



# Xanadu Enable AI

**Last updated: 06/13/2026**

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

ServiceNow, the ServiceNow logo, 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](http://www.servicenow.com/terms-of-use.html)

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

# Table of Contents

<b>Enable AI experiences.....</b>	<b>5</b>
Now Assist.....	10
Exploring Now Assist.....	13
Configuring Now Assist settings and features.....	120
Analyzing Now Assist performance.....	150
Now LLM Service updates.....	170
Now Assist reference.....	171
Knowledge Graph.....	176
Generative AI Controller.....	204
Now Assist AI agents.....	253
Exploring Now Assist AI agents.....	255
Configuring Now Assist AI agents.....	290
Create an AI agent.....	301
Create an agentic workflow.....	336
Examples of using AI agents.....	355
Now Assist AI agents reference.....	355
Now Assist Data Kit.....	374
Exploring Now Assist Data Kit.....	376
Configuring Now Assist Data Kit.....	377
Using Now Assist Data Kit.....	377
Now Assist Data Kit reference.....	381
Now Assist Skill Kit.....	381
Exploring Now Assist Skill Kit.....	383
Configuring Now Assist Skill Kit.....	393
Using Now Assist Skill Kit.....	397
Now Assist Skill Kit reference.....	410
Now Assist Readiness Evaluation.....	411
Explore.....	413
Configure.....	415
Use.....	417
Natural Language Understanding.....	425
Exploring Natural Language Understanding.....	429
Model management.....	442
Multilingual model management.....	536
Virtual Agent and NLU Workbench integration.....	551
NLU Workbench - Advanced Features.....	553
Natural Language Query.....	597
Exploring Natural Language Query.....	597
Using Natural Language Query.....	599

Configuring NLQ.....	601
Natural Language Query References.....	607
Predictive Intelligence.....	609
Explore Predictive Intelligence.....	610
Install Predictive Intelligence.....	613
Configure Predictive Intelligence.....	615
Creating and training solutions.....	648
Using Predictive Intelligence.....	684
Predictive Intelligence references.....	719
Document Intelligence.....	725
Exploring Document Intelligence.....	726
Configuring Document Intelligence.....	729
Integrating Document Intelligence with other applications.....	749
Using Document Intelligence.....	756
Monitoring Document Intelligence performance.....	777
Document Intelligence references.....	782
Now Assist in Document Intelligence.....	809
Exploring Now Assist in Document Intelligence.....	811
Configuring Now Assist in Document Intelligence.....	811
Using Now Assist in Document Intelligence.....	812
Now Assist in Document Intelligence reference.....	813
Task Intelligence.....	815
Exploring Task Intelligence.....	816
Configure Task Intelligence.....	819
Manage machine learning models with Task Intelligence.....	820
Task Intelligence references.....	825
Additional resources for AI products and solutions.....	827




# 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.

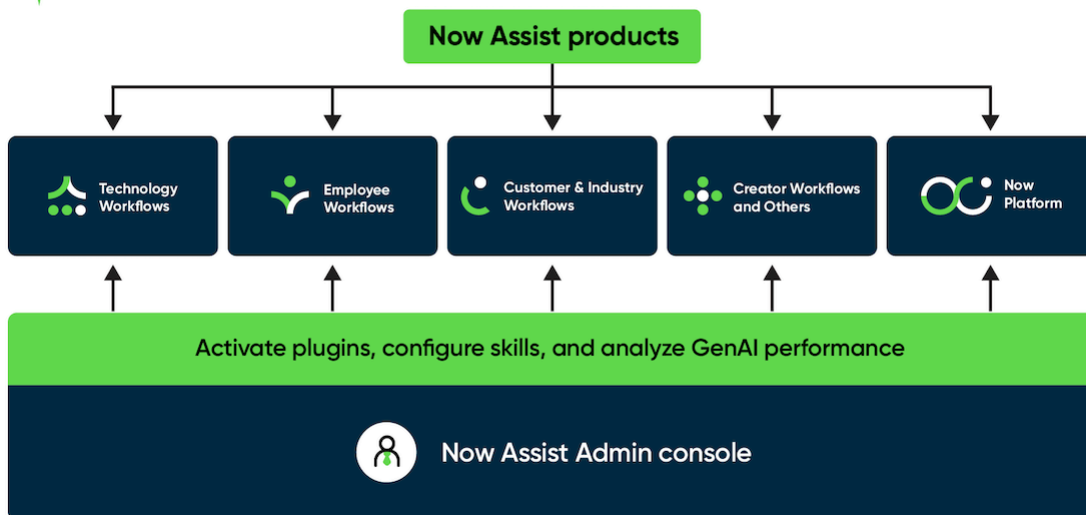
<p><b>Now Assist</b></p>  <p>Use a personalized generative AI assistant and skills to enhance your work on the ServiceNow AI Platform.</p>	<p><b>Now Assist AI agents</b></p>  <p>Build and assemble use cases that drive business outcomes</p>	<p><b>Now Assist Data Kit</b></p>  <p>Add datasets to a data catalog.</p>
<p><b>Now Assist Skill Kit</b></p>  <p>Create custom skills and prompts for your Now Assist experience.</p>	<p><b>AI Governance</b></p>  <p>Monitor, manage, and govern AI assets on the ServiceNow AI Platform.</p>	<p><b>Natural Language Understanding</b></p>  <p>Use Natural Language Understanding (NLU) to help your system work with natural human expression and intent.</p>
<p><b>Natural Language Query</b></p>  <p>Get data from your instance using plain language requests with this service on the ServiceNow AI Platform.</p>	<p><b>Predictive Intelligence</b></p>  <p>Develop ML-based solutions using Predictive Intelligence to create improved work experiences in your instance.</p>	<p><b>Document Intelligence</b></p>  <p>Automate the data extraction process for your documents and speed up your workflow with a powerful AI-based solution.</p>

<p><b>Now Assist in Document Intelligence</b></p>  <p>Use generative AI to analyze and extract information from documents</p>	<p><b>Task Intelligence</b></p>  <p>Create ML solutions with this AI tool, and track their impact on your business demands and outcomes.</p>	<p><b>Additional Resources</b></p>  <p>Upgrade your AI experience to the next level</p>
--	---	--

**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.

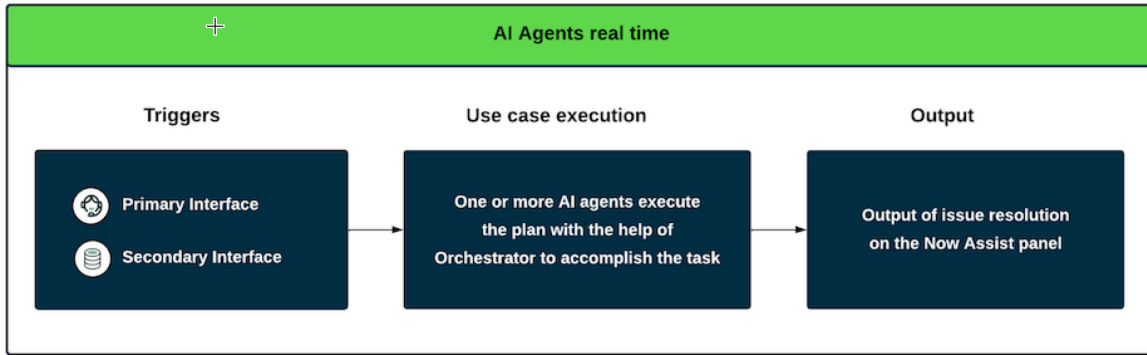
 **Now Assist framework**



**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.

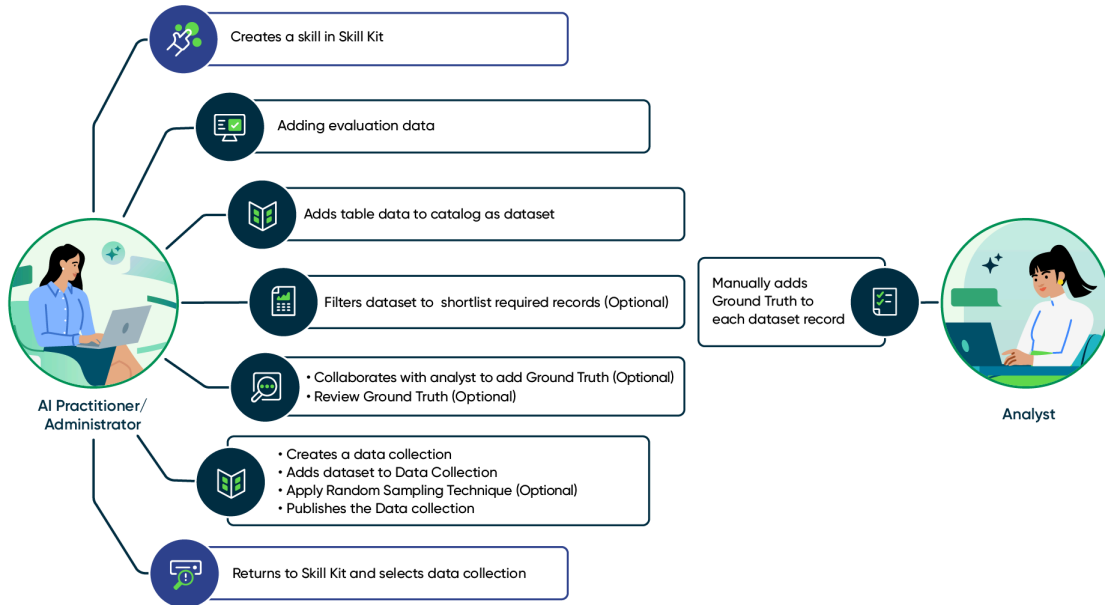
Use AI Agent Studio to create, manage, and test AI agents and use cases all in one place. The main Overview page has three sections that give you what you need to understand, begin, and continue developing AI agents and use cases.



## 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.

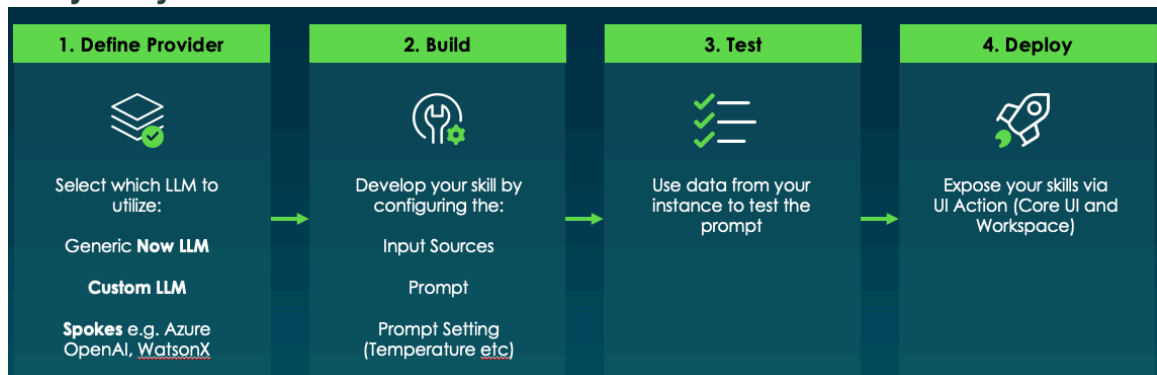
### User journey for Now Assist Data Kit



## 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 Governance

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

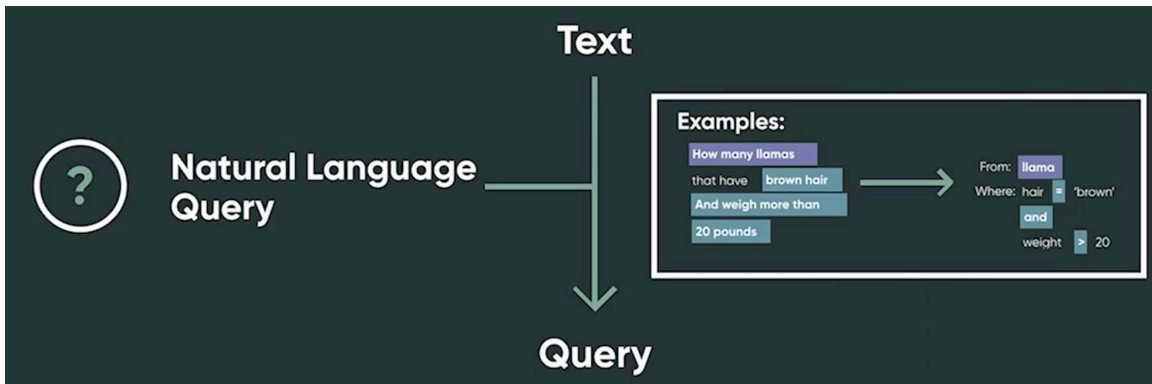
## Natural Language Understanding

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

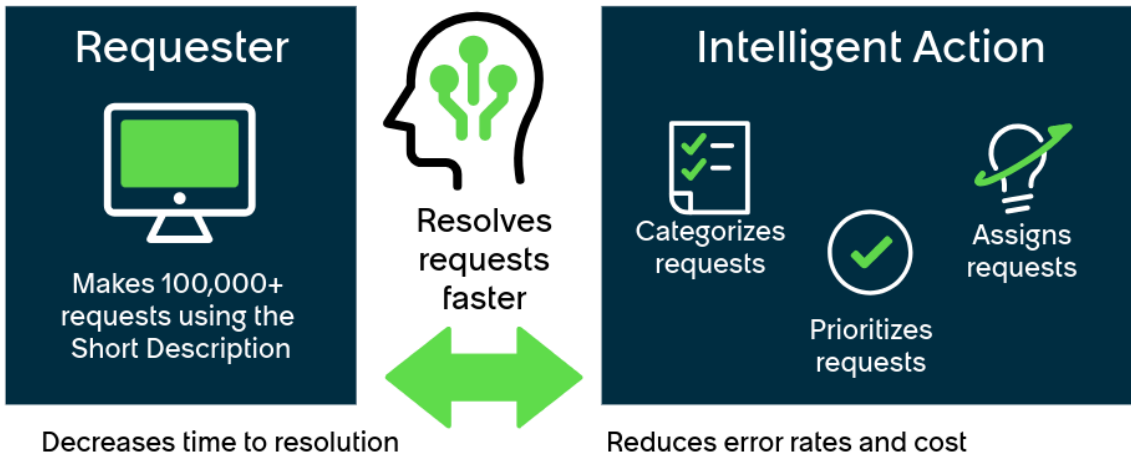
Transform natural-language questions into formal database queries with Natural Language Query (NLQ). Get data from your instance by using plain language requests in the supported languages American English, French, French Canadian, German, Japanese, and Spanish. NLQ is consumed by several other applications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



## Predictive Intelligence

Train your system to resolve requests swiftly with the Predictive Intelligence Workbench. There are four separate machine-learning frameworks that support all NLQ languages along with Dutch, Italian, and Brazilian Portuguese. These frameworks help your AI to handle a range of business challenges. Reduce task-handling times and interaction or error counts, solve issues shared between similar records, discover patterns and address records collectively, predict numeric outputs based on regression analysis, and train several other ML solutions.

# Predictive Intelligence



## 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.

PC Invoice 8323A In Progress All changes saved Submit

DegasFromagerieInvoice99PDF\_1.jpeg

Anytown, CA 94556  
Phone: (916) 555-0217

COMMENTS OR SPECIAL INSTRUCTIONS:  
Thank you for your order. Please remit payment immediately.

SALESPERSON	P.O. NUMBER	REQUISITIONER	SHIPPED VIA	F.O.B. POINT	TERMS
smc	Invoice Items (1/7)	4322	M J Grocer	Fedex	

Quant.	Description	Unit Price	Total
1 (18)	Stilton Blue round, 16 oz	6.77	121.86
2 (6)	Bue Brie, 12 oz	5.49	32.94
3 (11)	Swiss Raclette, 8 oz	4.09	44.99
4 (25)	Sharp milk feta, Greek, 8 oz	3.58	89.54
5 (1)	Gorgonzola cheese, 10 oz	6.03	6.03

SUBTOTAL 335.46  
SALES TAX 0  
SHIPPING & HANDLING 25.00  
TOTAL DUE 360.46

Make all checks payable to Degas Fromagerie  
If you have any questions concerning this invoice, contact Susan Anyone, (914) 555-1889 x44, sanyone@example.com

Invoice Items 5 rows All (20) To review (7) Reviewed (13) More

Id	Quantity	Description	Unit Price	Total	Actions
1	18	Stilton Blue round, 16 o	6.77	121.86	⋮
2	6	Bue Brie, 12 oz	5.49	32.94	⋮

Actions on selected rows: Mark as reviewed, Mark as unreviewed, Clear values, Delete rows

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.

## Applications and features



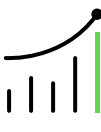
- Natural Language Understanding (NLU)
- Natural Language Query (NLQ)
- Predictive Intelligence
- Task Intelligence
- Document Intelligence



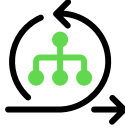
## Now Assist

ServiceNow<sup>®</sup> Now Assist uses generative AI that is designed to enhance user productivity and efficiency through conversation and proactive experiences.

[https://player.vimeo.com/video/1022777146?h=6a259700db&badge=0&autoplay=0&player\\_id=0&app\\_id=58479](https://player.vimeo.com/video/1022777146?h=6a259700db&badge=0&autoplay=0&player_id=0&app_id=58479)

## Get started

<p style="text-align: center;"><b>Explore</b></p>  <p style="text-align: center;">Learn more about Now Assist on the ServiceNow AI Platform.</p>	<p style="text-align: center;"><b>Configure</b></p>  <p style="text-align: center;">Explore, configure, use, and analyze generative artificial intelligence (AI) features and skills.</p>	<p style="text-align: center;"><b>Analyze</b></p>  <p style="text-align: center;">Analyze and monitor Now Assist performance.</p>
---	--	--

<p><b>Now LLM Service model updates</b></p>  <p>Learn about the Now LLM Service models.</p>	<p><b>Reference</b></p>  <p>Learn about user roles, data usage policy, and domain separation in Now Assist.</p>	<p><b>Generative AI Controller</b></p>  <p>Customize your generative AI experience on the ServiceNow AI Platform with the Generative AI Controller.</p>
--	--	--

**i 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.	<ul style="list-style-type: none"> <li>• <a href="#">Now Assist for Configuration Management Database (CMDB)</a></li> <li>• <a href="#">Now Assist for IT Operations Management (ITOM)</a></li> <li>• <a href="#">Now Assist for IT Service Management (ITSM)</a></li> <li>• <a href="#">Now Assist for Security Incident Response</a></li> <li>• <a href="#">Now Assist for Strategic Portfolio Management (SPM)</a></li> </ul>

**Now Assist products (continued)**

Workflow	Business areas	Available products
Customer	The Customer workflow includes applications that support customer service, including field service, financial services, telecommunications and media, and the public service sector.	<ul style="list-style-type: none"> <li>• <a href="#">Now Assist for Customer Service Management (CSM)</a> </li> <li>• <a href="#">Now Assist for Field Service Management (FSM)</a> </li> <li>• <a href="#">Now Assist for Financial Services Operations (FSO)</a> </li> <li>• <a href="#">Now Assist for PSDS</a> </li> <li>• <a href="#">Now Assist for Telecommunications, Media and Technology (TMT)</a> </li> </ul>
Employee	The Employee workflow supports HR Service Delivery features.	<ul style="list-style-type: none"> <li>• <a href="#">Now Assist for Health and Safety</a> </li> <li>• <a href="#">Now Assist for HR Service Delivery (HRSD)</a> </li> <li>• <a href="#">Now Assist for Legal Service Delivery (LSD)</a> </li> </ul>
Creator	<p>The Creator workflow supports a variety of Platform tools and builders, including the following:</p> <ul style="list-style-type: none"> <li>• App Engine Studio</li> <li>• ServiceNow AI Platform scripting</li> <li>• Platform Analytics</li> <li>• RPA Hub</li> <li>• Service Catalog</li> <li>• Workflow Studio</li> </ul>	<a href="#">Now Assist for Creator</a> 
Finance & Supply Chain	The Finance & Supply Chain workflow supports purchase requisitions, sourcing requests, and request for products or services.	<ul style="list-style-type: none"> <li>• <a href="#">Now Assist for Accounts Payable Operations (APO)</a> </li> <li>• <a href="#">Now Assist for Supplier Lifecycle Operations (SLO)</a> </li> <li>• <a href="#">Now Assist for Sourcing and Procurement Operations (SPO)</a> </li> </ul>

For a complete list of Now Assist skills and their workflows, see [Now Assist skills](#).

**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. Further, 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

Learn how Now Assist brings generative AI capabilities to the ServiceNow AI Platform. With Now Assist, 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 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.	<ul style="list-style-type: none"> <li>• <a href="#">Now Assist for Configuration Management Database (CMDB)</a></li> <li>• <a href="#">Now Assist for IT Operations Management (ITOM)</a></li> <li>• <a href="#">Now Assist for IT Service Management (ITSM)</a></li> <li>• <a href="#">Now Assist for Security Incident Response</a></li> <li>• <a href="#">Now Assist for Strategic Portfolio Management (SPM)</a></li> </ul>

**Now Assist products (continued)**

Workflow	Business areas	Available products
Customer	The Customer workflow includes applications that support customer service, including field service, financial services, telecommunications and media, and the public service sector.	<ul style="list-style-type: none"> <li>• <a href="#">Now Assist for Customer Service Management (CSM)</a> </li> <li>• <a href="#">Now Assist for Field Service Management (FSM)</a> </li> <li>• <a href="#">Now Assist for Financial Services Operations (FSO)</a> </li> <li>• <a href="#">Now Assist for PSDS</a> </li> <li>• <a href="#">Now Assist for Telecommunications, Media and Technology (TMT)</a> </li> </ul>
Employee	The Employee workflow supports HR Service Delivery features.	<ul style="list-style-type: none"> <li>• <a href="#">Now Assist for Health and Safety</a> </li> <li>• <a href="#">Now Assist for HR Service Delivery (HRSD)</a> </li> <li>• <a href="#">Now Assist for Legal Service Delivery (LSD)</a> </li> </ul>
Creator	<p>The Creator workflow supports a variety of Platform tools and builders, including the following:</p> <ul style="list-style-type: none"> <li>• App Engine Studio</li> <li>• ServiceNow AI Platform scripting</li> <li>• Platform Analytics</li> <li>• RPA Hub</li> <li>• Service Catalog</li> <li>• Workflow Studio</li> </ul>	<a href="#">Now Assist for Creator</a> 
Finance & Supply Chain	The Finance & Supply Chain workflow supports purchase requisitions, sourcing requests, and request for products or services.	<ul style="list-style-type: none"> <li>• <a href="#">Now Assist for Accounts Payable Operations (APO)</a> </li> <li>• <a href="#">Now Assist for Supplier Lifecycle Operations (SLO)</a> </li> <li>• <a href="#">Now Assist for Sourcing and Procurement Operations (SPO)</a> </li> </ul>

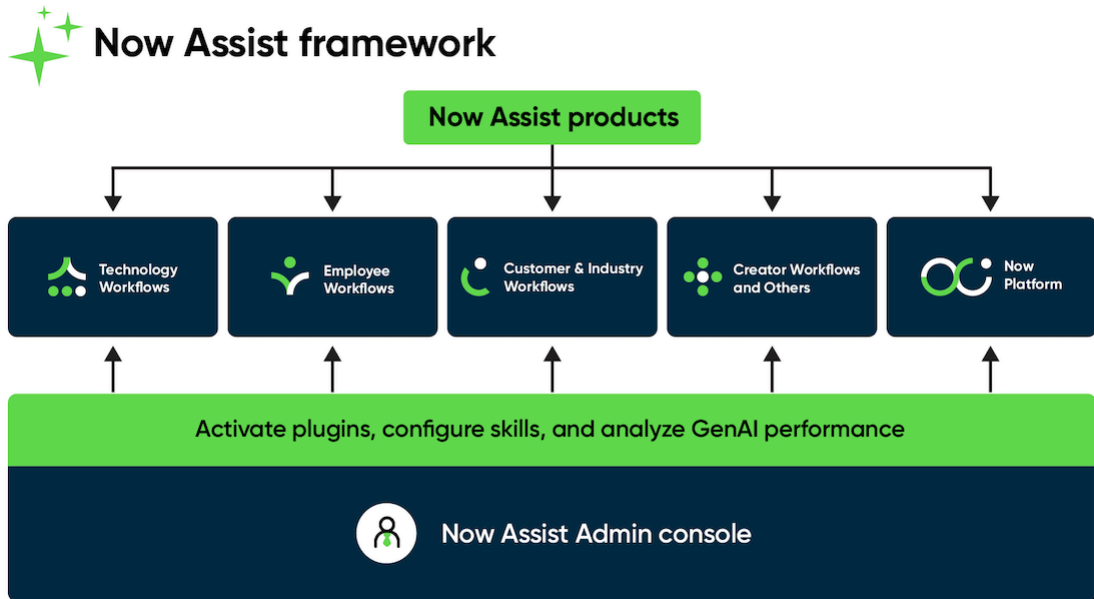
Now Assist products include some or all of the following foundational platform tools for Now Assist. For more information, see [Now Assist skills in the Platform workflow](#).

- Administrators install plugins, manage skills, and analyze usage and performance with the [Now Assist Admin console](#).
- 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 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.	<a href="#">Now Assist in AI Search</a>	Everyone
Install and configure Now Assist applications and the skills they provide.	<a href="#">Now Assist Admin console</a>	Administrators
Choose which skills to turn on, and which users can access them.	<a href="#">Now Assist Admin console</a>	Administrators
Monitor the usage and performance of your Now Assist implementation.	<a href="#">Analyzing Now Assist performance</a>	Administrators
Access generative AI skills in context through a user-friendly interface.	<a href="#">Now Assist panel</a>	Everyone
Use Now Assist skills on mobile devices.	<a href="#">Now Assist for Mobile</a>	Everyone
Customize your workflows and use your own third-party LLM license.	<a href="#">Generative AI Controller</a>	Administrators or developers

Benefit	Feature	Users
Use Now Assist in other platform features.	<a href="#">Now Assist skills in the Platform workflow</a>	Administrators or developers
Monitor Now Assist consumption on your instance.	<a href="#">Monitoring Now Assist usage in Subscription Management</a> 	Administrators

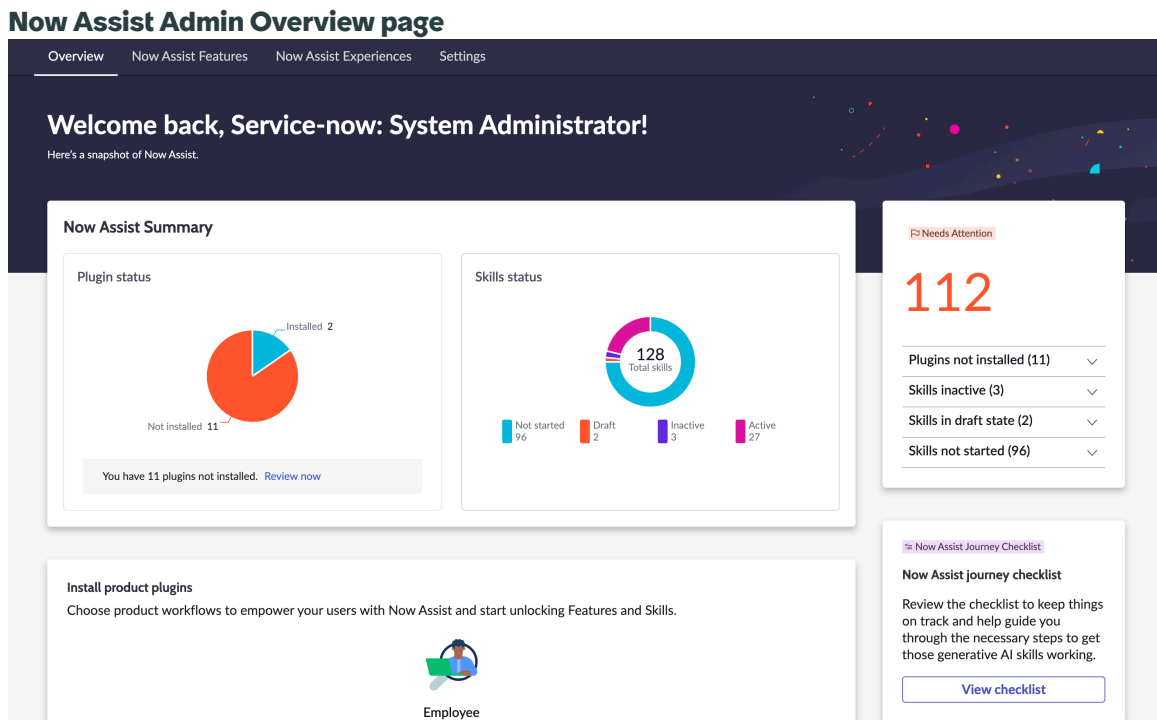
## Now Assist Admin console

The Now Assist Admin console provides you with quick and effortless access to the important information that you need to set up, configure, and monitor Now Assist applications and features.

## Now Assist Admin overview

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.



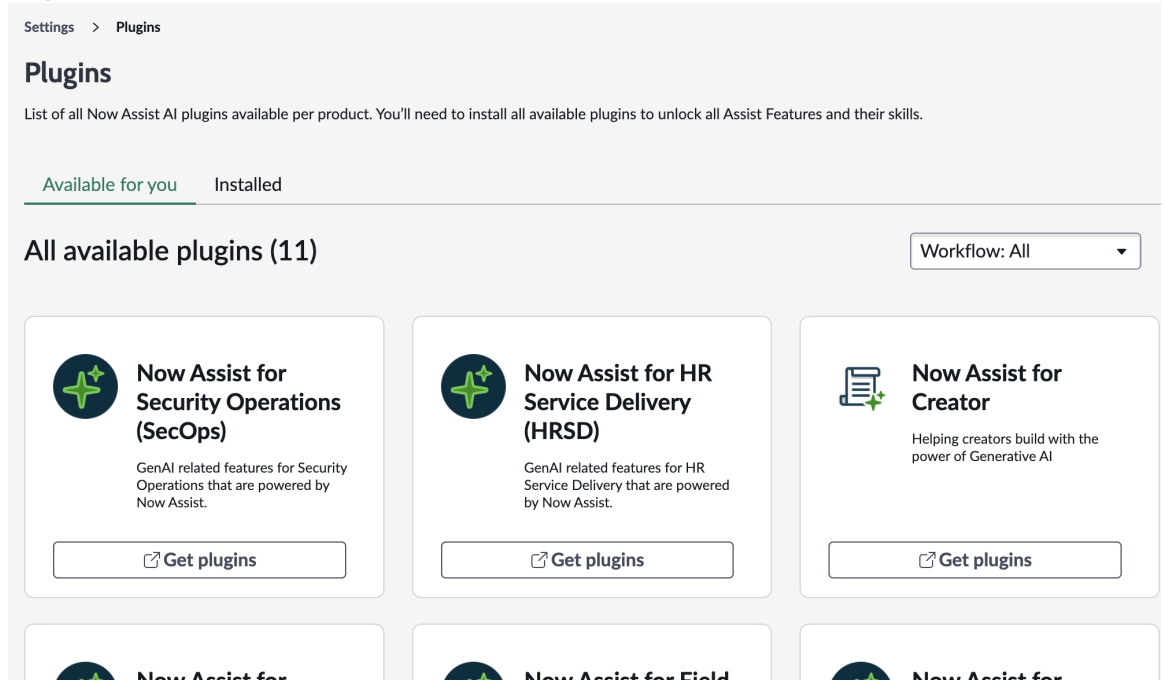
## Now Assist Admin workflow

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.

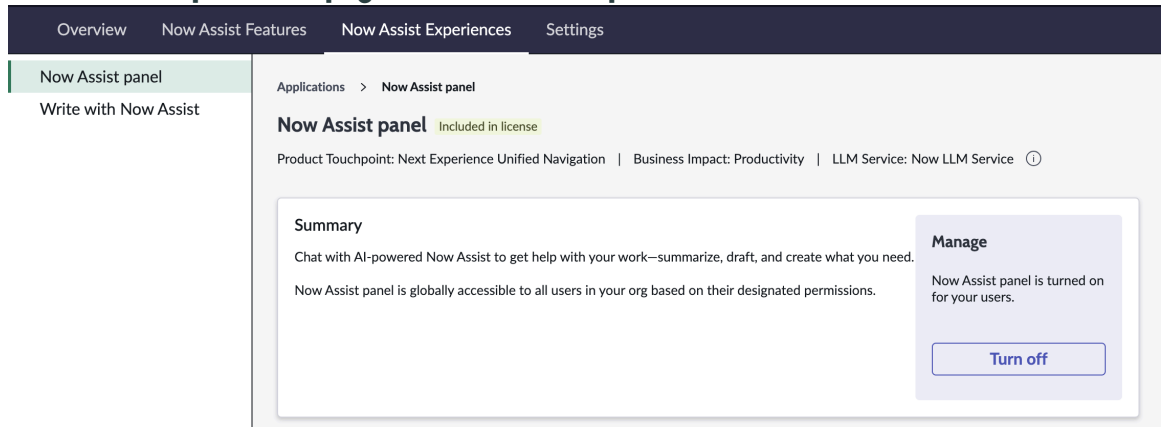
The Now Assist panel integrates the Now Assist skills into the Next Experience UI. By turning on the Now Assist panel directly from the Now Assist Admin console, you enable agents to access skills from anywhere on the ServiceNow AI Platform.

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.

### 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 Admin Features page to explore the skills that are available with your installed plugins. By selecting the **View details** button, you can

see more information about each skill. The following diagram shows the available features and skills in the Technology workflow.

**Available Now Assist features and skills in the Technology workflow**

**Now Assist skills for Technology**

Explore Now Assist skills for Technology products.

