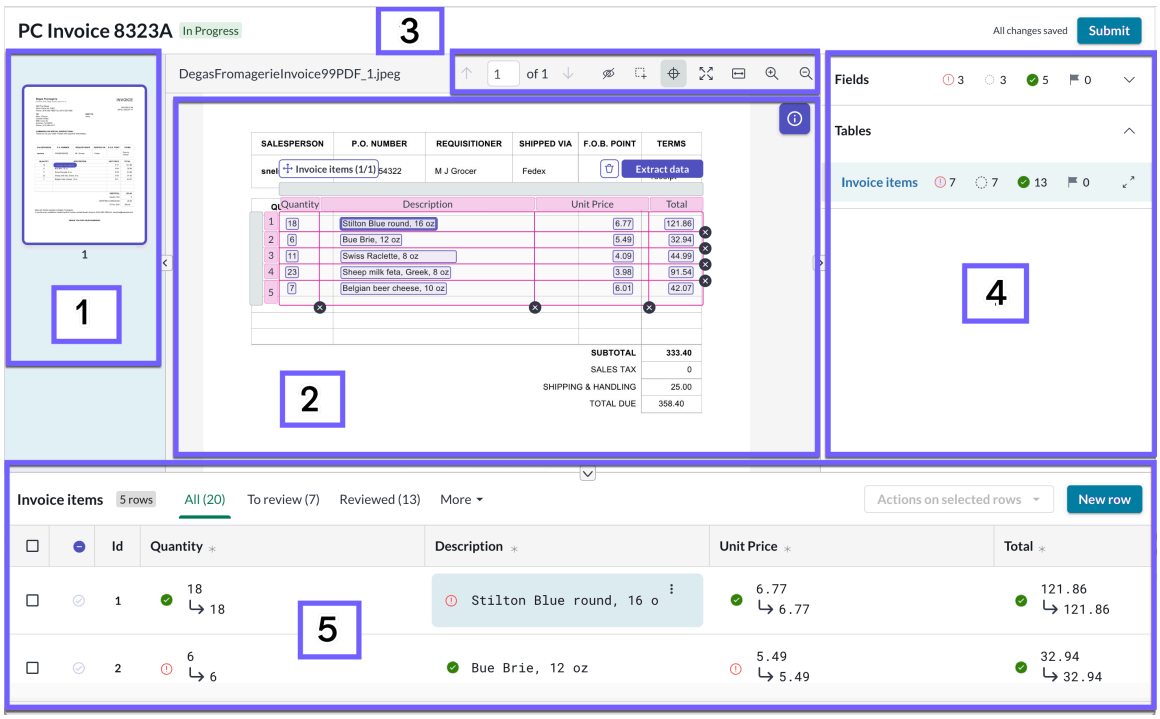
- Create a document task for a document data extraction use case, upload a document, and process the task.
- After DocIntel has processed the task, you can begin using the workspace. See Complete a document task.

Layout of the Document Intelligence workspace

The following illustration shows the Document Intelligence workspace for a document task. The workspace includes the following areas:

- 1 - Thumbnail panel
- 2 - Document image panel
- 3 - Document controls
- 4 - Document fields panel
- 5 - Table panel

Note: In this view, the document fields panel is expanded so that the fields are visible.

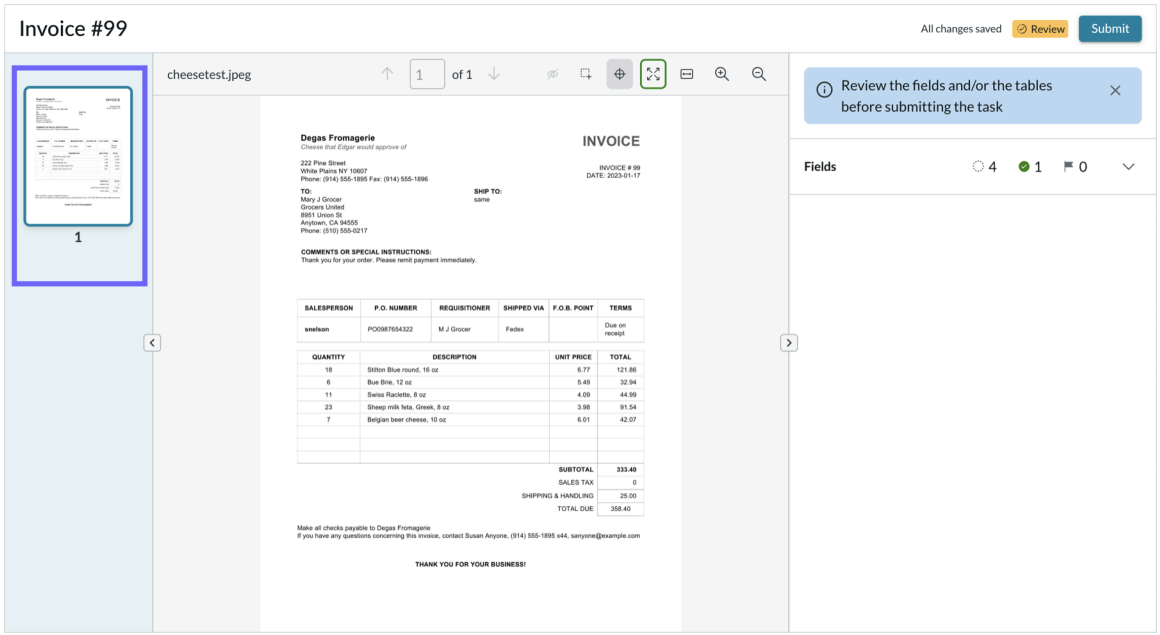


Thumbnail panel

In the thumbnail panel of the workspace, you can select one page from a multiple-page document. The selected page is displayed in the document image panel. Selecting a page doesn't affect what is displayed in the document fields panel.

The following image shows a document with two pages in the thumbnail panel.

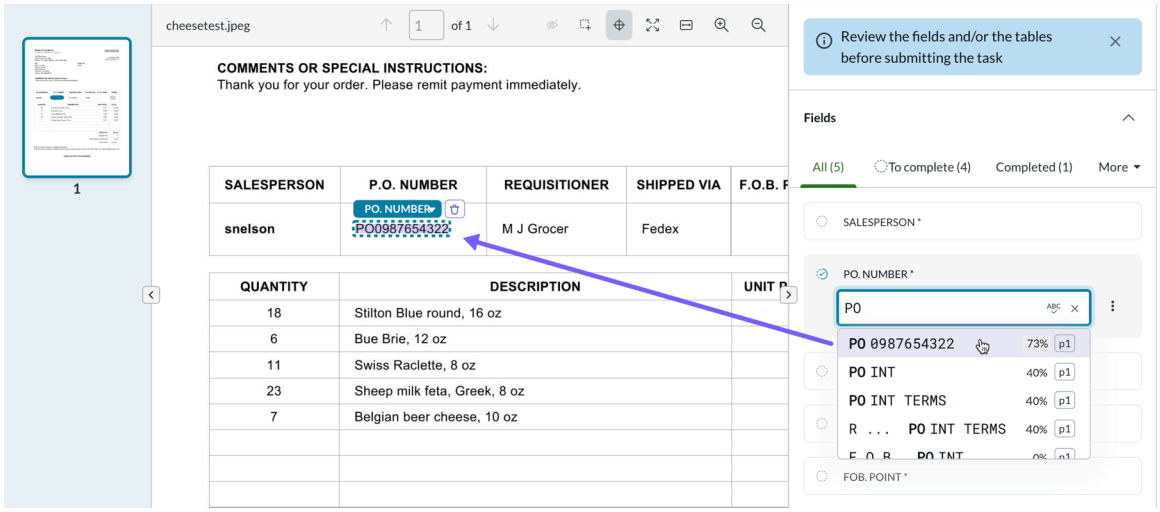
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Document image panel

The document image panel displays the page selected in the thumbnail panel.

As you move through the recommendations in the fields in the document fields panel, they'll be highlighted in the document image panel to help you select the correct option.



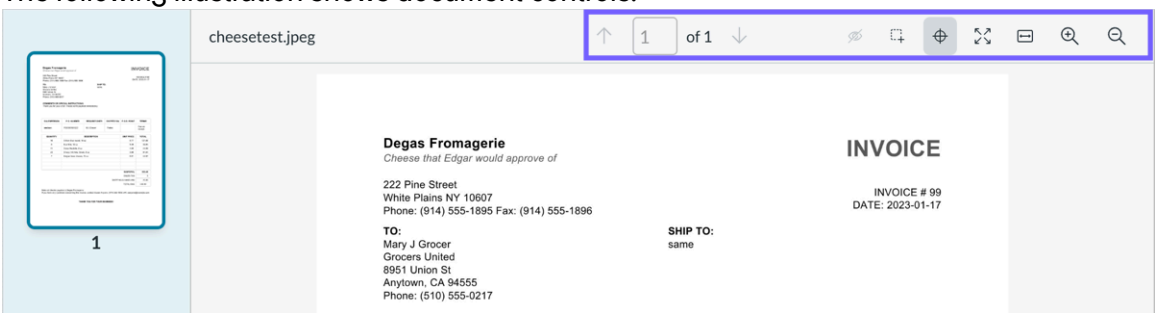
You can also extract information directly from the document image using the draw tool. For more details, see Extract fields using the draw tool.

Document controls

When you're reviewing a document for extraction, you can use various controls to maximize the viewing area, zoom, or focus on the areas that you need. You can also extract information directly from the document image using the draw tool. For more details, see Extract fields using the draw tool.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The following illustration shows document controls.



Document fields panel

The document fields panel enables you to open items for review, including viewing the AI's recommendations. You can also flag fields or mark fields as missing in the document.

Note: An asterisk indicates a required field.

The following illustration shows the different features of the document fields panel.

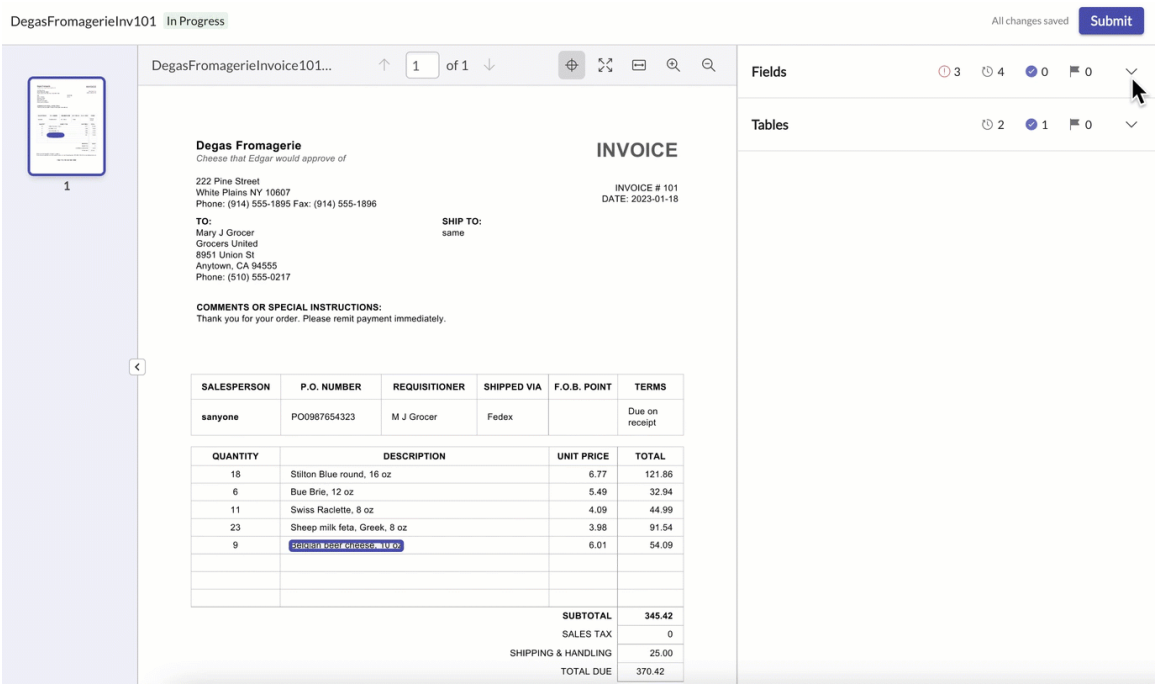


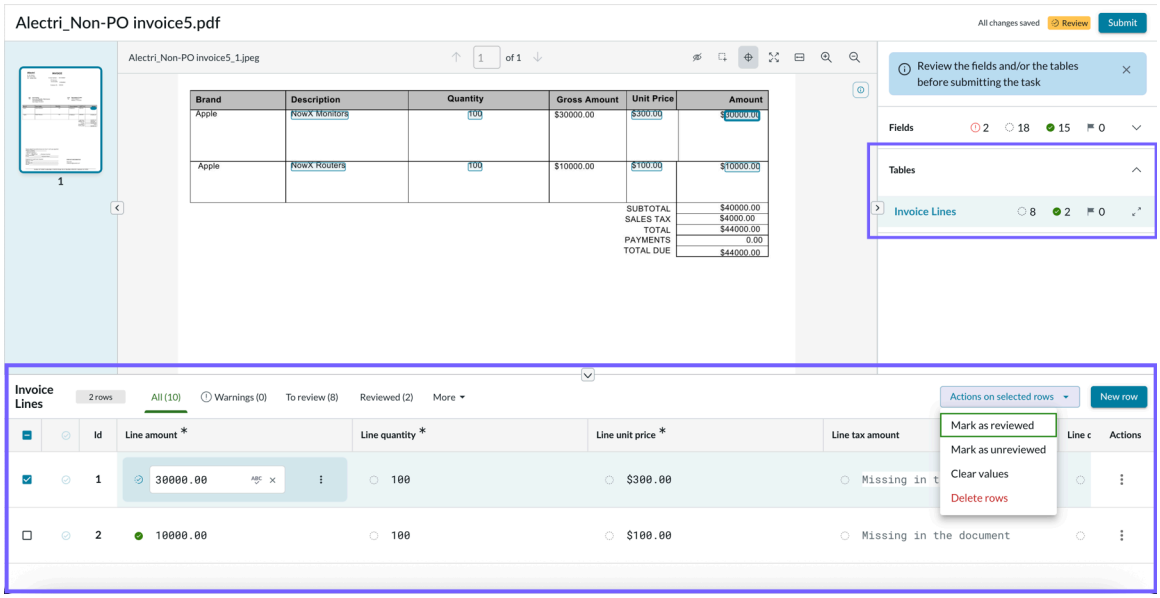
Table panel

The table panel enables you to open table rows for review, including viewing the AI's recommendations. You can also flag fields or mark fields as missing in the document. Other table controls enable you to insert rows and resize columns as needed.

Note: An asterisk in the column heading indicates a required field.

The following illustration shows the different features of the table panel.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Extract single fields

Extract text and number fields from your document in the document fields panel.

Before you begin

Role required: sn_docintel.extraction_agent

About this task

Use the following steps to extract single fields from a document.

If the fields are already auto-filled with values, you can review them to ensure they are correct or adjust the fields as needed.

Tip: You can also use the draw tool to easily extract fields directly on the document image. For more information, see Extract fields using the draw tool.

Procedure

1. In the document fields panel, expand **Fields**.
2. Select a field.
3. Begin typing in the field and select a recommendation from the drop-down list.

   The recommendation with the highest confidence score displays at the top of the list.

   Tip: As you move through the recommendations in the drop-down list, they're highlighted in the document to help you select the correct option.

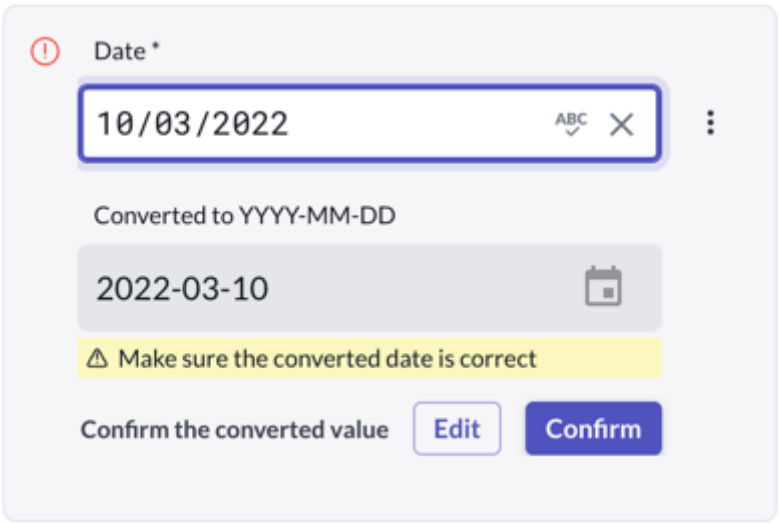
4. **Optional:** If needed, filter or flag the recommendations
   - To filter the list to show only a recommendation that matches exactly what you type, select the exact match mode icon ( ) in the field.
   - To flag the field for later attention, select **Flag for follow-up** in the field options menu ( ).
   - If you can't find an appropriate match in the document, select **Missing in the document** in the field options menu. To undo, select the **Edit** icon ( ) in the field.

   The field will be marked as complete when you move to the next field.
5. **Optional:** If needed, confirm or correct any field value conversions.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

(Optional) Some field types convert the extracted value into a standard format. See Data normalization.

When the extracted value is ambiguous in a document, DocIntel interprets it as defined in the field configuration. A note prompts you to confirm or edit the converted value.



- Select **Confirm** in the converted value field to confirm that the converted value is accurate.
- Select **Edit**, enter the updated value, and select **Save** to correct the conversion.

Extract check box fields

Extract check box fields from your document in the document fields panel.

Before you begin

Role required: sn_docintel.extraction_agent

About this task

Use the following steps to extract check box fields from a document.

If the check box fields are already auto-filled, you can review them to ensure they are correct or adjust the fields as needed.

Procedure

1. In the document fields panel, expand **Fields**.
2. Select a field.
   A list of check boxes displays if the field is designated as a check box field type.
3. Select the check boxes that apply to the field group.
   For each check box selected, DocIntel may provide one or more potential matches in the document.

   a. Select a check box.

   b. If there's more than one recommendation on the document image that matches the check box, select the one that applies.

      You can select or deselect a recommendation on the document at any time.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tip:

- Press **Tab** on the keyboard to move through the recommendations.
- Press **Enter** or **Return** to select one.
- To deselect a recommendation, select another one.

c. If there are no recommendations found, or if none of the recommendations are correct, select **Show all check box recommendations** in the check box options menu ( ) and choose from any other options in the document.

Tip: Flag the check box for later attention by selecting **Flag for follow-up** in the check box option menu.

d. If you can't find an appropriate match for the check box in the document, select **No match found** in the check box option menu.

e. Repeat these steps for each check box that applies to the field group.

4. If you're in the Recommendation mode, ensure that all fields are marked complete.
If you're in the Auto-fill mode, mark the fields as reviewed.

Extract table fields

Extract table fields from your document in the table panel.

Before you begin

Role required: sn_docintel.extraction_agent

About this task

Use the following steps to extract table fields from a document.

If the fields are already auto-filled with values, you can review them to ensure they're correct or adjust the fields as needed.

Tip: You can also use the draw tool to easily extract tables directly on the document image. For more information, see Extract fields using the draw tool.

Procedure

1. In the document fields panel, expand **Tables**.
   The Tables section only displays fields assigned to a table field group.
2. Select a table.
   The table panel displays.
3. Select the recommendation that applies to each table field.
   - To select a recommendation, begin typing in the field and select a recommendation from the drop-down list.

Tip: As you move through the recommendations in the drop-down list, they are highlighted in the document to help you select the correct option.

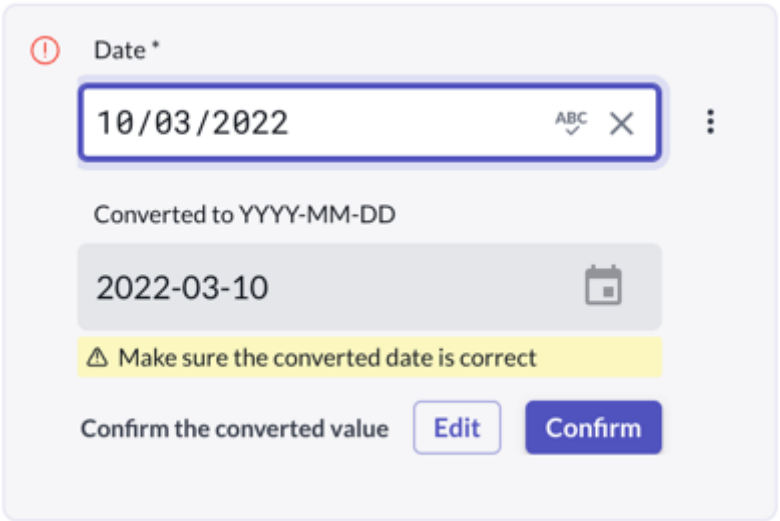
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- To flag the field for later attention, select **Flag for follow-up** in the field options menu ( ).
- If you can't find an appropriate match in the document, select **Missing in the document** in the field options menu. To undo, select the **Edit** icon ( ) in the field.

4. **Optional:** If needed, confirm or correct any field value conversions.

   (Optional) Some field types convert the extracted value into a standard format. See Data normalization.

   When the extracted value is ambiguous in a document, DocIntel interprets it as defined in the field configuration. A note prompts you to confirm or edit the converted value.



   - Select **Confirm** in the converted value field to confirm that the converted value is accurate.
   - Select **Edit**, enter the updated value, and select **Save** to correct the conversion.

5. Adjust the table rows as needed.

   Warning: If you have a grid on the document image, changes made directly to the table rows can't be synced to the grid(s). The grid(s) will be removed to avoid conflicting data.

   - To add a row, select **New row**.
   - To clear all field values in the row, select **Clear row values** in the row options menu ( ).
   - To insert a row, select **Insert row above** or **Insert row below** in the row options menu.
   - To delete a row, select **Delete row** in the row options menu.
   - To mark a row as reviewed, select **Mark row as reviewed** in the row options menu.
   - To make changes to multiple rows, select the check box in the first column of each row and select an action from the **Action on selected rows** list.

6. Review and complete the rows.
   - If you filled in the fields, ensure that all rows are marked as complete.
   - If the fields are auto-filled with values, then all rows are marked as reviewed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Extract fields using the draw tool

Use the draw tool to extract information directly from text and tables on the document image.

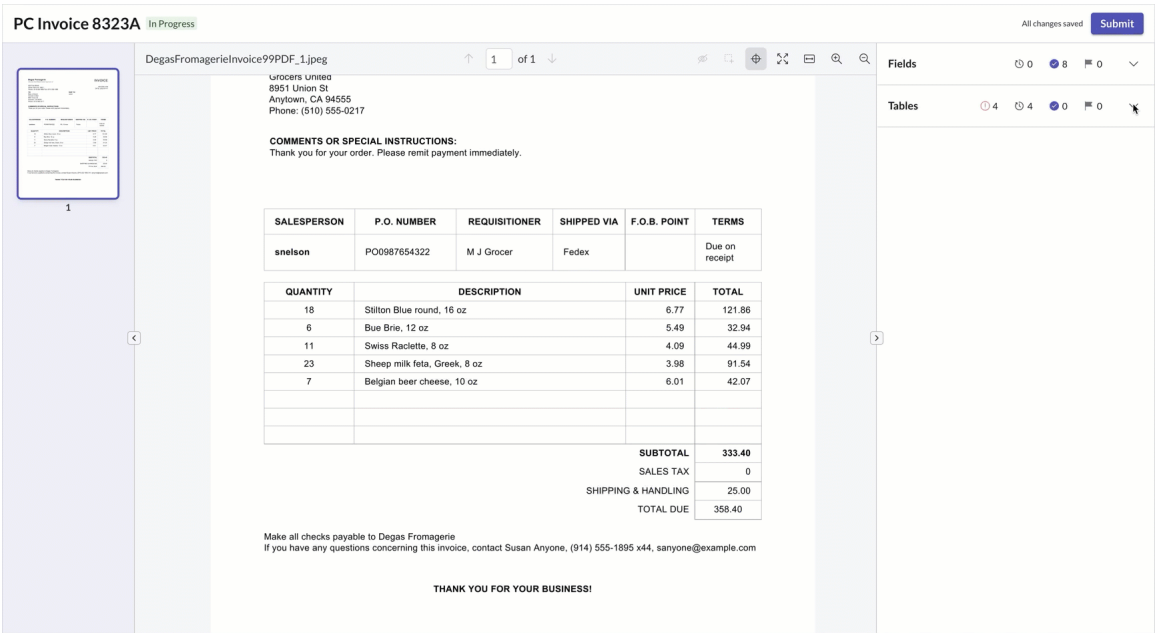
Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Overview of the draw tool

When viewing a processed document in the Document Intelligence workspace, you can use the draw tool to extract and review fields from the document image.

Especially in cases when there are no recommendations yet or incorrect recommendations, the draw tool is useful for selecting the specific area on the document you want to extract.

Any extracted table fields should already appear in a grid over the table. An extracted single field appears as a labeled box over the text. You can review the extracted fields right from the grid or box and make adjustments as needed. If the table fields are empty or there is no grid, you can create one or more grids over the table to extract the fields.



Extract a single field using the draw tool

Use the draw tool to extract a single field directly from the document image.

Before you begin

Role required: sn_docintel.extraction_agent

About this task

Use the following steps to extract a single field from a document using the draw tool. The extracted data is used to fill the related field in the document fields panel.

If the field is auto-filled, a box already appears over the extracted text. You can adjust the field as needed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In the document fields panel, expand **Fields**.
2. Select a field.

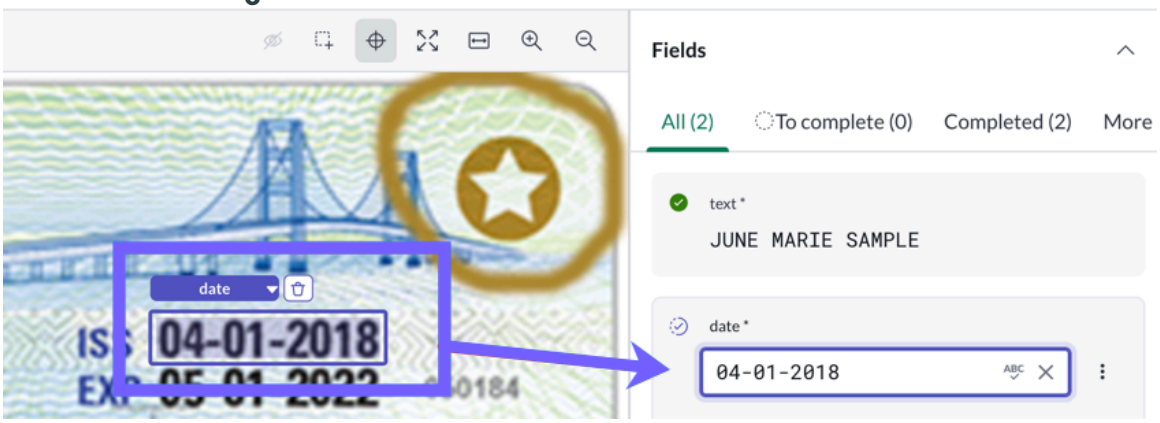
   The Draw Tool icon ( ) is enabled in the document controls toolbar.

   If the field is already auto-filled and the box appears over the field, or if you've already drawn a box, proceed to step 5.

3. Select the Draw Tool icon ( ).
4. Draw a box over the area on the document image that you want to extract.

   An editable box appears over the text. The field name is displayed next to the box. The extracted text in the box appears highlighted.

   Box drawn on a single field

   The extracted text fills the related field in the document fields panel.

   Tip: Press **Enter** or **Return** in the filled field to move to the next field and draw a new box.

5. **Optional:** Adjust the box as needed.

Option	Description
Resize the box	(Optional) Select the box. Drag the box edge or corner to the desired position. The updated text in the box is extracted and fills the related field in the document fields panel.
Remove the box	(Optional) Select the Remove the box icon ( ). The box is deleted and the extracted value is removed from the field.
Extract to a different field	(Optional) Select the field name next to the box and choose another option from the list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. **Optional:** Edit the extracted text by entering the changes in the field.
   DocIntel may not detect all the text within the selected area. In this case, you may need to enter the information directly in the field.

Extract a table using the draw tool

Use the draw tool to extract a table directly from the document image.

Before you begin

Role required: sn_docintel.extraction_agent

About this task

Use the following steps to extract a table from a document using the draw tool. The extracted data is used to fill the related table fields.

If the table fields are auto-filled, a draw tool grid is already available based on the extracted fields. You can adjust the fields as needed. For more information, see Adjust the draw tool grid.

You can select multiple sections of a table throughout the pages of a document.

Procedure

1. In the document fields panel, expand the Tables section.
   The Tables section only displays fields assigned to a table field group.
2. Select a table in the Tables section.

   The table panel is displayed, and the Draw Tool icon ( ) is enabled in the document controls toolbar.

   If the table fields are already auto-filled and the grid appears over the table, or if you've already drawn a grid, proceed to step 5.

   Tip: Collapse the Table panel using the Collapse table panel icon ( ), if needed.

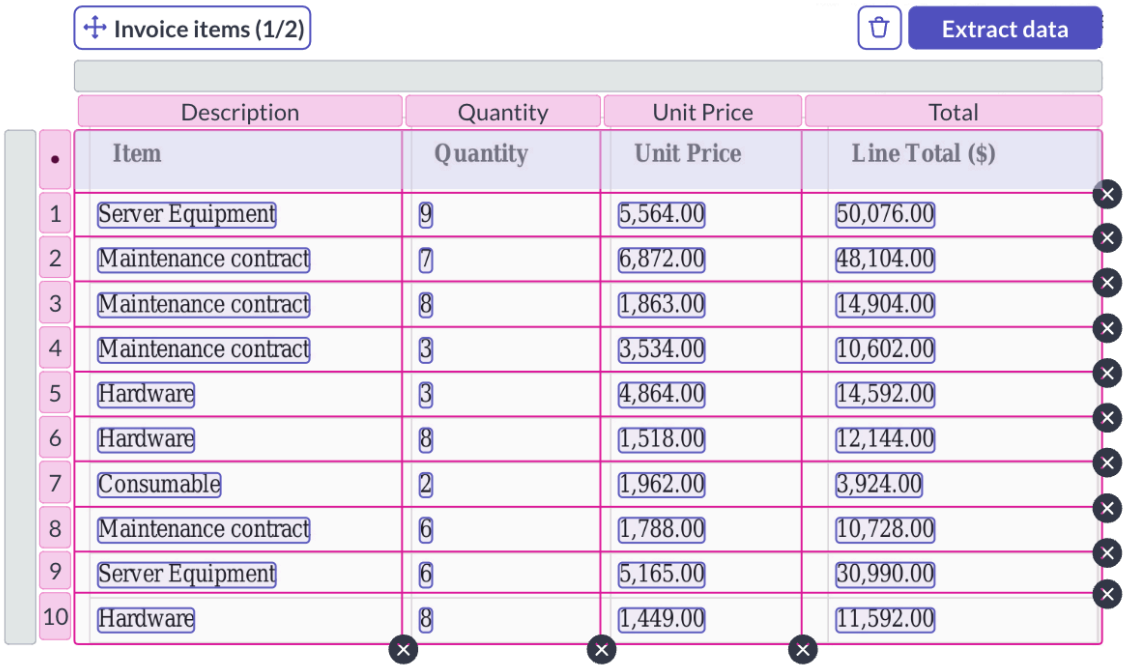
3. Select the Draw Tool icon ( ).
4. Draw a box over the area on the document image that you want to extract.

   Tip: Include the header row of the table in the selection area.

   An editable grid appears over the table. The grid defines the table cells in the area that you selected. The row numbers in the grid correspond to the rows in the Table panel.

   The text in the cell appears surrounded by a dotted-line box before it is extracted to the table fields. After it is extracted, it appears in a solid-line box.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



5. Adjust the grid as needed.

   For more information, see Adjust the draw tool grid.

6. **Optional:** To select another section of the table on the same or a different page of the document:

   a. Select the Draw Tool icon ( ).

   b. Draw a box over the other area on the document image that you want to extract.

      Note: It isn't possible to draw a new grid over an existing one.

      The grid number beside the table name shows the updated grid number. For example, "Line items (2/2)" shows you are in the second out of two grids used to extract the Line items table fields.

   c. Adjust the grid as needed.

7. **Optional:** To hide the grid(s), select the Hide/show the grid(s) icon ( ).

   (Optional) A partly-transparent box shows where a hidden grid is located on the document.

   Select the hidden grid or the Hide/show the grid(s) icon to show the grid.

8. Extract the data.

   - Select **Extract data**.

     This button displays if no data has been extracted yet.

   - Select **Update fields**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This button displays if the table fields are already filled and you have made changes that will overwrite some of them.

The **Update fields** button shows the number of fields to be updated.

The data in the grid(s) are extracted and used to fill the related table fields.

If there are any empty cells or missing columns in the grid(s), the related table fields are set to **Missing in the document**.

Adjust the draw tool grid

Adjust the draw tool grid to better fit the information that you want to extract from the document image.

Before you begin

- These steps apply to tables that have a draw tool grid on the document image. For more information, see Extract a table using the draw tool.
- Role required: sn_docintel.extraction_agent

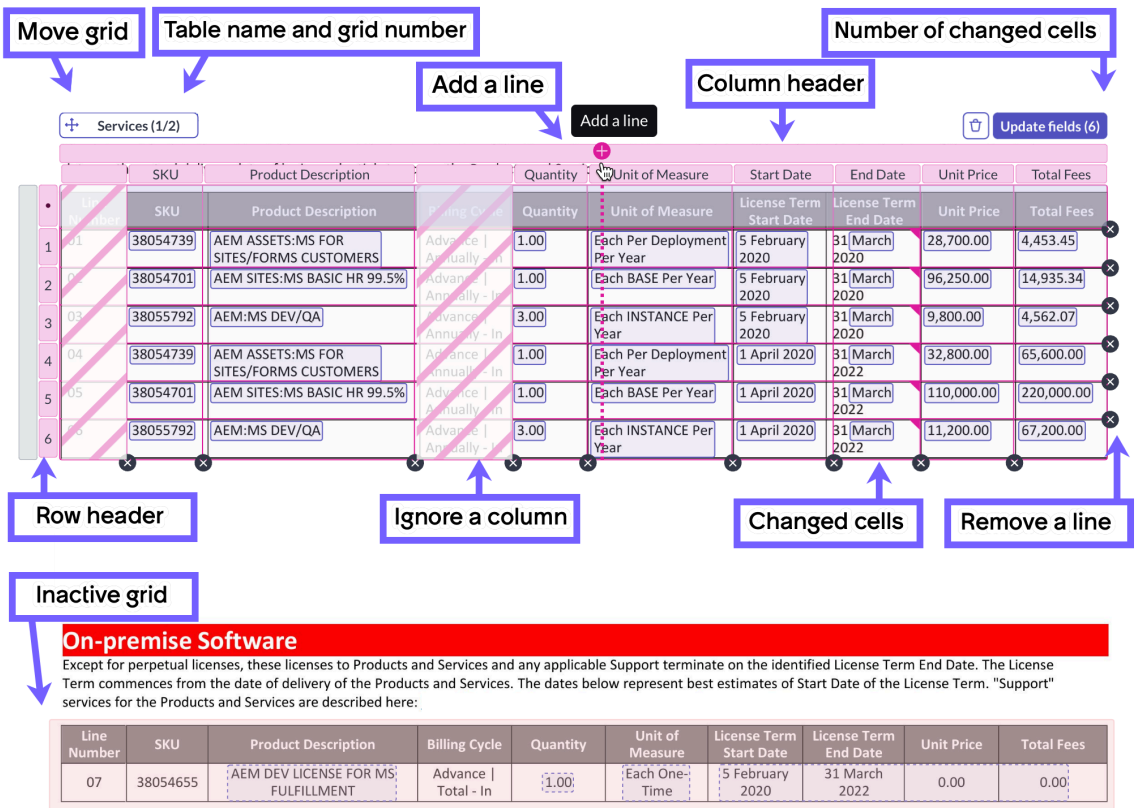
About this task

Use the following steps to make changes to a grid. You can then extract the updates to the related table fields.

Procedure

1. On the document image, select the grid that you want to change.

   If you have more than one grid on the document, only one displays editing tools. The other grid appears as an overlay until you select it for editing.



	SKU	Product Description	Billing Cycle	Quantity	Unit of Measure	License Term Start Date	License Term End Date	Unit Price	Total Fees
1	38054739	AEM ASSETS:MS FOR SITES/FORMS CUSTOMERS	Advance | Annually - In	1.00	Each Per Deployment Per Year	5 February 2020	31 March 2020	28,700.00	4,453.45
2	38054701	AEM SITES:MS BASIC HR 99.5%	Advance | Annually - In	1.00	Each BASE Per Year	5 February 2020	31 March 2020	96,250.00	14,935.34
3	38055792	AEM:MS DEV/QA	Advance | Annually - In	3.00	Each INSTANCE Per Year	5 February 2020	31 March 2020	9,800.00	4,562.07
4	38054739	AEM ASSETS:MS FOR SITES/FORMS CUSTOMERS	Advance | Annually - In	1.00	Each Per Deployment Per Year	1 April 2020	31 March 2022	32,800.00	65,600.00
5	38054701	AEM SITES:MS BASIC HR 99.5%	Advance | Annually - In	1.00	Each BASE Per Year	1 April 2020	31 March 2022	110,000.00	220,000.00
6	38055792	AEM:MS DEV/QA	Advance | Annually - In	3.00	Each INSTANCE Per Year	1 April 2020	31 March 2022	11,200.00	67,200.00

Line Number	SKU	Product Description	Billing Cycle	Quantity	Unit of Measure	License Term Start Date	License Term End Date	Unit Price	Total Fees
07	38054655	AEM DEV LICENSE FOR MS FULFILLMENT	Advance | Total - In	1.00	Each One-Time	5 February 2020	31 March 2022	0.00	0.00

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Adjust the grid as needed.

   Adjustments are auto-saved.

   Any adjustments that produce a change to previous extractions are indicated in the cell corner.

   The text in the cell appears surrounded by a solid-line box if it's already extracted. It appears in a dotted-line box if it isn't extracted yet.

Option	Description
Move the grid	Drag the Move icon ( ) to the desired posi tion.
Change a column header	Select the column header and select the col umn where the data will be extracted to in the table fields.
Make the first row a header row	Select the row header and check the **This is a header row** check box.
Move a column or row border	Drag the border to the desired position.
Add a new column or row	Go to the outside border of the grid and se lect the Add a line icon ( ). Adding a line splits the affected cells so they can be extracted to separate fields.
Remove a column or row border	Select the Remove a line icon ( ) at the end of the border outside the grid. Removing a line merges the affected cells so they can be extracted to the same fields.
Remove the grid	Select the Remove the grid icon ( ). This action only removes the grid, not the ex tracted data.
Ignore a table column or row	**a.** Select the column or row heading. **b.** Select the **Do not extract this column/row** check box. **c.** Select **Save**.