Select product  
ITSM

**Chat**

All chat-related functionality for workspaces and Now Assist Panel.

All Chat skills

Chat summarization Active

[View details](#)

**Incident**

All incident-related functionality for workspaces and Now Assist Panel.

All Incident skills

Incident summarization Active

Resolution notes generation Active

[View details](#)

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.

## Chat summarization activation in Now Assist Admin

Chat summarization CSM
Exit

- Define trigger ●
- Choose Input ●
- Select display ●
- Review and activate ●

### Define your trigger

Choose when and how the skill will be triggered. [Explain this](#)

**Virtual Agent to Live Agent handoff**  
A chat summary will be created when the conversation moves from a virtual agent to a live agent

**Live Agent to Live Agent handoff**  
A chat summary will be created when the conversation moves from a live agent to a live agent

**Quick action**  
Agents can use the summarize quick action to summarize the chat

**Chat wrap-up**  
The chat summary field will auto-populate after the conversation ends

Back Save and continue

### 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

Overview Now Assist Features Now Assist Experiences Settings

- Plugins
- Language and region
- Account

Settings > Account

### Account

View your account license information and manage data sharing.

**Account details**

Version 7.0.2

Your license includes: All included in license

- ✔

**Now Assist panel**

Empower users with a personal generative AI assistant that can help generate, draft, and summarize content. Requires a product plugin to install.
- ✔

**Now Assist skills**

Unlock Now Assist skills you can configure to improve user productivity. Requires product plugins to install.

Currently on

**Data sharing: Your data improves the experience**

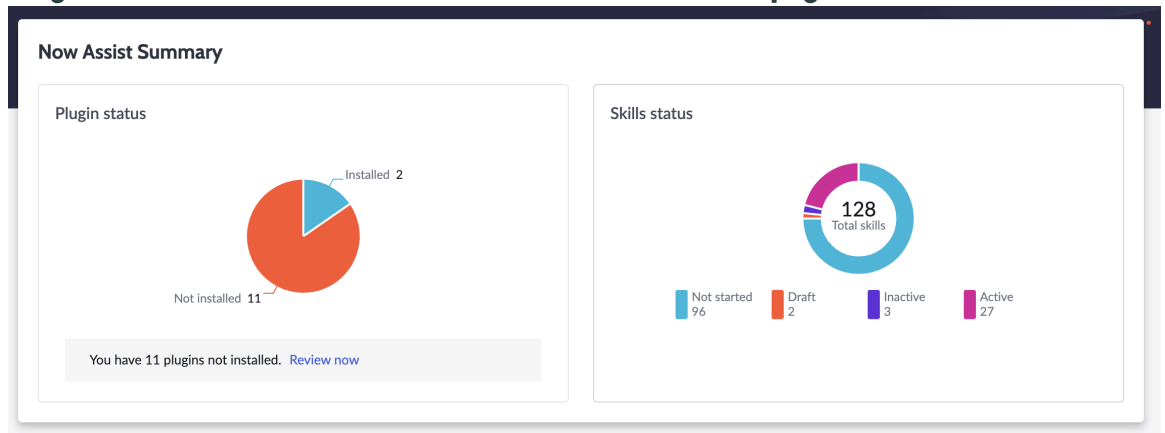
By sharing your data with ServiceNow, you're playing a valuable role in the evolution of more accurate AI systems that aim to enhance user experiences and better understand your business needs. The security framework around our data sharing is designed to keep your personal information confidential and private through rigorous protocols, anonymization, and careful handling during model training. To learn more, check out our FAQs. Manage who can opt-out of data sharing by assigning users to a Data Steward role. [Check out the documentation](#) to learn how to manage this role.

Opt out

### 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.

## 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](#) the overall health and performance of Now Assist.

## 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 or generate resolution notes so that you can get the context of this information more quickly.


Before using the Now Assist panel:

- Next Experience must be enabled on your instance. For more information, see [Considerations for activating Next Experience](#) .
- You must have the Now Assist panel user (now\_assist\_panel\_user) role.
- Your role must be applied to at least one active Now Assist skill to use skills on the Now Assist panel.
- You must activate the Now Assist panel. For more information, see [Turn on the Now Assist panel](#).

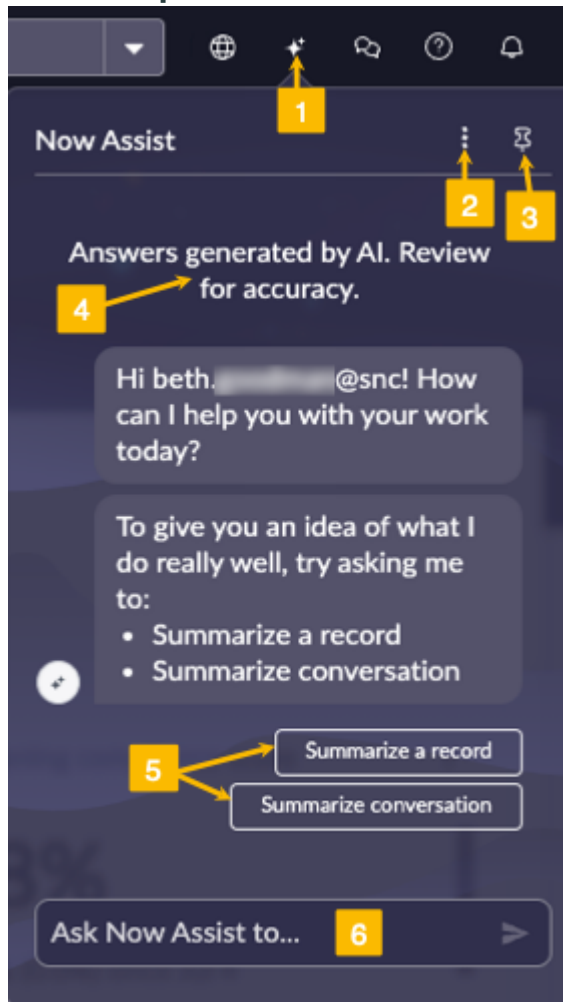
## Now Assist Overview

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.


Let's get started by selecting the Now Assist icon  to display the Now Assist panel.

### Now Assist panel



The Now Assist panel includes:

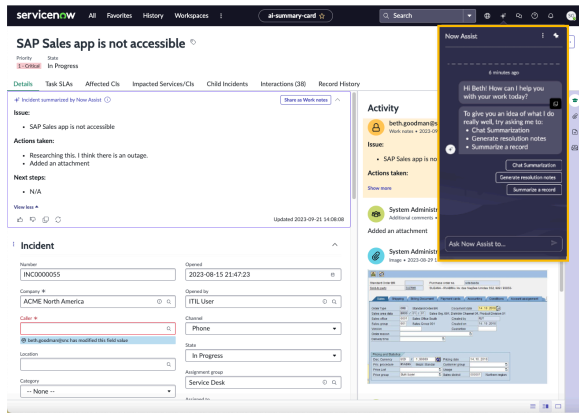
Item number	Description
1 - ✨	<p>Now Assist icon</p> <p>Displays the Now Assist panel.</p> <p>If a number in a square appears, it indicates the number of messages you have missed when the Now Assist panel was closed. </p>
2 - ⋮	<p>Reset icon</p> <p>Select the reset icon if you think that Now Assist doesn't understand what you want and you'd like to restart the conversation. If you reset the conversation, the existing messages are deleted, and a new conversation is started. You won't be able to access the previous messages after you reset the conversation.</p>

Item number	Description
3 - 	Pushpin icon  Permanently pin the Now Assist panel to the screen.
4 - Now Assist message	Indicates that the answers are generated by AI.
5 - Option buttons	Option buttons
6 - Ask Now Assist to..... field	Field where the user can enter actions.

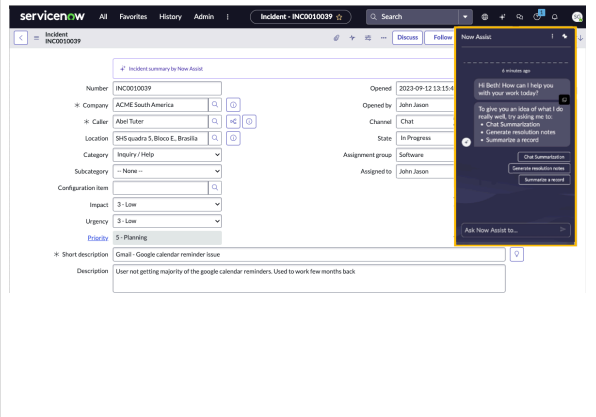
Now Assist panel is available on Next Experience and ServiceNow Studio. The following screenshots show the Now Assist panel in a workspace and on Core UI screens under Next Experience.


### Now Assist panel

#### Next Experience



#### Core UI



If you think that Now Assist doesn't understand what you want and you'd like to restart the conversation, select the reset icon  from the Now Assist panel. If you reset the conversation, the existing messages are deleted, and a new conversation is started. You won't be able to access the previous messages after you reset the conversation.

### 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 the **Ask Now Assist to** field field.

**Note:** You can also generate a chat summarization by using the `/summarize` quick action in Agent Chat.

For more information about the chat summarization, see [Chat summarization](#).

## 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.

For more information about the case or incident summarization, see [Record summarization](#).

## 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.

For more information about generating resolution notes, see [Resolution notes generation](#).

## 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 menu actions. The Now Assist context menu helps you summarize, create or modify existing documentation.

### 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.




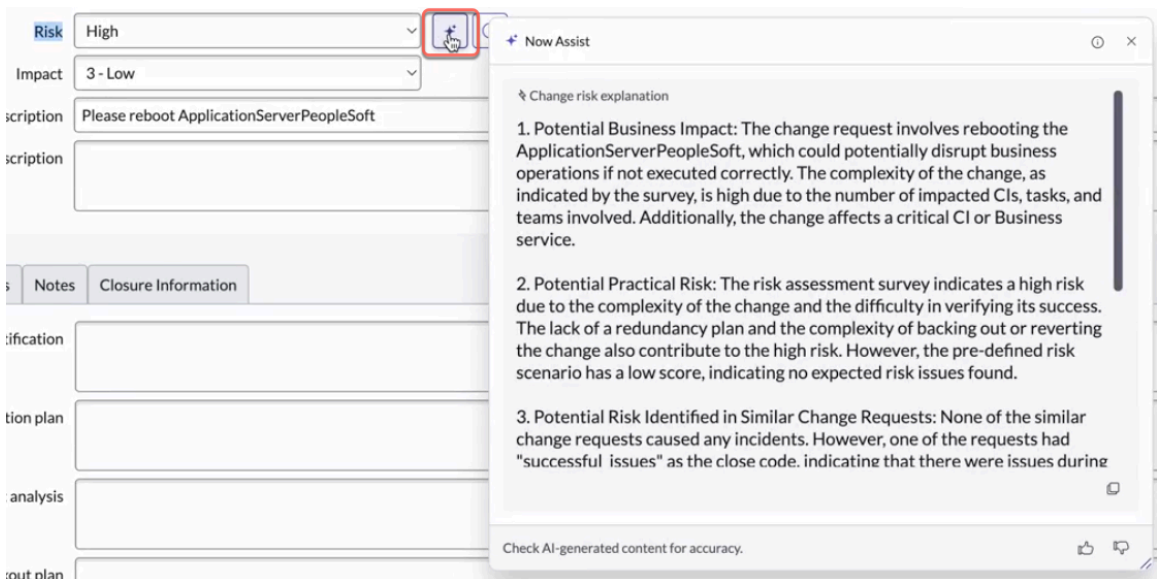
### Change request risk explanation using the Now Assist context menu

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 dialog box in the Record information section.



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 dialog box.




**Note:** The risk explanation that is presented in the dialog box is assessed and calculated on the change request form.

For more information about risk assessment and calculation, see [Risk assessment](#).

### Content editing using Now Assist content menu

The Now Assist context menu enables generative AI assisted content editing capabilities for user authoring and providing resolution notes in workspaces and UI16.

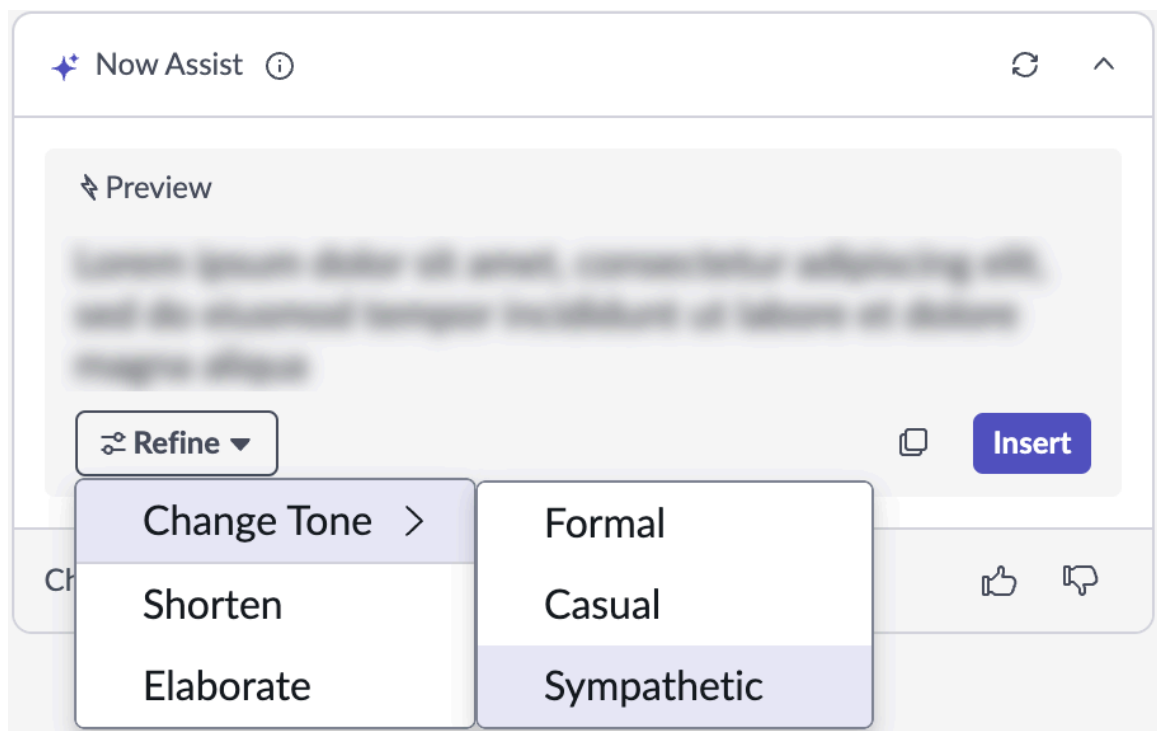
When you select some content, you'll see the Now Assist  icon pop-up and float along with your mouse device.

When you hover over the Now Assist icon or select **Refine** menu, you will 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.
- **Change Tone:** Generative AI changes the tone of the selected text to sound formal, casual, or Sympathetic.

If you select **Elaborate** or **Shorten**, generative AI provides alternative content option. You can select **Insert** button to update the content and publish. For more information, see [Edit an article using the Now Assist context menu](#) .

If you select **Change Tone**, Now Assist context menu allows you to further choose a preferred tone between **Formal**, **Casual**, and **Sympathetic**. Select **Insert** to add the refined text.

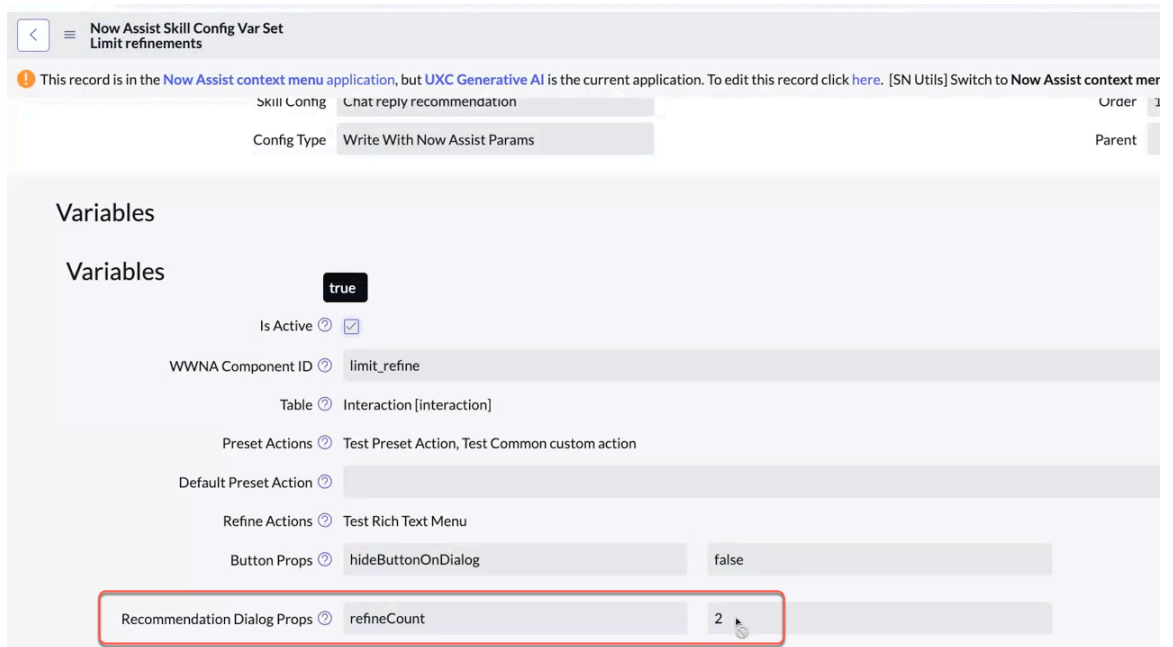


**Limit the number of content refinement calls using the Now Assist context menu**

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.

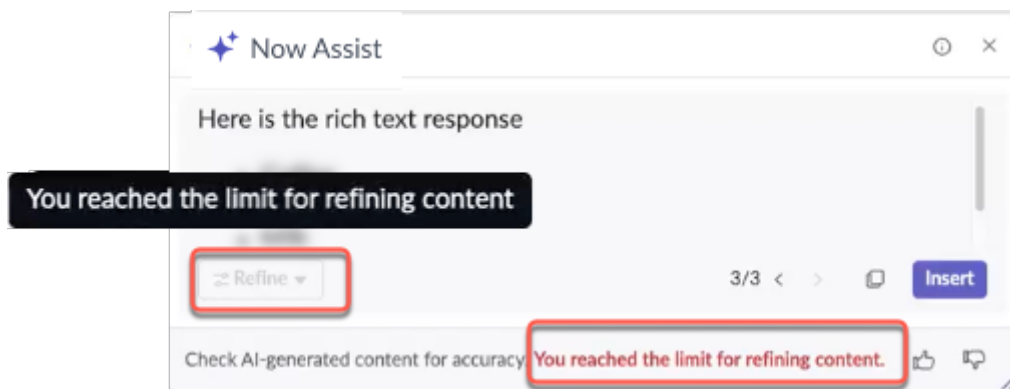
You can configure the refinement calls limitation at the **Recommendations Dialog Props** field in the **Limit refinements** record using the *refineCount* property.

**Note:** The **Limit refinements** record is available in the Now Assist config var set record in the Now Assist Skill context menu application.

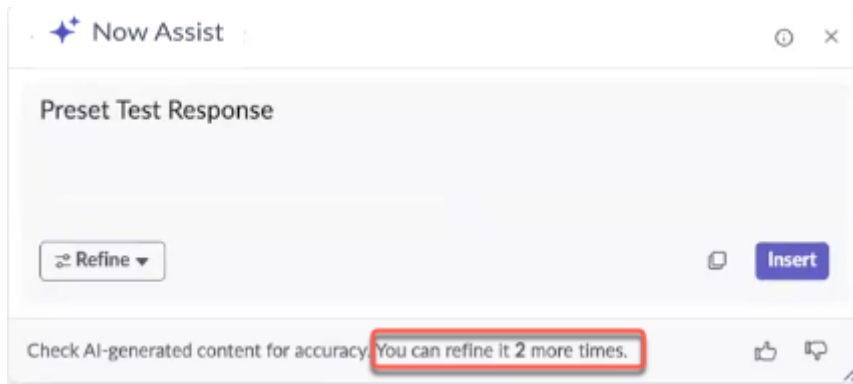


You can set the refineCount as follows:

- 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.



- 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.




## Related topics

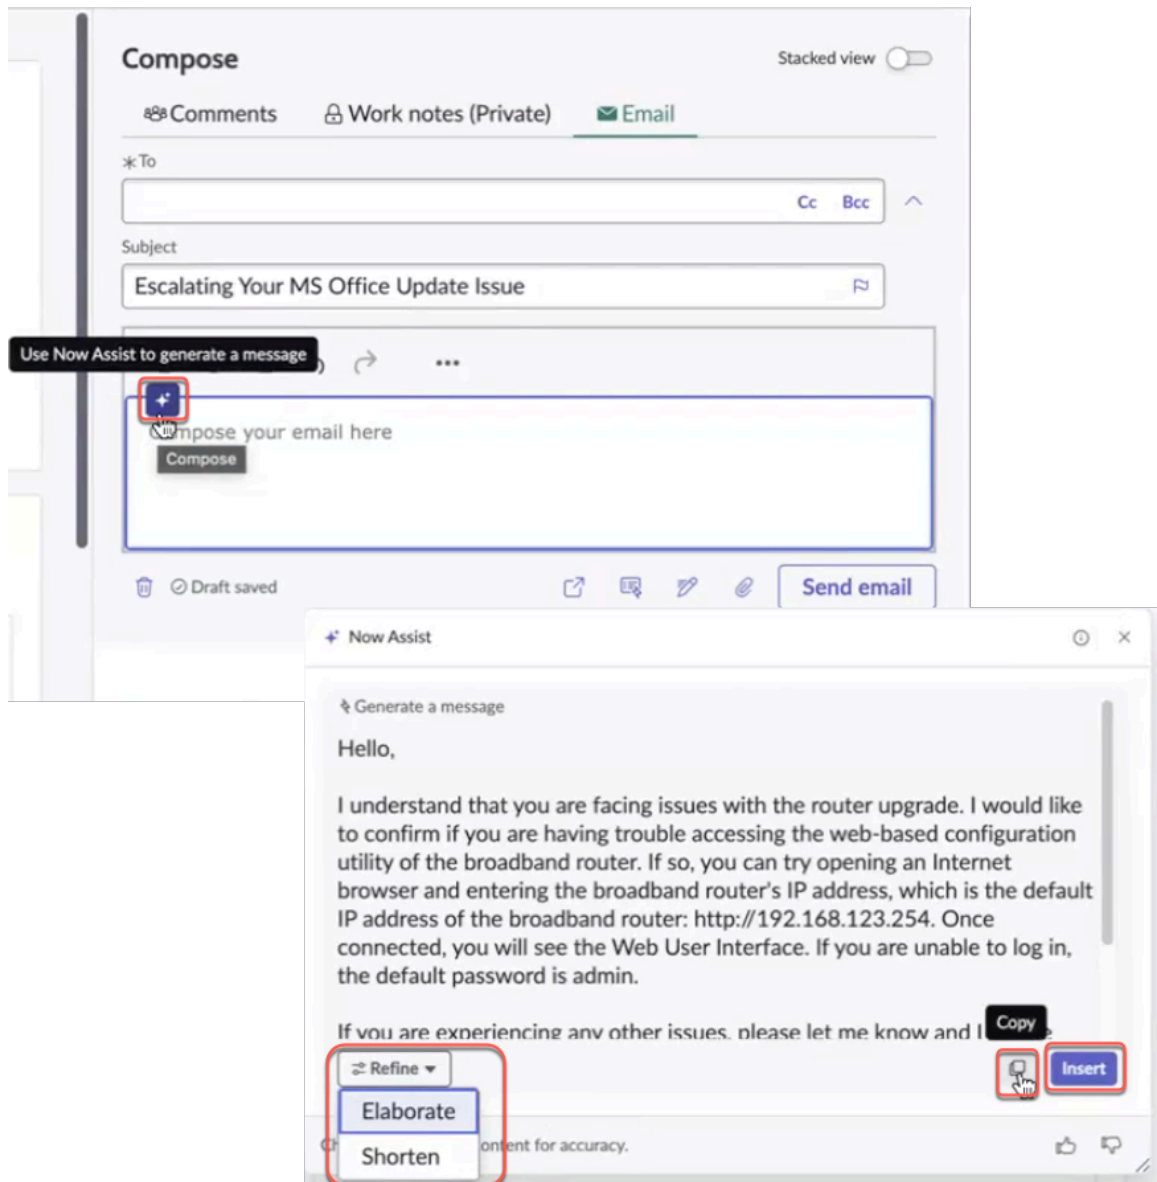
[Edit an article using the Now Assist context menu](#) 

## 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

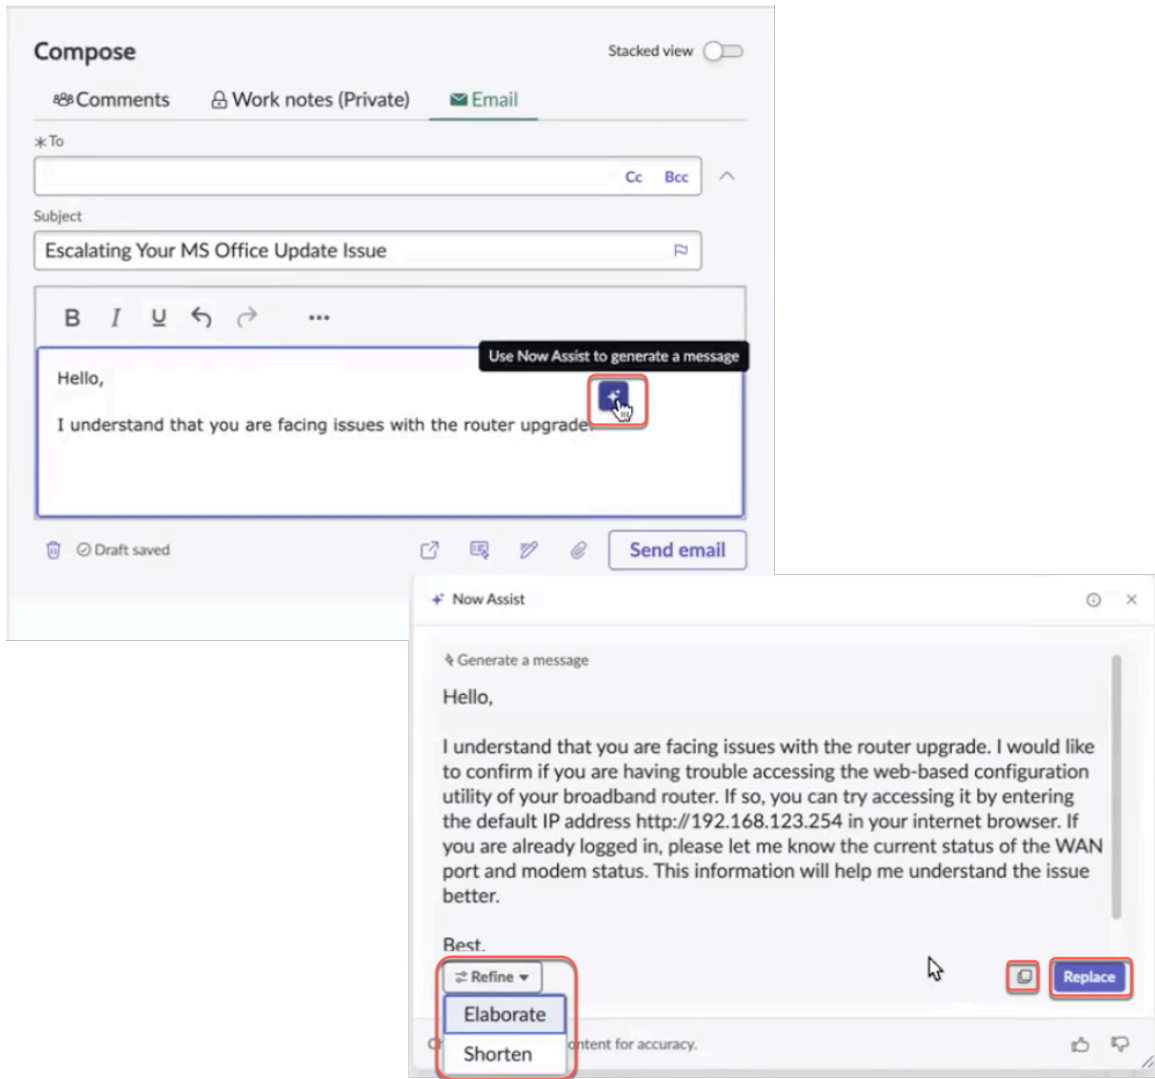
While you compose a new email, you see the Now Assist context menu icon () that displays the message Use Now Assist to generate a message. When you select the icon, generative AI generates a message recommendation.



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.

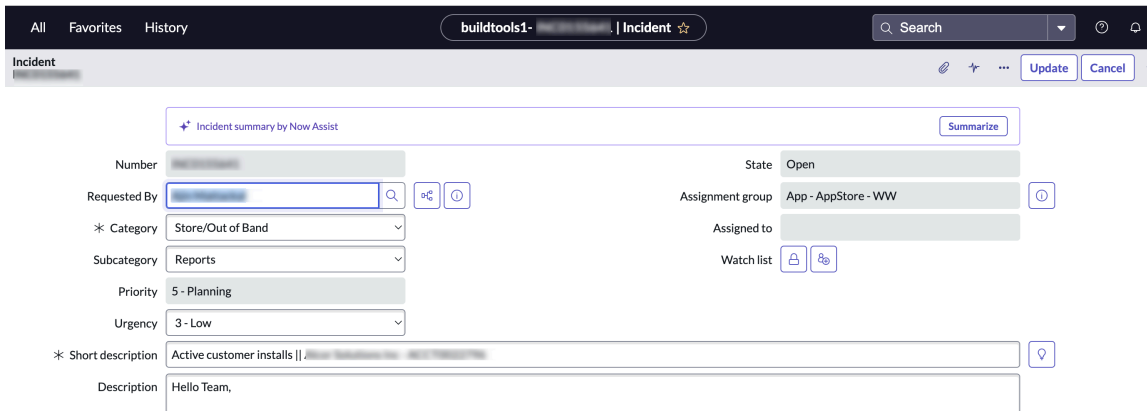


You can further refine the generated text by selecting **Elaborate** or **Shorten** context menu options on the Now Assist admin model. You can also copy the new text or select **Replace** to replace the generated text into the email body.

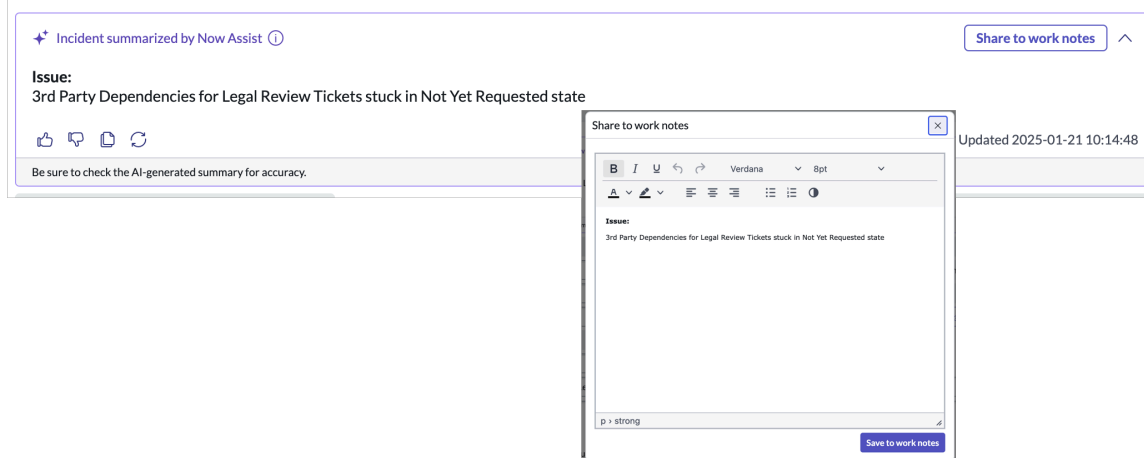
### Summarize records with the Now Assist context menu

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 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.



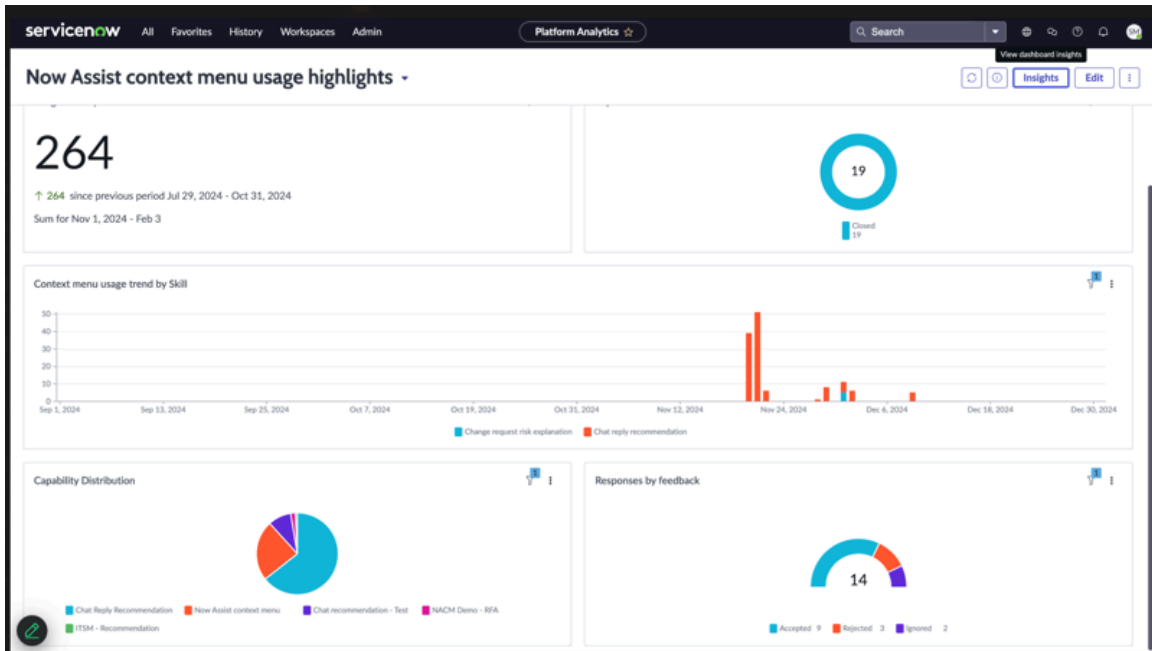
To provide feedback, use the feedback icons (👍) or (👎). If you're not satisfied with the provided summary, use the regenerate icon (🔄) to regenerate the summary. You can also copy the summary to the clipboard using the copy icon (📄).

### 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.

## Now Assist context menu usage highlights ▾

Date

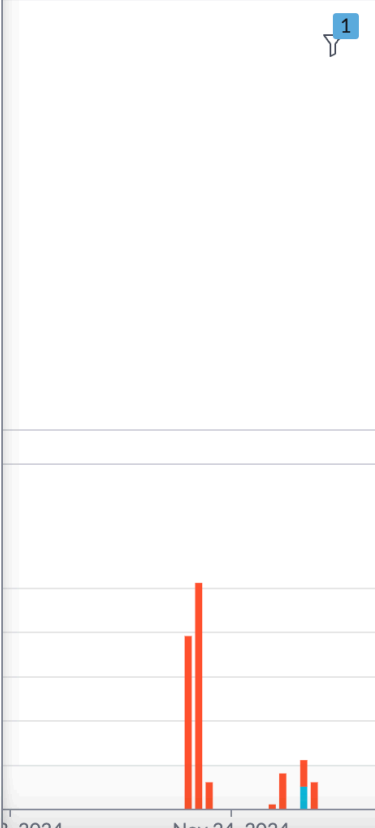
 Last 6 months ▾

Custom range < Mar ▾ 2025 ▾ >

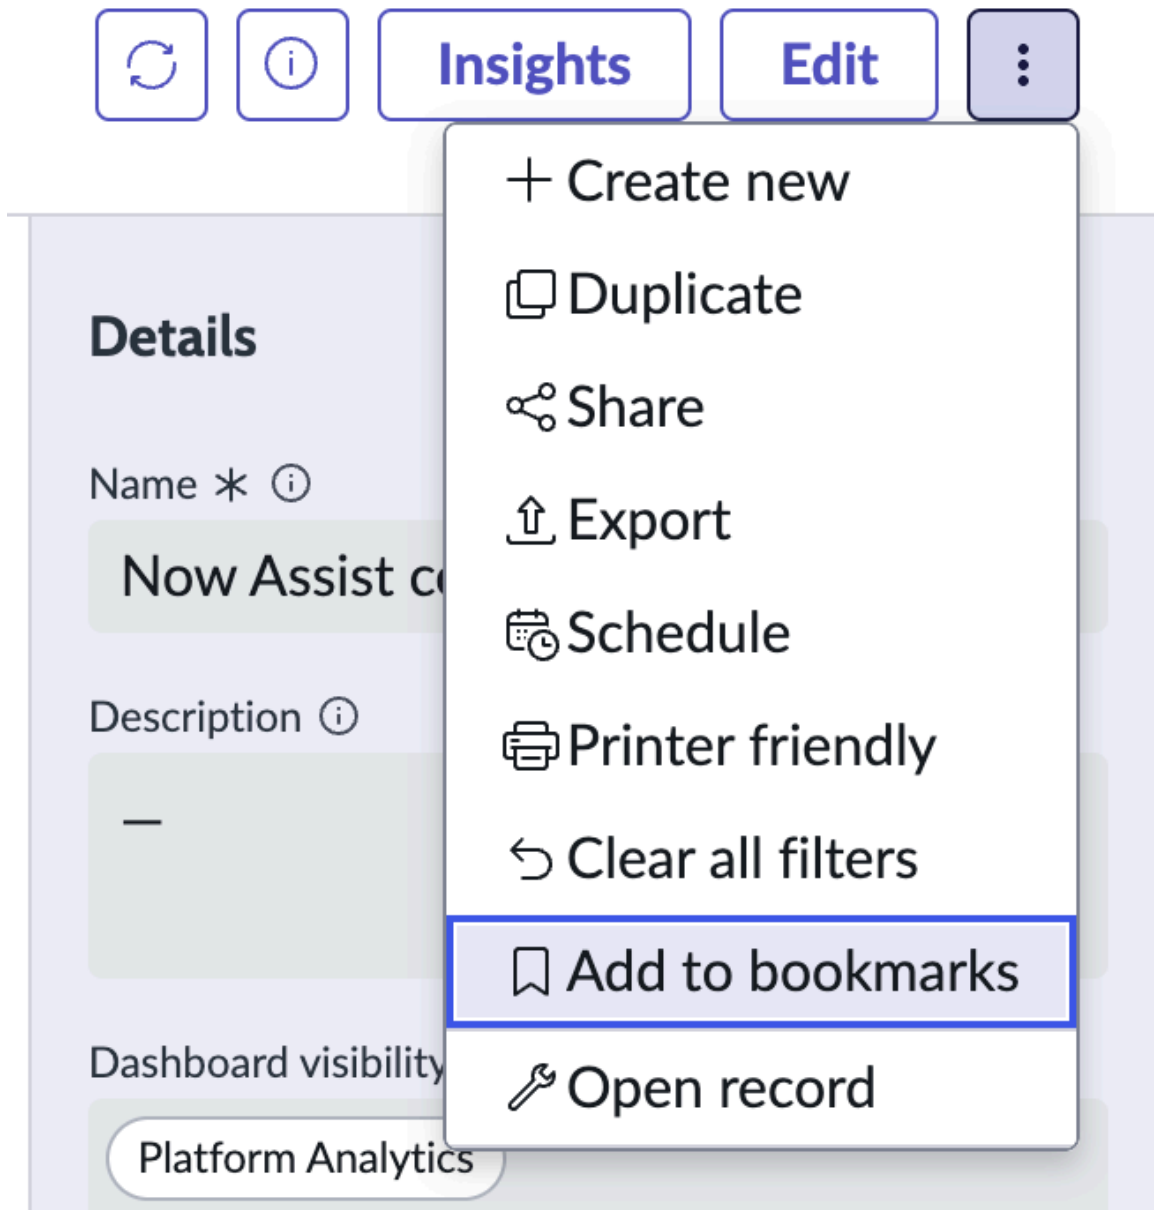
S	M	T	W	T	F	S
23	24	25	26	27	28	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31	1	2	3	4	5

Start date

End date



You have an option to refresh, duplicate, edit, create, and export a dashboard.




You can also select the info icon (i) to view the dashboard-related information.

  **Insights** **Edit** 

## Details



Name \* 

Now Assist context menu usage high...

Description 


—

Dashboard visibility \*

Platform Analytics

Dashboard category

—



Owner



## 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.

**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 settings** tab, 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.

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 tab.

11. Select options for the following fields on the Configure experience tab:

- 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 tab.

14. Select an option from the **Select a record to test the configurations** drop-down menu.

15. Select **Preview** and **Done**.

## 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.

### Available skills by workflow

Workflow	Product	Available skills
Technology	<a href="#">Now Assist for Configuration Management Database (CMDB)</a>	<ul style="list-style-type: none"> <li>• Configuration item (CI) summarization</li> <li>• Manage duplicate configuration items (CIs)</li> </ul>
Technology	<a href="#">Now Assist for IT Operations Management (ITOM)</a>	<ul style="list-style-type: none"> <li>• Alert analysis</li> <li>• Alert investigation</li> </ul>
Technology	<a href="#">Now Assist for IT Service Management (ITSM)</a>	<ul style="list-style-type: none"> <li>• Change request risk explanation</li> <li>• Change request summarization</li> <li>• Chat reply recommendation</li> <li>• Chat summarization</li> <li>• Incident assist</li> <li>• Incident summarization</li> <li>• KB generation</li> <li>• Resolution notes generation</li> <li>• Sidebar discussion summarization</li> </ul>
Technology	<a href="#">Now Assist for Security Incident Response</a>	<ul style="list-style-type: none"> <li>• Post incident analysis</li> <li>• Security incident recommended actions</li> <li>• Security incident summarization</li> <li>• Resolution notes generation</li> <li>• Correlation insights </li> </ul>
Technology	<a href="#">Service Graph Connectors</a>	Service Graph Connector diagnosis
Technology	<a href="#">Now Assist for Strategic Portfolio Management (SPM)</a>	<ul style="list-style-type: none"> <li>• Email project summary</li> <li>• Feedback summarization</li> <li>• Multi feedback summarization</li> <li>• Project Gen AI docs</li> </ul>



Available skills by workflow (continued)

Workflow	Product	Available skills
		<ul style="list-style-type: none"> <li>• Planning item Gen AI docs</li> <li>• EAP teams Gen AI docs</li> </ul>
Customer	Now Assist for Customer Service Management (CSM) <a href="#">↗</a>	<ul style="list-style-type: none"> <li>• Case summarization</li> <li>• Chat reply recommendation</li> <li>• Chat summarization</li> <li>• Email response skill</li> <li>• KB generation</li> <li>• Resolution notes generation</li> <li>• Sidebar discussion summarization</li> <li>• Voice call summarization</li> </ul>
Customer	Now Assist for Field Service Management (FSM) <a href="#">↗</a>	<ul style="list-style-type: none"> <li>• KB generation</li> <li>• Sidebar discussion summarization</li> <li>• Work order task closure summarization</li> </ul>
Customer	Now Assist for Financial Services Operations (FSO) <a href="#">↗</a>	<ul style="list-style-type: none"> <li>• Claim case summarization</li> <li>• Dispute case summarization</li> </ul>
Customer	Now Assist for PSDS <a href="#">↗</a>	<ul style="list-style-type: none"> <li>• Chat summarization</li> <li>• Government case summarization</li> <li>• Resolution notes generation</li> </ul>
Customer	Now Assist for Telecommunications, Media and Technology (TMT) <a href="#">↗</a>	<ul style="list-style-type: none"> <li>• Resolution notes generation</li> <li>• Service problem case summarization</li> <li>• Test summarization</li> </ul>
Employee	Now Assist for Health and Safety <a href="#">↗</a>	Health and Safety incident summarization
Employee	Now Assist for HR Service Delivery (HRSD) <a href="#">↗</a>	<ul style="list-style-type: none"> <li>• Chat summarization</li> <li>• Sidebar discussion summarization</li> <li>• Chat reply recommendation</li> <li>• KB generation</li> <li>• Email response skill</li> </ul>

Available skills by workflow (continued)

Workflow	Product	Available skills
		<ul style="list-style-type: none"> <li>• Case summarization</li> <li>• Resolution notes generation</li> </ul>
Employee	Now Assist for Legal Service Delivery (LSD) <a href="#">↗</a>	<ul style="list-style-type: none"> <li>• Legal matter summarization</li> <li>• Legal request summarization</li> </ul>
Employee	Now Assist in Contract Management <a href="#">↗</a>	<ul style="list-style-type: none"> <li>• Contract analysis</li> <li>• Contract metadata extraction</li> </ul>
Creator	Now Assist for Creator <a href="#">↗</a>	<ul style="list-style-type: none"> <li>• App generation</li> <li>• Catalog item generation</li> <li>• Intelligent code recommendations</li> <li>• Data visualization generation</li> <li>• Flow generation</li> <li>• Flow recommendations</li> <li>• Playbook generation</li> <li>• Robotic Process Automation (RPA) bot generation <a href="#">↗</a></li> <li>• Refine content</li> <li>• Spoke generation</li> <li>• UI generation</li> </ul>
Platform	Now Assist skills in the Platform workflow Platform	<p><b>Now Assist conversational experience</b></p> <ul style="list-style-type: none"> <li>• Now Assist Q&amp;A Genius Results</li> <li>• Now Assist Multi-Turn Catalog Ordering</li> <li>• Now Assist Topics</li> <li>• Subflows and actions</li> </ul> <p><b>Knowledge</b></p> <p>Knowledge content recommendation</p> <p><b>Platform skills</b></p> <p>Navigation</p>
Finance & Supply Chain	Now Assist for Accounts Payable Operations (APO) <a href="#">↗</a>	Invoice case summarization

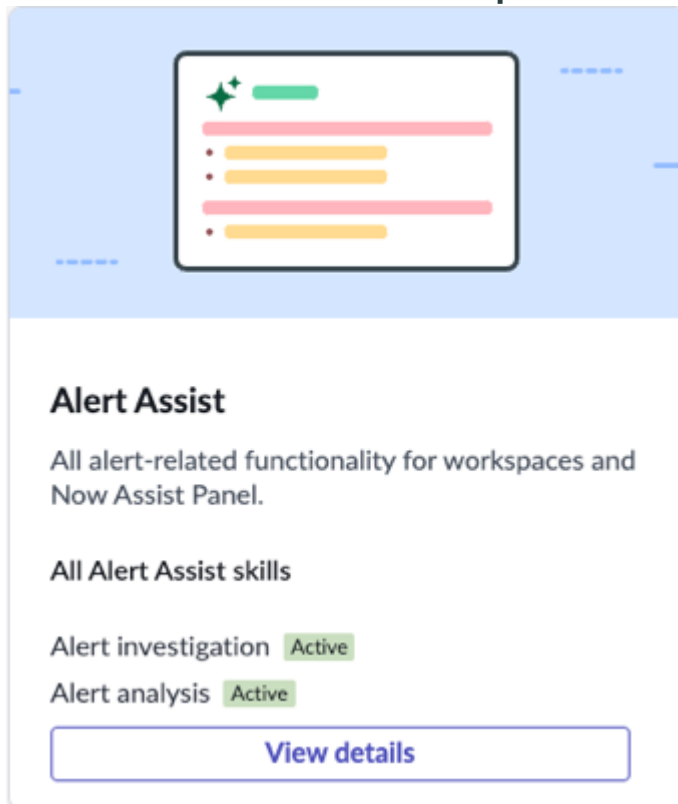
**Available skills by workflow (continued)**

Workflow	Product	Available skills
Finance & Supply Chain	Now Assist for Supplier Lifecycle Operations (SLO) 	Supplier case summarization
Finance & Supply Chain	Now Assist for Sourcing and Procurement Operations (SPO) 	Fulfiller summarization for Sourcing and Procurement Operations

**Now Assist analysis skills**

These skills provide generative AI capabilities for analyzing data on the ServiceNow AI Platform, including security incident analysis, alerts in Event Management, and legal contracts.

**Alert Assist tile for Now Assist for IT Operations Management (ITOM)**



## Security Incident skill tile with Post incident analysis skill



The image shows a skill tile for 'Security Incident' in ServiceNow. At the top, there is a blue header area containing a white box with a green star icon and several horizontal bars in pink and yellow. Below this, the main content area is white. It features the title 'Security Incident' in bold, followed by a description: 'All security incident-related functionality for workspaces and Now Assist Panel.' Underneath, there is a section titled 'All Security Incident skills' which lists four skills with their status: 'Security Incident Recommended Actions' (Active), 'Post incident analysis' (Not started), 'Security incident summarization' (Active), and 'Resolution notes generation' (Active). At the bottom of the tile is a blue button labeled 'View details'.

### Security Incident

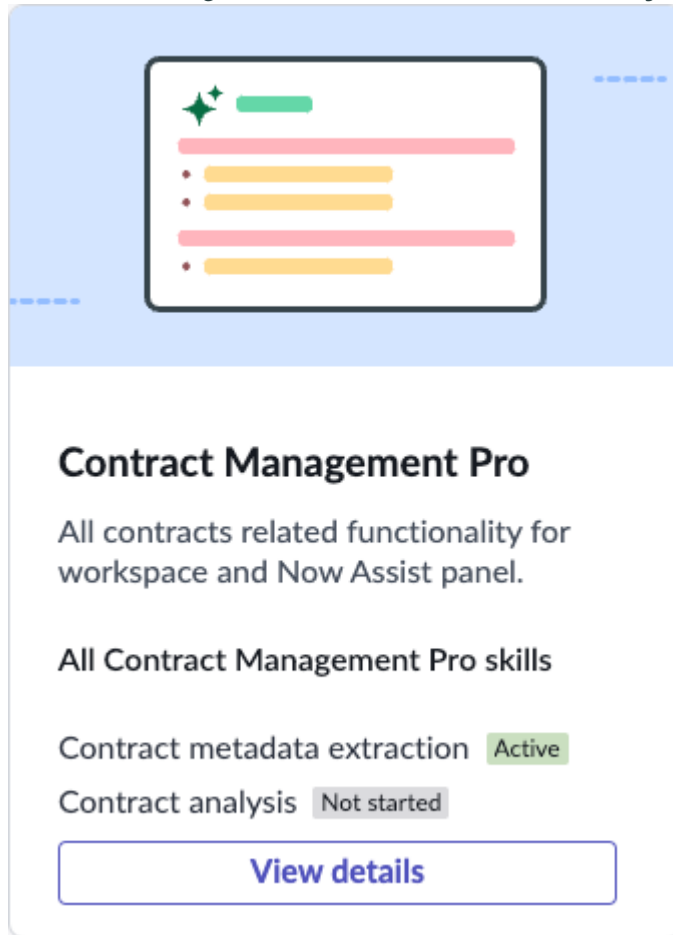
All security incident-related functionality for workspaces and Now Assist Panel.

#### All Security Incident skills

- Security Incident Recommended Actions Active
- Post incident analysis Not started
- Security incident summarization Active
- Resolution notes generation Active

[View details](#)

**Contract Management skill tile with Contract analysis skill**



**Alert analysis**

Use this skill to provide simplified generative AI alert analyses for the ServiceNow Event Management application.

Alert analyses include a human-readable brief of the alert, as well as technical information to help you investigate the alert more effectively. In Event Management, alert analysis is a crucial feature that completely transforms the alert triage process and can help achieve faster, seamless alert management. This skill uses a large language model (LLM) to generate human-readable alert titles and detailed descriptions based on enriched and actionable alert data. The AI-generated alert information facilitates quick alert triage and effective analysis for your team, which can lead to a substantial reduction in resolution time.

For more information, see [Using Now Assist for IT Operations Management \(ITOM\)](#) [↗](#).

**Availability**

This skill is available in the workflow and product listed below.

**Now Assist products and workflows**

Workflow	Product
Technology	<a href="#">Now Assist for IT Operations Management (ITOM)</a> <a href="#">↗</a>

**Alert investigation**

This skill enables Now Assist to investigate past related incidents.

Now Assist queries historical records to find past incidents on the same or related CIs, and analyzes their frequency, criticality, work notes, and resolution. It presents a summary of the most relevant related incidents in the Now Assist panel. The summary includes resolution strategies used, as well as contact details for individuals or teams who resolved the incidents and could assist with the current alert. The information provided helps you better understand the context and significance of the current alert and work more efficiently toward resolving it.

For more information, see .

## Availability

This skill is available in the workflow and product listed below.

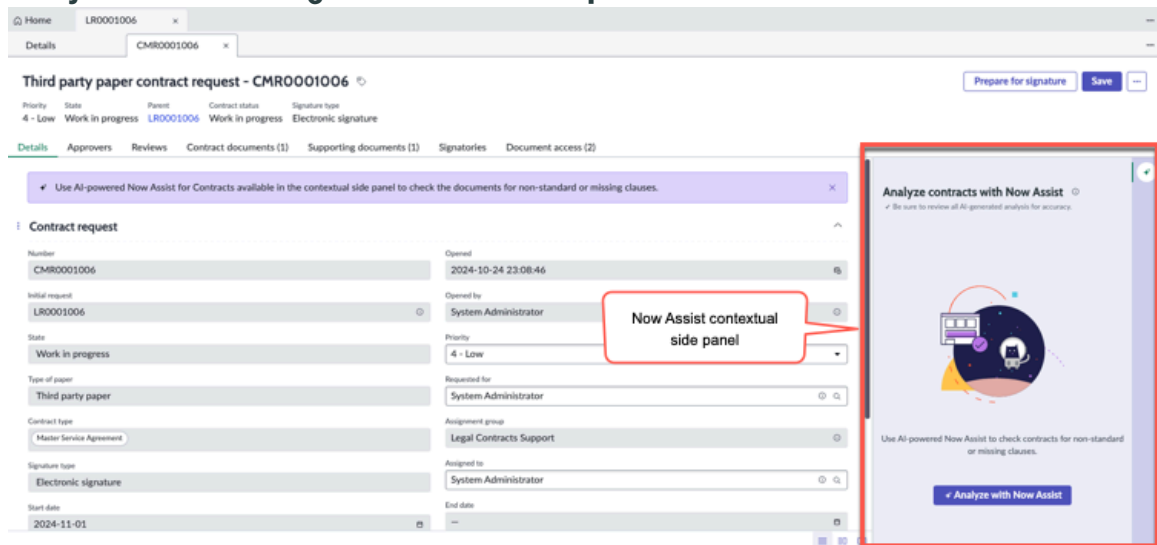
### Now Assist products and workflows

Workflow	Product
Technology	<a href="#">Now Assist for IT Operations Management (ITOM)</a>

## Contract analysis

This skill uses generative AI to analyze contracts and identify nonstandard and missing clauses.

### Analyze a contract using the Now Assist side panel



Once generated, a fulfiller can review the Now Assist recommendations and make a revision.

## Generating a contract analysis

You can generate a contract analysis in the following product.

[Contract analysis using Now Assist in Contract Management](#)

## Availability

This skill is available in the workflows and products listed below.

### Now Assist products and workflows

Workflow	Product
Employee	<p><a href="#">Now Assist in Contract Management</a> is available in the following products:</p> <ul style="list-style-type: none"> <li>• <a href="#">Now Assist for Legal Service Delivery (LSD)</a></li> <li>• <a href="#">Now Assist for Sourcing and Procurement Operations (SPO)</a></li> </ul>

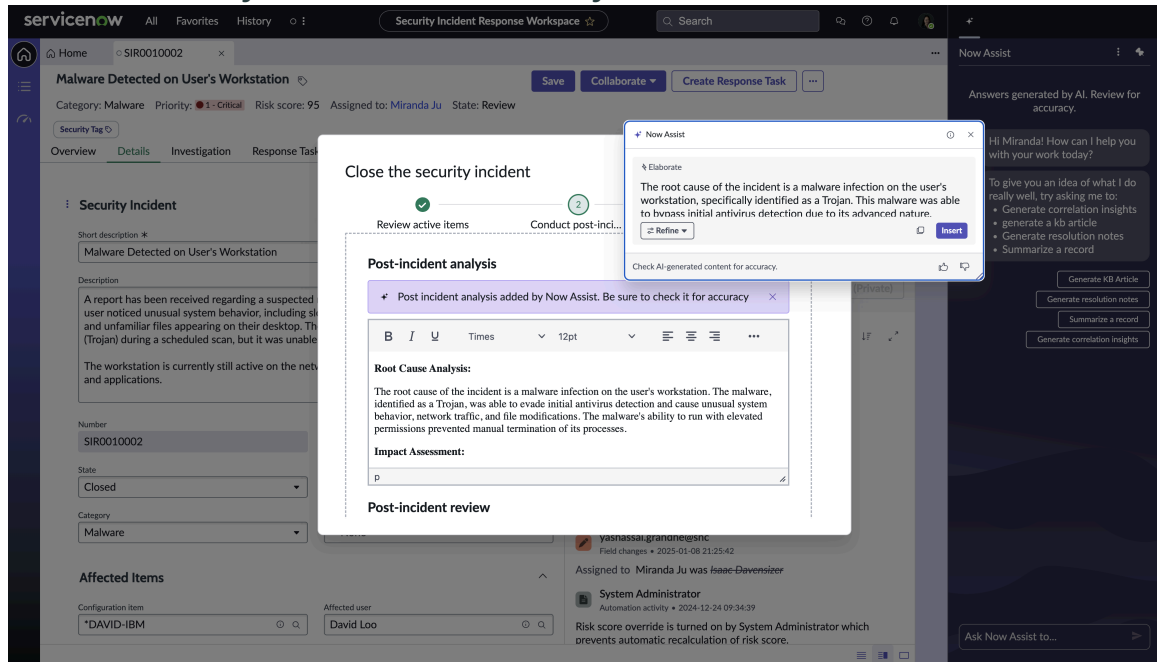
### Post incident analysis

Generate a post-incident analysis that includes a root cause analysis, impact assessment, and lessons learned within the workflow of closing a security incident.

Security managers and analysts can generate post-incident analysis data from within the security incident workflow from the following locations:

- The Security Incident Response Workspace.
- Security incident records in UI16.

### Post-incident analysis in the Close the security incident modal



For more information, see [Generate a post-incident analysis for a security incident with Now Assist for Security Incident Response](#).

### Availability

This skill is available in the workflow and product listed below.

### Now Assist products and workflows

Workflow	Product
Technology	<a href="#">Now Assist for Security Incident Response</a>

## Sentiment Analysis

Apply Now Assist Sentiment Analysis to each customer interaction to assess user satisfaction and facilitate informed decision-making. The customer communications can be related to cases, records, or problems. Use this skill to prioritise this data based on the overall sentiment and reasoning, by providing generative AI analysis to the fulfillers and enabling them identify cases that require urgent attention.

Now Assist Sentiment Analysis uses large language model (LLM) to produce sentiment evaluations for new or active case/problem records, and it also illustrates sentiment trends over time. Fulfillers are provided with a summary to comprehend the reasons behind the sentiments without reading all the notes and comments. This enables them to prioritize and focus on urgent cases and issues based on the severity of the sentiments, and improve the overall customer experience.

## Now Assist Sentiment Analysis refresh

The skill saves the analysis, trends and reasoning on the record and doesn't refresh with page refreshes. LLM calls are made for sentiment refresh when there is a relevant change on that specific record or updates to the input fields defined in the input prompt. The input fields are defined in Now Assist Admin, as parameters for analysing the customer interaction. Now Assist Sentiment Analysis will apply only to records that meet the filter conditions specified in the input templates, as well as those defined in the availability section of Now Assist Admin.



## Now Assist Sentiment Analysis in list view

The saved values for sentiment, sentiment trends and the underlying reasonings will also be available in the list view.

While you're in a list view, if any underlying records have been updated and a new LLM call is triggered, you will receive a notification prompting you to refresh the page. This helps you to stay informed and engaged with the latest information. Although the list can be long, the visual notification will focus only on the records currently visible on your page, which helps you maintain relevance as you work through the sentiment evaluations.

## Generating Now Assist Sentiment Analysis



You can generate sentiment analysis for the following products.

- [Analyze sentiments in Now Assist for Customer Service Management \(CSM\)](#) 
- [Analyze sentiments in Now Assist for IT Service Management \(ITSM\)](#) 

## Availability

This skill is available in the workflow and product listed below.

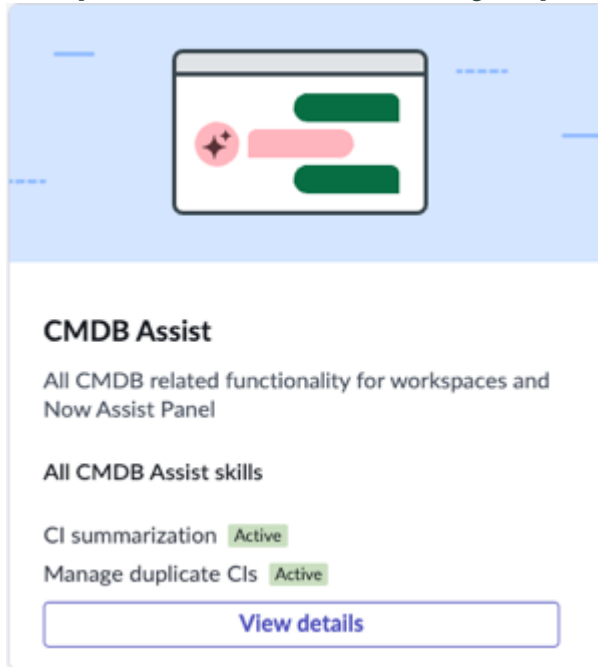
### Now Assist products and workflows

Workflow	Product
Technology	<a href="#">Now Assist for IT Service Management (ITSM)</a> 
Customer	<a href="#">Now Assist for Customer Service Management (CSM)</a> 

## Now Assist conversational solutions skills

Conversational skills use Now Assist in Virtual Agent to help users accomplish tasks in a natural language chat.

### Example CMDB skills tile with Manage duplicate CIs conversational skill



## Agile story generation

Breakdown epics and features into stories using Now Assist in the Enterprise Agile Planning (EAP) workspace. Product managers or scrum leads can save time and reduce functional gaps in stories, thereby improving overall quality and productivity.

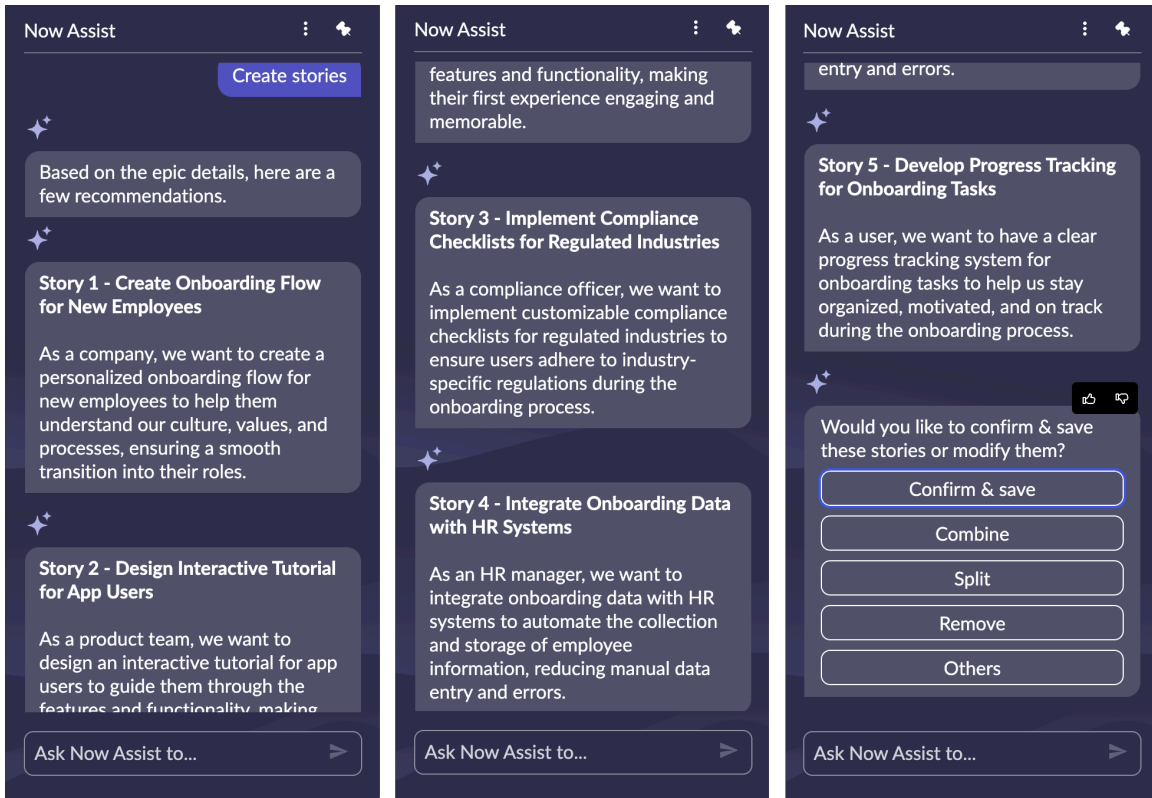
[https://player.vimeo.com/video/1062571498?h=68874331c2&badge=0&autoplay=0&player\\_id=0&app\\_id=58479](https://player.vimeo.com/video/1062571498?h=68874331c2&badge=0&autoplay=0&player_id=0&app_id=58479)

Product managers and scrum leaders typically spend time and effort in creating detailed stories for their epics and features. More often than not, stories that are created manually tend not be detailed enough, lacking context and complete information. This lack of information impacts the understanding of the requirements by the scrum teams, thereby affecting the quality of the product. Also, manually creating detailed individual stories for each epic takes a significant amount of time.

The Now Assist for Strategic Portfolio Management (SPM) application includes the skills of generative AI that enable product managers and scrum leads automatically generate context-based stories for their work items. Through the Agile story generation skill, Now Assist provides story recommendations for an epic or a feature using the name, description, Docs content, and any existing stories. Based on the recommendations, you can ask Now Assist to perform one of the following:

- Split a story recommendation into multiple stories.
- Combine multiple story recommendations into one story.
- Generate stories according to its original recommendations.
- Remove any story recommendation.

- Suggest modification of details within the recommendation such as story title, persona, and description.
- Suggest generating a new story recommendation.



Thus, using Now Assist to help generate stories for epics and features, product managers can reduce time spent in creating high-quality stories, while improving the productivity of the overall scrum team.

For information on how to use this skill, see [Generate stories from work items in EAP using Now Assist for SPM](#).

### Availability

This skill is available in the workflow and product listed below.