Review a field on the document image

Use the draw tool to review auto-filled fields directly from the document image panel.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- These steps apply to tables that have a draw tool grid or single fields that have a draw tool box on the document. For more information, see Extract a table using the draw tool or Extract a single field using the draw tool.
- Role required: sn_docintel.extraction_agent

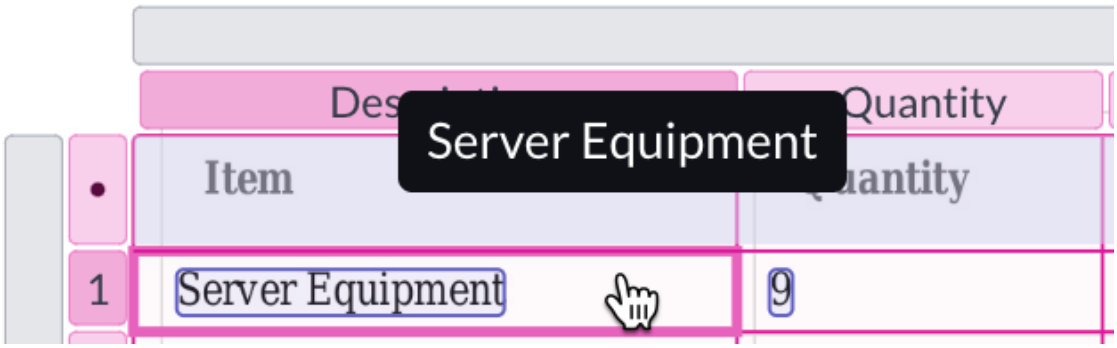
Procedure

1. On the document image, point to a grid cell or field box you want to review.

   For a single field, an editable box appears over the text. The field name is displayed next to the box.



   For a table field, the text in the grid cell appears surrounded by a box if the value is extracted.



2. Verify the value that appears in the field's tooltip.
   - If the value is accurate, move on to review the next field.
   - If the value must be changed, edit the field.

     See Edit a field on the document image.

3. Repeat steps 1 and 2 for each field that you want to review.

Edit a field on the document image

Edit the fields on the document image to make sure that the correct information is extracted.

Before you begin

- These steps apply to tables that have a draw tool grid or single fields that have a draw tool box on the document. For more information, see Extract a table using the draw tool or Extract a single field using the draw tool.
- Role required: sn_docintel.extraction_agent

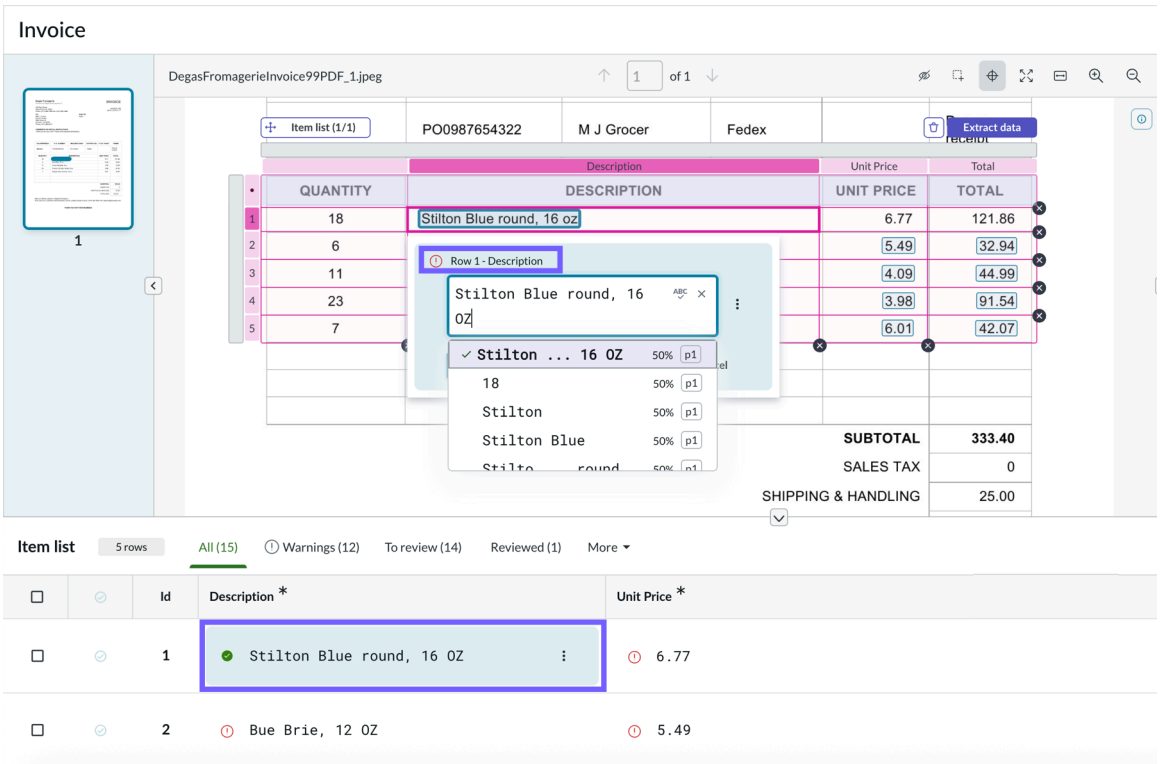
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Use the following steps to change to a field by editing the recommendation on the document image. The text detected by DocIntel appears surrounded by a box.

Procedure

1. On the document image, select the text box or the grid cell that you want to edit.
   A dialog box appears, showing the related field.
2. Enter a value in the text box or select another recommendation from the list.



3. Select **Save**.

Classify documents using the Document Intelligence workspace

Use the Document Intelligence workspace to label your documents. The workspace enables you to train the AI model by providing direct input and by validating or correcting the recommendations provided by DocIntel.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Overview of the Document Intelligence workspace

The Document Intelligence (DocIntel) workspace provides features that enable you to quickly apply classes or categories to the attached documents and each of their pages.

To get started with the Document Intelligence workspace, complete the following:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

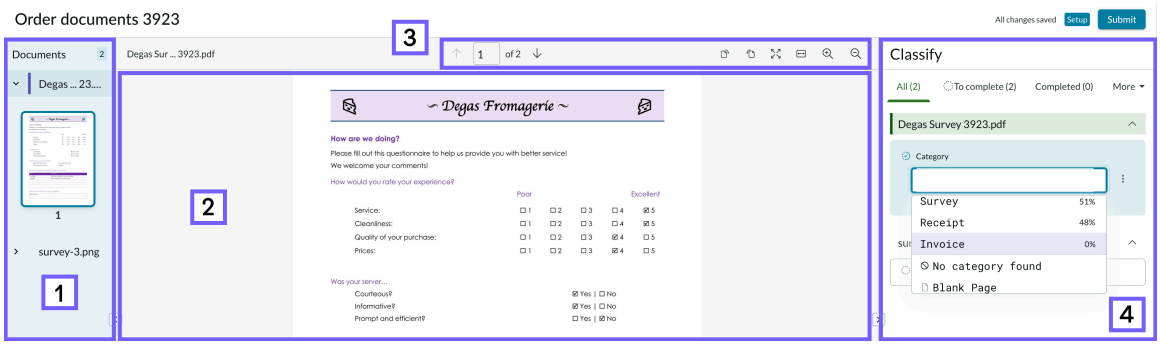
- Create a document task for a document classification use case, upload a document, and process the task.
- After DocIntel has processed the task, you can begin using the workspace. See Complete a document task.

Layout of the Document Intelligence workspace

The following image shows the Document Intelligence workspace for a document task. The workspace includes the following areas:

- 1 - Thumbnail panel
- 2 - Document image panel
- 3 - Document controls
- 4 - Classify panel

Note: In this view, the documents in the Classify panel are expanded so that the categories are visible.



Thumbnail panel

In the thumbnail panel of the workspace, you can select one page from one or more multiple-page documents. The selected page is displayed in the document image panel. Selecting a page doesn't affect what is displayed in the Classify panel.

Document image panel

The document image panel displays the page selected in the thumbnail panel.

Document controls

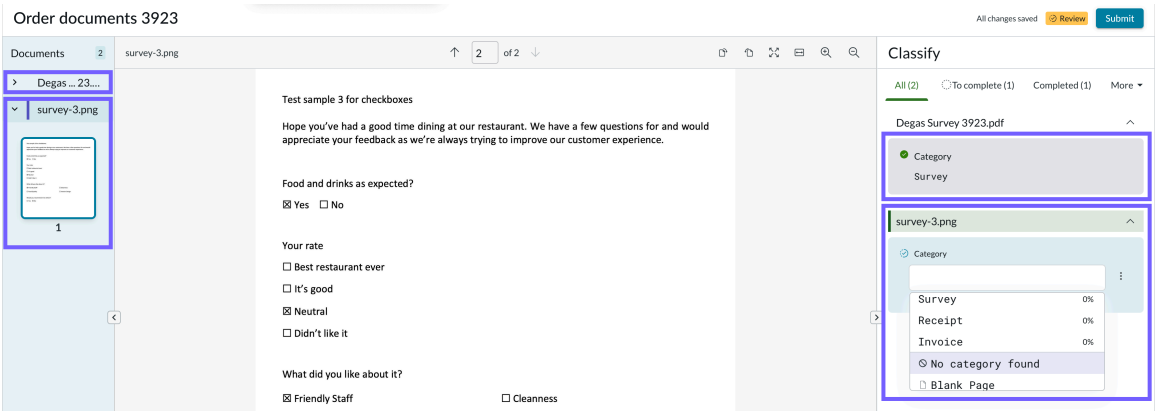
When you're reviewing a document for classification, you can use various controls to rotate, maximize the viewing area, zoom, or focus on the areas that you need.

Classify panel

The Classify panel enables you to open each document and page classification for review, including viewing the AI's recommendations.

The following image shows the Classify panel with a category field for each attachment in the document task, enabling you to enter or review the category.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Classify documents and document pages

Classify your documents in the Classify panel.

Before you begin

Role required: sn_docintel.extraction_agent

About this task

Use the following steps to apply categories to document attachments and each of their pages.

In this context, a document page is the individual page object of the digital document file, as displayed in the thumbnail panel of the workspace. A document file may contain multiple logical groupings of documents (like a PDF file with several invoices, each spanning several pages). Based on the contents, apply the categorization to the document file and the file pages.

If the categories are already auto-filled with values, you can review them to ensure they are correct or adjust them as needed.

Procedure

1. In the Classify panel, review any auto-filled categories for each document to ensure they're correct.
2. Select a category for a document by selecting the **Category** field.
3. Select an appropriate recommendation for the document using the recommendation list.

    - Begin typing in the **Category** field and select a recommendation from the list.

      The recommendation with the highest confidence score displays at the top of the list.

    - Select **Mixed categories** for a multi-page document that includes more than one document type.

      Otherwise, if you select a category for a multi-page document, that category is applied to all of the document's pages.

    - Select **Blank Page** for a blank page in a document.

    - Select **No category found** if you can't find a correct match.

      The document will be marked as complete when you move to the next **Category** field.

4. Select a recommendation for each page of a Mixed categories document.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The page will be marked as complete when you move to the next page.

5. **Optional:** If needed, flag the documents or pages for later attention by selecting **Flag for follow-up** in the options menu from the options icon ( ).

Monitoring Document Intelligence performance

Track document extraction performance in Document Intelligence to understand its usage and effectiveness.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

View reports on the Document Intelligence Admin home page

Monitor document extraction performance in the Admin experience.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Before you begin

- Ensure that the Document Intelligence application (sn_docintel) and Document Intelligence Admin (com.snc.docintel_admin) ServiceNow® Store application is installed and active. For more information, see Install Document Intelligence.
- Have an active use case with multiple completed document tasks. For more information, see Set up document extraction use cases.
- Role required: sn_docintel.admin, sn_docintel.manager, or admin.

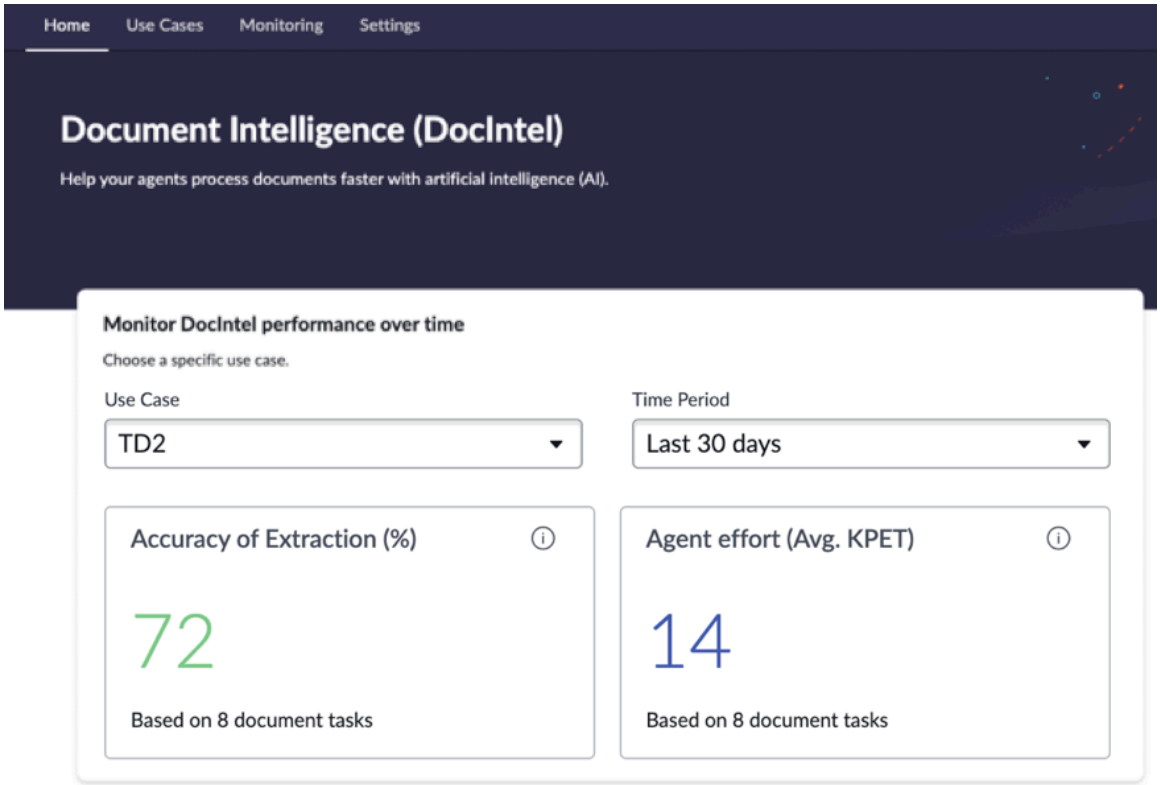
About this task

You can review the value of your Document Intelligence (DocIntel) implementation when you open the Document Intelligence Admin experience home page.

The **Monitor DocIntel performance over time** section displays the following measurements:

- The Accuracy of Extraction widget shows the average extraction accuracy per time period for the selected use case. Accuracy is defined as the number of times that the AI's top recommendation is the correct answer.
- The Agent effort widget shows the number of keystrokes that your agents need to perform in order to extract all field values for a document task. This measurement is an average per document task.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Home**.
2. Expand the **Use Case** list and select your use case.
3. Expand the **Time Period** list and select a date range.
4. Review the results displayed in the **Accuracy of Extraction** and the **Agent effort** widgets.

Document Intelligence monitoring dashboard

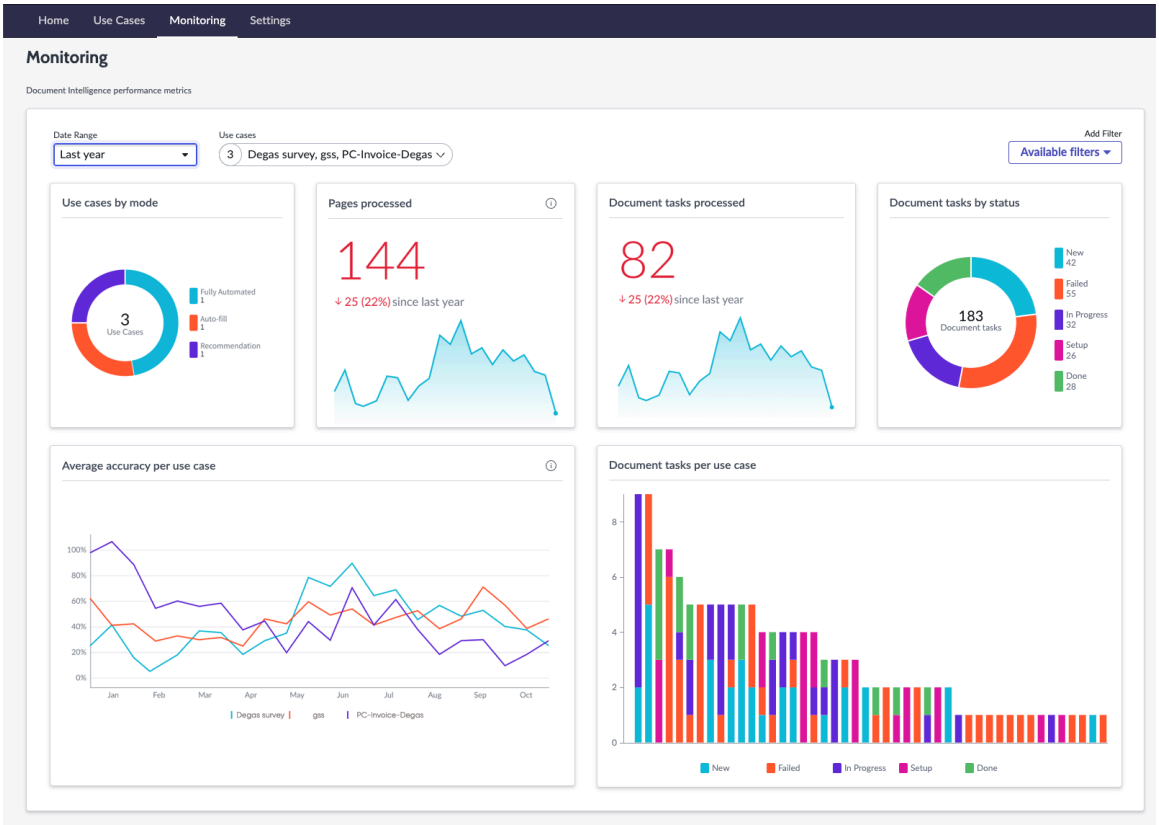
Monitor the overall performance of Document Intelligence over time in the Document Intelligence monitoring dashboard.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Overview of the Document Intelligence monitoring dashboard

The Document Intelligence monitoring dashboard provides a high-level overview of your Document Intelligence usage and value.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



The data visualizations show document extraction activity in your instance. For example:

- Active use cases
- Active document tasks
- Processed pages
- Processed document tasks
- Accuracy of the DocIntel recommendations

This dashboard provides useful answers to the following questions:

- How many documents are processed using DocIntel?
- How much of the document extraction is automated?
- How has DocIntel accuracy progressed over time?

View the Document Intelligence monitoring dashboard

Access the Document Intelligence monitoring dashboard on the Monitoring screen of the Admin experience.

Before you begin

- Ensure that the Document Intelligence application (sn_docintel) and Document Intelligence Admin (com.snc.docintel_admin) ServiceNow® Store application is installed and active. For more information, see Install Document Intelligence.
- Have an active use case with multiple completed document tasks. For more information, see Set up document extraction use cases.
- Role required: sn_docintel.admin, sn_docintel.manager, or admin.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Monitoring**.
2. **Optional:** Select a date range in the **Date range** filter to show only the data within the selected dates.
3. **Optional:** Select one or more use cases in the **Use cases** filter to show only the data that applies to the selected use cases.
4. **Optional:** Select an object within a chart to see additional details.

Data visualizations in the Document Intelligence monitoring dashboard

The Document Intelligence monitoring dashboard uses data visualizations to display your Document Intelligence (DocIntel) usage and performance data.

The following table describes the data visualizations shown on the Monitoring screen.

Data visualizations showing Document Intelligence usage and performance

Name	Description
Use cases by mode	The number of document extraction use cases in this instance, grouped by extraction mode.
Pages processed	The number of pages processed for data extraction using DocIntel in this instance over the selected date range.
Document tasks processed	The number of DocIntel document tasks processed in this instance over the selected date range.
Document tasks by status	The number of DocIntel document tasks in this instance, grouped by status.
Average accuracy per use case	The percentage of times that the top recommendation from DocIntel is the correct field value, based on the tasks completed for a use case.
Document tasks per use case	The number of DocIntel document tasks for each use case, grouped by status.

Use case performance dashboard

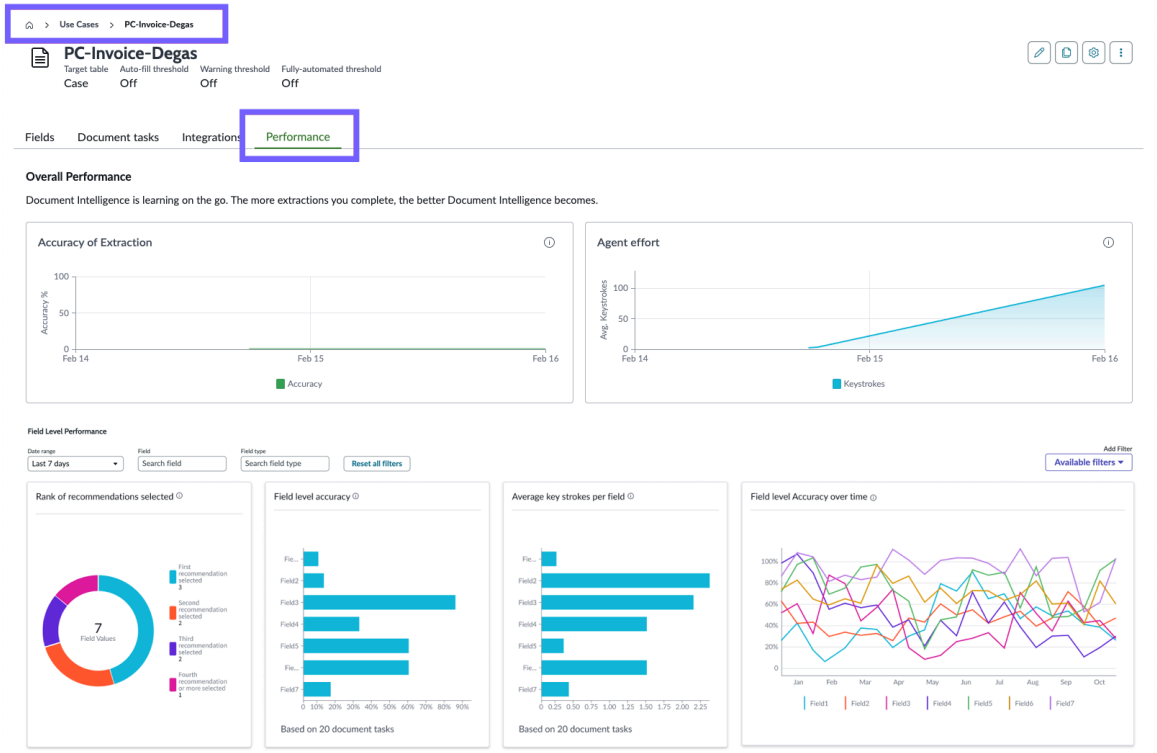
Monitor Document Intelligence (DocIntel) performance at the use case and field levels in the use case performance dashboard.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Overview of the use case performance dashboard

The use case performance dashboard provides quality metrics for the use case.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



The data visualizations show document extraction activity for the use case. For example:

- Agent effort to process document tasks
- Accuracy of the DocIntel recommendations

This dashboard provides useful answers to the following questions:

- How much effort is needed to complete document tasks for this use case?
- How has DocIntel accuracy progressed for this use case?

View the use case performance dashboard

Access the use case performance dashboard in the Performance tab on the use case screen.

Before you begin

- Ensure that the Document Intelligence application (sn_docintel) and Document Intelligence Admin (com.snc.docintel_admin) ServiceNow® Store application is installed and active. For more information, see Install Document Intelligence.
- Have an active use case with multiple completed document tasks. For more information, see Set up document extraction use cases.
- Role required: sn_docintel.admin, sn_docintel.manager, or admin.

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Use Cases**.
2. Select the **Performance** tab.
3. **Optional:** Select a date range in the **Date range** filter to show only the data within the selected dates.
4. **Optional:** Select one or more fields in the **Fields** filter to show only the data that applies to the selected fields.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. **Optional:** Select one or more field types in the **Field type** filter to show only the data that applies to the selected field types.
6. **Optional:** Select an object within a data visualization to see additional details.

Data visualizations in the use case performance dashboard

The use case performance dashboard provides data visualizations to display your usage and performance data for the use case.

The following table describes the data visualizations shown on the **Performance** tab.

Data visualizations showing Document Intelligence use case activity

Name	Description
Accuracy of extraction	The average extraction accuracy per time period for the use case. Accuracy is defined as the number of times that the AI's top recommendation is the correct answer.
Agent effort	The number of keystrokes your agents need to press in order to extract all field values for a document task. This measurement is an average per document task.
Rank of recommendations selected	The number of times that DocIntel recommendations from were selected by an agent when completing a document task, grouped by the order presented.
Field level accuracy	The percentage of times that the top recommendation from DocIntel is the correct value for a field.
Average key strokes per field	The average number of keystrokes an agent performs to extract a value for a field.
Field level accuracy over time	The percentage of times over the selected date range that the top recommendation from DocIntel is the correct value for a field.

Document Intelligence references

The following topics provide additional information about the features and properties installed with Document Intelligence.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Components installed with Document Intelligence

Several types of components are installed with activation of the Document Intelligence plugin, including tables and user roles.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Roles installed with Document Intelligence

For more information on roles, see Document Intelligence roles.

Role title [name]	Description	Contains roles
DocIntel Admin [sn_docintel.admin]	Has full access to the Document Intelligence application, except for modifying a subset of system properties, and the billing and internal tables.	• platform_ml_di.admin • action_designer • flow_designer • sn_ace.ace_user • canvas_user • usage_admin
DocIntel Viewer [sn_docintel.viewer]	Has view-only access on Document Intelligence document tasks that they are authorized to view.	• snc_read_only • platform_ml_di.viewer
DocIntel Extraction Agent [sn_docintel.extraction_agent]	Extracts data and text from documents using the Document Intelligence workspace.	platform_ml_di.extraction_agent
DocIntel Creation Agent [sn_docintel.creation_agent]	Extracts information from documents using the Document Intelligence workspace. Also enables users to create Document Intelligence document tasks and submit them for processing.	platform_ml_di.creation_agent
DocIntel Manager [sn_docintel.manager]	Creates and edits use cases, fields, field groups, and document tasks. Views, measures, and analyzes the usage and effectiveness of Document Intelligence using the Platform Document Intelligence Usage dashboard. Grants access to submit document tasks and interact with the Document Intelligence workspace.	• platform_ml_di.manager • action_designer • flow_designer • sn_ace.ace_user • canvas_user • usage_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tables installed with Document Intelligence

Note: Starting in Document Intelligence 3.0, DocIntel uses ServiceNow AI Platform database tables (sys_di_) in place of the scoped application tables (di_). See Upgrade to Document Intelligence 3.0 or later from version 2.4 or earlier.

Table	Description
Billable Event [sys_di_billable_event]	Contains all the billable events for the instance. A billable event corresponds to pages that have been processed using Document Intelligence.
Recommendation Meta Info [sys_di_candidate_meta_info]	[Internal table] Contains the data returned by the AI for field group extraction and for missing value prediction.
Recommendation Score [sys_di_candidate_score]	[Internal table] Contains the scores for each recommendation calculated by the AI. There's one record per field, for each page of a document.
Field Value [sys_di_extracted_value]	Contains all the data extracted in the instance, across all document tasks.
Image [sys_di_image]	[Internal table] Contains information about each page of a document. There's one record per document page.
Integration Setup [sys_di_integration_setup]	Contains DocIntel use case integrations.
Field [sys_di_key]	Contains all the fields created in the instance, across all use cases. A field corresponds to a recommendation to be extracted from documents.
Field Group [sys_di_key_group]	Contains all the field groups created in the instance, across all use cases. A field group is usually created to help extract data from tables and lists.
Lock [sys_di_lock]	Contains an index of locked task solution definitions. Used to improve the performance of scheduled jobs.
DocIntel Aggregated Metrics [sys_di_metrics_aggregated]	[Internal table] Contains aggregated metrics. Aggregation may happen multiple times per day, based on system properties.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
DocIntel Daily Metrics [sys_di_metrics_daily]	[Internal table] Contains the metrics aggregated daily. There's one record per aggregated metric.
DocIntel Metrics Job Log [sys_di_metrics_job_log]	[Internal table] Contains a log of the metrics daily aggregation jobs that happened in the instance.
DocIntel Raw Metrics [sys_di_metrics_raw]	[Internal table] Contains a list of raw metrics collected by Document Intelligence. The records are deleted when metrics aggregation happens.
DocIntel OCR Input [sys_di_ocr_input]	[Internal table] Contains image data to be sent to the OCR module for processing. The records are deleted when processing is complete.
PDF [sys_di_pdf]	[Internal table] Contains a list of PDF files stored across tasks.
DocIntel PDF Input [sys_di_pdf_input]	[Internal table] Contains PDF data to be sent to the PDF module for processing. The records are deleted when processing is complete.
DocIntel Prediction Input [sys_di_prediction_input]	[Internal table] Contains the data needed to make suggestions across all fields for a given document task. There's one record per document page.
Document Task [sys_di_task]	Contains all the document tasks created in the instance, across all use cases. A document task contains one or more documents from which fields must be extracted.
Task Solution Definition [sys_di_task_def_solution_def]	Contains the solution definitions related to DocIntel use cases.
DocIntel Use Case [sys_di_task_definition]	Contains all the document processing use cases created in the instance. A use case defines what and how data should be extracted from a set of documents.
DocIntel Training Input [sys_di_training_input]	[Internal table] Contains the data needed to improve the AI models.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Confidence scores

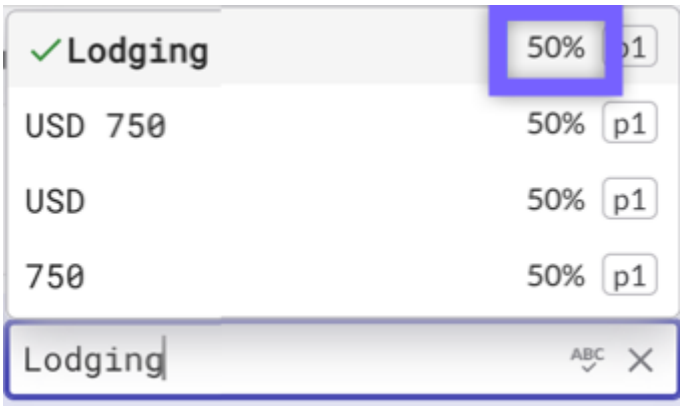
A confidence score is a measurement (percentage) of how reliable DocIntel is in providing a recommendation for a field. The higher the score, the more reliable the recommendation.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Confidence scores increase as the AI model is trained through user input when processing document tasks.

A score is displayed next to each recommendation.

Confidence score



In some areas, the score is color-coded with a different color for each confidence score range.

- Green indicates high confidence (76%-100%)
- Yellow indicates medium confidence (50%-75%)
- Red indicates low confidence (0%-49%)

To configure confidence score thresholds, see Configure data extraction modes.

Data extraction modes

Extraction modes determine how the data is extracted in the document task and how the task is processed. The mode changes the behavior of the fields in the Document Intelligence workspace.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

DocIntel uses the following extraction modes.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Extraction mode	Description
Recommendation	DocIntel provides recommendations for the fields in the Document Intelligence workspace. Choose the recommendation or manually enter the value. All fields must be reviewed. Recommendations are ordered based on how confident the AI is in the prediction. As DocIntel continues processing your documents, recommendations improve over time.
Auto-fill	DocIntel auto-fills the fields in the Document Intelligence workspace. All fields must be reviewed. Auto-fill works only if the AI has enough confidence to make the prediction. You can change the confidence threshold by updating the **Auto-fill threshold** field in the use case.
Fully automated (Straight-through processing)	DocIntel automatically extracts the data for all fields and processes the document task if the confidence scores for all required fields are above the defined confidence threshold. Fields don't need to be reviewed. **Note:** For Fully automated mode to process tasks successfully, the use case must have at least one field designated as required. DocIntel becomes more confident over time, as it processes more and more documents. Choose Fully automated mode for frequently processed documents or if you're confident in the system.

Data normalization

Certain types of data extracted from documents are converted into a standard format so that they appear the same across all fields.

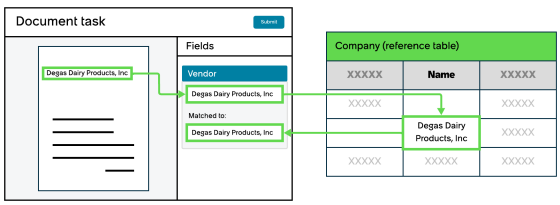
Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

This process increases the usefulness of the data by enabling it to be grouped and analyzed more easily. It also supports integration with other applications on the ServiceNow AI Platform.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field types

The following field types are converted to support data normalization:

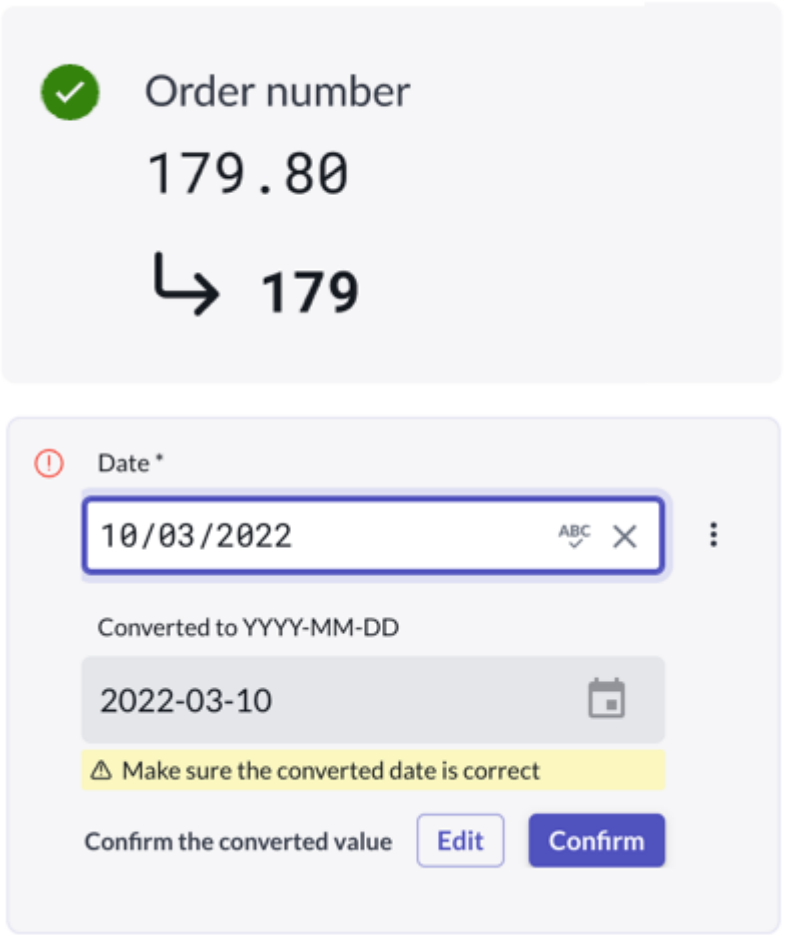
Field type	Description
Date	Standard date format. For example, YYYY-MM-DD.
Reference field	A field that uses a field in another table as a standard. DocIntel matches the extracted data to the standard. For example, a use case has a reference field called **Vendor** that points to the Name column in the Company table as the reference. When processing a document task, DocIntel extracts “Degas Dairy Products, Inc” from the document and fills the **Vendor** field with that value. DocIntel compares the value to the company names in the reference table and finds “Degas Dairy Products, Inc” as a match. In the document task, “Degas Dairy Products, Inc” is matched to “Degas Dairy Products, Inc” in the reference. 
Integer	Whole number. For example, 12.
Decimal	Number with up to two decimal places. For example, 12.5 or 12.55.
Floating point number	Number with up to seven decimal places. For example, 12.0 to 12.0000000.

To set the field type, see Create a field for data extraction.

Display

A completed data extraction field shows the converted value next to it.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



You can adjust the converted date value by selecting **Edit**.

Note: In some cases, the data extracted from the document may not be in a valid format to be converted. For example, if DocIntel read the letter O instead of a number 0 in a date field (11.12.2o23), then it would not be converted. In this case, edit the field to the correct format.

Ambiguous data

If there is data in a document that can be understood in more than one way, DocIntel interprets that value based on the default selected for it in the use case configuration. DocIntel must interpret an ambiguous value in order to accurately convert it to the normalized format.

For example, a use case has a **Date** field, and `Month first` is selected as the default order to interpret ambiguous dates. When a document containing the date 1/2/2024 is processed for the use case, DocIntel interprets that date as January 2, not February 1, when it extracts that value and converts it.

In such cases, the user completing a document task may need to confirm or correct the conversion of ambiguous values. Depending on the field's configuration in the use case, automated document processing may be interrupted to ensure the conversion is accurate.

Document field statuses

The following is a list of the statuses for fields in DocIntel document tasks. These statuses apply to fields for both document classification and data extraction.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Statuses for document fields

Status	Icon	Description
To complete		The field must be filled in with a value or marked as missing in the document.
Completed		The field is filled in with a value or marked as missing in the document.
To review		The field is auto-filled and needs review by a user.
Reviewed		The auto-filled field has been reviewed by a user.
In progress		The field is in the process of being filled in or reviewed.
Needs attention		The auto-filled field has triggered a warning. The causes for a warning include: • A required field must be filled in or marked as missing in the document. • An extracted value has a low confidence score and should be reviewed. • An extracted value is ambiguous and should be verified.

Document Intelligence forms

Use forms to view and update Document Intelligence information.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Check box list form

The Check box list form enables you to define a check box list for extraction.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

The Check box list form includes the following fields.

Check box list form

Field	Description
Check box list	
Check box list name	The name for the check box list as it appears in the Document Intelligence workspace.
Target table	The table that stores the document processing results for the check box list.
Parent mapping to field	Field on the target table you want to align this check box list with. **Note:** You must first select a target table.
Check boxes	
Check box title	The name for the check box as it appears in the Document Intelligence workspace.
Select target field	Field on the target table that you want to align this field with. This field is used for integration with other applications. See Integrate with a custom application or workflow.
This field is required for extraction	Option to make a field required. Required fields can't be left unreviewed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Check box list form (continued)

Field	Description
	Required fields affect how document tasks are processed in the Fully automated extraction mode. For more information, see Configure data extraction modes.
New check box	Option to add a check box to the list. Use the reorder icon (⠿) to reorder a check box in the list.
Create multiple check box lists	Option to keep the pop-up window displayed on the screen. Enable this option if you're adding more than one check box list to the use case.

Single field form

The Single field form enables you to define a single field for extraction.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

The Single field form includes the following fields.

Single field form

Field	Description
Single field name	The name for the field as it appears in the Document Intelligence workspace.
Type	The type of the field. For example, a text or date field. Some field types convert the extracted value into a standard format. For more information, see Data normalization.
Converted date format	The date format that the extracted value is converted to for data normalization. This field is available when the **Type** field is set to `Date`.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Single field form (continued)

Field	Description
Target table	The table that stores the document processing results for this use case.
Target field	Field on the target table that you want to align this field with. **Note:** The use case must have a target table selected. This field is used for integration with other applications. For more information, see Integrate with a custom application or workflow.
When the date is ambiguous in a document, DocIntel will interpret it in the following order	The default interpretation of the date format. For example, if you select `Month first` in this field, DocIntel will interpret an ambiguous date like 1/2/2024 as January 2 when it extracts that value from a document. If you select `Day first`, it will interpret it as February 1. This field is available when the **Type** field is set to `Date`.
When the number is ambiguous in a document, DocIntel will interpret it as	The default interpretation of the number format. For example, if you select `1,00` in this field, DocIntel will interpret an ambiguous number like 5 as 5,00 instead of 5.00 when it extracts that value from a document. This field is available when the **Type** field is set to `Integer`, `Decimal`, or `Float`.
Reference table	The table that stores the reference column. It's automatically populated based on the selected target field. This field displays when the **Type** field is set to `Reference field`.
Reference column	The column in the reference table that contains the referenced data. DocIntel uses the reference column to find data that matches the extracted field value

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Single field form (continued)

Field	Description
	when processing a document task. The field value is then converted to the format of the reference. For more information on converted values, see Data normalization. This field is available when the **Type** field is set to `Reference field`.
Distinguisher(s)	Additional columns in the reference table that help the user to distinguish between similar records. This field is available when the **Type** field is set to `Reference field`.
This single field is required for extraction	Option to make the field required. Required fields can't be left empty or unreviewed. They also can't contain ambiguous values. An ambiguous value is a field entry that can be interpreted in more than one way. If it's a reference field type, the required field must have a valid, exact match. By default, DocIntel uses the first matched record. Required fields affect how document tasks are processed in the Fully automated extraction mode. For more information, see Configure data extraction modes.
Select one of these options for cases when: • the date/number is ambiguous in the document • there are multiple reference matches in the document	Option for agent review in situations when DocIntel encounters an ambiguous value in a required field. In such cases, the selected default interpretation will apply to the extracted value. The option is whether to interrupt full automation of document tasks to allow agents to verify the interpreted values. Otherwise, continue automatic processing of document tasks without the agent review.
Create multiple single fields	Option to keep the pop-up window displayed on the screen. Enable this field if you're adding more than one single field to the use case.

Single field group form

The Single field group form enables you to define a related group of single fields for extraction.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

The Single field group form includes the following fields.

Single field group form

Field	Description
Field group	
Field group name	The name for the field group as it appears in the Document Intelligence workspace.
Target table	The table that stores the document processing results for the fields.
Parent mapping to field	Field on the target table that you want to align this field group with. **Note:** You must first select a target table.
This field group is required for extraction	Option to make the field as required. Required fields can't be left empty or unreviewed. They also can't contain ambiguous values. An ambiguous value is a field entry that can be interpreted in more than one way. If it's a reference field type, the required field must have a valid, exact match. By default, DocIntel uses the first matched record. Required fields affect how document tasks are processed in the Fully automated extraction mode. For more information, see Configure data extraction modes.
Fields	
Field name	The name for the field as it appears in the Document Intelligence workspace.
Type	The type of the field. For example, a text or date field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Single field group form (continued)

Field	Description
	Some field types convert the extracted value into a standard format. For more information, see Data normalization.
Select Target Field	Field on the target table that you want to align this field with. This field is used for integration with other applications. For more information, see Integrate with a custom application or workflow.
When the date is ambiguous in a document, DocIntel will interpret it in the following order	The default interpretation of the date format. For example, if you select `Month first` in this field, DocIntel will interpret an ambiguous date like 1/2/2024 as January 2 when it extracts that value from a document. If you select `Day first`, it will interpret it as February 1. This field is available when the **Type** field is set to `Date`.
When the number is ambiguous in a document, DocIntel will interpret it as	The default interpretation of the number format. For example, if you select 1,00 in this field, DocIntel will interpret an ambiguous number like 5 as 5,00 instead of 5.00 when it extracts that value from a document. This field is available when the **Type** field is set to `Integer`, `Decimal`, or `Float`.
Reference table	The table that stores the reference column. It's automatically populated based on the selected target field. This field displays when the **Type** field is set to `Reference field`.
Reference column	The column in the reference table that contains the referenced data. DocIntel uses the reference column to find data that matches the extracted field value when processing a document task. The field value is then converted to the format of the

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Single field group form (continued)

Field	Description
	reference. For more information on converted values, see Data normalization. This field is available when the **Type** field is set to `Reference field`.
Distinguisher(s)	Additional columns in the reference table that help the user to distinguish between similar records. This field is available when the **Type** field is set to `Reference field`.
This field is required for extraction	Option to make a field required. Required fields can't be left empty or unreviewed. They also can't contain ambiguous values. An ambiguous value is a field entry that can be interpreted in more than one way. If it's a reference field type, the required field must have a valid, exact match. By default, DocIntel uses the first matched record. Required fields affect how document tasks are processed in the Fully automated extraction mode. For more information, see Configure data extraction modes.
Select one of these options for cases when: • the date/number is ambiguous in the document • there are multiple reference matches in the document	Option for agent review in situations when DocIntel encounters an ambiguous value in a required field. In such cases, the selected default interpretation will apply to the extracted value. The option is whether to interrupt full automation of document tasks to allow agents to verify the interpreted values. Otherwise, continue automatic processing of document tasks without the agent review.
New single field	Option to add a field to the group. Use the reorder icon (⠿) to reorder a field in the group.
Create multiple field groups	Option to keep the pop-up window displayed on the screen. Enable this option if you're

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Single field group form (continued)

Field	Description
	adding more than one single field group to the use case.

Table form

The Table form enables you to define a table for extraction.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

The Table form includes the following fields.

Table form

Field	Description
Table	
Table name	The name for the table as it appears in the Document Intelligence workspace.
Target table	The table that stores the document processing results for these table fields.
Parent mapping to field	Field on the target table that you want to align this table with. **Note:** You must first select a target table.
This table is required for extraction	Option to make the table fields required. Required table fields can't be left empty or unreviewed. They also can't contain ambiguous values. An ambiguous value is a field entry that can be interpreted in more than one way. Required fields affect how document tasks are processed in the Fully automated extraction mode. For more information, see Configure data extraction modes.
Columns	

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table form (continued)

Field	Description
Column title	Name of the column header in the table.
Type	The type of the field in the table column. For example, a text or date field. Some field types convert the extracted value into a standard format. For more information, see Data normalization.
Select target field	Field on the target table that you want to align this field with. This field is used for integration with other applications. For more information, see Integrate with a custom application or workflow.
When the date is ambiguous in a document, DocIntel will interpret it in the following order	The default interpretation of the date format. For example, if you select `Month first` in this field, DocIntel will interpret an ambiguous date like 1/2/2024 as January 2 when it extracts that value from a document. If you select `Day first`, it will interpret it as February 1. This field is available when the **Type** field is set to `Date`.
When the number is ambiguous in a document, DocIntel will interpret it as	The default interpretation of the number format. For example, if you select `1,00` in this field, DocIntel will interpret an ambiguous number like 5 as 5,00 instead of 5.00 when it extracts that value from a document. This field is available when the **Type** field is set to `Integer`, `Decimal`, or `Float`.
Reference table	The table that stores the reference column. It's automatically populated based on the selected target field. This field displays when the **Type** field is set to `Reference field`.
Reference column	The column in the reference table that contains the referenced data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table form (continued)

Field	Description
	DocIntel uses the reference column to find data that matches the extracted field value when processing a document task. The field value is then converted to the format of the reference. For more information on converted values, see Data normalization. This field is available when the **Type** field is set to `Reference field`.
Distinguisher(s)	Additional columns in the reference table that help the user to distinguish between similar records. This field is available when the **Type** field is set to `Reference field`.
This field is required for extraction	Indicates whether a column required. It's automatically selected or cleared based on whether the table is required. Required column fields can't be left empty or unreviewed. They also can't contain ambiguous values. An ambiguous value is a field entry that can be interpreted in more than one way. If it's a reference field type, the required field must have a valid, exact match. By default, DocIntel uses the first matched record. Required fields affect how document tasks are processed in the Fully automated extraction mode. For more information, see Configure data extraction modes.
Select one of these options for cases when: • the date/number is ambiguous in the document • there are multiple reference matches in the document	Option for agent review in situations when DocIntel encounters an ambiguous value in a required field. In such cases, the selected default interpretation will apply to the extracted value. The option is whether to interrupt full automation of document tasks to allow agents to verify the interpreted values. Otherwise, continue automatic processing of document tasks without the agent review.
New column	Option to add a column to the table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table form (continued)

Field	Description
	Use the reorder icon (⠿) to reorder a column in the table.
Create multiple tables	Option to keep the pop-up window displayed on the screen. Enable this option If you're adding more than one table to the use case.

Document Intelligence properties

Document Intelligence (DocIntel) system properties control the behavior of the Document Intelligence application.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

The DocIntel Manager (sn_docintel.manager) can modify the system properties for DocIntel.

Note:

Document Intelligence 3.0 includes new and updated system properties. For more detail about the upgrade, see Upgrade to Document Intelligence 3.0 or later from version 2.4 or earlier.

Document Intelligence 2.4- system properties

Property name	Description	Values
sn_docintel.default_field_sidebar_width	The default sidebar width for the document fields panel in the Document Intelligence workspace.	Allowed: Format of `[integer]px` Default: 416px
sn_docintel.default_image_fit	The default image fit for the document panel viewer in the Document Intelligence workspace.	Allowed: fit_to_page, fit_to_width Default: fit_to_page
sn_docintel.default_thumbnail_sidebar_width	The default thumbnail sidebar width for the navigation panel in the Document Intelligence workspace.	Allowed: Format of `[integer]px`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property name	Description	Values
		Default: 167px
sn_docintel.field_sidebar_layout_position	The layout position for the document fields panel in the Document Intelligence workspace, in relation to the document fields panel sidebar.	Allowed: right, left Default: right
sn_docintel.show_exact_match_option	Show the exact match option for each field in the document fields panel in the Document Intelligence workspace.	Allowed: true, false Default: true
sn_docintel.show_candidate_score	Show confidence scores on the recommendations selection menu in the document fields panel in the Document Intelligence workspace.	Allowed: true, false Default: true
sn_docintel.warning_score_threshold	Threshold used to show the warning icon for low-score recommendations in the document fields panel in the Document Intelligence workspace.	Allowed: Number between 0.0 and 1.0 Default: 0.7
sn_docintel.autofill_threshold	Minimum score threshold required to auto-fill recommendations.	Allowed: Number between 0.0 and 1.0 Default: 0.0
sn_docintel.straight_through_processing_threshold	Minimum score threshold required for straight-through processing of a document task.	Allowed: Number between 0.0 and 1.0 Default: 1.0

Document Intelligence 3.0+ system properties

Property name	Description	Values
glide.platform_ml_di.doc_classifier.days_between_trainings	Document classifier property	Default: 30

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property name	Description	Values
	Minimum number of days between trainings for a given document classifier use case.	
glide.platform_ml_di.doc_data_extractor.warning_score_threshold	Document extraction property Threshold used to show the warning icon for low-score recommendations in the document fields panel in the Document Intelligence workspace.	Allowed: Number between 0.0 and 1.0 Default: 0.7
glide.platform_ml_di.doc_data_extractor.straight_through_processing_threshold	Document extraction property Minimum score threshold required for straight-through processing of a document task.	Allowed: Number between 0.0 and 1.0 Default: 1.0
glide.platform_ml_di.doc_data_extractor.show_exact_match_option	Document extraction property Show the exact match option for each field in the document fields panel in the Document Intelligence workspace.	Allowed: true, false Default: true
glide.platform_ml_di.doc_data_extractor.show_candidate_score	Document extraction property Show confidence scores on the recommendations selection menu in the document fields panel in the Document Intelligence workspace.	Allowed: true, false Default: true
glide.platform_ml_di.doc_data_extractor.field_sidebar_layout_position	Document extraction property The layout position for the document fields panel in the Document Intelligence workspace, in relation to the document fields panel sidebar.	Allowed: right, left Default: right
glide.platform_ml_di.doc_data_extractor.default_thumbnail_sidebar_width	Document extraction property The default thumbnail sidebar width for the navigation panel in the Document Intelligence workspace.	Allowed: Format of [integer]px Default: 167px
glide.platform_ml_di.doc_data_extractor.default_image_fit	Document extraction property The default image fit for the document panel viewer in the Document Intelligence workspace.	Allowed: fit_to_page, fit_to_width Default: fit_to_page

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property name	Description	Values
glide.platform_ml_di.doc_data_extractor.default_field_sidebar_width	Document extraction property The default sidebar width for the document fields panel in the Document Intelligence workspace.	Allowed: Format of [integer]px Default: 416px
glide.platform_ml_di.doc_data_extractor.autofill_threshold	Document extraction property Minimum score threshold required to auto-fill recommendations.	Allowed: Number between 0.0 and 1.0 Default: 0.01
glide.platform_ml_di.doc_data_extractor.draw_tool_enable	This property is used to enable or disable draw tool features for table extraction.	Allowed: true, false Default: true

Document Intelligence roles

Document Intelligence is installed with these roles.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

To learn more about managing per-user subscriptions, see Managing per-user subscriptions in Subscription Management and contact your account representative.

DocIntel Admin [sn_docintel.admin]

Has full access to the Document Intelligence application, except for modifying a subset of system properties, and the billing and internal tables.

Contains Roles

List of roles contained within the role.

- platform_ml_di.admin
- action_designer
- flow_designer
- sn_ace.ace_user
- canvas_user
- usage_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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.

DocIntel Viewer [sn_docintel.viewer]

Has view-only access on Document Intelligence document tasks that they are authorized to view.

Contains Roles

List of roles contained within the role.

- snc_read_only
- platform_ml_di.viewer

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.

DocIntel Extraction Agent [sn_docintel.extraction_agent]

Extracts information from documents using the Document Intelligence workspace.

Contains Roles

List of roles contained within the role: platform_ml_di.extraction_agent.

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.

DocIntel Creation Agent [sn_docintel.creation_agent]

Extracts information from documents using the Document Intelligence workspace. Also enables users to create Document Intelligence document tasks and submit them for processing.

Contains Roles

List of roles contained within the role: platform_ml_di.creation_agent.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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.

DocIntel Manager [sn_docintel.manager]

Creates and edits use cases, fields, field groups, and document tasks. Views, measures, and analyzes the usage and effectiveness of Document Intelligence using the Platform Document Intelligence Usage dashboard. Grants access to submit document tasks and interact with the Document Intelligence workspace.

Contains Roles

List of roles contained within the role.

- platform_ml_di.manager
- action_designer
- flow_designer
- sn_ace.ace_user
- canvas_user
- usage_admin

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.

Document Intelligence terminology

Before getting started with Document Intelligence (DocIntel), it's important to understand some key concepts used in the application.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Updated terminology

As of Document Intelligence v2.3, DocIntel features include updated terminology.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Terminology updates

New term	Old term
Document task	Task, extraction task
Field	Key, attribute
Field group	Key group
Field value	Extracted value
Integration	Integration setup
Recommendation	Candidate, suggestion
Fully automated	Straight through processing
Use case	Task definition

classification

In Document Intelligence, the process of categorizing documents and document pages based on their type.

confidence score

A numerical value assigned to a recommendation by Document Intelligence indicating its certainty about the extracted information. The higher the score, the more reliable the recommendation.

document class

A field used to apply a category or label to a document and to pages within a document in Document Intelligence.

For example, for an identity document use case, the classes might be passport, driver's license, birth certificate, and the like.

document task

A document processing activity in Document Intelligence. It includes the information that you want to extract from the document or documents.

extraction

In Document Intelligence, the process of identifying relevant information in a document and using it as a basis for the AI to recommend a field value.

field

A single piece of information to extract from a document in Document Intelligence. For example, the date on a document.

Fields are sometimes called keys or attributes.

field group

A group of fields that belong together in Document Intelligence. Field groups are used to extract information from lists and tables.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, a group named "item" contains description, quantity, and unit price fields from a purchase order.

field value

The final output of the Document Intelligence application. The output contains the values for the specified fields that were extracted for a given document task.

recommendation

A bit of text found on a document. The recommendation includes details about its location in the document, such as the page of a document and the specific location on that page.

The recommendations that the AI provides are sorted based on how likely the AI believes a given recommendation to be the correct value for the current field.

unstructured document

A document that mainly contains free-form textual information and does not conform to a specific format or structure.

use case

A use case is a set of configurations used to define the structure of a type of document that you want to process. It's made up of the use case record and its related fields, field groups, integrations, flows, and all the related machine learning (ML) models. The use case also includes the mode for how the extraction should occur.

Document task statuses

The following is a list of the statuses for DocIntel document tasks. These statuses apply to tasks for both document classification and data extraction.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Statuses for document tasks

Status	Description
Setup	The task has been created and is getting prepared for processing.
New	The task has been initiated by an admin or by a user who has selected **Process Task**. The documents for the task are being processed and the fields aren't yet populated.
In Progress	The documents for the task are processed. The fields are populated through auto-fill or by a user, who reviews and updates the predictions generated by the AI as needed.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Statuses for document tasks (continued)

Status	Description
Done	The task is completed after having been automatically processed in Fully Automated mode or after a user has reviewed the fields and selected **Submit**.
Failed	The task encountered an error during its processing and failed to process.

Domain separation and Document Intelligence

This domain separation overview relates to Document Intelligence. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can then control several aspects of this separation, including which users can see and access data.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

Support level

Support: Basic.

How domain separation works with Document Intelligence

Follow these steps to achieve domain separation:

- Create a user with the required sn_docintel.admin roles in the respective domain.
- Replicate the following for every domain:
  - Document tasks
  - Fields

Languages supported by Document Intelligence

The Document Intelligence application provides support for documents in different languages.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

The language model configured for a use case allows Document Intelligence to detect data in documents, make predictions for field values, and further train the model in the selected languages.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Select the language model when creating a document extraction use case. For more information, see Create a document extraction use case.

The following table lists the languages supported for Document Intelligence along with the models used to support them.

Document Intelligence languages and models

Model	Languages
Default	English, French, German, Spanish, Portuguese, Dutch, Italian, Czech, Danish, Finnish, Norwegian, and Swedish.
OCR for Japanese and Chinese	Chinese (simplified) and Japanese

Limitations in Document Intelligence

There are several important limitations to be aware of when you're using Document Intelligence.

Important: Starting with the Zurich release, Document Intelligence is being prepared for future deprecation. It will be hidden and no longer activated on new instances but will continue to be supported. For details, see the Deprecation Process article [KB0867184] in the Now Support Knowledge Base. Instead, you can extract information from documents using the Now Assist in Document Intelligence application. For more information, see Now Assist in Document Intelligence.

For a list of the important limitations in Now Assist in Document Intelligence that differ from the limitations in Document Intelligence, see Limitations in Now Assist in Document Intelligence.

The following table is a list of the important limitations in Document Intelligence.

Limits in Document Intelligence

Limit	Description
File formats	The supported file formats are JPEG, PNG, and PDF.
File size limits	The file size limit is 10 MB.
Page count limit per document task	Document Intelligence supports a 10-page count limit for JPEG and PNG. For PDFs, the page count limit is 25.
Supported languages	For information on supported languages, see Languages supported by Document Intelligence.
Document rotation	Document Intelligence supports rotating in 90-degree increments.
Text alignment	The text must be aligned horizontally within the document.
Minimum character size	The minimum character size is 15 pixels.
Character type	Document Intelligence supports only printed character types in a document.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Limits in Document Intelligence (continued)

Limit	Description
Character set	Document Intelligence detects the following characters: • a-z, A-Z, 0–9 • á à â ä ã å ß ç é è ê ë î ï í ñ ó ò ô ö õ ú ù ü û œ • Á À Â Ä Ã Å ß Ç É È Ê Ë Î Ï Í Ñ Ó Ò Ô Ö Õ Ú Ù Ü Û Œ • , : ; . ' \ " ! ? ¿ ¡ + - * ( ) [ ] } { & % @ # / | ~ ^ < > ` ± = _ $ £ ¢ €
Sync/async operations	Async
Maximum number of document tasks processed per instance per day	The maximum number of document tasks processed per instance per day is 2000.
Maximum number of fields per use case	The maximum number of fields per document extraction use case is 50. The maximum number of fields (categories) per document classification use case is 30.

Task Intelligence

Task Intelligence uses machine learning to train solutions with your data and achieve important outcomes.

Overview of Task Intelligence

Task Intelligence enables you to set up machine learning models to automate task creation, triage, and investigation. You can then track and assess the performance of your solutions to optimize their potential to help you work more efficiently and lowering the mean time to resolve (MTTR).

Get started

Explore	Install	Configure
		
Explore Task Intelligence and machine learning models.	Install Task Intelligence and related applications.	Configure machine learning models.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create models	Analyze models	References
		
Create and edit Task Intelligence models.	Monitor and analyze model performance.	References for components installed with Task Intelligence and language support.

Troubleshoot and get help

- Search the Now Community portal for AI & Intelligence
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Task Intelligence

Learn more about Task Intelligence and how machine learning models can learn from your data to make predictions and achieve important outcomes.

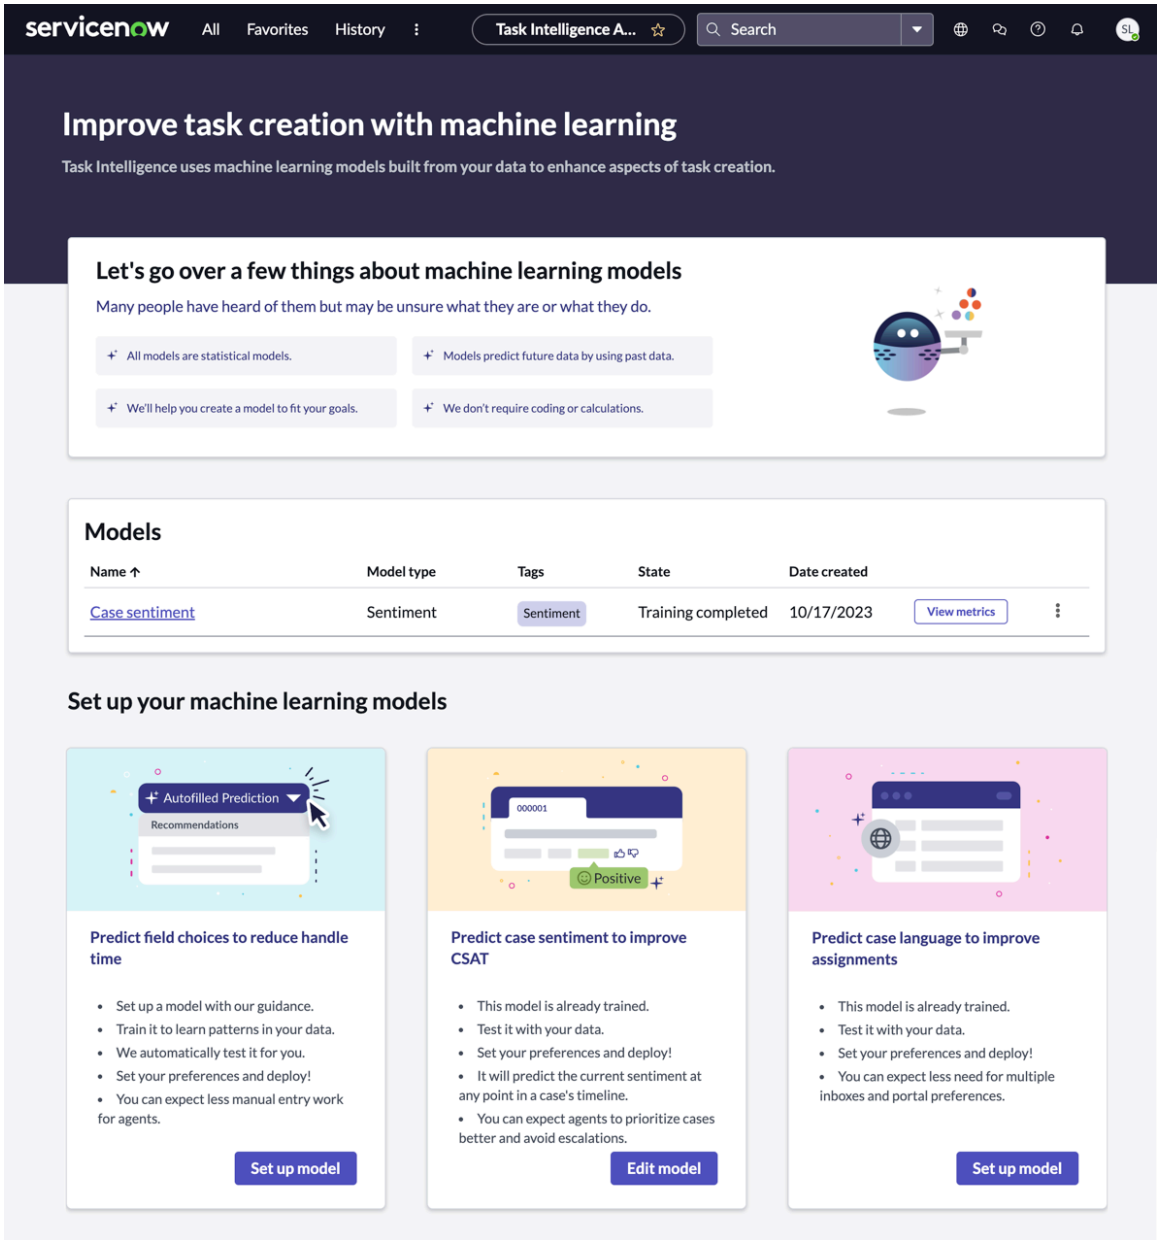
Task Intelligence overview

With Task Intelligence, you can easily set up machine learning solutions that interact with your data. Then, track how the solutions are impacting creation, deflection, triage, remediation, and optimization moments. With automated task creation, triage, and investigation, agents spend less time receiving and preparing information, so tasks can proceed faster. These models can help you work more efficiently, helping to lower the mean time to resolve (MTTR) of tasks.

With Task Intelligence, agents:

- Solve issues faster for better service and experiences
- Focus on meaningful, high-value work
- Reduce error rates and drive down costs
- Improve time to value with intuitive experiences

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



The Admin Console supports the following model types:

Model	Application	Description
Incident field prediction	Task Intelligence for ITSM	Create an incident field prediction model
Similarity model management	Task Intelligence for ITSM	Create a similar records prediction model in Task Intelligence for ITSM
Case field prediction	Task Intelligence for CSM	Predicts output fields on case forms, including categorization. Record categorization

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Model	Application	Description
Sentiment prediction	Task Intelligence for CSM	Predicts current and trending user sentiment. **Note:** Model is pre-trained but can be tested and edited. Only one sentiment model can be configured.
Language detection	Task Intelligence for CSM	Detects language which language is being used. **Note:** Model is pre-trained but can be tested and edited. Only one language detection model can be configured.

Task Intelligence workflow

The Task Intelligence Admin Console uses machine learning models. The models are statistical and can predict future data by training with your past data.

With the Admin Console, create, configure, train, test, and deploy predictive models used for automation in other ServiceNow features and applications, such as Customer Service Management (CSM).

Training a machine learning model is when the model learns patterns in past data to make predictions for new data. Models are trained using a lot of data so that they can learn patterns and the large data set makes the learned patterns statistically significant. By answering questions about your information systems, business process, and service operations, the system actively learns from your responses.

Task Intelligence benefits

The Task Intelligence Admin Console provides admins with a no-code experience for deploying Task Intelligence solutions. The seamless experience helps you automate and optimize task creation, deflection, triaging, and resolution.

Features are implemented by two applications: Task Intelligence for Customer Service and Task Intelligence for ITSM.

Benefit	Feature	Users
Predict form field values to auto-populate or recommend as suggestions	Create a field prediction model	Admins, agents
Categorize emails and cases based on language and attachment content	Record categorization	Agents
Predict incident category and priority based on previous incidents to lower time to resolution	Create an incident prediction model	Agents
Analyze initial and ongoing sentiment during customer service cases	Sentiment Analysis	Agents
Identify the language used to create customer service cases	Language detection	Agents

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Benefit	Feature	Users
Analyze and assess the performance of models	Task Intelligence Analytics and Monitoring	Admins

Configure Task Intelligence

Configure and install the Task Intelligence Admin Console and its related applications.

Configuration overview

The Task Intelligence Admin Console is installed along with Task Intelligence for CSM or Task Intelligence for ITSM. See Install Task Intelligence Admin Console for more details.

Configure Task Intelligence for ITSM

Configure Task Intelligence for ITSM so that agents can use the application to fulfill their requirements efficiently.

Configure Task Intelligence for Customer Service

Install the Task Intelligence for Customer Service application and configure the different features: case categorization, language detection, sentiment analysis, and Document Intelligence for Customer Service.

Install Task Intelligence Admin Console

You can install the Task Intelligence application (sn_ti_admin) 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

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the Task Intelligence Admin Console application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.

Role required: admin

About this task

The following items are installed with Task Intelligence Admin Console:

- Roles
- Tables

For more information, see Task Intelligence roles.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Task Intelligence Admin Console application (sn_ti_admin) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you cannot find the application, you might have to request it from the ServiceNow Store.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In the list next to the **Install** button, the versions that are available to you are displayed.

3. Select a version from the list and select **Install**.

   In the Review Installation Details dialog box, any dependencies installed with your application are listed.

4. If you're 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.
   Demo data are the sample records that describe application features for common use cases. Load the demo data when you first install the application on a development or test instance.
6. Select **Install**.

Manage machine learning models with Task Intelligence

Create, edit, train, and assess machine learning models to make predictions for customer service cases and incidents.

Managing Task Intelligence for ITSM models

Use the machine learning capabilities of Task Intelligence for ITSM to predict and recommend incident field information and help with incident categorization.

Use Task Intelligence for Customer Service

Use the Task Intelligence features to create field prediction and sentiment models for cases, and review Task Intelligence analytics and prediction history.

Create a Task Intelligence model

Create machine learning models to predict field values, analyze case sentiment, or detect case language.

Create a model to predict incident fields

Create and deploy incident-prediction, solution-based models to predict incident fields for new IT service incidents.

Create a model to predict record fields

Create and train a model to predict fields for customer service cases.

Edit a Task Intelligence model

Retrain your machine learning models or edit what they predict to better align with your business goals.

Edit a model to predict incident fields

Edit the incident prediction model that has already been trained and deployed. Change the model configurations, view the updated training results, and redeploy the model.

Create and edit a case field prediction model

Create, train, and edit a model to predict fields for customer service cases.

Create a model to predict case sentiment

Edit and test the pre-trained sentiment model to predict sentiment for customer service cases.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a model to detect case language

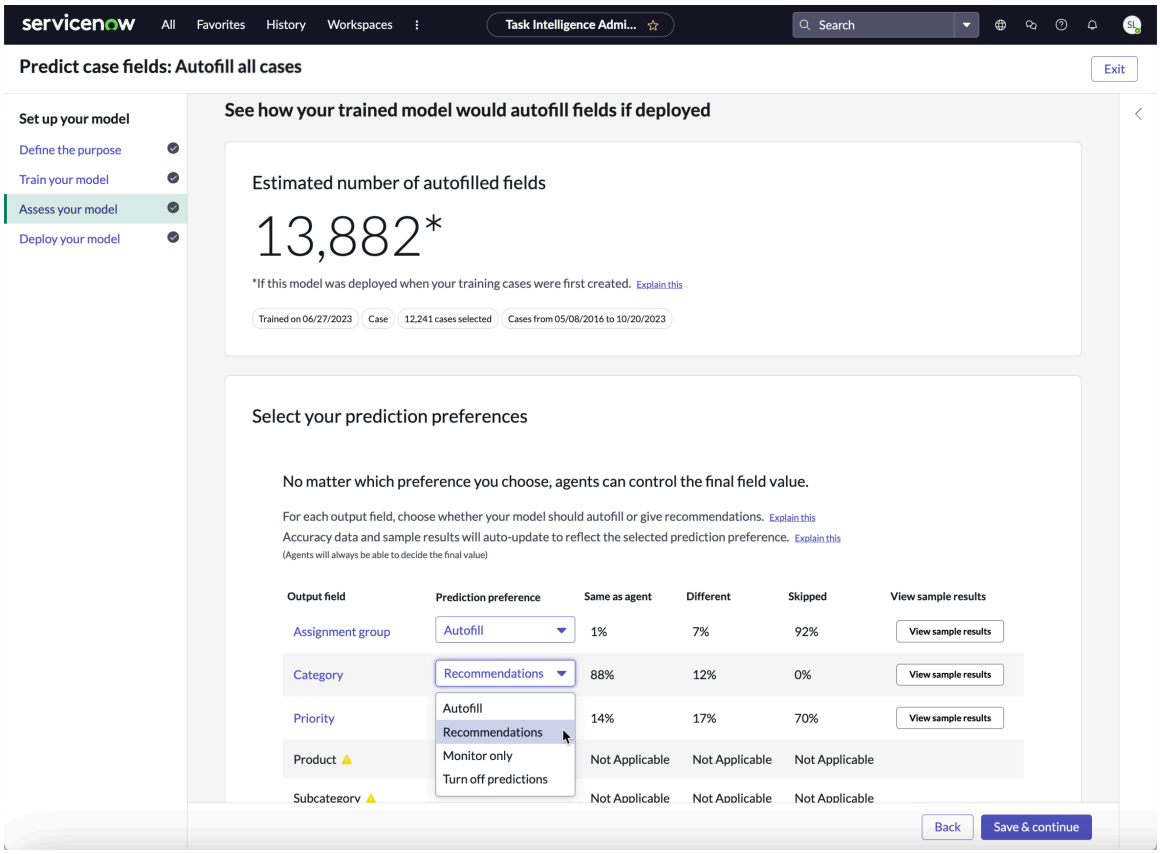
Edit and test the pre-trained model to detect the language used to create customer service cases.

Assess a Task Intelligence model

Assessing a machine learning model's performance helps determine how to use and train the model to achieve your desired outcomes.

The Assess Your Model screen allows you to evaluate the model's performance. After training or retraining your model, the Assess Your Model screen will show an estimate of the model's average performance against your most recent data.

Note: It is normal to see some variability in model performance from day to day. Performance tends to average out to the estimated performance over time.



The Assess Your Model screen also allows you to view example predictions on a sample of records. These examples demonstrate the predictions but do not necessarily reflect the quality or average performance of the model. The estimates provided on the Assess Your Model screen as well as the reports on the **Monitoring** page are calculated from a much larger number of cases.

You can also use the Assess Your Model screen to choose one of the following preferences for each field:

- **Autofill** the predicted value in the field.
- **Recommendations** are provided for the predicted value in the field.
- **Monitor only** and run the prediction model for the field only in the background.
- **Turn off predictions** for the field.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Monitoring mode

Monitoring mode allows you to monitor the performance of a model at the field level without the predictions being applied to records. The model runs in the background only and can be trained and retrained until you are satisfied with its performance. You can set the model fields to Monitoring mode from the Assess Your Model screen when editing your model.

You can view and track model performance with the Task Intelligence Admin Console. See Task Intelligence Analytics and Monitoring.

Assess field-level accuracy for multi-output models

Evaluate the performance of your model by field.

Before you begin

Role required: admin

About this task

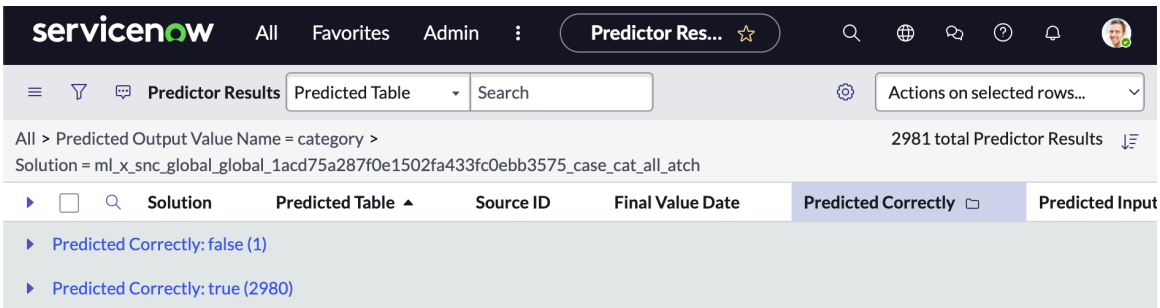
After you have created and tested your model, assess its performance to evaluate how well it's predicting what you want it to predict. For multi-output models that predict multiple fields, you may want to assess the performance of each field individually.

If you have not already created a model, see either Create a model to predict record fields or Create a model to predict incident fields to make one.

Your model must be in Monitoring mode. You can set Monitoring mode on the Assess Your Model screen. For more details, see the Set your preferences step in Create a case field prediction model or Create a model to predict incident fields model.

Procedure

1. Navigate to the **Predictor Results** (ml_predictor_results_task) by entering `ml_predictor_results_task.list` in the Application navigator.
2. Filter the **Predicted Output Value Name** for the name of the field you want to assess, such as "product" or "category."
3. Group the list by **Predicted Correctly** by selecting the list controls icon in the top left corner of the screen.