#### Now Assist product and workflow

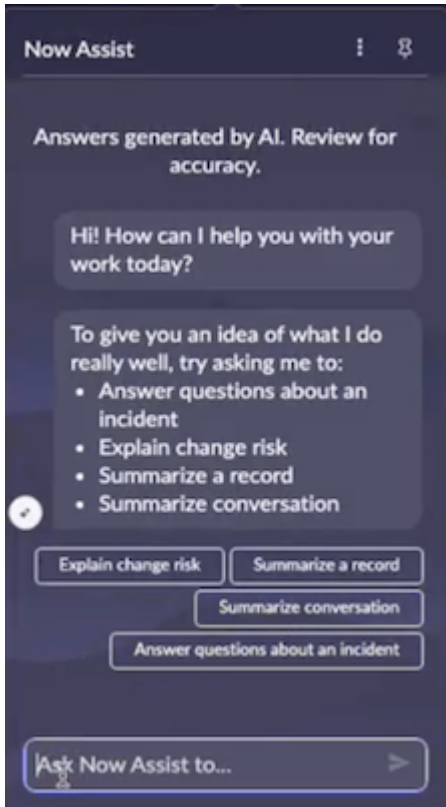
Workflow	Product
Technology	<ul style="list-style-type: none"> <li>• <a href="#">Now Assist for Strategic Portfolio Management (SPM)</a></li> <li>• <a href="#">Enterprise Agile Planning in Strategic Planning</a></li> </ul>

The Agile story generation skill is supported starting with the Xanadu Patch 9 release and with Strategic Planning v4.5.0. If you are on earlier versions of the Xanadu release, you can activate the Story generation skill, which generates stories from Epics only.

### Incident assist

Use this conversational skill to obtain common incident-related information within the incident record by asking questions in the Now Assist panel.

You can ask questions about an incident by using the Now Assist panel in Core UI and Service Operations Workspace for ITSM.



For more information, see [Ask questions about an incident by using the Now Assist panel](#).

## Availability

This skill is available in the workflows and products listed below.

### Now Assist products and workflows

Workflow	Product
Technology	<a href="#">Now Assist for IT Service Management (ITSM)</a>

### Manage duplicate configuration items (CIs)

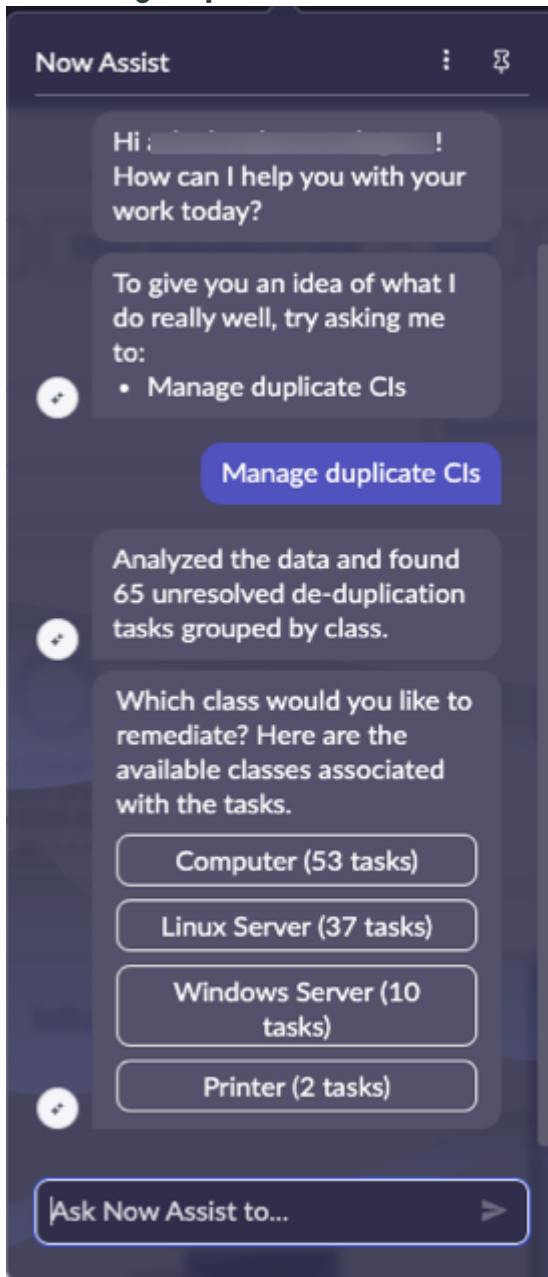
Use Now Assist to review de-duplication tasks, create de-duplication templates, assign tasks to templates, and run the templates. In every step and choice that you make, the manage duplication CIs skill directs you to the next step according to your selection.

In CMDB Workspace you can open the Now Assist panel, and then select **Manage duplicates**. Continue to review responses on the Now Assist panel and to select next actions to manage and remediate duplicate CIs.

As you interact with the skill, it guides you to the next step according to your selection. You can review the root cause analysis for the de-duplication tasks. Root cause analysis groups the de-duplication tasks by IRE identification rules (including criterion attributes) and discovery sources. Use that information to prevent further generation of duplicate CIs.

For more information, see [Using Now Assist for CMDB](#).

The Manage duplicate CIs skill in the Now Assist panel



**Availability**

This skill is available in the workflows and products listed below.

**Now Assist products and workflows**

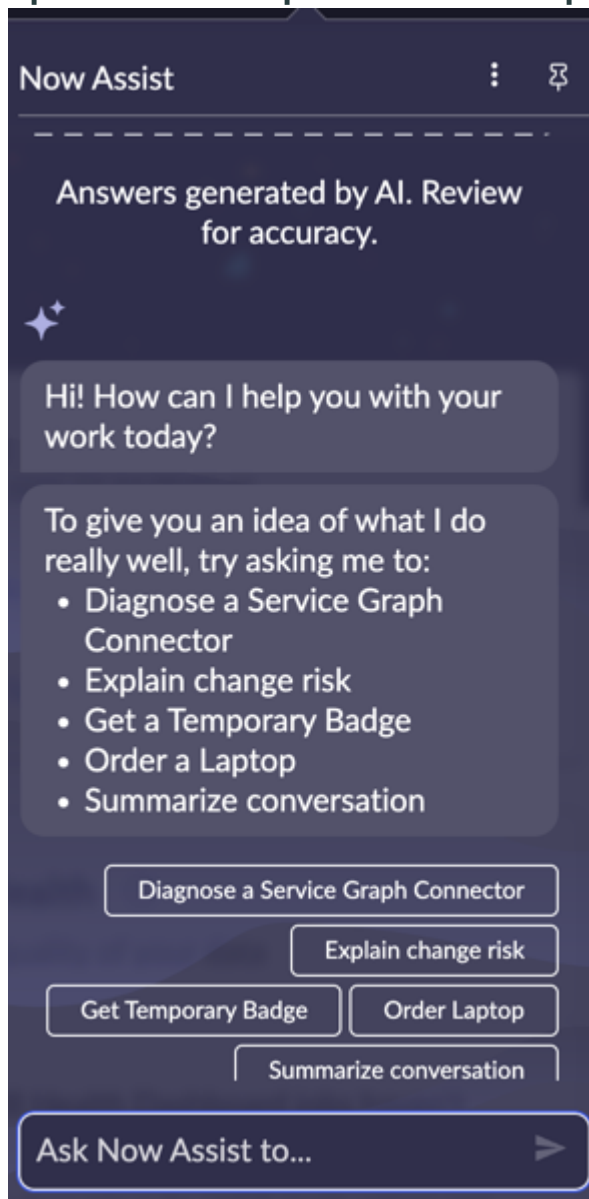
Workflow	Product
Technology	<a href="#">Now Assist for Configuration Management Database (CMDB)</a>

**Service Graph Connector diagnosis**

Use this skill to diagnose a failed import set that is associated with a Service Graph Connector. Now Assist returns a summary of the errors with the recommendations to resolve the issues.

The Service Graph Connector diagnosis skill enables you to identify and resolve issues faster, reducing manual effort and debugging time.

Open the Now Assist panel in CMDB Workspace




For more information, see [Use Now Assist to diagnose a Service Graph Connector issue](#) .

**Availability**

This skill is available in the workflow and product listed below.

**Now Assist products and workflows**

Workflow	Product
Technology	<a href="#">Service Graph Connectors</a> 

**Task generation in Collaborative Work Management**

Save time and significantly reduce manual effort by generating tasks from the context of a Doc page and add them to the Board in Collaborative Work Management (CWM) using Now Assist.

[https://player.vimeo.com/video/1079861522?h=a6bccdb612&badge=0&autoplay=0&player\\_id=0&app\\_id=58479](https://player.vimeo.com/video/1079861522?h=a6bccdb612&badge=0&autoplay=0&player_id=0&app_id=58479)

Creating tasks with detailed descriptions for your CWM Board requires significant time and manual effort. If the tasks aren't detailed enough, it can lead to confusion and misalignment within the team, affecting their understanding of the expected outcomes. To avoid this manual effort and improve time to value, Now Assist can generate tasks for your Board using the information in your Docs. This way, you can ensure clear and comprehensive task descriptions, allowing you to focus more on execution and less on the administrative work.



Based on the recommendations, you can ask Now Assist to perform one of the following:

- Generate tasks according to the initial recommendations.
- Split a task recommendation into multiple tasks.
- Combine multiple task recommendations into one task.
- Remove any task recommendation.

## Availability

This skill is available in the workflow and product listed below.

### Now Assist product and workflow

Workflow	Product
Technology	<ul style="list-style-type: none"> <li>• <a href="#">Now Assist for Collaborative Work Management (CWM)</a> </li> <li>• <a href="#">Collaborative Work Management</a> </li> </ul>

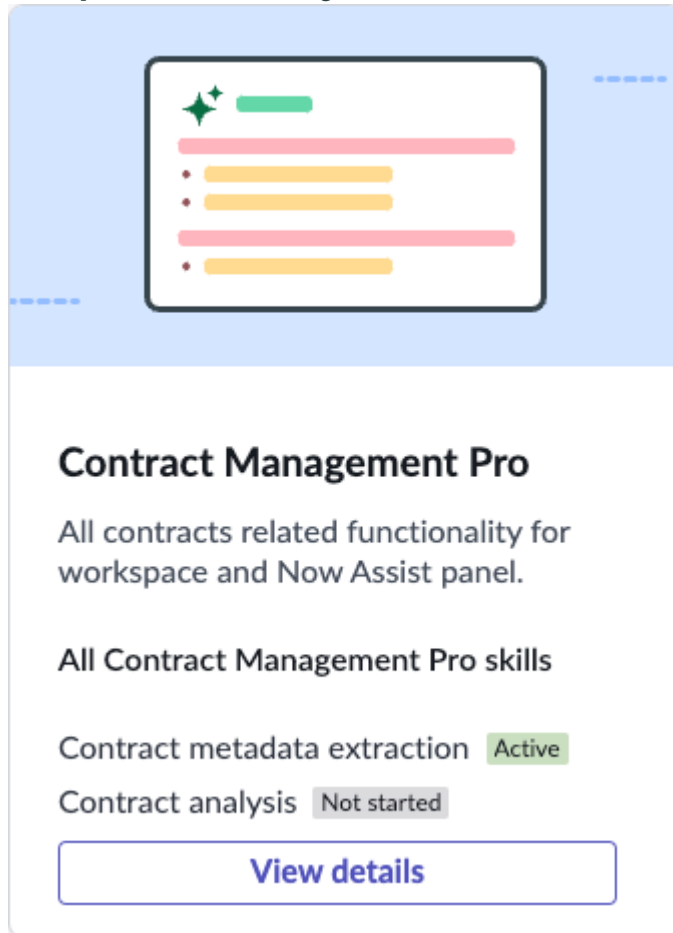
The Task generation skill is supported starting with the Xanadu Patch 9 release and Collaborative Work Management v6.0.0.

### Now Assist content management skills

These Now Assist skills provide generative AI capabilities for chat features on the ServiceNow AI Platform.

The Contracts skills in the Employee workflow area are available when you install Now Assist for Legal Service Delivery (LSD). These skills work in Contract Management Pro.

### Example Contract Management Pro skills tile in Now Assist Admin



#### Contract metadata extraction

Use this Now Assist in Contract Management Pro skill to extract metadata from a signed contract and add the information to the mapped fields in the contract repository. You can review the extracted information and make necessary corrections on the DocIntel viewer.

Signed contracts contain key metadata that are useful for administering the contracts. For example, the start date, term, and payment terms are important. Often this metadata is identified manually, which is time-consuming and can result in missing or erroneous data.

Once configured, the contract metadata extraction skill sends signed contracts with the required metadata elements to Document Intelligence for extraction. This metadata is populated in the contract record and sent to the contract administrator for approval.

#### Extracting contract metadata

You can extract contract metadata in the following product.

[Metadata extraction using Now Assist in Contract Management](#) 

#### Availability

This skill is available in the workflow and product listed below.

### Now Assist products and workflows

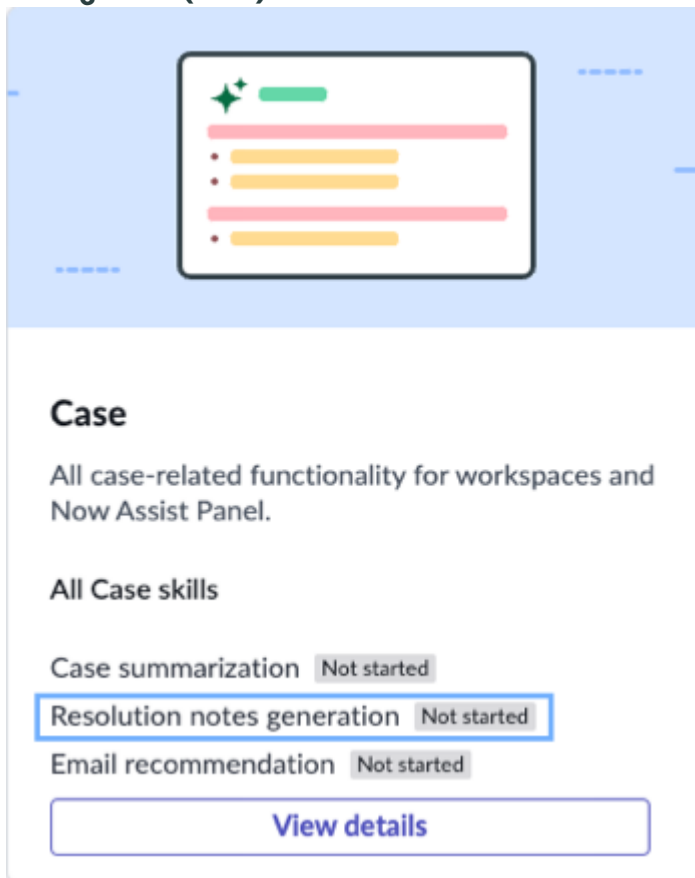
Workflow	Product
Employee	<p><a href="#">Now Assist in Contract Management</a> is available in the following products:</p> <ul style="list-style-type: none"> <li>• <a href="#">Now Assist for Legal Service Delivery (LSD)</a></li> <li>• <a href="#">Now Assist for Sourcing and Procurement Operations (SPO)</a></li> </ul>

### Now Assist content generation skills

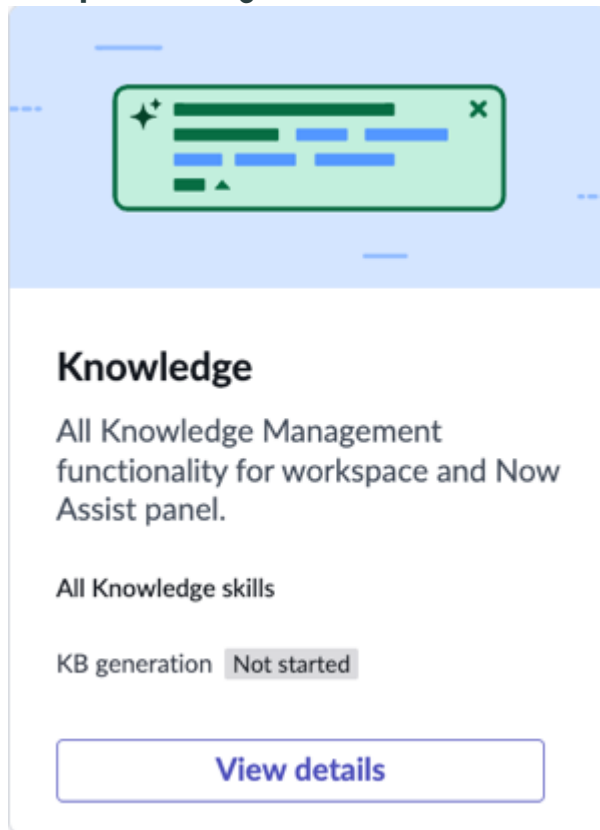
These Now Assist skills can generate resolution notes and knowledge articles.

#### Example skill tiles with content recommendations and resolution notes generation

#### Example resolution notes generation skill in Now Assist for Customer Service Management (CSM)



### Example Knowledge tile in Now Assist Admin




#### KB generation

You can quickly write drafts of knowledge articles based on cases, incidents, or work order tasks with Now Assist. Generating article content with artificial intelligence (AI) enables agents to write efficiently as they address common user concerns.

[https://player.vimeo.com/video/1009282001?h=fc84a7cd87&badge=0&autopause=0&player\\_id=0&app\\_id=58479](https://player.vimeo.com/video/1009282001?h=fc84a7cd87&badge=0&autopause=0&player_id=0&app_id=58479)

#### KB generation overview

Now Assist can create drafts of knowledge articles on how to resolve a case, an incident, or a work order task for agents to review and edit before publishing. Articles can be created in CSM Configurable Workspace, Service Operations Workspace for ITSM, Agent Workspace for HR Case Management, CSM/FSM Configurable Workspace, classic environment, or in the Now Assist panel. As a prerequisite to creating knowledge articles, it is recommended to activate the appropriate plugins for the respective workspaces. For example, to create a knowledge article from an incident, activate the [Activate KCS Integration for Incident Management](#)  plugin (com.snc.incident.knowledge).

The following fields are used as inputs:

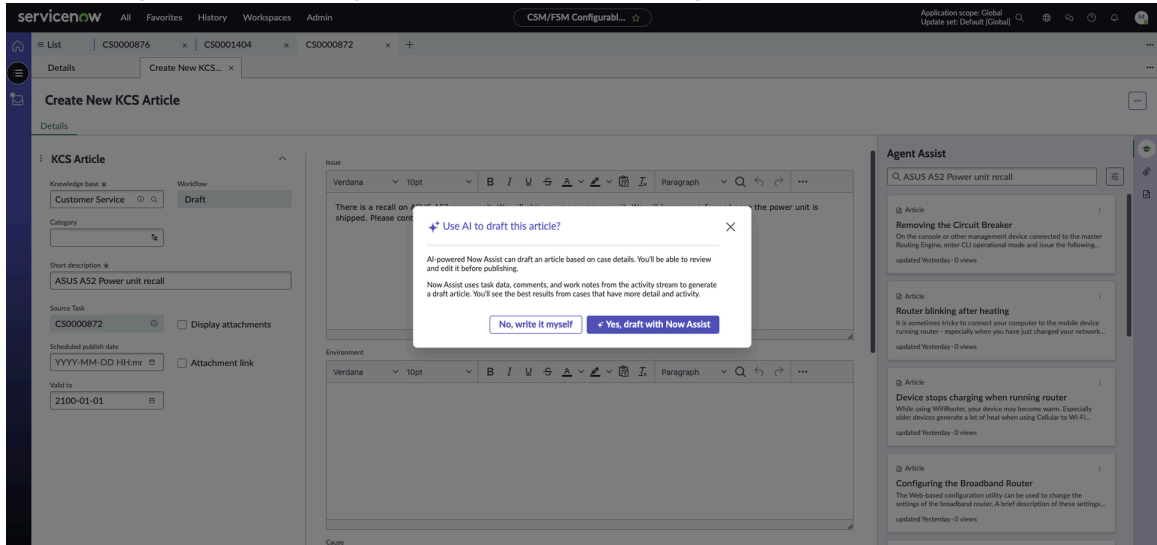
- Short description
- Description
- Resolution notes
- Close notes
- Work notes
- Comments

**Note:** Now LLM Service is the provider for this Now Assist skill.

Once a Now Assist application is installed, the agent will have the option to use Now Assist to generate a knowledge article on cases that are resolved and closed, incidents that are resolved, as well as work order tasks that are closed complete or closed incomplete.

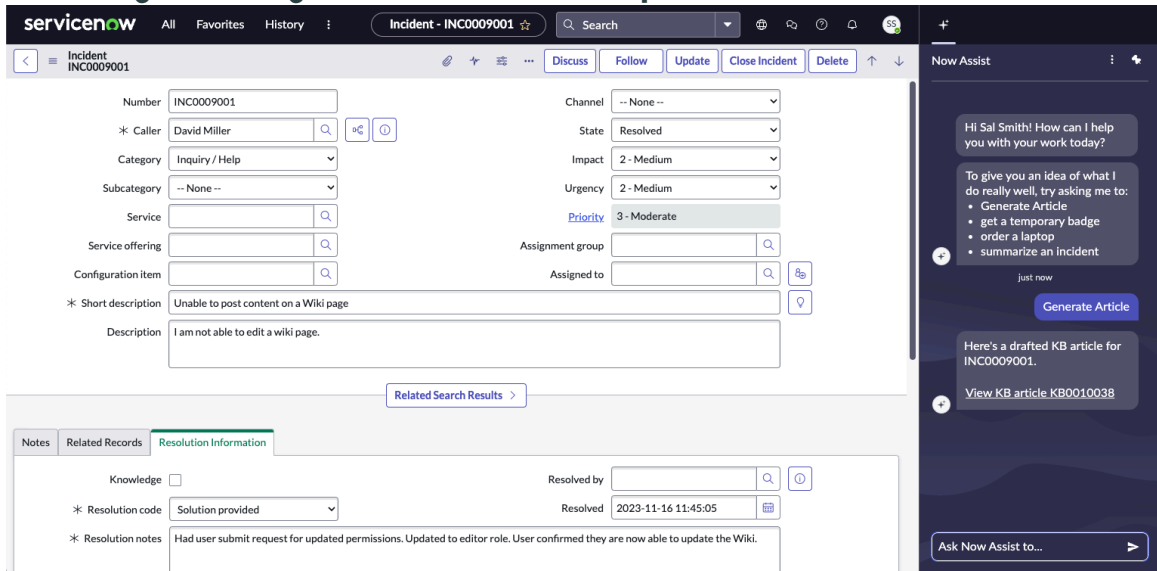
In this example, an agent working in the CSM Configurable Workspace has the option to use Now Assist to draft an article.

### Generating a new knowledge article in the CSM Configurable Workspace



In this example, an agent has used the Now Assist panel to start a draft from a resolved incident in the classic environment. After an article draft is generated, the agent clicks a link to review, edit, and publish the article.

### Generating a knowledge article in the Now Assist panel



## Generating knowledge articles

You can generate knowledge articles in the following products.

- [Generate a knowledge article from the CSM Configurable Workspace and classic environment with Now Assist](#)
- [Generate a knowledge article from the Service Operations Workspace for ITSM and classic environment by using Now Assist](#)
- [Generate a knowledge article from HR Agent Workspace with Now Assist](#)
- [Generate a knowledge article from multiple cases](#)
- [Generate a knowledge article from the CSM/FSM Configurable Workspace and classic environment with Now Assist](#)
- [Generate a knowledge article from the classic environment with Now Assist](#)
- [Generate a knowledge article from the Now Assist panel](#)

## Availability

This skill is available in the workflows and products listed below.

### Now Assist products and workflows

Workflow	Product
Technology	<a href="#">Now Assist for IT Service Management (ITSM)</a>
Customer	<a href="#">Now Assist for Customer Service Management (CSM)</a>
	<a href="#">Now Assist for Field Service Management (FSM)</a>
Employee	<a href="#">Now Assist for HR Service Delivery (HRSD)</a>

### Resolution notes generation

You can quickly learn the details of how an interaction was resolved by generating and reading the Now Assist resolution notes.

Generating the resolution notes creates a condensed version of the incident's resolution notes that are generated by Now Assist. The generated resolution notes can provide information to other agents who might encounter similar interactions.

**Note:** Now LLM Service is the provider for this Now Assist skill.

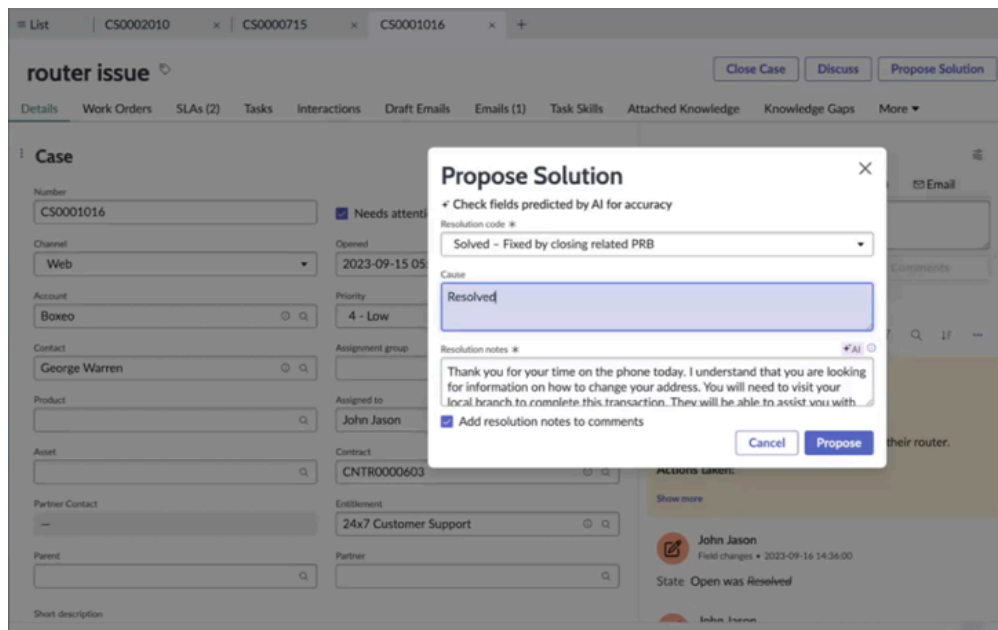
You can use the Now Assist icon to perform these actions:

- Generate resolution notes that are based on the information in the incident.
- Refine the resolution notes by elaborating or shortening them.

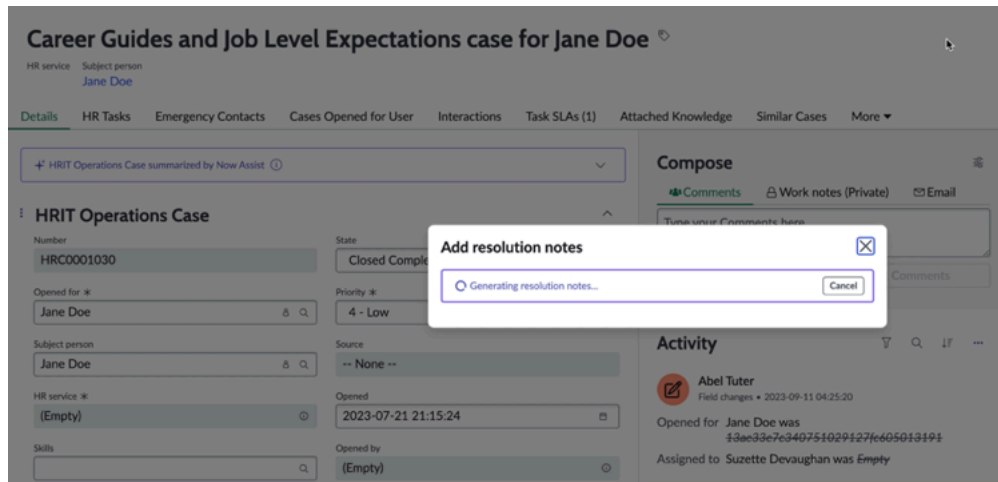
Once the resolution notes are generated, they can be posted by the user closing the ticket. When the notes are posted, they are attributed to the user that reviewed and accepted the AI-generated content.

The following examples show Now Assist generated resolution notes in different products.

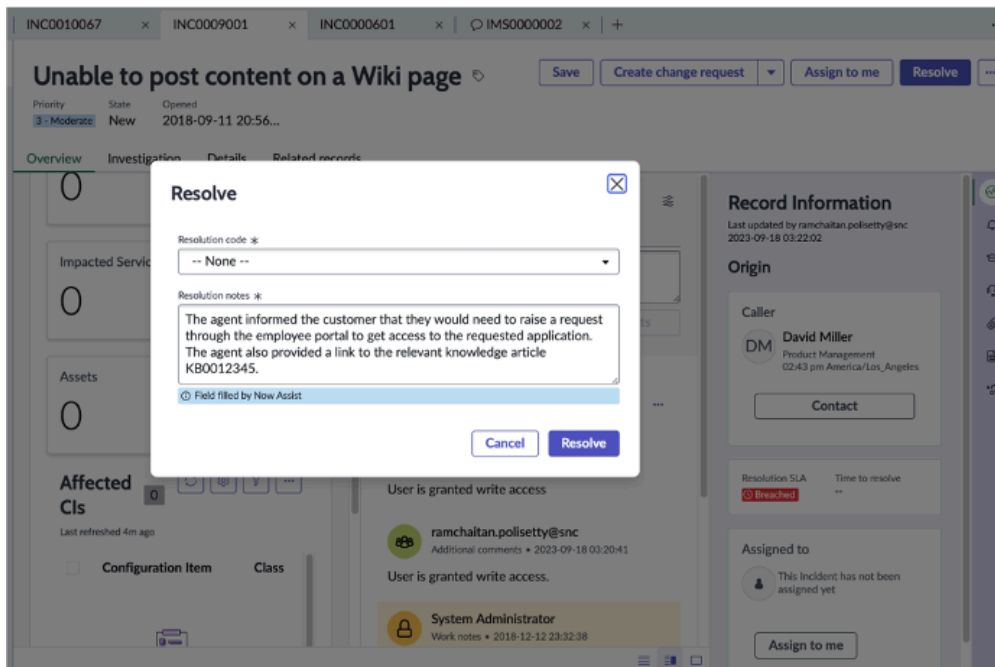
### Generated resolution notes in Now Assist for CSM



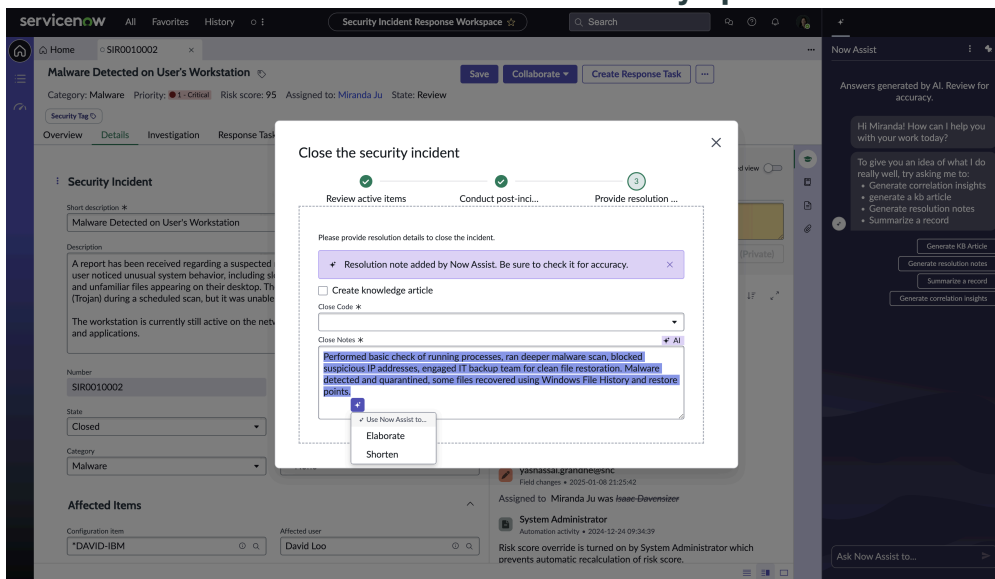
### Generated resolution notes in Now Assist for HRSD



## Generated resolution notes in Now Assist for ITSM



## Generated resolution notes in Now Assist for Security Operations



## Generating resolution notes

You can generate resolution notes for the following products.

- Generate the resolution notes for a case by using Now Assist for Customer Service Management (CSM) [🔗](#)
- Generate the resolution notes for a case by using the Now Assist for HR Service Delivery (HRSD) [🔗](#)
- Generate the resolution notes for an incident by using Now Assist for IT Service Management (ITSM) [🔗](#)
- Generate closure notes for a security incident with Now Assist for Security Incident Response [🔗](#)
- Generate the resolution notes for a service problem case using Now Assist for Telecommunications, Media and Technology (TMT) [🔗](#)

## Availability

This skill is available in the workflows and products listed below.

### Now Assist products and workflows

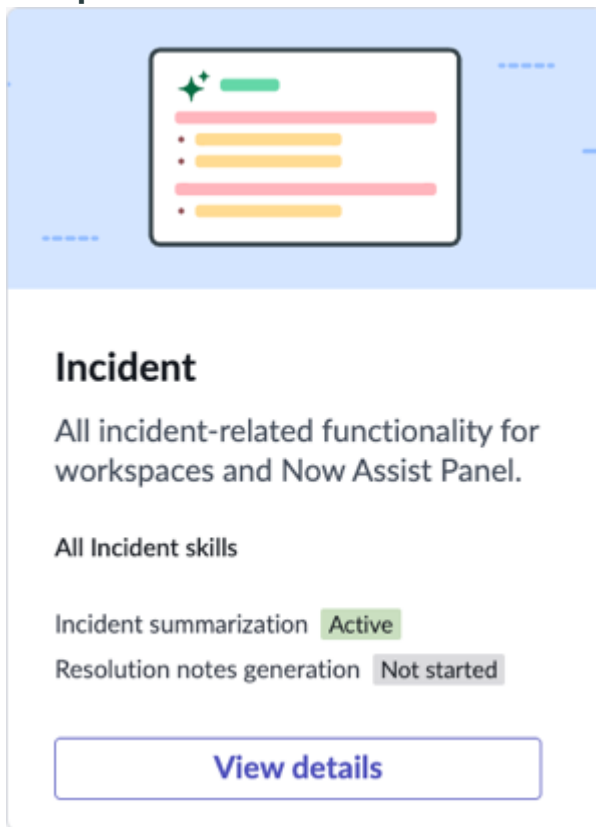
Workflow	Product
Technology	<a href="#">Now Assist for IT Service Management (ITSM)</a> ↗
	<a href="#">Now Assist for Security Incident Response</a> ↗
Customer	<a href="#">Now Assist for Customer Service Management (CSM)</a> ↗
	<a href="#">Now Assist for Telecommunications, Media and Technology (TMT)</a> ↗
Employee	<a href="#">Now Assist for HR Service Delivery (HRSD)</a> ↗

### Now Assist content and record summary skills

Now Assist skills can provide generative AI summarization capabilities for cases, incidents, and other records on the ServiceNow AI Platform.

**Note:** If you're looking to add summarization to custom or child tables, you can create a new skill in Now Assist Skill Kit. See [Create a skill](#).

### Example Incident skills tile in Now Assist Admin



## Example summarization skill tile for Now Assist for Sourcing and Procurement Operations (SPO)



### Summarization for fulfillers

Summarize requests and their related records to quickly update fulfillers on their progress and open action items.

#### All Summarization for fulfillers skills

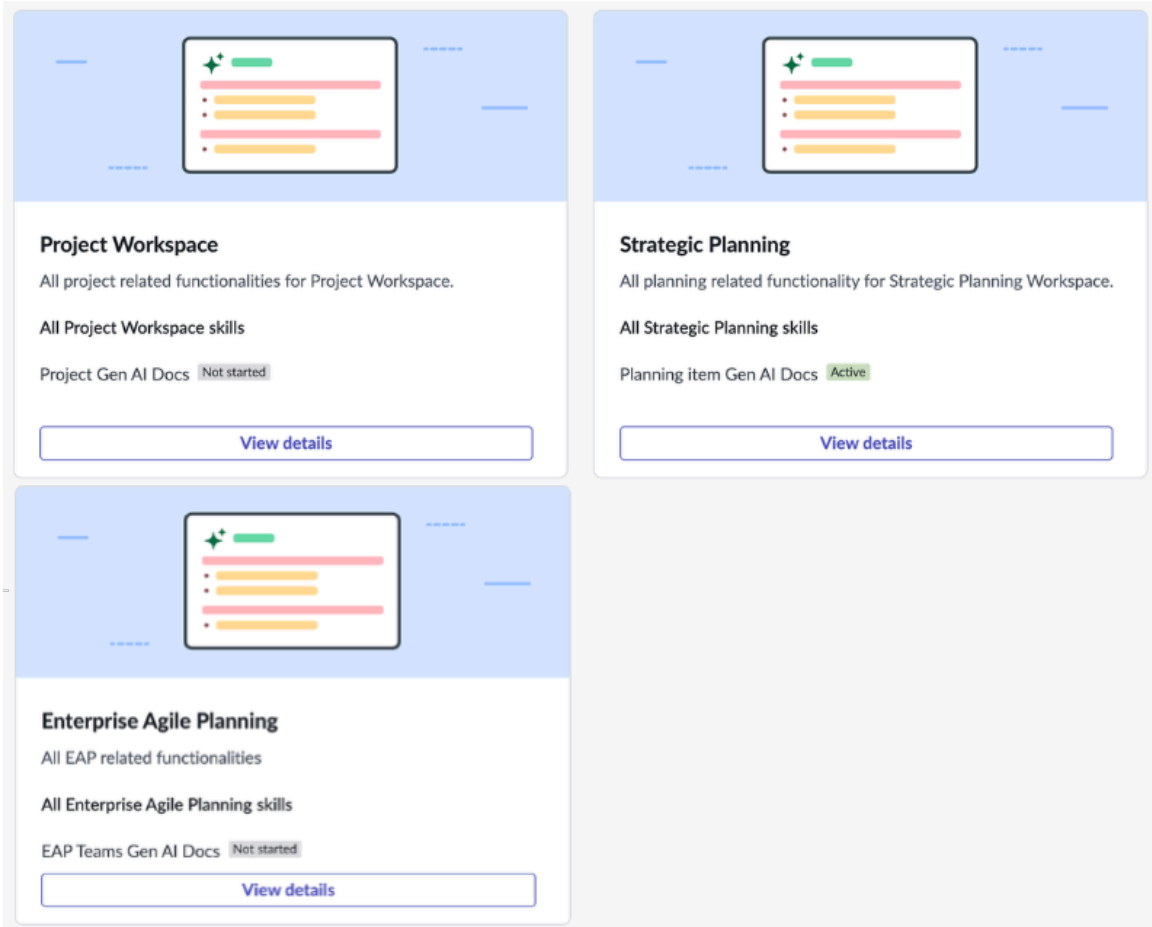
Sourcing request summarization Active

Purchase requisition summarization Active

Procurement case summarization Active

[View details](#)

## Example content summary skill tiles for Now Assist for Strategic Portfolio Management (SPM)



### Record summarization

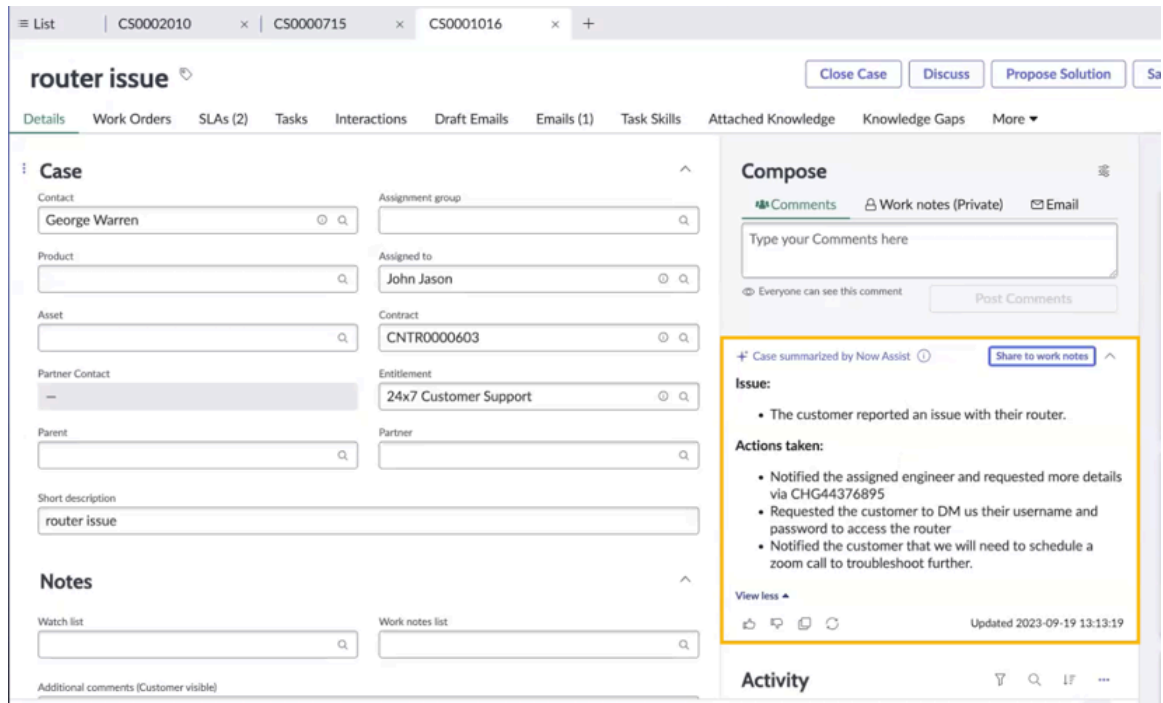
You can quickly learn the details of a record from record summarization generated by Now Assist. By reading a Record summarization, you can get details about a case or incident and can save your requester from having to repeat the same information that the requester has already provided.

Now Assist can generate a summary of records only if the input data contains at least 50 words in the relevant fields. If you have the latest version, there is no minimum word count requirement.

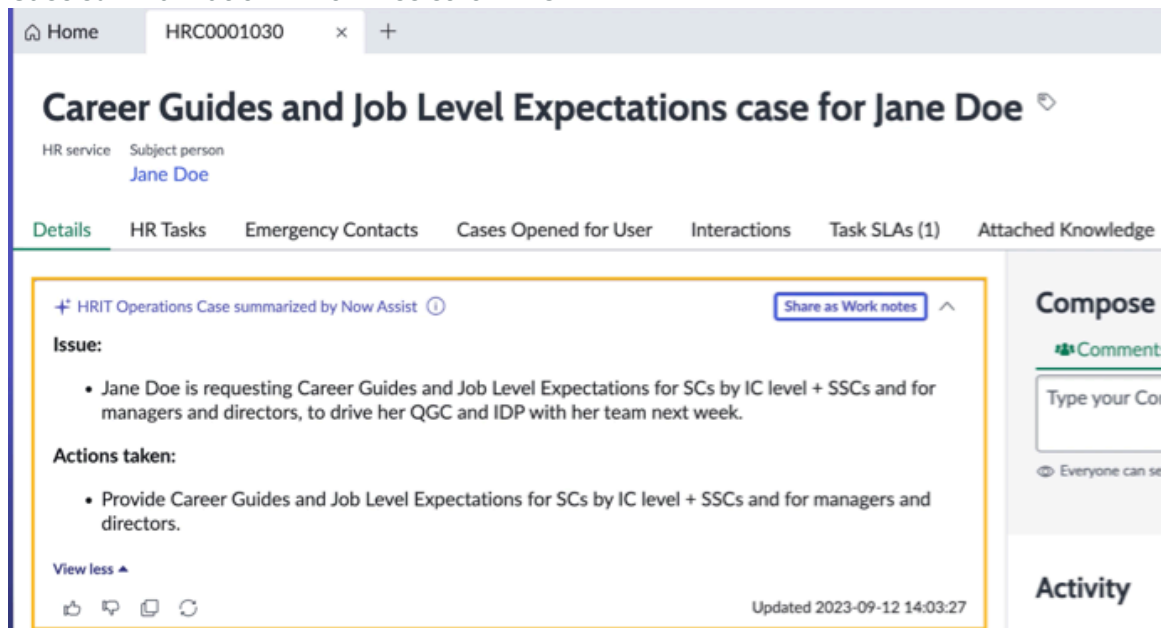
**Note:** The ServiceNow Large Language Model (Now LLM Service) is the provider for this Now Assist skill.

In the following examples Now Assist generated a case/incident summarization in Now Assist for Customer Service Management (CSM), Now Assist for HR Service Delivery (HRSD), and Now Assist for IT Service Management (ITSM). The incident summarization is displayed in the Active Chat pane.

### Case summarization in Now Assist for CSM



### Case summarization in Now Assist for HRSD

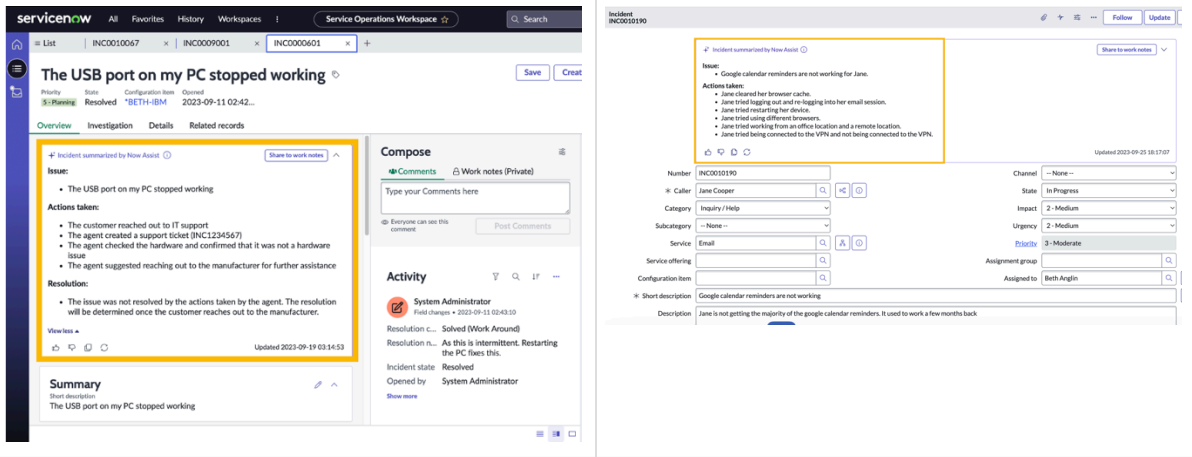


In the Now Assist for ITSM example, Now Assist generated an incident summarization that can be displayed in Next Experience or Core UI.

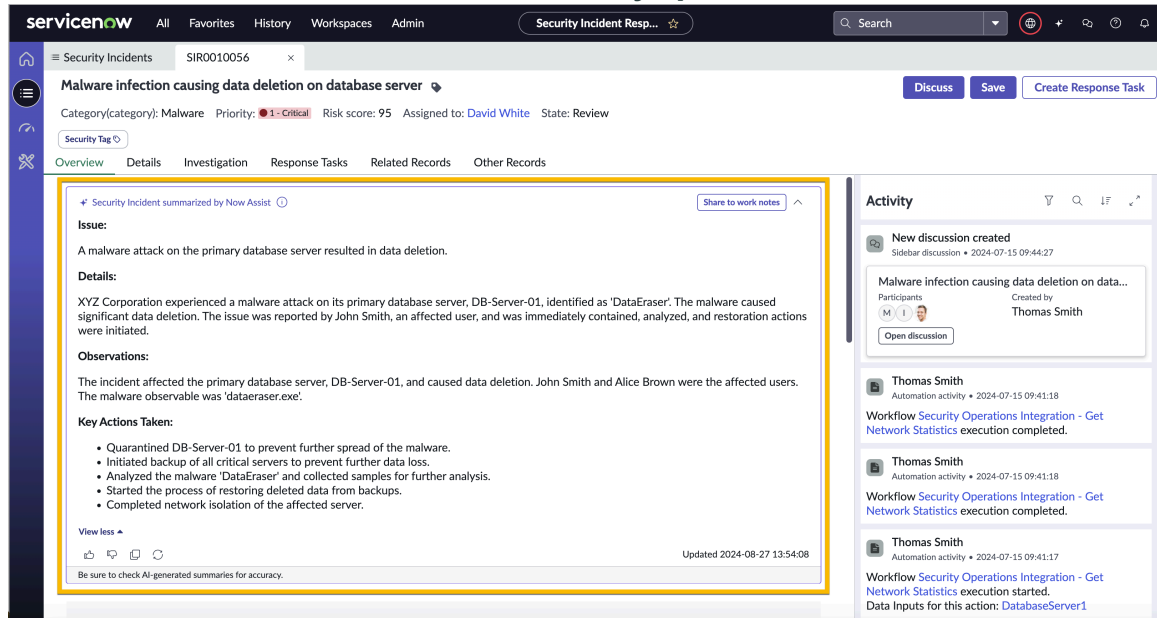
### Incident summarization in Now Assist for ITSM

Next Experience	Core UI
-----------------	---------

## Incident summarization in Now Assist for ITSM (continued)



## Incident summarization in Now Assist for Security Operations



## Generating a case or incident summarization

You can generate a case or incident summarization for the following products:

- Summarize a case by using Now Assist for Customer Service Management (CSM) [🔗](#)
- Summarize a case by using Now Assist for Financial Services Operations (FSO) [🔗](#)
- Summarize a safety incident by using Now Assist for Health and Safety [🔗](#)
- Summarize a case by using Now Assist for HR Service Delivery (HRSD) [🔗](#)
- Summarize an incident by using Now Assist for IT Service Management (ITSM) [🔗](#)
- Summarize a legal request or matter by using Now Assist for Legal Service Delivery (LSD) [🔗](#)
- Summarize a security incident with Now Assist for Security Incident Response [🔗](#)
- Summarize a service problem case using Now Assist for Telecommunications, Media and Technology (TMT) [🔗](#)

## Availability

This skill is available in the workflows and products listed below.

### Now Assist products and workflows

Workflow	Product
Technology	<a href="#">Now Assist for IT Service Management (ITSM)</a> <a href="#">Now Assist for Security Incident Response</a>
Customer	<a href="#">Now Assist for Customer Service Management (CSM)</a> <a href="#">Now Assist for Financial Services Operations (FSO)</a> <a href="#">Now Assist for PSDS</a> <a href="#">Now Assist for Telecommunications, Media and Technology (TMT)</a>
Employee	<a href="#">Now Assist for Health and Safety</a> <a href="#">Now Assist for HR Service Delivery (HRSD)</a> <a href="#">Now Assist for Legal Service Delivery (LSD)</a>
Finance & Supply Chain	<a href="#">Now Assist for Accounts Payable Operations (APO)</a> <a href="#">Now Assist for Supplier Lifecycle Operations (SLO)</a> <a href="#">Now Assist for Sourcing and Procurement Operations (SPO)</a>

### Change request summarization

Use the change request summarization skill to quickly capture important details of the change request, including the current status.

A change request summary provides you with a concise summary of a change request. The summary is based on the change request state and is generated from the information in the related table.

### Change request summary in Now Assist for IT Service Management (ITSM)

Change Request summarized by Now Assist

**Objective:**  
Upgrade Service Operations Workspace (ITSM Applications plugin from version 2.0.1 to 3.1.0. Justification not provided.

**Plan:**  
Test plan: Validate the upgrade changes (CTASK0164389). Backout plan: Fix forward, create a new change for modified steps in case of issues. Implementation plan: Upgrade the plugin from 2.0.1 to 3.1.0 (CTASK0164387).

**Risk:**  
Low risk involved in the change request. The conflict status of this change request is No Conflict. There are no scheduling conflicts associated with this change request.

**Affected CIs:**  
No configuration items are affected.

**Impacted services:**  
No services are impacted.

Updated 2024-10-15 12:25:40

Be sure to check the AI-generated summary for accuracy.

Workflow: New → **Assess** → Authorize → Scheduled → Implement → Review → Closed → Canceled

Number: CHG0502442  
 Requested by:   
 Category: Internal Application  
 Service:

Model:   
 Type: Normal  
 State: Assess  
 Change is waiting for approval

## Generating a change request summary

You can generate a change request summary in the following products.

[Summarize a change request by using Now Assist for IT Service Management \(ITSM\)](#) 

### Availability

This skill is available in the workflow and product listed below.

#### Now Assist products and workflows

Workflow	Product
Technology	<a href="#">Now Assist for IT Service Management (ITSM)</a> 

### Chat summarization

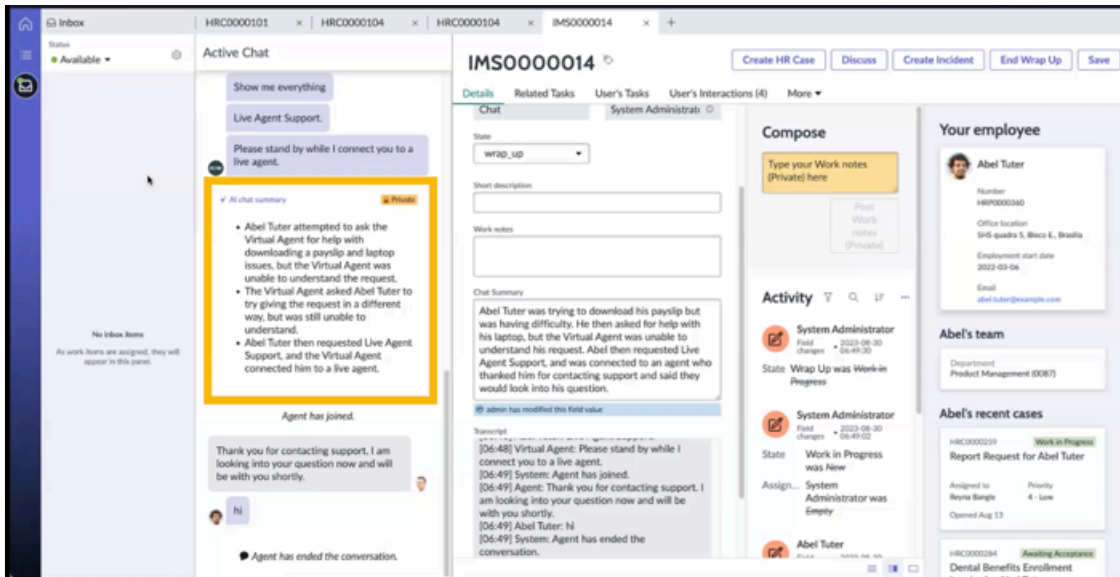
You can quickly learn the details of a chat from a chat summarization generated by Now Assist. By reading a chat summarization, you can get enough details about the chat and can save your requester from having to repeat the same information that the requester already provided to Virtual Agent.

The chat summarization skill generates a condensed version of the conversation between live agents or between a requester and Virtual Agent. Chat summarizations are generated in these situations:

- When the conversation ends.
- When the agent uses the `/summarize` quick action in Agent Chat.
- When the user selects **Chat Summarization** or enters `summarize chat` in the **Ask Now Assist to** field on the Now Assist panel.
- When an interaction is transferred from Virtual Agent to a live agent and the conversation is at least six lines long. The chat summarization displays in a summary card in the conversation and also populates the interaction's summarization and short description fields.
- A live agent transfers an interaction to a different live agent.

 **Note:** Now LLM Service is the provider for this Now Assist skill.

In the following example, Now Assist generated a chat summarization in Now Assist for HR Service Delivery (HRSD).



## Generating a chat summarization

You can generate a chat summarization in the following products.

- Summarize a chat conversation by using Now Assist for Customer Service Management (CSM) [🔗](#)
- Summarize a chat conversation by using Now Assist for HR Service Delivery (HRSD) [🔗](#)
- Summarize a chat conversation by using Now Assist for IT Service Management (ITSM) [🔗](#)

## Availability

This skill is available in the workflows and products listed below.

### Now Assist products and workflows

Workflow	Product
Technology	<a href="#">Now Assist for IT Service Management (ITSM) <a href="#">🔗</a></a>
Customer	<a href="#">Now Assist for Customer Service Management (CSM) <a href="#">🔗</a></a>
Employee	<a href="#">Now Assist for HR Service Delivery (HRSD) <a href="#">🔗</a></a>

## Configuration item (CI) summarization

Configuration items (CIs) are a central element in numerous features such as Incident Management and Event Management, and they are associated with the respective incidents and alerts. This skill gathers all of those incidents, alerts, security vulnerabilities, and associated records into a single summary.

You can use Now Assist to create a summary of a configuration item and all of the information that is associated with it.

## A configuration item with a Now Assist summary

< Computer \*BETH-IBM
Open in CMDB Workspace Update Delete

**Computer summarized by Now Assist**

**Summary:**  
The configuration item BETH-IBM is of the class Computer, with the most recent discovery date and time being 2007-04-12 16:17:45. There are no associated incidents, problems, or change requests. No services are linked to this CI. The most recent discovery source that updated the CI is unknown.

Updated 2024-10-08 10:59:00

Be sure to check the AI-generated summary for accuracy.

Name

Asset tag

Manufacturer

Asset

Company

Serial number

Model ID

Assigned to

For more information, see [Now Assist for Configuration Management Database \(CMDB\)](#) and [Configuration Management Database \(CMDB\)](#).

### Availability

This skill is available in the workflows and products listed below.

#### Now Assist products and workflows

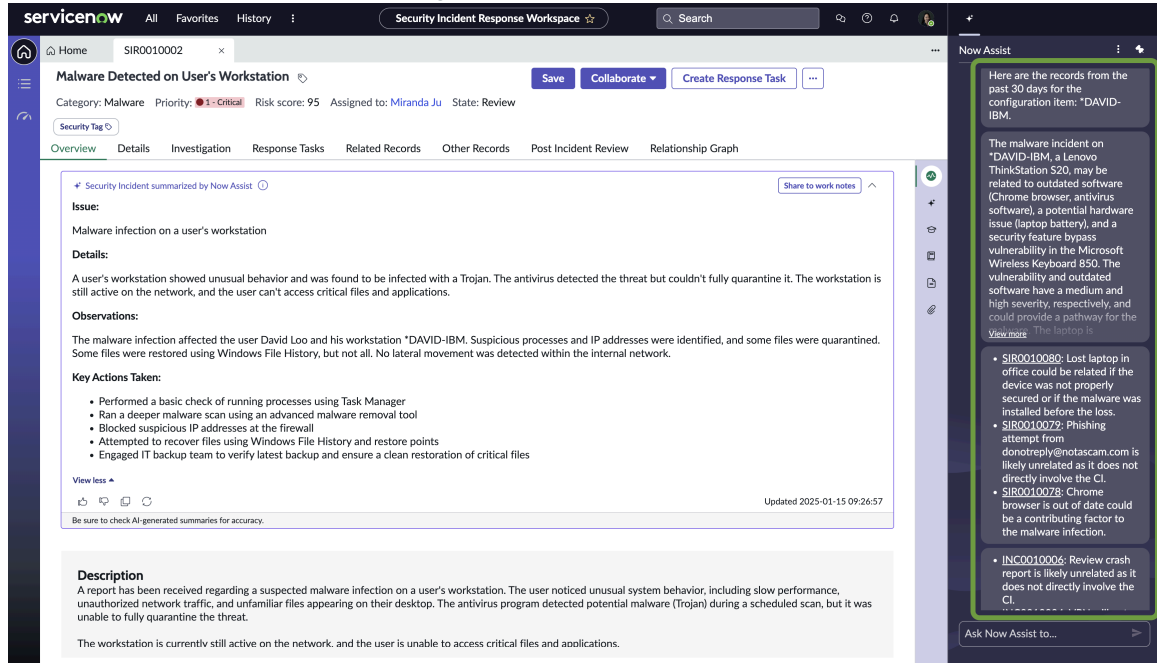
Workflow	Product
Technology	<a href="#">Now Assist for Configuration Management Database (CMDB)</a>

### Correlation insights for records in Security Operations

Generate correlation insights from the Now Assist panel to help you connect past events to the security incidents you are working on.

You can use the related information you generate from existing security incident (SIR), incident (INC), change request (CHG), problem (PRB), and vulnerable item (VIT) records to avoid duplicating your investigation into the security incident you are working on. Get information about affected users, configuration items, and observables that already exists to help you resolve new security incidents more quickly.

## Returned results for correlation insights generated for a configuration item from the Now Assist panel for a security incident



### Generating correlation insights

You can generate correlation insights with [Correlation insights with Now Assist for Security Incident Response](#).

### Availability

This skill is available in the workflow and product listed below.

#### Now Assist products and workflows

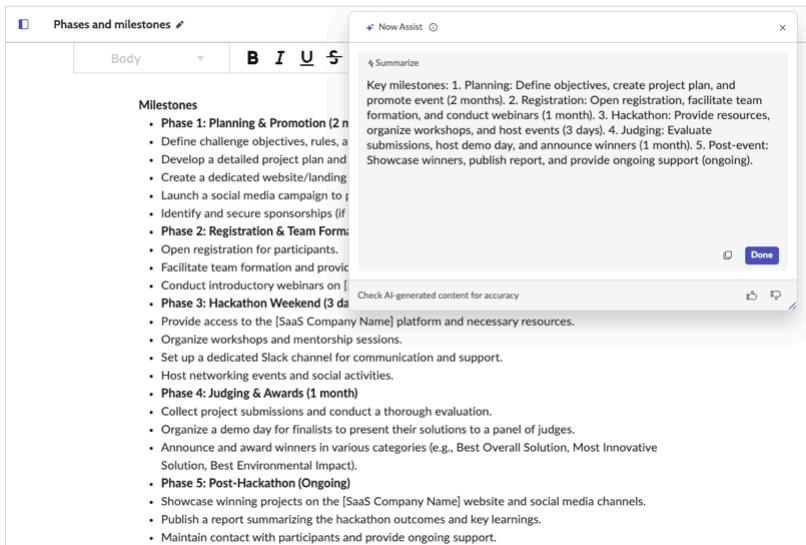
Workflow	Product
Technology	<a href="#">Now Assist for Security Incident Response</a>

### Docs summarization in Collaborative Work Management

Gain insights into the page content by summarizing it or improve content quality by refining it in Docs for Collaborative Work Management (CWM).

Quickly summarize the content of a page in CWM Docs. Whether you're reviewing long documents or preparing for meetings, Doc summarization skill helps you stay informed and efficient.

With the Doc summarization skill, you can also elaborate and shorten selected content, resulting in improved content quality within the CWM Doc pages. Quickly shorten lengthy paragraphs, paraphrase bullet points, or expand on key points ensuring that your content is tailored to specific needs.



The Doc summarization skill is supported starting with the Xanadu Patch 9 release and Collaborative Work Management v6.0.0.

### Feedback and multi feedback summarization

Summarize one or multiple feedbacks received using the feedback summarization or multi feedback summarization that is generated by Now Assist. By reading a feedback summary, your product managers can get enough details to improve the product features, usability, and performance.

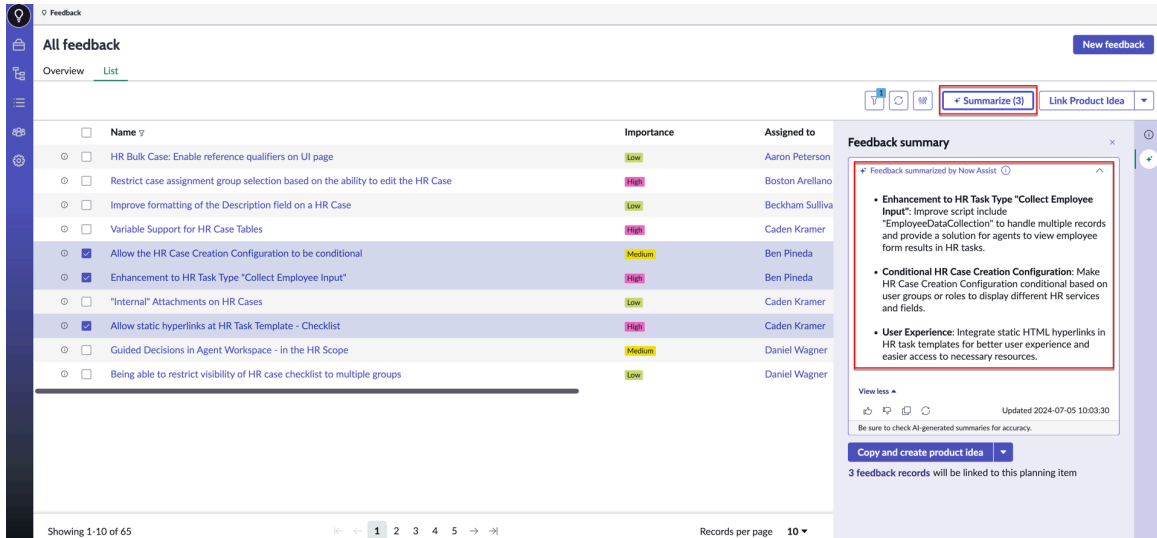
[https://player.vimeo.com/video/1016235687?h=c41764ac08&badge=0&autoplay=0&player\\_id=0&app\\_id=58479](https://player.vimeo.com/video/1016235687?h=c41764ac08&badge=0&autoplay=0&player_id=0&app_id=58479)

The Now Assist for Strategic Portfolio Management (SPM) application includes the skills and features of generative AI that enable product managers to generate a summary from the customer feedback. Your product managers can generate a summary from the name and description of one or multiple feedback record so that they can quickly understand the feedback context. With feedback summarization or multi feedback summarization skill, your product managers can eliminate the need to analyze feedback manually.

**Note:** Now LLM Service is the default provider for feedback summarization or multi feedback summarization Now Assist skill. You can also choose the Azure OpenAI AI provider. For more information, see [Configuring Generative AI Controller](#).

In the following example, Now Assist generated a multi feedback summarization in Now Assist for Strategic Portfolio Management (SPM).

## Multi feedback summarization skill



Now Assist can generate a summary from the feedback only if the feedback has at least 60 words in the fields that are used for the input data. The 60-word minimum optimizes the experience by ensuring that there’s enough information to make a summary. For more information, see [Summarize the feedback by using Now Assist for Strategic Portfolio Management \(SPM\)](#).

### Availability

This skill is available in the workflow and product listed below.

#### Now Assist product and workflow

Workflow	Product
Technology	<a href="#">Now Assist for Strategic Portfolio Management (SPM)</a>

### Fulfiller summarization for Sourcing and Procurement Operations

Summarize requests and their related records to quickly update fulfillers on their progress and open action items.

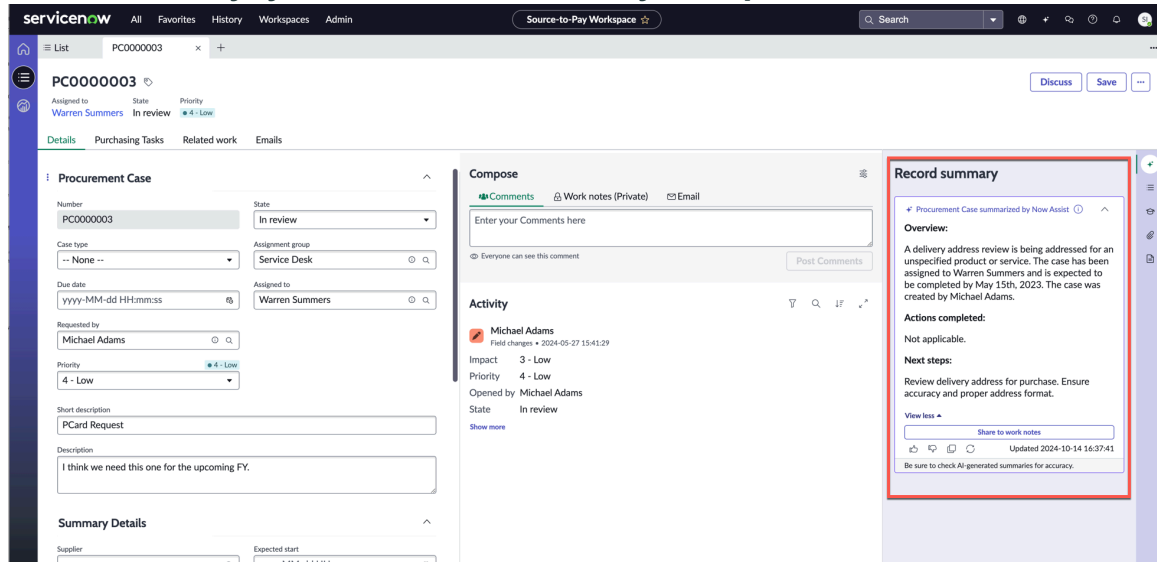
These skills provide a fulfiller with a summary of procurement-related records. This enables them to get visibility into the status of the record in the procurement process and consequently enables them to take the necessary actions quickly.