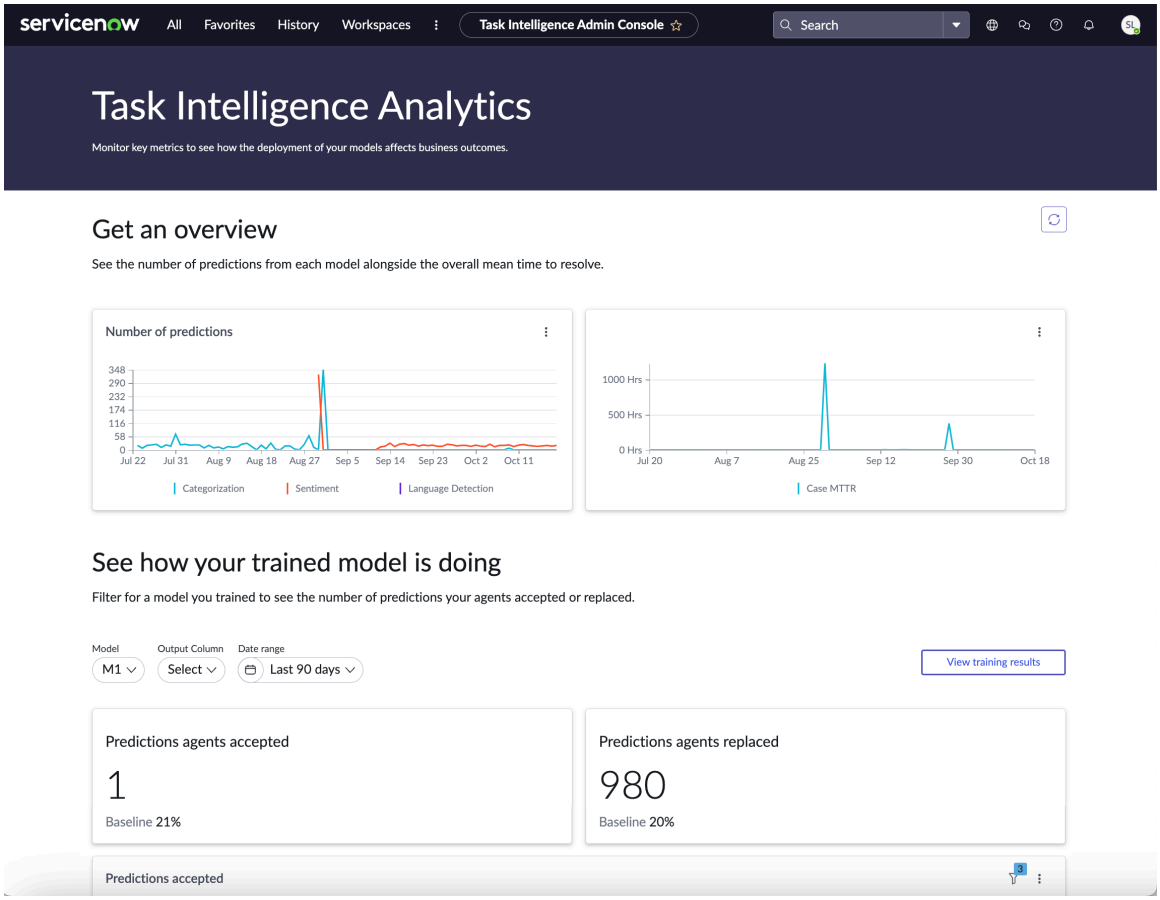


4. Divide the number of records where **Predicted Correctly** is **true** by the total number of records. This represents your accuracy for the specific field.
5. Follow steps 2-4 for each field in which you would like to calculate the accuracy.
6. **Optional:** View the performance of your model on the Monitoring screen.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Navigate to **Task Intelligence for Customer Service** > **Monitoring** in the Application navigator.

(Optional) The Monitoring screen allows you to select the model, field, and date range to display in the charts. Accuracy is measured based on whether the top autofilled value or the top 3 recommendations are correct.



b. Select a model.

c. Select an output column for the field.

d. Select a date range.

What to do next

If the accuracy of each field is acceptable, transition your model from monitoring mode to real-time predictions and deploy. If the accuracy of a specific field is not acceptable, you can remove that output field from your model, retrain, and deploy. See Edit a Task Intelligence model for more information on editing your model.

Task Intelligence Analytics and Monitoring

Use the Task Intelligence Admin Console to monitor, analyze, and assess the performance of your machine learning model.

Task Intelligence for ITSM Analytics and Monitoring

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View the impacts of your Task Intelligence models. Monitor model performance overtime, track business value, and view what predictions your agents did and didn't use.

View Analytics for Task Intelligence for CSM

View Task Intelligence for Customer Service analytics and see how your machine learning models are working.

View Task Intelligence prediction history

View details about historical predictions in the Predictor Results for CSM models.

Export a Task Intelligence model

Export a Task Intelligence model to another instance.

Before you begin
Role required: admin or tia_admin

Procedure

1. Navigate to **All** > **Task Intelligence for CSM** > **Setup** or **All** > **Task Intelligence for ITSM** > **Setup**, depending on which application and context you want to export for.
2. In the **Models** list, click the more options menu of the model you want to export.
3. Click **Export model**.
   A new tab opens with authentication steps.
4. Click **Allow**.

Result
The XML file downloads to your browser.

What to do next
You can use the XML file to easily transfer the model to another instance using update sets. To learn more, see System update sets.

Task Intelligence references

Reference topics provide additional information about components installed and language detection details.

Task Intelligence roles

Reference for user roles that are installed with activation of the sn_ti_admin plugin.

Reference for Task Intelligence for ITSM

Reference topics provide additional information about components installed with Task Intelligence for ITSM.

View Task Intelligence Analytics

View Task Intelligence for Customer Service analytics and see how your machine learning models are working.

Languages supported by Task Intelligence

The language detection feature for Task Intelligence for CSM supports 22 different languages.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Task Intelligence roles

Task Intelligence is installed with these roles.

To learn more about managing per-user subscriptions, see Managing per-user subscriptions in Subscription Management and contact your account representative.

Task Intelligence Admin [tia_admin]

Admin user with rights to all actions in the application. Creates, edits, and monitors all Task Intelligence models.

Contains Roles

List of roles contained within the role.

- ml_admin
- platform_ml_read

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.

Task Intelligence User [tia_user]

Views and monitors Task Intelligence models. Requires additional accesses to see training.

Contains Roles

List of roles contained within the role: ml_report_user.

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.

Task Intelligence Analyst [tia_analyst]

Monitors Task Intelligence models through the Analytics dashboard.

Contains Roles

List of roles contained within the role: ml_admin.

Groups

List of groups this role is assigned to by default.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

None.

Special considerations

Important: Avoid granting an admin role when more specialized roles are available.

Languages supported by Task Intelligence

The Task Intelligence application provides support for predictions in different languages.

Each of the supported languages has an associated two-character code. For example, `English - en`.

The following table lists the languages supported for Task Intelligence along with the models they support.

Note: Only languages included in the training dataset are supported.

Language name	Language code	Incident fields (ITSM)	Case fields (CSM)	Sentiment prediction (CSM)	Language detection (CSM)
Arabic	ar				X
Brazilian Portuguese	pb	X			
Chinese (simplified)	zh	X			X
Czech	cs				X
Danish	da	X			
Dutch	nl	X			X
English	en	X	X	X	X
Estonian	et				X
Finnish	fi	X			X
French	fr	X	X		X
French Canadian	fq	X			
German	de	X	X		X
Hebrew	he				X
Hungarian	hu				X
Italian	it	X			X
Japanese	ja	X			X
Korean	ko	X			X
Norwegian	nb	X			X
Polish	pl	X			X
Portuguese	pt	X			X

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Language name	Language code	Incident fields (ITSM)	Case fields (CSM)	Sentiment prediction (CSM)	Language detection (CSM)
Russian	ru				X
Spanish	es	X	X		X
Swedish	sv	X			X
Thai	th				X
Turkish	tr				X

ServiceNow AI Lens

With ServiceNow AI Lens, which is a ServiceNow® Now Assist application, you can use generative AI to scan, extract, comprehend, and synthesize data to optimize your workflows.

Get started

Explore



Learn about ServiceNow AI Lens concepts and features.

Configure



Configure the ServiceNow AI Lens application to get started.

Use



Capture screens to analyze, comprehend, and use data.

Reference



Get details about ServiceNow AI Lens components, such as roles, tables, and limitations.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important:

- 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.
- Don't scan any personally identifiable information, such as medical reports, financial reports, or other sensitive data, when using ServiceNow AI Lens as you don't want to expose the large language model (LLM) to any sensitive information.
- ServiceNow AI Lens currently uses Azure OpenAI GPT-4o LLM.

Troubleshoot and get help

- Workflow Automation product on the ServiceNow Community
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

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. In addition, this application will collect information about scripts (and associated script records) in which Now Assist for code generation is called. Customers can opt out of future data collection at any time, as described in the Now Assist Opt-Out page.

For more information, see the Now Assist documentation.

Exploring ServiceNow AI Lens

ServiceNow AI Lens is a ServiceNow® Now Assist application that uses generative AI to scan images, extract information, understand visual data, and help improve your workflows.

ServiceNow AI Lens overview

The ServiceNow AI Lens application extracts and analyzes data from the screenshots of documents such as a scanned handwritten note or an email that you capture . ServiceNow AI Lens is integrated with the ServiceNow® forms so that you can create or update forms on the ServiceNow instance by auto-filling data.

By connecting to a ServiceNow instance, you can also use ServiceNow AI Lens in a standalone mode to see a preview of the extracted data or get insights about your use cases. For more information, see ServiceNow AI Lens features.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can leverage Lens actions to customize the Lens behavior by providing default context, trigger options, and transform response logic for a table. When Lens is triggered to auto-fill a form on this table, it uses the configuration defined in the Lens action to auto-fill the form with the transformed LLM response. You can also define context, trigger options, and post processing instructions in a Lens action, which is used when Lens is triggered from desktop to show preview of extracted data and post process the LLM response. For more information, see Define ServiceNow AI Lens behavior with Lens actions.

Extract data from documents such as an image, a scanned handwritten note, web page, or a Microsoft Word document, and autofill catalog item forms in the Service Portal. For example, autofill a new vendor onboarding form by extracting key details such as vendor name, address, contact email, and banking information from multiple artifacts, including Excel files, emails, images, and PDF documents. For more information, see Autofill catalog item forms on Service Portal using ServiceNow AI Lens.

Trigger ServiceNow AI Lens from the Now Mobile® application to extract data from documents such as an image, a scanned handwritten note, web page, or a Microsoft Word document, and auto-fill fields in a form. For more information, see Trigger ServiceNow AI Lens from Now Mobile® and autofill forms.

ServiceNow AI Lens is a desktop application that you can download for both Windows and macOS operating systems.

Benefits of ServiceNow AI Lens

ServiceNow AI Lens can do the following tasks automatically for you:

- Fill in a form without having to do it manually. ServiceNow AI Lens can scan and extract the visual data to auto-fill in your forms. For example, it can scan the invoice receipt to auto-fill the fields on the expense submission form.
- Scan and identify the information from handwritten notes or images. For example, ServiceNow AI Lens can scan a sequence diagram to generate a script.
- Search for specific information from images. For example, ServiceNow AI Lens can look for a specific clause in a scanned legal contract.
- Process the screenshots of multiple images at once. For example, ServiceNow AI Lens can extract the data from multiple screenshots of an error to get insights on the data at one time.

ServiceNow AI Lens skill

Now Assist for Platform includes the AI-powered ServiceNow AI Lens skill to read, understand, respond, and act on visual data such as hand-written texts, images, and websites and take powerful actions to boost productivity.

Admins can enable the ServiceNow AI Lens skill in the Now Assist Admin console. Once enabled, users can perform the following actions:

- Create or update ServiceNow records by auto-filling forms using the extracted data.
- Preview the extracted data.
- Provide instructions to extract data in a specific way or to get a specific output.
- Capture multiple screenshots to gather insights from multiple images.

To view the procedure to enable the ServiceNow AI Lens skill, see Activate the ServiceNow AI Lens skill.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The ServiceNow AI Lens application requires a Workflow Data Fabric (previously known as Automation Engine) license and a Now Assist license.

Available LLMs

You can use Azure OpenAI, Google Gemini, or Anthropic Claude on AWS as the AI model provider for ServiceNow AI Lens skill. You can set skill-level preferences in the Now Assist Admin console. For more information, see Large language models on the ServiceNow AI Platform.

What to explore next

To learn how to configure and use ServiceNow AI Lens, see:

- Configure ServiceNow AI Lens
- Using ServiceNow AI Lens
- ServiceNow AI Lens reference

ServiceNow AI Lens features

Learn about the various ServiceNow AI Lens features to help you get started with analyzing and gathering insights from the visual data.

ServiceNow AI Lens offers the following features:

Capture data

Scan and get actionable insights from visual data such as images, scanned handwritten notes, Excel sheets, web pages, or any document that gives visual data.

For example, ServiceNow AI Lens can scan an email and auto-fill the fields on the Incident form.

Provide instructions to analyze data in a specific way

Provide instructions to analyze the data in the screenshots or files in a way that you want after selecting the Edit ( ) icon and enter additional instructions. For example, you can capture screenshots or upload copies of all the expense bills of all your business trips in the last six months and provide instructions to arrange the expense bill amounts date-wise and expense head-wise.

Capture multiple screenshots

Capture multiple screenshots to analyze and gather data from various images.

For example, ServiceNow AI Lens can scan multiple bills so that you can submit a consolidated invoice report.

Launch from ServiceNow instance

Launch ServiceNow AI Lens from a ServiceNow instance to fill the extracted data automatically into the form fields while you're creating or updating records.

Use Lens actions to customize Lens behavior

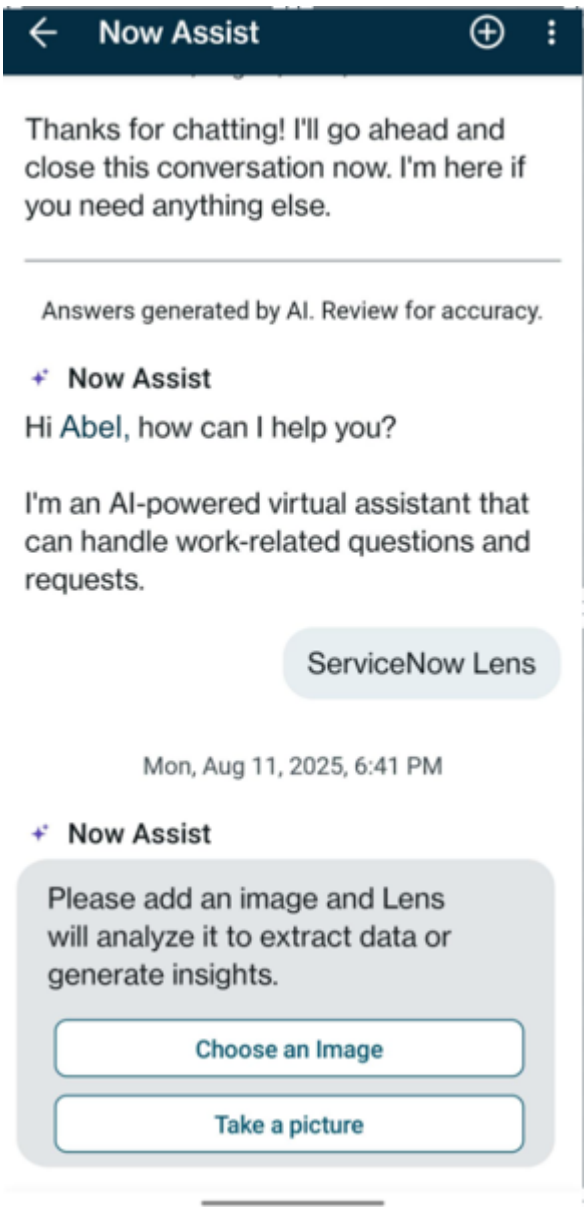
You can use Lens actions that define Lens behavior depending on how ServiceNow AI Lens is triggered and what context is set. You can define default instructions, trigger options, custom context, transform response logic, and post processing instructions for the Lens execution. With Lens actions, you can customize how a classic form auto-filled using ServiceNow AI Lens or what data must be extracted for preview and post process the LLM response. For more information, see Customize ServiceNow AI Lens behavior by creating a Lens action.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, you can define a Lens action that is used when Lens is triggered from an instance to populate a form of a particular table. You can also define form fields that must be populated.

ServiceNow AI Lens in Virtual Agent

Trigger ServiceNow AI Lens from a Virtual Agent conversation by using ServiceNow AI Lens topic in Virtual Agent. For more information, see Enabling ServiceNow AI Lens in Virtual Agent and Extract and analyze data using ServiceNow AI Lens from Virtual Agent.



Use as a back-end service

As part of your integration logic, configure a Lens action as one of the steps to invoke a ServiceNow AI Lens service from any part of the ServiceNow AI Platform, such as a workspace form or a portal.

For your own use cases, you can leverage Lens actions for using Lens as a service in the following ways:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Client service: Launch the Lens client to scan and analyze images. For example, you can configure a UI action on a workspace form or a service portal to launch the Lens client. For more information, see KB2492775.
- Server service: Use Lens as a service to analyze the provided images without launching the Lens client. For example, you can call Lens service from a Script Include or flow action to get the LLM response after analyzing the visual data. For more information, see Script include - AILensActionService.

Attach images to a record

View the captured images attached to the record that is auto-filled using ServiceNow AI Lens. You can later refer the attached images to understand the source of the record data. For more information, see Customize ServiceNow AI Lens behavior by creating a Lens action.

Use as a standalone application

Use ServiceNow AI Lens as a standalone application for previewing extracted data or getting insights. You must connect to a ServiceNow instance before you start using ServiceNow AI Lens. For more information, see Extract and analyze data with ServiceNow AI Lens desktop app.

Autofill form fields in the Now Mobile® application

Trigger ServiceNow AI Lens from Now Mobile® application to extract data from artifacts and autofill fields in a form on the Now Mobile® application.. For more information on creating or updating records in the Now Mobile® application, see ServiceNow AI Lens for mobile. You can add a ServiceNow AI Lens button to input form screens to allow users to quickly extract data from attachments and complete forms faster. For more information, see Configure Lens launcher button.

Autofill Catalog Item forms in Service Portal

Autofill Catalog Item form fields by triggering ServiceNow AI Lens from Service Portal. ServiceNow AI Lens extracts data from one or more artifacts and auto fills the relevant fields in the form.. For more information on using ServiceNow AI Lens in the Service Portal, see Autofill catalog item form in the Service Portal.

ServiceNow AI Lens desktop app landing page overview

Learn about the ServiceNow AI Lens desktop application landing page.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ServiceNow AI Lens landing page

No.	Feature	Description
1	Side panel	Side panel to view user profile and instance details, getting help options, and option to log out from the current session.
2	Action	Preconfigured Lens actions that define behavior of ServiceNow AI Lens. Only the Lens actions that are configured to be triggered from the desktop and are assigned to you are visible on the landing page. If there are no Lens actions assigned to you, you will not see any Lens action on the landing page.
4	Action card options	• Pin and unpin - Pin option to make the Lens actions favorite on the landing page. The last pinned Lens action appears first in the list. • Details - Shows Lens action details, such as Name, Description, Default additional directions, and Table.

Supporting information for ServiceNow AI Lens

Get a quick overview of the important information that is related to the ServiceNow AI Lens application.

Supported versions

ServiceNow AI Lens is supported from Yokohama Patch 6 and Zurich Patch 1.

Supported user interfaces

To scan artifacts such as images, scanned handwritten notes, Excel sheets, web pages using Now Assist, enable the ServiceNow AI Lens application and then activate the ServiceNow AI Lens skill. For more information, see Configure ServiceNow AI Lens.

The ServiceNow AI Lens application includes the skills that are listed in the following table.

ServiceNow AI Lens skill supported interfaces

Interface	Skill
• Core UI • Desktop	ServiceNow AI Lens skill

Licensing requirements

The ServiceNow AI Lens application requires a ServiceNow Pro Plus or Enterprise Plus license.

Application information

Enable the ServiceNow AI Lens (com.sn_ai_lens) app plugin to use the ServiceNow AI Lens skill.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This store app has the following dependencies:

- ServiceNow AI Lens Core (sn_app_lens_core)
- Now Assist for Platform (sn_genai_platform)

ServiceNow AI Lens skill

This skill provides generative AI capabilities to read, understand, respond, and act on visual data such as hand-written texts, images, and websites and take powerful actions to boost productivity.

This skill is a Platform skill that is enabled by the admin from the ServiceNow AI Lens skill card and enables users to perform the following actions:

- Create or update ServiceNow records by auto-filling forms using the extracted data.
- Preview the extracted data.
- Provide instructions to extract data in a specific way or to get a specific output.
- Gather insights from multiple images.

Availability

This skill is available in the workflow and product listed below.

Now Assist products and workflows

Workflow	Product
Platform	ServiceNow AI Lens

Licensing requirements

The ServiceNow AI Lens application requires a ServiceNow Pro Plus or Enterprise Plus license.

Configure ServiceNow AI Lens

If you have the admin role, you can enable the ServiceNow AI Lens application (sn_ai_lens) to extract and comprehend data for various workflows.

Before you begin
To get started with ServiceNow AI Lens, you must have:

- A ServiceNow Pro Plus or Enterprise Plus license.
- An instance on Yokohama Patch 6+ or Zurich Patch 1+.

Role required: admin

About this task
ServiceNow AI Lens isn't a standalone application that you can install directly. To enable ServiceNow AI Lens on your instance, you must install other Now Assist applications, such as Now Assist for IT Service Management (ITSM) or Now Assist for Customer Service Management (CSM).

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Review the ServiceNow AI Lens application listing in ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.
- Refer to the Now Assist Overview to see the prebuilt skills that exist and are currently available for purchase in ServiceNow Now Assist packages and to review the number of assists consumed.
- Perform these steps in your ServiceNow instance.

By enabling the ServiceNow AI Lens application (sn_ai_lens), you can view the ServiceNow AI Lens and download modules in your ServiceNow instance.

The following items are installed with ServiceNow AI Lens:

- Roles
- Tables

For more information, see Components installed with ServiceNow AI Lens.

Procedure

1. Navigate to **All** > **System Definition** > **Plugins**.
2. In the search field, under the heading **Search your licensed applications and plugins**, enter the name of the Now Assist application that you want to install.
   For example, you can enter Now Assist for IT Service Management (ITSM).
3. Press **Enter**.
   The Now Assist application listing appears.
4. Select the Now Assist application card.
5. Select **Install**.

Result
The ServiceNow AI Lens and ServiceNow AI Lens Core applications are installed.

What to do next
Turn on the ServiceNow AI Lens skill to extract, comprehend, and synthesize visual data to optimize your workflows. For more information, see Activate the ServiceNow AI Lens skill.

Activate the ServiceNow AI Lens skill

Activate the ServiceNow AI Lens skill to start using generative AI to scan and analyze visual data such as scan artifacts such as images, scanned handwritten notes, Excel sheets, web pages and optimize your workflows.

Before you begin

- Confirm that the ServiceNow AI Lens application is installed. For more information, see Configure ServiceNow AI Lens.
- Perform this task in your ServiceNow instance.
- Role required: admin

Procedure

1. Navigate to **All** > *Now Assist Admin* > **Now Assist Skills**.
2. In the workflow list, select **Platform**.
3. In the Search using keywords field, enter ServiceNow AI Lens, and press **Enter**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. On the tile for the ServiceNow AI Lens skill, select **Turn on**.
5. On the Successfully activated dialog box, select **Back to skills**.
6. On the ServiceNow AI Lens card, verify that the ServiceNow AI Lens skill is activated.

Download the ServiceNow AI Lens installer

Download the ServiceNow AI Lens installer so that you can install the ServiceNow AI Lens on your system for scanning visual data.

Before you begin

Role required: admin

Procedure

1. Navigate to **All** > **ServiceNow AI Lens** > **Downloads**.
2. In the ServiceNow AI Lens Downloads page, do any of the following actions to download the required application for macOS or Windows:
   - Select the download icon ( ).
   - Select the copy link ( ). In a browser, right-click and select the **paste and go** option.

   PORT TO AUSTRALIA

   ServiceNow AI Lens Downloads

   A dialog box might prompt you to save or open the file.

   Note: Depending on your browser setting, the browser might automatically save the file to your Downloads folder.
3. Execute the application installation file (.msi file) or (.pkg file).
4. Follow the on-screen instructions to install the application.
5. From the Privacy & Security settings on your macOS system, grant ServiceNow AI Lens permission to record your screen.

What to do next

Use ServiceNow AI Lens to capture screens, extract and comprehend data, and take further actions. For more information, see Using ServiceNow AI Lens.

Define default instructions for ServiceNow AI Lens

Create a system property to define default instructions for ServiceNow AI Lens execution on a specific form.

Before you begin

Role required: admin

About this task

You can define default instructions for specific forms to create or update records using ServiceNow AI Lens in a ServiceNow instance. This helps streamline the request process by eliminating the need for users to manually enter the same instructions each time. The defined instructions are automatically included in all ServiceNow AI Lens requests for that form or table.

When users provide additional instructions through the ServiceNow AI Lens UI, those instructions are appended to the default instructions. If there is a conflict between user input,

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

application-generated instructions, and system-defined defaults, ServiceNow AI Lens instructs the LLM to prioritize the user-provided input.

Procedure

1. In the filter navigator, enter *sys_properties.list* and press **Enter**.
2. Select **New**.
3. On the System Property form, fill in the details.

Default instructions system property

Field	Description
Name	sn_lens.prompt.<tablename> Replace <tablename> with the internal name of the table. For example, sn_lens.prompt.incident.
Value	Enter the default instructions. For example - Suggest a suitable start date for this incident based on the incidence type and urgency.

4. Select **Submit**.

Define ServiceNow AI Lens behavior with Lens actions

As a Lens admin, you can create Lens actions in ServiceNow AI Lens to customize Lens behavior by providing default instructions, configuring context, and more.

Use Lens actions to define default instructions, trigger options, custom context, and post processing instructions for the Lens response. With Lens actions, you can achieve the following goals:

- Trigger ServiceNow AI Lens from a desktop, ServiceNow instance, or Virtual Agent (from a mobile device or a portal)
- Auto-fill a form on the ServiceNow instance
- View the gathered insights in an editable preview
- Trigger flows, subflows, or actions after Lens execution
- Invoke AI agents
- Run ServiceNow AI Lens as a back-end service

Whenever ServiceNow AI Lens is launched from an instance, it checks if any active Lens action is available for the table. If available, it uses the logic defined in Lens action during execution. The Lens action can only be used by the users or groups that are assigned to the Lens Actions record.

Refer to the following examples to understand the different use cases of setting up the Lens actions.

- Example: Scan documents to auto-fill the Insurance Request form
- Example: Extract structured data from vendor agreements
- Example: Auto-fill user records on a user table

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Customize ServiceNow AI Lens behavior by creating a Lens action

Create a Lens action to define custom context, trigger options, and default instructions for the ServiceNow AI Lens execution.

Before you begin

Role required: lens_admin

About this task

Learn about the various Lens trigger options for ServiceNow AI Lens.

Lens trigger options

Goal	Trigger From	Trigger For
Auto-fill a form with a custom context and default instructions when triggered from the instance. When ServiceNow AI Lens is triggered from the instance in the context of the Lens action, the default instructions are always applied and specific form fields are extracted. To view the procedure to define default instructions, see Define default instructions for ServiceNow AI Lens.	Instance	Form
Auto-fill a form with a custom context and default instructions when triggered from the ServiceNow AI Lens desktop app. You can trigger ServiceNow AI Lens from the desktop app directly in context of the Lens action instead of launching Lens from the instance. The default instructions are always applied and specific form fields are extracted.	Desktop	Form
With the help of a client script or server script, run ServiceNow AI Lens as a back-end service.	Instance	Service
View the gathered insights in the preview window of ServiceNow AI Lens by providing custom context and perform post processing steps.	Desktop	Others

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When ServiceNow AI Lens is triggered from the instance or desktop to auto-fill a form, it uses an active Lens action associated with the target table.

For example, if a user selects **Create with Lens** from Incident table list, ServiceNow AI Lens checks for an active Lens action record for the Incident table where **Trigger From** is set to Instance and **Trigger For** is set to Form. If a matching action record exists, Lens executes in the context of that action.

Procedure

1. Navigate to **All** > **ServiceNow AI Lens** > **Lens Actions**.
2. Select **New**.
3. Enter a name and description for the Lens action in the Name and Description fields respectively.
4. In the **Trigger From** field, select one of the following options.

Option	Description
Desktop	Option to specify that this action must be used when ServiceNow AI Lens is triggered from the ServiceNow AI Lens desktop app.
Instance	Option to specify that this action must be used when ServiceNow AI Lens is triggered from a ServiceNow instance.

5. To save the changes, right-click the form header and select **Save**.
6. Set up the Lens action by filling the fields in the Lens Actions form, under the Customize heading.

Lens Actions form

Field	Description
Trigger For	○ **Form** - Lens action is used when ServiceNow AI Lens is triggered to auto-fill a form. Apply a transformation logic to the Lens response before the form is auto-filled by using the #unique_867_Connect_42_transform-response option. This option is available when you select **Desktop** or **Instance** from the Trigger From list. ○ **Service** - Lens action is used when ServiceNow AI Lens is used as a service. Select this option to trigger Lens from an instance to auto-populate a **Workspace form**, or to trigger it using Virtual Agent or Script Include. This option is available when you select **Instance** from the Trigger From list. ○ **Others** - Lens action is used when ServiceNow AI Lens is triggered to display a preview of extracted data and perform post processing steps, if applicable. This option is available when you select **Desktop** from the Trigger From list.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Order	Execution priority of the Lens action for a given table: ○ **1**: Highest priority ○ **999999**: Least priority The default value is set to 100. For example, if any user is assigned to two Lens action records for the same table, the Lens action with the least number is used during Lens execution.
Table	Table to which this Lens action applies. Selected table provides context during Lens execution when specific fields are selected or script is provided. For example, Incident table.
Request Context	Option to provide enhanced context during Lens execution by selecting specific fields or providing custom schema via script. ○ **Supported Fields** - Selected fields for providing enhanced context for Lens execution. ○ **Advanced** - Script to generate a custom schema object for providing enhanced context for Lens execution.
Fields	Table fields that you can select to provide enhanced context for Lens execution. Select fields to help Lens determine what information to extract from images. Only fields supported by Lens are available for selection. For more information, see Field types supported. This field appears only when **Supported Fields** is selected from the **Request Context** choice field.

Pre Processing Script	Script to generate a custom schema object to provide enhanced context for Lens execution. Custom schema help Lens determine what information to extract from images.
	This field appears only when **Advanced** is selected from the **Request Context** choice field.
	Example sample script:
	`function preProcessRequest(lensActionRecord, inputJSON) {`
	`  // Initialize the schema object`
	`  var schemaObj = {};`
	`  /*`
	`  This is a utility method to generate schema object,`
	`  users can pass tableName and formFields as an array of fields for the given table to generate custom schema that is passed to the LLM.`
	`  //Return schemaObj adhering to the structure defined in example`
	`  // tableName = 'incident'`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description

	<pre> tableName = inputJSON.tableName;
	
	    //formFields - array of fields for the given table.
	    //formFields = ['short_description', 'comments', 'impact'];
	
	    fieldsToProcess = inputJSON.formFields;
	    schemaObj = new global.AILensSchemaUtil().generateSchema(tableName, fieldsToProcess);
	
	    */
	    return schemaObj;
	}
	var schemaObj = preProcessRequest(lensActionRecord, inputJSON);</pre>
	Parameters accepted by the `preProcessRequest` method:
	○ **lensActionRecord** (Input) - GlideRecord of the Lens action that is executed.
	○ **inputJSON** (Input) - Optional JSON object that users can pass through service APIs.
	○ **schemaObj** (Output) - Returned JSON object of custom schema.
	
	Example of the returned schemaObj object:
	<pre>schemaObj = {
	  "schemaPayload": {
	    "type": "json_schema",
	    "json_schema": {
	      "name": "schema_extraction",
	      "schema": {
	        "type": "object",
	        "properties": {
	          "short_description": {
	            "type": "string",
	            "label": "Short description"
	          },
	          "comments": {
	            "type": "string",
	            "label": "Additional comments"
	          },
	          "impact": {
	            "type": "choice",
	            "label": "Impact",
	            "choiceValues": ["1", "2", "3"],
	            "choiceLabels": {
	              "1 - High": "1",
	              "2 - Medium": "2",
	              "3 - Low": "3"
	            }
	          }</pre>

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	<pre> }, "required": ["short_description", "comments", "impact"], "additionalProperties": false } } }, "fieldLabels": ["short_description", "comments", "impact"] //Optional key }</pre>
Default Additional Directions	Default instructions that are supplied to Now Assist in addition to the user-provided instructions. For example, `Summarize this paper by extracting the title, author name, publication date in the dd-mm-yyyy format, abstract, and keywords`
Transform Response	Option to specify the transformation script for transforming the Lens response before auto-filling the form. This field appears only when **Form** is selected from the **Trigger For** choice field.

Transformation Script	Script that transforms the Lens response before populating the form.
	
	This field appears only when the **Transform Response** option is selected.
	
	Sample script:
	<pre>function
	 transformResponseScript(lensActionRecord,
	 lensResponse) {
	   // Initialize the schema object
	    var transformResponse = {};
	
	    //Add your code here
	
	    //Return schemaObj adhering to the
	 structure defined in example
	    return transformResponse;
	}
	var transformResponse =
	 transformResponseScript(lensActionRecord,
	 lensResponse);</pre>
	Parameters accepted by the `transformResponseScript` method:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	○ **lensActionRecord** (Input) - GlideRecord of the Lens action that is executed. ○ **lensResponse** (Input) - JSON object of the Lens response provided after Lens execution. ○ **transformResponse** (Output) - Returned JSON object of transformed schema. Example of the returned transformResponse object: `transformResponse = { "short_description": "Service Degradation Error in Order Processing System", "description": "The Order Processing API v2.1 encountered a service degradation issue in the Production environment", "urgency": "1" }`
Enable Post Processing	Option to specify advanced processing instructions after Lens execution is complete. This field appears only when **Others** or **Service** is selected from the **Trigger For** field.
Post Processing Script	Script that is run asynchronously after Lens execution to further process the Lens response, such as triggering a flow, subflow, actions, or AI Agent. This field appears only when the **Enable Post Processing** option is selected. `function postProcessResponse(lensActionRecord, lensResponse, ) { //Add your code here } postProcessResponse(lensActionRecord, lensResponse, );` Parameters accepted by the `postProcessResponse` method: ○ **lensActionRecord** (Input) - GlideRecord of the Lens action that is executed. ○ **lensResponse** (Input) - JSON object of the Lens response provided after Lens execution.
Wait for processed response	Option to wait for ServiceNow AI Lens to display the output of post-processing on the Preview window. If enabled, it waits for the post processing output. Else, it doesn't wait for post processing output but post processing continues in the background.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Max wait time (seconds)	The maximum time ServiceNow AI Lens waits to display the output of post-processing. If the output doesn't appear within this duration, the session times out while the post-processing continues in the background until completion.
Attach Captured Images To Record	Option to attach the captured images to the record that is auto-filled using ServiceNow AI Lens. This option is enabled by default.

7. To assign users or groups to the Lens action, under the Assigned Users and Groups or Roles heading, select **New**, and then use any one of the following methods.

Method	Steps
Assign one or more users to the Lens action	**a.** To assign yourself, select the Add Me icon ( ). **b.** To assign one or more users, select the Un lock Users ( ) icon. **c.** Select the Look up icon ( ) and then se lect a user. **d.** Repeat the step to add more users. **e.** Select the Lock Users icon ( ).
Assign one or more groups to the Lens ac tion	**a.** Select the Unlock Groups ( ) icon. **b.** Select the Look up icon ( ) and then se lect a group. Repeat the step to add more groups. **c.** Select the Lock Groups icon ( ).

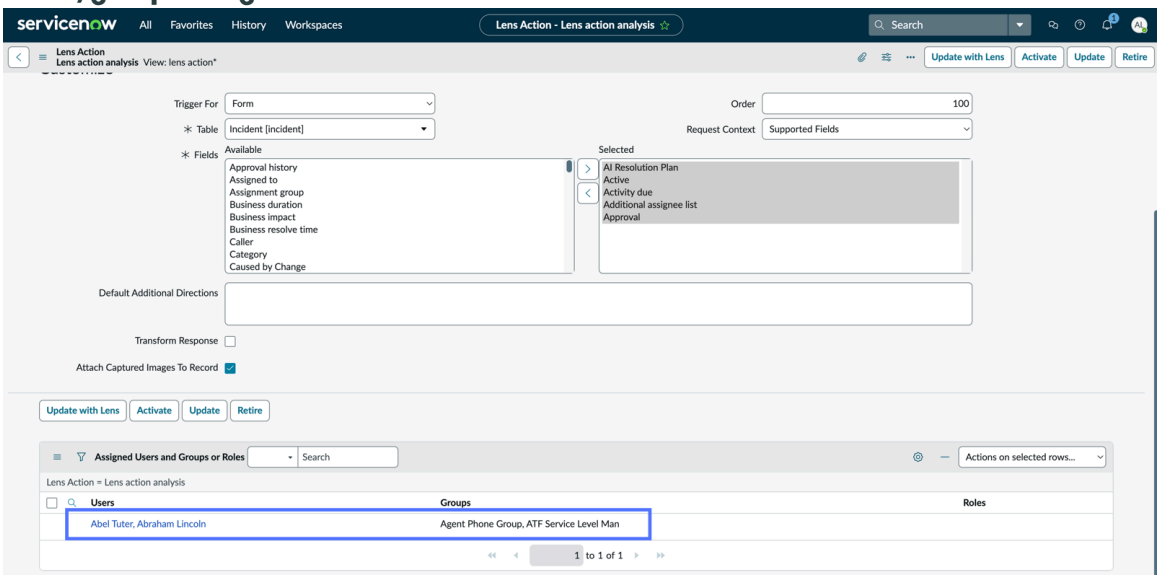
You can assign any one of the following combinations to the Lens action:

- Assign yourself and one or more groups to the Lens action.
- Assign one or more users and groups to the Lens action

8. Select **Submit**.
The users or groups are assigned to the Lens action.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Users, groups assigned to Lens action



9. Select **Update**.
10. **Optional:** Update the assignment of users or groups to the Lens action.

Assignment type	Steps
Users	**a.** Under the heading Assigned Users and Groups or Roles, select the user. **b.** To remove a user, select a user and then se lect () **c.** Assign one or more new users. For more in formation, see #unique_867_Connect_42_assign-users-to-lens-action
Groups	**a.** Under the heading Assigned Users and Groups or Roles, select the group. **b.** To remove a group, select a group and then select () **c.** Assign one or more new groups. For more information, see #unique_867_Connect_42_assign-groups-to-lens-action

11. Select **Activate** for activating the Lens action.
    The Lens action will be available on the ServiceNow AI Lens desktop application landing page.