The fulfillers can get a summary of the following procurement records using the Source-to-Pay Workspace.

- Sourcing request summarization
- Purchase requisition summarization
- Procurement case summarization

For more information, see [Summarize a record by using Now Assist for Sourcing and Procurement Operations \(SPO\) in Source-to-Pay Workspace](#).

## A record summary by Now Assist in Source-to-Pay Workspace



### Availability

This skill is available in the workflow and product as shown in the following table.

#### Now Assist product and workflow

Workflow	Product
Finance & Supply Chain	<a href="#">Now Assist for Sourcing and Procurement Operations (SPO)</a>

### Gen AI Docs

Summarize content in documents using Gen AI Docs that is generated by Now Assist. By reading a document summary, you can get enough details to improve the features, usability, and performance of a product or project.

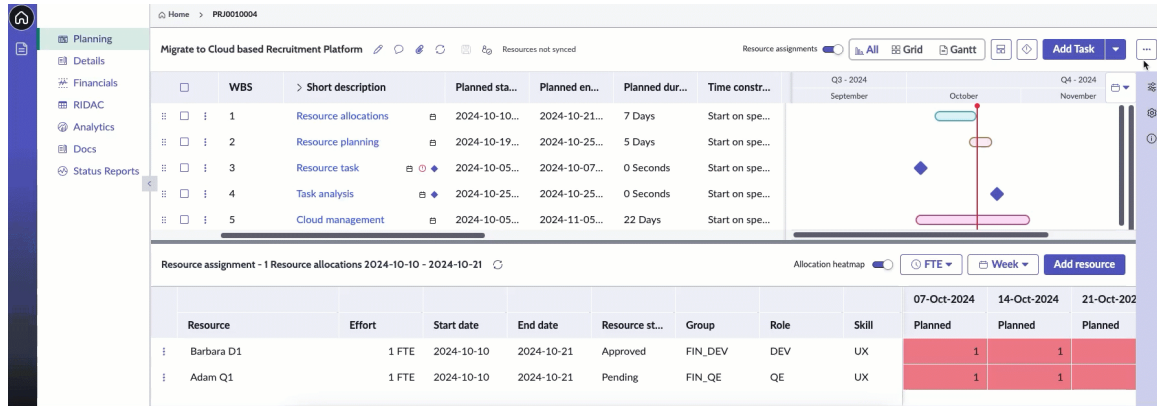
The Now Assist for Strategic Portfolio Management (SPM) application includes the skills and features of generative AI that enable you to generate a summary from the documents. With the Gen AI Docs skills, you can eliminate the need to analyze documents manually. The following skills are available:

- EAP Teams Gen AI Docs
- Planning item Gen AI Docs
- Project Gen AI Docs

**Note:** Now LLM Service is the default provider for these Now Assist skills. You can also choose the Azure OpenAI AI provider. For more information, see [Configuring Generative AI Controller](#).

In the following example, Now Assist generated a document summary in Now Assist for Strategic Portfolio Management (SPM).

## Gen AI Docs skill



For more information, see the following topics:

- [Generate the summary for selected or complete content with Project doc summarization skill in Project Workspace](#)
- [Generate the summary for selected or complete content with Planning item doc summarization skill in Strategic Planning](#)

### Availability

This skill is available in the workflow and product as shown in the following table.

#### Now Assist product and workflow

Workflow	Product
Technology	<a href="#">Now Assist for Strategic Portfolio Management (SPM)</a>

### Email project summary

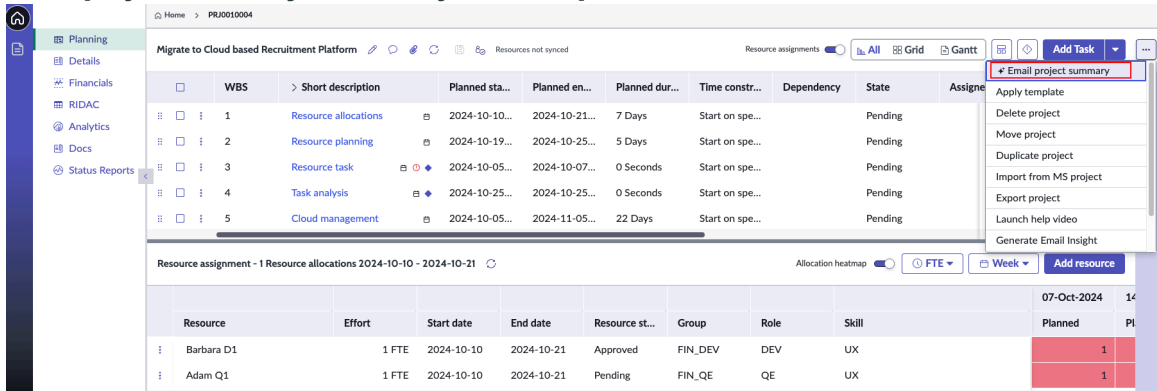
Schedule the project summary emails to help Project Managers track key elements such as milestones, resource assignments, tasks, financials, RIDAC (Risk, Issue, Decision, Action, Change request), and overall project metrics.

The Now Assist for Strategic Portfolio Management (SPM) application includes the skills and features of generative AI that enable Project Managers to generate a summary from the project. Your Project Managers can use the email project summary skill to analyze project changes and prioritize them based on their impact and urgency.

**Note:** Now LLM Service is the default provider for the email project summary Now Assist skill. You can also choose the Azure OpenAI AI provider. For more information, see [Configuring Generative AI Controller](#).

In the following example, the Now Assist generated a summary of the project in Project Workspace.

## Email project summary skill in Project Workspace



For more information, see [Schedule the project summary emails with Email project summary skill](#).

## Availability

This skill is available in the workflow and product as shown in the following table.

### Now Assist product and workflow

Workflow	Product
Technology	<a href="#">Now Assist for Strategic Portfolio Management (SPM)</a>

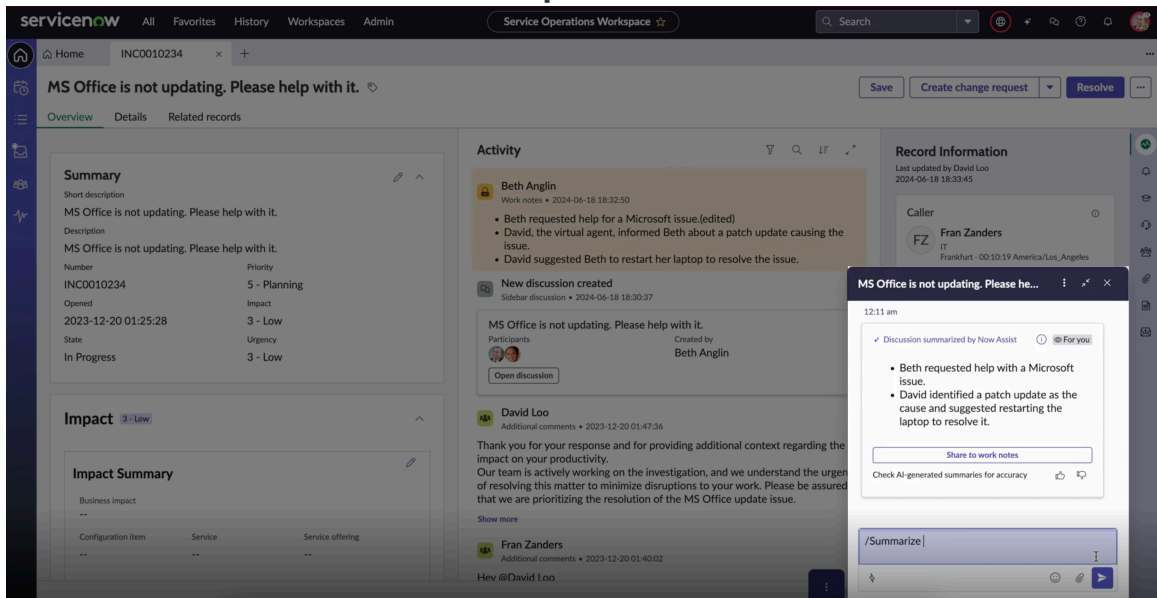
## Sidebar discussion summarization

Generate a summary of Sidebar discussions between agents, requesters, and subject matter experts using the Sidebar discussion summarization skill.

You can use Sidebar discussion summarization for any task table you've enabled Sidebar on.




Summarizing Sidebar chats integrated with Microsoft Teams is also supported. For more information, see [Sidebar and Microsoft Teams](#).

## Summarize a Sidebar discussion in Workspace







## Generating a Sidebar discussion summary

You can generate a chat reply recommendation in the following products.

- [Summarize a Sidebar discussion by using Now Assist for Customer Service Management \(CSM\)](#) 
- [Summarize a Sidebar discussion by using Now Assist for IT Service Management \(ITSM\)](#) 
- [Summarize a Sidebar discussion by using Now Assist for HR Service Delivery \(HRSD\)](#) 

## Availability

### Now Assist products and workflows

Workflow	Product
Technology	<a href="#">Now Assist for IT Service Management (ITSM)</a> 
Customer	<a href="#">Now Assist for Customer Service Management (CSM)</a> 
	<a href="#">Now Assist for Field Service Management (FSM)</a> 
Employee	<a href="#">Now Assist for HR Service Delivery (HRSD)</a> 

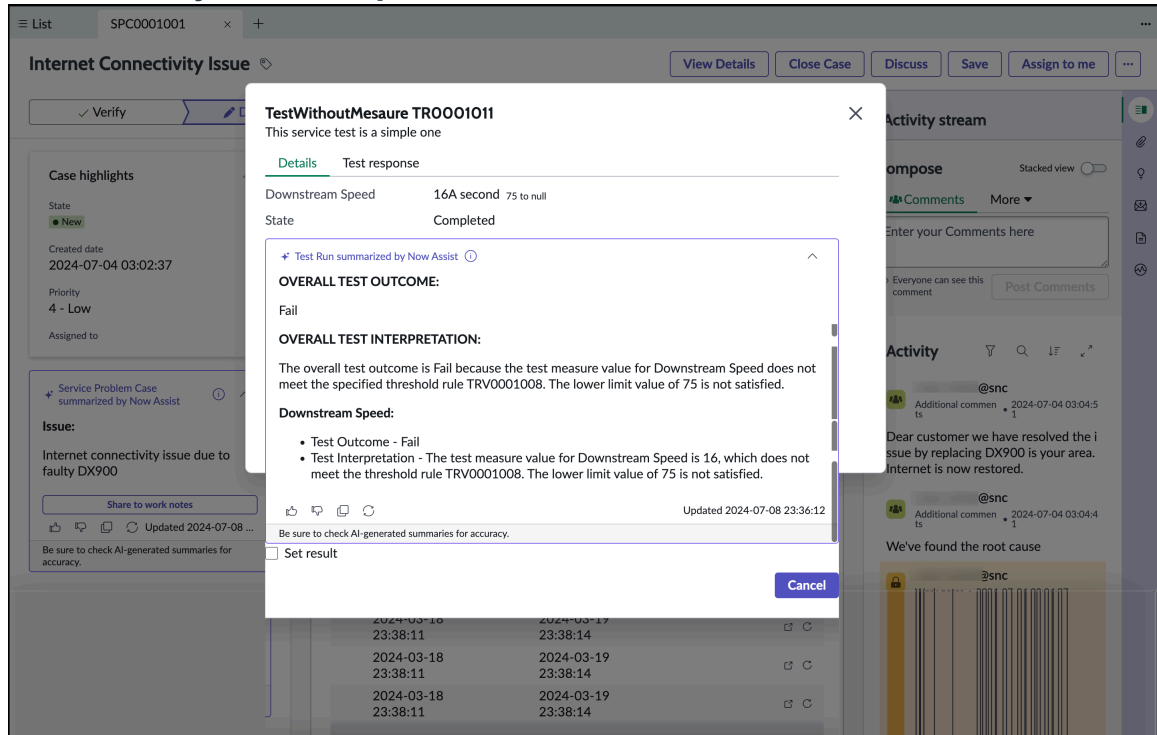
### Test summarization

Use this skill to create a test run summary after the test is executed. It includes the main points covered during the test execution, including the test output, test interpretation, and other defined test parameters.

The test summarization skill provides you with a concise summary of the test executed for a service problem, including the test outcome, test interpretation, and other parameters configured for the specific test definition. With this skill, you can generate the test summary so that you can analyze the root cause of the problem.

The test summarization skill is available in CSM/FSM Configurable Workspace and in Core UI.

## A test summary in the Workspace window



For more information, see [Summarize test for a service problem case using Now Assist for Telecommunications, Media and Technology \(TMT\)](#).

### Availability

This skill is available in the workflow and product as shown in the following table.

#### Now Assist product and workflow

Workflow	Product
Customer	<a href="#">Now Assist for Telecommunications, Media and Technology (TMT)</a>

### Work order task closure summarization

Field service agents can create detailed notes and close work order tasks faster with help from Now Assist.

Now Assist can generate a work order task summary only if the task has at least 50 words in the fields that are used for the input data. The 50-word minimum optimizes the agent experience by ensuring that there’s enough information to summarize.

**Note:** Now LLM Service is the provider for this Now Assist skill.

In the following example, Now Assist generated a work order task summarization that is displayed in the Mobile app.

**Example closure notes for a work order task on a mobile device**

Cancel
Close complete

✦ Closure notes generated by Now Assist
Learn more

✦ Closure notes ✦ AI

I visited ACME Airlines to complete a task of replacing HP Indigo Digital Press, L10546 on March 1st. The task was completed successfully and took 34 minutes to complete.

Steps taken:

- Obtained a new printer, HP Indigo Digital Press L9862, from the inventory room
- Replaced the old printer.

Create a follow-up task

Submit

For more information, see [Generate work order task closure summaries in the Mobile Agent application](#) and [Generate a work order task summary](#).

**Availability**

This skill is available in the workflow and product listed below.

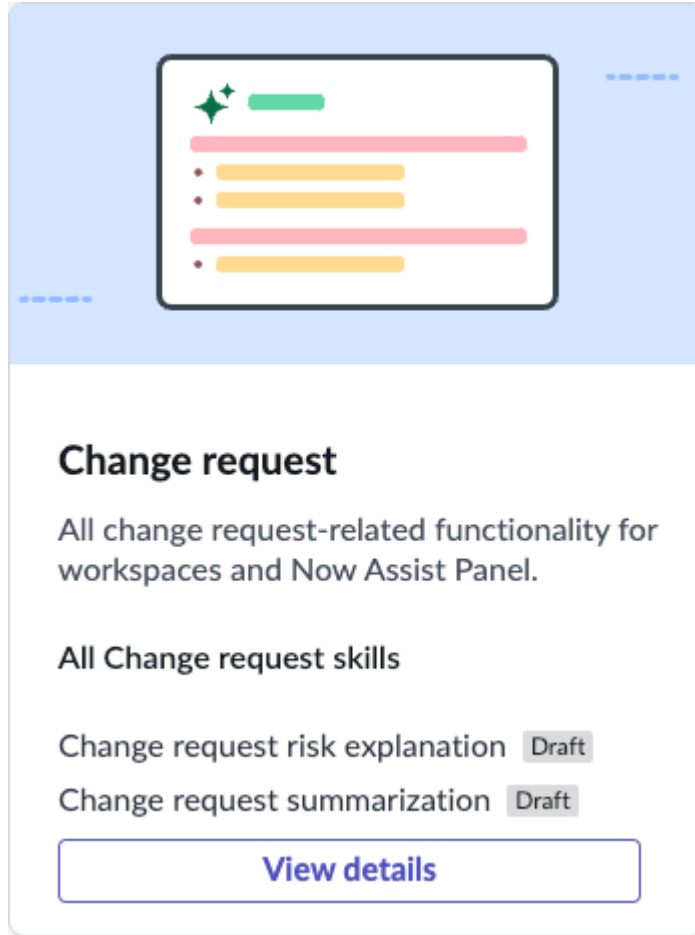
**Now Assist products and workflows**

Workflow	Product
Customer	<a href="#">Now Assist for Field Service Management (FSM)</a>

**Now Assist explanation skills**

These Now Assist skills provide generative AI capabilities for understanding table values on the ServiceNow AI Platform.

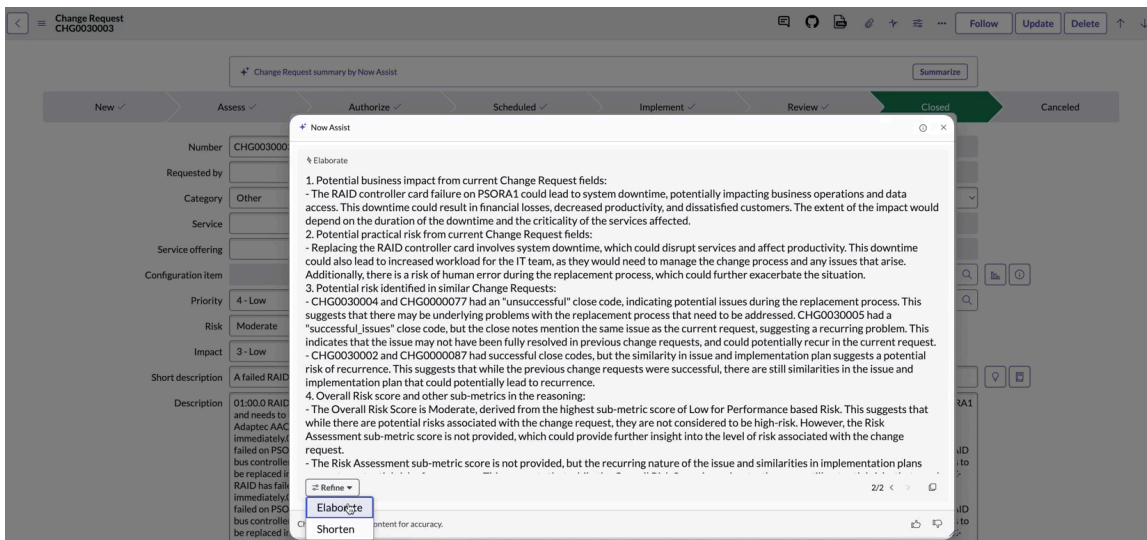
### Example change request skills tile in Now Assist Admin



### Change request risk explanation

Use this skill to show the explanation of the calculated risk rating for a change request.

Use the Now Assist icon ✨ on the Risk field of the change request to calculate the risk. Results are low, medium, or high.




For more information, see [Explain the risk of a change request by using Now Assist for IT Service Management \(ITSM\)](#).

## Availability

This skill is available in the workflow and product listed below.

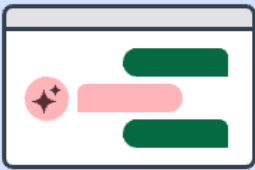
### Now Assist products and workflows

Workflow	Product
Technology	<a href="#">Now Assist for IT Service Management (ITSM)</a> 

### Now Assist recommendation skills

These skills provide generative AI recommendations for chat and email replies, as well as recommended actions for security incidents.

### Example reply recommendation skills in Now Assist for Customer Service Management (CSM)



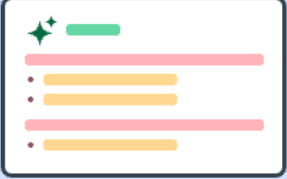
**Chat**

All chat-related functionality for workspaces and Now Assist Panel.

All Chat skills

- Sidebar summarization Not started
- Chat summarization Active
- Chat recommendation** Not started

[View details](#)



**Case**

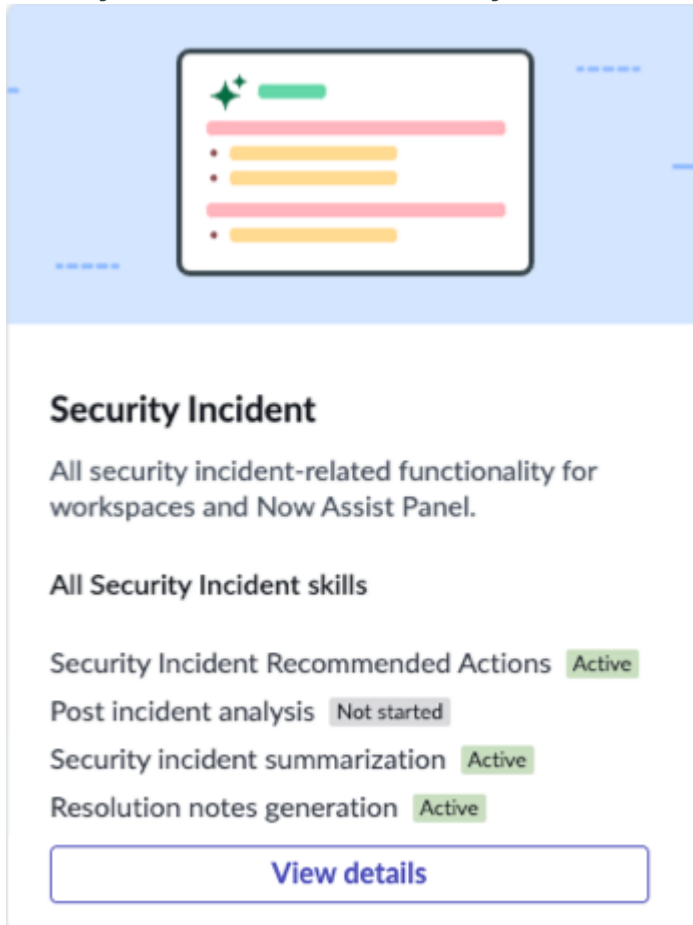
All case-related functionality for workspaces and Now Assist Panel.

All Case skills

- Case summarization Not started
- Resolution notes generation Not started
- Email recommendation** Not started

[View details](#)

## Security Incident skill tile with Security Incident Recommended Actions skill



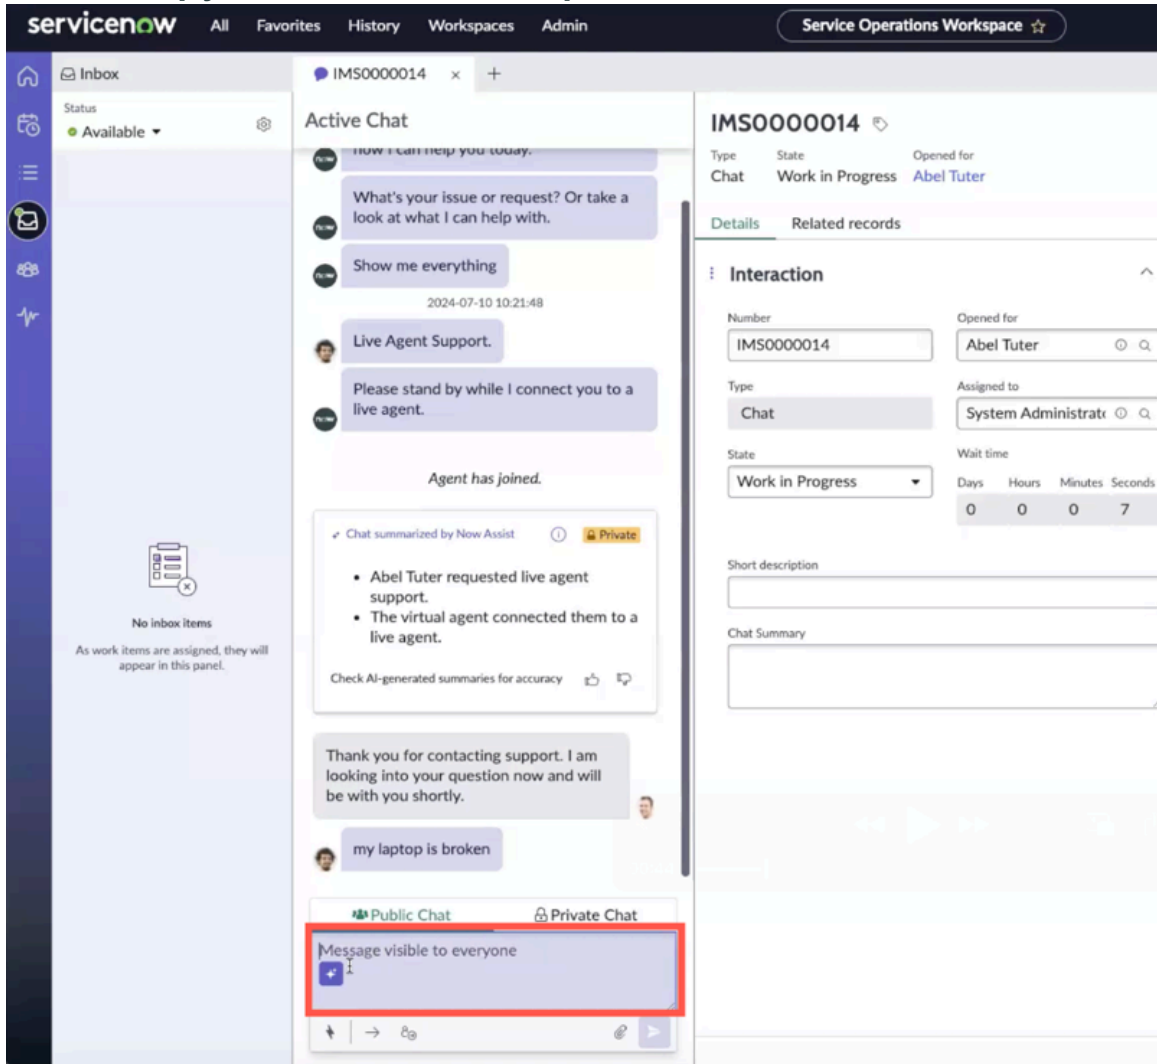
### Chat reply recommendation

Generate a reply based on the context of the chat conversation using Now Assist. Chat reply recommendations provide agents with quick replies to common questions.

The chat reply recommendation skill displays a pop-up window that an agent can use to generate a recommendation with Write with Now Assist, and then review it before sending it as a reply. In the chat message window, either type a response, or leave it blank, and then select the Now Assist icon ✨. Reply recommendations are generated based on KB articles, similar chats, and the context of the conversation.

**Note:** The system performs indexing operations on a few tables, including interaction, incident, and case, whenever the chat reply recommendation functionality is activated.

## Generate a reply recommendation in Workspace



### Generating a chat reply recommendation

You can generate a chat reply recommendation in the following products.

- [Generate a chat reply recommendation by using Now Assist for Customer Service Management \(CSM\)](#)
- [Generate a chat reply recommendation by using Now Assist for IT Service Management \(ITSM\)](#)
- [Generate a chat reply recommendation by using Now Assist for HR Service Delivery \(HRSD\)](#)

### Availability

This skill is available in the workflows and products listed below.

#### Now Assist products and workflows

Workflow	Product
Technology	<a href="#">Now Assist for IT Service Management (ITSM)</a>
Customer	<a href="#">Now Assist for Customer Service Management (CSM)</a>

## Now Assist products and workflows (continued)

Workflow	Product
Employee	<a href="#">Now Assist for HR Service Delivery (HRSD)</a>

### Email response skill

Generate responses for email and suggest templates based on the context of the conversation using Now Assist. This skill can help your agents to draft emails quickly and stay focused.

You can generate email response recommendations in New, Forward, and Reply scenarios. The email response skill displays a pop-up window that an agent can use to generate a reply recommendation or compose a new email or complete a draft. This skill also recommends the most relevant templates from the list of templates based on the case context. In the email message window, either type a response, or leave it blank, and then select the Now Assist icon

The following actions are available from the Now Assist context menu:

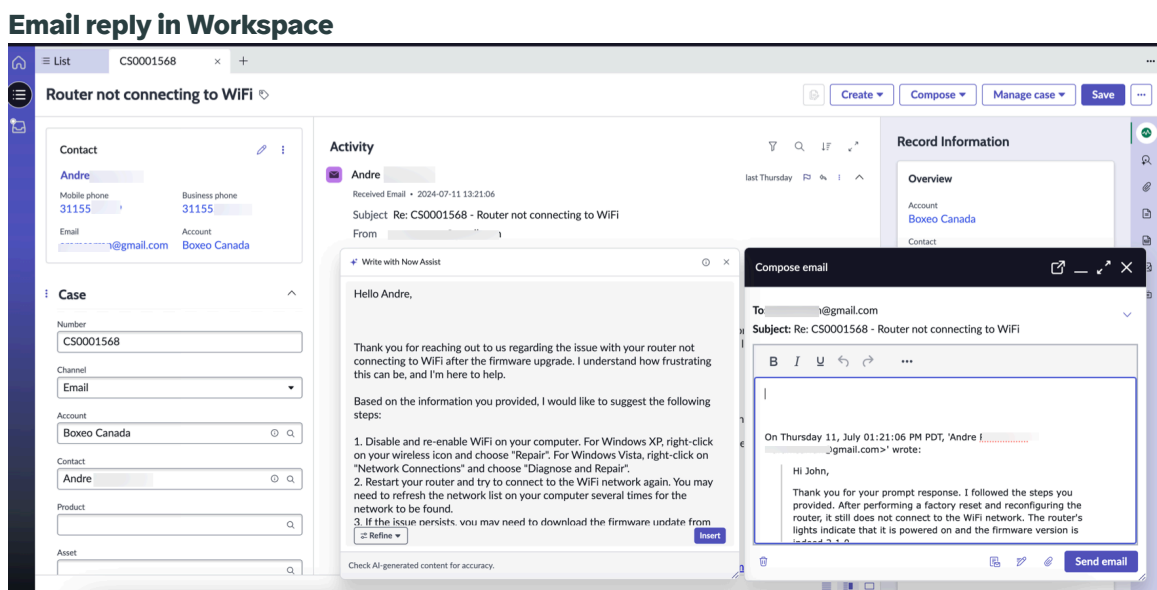
- Generate a recommended email response that is based on the case and email context.
- Generate email responses for new, forwarded, and finishing draft mails.
- Refine the recommendation by elaborating or shortening the response.
- Availability of email template recommendations while composing an email.

### Note:

- Reply recommendations are generated based on KB articles, similar chats, and the context of the conversation.
- The email recommendation skill is on the Case skill card in the Customer group. For more info, see [Now Assist for Customer Service Management \(CSM\)](#) .
- Currently, Generative AI is only supported for the English language.

## Generating an email response by Now Assist

You can generate an email response in Workspace:



To learn more, see [Generate an email response by using Now Assist for Customer Service Management \(CSM\)](#) 

## Email template recommendations by Now Assist


You can view the email templates recommended by Now Assist in Workspace. You select Apply templates in the Compose email screen to view the Email templates, Response templates, and Quick messages. Now Assist recommends you the templates to select and insert it to your email.

You can also search for response templates by using the keyboard shortcut keys. In the Compose email screen, type `/r` to view the response templates recommended by Now Assist.

### Availability

This skill is available in the workflows and products listed below.

#### Now Assist products and workflows

Workflow	Product
Customer	<a href="#">Now Assist for Customer Service Management (CSM)</a> 

### Major Incident Management email content recommendation

Draft an email using well-structured templates to communicate with stakeholders about a major incident. Use different GenAI variables in the template types to communicate efficiently to requesters as per the context.



You can compose email content to communicate about a major incident scenario as a part of incident management. The major incident management email content recommendation skill provides well-structured templates where you can fill the necessary fields values and create email content using an AI-generated response.

This skill recommends the most relevant templates from the list of templates based on the case context. The context comprises different factors like previous emails, targeted email recipients, timelines and latest developments since last communication. Select a template that fits the context and refine your choices until you find one that clearly communicates your email response.

In a scenario where the content isn't generated as required, the skill identifies and displays the potential errors.

Use Now Assist context menu to refine the recommended content by elaborating, shortening or changing the tone of the response.

### Composing an email content recommendation by for major incident management

You can compose a major incident management email content recommendation for [Now Assist for IT Service Management \(ITSM\)](#) . To learn more, see [Generate a Major Incident email content recommendation by using Now Assist for IT Service Management \(ITSM\)](#) .

### Availability

This skill is available in the workflows and products listed below.

### Now Assist products and workflows

Workflow	Product
IT Service Management	<a href="#">Now Assist for IT Service Management (ITSM)</a>

### Security incident recommended actions

View generated remediation steps that might solve security incidents without having to check multiple sources for information.

In Security Incident Response Workspace, you can use the Now Assist icon in the sidebar to generate recommended actions from existing security incidents and knowledge articles. Your analysts can create a response task from the recommended actions.

Generated recommended actions are displayed in cards. Up to four references for the actions are displayed at the top. These references can be any combination of knowledge articles (KB)s or security incidents (SIR#).