Example: Scan documents to auto-fill the Insurance Request form

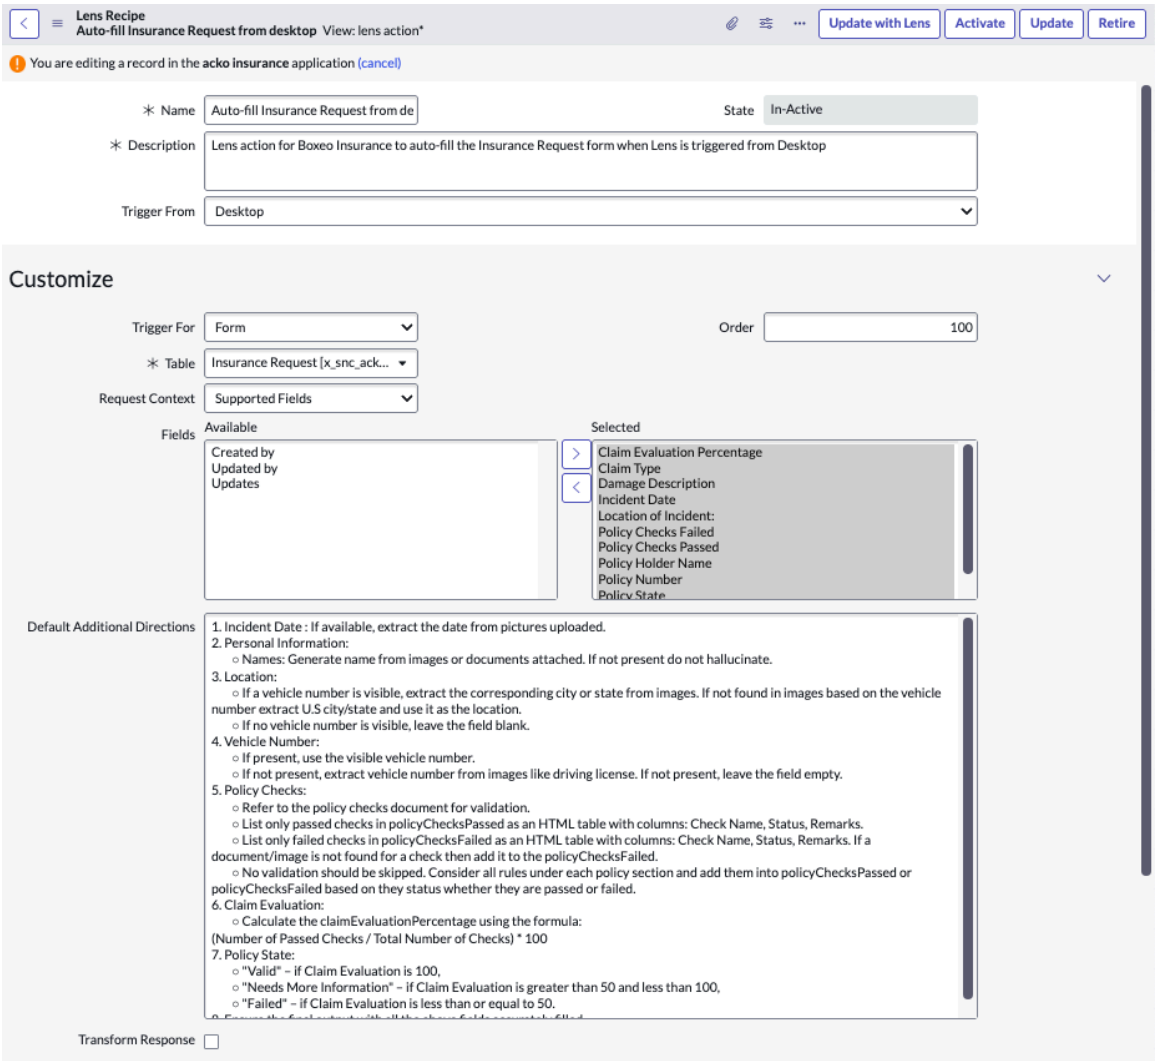
Enable customers to scan documents related to car accident on their desktop to auto-fill the Insurance Request form on the ServiceNow instance.

Before you begin
Role required: lens_admin

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

You can define in the Lens action that ServiceNow AI Lens must be triggered from the desktop to auto-fill the Insurance Request form. Lens action includes the fields that must be updated with the extracted data.



Procedure

1. Navigate to **All** > **ServiceNow Lens** > **Lens Actions**.
2. In the **Name** field, enter `Auto-fill Insurance Request from desktop`.
3. In the **Description** field, enter `Lens action for Boxeo Insurance to auto-fill the Insurance Request form when Lens is triggered from Desktop`.
4. From the **Trigger From** option, select Desktop.
5. In the form header, right-click and select **Save**.
6. From the **Trigger For** option, select Form.
7. In the **Table** field, select the Insurance Request table.
8. In the **Request Context** field, select Supported Fields to provide context.
9. From the **Fields** field, move fields to the Selected list.
   `policy_number, location_of_incident, policy_holder_name, vehicle_number, policy_state, policy_checks_failed,`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
policy_checks_passed, incident_date_time, damage_description,
claim_type, claim_evaluation_percentage
```

10. In the **Default Additional Directions** field, enter the following directions for the Now Assist.

    Example

```
1. Incident Date : If available, extract the date from pictures
 uploaded.
2. Personal Information:
 ○ Names: Generate name from images or documents attached. If
 not present do not hallucinate.
3. Location:
 ○ If a vehicle number is visible, extract the corresponding
 city or state from images. If not found in images based on
 the vehicle number extract U.S city/state and use it as the
 location.
 ○ If no vehicle number is visible, leave the field blank.
4. Vehicle Number:
 ○ If present, use the visible vehicle number.
 ○ If not present, extract vehicle number from images like
 driving license. If not present, leave the field empty.
5. Policy Checks:
 ○ Refer to the policy checks document for validation.
 ○ List only passed checks in policyChecksPassed as an HTML
 table with columns: Check Name, Status, Remarks.
 ○ List only failed checks in policyChecksFailed as an
 HTML table with columns: Check Name, Status, Remarks. If a
 document/image is not found for a check then add it to the
 policyChecksFailed.
 ○ No validation should be skipped. Consider all rules under
 each policy section and add them into policyChecksPassed
 or policyChecksFailed based on they status whether they are
 passed or failed.
6. Claim Evaluation:
 ○ Calculate the claimEvaluationPercentage using the formula:
(Number of Passed Checks / Total Number of Checks) * 100
7. Policy State:
 ○ "Valid" - if Claim Evaluation is 100,
 ○ "Needs More Information" - if Claim Evaluation is greater
 than 50 and less than 100,
 ○ "Failed" - if Claim Evaluation is less than or equal to 50.
8. Ensure the final output with all the above fields accurately
 filled.
```

11. Add users or groups who can use this Lens action.

    a. In the Assigned Users and Groups related list, select **New**.

    b. Select the required users or group of users that have the lens_user role.
    Only the assigned users or group of users can leverage the Lens action during Lens execution.

    c. Select **Submit**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

12. Select **Update**.
13. Select **Activate** for activating the Lens action.

Example: Extract structured data from vendor agreements

Enable compliance analysts to digitize and review several vendor agreements. Instead of manually entering details, they can launch ServiceNow AI Lens directly from the desktop, scan agreements, and adjust extracted values in the preview window before saving.

Before you begin
Role required: lens_admin

About this task
You can define in the Lens action that ServiceNow AI Lens must be triggered from the desktop to scan the documents saved on the user system. With the post processing option, you can edit the extracted information in the preview window before saving and trigger post processing to create related list records.

Procedure
1. Navigate to **All** > **ServiceNow Lens** > **Lens Actions**.
2. In the **Name** field, enter `Extract data from vendor agreements`.
3. In the **Description** field, enter `Digitize and review several vendor agreements`.
4. From the **Trigger From** option, select Desktop.
5. In the form header, right-click and select **Save**.
6. From the **Trigger For** option, select Others.
7. In the **Request Context** field, select Advanced.
8. In the **Pre Processing Script** field, enter the script to generate custom schema object for providing enhanced context.

   Example
   Here is the example of returned custom schema object:

```
schemaObj =  {
  "schemaPayload": {
    "type": "json_schema",
    "json_schema": {
      "name": "schema_extraction",
      "schema": {
        "type": "object",
        "properties": {
          "vendorName": {
            "type": "string",
            "label": "Vendor Name"
          },
          "contractTitle": {
            "type": "string",
            "label": "Contract Title"
          },
          "contractReferenceNumber": {
            "type": "string",
            "label": "Contract Reference Number"
          },
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
          "contractStartDate": {
            "type": "string",
            "label": "Contract Start Date"
          },
          "contractEndDate": {
            "type": "string",
            "label": "Contract End Date"
          },
          "renewalTerminationClauses": {
            "type": "string",
            "label": "Renewal / Termination Clauses"
          }
        },
        "required": [
          "vendorName",
          "contractTitle",
          "contractReferenceNumber",
          "contractStartDate",
          "contractEndDate"
        ],
        "additionalProperties": false
      }
    }
  },
  "fieldLabels": ["vendorName", "contractTitle",
 "contractReferenceNumber", "contractStartDate",
 "contractEndDate"]  //Optional key
}
```

9. In the **Default Additional Directions** field, enter directions for the Now Assist, such as `Leave the fields empty if you do not find the relevant information. Do not hallucinate. Enter the date in the DD-MM-YYYY.`
10. Select the **Enable Post Processing** option.
11. In the **Post Processing Script** field, enter the script to provide the post processing logic for the Lens response after the Lens execution is complete.

    Example
    You can provide a post processing logic to pass the data extracted to a subflow for creating related list records.
12. **Optional:** Select the **Wait for processed response** field to enable the ServiceNow AI Lens to wait for ServiceNow AI Lens to display the output of post-processing on the Preview window. If enabled, it waits for the post processing output. Else, it doesn't wait for post processing output but post processing continues in the background.
13. **Optional:** In the **Max wait time (seconds)** field, specify the maximum time ServiceNow AI Lens waits to display the output of post-processing.

    Note: If the output doesn't appear within that duration, the session times out while the post-processing continues in the background until completion.

14. Add users or groups who can use this Lens action.

    a. In the Assigned Users and Groups related list, select **New**.

    b. Select the required users or group of users that have the lens_user role.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Only the assigned users or group of users can leverage the Lens action during Lens execution.

c. Select **Submit**.

15. Select **Update**.
16. Select **Activate** for activating the Lens action.

Example: Auto-fill user records on a user table

Enable service agents to create user records by extracting specific information from identity documents using ServiceNow AI Lens, directly from the ServiceNow instance.

Before you begin
Role required: lens_admin

About this task
A service agent must onboard multiple users by scanning their official ID cards (for example, government-issued ID). Instead of manually entering each user's details, the agents can trigger ServiceNow AI Lens from the instance to automate data extraction and auto-fill specific fields defined in the Lens action record.

You can define in the Lens action that ServiceNow AI Lens must be triggered from the instance to create user records and provide default directions.

Procedure
1. Navigate to **All** > **ServiceNow Lens** > **Lens Actions**.
2. Select **New**.
3. In the **Name** field, enter `Auto-fill the user records`.
4. In the **Description** field, enter `Lens action to auto-create the user records on the ServiceNow instance`.
5. From the **Trigger From** option, select Instance.
6. In the form header, right-click and select **Save**.
7. Confirm that the Form option is selected in the **Trigger For** field.
   The Form option is selected by default.
8. In the **Table** field, select the User [sys_user] table.
9. Confirm that Supported Fields is selected in the **Request Context** field.
   The Supported Fields value is selected by default, and it provides context.
10. From the **Fields** field, move the fields to the Selected list.
    `userID, first_name, last_name, title, and date_of_birth`
11. In the **Default Additional Directions** field, enter `User ID must always be in lower case and uses the firstname.lastname format. Enter Date of birth in the dd-mm-yyyy format.`
12. Add users or groups who can use this Lens action.

    a. In the Assigned Users and Groups related list, select **New**.

    b. Select the required users or group of users that have the lens_user role.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Only the assigned users or group of users can leverage the Lens action during Lens execution.

c. Select **Submit**.

13. Select **Update**.
14. Select **Activate** for activating the Lens action.

Enabling ServiceNow AI Lens in Virtual Agent

You can enable ServiceNow AI Lens in Virtual Agent so that ServiceNow AI Lens can be triggered from Virtual Agent conversation.

By enabling ServiceNow AI Lens in Virtual Agent, you can trigger AI-powered ServiceNow AI Lens to scan an image for getting insights, recommendation, resolution steps, and more.

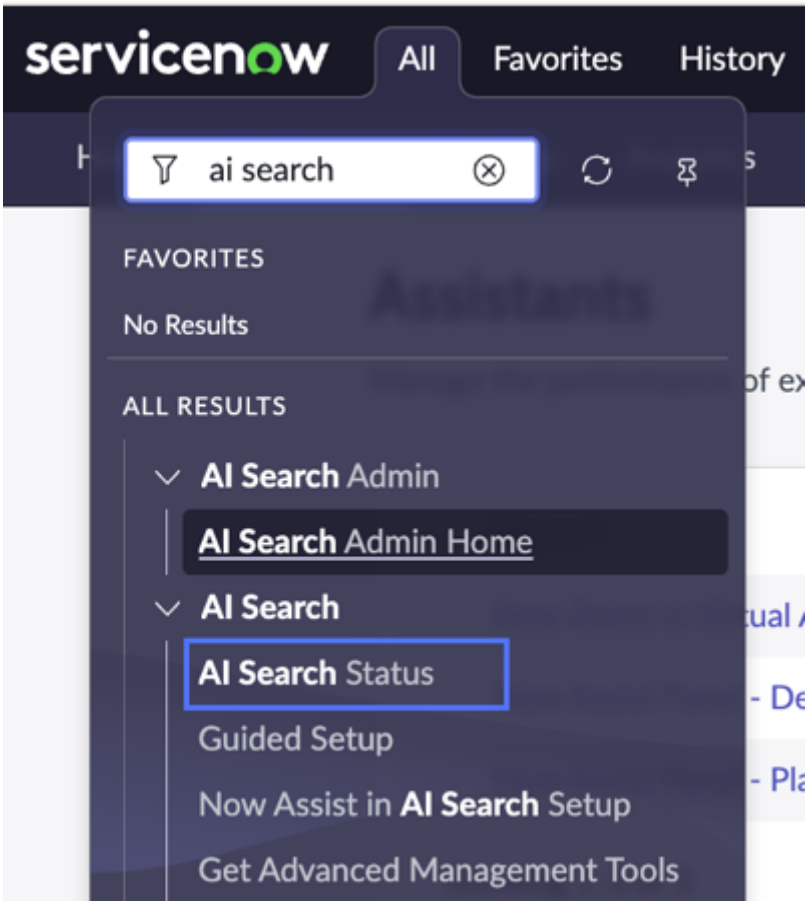
Enable AI Search

Enable the AI Search application on your instance so that you can use ServiceNow AI Lens skill to gather insights from an image using generative AI.

Before you begin
Role required: ais_admin or admin

Procedure

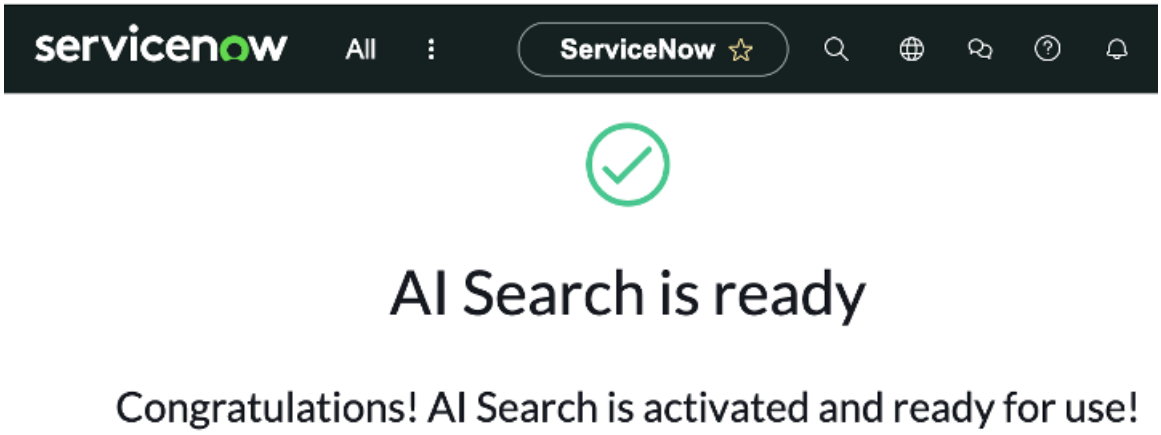
1. Navigate to **All** > **AI Search** > **AI Search Status**.



2. If AI search is not enabled, select **Request AI Search**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

After AI Search is activated in your instance, you will see the following success message.



Enable the default Now Assist assistant for ServiceNow AI Lens

Enable the default Now Assist in Virtual Agent assistant so that you can use ServiceNow AI Lens in Virtual Agent.

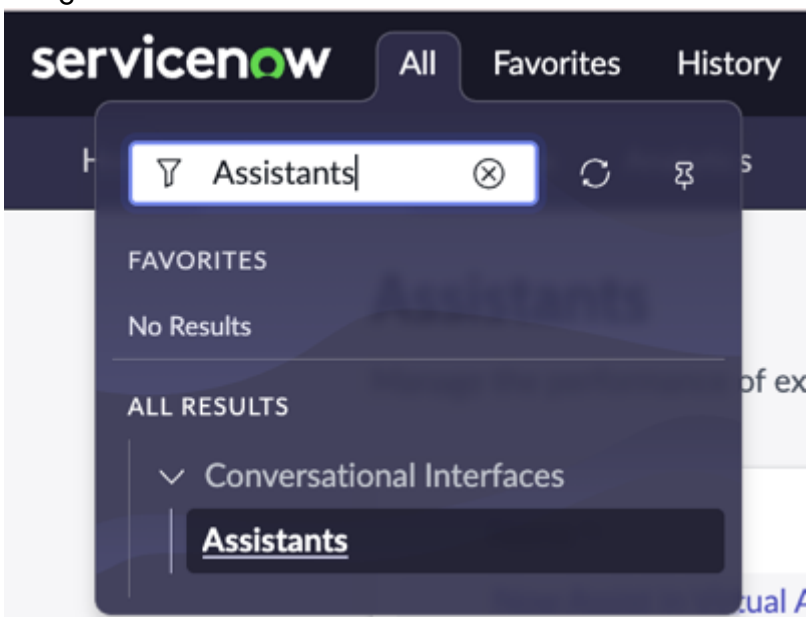
Before you begin

If you have another assistant configured for your specific use case or industry, you can skip this procedure and follow Enable other assistants in ServiceNow AI Lens topic.

Role required: virtual_agent_admin or admin

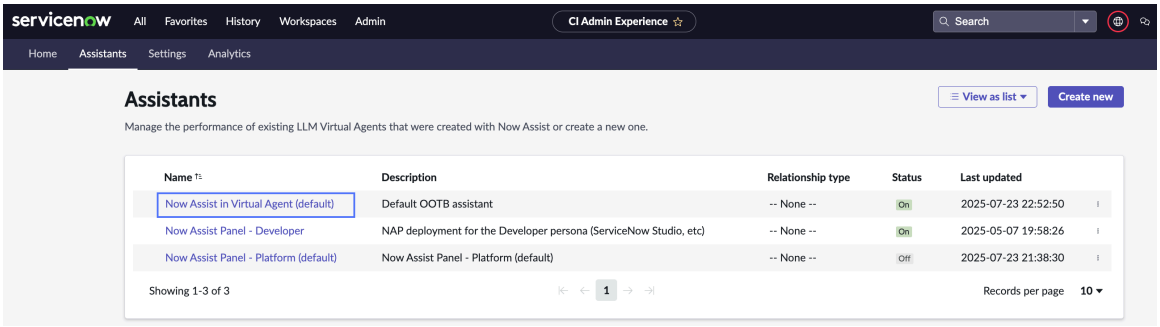
Procedure

1. Navigate to **All** > **Conversational Interfaces** > **Assistants**.

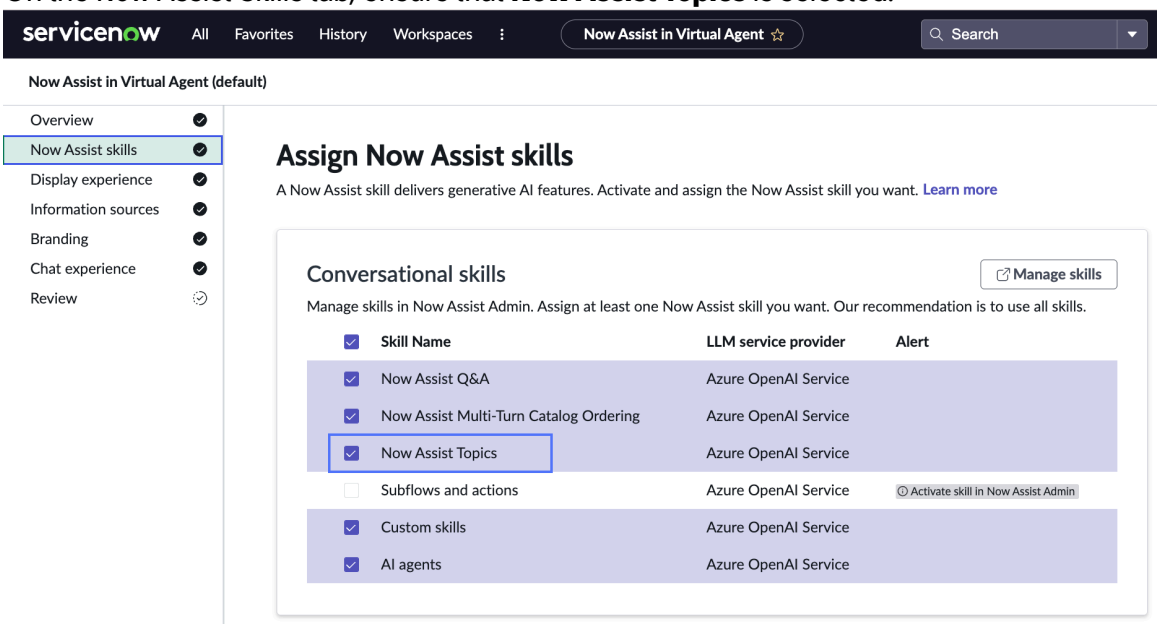


2. Select the **Now Assist in Virtual Agent (default)** assistant.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



3. On the Overview tab, select **Save and continue**.
4. On the Now Assist Skills tab, ensure that **Now Assist Topics** is selected.

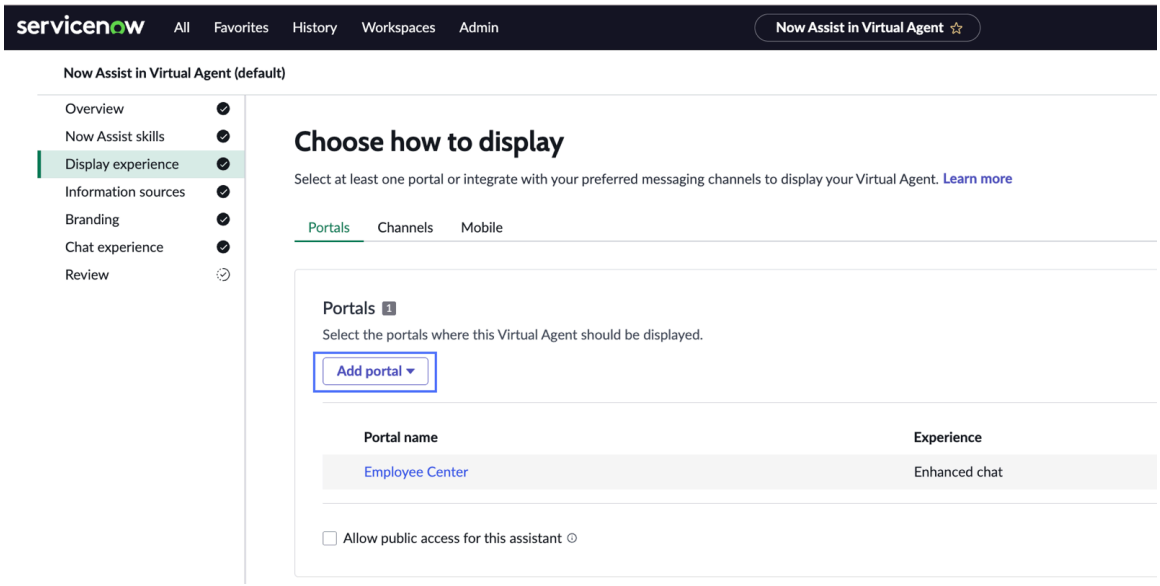


5. Select **Save and continue**.
6. Display your assistant on a portal.

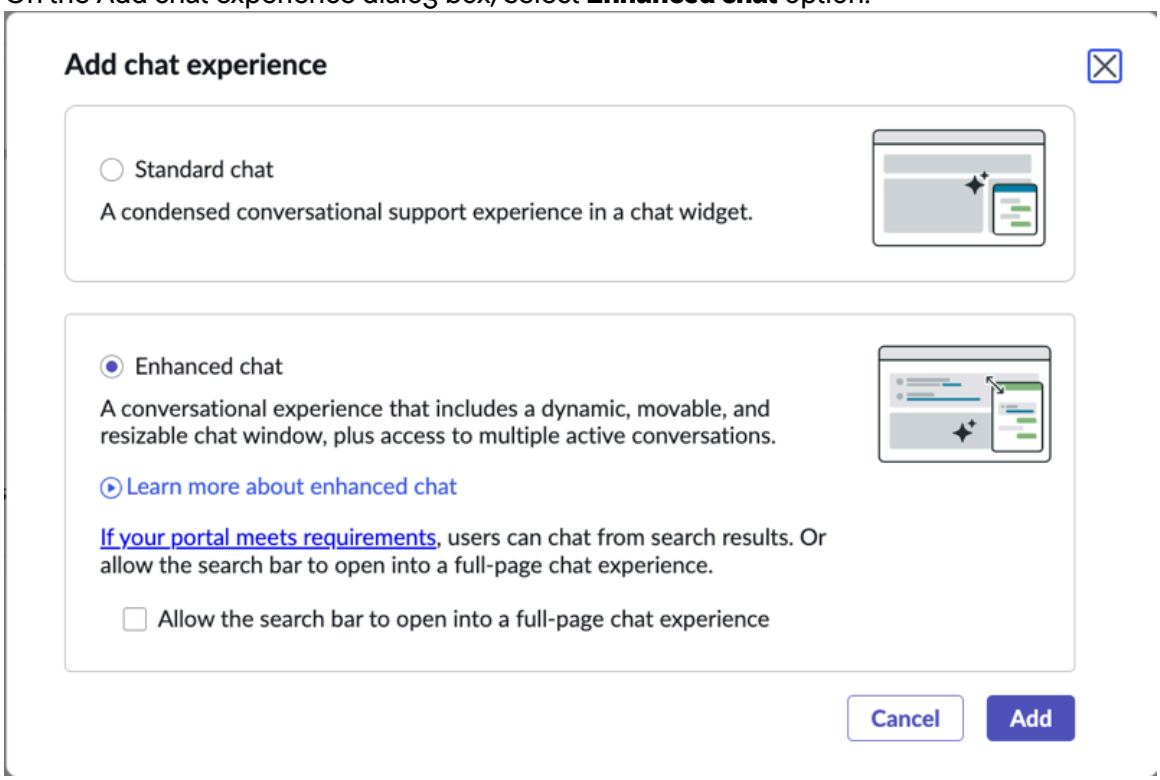
   a. On the Display experience tab, select the Portals tab.

   b. Select the portal for which you want to enable the Now Assist assistant from the **Add portal** drop-down.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



c. On the Add chat experience dialog box, select **Enhanced chat** option.



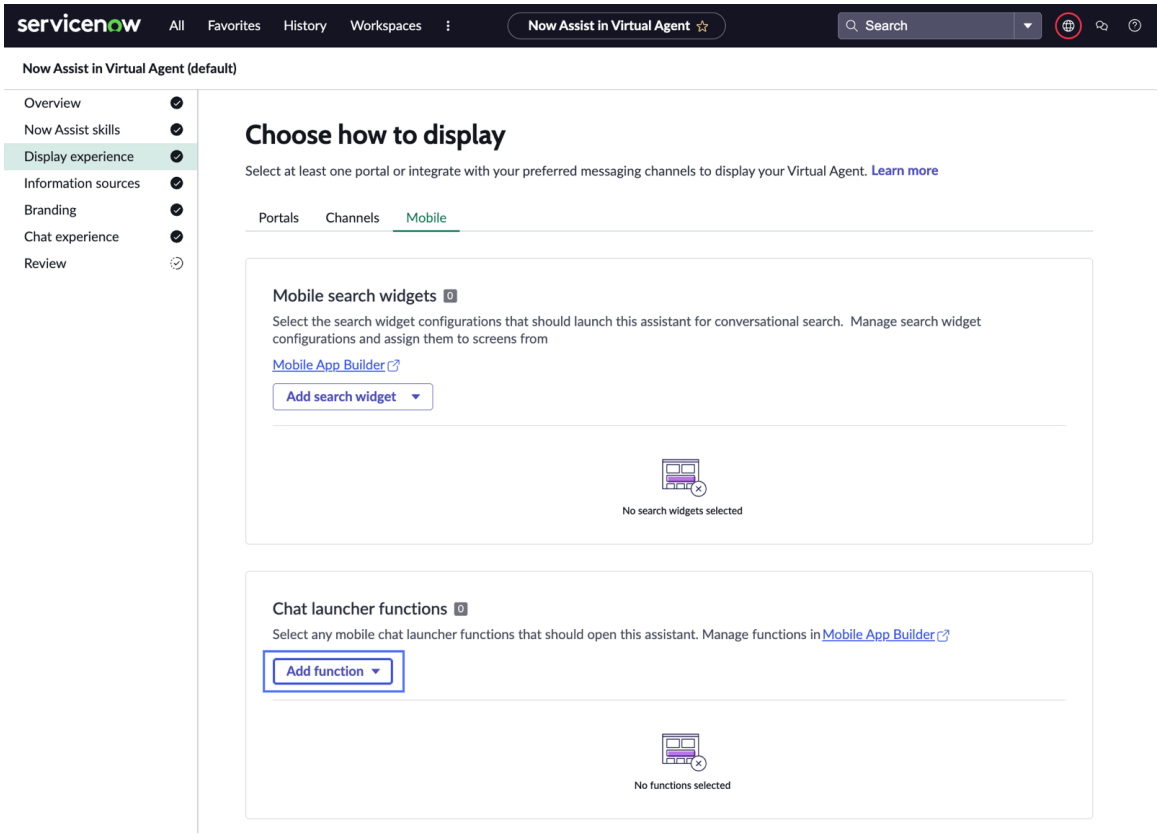
d. Select **Add**.

7. Display your assistant on a mobile device.

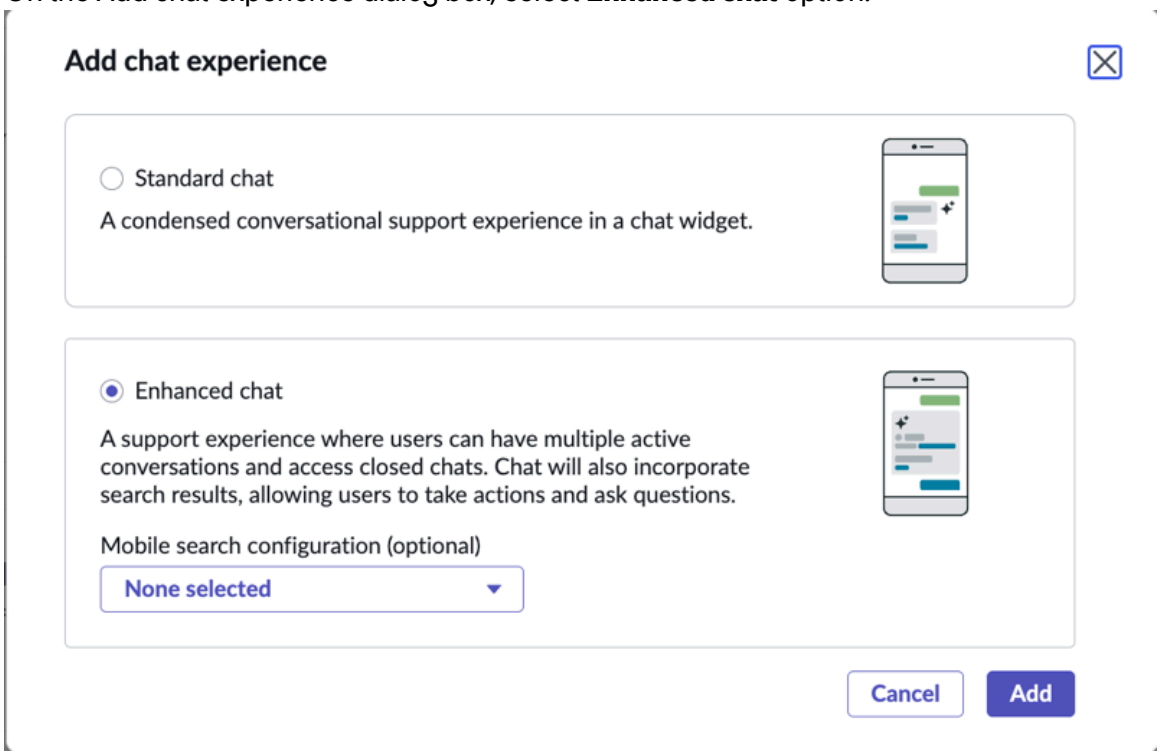
a. On the Display experience tab, select the Mobile tab.

b. Under the Chat launcher functions section, select Agent Chat from the **Add function** drop-down.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



c. On the Add chat experience dialog box, select **Enhanced chat** option.



d. Select **Add**.

8. On the Information sources tab, select **Save and continue**.

9. On the Branding tab, select **Save and continue**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

10. On the Chat experiences tab, select **Save and continue**.
11. On the Review tab, select **Done**.

Enable other assistants in ServiceNow AI Lens topic

Enable an assistant other than the default Now Assist in Virtual Agent assistant in ServiceNow AI Lens topic.

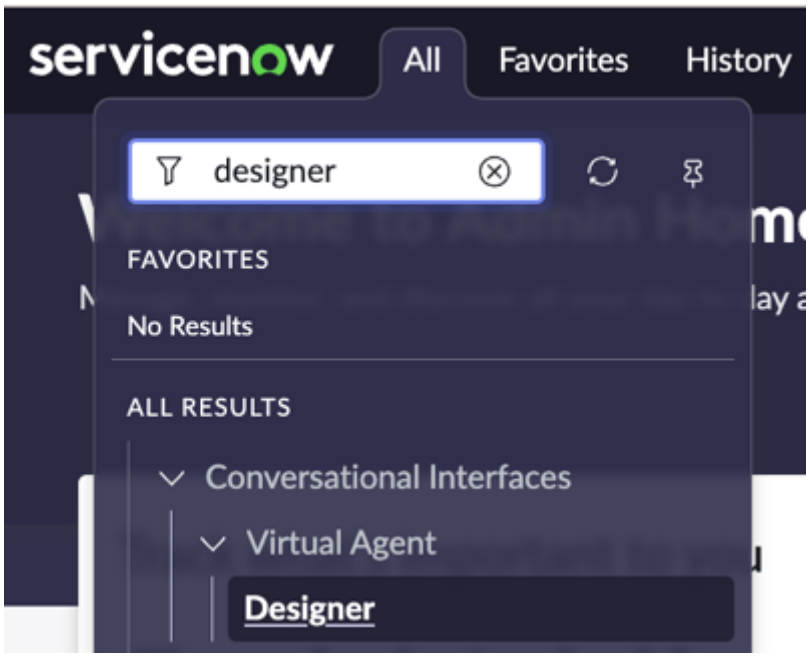
Before you begin

Skip this procedure, if you have enabled the default Now Assist assistant in the previous procedure.

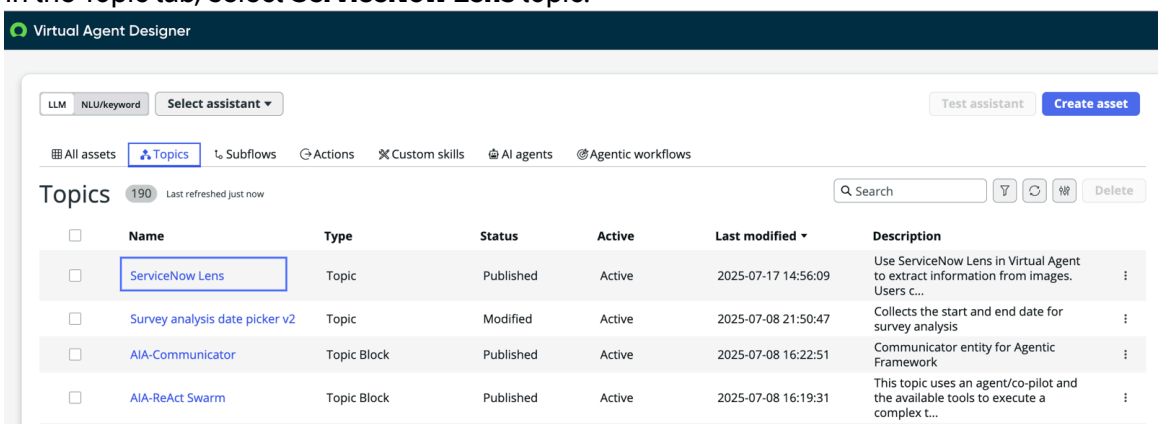
Role required: virtual_agent_admin or admin

Procedure

1. Navigate to **All** > **Conversational Interfaces** > **Virtual Agent** > **Designer**.
   Virtual Agent Designer opens in a new tab.

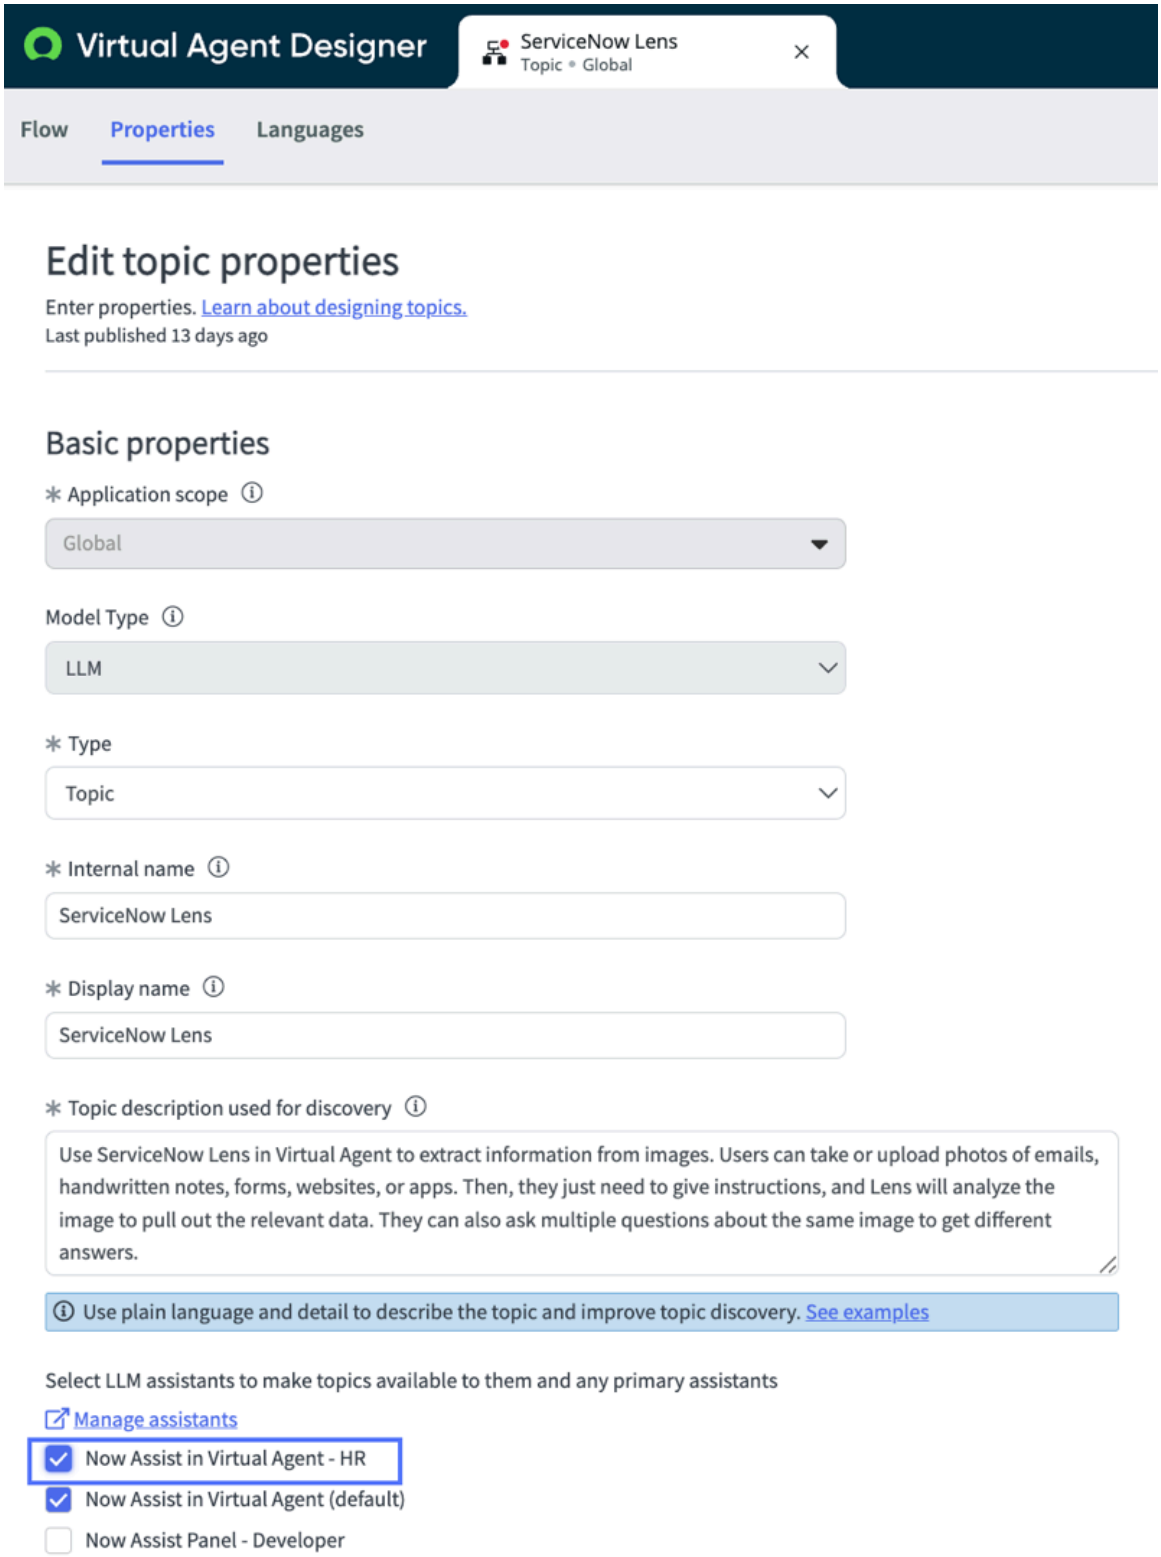


2. In the Topic tab, select **ServiceNow Lens** topic.



3. In the Select LLM assistants section, select your assistant.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. Publish the **ServiceNow Lens** topic.

(Optional) Enable Agent chat function for a mobile app

Enable the chat function in Virtual Agent for a mobile device for ServiceNow AI Lens.

Before you begin

Role required: admin or delegated developer

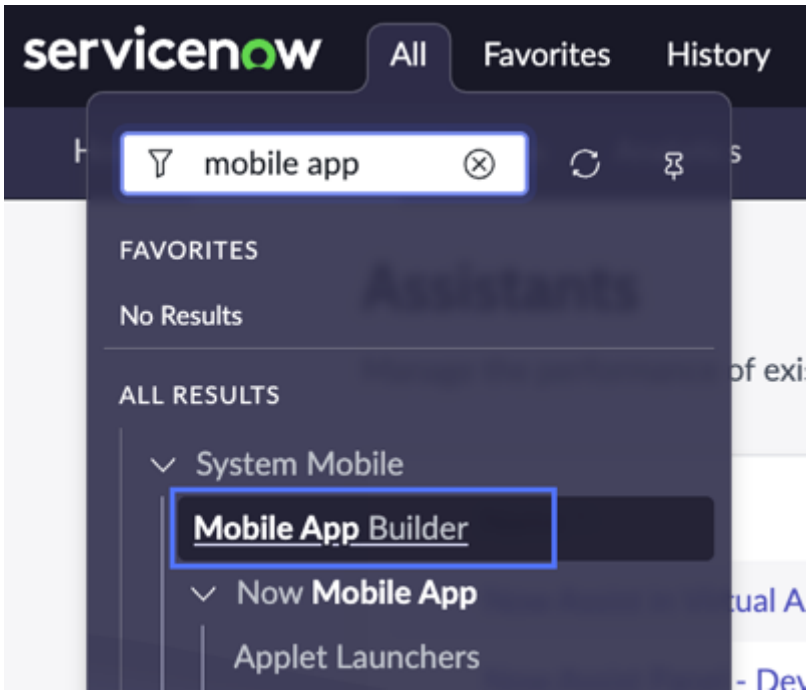
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Do this procedure when you do not see the Now Assist ✦ icon in the Now Mobile app.

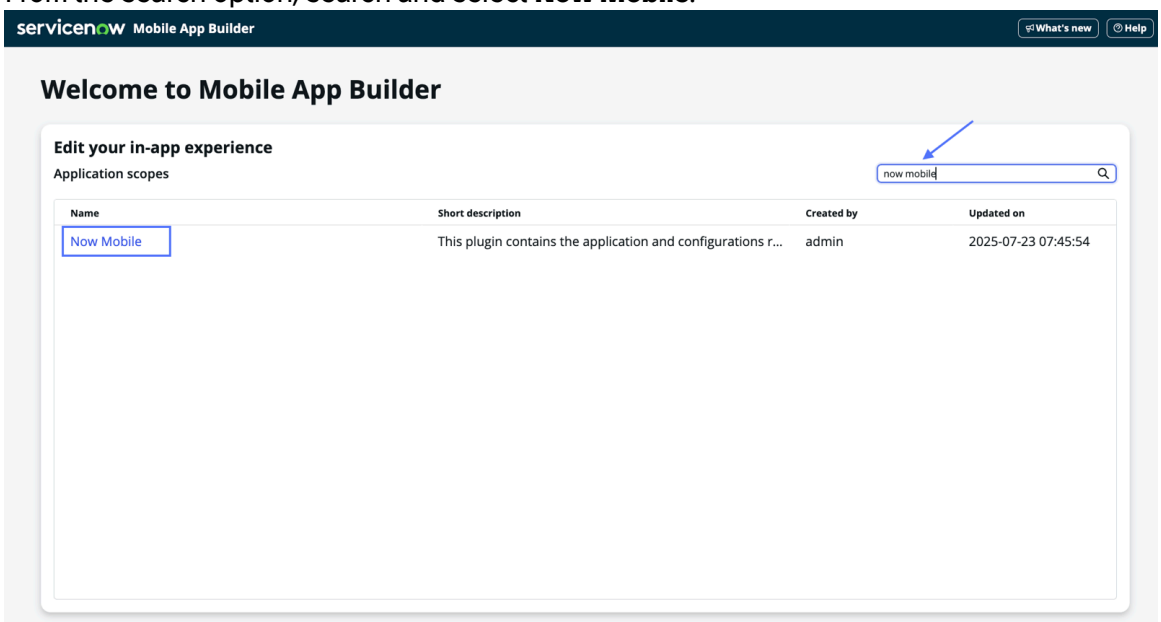
Procedure

1. Navigate to **All** > **System Mobile** > **Mobile App Builder**.



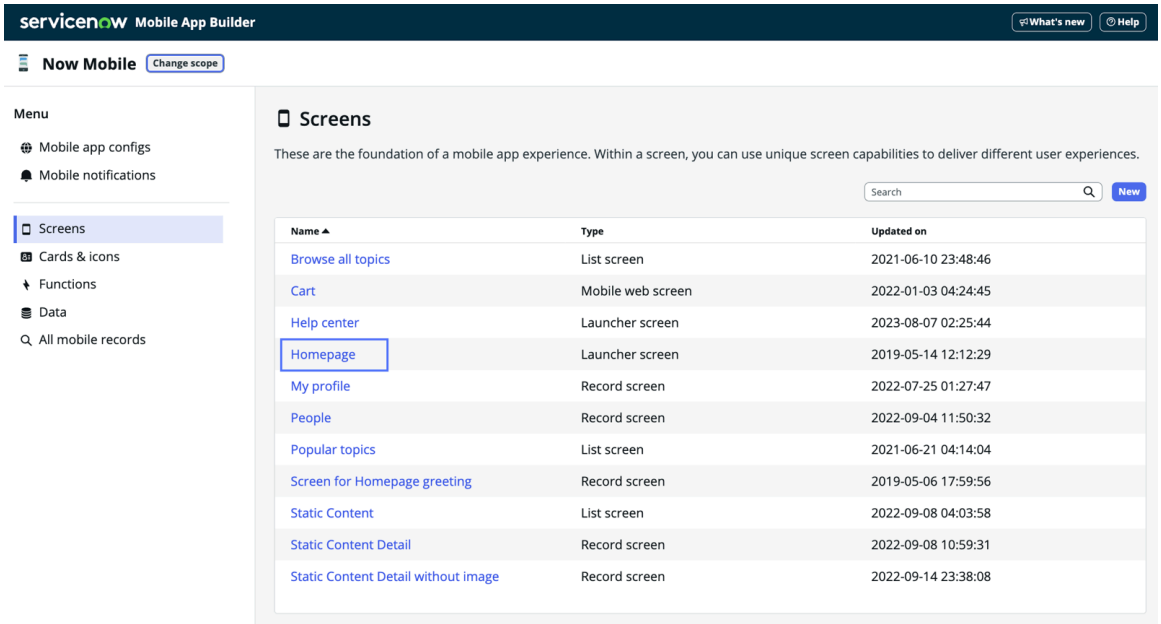
The Mobile App Builder opens in a new tab.

2. From the search option, search and select **Now Mobile**.

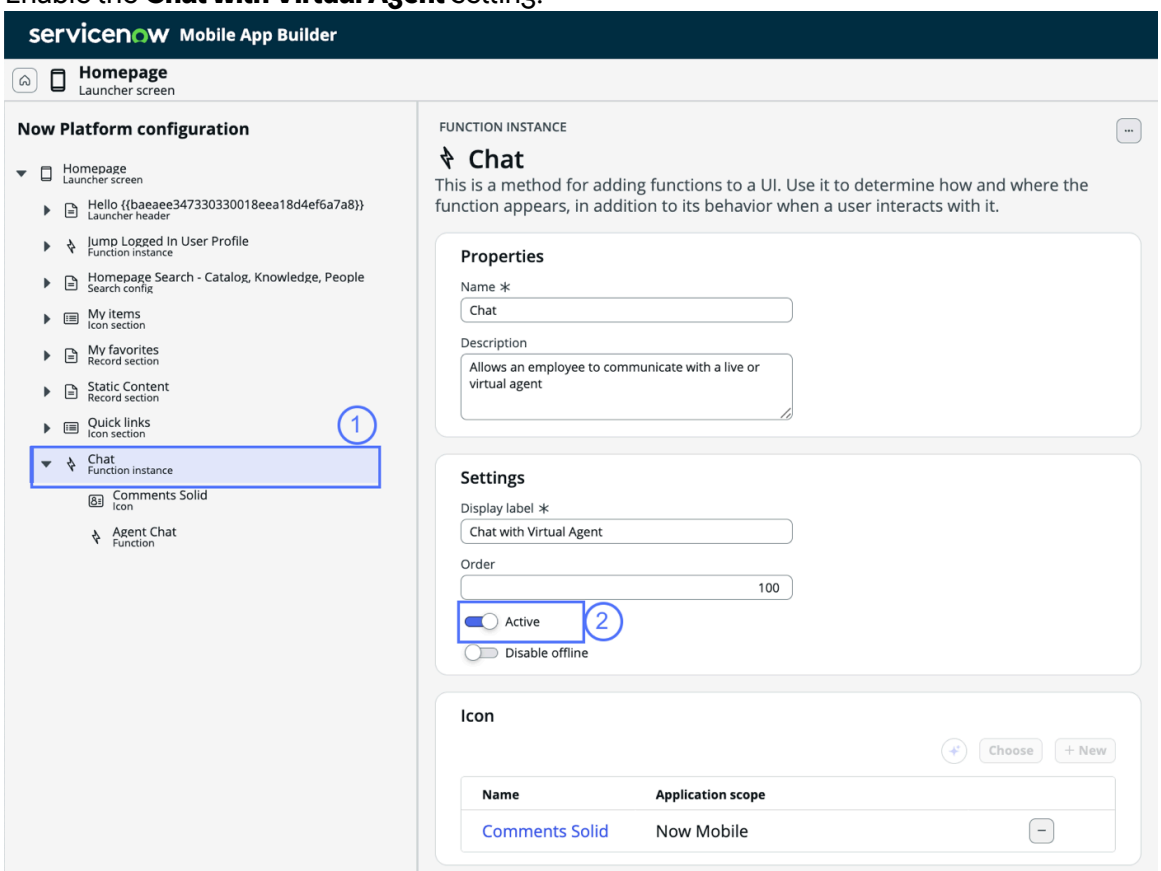


3. From the list of screens, select **Homepage**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. From the Now Platform configuration list, select the **Chat** function.
5. Enable the **Chat with Virtual Agent** setting.



Using ServiceNow AI Lens

If you have a lens_user role, you can use the ServiceNow AI Lens application to capture screens , extract and comprehend data from the captured screenshots, auto-fill form fields, or preview the extracted data.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Using ServiceNow AI Lens overview

Use ServiceNow AI Lens to do the following tasks:

- Create a record in the ServiceNow instance by auto-filling data that is scanned from captured images. For more information, see Create a record in an instance by using ServiceNow AI Lens.
- Update an existing record in the ServiceNow instance by auto-filling data that is scanned from captured images . For more information, see Update a record in an instance by using ServiceNow AI Lens.
- Trigger ServiceNow AI Lens in the context of Lens actions to auto-fill a form on the ServiceNow instance or view preview of extracted data. For more information, see Trigger ServiceNow AI Lens from the desktop app.
- Autofill the fields in a catalog item request form in the Service Portal. Launch ServiceNow AI Lens from the catalog item request form, then scan and extract data from artifacts such as Excel files, emails, or images of handwritten notes. For more information, see Autofill catalog item form in the Service Portal.
- Scan visual data, such as handwritten texts and notes, images, applications, and websites. You can use ServiceNow AI Lens as a standalone application to preview the generated outcome according to the instructions that you provide. For more information, see Extract and analyze data with ServiceNow AI Lens desktop app.
- Trigger ServiceNow AI Lens in the context of Lens actions from a Virtual Agent conversation. For more information, see Extract and analyze data using ServiceNow AI Lens from Virtual Agent.

Create a record in an instance by using ServiceNow AI Lens

Create a record in the ServiceNow instance by auto-filling the form fields with data that is extracted from one or more documents.

Before you begin

You must perform this task in Core UI.

To access the ServiceNow AI Lens functionality, perform the following steps:

- Install ServiceNow AI Lens on your ServiceNow instance. For more information, see Install the ServiceNow Lens in the ServiceNow instance.
- Turn on the ServiceNow AI Lens skill to add the generative AI capability. For more information, see Activate the ServiceNow AI Lens skill.
- Download the ServiceNow AI Lens installer to scan your desktop screen. For more information, see Download the ServiceNow AI Lens installer.

Verify that ServiceNow AI Lens has access to record the screen on your system. For more information, see Providing permission to ServiceNow AI Lens.

Use the **Create with Lens** button to launch the ServiceNow AI Lens scanner window and scan the data from documents, and auto-fill a form on the instance. You must have create and update access privileges to see the **Create with Lens** button in the list view of any table.

Don't scan any personally identifiable information, such as medical reports, financial reports, or other sensitive data, when using ServiceNow AI Lens as you don't want to expose the large language model (LLM) to any sensitive information.

Role required: lens_user

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. In your ServiceNow instance, navigate to the list view of any table.
   Example of a table: Incidents
2. Select **Create with Lens**.



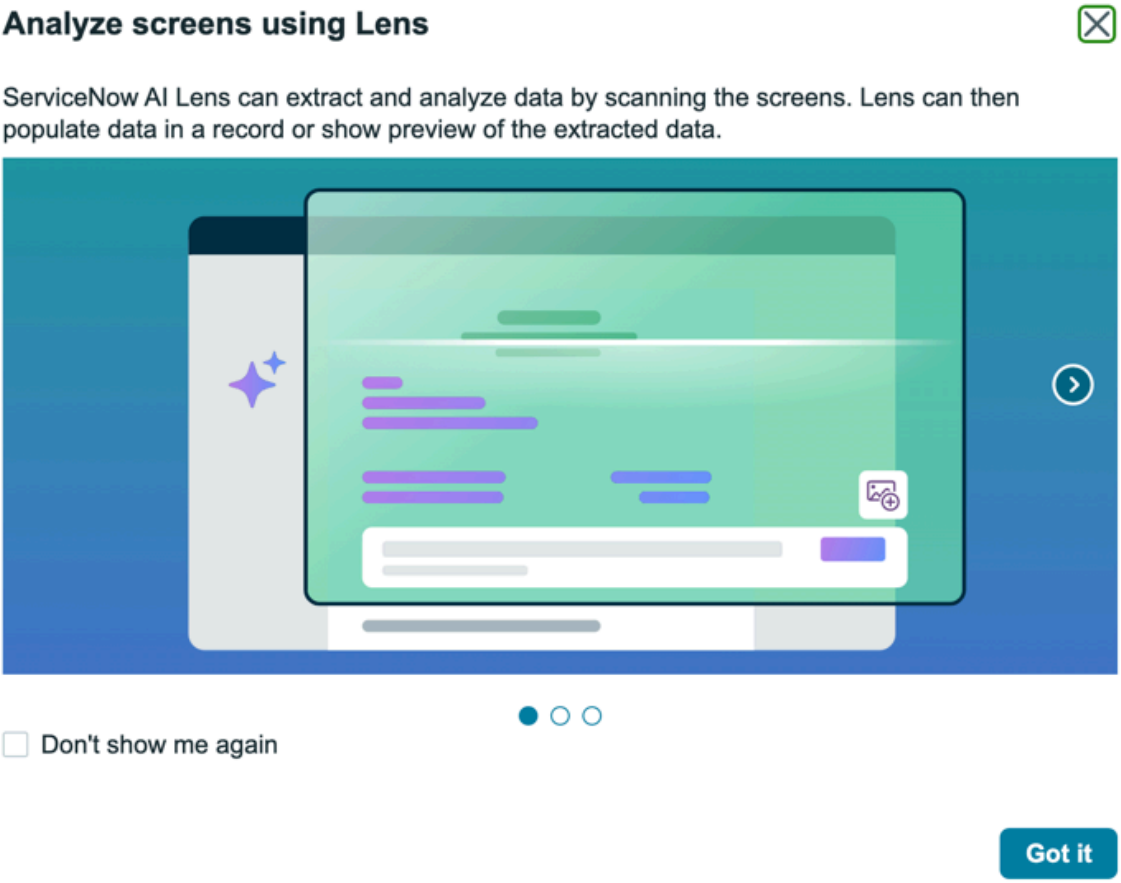
> **Note:** On sub-production instances, you can control which tables display the **Create with Lens** button using the following system properties:
>
> - `sn_app_lens_core.show_lens_action_on_all_tables`: Primary toggle. Set to true (default) to show Lens actions on all tables, or false to restrict them to a defined list of tables.
> - `sn_app_lens_core.lens_inclusion_table_list`: Comma-separated list of tables on which Lens actions should appear. Active only when the primary toggle is set to false.
> - `sn_app_lens_core.lens_exclusion_table_list`: Comma-separated list of tables on which Lens actions should be hidden, even if the primary toggle is set to true.

3. In the ServiceNow AI Lens.app dialog box, select Open ServiceNow AI Lens.app.

> **Tip:** This confirmation dialog appears when you select **Create with Lens** for the first time. You can make this a one-time step by selecting **Always allow <instance-name.service-now.com> to open links of this type in the associated app** before selecting Open ServiceNow AI Lens.app.

4. **Optional:** On the onboarding journey widget, complete the onboarding and select **Got it**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



If you launch the ServiceNow AI Lens for the first time, the onboarding journey widget appears. You can select **Don't show me again** to hide the widget the next time you launch ServiceNow AI Lens.

5. On your system, open an artifact that you want to scan.
   An artifact can be an image, scanned or handwritten note, website, or application.
6. Place the ServiceNow AI Lens scanner window on the top of the artifact.
   You can resize the ServiceNow AI Lens scanner window by dragging its borders.
7. **Optional:** Extract the data from multiple screenshots.

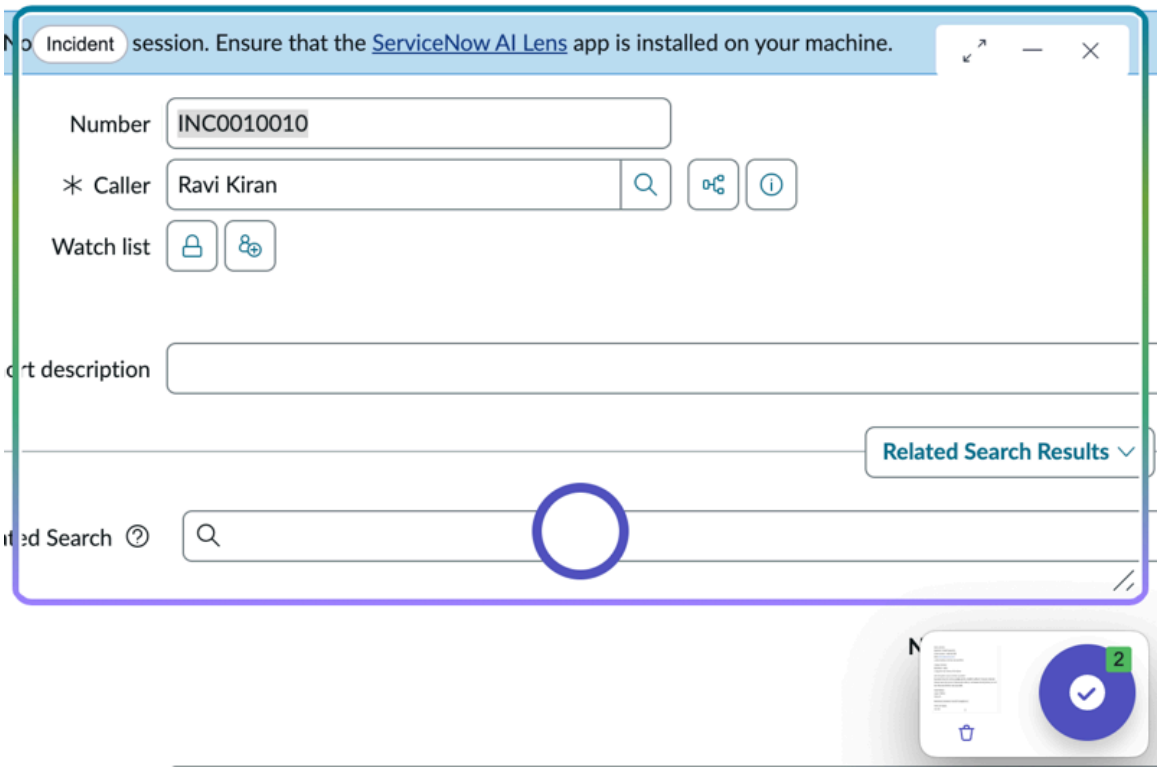
   a. Select the **Multi-capture** button .

   b. Continue to take the required screenshots by selecting the **Capture** button .
   You can capture up to 10 images to ServiceNow AI Lens.

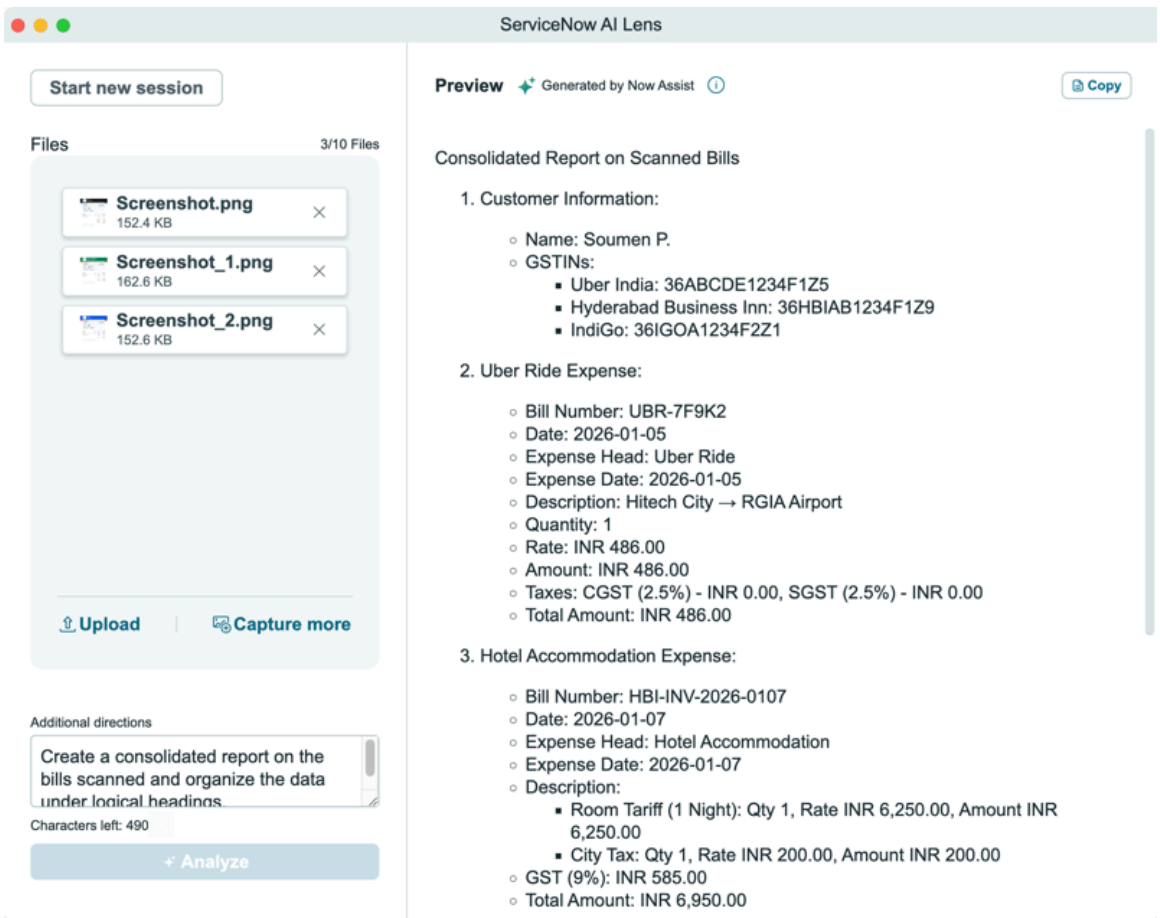
   You can take multiple screenshots of the same artifact or different artifacts.

   c. When you're finished taking screenshots, select **Done** .

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



The window shows the thumbnails of the images that ServiceNow Lens captured.



8. In the ServiceNow instance, review the text that is auto-filled by Now Assist into your record.
The fields that are auto-filled by Now Assist are highlighted with the Sparkle icon ✦.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Only the fields that are supported by ServiceNow AI Lens get auto-populated with the extracted data. If you don't have any supported fields in your form, then ServiceNow AI Lens won't update the record. For more information about the supported fields, see Field types supported.

Option	Action
If the auto-filled text looks good	Save the record by selecting **Save**.
If the auto-filled text requires changes	Do one of the following actions: ○ Manually adjust the information in the fields and save the record. ○ In the ServiceNow AI Lens window, provide different instructions or take more screen shots and select **Analyze** so that Service Now AI Lens can extract, comprehend the data again, and auto-fill the data into the record. Save the record by selecting **Save**. You can analyze the artifacts as many times as needed without reloading the form.

Related topics

Supporting information for ServiceNow AI Lens

ServiceNow AI Lens limitations

Update a record in an instance by using ServiceNow AI Lens

Extract and analyze data with ServiceNow AI Lens desktop app

Update a record in an instance by using ServiceNow AI Lens

Update a record in the ServiceNow instance by auto-filling the form fields with data that is extracted from one or more documents.

Before you begin

You must perform this task in Core UI.

To access the ServiceNow AI Lens functionality, perform the following steps:

- Install ServiceNow AI Lens on your ServiceNow instance. For more information, see Install the ServiceNow Lens in the ServiceNow instance.
- Turn on the ServiceNow AI Lens skill to add the generative AI capability. For more information, see Activate the ServiceNow AI Lens skill.
- Download the ServiceNow AI Lens installer to scan your desktop screen. For more information, see Download the ServiceNow AI Lens installer.

Verify that ServiceNow AI Lens has access to record the screen on your system. For more information, see Providing permission to ServiceNow AI Lens.

Use the **Update with Lens** button to launch the ServiceNow AI Lens scanner window and scan the data from documents, and update a form on the instance. You must have create and update access privileges to see the **Create with Lens** button in the list view of any table.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Don't scan any personally identifiable information, such as medical reports, financial reports, or other sensitive data, when using ServiceNow AI Lens as you don't want to expose the large language model (LLM) to any sensitive information.

Role required: lens_user

Procedure

1. On your ServiceNow instance, navigate to the list view of a table.
   Example of a table: Incidents
2. Select the record that you want to update.
3. Select **Update with Lens**.



Note: On sub-production instances, you can control which tables display the **Update with Lens** button using the following system properties:

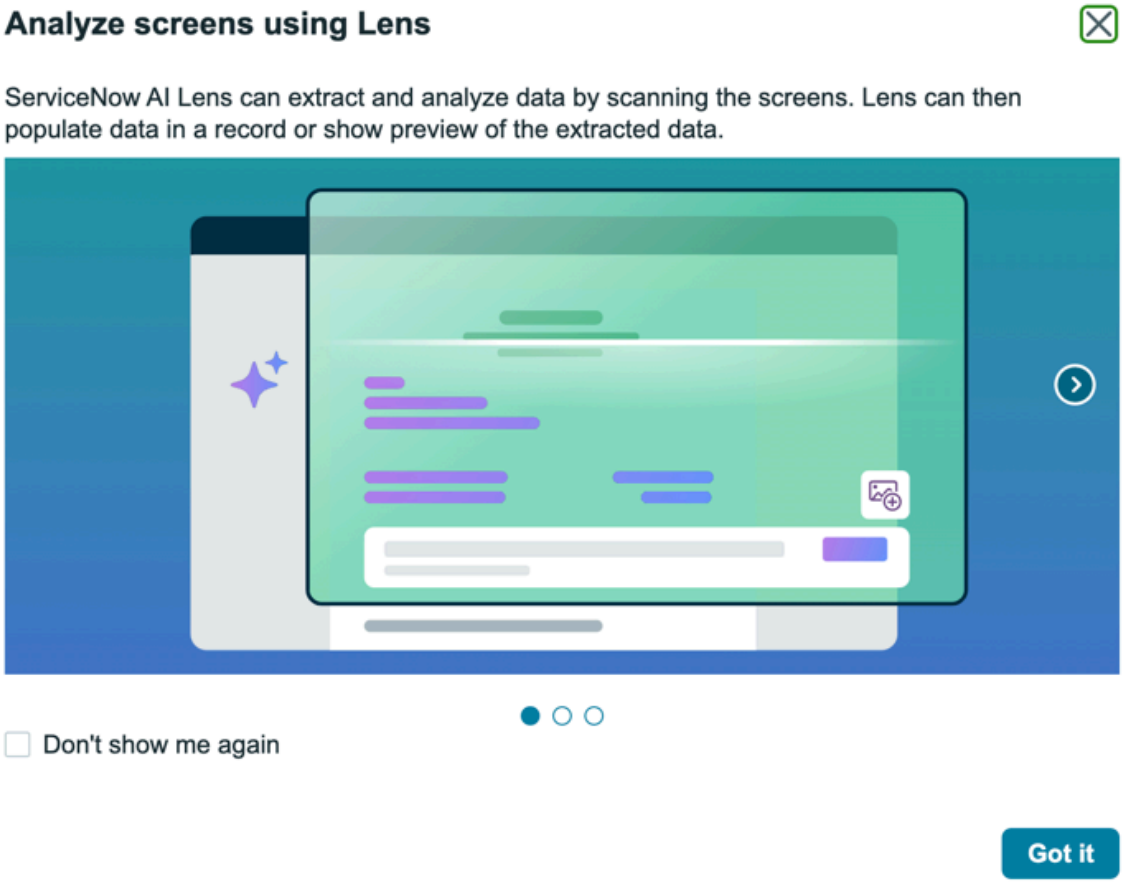
- `sn_app_lens_core.show_lens_action_on_all_tables`: Primary toggle. Set to true (default) to show Lens actions on all tables, or false to restrict them to a defined list of tables.
- `sn_app_lens_core.lens_inclusion_table_list`: Comma-separated list of tables on which Lens actions should appear. Active only when the primary toggle is set to false.
- `sn_app_lens_core.lens_exclusion_table_list`: Comma-separated list of tables on which Lens actions should be hidden, even if the primary toggle is set to true.

4. In the ServiceNow AI Lens.app dialog box, select Open ServiceNow AI Lens.app.

Tip: This confirmation dialog appears when you select **Create with Lens** for the first time. You can make this a one-time step by selecting **Always allow <instance-name.service-now.com> to open links of this type in the associated app** before selecting Open ServiceNow AI Lens.app.

5. **Optional:** On the onboarding journey widget, complete the onboarding and select **Got it**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



If you launch the ServiceNow AI Lens for the first time, the onboarding journey widget appears. You can select **Don't show me again** to hide the widget the next time you launch ServiceNow AI Lens.

6. On your system, open an artifact that you want to scan.
   An artifact can be an image, scanned or handwritten note, website, or application.
7. Place the ServiceNow AI Lens scanner window on the top of the artifact.
   You can resize the ServiceNow AI Lens scanner window by dragging its borders.
8. **Optional:** Extract the data from multiple screenshots.

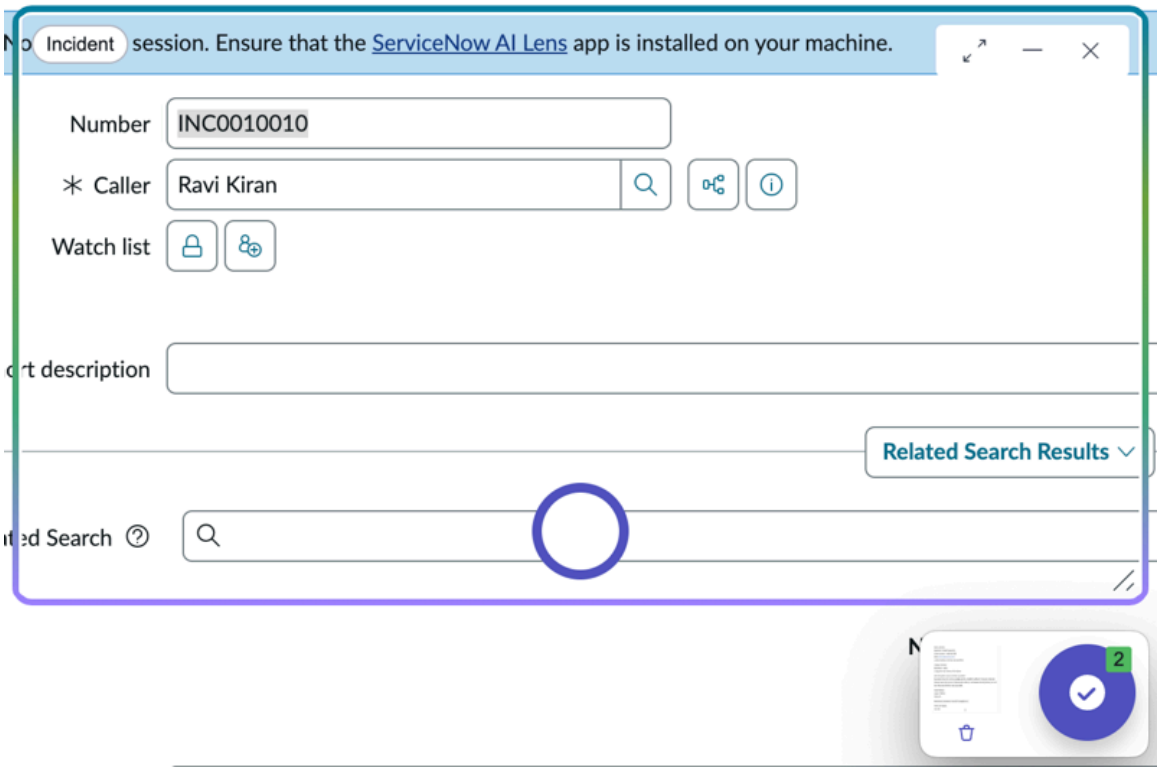
   a. Select the **Multi-capture** button .

   b. Continue to take the required screenshots by selecting the **Capture** button .
   You can capture up to 10 images to ServiceNow AI Lens.

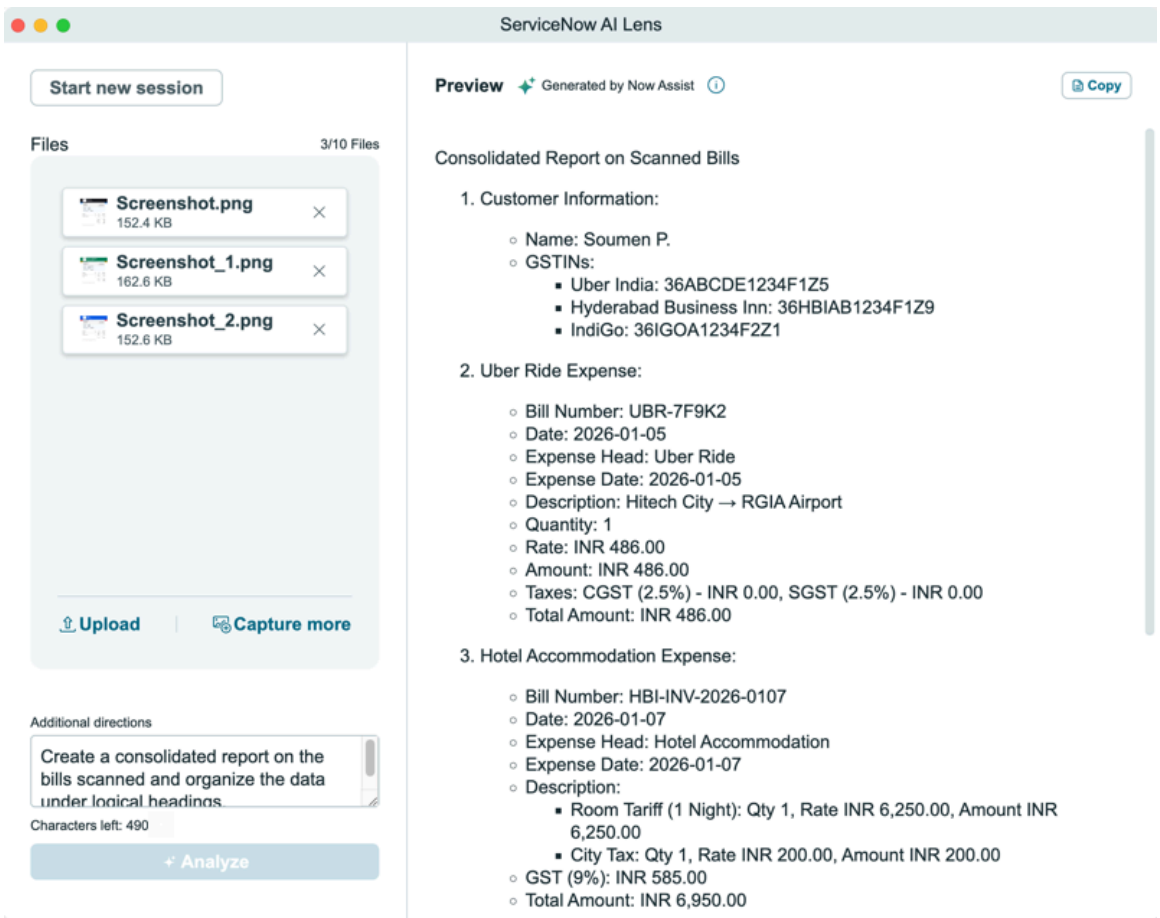
   You can take multiple screenshots of the same artifact or different artifacts.

   c. When you're finished taking screenshots, select **Done** .

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



The window shows the thumbnails of the images that ServiceNow Lens captured.



9. In the ServiceNow instance, review the text that is auto-filled by Now Assist into your record.
The fields that are auto-filled by Now Assist are highlighted with the Sparkle icon ✦.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Only the fields that are supported by ServiceNow AI Lens get auto-populated with the extracted data. If you don't have any supported fields in your form, then ServiceNow AI Lens won't update the record. For more information about the supported fields, see Field types supported.

Option	Action
If the auto-filled text looks good	Save the record by selecting **Save**.
If the auto-filled text requires changes	Do one of the following actions: ○ Manually adjust the information in the fields and save the record. ○ In the ServiceNow AI Lens window, provide different instructions or take more screen shots and select **Analyze** so that Service Now AI Lens can extract, comprehend the data again, and auto-fill the data into the record. Save the record by selecting **Save**. You can analyze the artifacts as many times as needed without reloading the form.

Related topics

Supporting information for ServiceNow AI Lens

ServiceNow AI Lens limitations

Create a record in an instance by using ServiceNow AI Lens

Extract and analyze data with ServiceNow AI Lens desktop app

Autofill catalog item form in the Service Portal

Use ServiceNow AI Lens to extract data from artifacts and autofill catalog item forms in Service Portal. For example, autofill a new vendor onboarding form by extracting key details such as vendor name, address, contact email, and banking information from multiple artifacts, including Excel files, emails, images, and PDF documents.

Before you begin

Role required: lens_user

Procedure

1. Navigate to the Service Portal.
   The URL is:

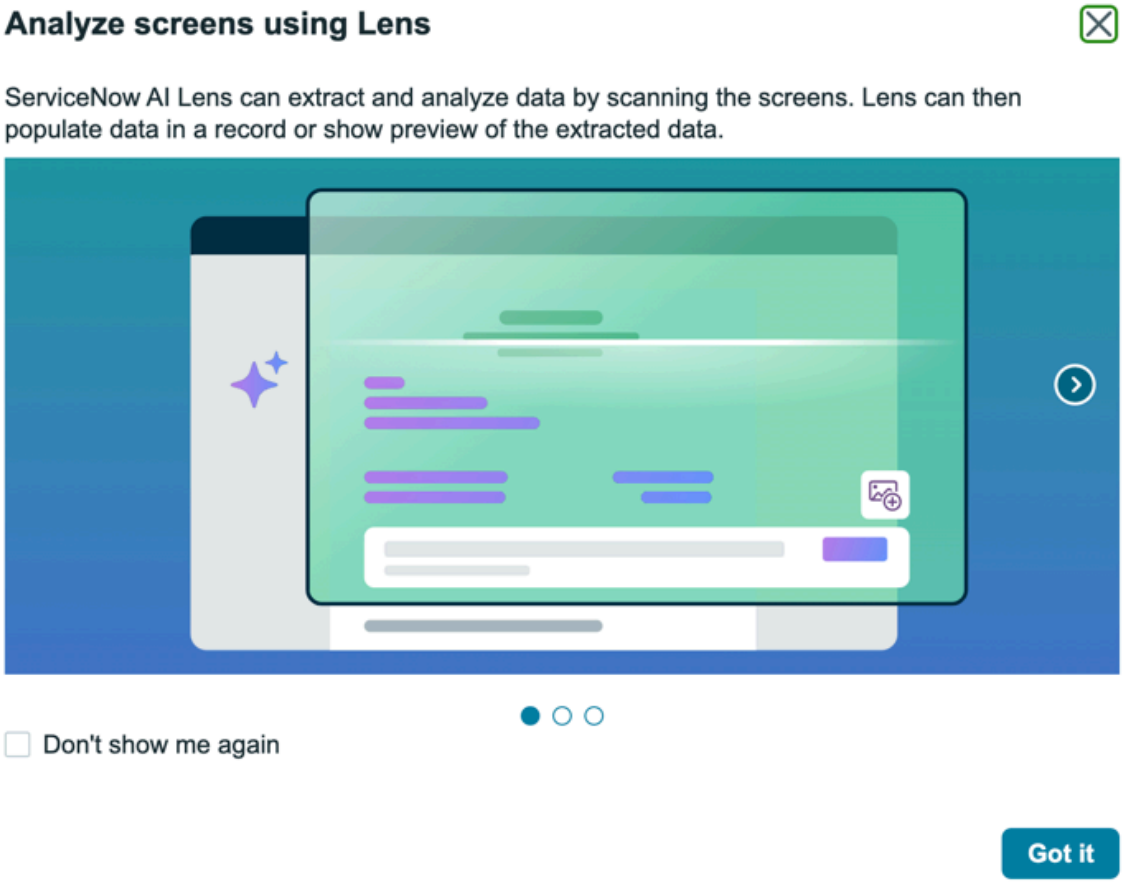
   ```
   https://<instance-name>.service-now.com/sp
   ```

2. Navigate to the page that lists the catalog item that you need.
3. Select the item category and the item from the catalog.
   For example, you could select New Vendor Registration.
4. Select **Fill with Lens**.
5. In the ServiceNow AI Lens.app dialog box, select Open ServiceNow AI Lens.app.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tip: This confirmation dialog appears when you select **Fill with Lens** for the first time. You can make this a one-time step by selecting **Always allow <instance-name.service-now.com> to open links of this type in the associated app** before selecting Open ServiceNow AI Lens.app.

6. **Optional:** On the onboarding journey widget, complete the onboarding and select **Got it**.

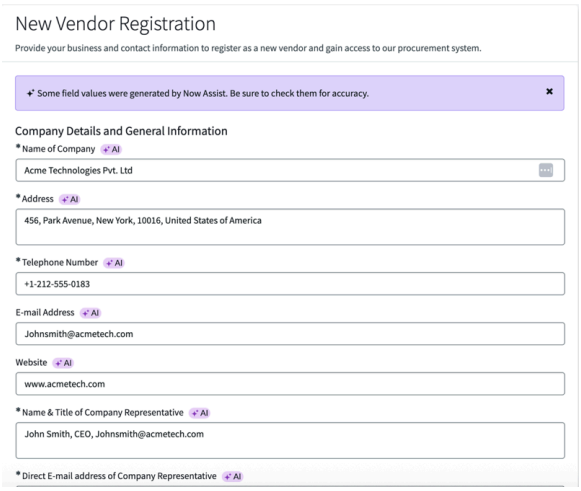


If you launch the ServiceNow AI Lens for the first time, the onboarding journey widget appears. You can select **Don't show me again** to hide the widget the next time you launch ServiceNow AI Lens.

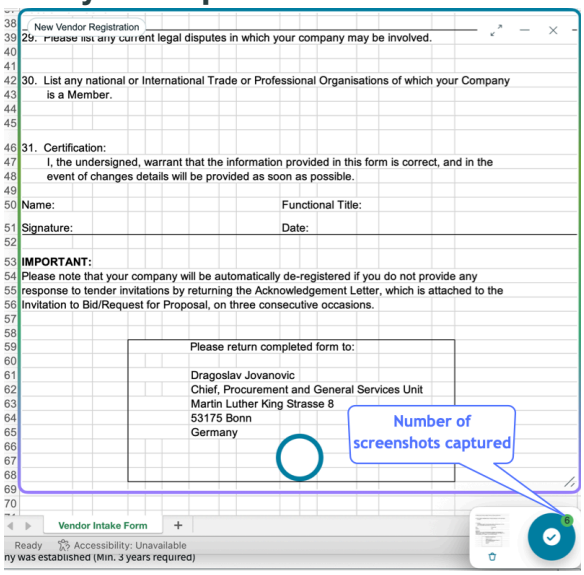
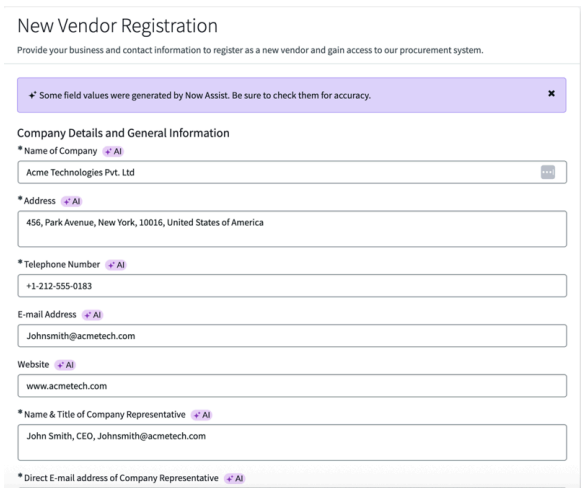
7. On your system, open an artifact that you want to scan.
An artifact can be an image, scanned or handwritten note, website, or application.
8. Place the ServiceNow AI Lens scanner window on the top of the document.
You can resize the ServiceNow AI Lens scanner window by dragging the scanner window borders.
9. Perform any one of the set of steps given below.

Method	Steps
Extract data from a single screenshot and auto-fill the form	**a.** (Optional step). To provide instructions to extract the data from the document in a specific way, select the Edit icon ( ) and enter instructions to analyze. The default character limit is 500. Users with the admin role can increase this lim

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Method	Steps
	it to up to 5000 characters by navigating to the `sn_lens_user_prompt_max_length` system property. **b.** Select **Analyze**. ServiceNow AI Lens confirms that the cata log item form is filled. **New Vendor Registration form fields auto-filled** 
Extract data from multiple screenshots and auto-fill the form	**a.** Select the Multi-capture icon ( ), and then place the scanner window over the document that you want to scan. You can resize the scanner window by drag ging its borders. **b.** (Optional step). To provide instructions to extract the data from the document in a spe cific way, select the Edit icon ( ), and en ter the instructions. The default character limit is 500. Users with the admin role can increase this lim it to up to 5000 characters by navigating to the `sn_lens_user_prompt_max_length` system property. **c.** Select the Capture icon ( ). The first screenshot is captured. **d.** Place the ServiceNow AI Lens scanner win dow over the page of another or the same document and then select the Capture icon ( ). The second screenshot is captured.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Method	Steps
	Repeat the step to capture more screen shots, if required. **View the number of screen shots you've captured**  **e.** (Optional step) To remove a screenshot that you had captured, select the Delete icon. **f.** To complete the capture, select the Done icon ( ). The ServiceNow AI Lens preview window displays the screenshots that you've cap tured. **g.** Select **Analyze**. The catalog item form is auto-filled. **New Vendor Registration form fields auto-filled** 

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

10. In the catalog item request form, confirm that the catalog item form fields are correctly filled. The fields that are auto-filled by Now Assist are highlighted with the Sparkle icon AI.

    Only the field types that are supported by ServiceNow AI Lens get auto-populated with the extracted data. If the form doesn't have field types that are supported, then ServiceNow AI Lens won't update the record. For more information about the supported fields, see Field types supported.

Option	Action
If the auto-filled text looks good	Save the record by selecting **Submit**.
If the auto-filled text requires changes	Do one of the following actions: ○ Manually adjust the information in the fields and save the record. ○ In the ServiceNow AI Lens window, provide different instructions or take more screen shots and select **Analyze** so that Service Now AI Lens can extract, comprehend the data again, and auto-fill the data into the record. Save the record by selecting **Sub mit**. You can analyze the artifacts as many times as needed without reloading the form.

Trigger ServiceNow AI Lens from the desktop app

Trigger ServiceNow AI Lens from the desktop app by using a Lens action to preview the extracted data and initiate post processing or auto-fill a form.

Before you begin

To access the ServiceNow AI Lens functionality, perform the following steps:

- Install ServiceNow AI Lens on your ServiceNow instance. For more information, see Install the ServiceNow Lens in the ServiceNow instance.
- Turn on the ServiceNow AI Lens skill to add the generative AI capability. For more information, see Activate the ServiceNow AI Lens skill.
- Download the ServiceNow AI Lens installer to scan your desktop screen. For more information, see Download the ServiceNow AI Lens installer.

Verify that you've defined the Lens action for this purpose. For more information, see Define a Lens action.

Verify that ServiceNow AI Lens has access to record the screen on your system. For more information, see Providing permission to ServiceNow AI Lens.

You must have create and update access privileges for creating or updating a record using ServiceNow AI Lens.

Don't scan any personally identifiable information, such as medical reports, financial reports, or other sensitive data, when using ServiceNow AI Lens as you don't want to expose the large language model (LLM) to any sensitive information.

Role required: lens_user

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

By using Lens actions, you can perform one of the following tasks:

Fill a form

In the Lens action configuration, when **Trigger From** is Desktop and **Trigger For** is Form, and a table is selected, ServiceNow AI Lens can be triggered from the desktop to auto-fill a form. It first shows a preview of the data, and if the preview is editable, you can edit the fields before saving. After you select **Submit** in the preview window, it saves the form with the updated data on the ServiceNow instance.

Preview the extracted data

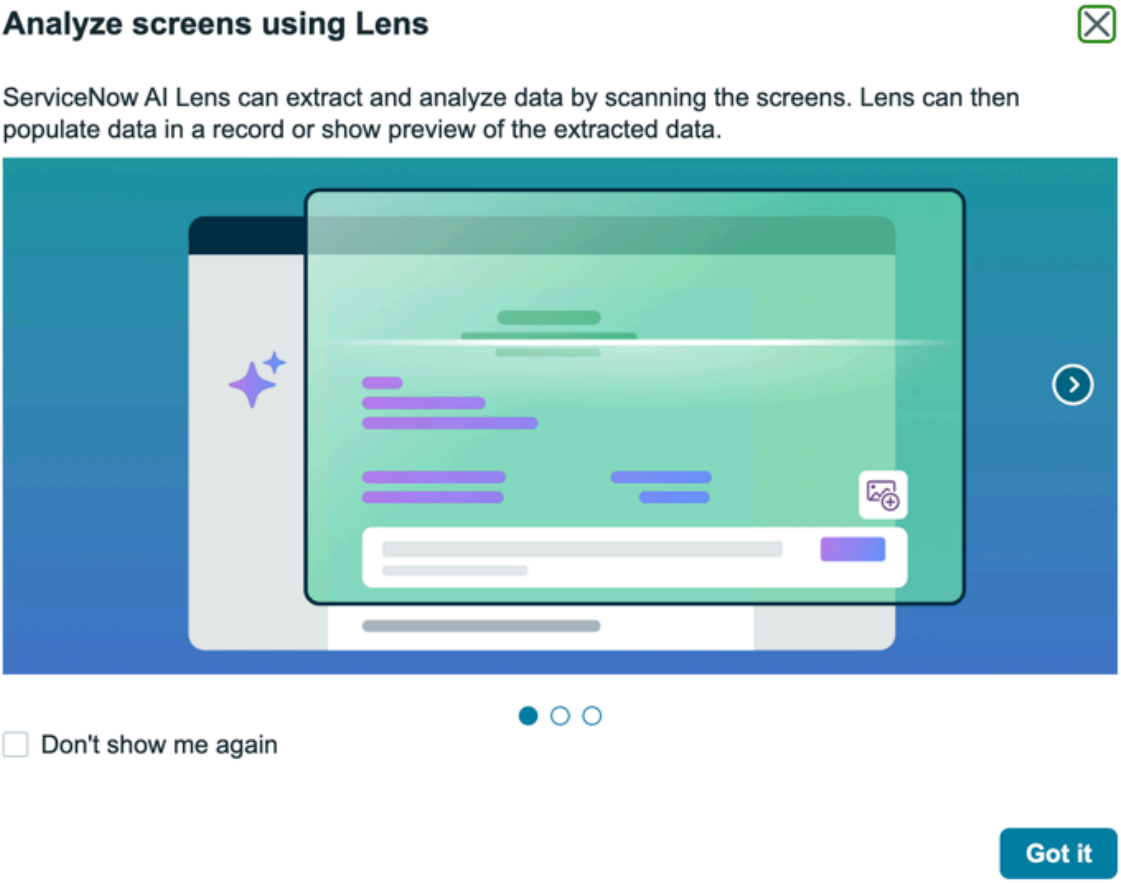
In the Lens action configuration, when **Trigger From** is Desktop and **Trigger For** is Others, ServiceNow AI Lens can be triggered from the desktop to show preview of the extracted data. If the preview is editable, you can edit the fields before saving. After you select **Submit and close preview** in the preview window, it saves the data and initiates post processing.

Procedure

1. From your system, launch the ServiceNow AI Lens application.
2. On the login page, in the **Instance URL** field, enter the ServiceNow instance URL.

   Example
   For example, `https://<instance name>.service-now.com`.
3. Select **Proceed**.
4. Log in to your ServiceNow account by entering your username and password.
   Your account must have the lens_user role.
5. **Optional:** On the onboarding journey widget, complete the onboarding and select **Got it**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



If you launch the ServiceNow AI Lens for the first time, the onboarding journey widget appears. You can select **Don't show me again** to hide the widget the next time you launch ServiceNow AI Lens.

6. Select a Lens action.
7. Select **Proceed with Lens**.
   Lens scanner opens in a separate window.

   Note: Use the scanner window to scan the data of the screenshot so you can submit it for analysis.

8. On your system, open an artifact that you want to scan.
   An artifact can be an image, scanned or handwritten note, website, or application.
9. Place the ServiceNow AI Lens scanner window on the top of the artifact.
   You can resize the ServiceNow AI Lens scanner window by dragging its borders.
10. Extract data from a single screen.

    a. Enter additional instructions to extract, analyze, and organize data by selecting the Edit icon (✎) and then enter the instructions.
    The default character limit is 500. Users with the admin role can increase this limit to up to 5000 characters by navigating to the `sn_lens_user_prompt_max_length` system property.

    Note: This is an optional step.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. Select **Analyze**.
The ServiceNow AI Lens preview window displays the extracted output in an editable form.

11. Optional: Extract data from multiple screenshots.

a. To provide additional instructions to extract, analyze, and organize data, select the Edit icon ( ) and then enter the instructions.
The default character limit is 500. Users with the admin role can increase this limit to up to 5000 characters by navigating to the `sn_lens_user_prompt_max_length` system property.

Note: This is an optional step.

b. Select the **Multi-capture** button .

Tip: Use the Upload file icon ( ) to upload one or more files. For more information, see #unique_871_Connect_42_upload-files-steps

c. Select the Capture icon ( )
The first screenshot is captured.

d. Place the ServiceNow AI Lens scanner window over another document or page and then select the Capture icon).
The second screenshot is captured.

(Optional) Repeat the step to capture more screenshots, if required.

Note:

- You can capture a total of 10 screenshots with the combined size of all captured screenshots not exceeding 10 MB.
- To enable the desktop app to send large screenshot data to the server, confirm that the following system properties are set exactly as shown below:

System property

Property name	Type	Recommended value
glide.rest.max_content_length	Integer	15
glide.rest.scripted.max_inbound_content_length_mb	Integer	15

For more information, see Configure system property

Note: You must have the admin role to set the system properties.

e. To complete the capture, select the Done icon ().

f. Select **Analyze**.
The data is extracted from the screenshots and the form that you can update and submit is shown.

12. Depending on the type of Lens action selected, perform the following steps.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Task	Steps
Filling form	**a.** (Optional) In case of editable preview, if necessary, edit the auto-filled field values before saving. **b.** On the form header in the Preview window, select **Submit** to save the filled form on the instance.
Previewing extracted data	**a.** (Optional) In case of editable preview, if necessary, edit the auto-filled field values before saving or copying the values. **b.** (Optional) Copy the previewed data by selecting the Copy icon .

13. End the current session by selecting **Start new session**.

Extract and analyze data with ServiceNow AI Lens desktop app

Extract and analyze data from one or more screenshots that you capture , and then preview the data analysis.

Before you begin

To access the ServiceNow AI Lens functionality, perform the following steps:

- Install ServiceNow AI Lens on your ServiceNow instance. For more information, see Install the ServiceNow Lens in the ServiceNow instance.
- Turn on the ServiceNow AI Lens skill to add the generative AI capability. For more information, see Activate the ServiceNow AI Lens skill.
- Download the ServiceNow AI Lens installer to scan your desktop screen. For more information, see Download the ServiceNow AI Lens installer.

Don't scan any personally identifiable information, such as medical reports, financial reports, or other sensitive data, when using ServiceNow AI Lens as you don't want to expose the large language model (LLM) to any sensitive information.

Verify that you have provided permission to ServiceNow AI Lens to record the screen on your system.

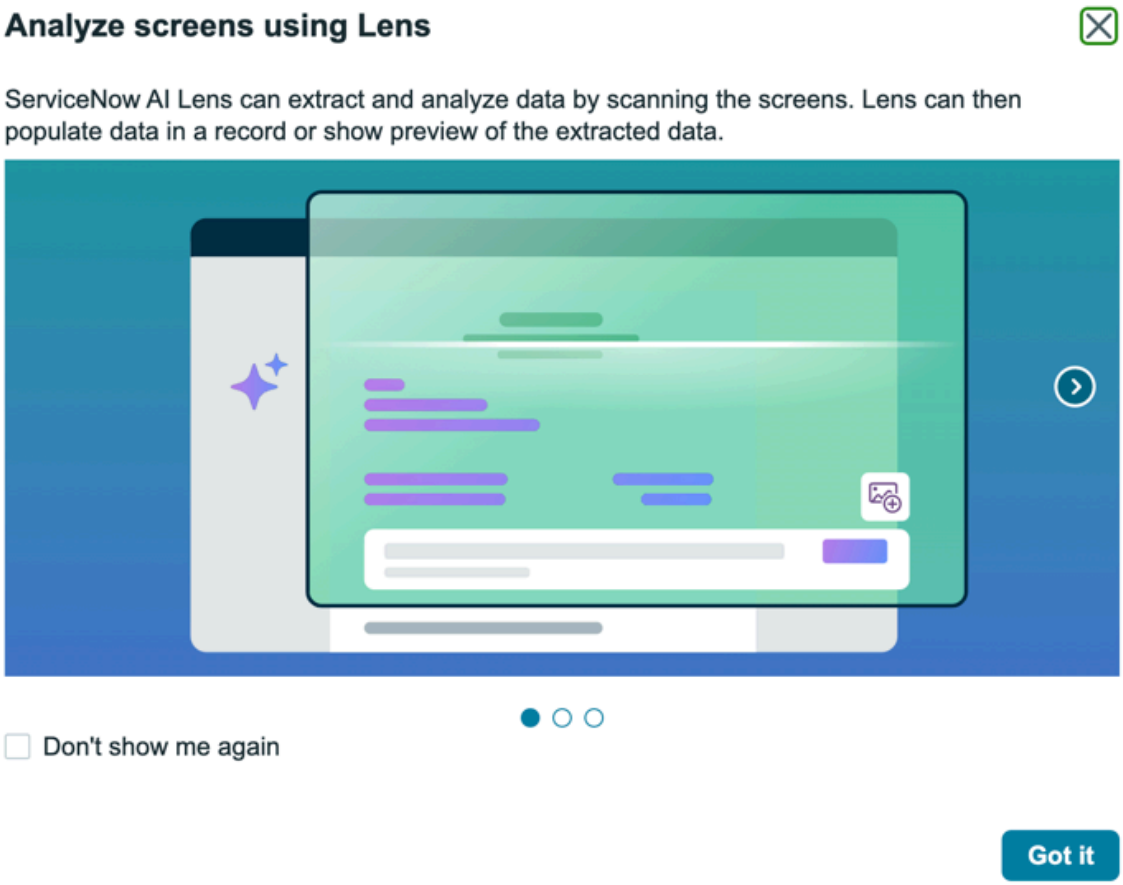
Role required: lens_user

Procedure

1. From your system, launch the ServiceNow AI Lens desktop application.
2. On the login page, in the **Instance URL** field, enter the ServiceNow instance URL.

   Example
   For example, `https://<instance name>.service-now.com`.
3. Select **Proceed**.
4. Log in to your ServiceNow account by entering your username and password.
   Your account must have the lens_user role.
5. **Optional:** On the onboarding journey widget, complete the onboarding and select **Got it**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



If you launch the ServiceNow AI Lens for the first time, the onboarding journey widget appears. You can select **Don't show me again** to hide the widget the next time you launch ServiceNow AI Lens.

6. Perform any one of the following methods to capture screenshots or upload files and analyze the data.

Method	Steps
Capture a single screenshot of a document and analyze the data.	**a.** Select **Capture**. The ServiceNow AI Lens scanner window is launched. **Note:** Use the scanner window to scan the data of the screenshot so you can submit it for analysis. **b.** Open a document that you want to scan. Example of a document: image, scanned handwritten note, Excel sheet, web page, or any document that gives visual data. **c.** Place the ServiceNow AI Lens scanner win dow over the document. **d.** (Optional step). To provide instructions to analyze the data in the screenshot in a spe

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Method	Steps
	cific way, select the Edit icon ( ), and en ter the instructions. The default character limit is 500. Users with the admin role can increase this lim it to up to 5000 characters by navigating to the `sn_lens_user_prompt_max_length` system property. **e.** To analyze the data in the screenshot that you captured, select **Analyze**. **f.** The preview of the analysis is displayed on the ServiceNow AI Lens preview window.
Capture multiple screenshots from one or more documents.	**a.** Select **Capture**. **b.** Open one or more documents from which you want to capture the screenshots. Example of a document: image, scanned handwritten note, Excel sheet, web page, or any document that gives visual data. **c.** Select the Multi-capture icon ( ), and then place the scanner window over the document that you want to scan. You can resize the scanner window by drag ging its borders. **d.** (Optional step). To provide instructions to analyze the data in the screenshots in a specific way, select the Instructions icon ( ), and enter the instructions. The default character limit is 500. Users with the admin role can increase this lim it to up to 5000 characters by navigating to the `sn_lens_user_prompt_max_length` system property. **e.** Select the Capture icon () The first screenshot is captured. **f.** Place the ServiceNow AI Lens scanner win dow over the page of another or the same document and then select the Capture icon (). The second screenshot is captured. Repeat the step to capture more screen shots, if required.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Method	Steps
	g. (Optional step) To remove a screenshot that you had captured, select the Delete icon. **Tip:** ▪ Repeat the step as many times to re move as many screenshots. ▪ The last screenshot that you had captured is removed at the first click. ▪ You can capture one or more new screenshots after removing the screenshots. **h.** To complete the capture, select the Done icon (). The ServiceNow AI Lens preview window dis plays the screenshots that you've captured.

7. Select **Analyze**.
   ServiceNow AI Lens displays the preview of the response
8. Select **Submit** to trigger post processing.
   The **Submit** button appears on the Preview window if the post-processing is enabled in the related ServiceNow AI Lens action.
9. **Optional:** Copy the previewed data by selecting the Copy icon .
10. End the current session by selecting **Start new session**.

Related topics

Supporting information for ServiceNow AI Lens

ServiceNow AI Lens limitations

Create a record in an instance by using ServiceNow AI Lens

Update a record in an instance by using ServiceNow AI Lens

Extract and analyze data using ServiceNow AI Lens from Virtual Agent

Extract and analyze data from an image using ServiceNow AI Lens from a Virtual Agent conversation on a mobile device or a portal. ServiceNow AI Lens can scan an image and then gather insights or provide recommendations as per your instructions.

Before you begin

To access the ServiceNow AI Lens functionality, perform the following steps:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Install ServiceNow AI Lens on your ServiceNow instance. For more information, see Install the ServiceNow Lens in the ServiceNow instance.
- Turn on the ServiceNow AI Lens skill to add the generative AI capability. For more information, see Activate the ServiceNow AI Lens skill.
- Enable ServiceNow AI Lens topic for Now Assist in Virtual Agent. For more information, see Enabling ServiceNow AI Lens in Virtual Agent.

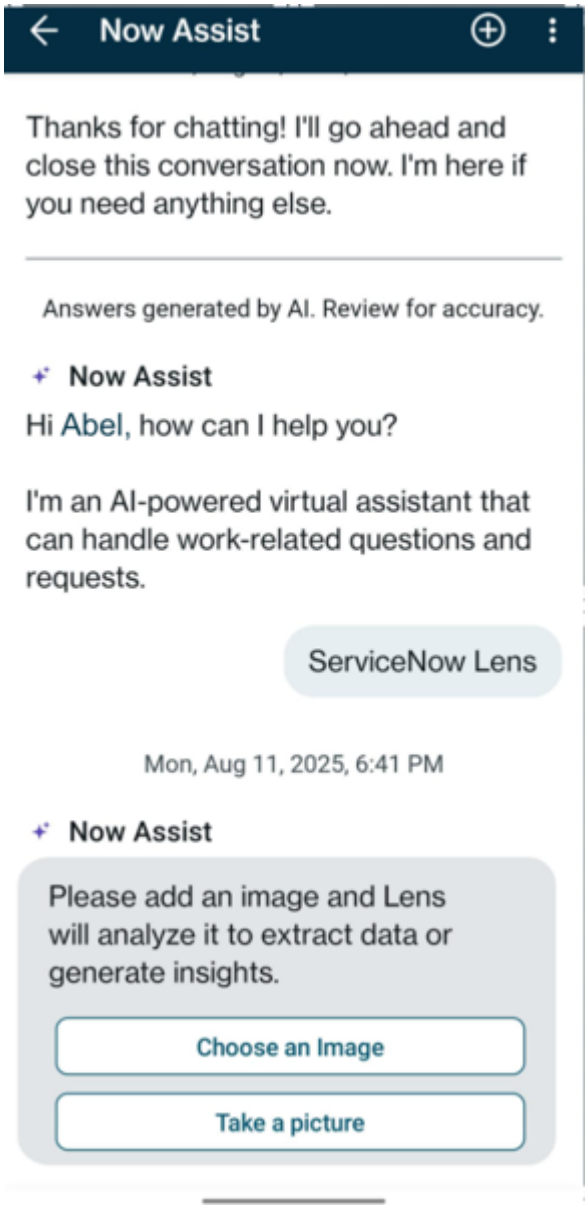
Don't scan any personally identifiable information, such as medical reports, financial reports, or other sensitive data, when using ServiceNow AI Lens as you don't want to expose the large language model (LLM) to any sensitive information.

Role required: lens_user

Procedure

1. Navigate to the home screen on your Now mobile app or a portal.
2. Select the Now Assist icon .
3. Select **Show all topics** and select ServiceNow AI Lens topic or type ServiceNow AI Lens. You would see a welcome message from ServiceNow AI Lens topic.
4. Select one of the following options:
   - Choose an image
   - Take a picture

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



5. Upload an image or take a picture.
6. Provide instructions on how to process the provided image.
   ServiceNow AI Lens can provide insights, recommendations, resolution steps, and more.
7. **Optional:** Ask follow-up questions.
8. When you are done, type Exit to leave the conversation.

ServiceNow AI Lens reference

Reference topics provide additional information about the roles and tables that are installed with the ServiceNow AI Lens application.

Components installed with ServiceNow AI Lens

Several types of components are installed with the activation of the ServiceNow AI Lens plugin, including user roles.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Roles installed

Role title [name]	Description	Contains roles
Lens user [lens_user]	Enables you to use ServiceNow AI Lens.	sn_nowassist_admin.user
Lens admin [lens_admin]	Enables you to configure lens actions.	lens_user

Field types supported

Learn about the field types that are currently supported with the ServiceNow AI Lens application.

The following field types are currently supported with ServiceNow AI Lens:

- name
- due_date
- float
- glide_time
- glide_utc_time
- html
- journal
- journals_input
- longint
- password
- password2
- simple_name_values
- table_name
- translated_html
- translated_text
- boolean
- choice
- decimal
- email
- glide_date
- glide_duration
- insert_timestamp
- integer
- ip_addr
- percent_complete
- phone_number_e164
- price
- script
- script_plain

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- string
- string_full_utf8
- url

For more information about these fields, see Field types reference.

Formatting for date, time, and table name fields

Field	Formatting description
glide_date	• Priority 1 - For a glide_date field type on any record, ServiceNow AI Lens uses the format specified in the date_format field on a user's record of a user running the Lens session. • Priority 2 - If no date_format is specified during user record creation or later, Lens uses the format specified in the system property `glide.sys.date_format`. • Priority 3 - If the system property is not available, the default format `YYYY-MM-DD` is used.
glide_time	• Priority 1 - For a glide_time field type on any record, ServiceNow AI Lens uses the format specified in the time_format field on a user's record of a user running the Lens session. • Priority 2 - If no time_format is specified during user record creation or later, Lens uses the format specified in the system property `glide.sys.time_format`. • Priority 3 - If the system property is not available, the default format `HH:mm:ss` is used.
glide_utc_time	The value is populated in the `HH:mm:ss` format.
glide_duration	Duration is populated in the `D HH:mm:ss` format. For example, the duration of 1 day, 2 hours, 30 minutes is populated as 1 02:30:00.
due_date	The value is populated using the formats of glide_date and glide_time, concatenated with a space.
table_name	The value is only populated when the exact internal name of table is identified.

ServiceNow AI Lens limitations

Be aware of a few limitations when you're using the ServiceNow AI Lens application.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Limitations with ServiceNow AI Lens

Limit	Description
Maximum number of images	Maximum number of images that can be uploaded to ServiceNow AI Lens is 10.
Maximum character limit	In the instructions field for the analysis, you can enter up to 500 characters.
Themes	When you change the theme on the ServiceNow instance, the theme of the ServiceNow AI Lens is automatically updated. But, if you have already started a ServiceNow AI Lens session and then you change the theme on the instance, the theme doesn't change automatically until you launch a new session. This behavior is standard for the ServiceNow AI Platform.
Language	When you change the language on the ServiceNow instance, the language of the ServiceNow AI Lens is automatically changed. But, if you have already started a ServiceNow AI Lens session and then you change the language on the instance, the language doesn't change automatically until you launch a new session. This behavior is standard for the ServiceNow AI Platform. When using ServiceNow AI Lens to scan non-English documents—such as invoices in Chinese— you may notice inconsistencies in translation, formatting, or data extraction. These discrepancies result from limitations in how the underlying LLM processes different locales.
Auto-fill	Pre-populated form fields, such as record number, can be overwritten by ServiceNow AI Lens. You must verify the auto-filled form fields before saving the form.
Visibility of the **Create with Lens** button	As a user with the lens_user role, if you don’t have the Write access for the first record of any list, the **Create with Lens** button doesn't appear in the list view of the table. You can sort the records to show a different record on top and reload the list to view the **Create with Lens** button.
Hallucinations	If the large language model (LLM) returns text that doesn’t exist in the source artifacts, the LLM might inject out-of-context information. You must verify that all information aligns with

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Limitations with ServiceNow AI Lens (continued)

Limit	Description
	the source material to ensure that there are no fabricated facts or unsupported conclusions.

Script include - AILensActionService

Use the *AILensActionService* script include together with Lens actions to leverage ServiceNow AI Lens as a service for extracting information from the provided images and getting answers to your questions.

This script include is part of the ServiceNow AI Lens (sn_ai_lens) store application and is located within the `sn_app_lens_core` scope.

This script include provides methods that enable the following:

- Calls Lens as a back-end service
- Analyzes and comprehends data from provided images
- Gets response from Now Assist as per provided directions
- Does not require ServiceNow AI Lens desktop app

AILensActionService - AILensActionService()

Creates an AILensActionService instance.

Parameters

Name	Type	Description
None		

The following example shows how to initialize *AILensActionService*.

```
var lensService = new sn_app_lens_core. AILensActionService()
```

AILensActionService - invokeLens(String lensActionId, String[] attachmentIds, String userPrompt, Object[] imageArr, Object inputJSON)

Invokes ServiceNow AI Lens as a service.

Parameters

Name	Type	Description
lensActionId	String	Sys_id of the Lens Actions record created for your use case, or you can select the out-of-the-box option that fits your requirements. Example: 842bfc8e37066210b97528c734924baf This parameter is mandatory.
attachmentIds	String[]	Array of sys_ids for existing image attachments. Example: `["0067e66f93f022108319f9ed1dba108b", "0000e8a42c9a7110f877137af4eab4b5"]`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameters (continued)

Name	Type	Description
		You must pass either `attachmentIds` or `imageArr` parameter.
userPrompt	String	An instruction or question for Now Assist to answer after analyzing the contents of the attachments. Example: `Analyze this production issue and create an incident ticket`
imageArr	Object[]	Array of objects containing name of the screenshot and base64 encoded image data. Example: `[ { name: "screenshot1.png", data: "iVBORw0KGgoAAAANSUhEUgAAAAEAAAABCAYAAAAfFcSJAAAADUlEQVR42mNkYPhfDwAChwGA60e6kgAAAABJRU5ErkJggg==" }, { name: "screenshot2.png", data: "iVBORw0KGgoAAAANSUhEUgAAAAEAAAABCAYAAAAfFcSJAAAADUlEQVR42mNk+M9Qz0AEYAJMgkU1f5kAAAAASUVORK5CYII=" } ];` You must pass either `attachmentIds` or `imageArr` parameter.
inputJSON	Object	Additional JSON input parameters that you want to pass in the pre-processing script of the Lens action. Example: `{ "type" : "object", "properties" : { "short_description" : { "type" : "string", "label" : "Short description" }, "description" : { "type" : "string", "label" : "Description" }, },`

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Parameters (continued)

Name	Type	Description
		`"required" : [ "short_description", "comments" ], }`

Returns

Type	Description
<object>	Returned success object `{ "status": "success", "lensResponse": "{\"short_description\":\"Service Degradation Error in Order Processing System\",\"description\":\"The Order Processing API v2.1 encountered a service degradation issue in the Production environment.\" }" }`
error	Returned error object `{ "status": "error", "error": { "errorType": "Execution Error", "message": "Detailed error message here" } }`

This example shows how to call Lens service from a script block.