### Recommended actions for an incident in Security Incident Response Workspace

The screenshot displays the 'Recommended actions' panel for a security incident titled 'virus spread'. The incident details include: Category: Malicious code act..., Priority: 1 - Critical, Risk score: 92, Assigned to: [blank], State: Contain. The interface shows a navigation bar with tabs for Overview, Details, Investigation, Response Tasks, Related Records, and More. The main content area is divided into several sections:

- Description:** a user was infected via email virus. Virus copied the email to the GAL and company is experiencing wide-spread virus spread.
- Business Impact:**
  - Configuration items:** No data available. There are no scores available for the indicator or for the selected criteria. Contact your ServiceNow administrator for more information.
  - Affected users:** No data available. There are no scores available for the indicator or for the selected criteria. Contact your ServiceNow administrator for more information.
- Threat Intelligence:**
  - Observables:**
    - By fi...:** No data available. There are no scores available for the indicator or for the selected criteria. Contact your ServiceNow administrator for more information.
    - By...:** No data available. There are no scores available for the indicator or for the selected criteria. Contact your ServiceNow administrator for more information.
- Recommended actions:**
  - Analyze the virus for indicators o...**
    - Retrieve the infected email and its attachments for further analysis.
    - Extract the virus from the email attachment using a secur...
  - Contain the spread of the virus b...**
    - Implement network-level blocking rules based on the identified IOCs such as IP addresses and domain names.
    - Use a...
  - Eradicate the virus by cleaning inf...**
    - Use a reputable antivirus solution to clean the infected systems and remove the virus.
    - Update all systems with the latest security...
  - Recover from the incident by rest...**
    - Identify all systems and data that have been affected by the virus.
    - Restore affected systems and data from secure backups that...

For more information, see [Generate recommended actions for a security incident with Now Assist for Security Incident Response](#).

## Availability

This skill is available in the workflow and product listed below.

### Now Assist products and workflows

Workflow	Product
Technology	<a href="#">Now Assist for Security Incident Response</a> 

### Now Assist skills in the Creator workflow


Now Assist for Creator includes a number of skills that can make developing on the ServiceNow AI Platform more efficient.

Example Now Assist for Creator tiles in Now Assist Admin

Now Assist Features > Creator

## Now Assist skills for Creator

Manage all Creator skills.




**Analytics**  
All analytics generated functionalities

All Analytics skills

Data visualization generation Not started

[View details](#)




**Flow**  
All flow-related functionality for creators.

All Flow skills

Flow generation Not started  
Flow recommendations Not started

[View details](#)

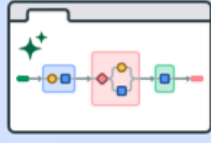


**Code**  
All code-related functionality for creators.

All Code skills

Intelligent code recommendations Active

[View details](#)




**Playbook**  
All playbook-related functionality for creators.

All Playbook skills

Playbook generation with images Not started  
Playbook recommendations Not started  
Playbook generation Not started

[View details](#)




**Spoke**  
All spoke related skills are available to enhance third-party integration experiences.

All Spoke skills

Spoke Generation Not started

[View details](#)




**UI Generation**  
All skills available for creating UI in UI Builder

All UI Generation skills

Experience Generation Not started

[View details](#)




**Service Catalog**  
All skills that are related to creating and maintaining Service Catalog.

All Service Catalog skills

Catalog item generation Not started

[View details](#)



**App Generation**  
Kickstart application generation for your business workflows.

All App Generation skills

App Generation Active

[View details](#)

For more information, see [Now Assist for Creator](#).

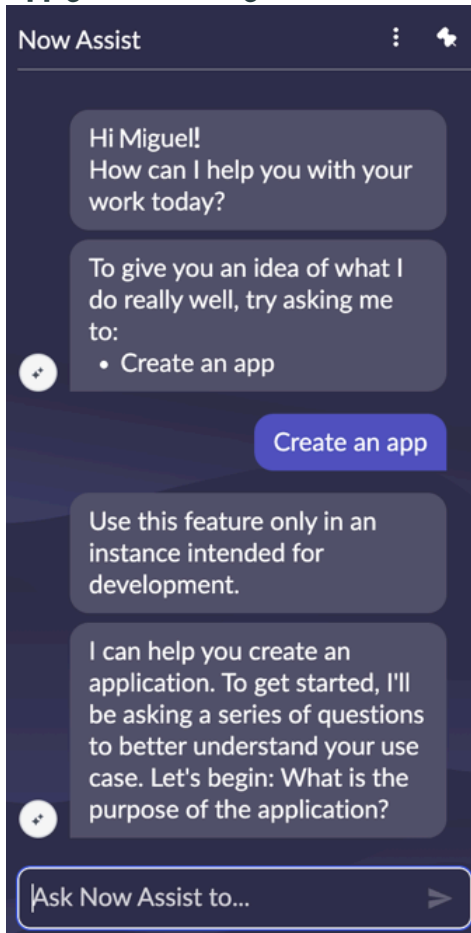
**App generation**

App generation enables developers to kick start application building through conversation.

The app generation skill lets your developers begin creating applications through a natural conversation with generative AI. They can describe their business process and engage in a back-and-forth conversation with Now Assist to develop an application for your organization. With this feature, your organization can expedite the initial development of a basic app that can then be customized.

**Note:** Your administrators and developers must be assigned with the admin and now.assist.creator roles to use app generation. For more information about using app generation, see [Now Assist for app generation in ServiceNow Studio](#) (Xanadu Store Release 2) or [Generate apps with Now Assist for Creator for use with AES](#) (Xanadu Store Release 1).

### App generation begins



### Conversation

Chat with Now Assist for Creator to specify the business processes you want in the application including the details on the objectives, users, workflows, and experiences.

### Refinement in Xanadu Store Release 2

Now Assist for Creator provides a summary of the application requirements based on the information collected during the conversation. Preview the application and check that it's accurate. If you want to make changes, stay in the conversation and keep editing. Now Assist continues to update the preview tab based on your comments and provides summaries until you choose to proceed with generating the application.

### Generation

Now Assist for Creator generates the application and associated components, including the tables, roles, access control lists (ACLs), and record producers.

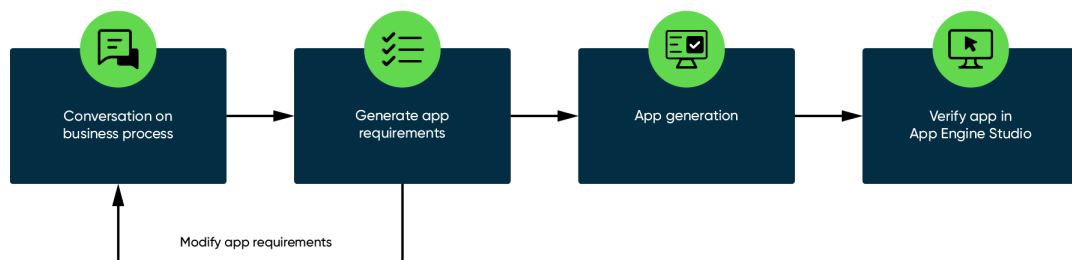
### Verify

Developers verify everything that is generated in ServiceNow Studio. The developer can modify the app to make it suit your organization's needs. For example, the application's functionality can be extended by adding flows and automation, script includes, business rules, and other features.

For more information, see [Now Assist for app generation in ServiceNow Studio](#).

In Xanadu Store Release 1, there are three distinct phases to building apps with Now Assist for Creator using the app generation skill and a fourth phase to verify the app in App Engine Studio.

### App generation workflow in Xanadu Store Release 1



### Conversation

Your developers can chat with Now Assist for Creator to specify the business processes they want in the application including the details on the objectives, users, workflows, and experiences.

### Refinement

Now Assist for Creator provides a summary of the application requirements based on the information collected during the conversation. The developer reviews each summary and if the requirements meet the application's needs, they move forward with generating the app. If the developer wants to make changes, they can stay in the conversation and keep editing. Now Assist continues to refine the application requirements based on their comments, and provides summaries until the developer chooses to proceed with generating the application.

### Generation

Now Assist for Creator generates the application and associated components, including the tables, roles, access control lists (ACLs), and record producers.

### Verify

Developers verify everything that is generated in App Engine Studio. The developer can modify the app to make it suit your organization's needs. For example, the application's functionality can be extended by adding flows and automation, script includes, business rules, and other features.

For more information, see [Now Assist for app generation in AES](#).

### Availability

This skill is available in the workflow and product listed below.

### Now Assist products and workflows

Workflow	Product
Creator	<a href="#">Now Assist for Creator</a>

### Catalog item generation

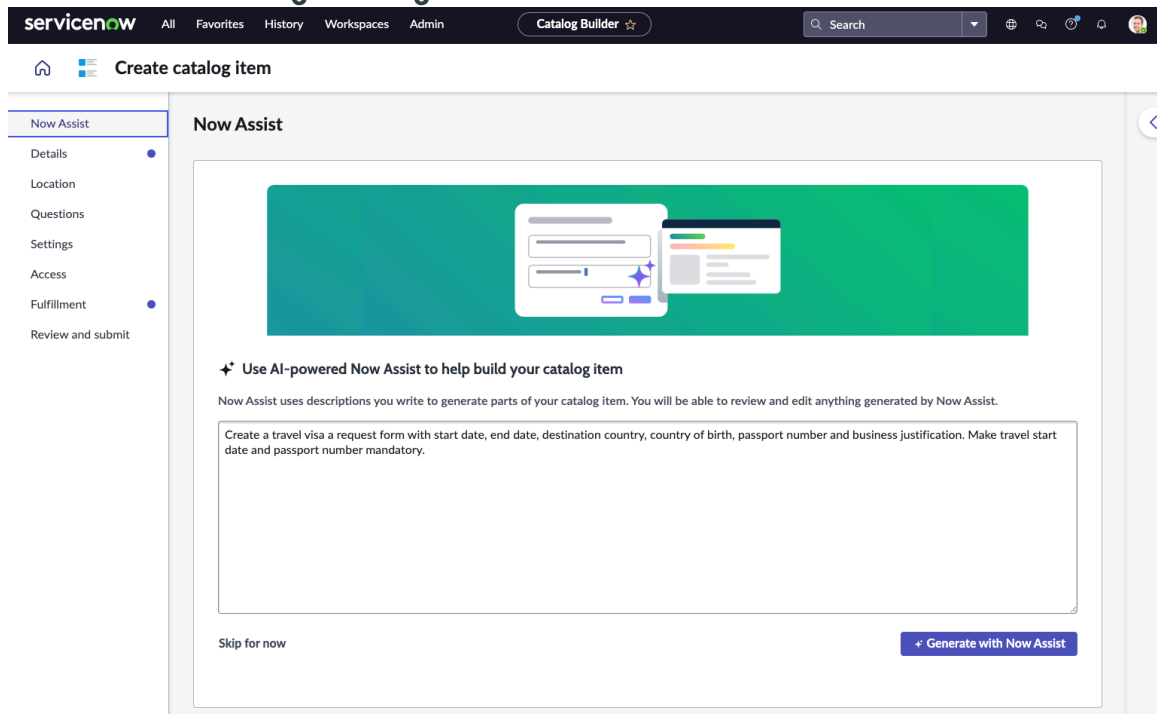
Use the Catalog item generation skill to create catalog items of your choice using Now Assist.

The Catalog item generation skill helps reduce the time needed to create catalog items. While creating a catalog item, you can describe the catalog item you want, and Now Assist generates the catalog item for you.

For more information, see [Now Assist in Catalog Builder](#).

**Note:** Verify that you have activated the skill before creating a catalog item by using Now Assist.

### Now Assist for creating a catalog item



### Availability

This skill is available in the following workflow and product.

### Now Assist products and workflows

Workflow	Product
Creator	<a href="#">Now Assist for Creator</a>

### Related topics

[Now Assist in Catalog Builder](#)

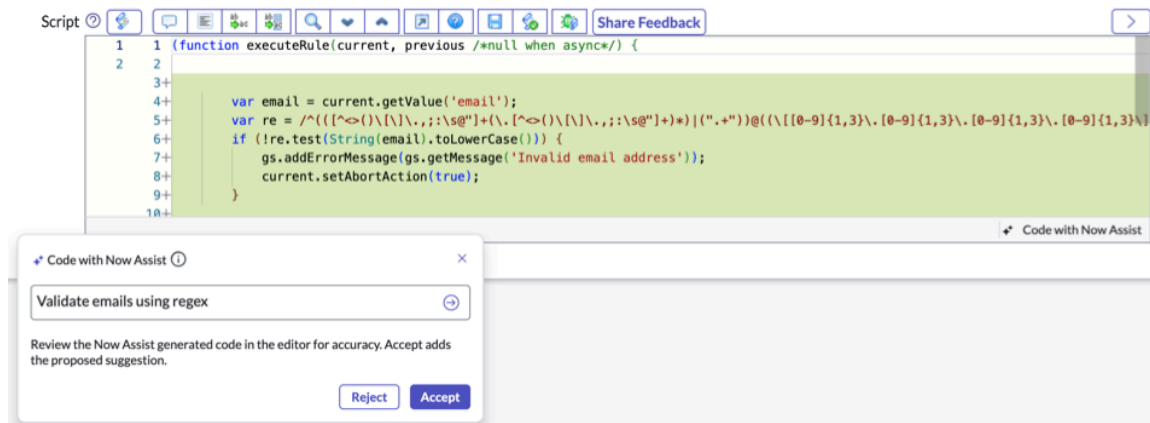
### Intelligent code recommendations

Code generation is designed to empower developers to write scripts quickly with AI-generated code based on text or code prompts and can improve time to value for applications.

Now Assist for Creator activates the code generation skill. With code generation, you provide text describing the code to generate and get code suggestions in the JavaScript editor on forms in the ServiceNow AI Platform and in Script steps in Workflow Studio. Developers with varying levels of experience in scripting on the ServiceNow AI Platform can benefit from using code generation to get started writing custom scripts or iterate on scripts more efficiently.

**Note:** Now LLM Service is currently the only provider for this Now Assist application's skills.

To generate code suggestions, you describe the goal of the code to generate in the Code with Now Assist dialog box. The code suggestion appears in the lines following your prompt but isn't added to your script until it's accepted.



Optionally, you can turn on code completion functionality to use code or single-shot prompts in script editors with Now Assist for code generation.

Developers must be assigned the now.assist.creator role to use code generation. To get started developing with AI-generated code, see [Now Assist for code generation](#).

### Availability

This skill is available in the workflow and product listed below.

#### Now Assist products and workflows

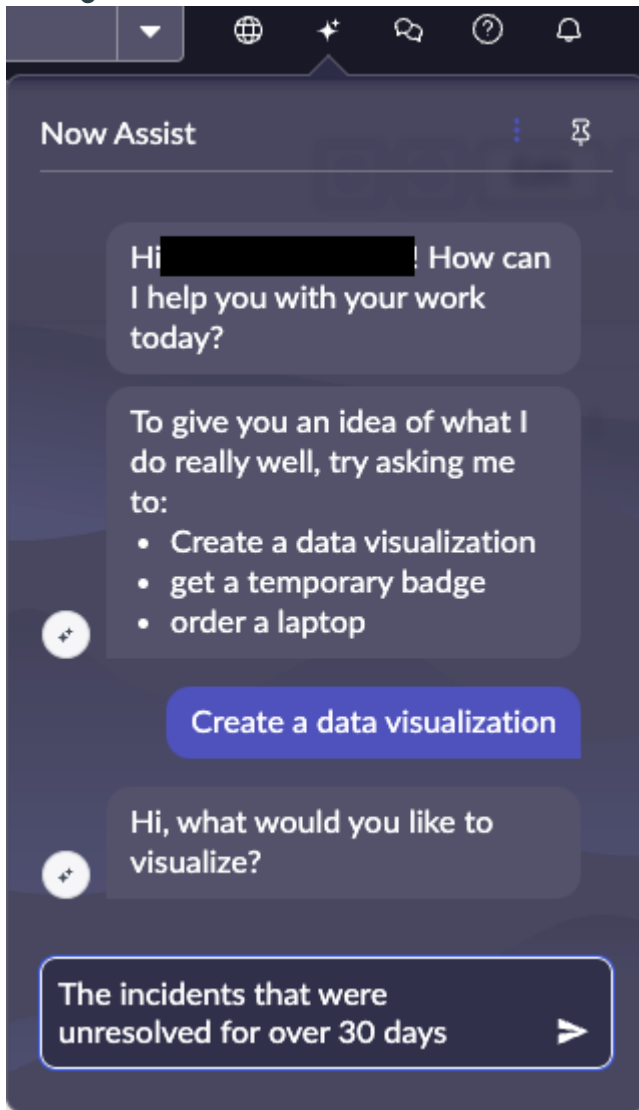
Workflow	Product
Creator	<a href="#">Now Assist for Creator</a>

### Data visualization generation

Data visualization generation lets users create Platform Analytics data charts through conversation.

The data visualization generation skill lets your users create relevant charts and scores based upon natural language queries and available data, leveraging Platform Analytics. The charts can show data from Platform Analytics table and indicator data sources. This skill simplifies the process of data visualization generation and configuration, potentially increasing efficiency.

### Asking Now Assist for a data visualization



## Request results: Visualization details

**Visualization details**

Visualization generated by Now Assist. Check it for accuracy.

Number	Resolved	Opened	Short description
INC0000601	2024-06-14 04:09:56	2024-06-11 02:42:59	The USB port on my PC stopp...
INC0000038	2024-06-15 12:56:13	2024-03-22 17:36:44	my PDF docs are all locked ...
INC0000036	2024-06-15 12:56:13	2024-03-22 17:33:49	Issue with networking
INC0000035	2024-06-15 12:56:13	2024-03-25 17:32:47	Reset my password
INC0000034	2024-06-06 09:12:01	2024-02-27 12:25:29	Does not look like a backup...

[View all](#)

Table  
Incident [incident]

Conditions  
Resolved > javascript:gs.daysAgo(30)

[Close](#)

**Now Assist**

Hi, what would you like to visualize?

1m ago

The incidents that were unresolved for over 30 days

Sure, I've created a list that matches your request.

[See details](#)

[Add to Dashboard](#)

This is a list visualization from a table source named Incident, filtered by records resolved more than 30 days ago.

Ask Now Assist to...

Because it is part of Now Assist for Creator, data visualization generation requires the `now.assist.creator` and `now_assist_panel_user` roles. You also need access to the data that you want to see. If you want to add the visualization to a dashboard, you need editing rights to that dashboard. If the user asks for a data visualization while a dashboard is open, and they have editing rights to that dashboard, they have the option to add the visualization to that dashboard.

For more information, see [Data visualizations in Platform Analytics](#) and [Analytics Generation](#).

## Availability

This skill is available in the workflow and product listed below.

### Now Assist products and workflows

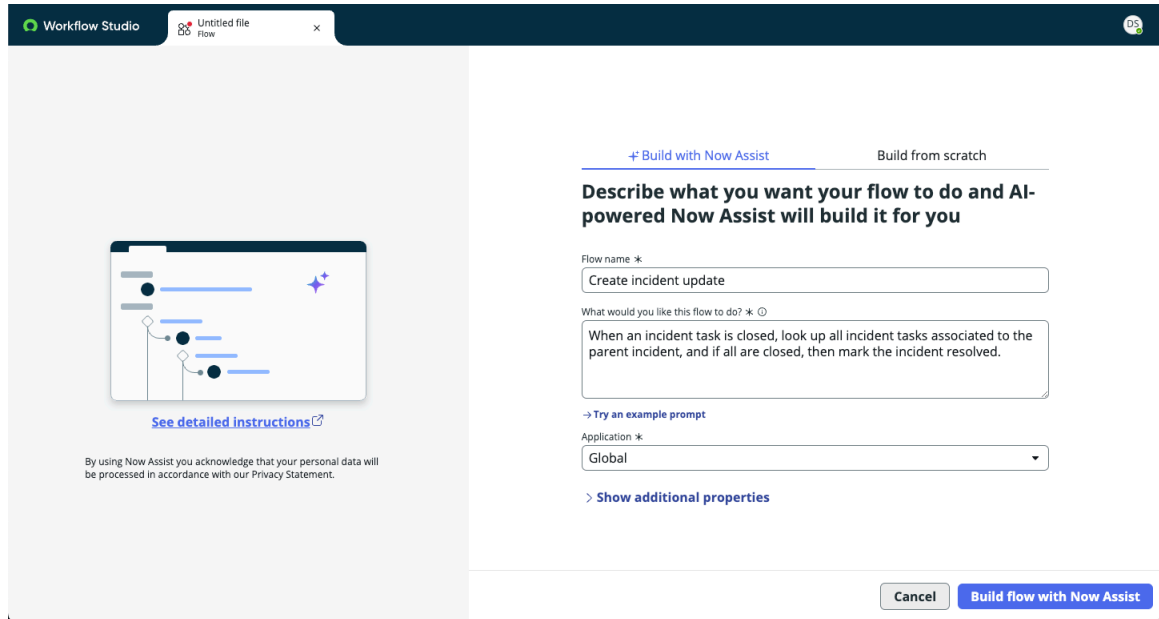
Workflow	Product
Creator	<a href="#">Now Assist for Creator</a>

## Flow generation

Use the flow generation skill in Workflow Studio to create new flows.

Create multi-step flow outlines with generative AI. Flow outlines require configuration to add input values and data references.

## Use Now Assist to create a flow



For more information, see [Flow generation](#).

## Availability

This skill is available in the workflow and product listed below.

### Now Assist products and workflows

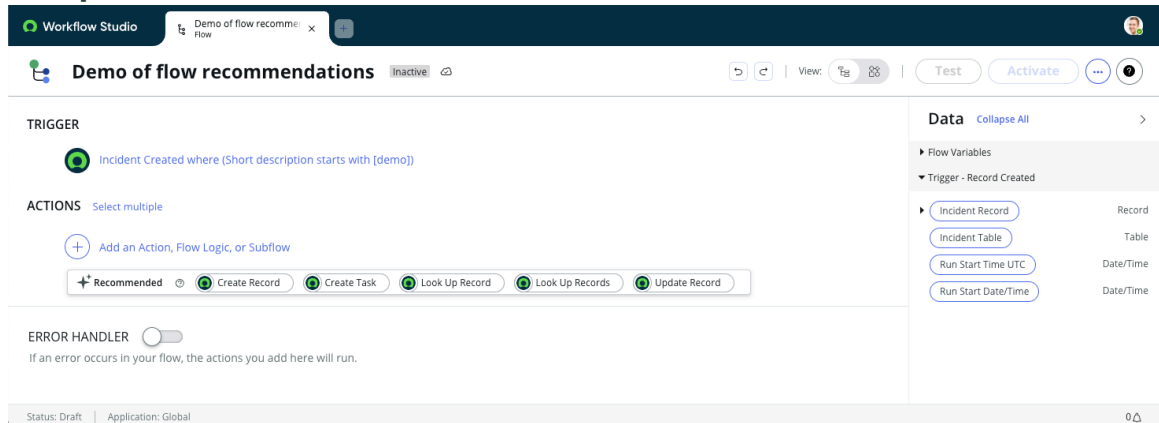
Workflow	Product
Creator	<a href="#">Now Assist for Creator</a>

## Flow recommendations

Generative AI recommendations can help you author flows by recommending possible components based on what you've entered.

Select the next component in your flow from a list of AI-generated recommendations. The system generates recommendations based on the current position in the flow and the flow component names listed before.

### Example flow recommendations



For more information, see [Flow recommendations](#).

## Availability

This skill is available in the workflow and product listed below.

### Now Assist products and workflows

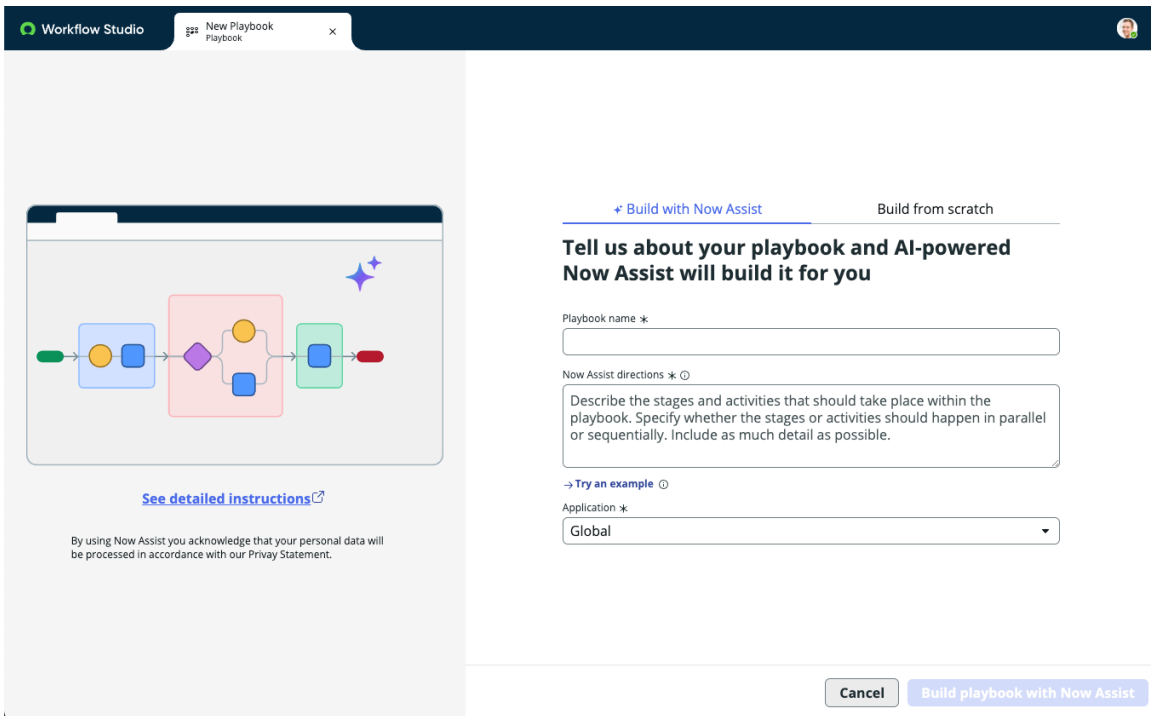
Workflow	Product
Creator	<a href="#">Now Assist for Creator</a>

## Playbook generation

Now Assist for Creator includes the playbook generation skill. Playbook generation gives generative AI capabilities to playbook authors, who can provide text directions to create a playbook outline with placeholder activities.

Playbook generation is a skill installed with the Now Assist for Creator (sn\_now\_creator) application. You can install this application from the [ServiceNow Store](#) website.

Access the playbook generation skill when creating a new playbook in Workflow Studio.



For more information, see [Playbook Assist](#).

**Note:** Playbook authors must be assigned the **now.assist.creator** role to use playbook generation. See [Playbook Assist roles](#) for more information.

## Availability

This skill is available in the workflow and product listed below.

### Now Assist products and workflows

Workflow	Product
Creator	<a href="#">Now Assist for Creator</a>

### Refine content

Use the refine content skill to shorten, elaborate, rephrase, or generate content in the fields by using Now Assist.

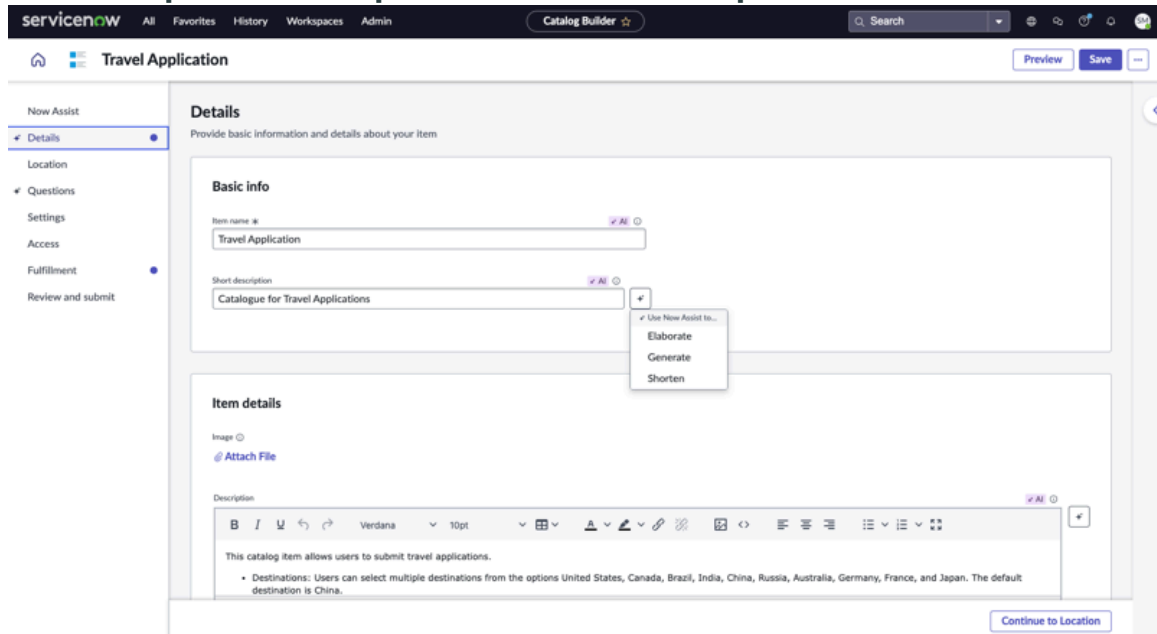
The refine content skill helps users shorten, elaborate, rephrase, or generate content in the **Short description** and **Description** fields in the Details step while creating or editing a catalog item.

You can also elaborate or shorten the content in the **Instructions** field and the **Conversational label** field in the Questions step.

For more information, see [Now Assist in Catalog Builder](#).

**Note:** Verify that you have activated the skill before refining the information in the fields for a catalog item.

### Short description and Description fields in the Details step



### Availability

This skill is available in the following workflow and product.

#### Now Assist products and workflows

Workflow	Product
Creator	<a href="#">Now Assist for Creator</a>

### Related topics

[Now Assist in Catalog Builder](#)

### Robotic Process Automation (RPA) bot generation

Use this skill to create automations, activities, and automation logic from text instructions and preview options.




RPA developer or RPA admin roles are required to use this skill.

Access RPA bot generation skills from the RPA Desktop Design Studio user interface.

Enable RPA bot generation to gain the following benefits:

- Build simple, brand-new automations quickly and efficiently.
- Easily add new activities to existing automations, ensuring modularity and scalability.
- Use in-line prompting to refine automation logic, whether starting from components or a empty design surface.



For more information, see the following topics:

- [Create an automation with Now Assist](#) 
- [Create an activity with Now Assist](#) 
- [Build an automation with Now Assist](#) 

## Licensing requirements

The Now Assist for RPA Hub application requires a Workflow Data Fabric (previously known as Automation Engine) license and a Now Assist for Creator license.

### Related topics

- [Robotic Process Automation \(RPA\) bot generation](#) 
- [Supporting information for Now Assist for RPA Hub](#) 

### Spoke generation

Generate spoke and create actions using Now Assist in Workflow Studio.

## Overview of spoke generation

Build custom spokes for the required third-party application by using its API documentation. Kickstart spoke creation by copying and pasting the required API documentation snippet in Workflow Studio. This feature is useful when the required third-party application doesn't have an OpenAPI specification or Postman collection. To use the feature, activate the spoke generation skill in the Now Assist for Creator feature.

## Create spoke using Now Assist

**BUILD INFO**

### How do you want to create your spoke?

Select the method by which you want to create your spoke.

**Create using API specification**

OpenAPI

Postman collection

**Create using documentation**

Now Assist

Using action designer

Cancel
Continue

## Use Now Assist to generate action preview

For more information, see [Use Now Assist to create spokes and build actions](#).

## Availability

This skill is available in the workflow and product listed below.

### Now Assist products and workflows

Workflow	Product
Creator	<a href="#">Now Assist for Creator</a>

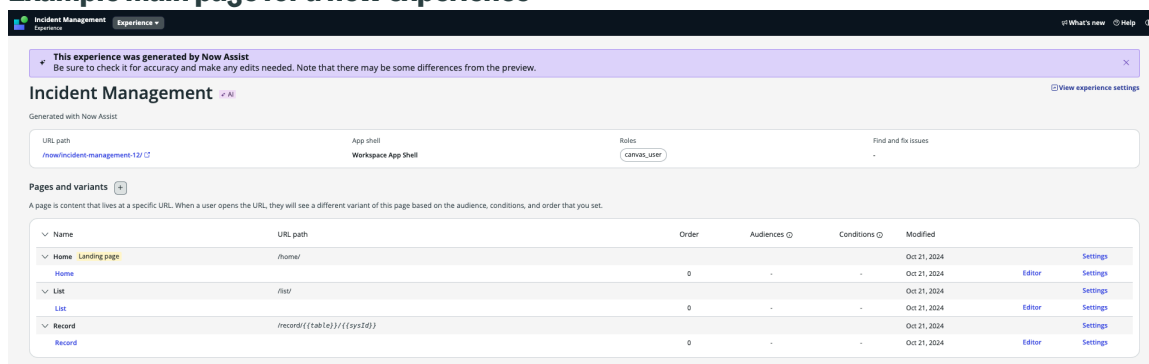
### UI generation

Use this generative AI skill to quickly create experiences by describing what you want using natural language.

To generate a UI experience, describe the table, chart type, and navigation your users will be working with. For best results, include as much detail as possible.

Once generated, you can review the generated experience and accept or reject it. Once accepted, you can modify the experience description to achieve the desired results.

#### Example main page for a new experience



For more information, see [Experience generation](#) .

### Availability

This skill is available in the following workflow and product.