```
var lensActionId = "cd6570cdf36a2210b9751f09f6968c42";
var attachmentIds = ["3fe930093b626210aba1fadc73e45a38",
 "0000e8a42c9a7110f877137af4eab4b5"];
var userPrompt = "Analyze this production issue and create an
 incident ticket";
var imageArr = [
    {
        name: "screenshot1.png",
        data:
 "iVBORw0KGgoAAAANSUhEUgAAAAEAAAABCAYAAAAfFcSJAAAADUlEQVR42mNkYP
hfDwAChwGA60e6kgAAAABJRU5ErkJggg=="
    },
    {
        name: "screenshot2.png",
        data:
 "iVBORw0KGgoAAAANSUhEUgAAAAEAAAABCAYAAAAfFcSJAAAADUlEQVR42mNk+M
9QzOAEYAJMgkU1f5kAAAAASUVORK5CYII="
    }
```

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
];
var inputJSON = {
      "type" : "object",
      "properties" : {
        "short_description" : {
          "type" : "string",
          "label" : "Short description"
        },
        "description" : {
          "type" : "string",
          "label" : "Description"
        },
      },
      "required" : [ "short_description", "comments" ],
   }
 // Call the method
var result = new sn_app_lens_core.
 AILensActionService().invokeLens(lensActionId, attachmentIds,
 userPrompt, imageArr, inputJSON);

// Handle the response
if (result.status === 'success') {
    var response = JSON.parse(result.lensResponse);
    gs.info("AI Lens Analysis Complete:");
    gs.info("Title:", response.short_description);
    gs.info("Description:", response.description);
} else {
    gs.error("Error occurred:", result.error.message);
}
```

Conversation Insights

The ServiceNow® Conversation Insights application is designed to deliver Inferred customer satisfaction (CSAT) scores, with underlying factors that can help improve CSAT actionability.

Get started

Get started with Conversation Insights to integrate Inferred CSAT scores and factors into your analytics dashboard.

Explore	Install	Reference
Learn more about Conversation Insights	Install Conversation Insights	Additional information about Conversation Insights

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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 Conversation Insights

Learn how Conversation Insights can help you to augment conversation insights with AI-based Inferred customer satisfaction (CSAT) scores and factors.

Conversation Insights overview

Conversation Insights is designed to deliver Inferred CSAT scores and explanatory factors for conversations in Virtual Agent and agentic workflows. It leverages AI to analyze conversations in real time, and provides actionable insights that help improve Virtual Agent and live agent interactions, and agentic workflows.

Inferred CSAT is a numerical score from 1 (least satisfied) to 5 (most satisfied). It's predicted entirely from conversation transcripts in real time without any input from the user. In addition to the CSAT score, the model also predicts CSAT factors that contributed to the CSAT score. The following CSAT factors are associated with the Inferred CSAT score.

- Resolution: Indicates whether the Virtual Agent or AI agent successfully resolved the user's issue without human intervention.
- Confusion: Indicates how often the Virtual Agent or AI agent misunderstood or failed to interpret the user's intent.
- Effort: Indicates the number of user turns or interactions required to reach a resolution.
- Empathy: Indicates how well the Virtual Agent or AI agent acknowledged and responded to the user's emotional tone.
- Next Steps: Captures whether the Virtual Agent or AI agent clearly communicated what the user should do next.
- Frustration: Flags signs of user dissatisfaction or repeated failed attempts during the interaction.
- Transfers and Escalations: Tracks how often the Virtual Agent or AI agent handed off the conversation to a human agent or another system.

Inferred CSAT scores and factors are calculated for each conversation. Conversational analytics applications can leverage the scores written to the Conversation Insights [sn_aci_insights] table to create custom dashboards and workflows. The AI Agent Analytics dashboard includes visualizations with Inferred CSAT scores and factors by default.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The data retention period for the Conversation Insights [sn_aci_insights] table is two years. For more information on creating custom dashboards, see Create a dashboard with the in-line editor and Select a table data source for a data visualization.

Conversation Insights are also supported by Voice Agents, Now Assist Panel and Now Assist Virtual Agent.

To view the conversation insights for Voice Agents, go to **Assistant Designer** > **Analytics** > **Sentiment** dashboard.

To view the conversation insights for Now Assist Panel or Now Assist Virtual Agent, go to **Platform Analytics** > **Dashboards** > **AI Engagement Analytics**.

Inferred CSAT segmented by intent

Use the additional insights introduced to report on Inferred CSAT segmented by intent. It gives better insights into why customers contact the brand and how well their needs are met. This helps companies quickly spot and fix issues with their top customer intents.

The score for Inferred CSAT segmented by intent is automatically generated when customers upgrade the Conversation Insights app. It requires a minimum sample size of 500 conversations to start and samples up to 2000 records every 15 minutes.

These metrics are available in the AI Agent Analytics dashboard, under the Insights tab, where you can find two new charts.

Conversation Insights workflow

The Conversation Insights workflow illustrates how each interaction, whether handled by Virtual Agent or an AI agent, is transformed into actionable insights. You can feed the insights directly into the dashboards for analysis and decision making. The Conversation Insights workflow shows the journey from conversations to insights on dashboards.

Conversation Insights workflow

1. Conversation sources
   - Agentic AI chats
   - Virtual Agent chats
2. Data aggregation

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Agentic AI and Virtual Agent interactions are captured in the Conversation table [sys_cs_conversation].
- The conversation transcript, including user query, agent response, timestamps, and metadata such as the session ID and channel type are also stored in the Conversation table for processing.

3. Insight generation
   - The model analyzes the conversation transcript.
   - Inferred CSAT scores are generated for CSAT factors such as Empathy, Resolution, Frustration, and so on.
4. Insights storage
   - Inferred CSAT scores and factors are stored in the Conversation Insights table [sn_aci_insights].
   - The Conversation Insights table acts as a structured repository for the extracted insights.
5. Dashboards
   - The insights are made available to create adhoc dashboards and workflows.
   - You can explore trends, performance metrics, and target improvements based on the Inferred CSAT scores.

Conversation Insights benefits

Problem	Solution
Traditional surveys often reflect extreme opinions and low response rates.	Inferred CSAT helps solve this problem by using AI to estimate CSAT score for conversations in real time, based on full conversation transcript. This CSAT score can help to eliminate bias and reduce the need for reliance on explicit survey feedback.
Post-interaction feedback delays insights resulting in lagging indicators.	CSAT scores are generated immediately after the interaction, enabling faster detection of issues and trends.
Lack of actionable insight behind CSAT scores.	CSAT factors like Resolution, Empathy, Effort, and so on explain user satisfaction or dissatisfaction, helping you to target improvements.

What to explore next

To learn more about configuring and using Conversation Insights, see:

- Install Conversation Insights
- Conversation Insights reference

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Install Conversation Insights

If you have the admin role, you can install the Conversation Insights application (sn_aci).The application includes demo data and installs related ServiceNow® Store applications and plugins, if they aren't already installed.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.
- Inferred CSAT is calculated by an internal service that communicates with the instance using mutual Transport Layer Security (mTLS). To enable this, the instance must be configured to support Application Delivery Controller-to-Application (ADC-to-APP) mTLS. This setup ensures strong authentication and encrypted data exchange between internal services and the instance, requiring both client and server to present and validate certificates during the TLS handshake. For more information on mTLS configuration, refer to steps 1 and 2 in the Prerequisites section of KB0993615.

Role required: admin

About this task

The following items are installed with Conversation Insights:

- Plugins
- Store applications
- Roles
- Scheduled jobs
- Conversation Insights [sn_aci_insights] table

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Conversation Insights application (sn_aci) using the filter criteria and search bar.

   You can search for the application by its name or ID. If you can't find the application, you might have to request it from the ServiceNow Store.

   A list of the versions available to you're displayed.
3. Select a version from the list and select **Install**.

   In the Review Installation Details dialog box, any dependencies installed with your application are listed.
4. If you're 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.
   Demo data are the sample records that describe application features for common use cases. Load the demo data when you first install the application on a development or test instance.
6. Select **Install**.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Conversation Insights reference

The following topics provide additional information about the features and properties installed with Conversation Insights.

Components installed with Conversation Insights

Installation of the Conversation Insights application also installs the Conversation Insights [sn_aci_insights] table.

Tables installed with Conversation Insights

The installation of Conversation Insights installs the Conversation Insights [sn_aci_insights] table. The data retention period for the Conversation Insights table is two years.

Conversation Insights table

Column	Description
Human Agent CSAT	Estimated Inferred customer satisfaction (CSAT) score based on the human agent's responses. The value range is 1–5 with 1 meaning that the user is extremely dissatisfied and 5 meaning that the user is extremely satisfied with the interaction.
Effort Score	A score of time and energy that the user had to put in during the interaction. The values are on a 3-point scale of low, medium, and high.
AI agent CSAT	Estimated CSAT score based on the AI agent's responses. The value range is 1–5 with 1 meaning that the user is extremely dissatisfied and 5 meaning that the user is extremely satisfied with the interaction.
Frustration	Estimated CSAT score that indicates whether the user expressed frustration during the interaction. The values are Yes or No.
Empathy	Estimated CSAT score that indicates empathy shown toward the user during the interaction. The values are on a 3-point scale of low, medium, and high.
Session CSAT	Estimated CSAT score calculated based on the overall session interactions of a user. The value range is 1–5 with 1 meaning that the user is extremely dissatisfied and 5 meaning that the user is extremely satisfied with the interaction.
Next Steps	Estimated CSAT score that indicates whether the agent recapped the outcome and provided clear instructions on what the user should do next. The values are on a 3-point scale of low, medium, and high.
CSAT processed	Boolean flag indicating whether CSAT data has been processed for the interaction record.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Conversation Insights table (continued)

Column	Description
Transfer and Escalation	Estimated CSAT score that indicates whether the conversation was transferred or escalated to a human agent. The values are Yes or No.
Closed At	Timestamp showing when the conversation was closed.
Conversation	Refers to the associated conversation record in the Conversation [sys_cs_conversation] table.
Resolved	Estimated CSAT score that indicates whether the issue was resolved. The values are Yes or No.
Confusion	Estimated CSAT score that indicates whether the user expressed confusion or the agent didn't understand the user's intent. The values are Yes or No.

Conversation Improvement Themes

The Conversation Improvement Themes application helps to transform Conversation Evaluators evaluations that are part of AI Control Tower into long-term performance insights by analyzing conversation quality data over time.

Get started

Explore  Learn more about Conversation Improvement Themes	Activate  Activate Conversation Improvement Themes	Use  Use Conversation Improvement Themes
Reference  Additional information about Conversation Improvement Themes		

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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.

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 Conversation Improvement Themes

The Conversation Improvement Themes application helps to transform conversation evaluations into long-term performance insights.

Conversation Improvement Themes overview

The Conversation Improvement Themes analyzes conversation quality data over time. The approach focuses on identifying recurring patterns linked to low conversation quality (and high) and categorizing user requests into actionable themes using metadata-driven classification powered by large language models (LLMs).

Note: Conversation Improvement Themes can use OpenAI, Gemini, or Claude. It's observed in certain instances that Gemini and Claude could consume a higher number of assists, compared to the OpenAI model, which is observed to have optimal use of assists.

How themes are created

To create themes, metadata like name, short description, and so on, from different content types like knowledge base articles, catalog items, virtual agent topics, and AI agents are analyzed. Theme generation is performed using an LLM-based categorization process that relies only on

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

metadata rather than full article content, which helps limit token usage, reducing processing overhead. For each created theme, definitions are created using the tagged content from the LLM-based classification process, which helps in tagging primary requests from conversations.

Approach:

1. Filter by User Satisfaction Score. For example, select conversations where evaluator results are associated with low user satisfaction score (<=1.5) or High Score >= 4.5.
2. Extract and standardize Primary requests:
   a. Parse the user's Primary request, removing noise from different writing styles or phrasing.
   b. Apply standardization so similar intents, for example, reset my password versus can't log in are grouped consistently.
3. Frequency analysis:
   a. Link each standardized request to a theme.
   b. Count how often each theme appears when the virtual agent underperforms or performs well.
   c. Highlight themes where the virtual agent handles requests poorly or well on the screen.

Activating Conversation Improvement Themes

Activate the Conversation Improvement Themes application to analyze conversation quality.

Before you begin

Role required: admin

Conversation evaluations have to be enabled in AI Control Tower. For more information, see .

Procedure

1. Navigate to **All** > **Now Assist Admin Console** > **Now Assist Skills**.
2. Within the Platform category, select the feature **Conversation Evaluator** and verify that all the skills shipped with the Conversation Evaluator application are activated.
3. After that, activate the skills shipped with this application.
   - Theme Extraction
   - Themes Update Definition
   - Themes Primary Request Tagging
4. Navigate to **All** > **System Definition** > **Scheduled Jobs**.

   Apply the filter: `Application = Conversation Improvement themes`. Activate all scheduled jobs.

   - Create thematic executions
   - Evaluations for themes [Daily job]
   - Evaluations for themes [Historic job]
   - Themes - Create Daily Primary Request Executions
   - Update Theme Effective and Ineffective Scores
5. For access to the application, assign the following roles to the users:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- sn_na_thematic.theme_admin
- sn_na_conv_eval.eval_admin

Note:

Only users who have access to the roles mentioned previously can access the data on the Conversation Improvements dashboard and the tables. But they won't be able to see live agent invocations.

User must have the following roles to access all the data on the themes dashboard:

- sn_na_thematic.theme_admin
- sn_na_conv_eval.eval_admin
- Interaction_admin

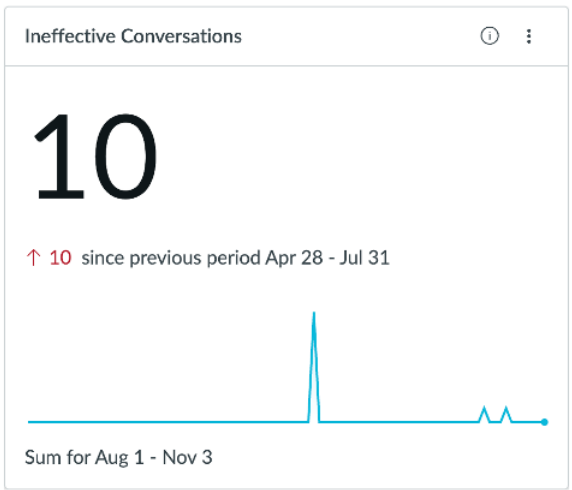
Using Conversation Improvement Themes

View performance insights using the widgets in Conversation Improvement Themes.

To access the application, navigate to **All** > **Conversation Improvement Themes** > **Overview**.

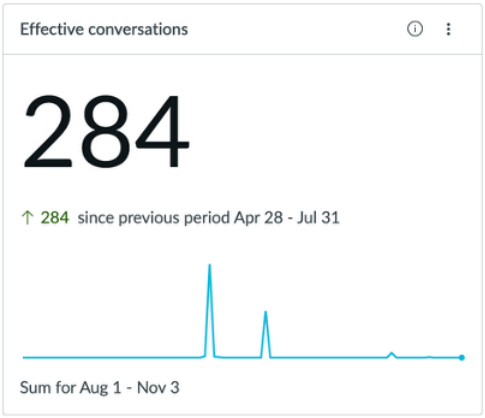
The Overview page has the following widgets:

- Ineffective Conversations: The number of conversations with Auto evaluated user satisfaction scores of less than or equal to 1.5.

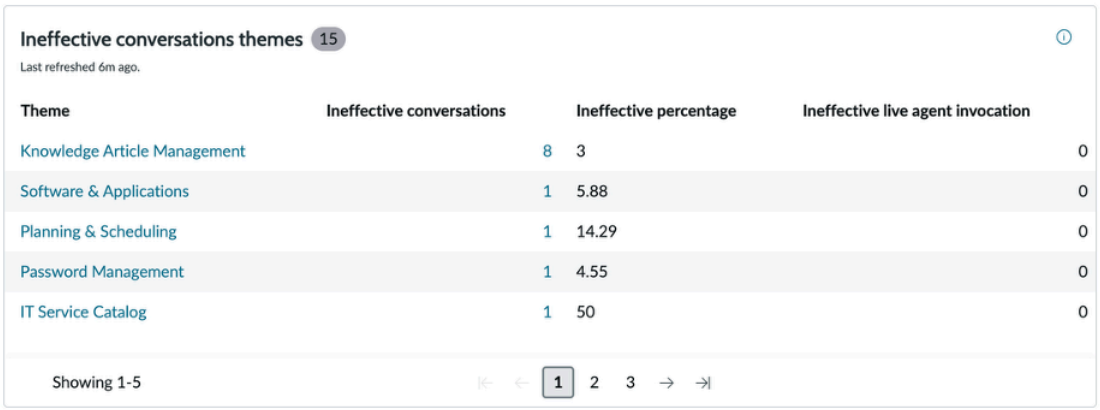


- Effective Conversations: The number of conversations with Auto evaluated user satisfaction scores of greater than or equal to 4.5.

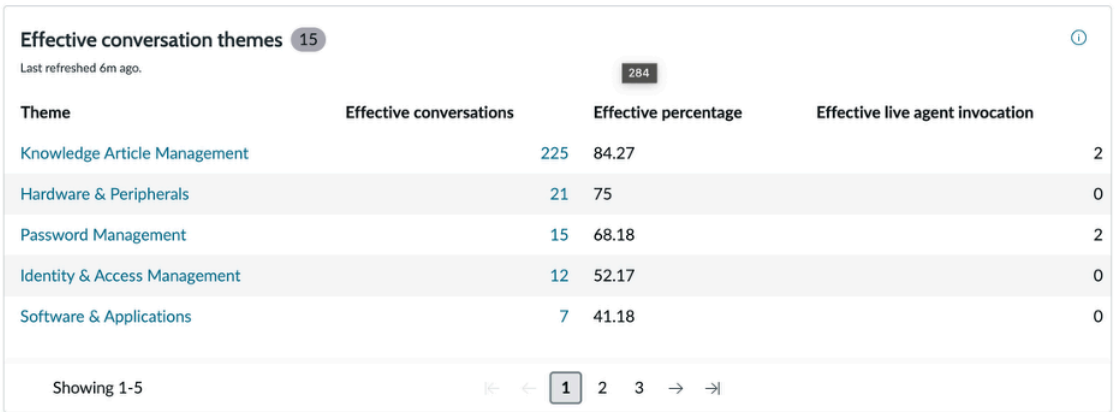
© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



- Ineffective conversation themes: The themes that have ineffective conversations.



- Effective conversation themes: The themes that have effective conversations.



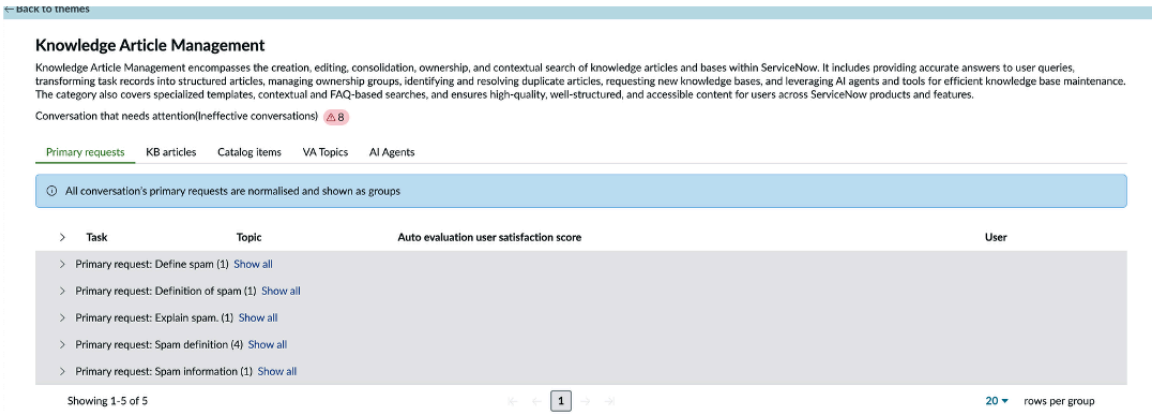
When you select any of themes under **Ineffective conversations themes** or **Effective conversations themes**, you're redirected to the theme.

For example, if you select Knowledge Article Management, it takes you to the page showing the theme name with its description and then shows the primary requests, knowledge articles, catalog items, VA topics, and AI Agents tagged to the theme. When you select on any of the tagged Primary requests, KB articles, Catalog Items, VA Topics, or AI Agents, you're redirected to these items.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note: You can only see the sections to which you have access.

For example, you need access to AI Control Tower to access primary requests, Knowledge Base to access KB articles, Service Catalog to access catalog items, VA administrator to access VA topics, AI Agent Studio to access AI agents.



Reference for Conversation Improvement Themes

The following topics provide additional information about the features and properties installed with Conversation Improvement Themes.

Components installed with Conversation Improvement Themes

Several types of components are part of Conversation Improvement Themes, including scheduled jobs, tables, system properties, and flows.

Roles added

Role	Description
sn_na_thematic.theme_admin	Primary role to access dashboard and tables installed with the application.

Scheduled jobs installed

Scheduled job	Description
Create thematic executions	This scheduled job runs once a month and creates Themes Execution Staging Records that create themes and their description and tags knowledge articles, catalog items, VA topics, AI Agents, and conversations to these themes.
Themes - Create Daily Primary Request Executions	Runs daily and creates Themes Execution Staging Records that tags conversations to the existing themes, and if necessary, creates new themes
Update Theme Effective and Ineffective Scores	Runs daily and populates number of effective conversations and number of ineffective conversations for every theme.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tables installed

Label	Name
Themes Configurations [sn_na_thematic_themes_configurations]	Contains various configurations for Conversation Improvement Themes. Configuration Name - Last scan date: Contains the last scanned date of the theme creation. After the first run, theme extraction and theme description update will only run on content updated or created after the last run date.
Themes [sn_na_thematic_themes]	Contains the themes created along with their descriptions.
Themes associations [sn_na_thematic_associates]	Contains the relationship between themes and different content types tagged to that theme.
Primary Requests [sn_na_thematic_primary_requests]	Contains all the primary requests extracted from evaluated chats.

System properties installed

Property	Description
sn_na_thematic.dataChunkSize	The default chunk size, used for Azure Open AI provider and as a backup. The default chunk size is 100. If you see Themes Execution staging record is either in a failed state or stuck on processing for a long time, reduce the chunk size and change the state of the Execution staging record to draft and then again to processing.
sn_na_thematic.fallbackDataChunkSize	If any of the theme skills are not part of the Azure Open AI provider, the fallback size is used to create execution staging records. If you see Themes Execution staging record is either in a failed state or stuck on processing for a long time, reduce the chunk size and change the state of the Execution staging record to draft and then again to processing. The default chunk size is 25. But it should be set differently for different LLMS. • Gemini: 25 • Claude: 10 • Now LLM: 1

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Business rules installed

Name	Description
Theme tagging	An asynchronous business rule that runs the state of the Themes Execution Staging record changes to processing. It invokes the Theme Tagging subflow.
Themes Update next record to processing	Runs state of the Themes Execution Staging record changes to Classification Complete, Complete or Failed.

Subflows installed

Flow	Description
Theme Tagging	Executes when the state of the Themes Execution Staging record changes to processing. The subflow creates themes, updates their description, and then tags knowledge articles, catalog items, VA Topics, AI agents, and conversations with the themes and moves the state of the Themes Execution staging record to complete.
Primary Theme Tagging	The subflow is called from Theme Tagging subflow. It tags conversations to the existing themes and if necessary, creates new themes.

Script includes installed

Script includes	Description
themesUtil	Primary utility function for Conversation Evaluator.

Troubleshooting for Conversation Improvement Themes

Troubleshooting steps for errors that might occur in Conversation Improvement Themes.

Scenario/Symptom: The dashboard is not showing any data, one week after the install.

Associated Error Codes/Warning: N/A.

Steps to Resolve:

Make sure all the skills and scheduled jobs are active by following the steps given in Activating Conversation Improvement Themes.

Scenario/Symptom: Execution staging records are stuck on Processing/Failed.

Associated Error Codes/Warning: Flow-related logs in syslog.

Steps to Resolve:

Increase flow logging (preferably in dev).

Environment guidance:

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Dev/Test: Set verbose logging to All configurations and runtime value.
- Prod: Use Runtime states and duration only. Don't enable All configurations and runtime value in production.

Steps:

1. Navigate to **Process Automation** > **Flow Administration** > **Properties**.
2. Set Level of reporting data generated by the flow engine to the appropriate level (see the guidance above).
3. Create a conversation through the Portal and end the conversation.
4. Navigate to **Process Automation** > **Flow Designer** > **Theme Tagging(subflow) or Primary Request Tagging(subflow)**.
5. Review flow executions. Identify any actions marked Failed or Error and address the cause.
6. If all actions are Complete and there is no other path indicating failure handling, open a Support ticket.

Scenario/Symptom: Execution staging records Failed when default model is Gemini/Claude/Now LLM.

Associated Error Codes/Warning: Skill Execution timed out.

Steps to Resolve:

This is a known issue with Gemini and Claude as they are slow thinking models, responses may be slow and sometimes can get timed out. To resolve this, reduce the chunk size in system properties by performing the following steps:

1. Navigate to **All** > **System Properties** (sys_properties.list) and search for **sn_na_thematic.dataChunkSize** and reduce it. The ideal chunk size for all the models is:
   - Azure Open AI: 100
   - Gemini: 25
   - Claude: 5
   - Now LLM: 1
2. Move the failed Execution staging record to processing.

Generative AI Logs: If an evaluation is not complete, navigate to sys_generative_ai_log and check for the following definitions to see if there are any errors:

- Theme Extraction
- Themes Update Definition
- Themes Primary Request Tagging

Outbound HTTP:

If no Gen AI logs are created, check the Outbound HTTP logs for any issues exist.

Sys Logs: Check sys logs for any errors related to passing with scope as sn_na_thematic.

Now Assist and Agentic AI glossary

Learn about the terms and concepts used in Now Assist and Now Assist AI agents.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

agents

An agent, in the context of AI, refers to a software entity that can perceive its environment, make decisions, and take actions to achieve specific goals. Agents can be simple or complex, ranging from basic reactive systems to sophisticated entities that learn and adapt over time. They are often used to automate tasks, interact with users, or manage complex processes.

agentic AI

Agentic AI refers to artificial intelligence systems that possess autonomy, enabling them to function independently and make decisions without continuous human oversight. These systems can perceive their environment, reason, learn from experience, and execute tasks to achieve specific goals. In ServiceNow AI Platform, agentic AI is powered by AI agents, agentic workflows, and Voice/Web agents.

agentic conversation

Agentic conversation refers to interactions between humans and AI systems where the AI exhibits autonomous behavior, making decisions and taking actions to achieve specific goals. In an agentic conversation, the AI agent can understand the context, reason, and respond accordingly, often without requiring explicit human direction. You can enable it in Virtual Agent, Now Assist panel, Voice/Web agents, and enhance it through Knowledge Graph, RAG, and Open Prompts.

agentic system

An agentic system is a type of software or AI that perceives its environment, makes decisions that are based on that perception, and takes actions to achieve specific goals, often with minimal human intervention. An agentic system can learn, adapt, and operate independently to solve problems or perform tasks.

agentic workflow

An agentic workflow is a structured sequence of tasks executed by one or more AI agents with minimal human intervention to fulfill a business objective. You can create and manage these workflows in AI Agent Studio by using triggers, tools, and evaluation plans, and deploy across ITSM, HR, and CSM use cases.

AI agents

AI agents are autonomous digital workers that use LLMs, tools, and workflows to complete tasks across ITSM, HR, and CSM on behalf of users. They can reason, plan, and act independently or collaboratively. You can create and manage them using AI Agent Studio and orchestrate them in AI Agent Orchestrator. These agents are embedded in Now Assist panels, Virtual Agent, and custom workflows.

The AI agent is a virtual worker that performs specific tasks leveraging tools from the ServiceNow AI Platform such as workflow, skill, script, knowledge base, etc.

AI compliance

AI compliance refers to the adherence to regulatory requirements, standards, and guidelines governing the development, deployment, and use of artificial intelligence systems. It involves ensuring that AI systems are designed and implemented in a way that meets legal, ethical, and organizational standards. This includes addressing issues

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

related to data privacy, bias, transparency, and accountability. You can manage it through audit trails, model life cycle management, and data masking by using AI Control Tower in ServiceNow AI Platform.

AI Control Tower

AI Control Tower is a centralized governance hub for managing AI models, agents, and compliance across the ServiceNow AI Platform. It provides visibility into model usage, performance, risk and compliance, and is integrated with Now Assist Admin console.

AI governance

AI governance refers to the set of policies, procedures, and standards that guide the development, deployment, and use of AI systems within an organization. It ensures that AI systems are aligned with organizational values, comply with regulatory requirements, and are used in a responsible and transparent manner. In ServiceNow AI Platform, you can manage AI governance through AI Control Tower, AI Trust Guard, and role-based access in Now Assist Admin console.

AI steward

An AI steward within ServiceNow AI Platform is a governance role responsible for overseeing the life cycle of AI assets, including models, data sets, and prompts. You can use AI stewards in AI Control Tower to coordinate governance activities between legal, security, compliance, and innovation teams for the responsible deployment of AI agents, skills, and models across the platform.

Citation

Small, interactable number next to AI-generated content in the Now Assist panel that cites the source of the information.

custom skill

A user-defined AI capability tailored to meet specific business needs. It is created and trained using relevant data and prompts.

dynamic translation

A feature in Now Assist that translates user input to English and then translates the AI-generated responses back to the user's preferred language.

External Content Connectors

External

External Content Connectors is a ServiceNow® Store application that enables AI Search to index and search content from external repositories like Atlassian Confluence Cloud and Microsoft SharePoint Online. It expands search capabilities to include external content while maintaining original access permissions.

generative AI

A technology that uses large language model (LLM) to generate text, summaries, and other content.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Generative AI Controller

A tool to customize the generative AI experience on the ServiceNow AI Platform, including integrating third-party large language model (LLM) licenses.

guardrail

A mechanism that monitors and evaluates user input and AI-generated response to ensure it is safe and appropriate.

Knowledge Graph

The Knowledge Graph organizes structured and unstructured enterprise data into a semantic network that AI agents can query. The ServiceNowKnowledge Graph application enhances the performance of Now AssistVirtual Agent, AI agents, and generative AI skills by leveraging structured and unstructured data from various sources. It provides a customized user experience by creating and managing Knowledge Graph schemas through a user-friendly interface. The application integrates with AI agents and Virtual Agent, offering prebuilt integrations for seamless functionality.

large language model (LLM)

A complex machine-learning model trained on large datasets to generate human-like responses to user input.

multi-turn catalog ordering

A feature in Now Assist that allows users to order catalog items from the service catalog through a conversational interface.

native translation

A feature in Now Assist that uses the LLMs provided by the Now LLM Service to generate responses in the user's preferred language. You can enable and configure it in the Now Assist Admin console.

Now Assist

Now Assist refers to generative AI experiences on the ServiceNow AI Platform. With Now Assist, you can improve productivity and efficiency in your organization with better self-service, faster answers and recommended actions, and empower users to search more effectively.

Now Assist Admin console

A console for administrators to quickly set up, configure, and monitor generative AI features and capabilities offered under Now Assist.

Now Assist Analytics

A dashboard that provides insights into the usage, value, and performance of Now Assist.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist context menu

A generative AI feature that assists agents in creating and editing emails and chat replies, available in the Next Experience for specific workflows.

Now Assist Guardian

A feature that monitors and evaluates content created with generative AI to protect users from harmful, offensive, or inappropriate content. It can log incidents and prevent such content from being displayed.

Now Assist panel

A conversational interface to summarize chat, case, or incident or generate resolution notes so that you can get the context of the information more quickly.

Now Assist skill

A predefined capability within Now Assist that uses generative AI to perform tasks such as generating summaries, resolution notes, and so on.

Now LLM Service

A component of the ServiceNow AI Platform that leverages large language models (LLM) to enhance generative AI capabilities. It supports various Now Assist skills, such as chat summarization, case summarization, and resolution notes generation.

offensive content detection

A guardrail in Now Assist Guardian that logs and blocks offensive, harmful, or inappropriate content.

Orchestrator

The AI agent orchestrator is a specialized agent that plans, coordinates, and manages the execution of tasks across multiple AI agents. It holds short-term memory, tracks goals, and ensures agents collaborate effectively. This component is central to Agentic AI architecture and is configured in AI Agent Studio. The orchestrator conducts the planning and leads a team of AI agents to address a given agentic workflow.

prompt injection detection

A security feature in Now Assist Guardian that identifies and prevents malicious inputs from being used to manipulate AI-generated responses.

Q&A genius result

Q&A Genius Results in Now Assist provide users with concise, actionable answers generated from relevant knowledge articles and other sources. These results are displayed as actionable cards, each including a link to the source knowledge article for further reference.

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Retrieval Augmented Generation

Retrieval Augmented Generation is a generative AI technique that enhances large language model (LLM) outputs by retrieving relevant enterprise content, such as knowledge articles or records and using it as context to generate more accurate and actionable responses. Within ServiceNow, RAG is implemented in products like Now Assist for Search, where it powers Genius Results cards by combining AI Search with generative responses. It is also configurable in AI Agent Studio, allowing admins to define RAG based tools that retrieve external or internal data to support autonomous or supervised AI agent workflows.

Additional resources for AI products and solutions

If you're looking for AI best practices, troubleshooting, or other implementation guidelines, select a feature or resource type to discover ServiceNow resources on other relevant websites.

Resource links

AI features or products	Resource type	Resources
AI Agents	Release Notes	Now Assist AI Agents Release Notes.
AI Agents	Getting started	Getting started with AI Agents. Identifying agentic workflows.
AI Agents	Best practices	Build custom skills Assess the AI agents' and agentic workflows' performance
AI Agents	Glossary	Now Assist and Agentic AI glossary
AI Agents	Training	AI Agents Essentials (2025 Q1 Store Release) AI Agent Readiness AI Essentials Initiating AI Agents Introduction to Agentic AI Now Assist AI Agents learning path
AI Control Tower	Release notes	AI Control Tower release notes
AI Control Tower	Getting started	Overview of Advanced AI and Data Products

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource links (continued)

AI features or products	Resource type	Resources
Document Intelligence	Release notes	Document Intelligence release notes
Document Intelligence	Getting started	Getting started with Document Intelligence Document Intelligence quick start guide (ServiceNow Community)
Document Intelligence	Tips and examples	Document Intelligence examples
Document Intelligence	Troubleshooting	Troubleshooting and best practices for Document Intelligence
Document Intelligence	FAQs	Document Intelligence FAQs
Document Intelligence	Training	Document Intelligence Overview Document Intelligence Fundamentals
Natural Language Understanding and Natural Language Query	Getting started	Getting started with NLU NLU quick start guide (ServiceNow Community)
Natural Language Understanding and Natural Language Query	Best practices	Best practices for NLU
Natural Language Understanding and Natural Language Query	Tips and examples	NLU examples
Natural Language Understanding and Natural Language Query	Troubleshooting	Troubleshooting for NLU
Natural Language Understanding and Natural Language Query	FAQs	NLU FAQs
Natural Language Understanding and Natural Language Query	Training	Natural Language Understanding (NLU) Fundamentals
Now Assist	Release notes	Now Assist release notes Now Assist for Creator release notes

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource links (continued)

AI features or products	Resource type	Resources
		Now Assist in Document Intelligence release notes Now Assist for Enterprise Architecture (EA) release notes Now Assist for Strategic Portfolio Management (SPM) release notes
Now Assist	Getting started	Getting Started with Now Assist (YouTube) Getting started with Now Assist for Code (YouTube) Getting Started with Now Assist for Creator (YouTube) Getting Started with Now Assist for CSM (YouTube) Getting Started with Now Assist for ITSM (YouTube) Getting Started with Now Assist for AI Search (YouTube) Now Assist quick start guide (ServiceNow Community) Now Assist for Code quick start guide (ServiceNow Community) Now Assist panel fundamentals (YouTube)
Now Assist	Best practices	Now Assist best practices (ServiceNow Community) Knowledge article best practices (ServiceNow Community)
Now Assist	Tips and examples	Now Assist conversational catalog request in practice (ServiceNow Community) Now Assist examples

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource links (continued)

AI features or products	Resource type	Resources
Now Assist	FAQs	Generative AI Controller FAQ Now Assist FAQ Now Assist in AI Search FAQ
Now Assist	Training	AI Builder Fundamentals Introduction to Generative AI Now Assist Essentials (November 2024 Store Release) Now Assist in Virtual Agent Implementation (November 2023 Store Release) Now Assist in Virtual Agent Implementation (August 2024 Store Release) Now Assist in Search Implementation (Nov 2023 Store Release) Now Assist in Search Implementation (Feb 2024 Store Release) Now Assist for Creator Implementer Now Assist for Field Service Management (FSM) Implementer Now Assist for Customer Service Management (CSM) Implementer Now Assist for HR Service Delivery (HRSD) Implementer Now Assist for IT Service Management (ITSM) Implementer
Predictive Intelligence	Release notes	Predictive Intelligence release notes
Predictive Intelligence	Getting started	Getting started with Predictive Intelligence Predictive Intelligence quick start guide (ServiceNow Community)
Predictive Intelligence	Best practices	Predictive Intelligence best practices

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource links (continued)

AI features or products	Resource type	Resources
Predictive Intelligence	Tips and examples	Predictive Intelligence examples
Predictive Intelligence	FAQs	Predictive Intelligence FAQs
Predictive Intelligence	Training	Predictive Intelligence Fundamentals
Task Intelligence	Getting started	Getting started with Task Intelligence Task Intelligence quick start guide (ServiceNow Community)
Task Intelligence	Best practices	Troubleshooting and best practices for Task Intelligence
Task Intelligence	FAQs	Task Intelligence FAQs
Task Intelligence	Training	Task Intelligence for ITSM - Quick Start Guide Task Intelligence for CSM - Quick Start Guide AI Academy Get Started with Task Intelligence -use AI to categorize cases AI Academy Task Intelligence for ITSM AI Academy Task Intelligence - Customer Service Management

© 2026 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.