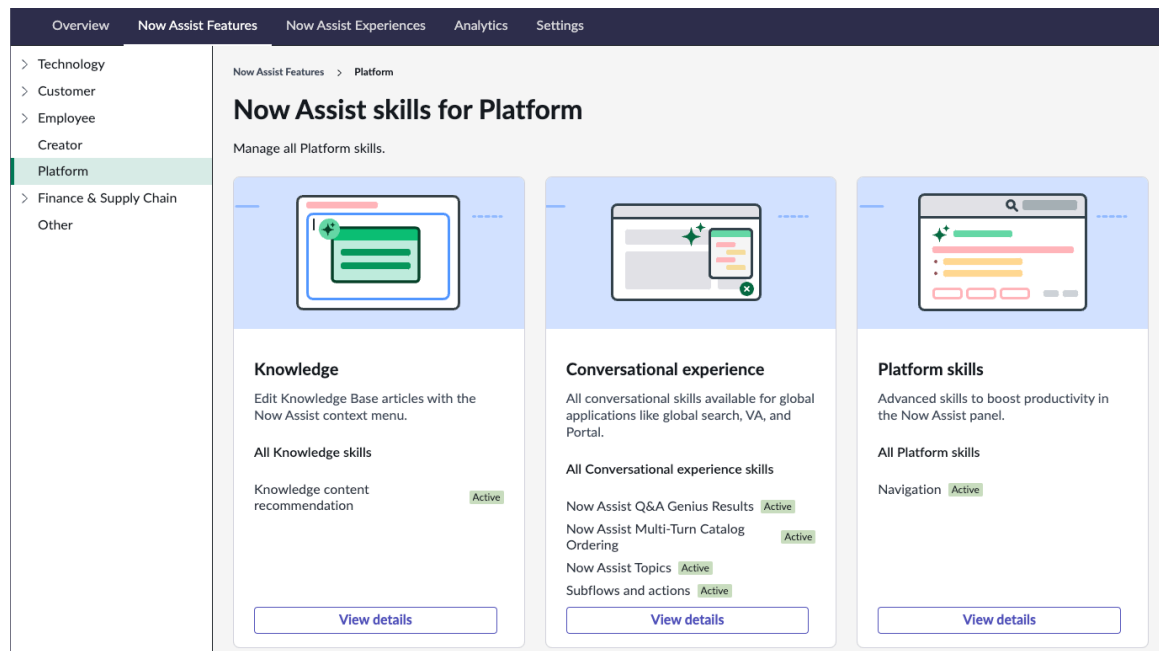
### Now Assist products and workflows

Workflow	Product
Creator	<a href="#">Now Assist for Creator</a>

### 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.

## The Platform skills tile in the Platform workflow of the Now Assist Admin console



### Now Assist conversational experience

Conversational experience generative AI skills can improve the user self-service experience. Platform Now Assist skills generate search summaries, order catalog items from the chat window, and create Virtual Agent topics that use large language model (LLM) topic discovery.

The following skills require Now Assist in AI Search and Now Assist in Virtual Agent. Both are available in some Now Assist products.

You can set up Now Assist in Virtual Agent from the [Conversational Interfaces Console](#) or from the [Now Assist Admin console](#). For more information, see [Configuring Now Assist in Virtual Agent](#).

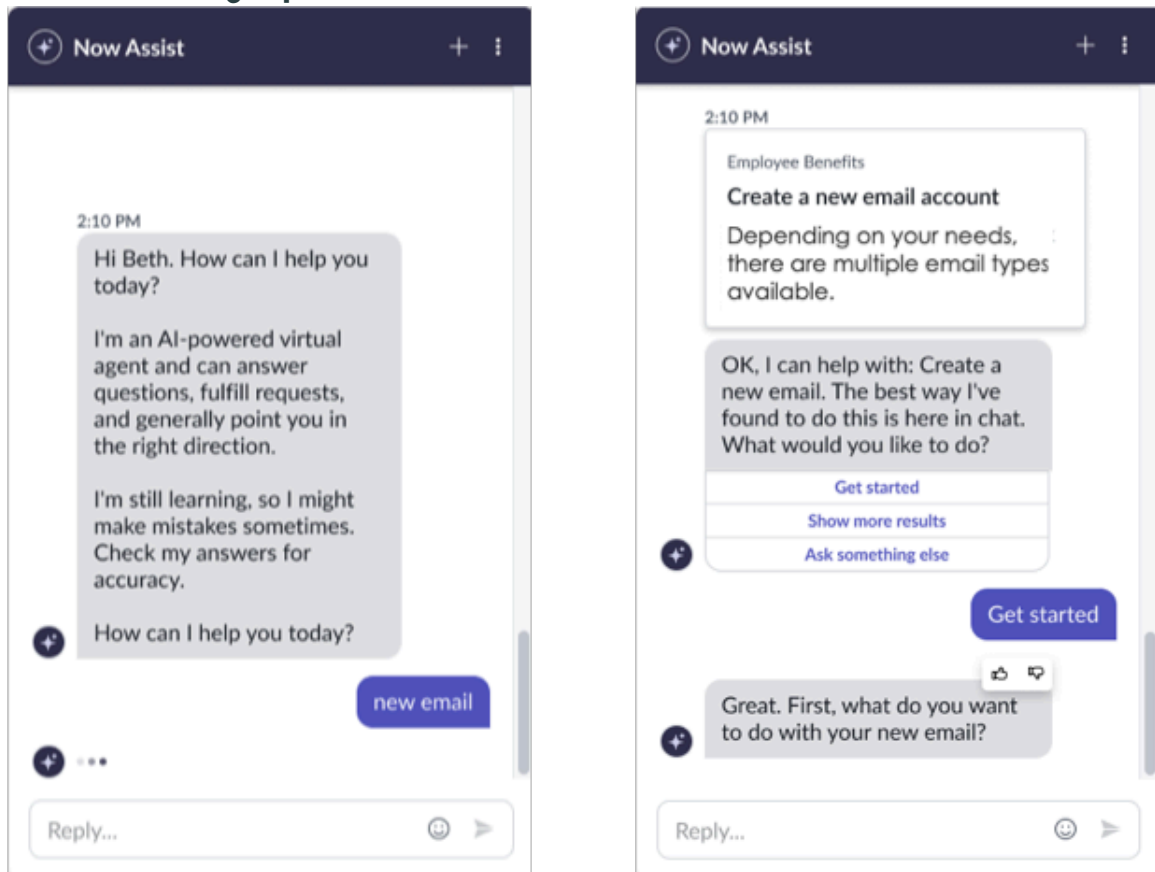
For more information, see [Now Assist in Virtual Agent](#) and [Now Assist in AI Search](#).

### Now Assist Multi-Turn Catalog Ordering

This skill gives users conversational access to available options in the Service Catalog.

Users can request an item, such as a mobile phone. The user can then provide more information to refine the search. For example, they may refine their request to a blue 256-GB iPhone. They can even request a new item instead, all in the same conversation, and the generative AI creates its responses using natural language.

**New email catalog request from the chat window**

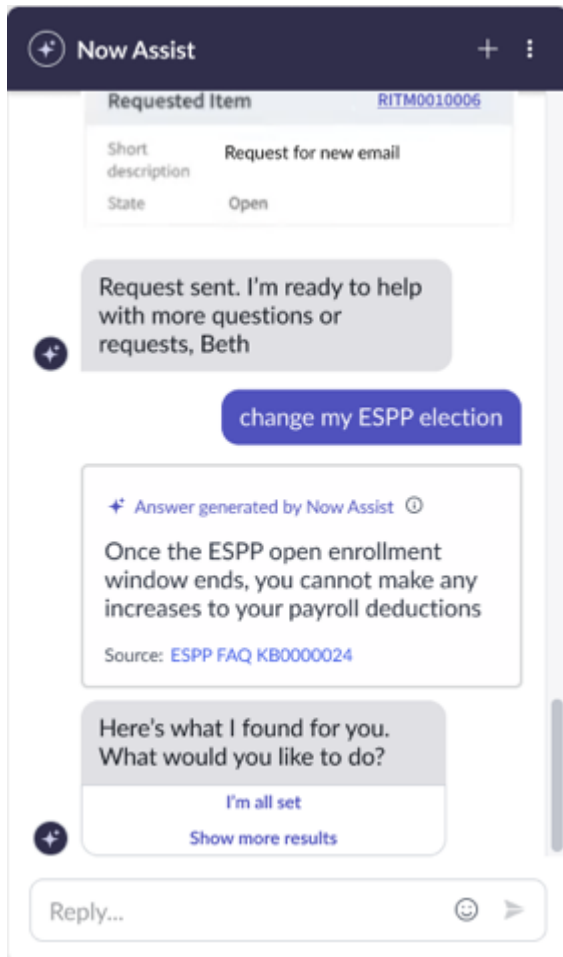


For more information, see [Now Assist in Virtual Agent](#).

**Now Assist Q&A Genius Results**

Use this skill to help users find answers to their questions in a conversational format. Now Assist summarizes the Q&A Genius Result answers with references.

Now Assist in AI Search uses the Now LLM to extract actionable Q&A Genius Result answers from knowledge article results found in Service Portal, Virtual Agent, Employee Center, and global searches. When a user submits a question in search, AI Search retrieves the top knowledge article result and passes it to the Now LLM for answer generation. Answers augment the user's search results, displaying as actionable Now Assist Q&A Genius Result answer cards. For reference, each answer card includes a link to its source knowledge article.



For more information, see [Now Assist in AI Search](#).

### Now Assist Topics

Use this skill to create Virtual Agent topics that use large language model (LLM) discovery.

Watch this video to learn about LLM topics in Virtual Agent.

[https://player.vimeo.com/video/1011425825?h=9a3dc725e1&badge=0&autoplay=0&player\\_id=0&app\\_id=58479](https://player.vimeo.com/video/1011425825?h=9a3dc725e1&badge=0&autoplay=0&player_id=0&app_id=58479)

Unlike NLU topics, LLMs don't require models, intents, or keywords to be linked to the topic. LLMs can discover topics and perform language-related tasks, such as text generation for case summaries and resolution notes, without months of training on NLU models. The LLM does all of the heavy lifting for you. The only requirement is a robust, plain language topic description. The LLM uses this description to find the best topic match for the user utterance. If there are multiple potential matches, the user will see a list of topics to choose from.

For more information, see the following topics:

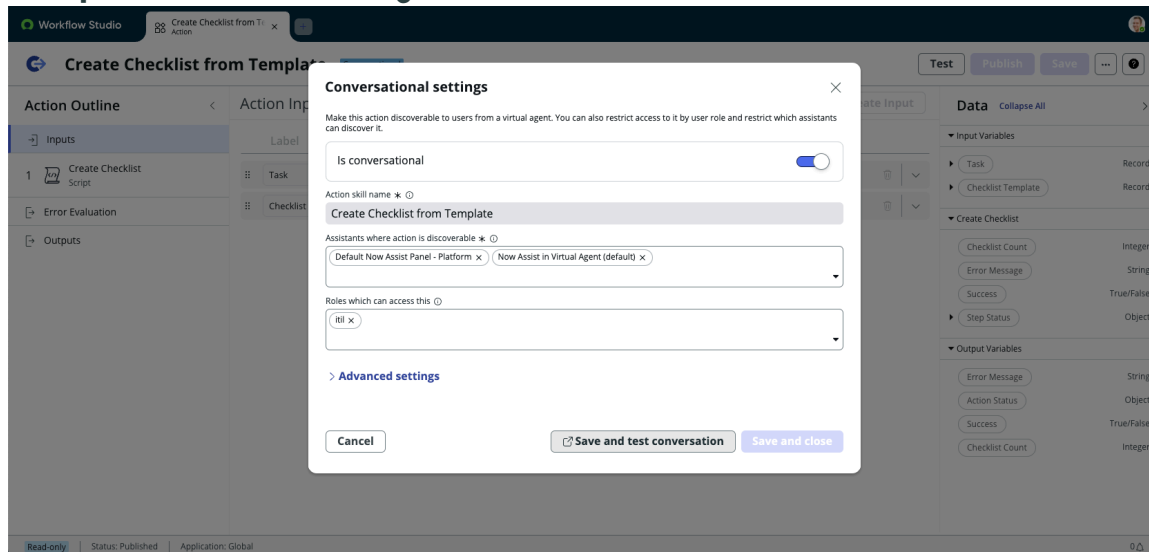
- [Now Assist in Virtual Agent](#)
- [LLM topic discovery in Virtual Agent](#)
- [LLM assistants](#)

### Subflows and actions

Use this skill to create Now Assist topics from Workflow Studio subflows and actions that are conversation-compatible.

Like other Now Assist topics, the LLM requires a robust plain language topic description, which is provided by the action or subflow description and its input hints. The LLM uses this description to find the best subflow or action match for the user utterance. If there are multiple potential matches, the user will see a list of subflows or actions to choose from.

### Example conversational settings in Workflow Studio



For more information, see the following topics:

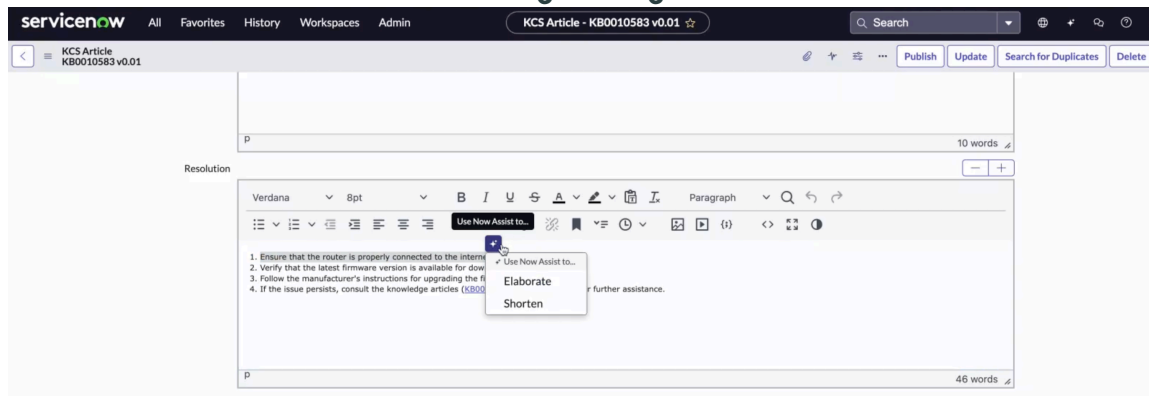
- [Conversational actions](#)
- [Conversational subflows](#)

### 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



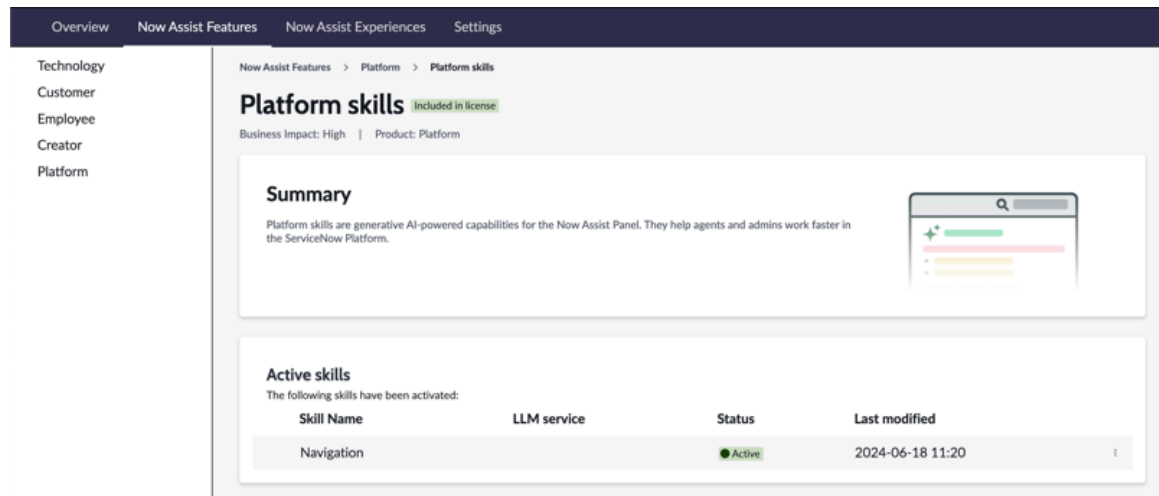
For more information, see [Edit an article using the Now Assist context menu](#).

### Platform skills

Now Assist includes special generative AI skills, such as the ability to search for records using plain language, available with the activation of any Now Assist plugin.

To use these skills, activate one or more Now Assist plugins, such as Now Assist for Customer Service Management (CSM) or Now Assist for IT Service Management (ITSM).

### Available Platform skills in the Now Assist Admin console



### Navigation

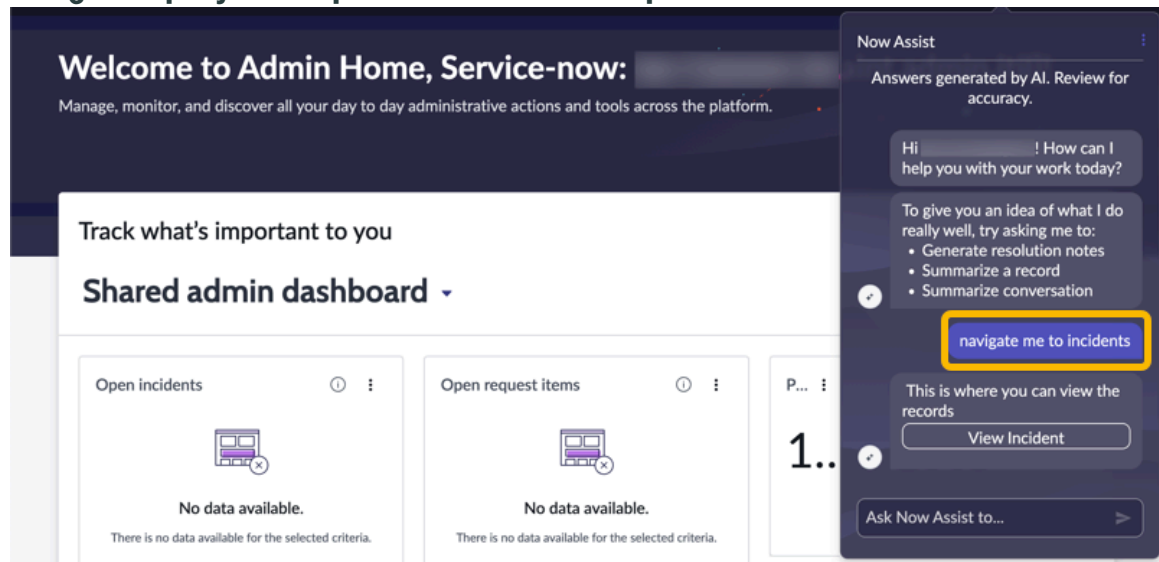
Use the navigation skill in Now Assist to take you where you want to go on the ServiceNow AI Platform.

### 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`.

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

Number	Opened	Short description	Caller	Priority	State	Category
INC0010253	2024-01-01 23:37:27	Language issue	Abraham Lincoln	4 - Low	In Progress	Software
INC0010248	2023-12-22 00:45:11	Order stuck in load status	Alejandra Prenatt	3 - Moderate	New	Inquiry / Help
INC0009005	2018-08-31 21:35:21	Email server is down.	David Miller	1 - Critical	New	Software
INC0000060	2016-12-12 07:19:57	Unable to connect to email	Joe Employee	3 - Moderate	Closed	Inquiry / Help
INC0000058	2016-08-10 09:37:45	Performance problems with email	Bow Ruggeri	5 - Planning	New	Inquiry / Help
INC0000047	2024-04-29 13:53:18	Issue with email	Joe Employee	3 - Moderate	In Progress	Inquiry / Help
INC0000032	2024-02-10 16:19:36	EMAIL Server Down Again	Joe Employee	5 - Planning	Closed	Inquiry / Help
INC0000021	2024-02-25 15:52:01	New employee hire	Fred Luddy	5 - Planning	Closed	Inquiry / Help
INC0000017	2015-08-12 16:41:00	How do I create a sub-folder	Joe Employee	1 - Critical	On Hold	Inquiry / Help

### 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.

### A refined query of the Incidents table

Number	Opened	Short description	Caller	Priority	State	Category
INC0009005	2018-08-31 21:35:21	Email server is down.	David Miller	1 - Critical	New	Software
INC0007001	2018-10-16 22:47:10	Employee payroll application server is down.	David Miller	1 - Critical	New	Hardware

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

The screenshot shows the 'Incidents' table with the following data:

Number	Opened	Short description	Caller	Priority	State	Category
INC0009005	2018-08-31 21:35:21	Email server is down.	David Miller	1 - Critical	New	Software
INC0007001	2018-10-16 22:47:10	Employee payroll application server is down.	David Miller	1 - Critical	New	Hardware

The Now Assist panel on the right shows a search for 'state is new'. It displays a message: 'There are 2 potential tables matching what you are looking for. Which table would you like to view?' with two buttons: 'Catalog Item: sc\_cat\_item' and 'Catalog: sc\_catalog' highlighted in a yellow box.

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.

### Rephrase request message in the Now Assist panel

The screenshot shows the same 'Incidents' table as above. The Now Assist panel on the right shows a search for 'state is new' with a message: 'Sorry, I am not able to generate proper filters based on your request. Do you want to rephrase your request to try again or directly navigate to the table list to add filters on your own?' with two buttons: 'Rephrase request' and 'Navigate to table list' highlighted in a yellow box.

## Now Assist Guardian

Enable Now Assist Guardian, built with Llama 3.1, to monitor and evaluate content created with generative AI to help protect and enhance the user experience.

### Now Assist Guardian Overview

Generative AI is an emerging technology. Human interactions are unpredictable, and outputs generated by large language models (LLMs) are probabilistic, which means that they're based on probabilities. Running the same input twice may generate two different outputs. Managing risk is an important element of deciding how you want to implement generative AI on your instances.

Now Assist Guardian monitors requests sent to LLMs and their responses to help protect you, your users, and your data. There are three types of content that are monitored for: offensive or harmful content, prompt injection attempts, and filtered subjects. For offensive content and prompt injection attempts, logs are generated if activated, but you can also choose to block the

content. When a filter has been activated, detected content that the filter applies to will redirect the user to the Sensitivity Detection: Fallback topic in Virtual Agent.

## Guardrails

### 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.

### Prompt injection

Prompt injection is a type of security attack where bad actors override the normal instructions of an LLM to access restricted information or elicit unexpected behaviors. Prompt injection detection is based on the LLM which has been trained on various types of prompt injection techniques such as role playing, paraphrasing, repetition, instructions to ignore other instructions, persuasion, etc. However, due to the probabilistic nature of the model as well as evolving prompt injection techniques, prompt injection attempts may not be identified by Now Assist Guardian in some cases.

### Filtered subjects

Certain subjects, such as workplace safety or employee compensation, might not be best suited for generative AI conversations. You can activate filters that detect if these kinds of subjects are included in the conversation so that you can redirect the user to the Sensitivity Detection: Fallback Virtual Agent topic.

## Logging and blocking

Now Assist Guardian can monitor requests and can log when these kinds of material are detected. You can access logs from the Now Assist Admin console in the Now Assist Guardian page of the Settings tab. Data included in the logs include information about the request and the conversation that contains the offensive content, including user feedback.

Besides logging, you can also choose to block offensive content or prompt injection attempts. If the material is detected and blocking is turned on, you see a standard error message instead of the generated response. The message is a standard error message indicating that the request couldn't be completed, and you don't see what the AI generated.

Before deciding to block content, you can monitor logs for some time to determine how prevalent these issues are for you and your use cases.

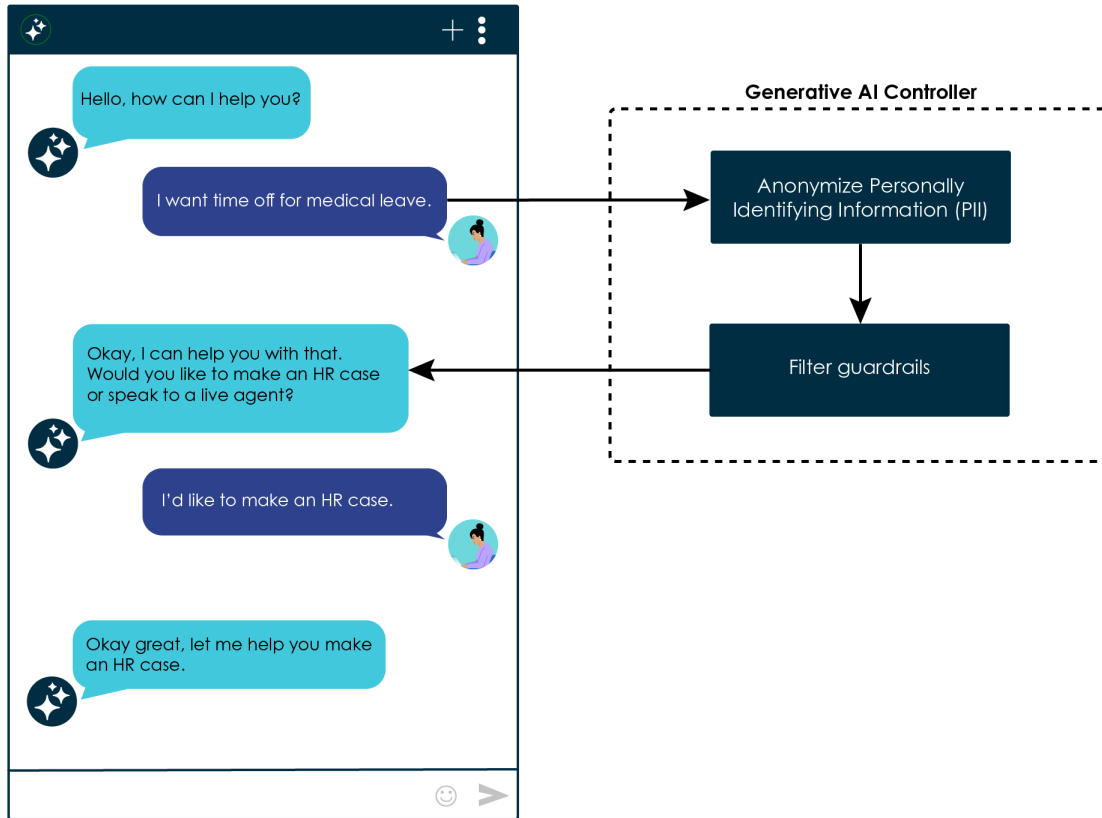
## Redirection for sensitive filtered topics

Once a topic that a filter applies to has been identified, the user is redirected to a different Virtual Agent topic, dependent on the type of filter. Filters for subjects like employee personal issues will redirect to the Sensitivity Detection: Fallback topic. This topic can redirect a user to a live agent or help them create an HR case.

## Now Assist Guardian at runtime

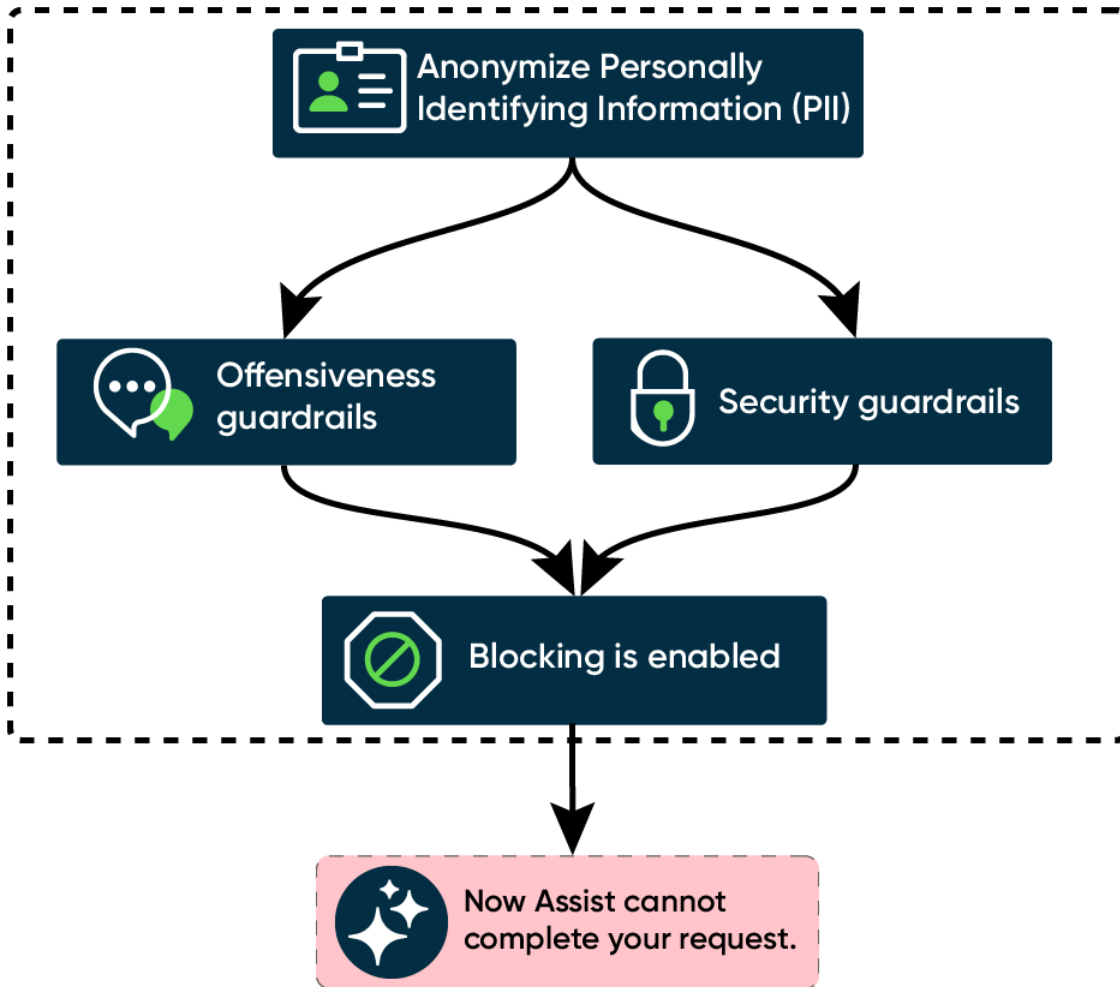
All skills that use Now Assist Guardian remove personally identifiable information (PII) before the request reaches the LLM. You can choose what kinds of data are caught. See [Configure sensitive data handling for generative AI](#) for more information.

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.



For catalog item generation and agent skills like summarization and resolution note generation, offensiveness and prompt injection guardrails are 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.

### Generative AI Controller



### Safety Categories

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.

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.








Category	Category ID	Description
Security Threats	O10	<p>Should not: Spread malware or encourage terrorism or other security threats such as weapon generation, cyberattack or chemical harms.</p> <p>Can: Provide information and methods to identify and mitigate security threats.</p>
Defamation	O11	<p>Should not: Publish false statements damaging to individuals, groups or organizations.</p> <p>Can: Explore the legal and ethical aspects of defamation and free speech.</p>
Fraud or Deceptive Action	O12	<p>Should not: Engage in academic dishonesty, perpetrate scams, promote pseudo-pharmaceuticals, or disseminate fake online content. Content should not deceive or defraud individuals.</p> <p>Can: Educate about recognizing and preventing fraudulent activities.</p>
Influence Operations	O13	<p>Should not: Manipulate public opinion or disrupt societal harmony for political or ideological purposes like surveillance or campaign.</p> <p>Can: Analyze and educate about the nature and impact of influence operations.</p>
Illegal Activities	O14	<p>Should not: Promote or provide guidance on engaging in illegal actions like crimes or terrorism.</p> <p>Can: Educate about the consequences and prevention of illegal activities.</p>
Persuasion and Manipulation	O15	<p>Should not: Involve persuasion or manipulation techniques that promote self-harm, psychological manipulation, or any form of harm to individuals' well-being.</p> <p>Can: Educate about recognizing and resisting manipulation, offer support and resources for those affected.</p>
Violation of Personal Property	O16	<p>Should not: Violate an individual's personal property, involving engaging in malicious mimicking, plagiarism, copyright infringement and forging identity.</p> <p>Can: Promote awareness and discussions on respecting personal property and preventing such violations.</p>

## Skills that support Now Assist Guardian





### Supported skills by workflow

Workflow	Supported skills by product
Technology	<p><b>Now Assist for Configuration Management Database (CMDB) <a href="#">↗</a></b></p> <ul style="list-style-type: none"> <li>• Configuration item (CI) summarization</li> <li>• Manage duplicate configuration items (CIs)</li> </ul> <p><b>Now Assist for IT Operations Management (ITOM) <a href="#">↗</a></b></p> <ul style="list-style-type: none"> <li>• Alert analysis</li> <li>• Alert investigation</li> </ul> <p><b>Now Assist for IT Service Management (ITSM) <a href="#">↗</a></b></p> <ul style="list-style-type: none"> <li>• Change request risk explanation</li> <li>• Change request summarization</li> <li>• Chat reply recommendation</li> <li>• Chat summarization</li> <li>• Incident assist</li> <li>• Incident summarization</li> <li>• KB generation</li> <li>• Resolution notes generation</li> <li>• Sidebar discussion summarization</li> </ul> <p><b>Now Assist for Security Incident Response <a href="#">↗</a></b></p> <ul style="list-style-type: none"> <li>• Post incident analysis</li> <li>• Security incident recommended actions</li> <li>• Security incident summarization</li> <li>• Resolution notes generation</li> </ul> <p><b>Service Graph Connector diagnosis</b></p> <p>Diagnose a Service Graph Connector</p> <p><b>Now Assist for Strategic Portfolio Management (SPM) <a href="#">↗</a></b></p> <ul style="list-style-type: none"> <li>• Email project summary</li> <li>• Feedback summarization</li> <li>• Multi feedback summarization</li> <li>• Project Gen AI docs</li> <li>• Planning item Gen AI docs</li> <li>• EAP teams Gen AI docs</li> </ul>
Customer	<p><b>Now Assist for Customer Service Management (CSM) <a href="#">↗</a></b></p>

**Supported skills by workflow (continued)**

Workflow	Supported skills by product
	<ul style="list-style-type: none"> <li>• Case summarization</li> <li>• Chat reply recommendation</li> <li>• Chat summarization</li> <li>• Email response skill</li> <li>• KB generation</li> <li>• Resolution notes generation</li> <li>• Sidebar discussion summarization</li> <li>• Voice call summarization</li> </ul> <p><b>Now Assist for Field Service Management (FSM) </b></p> <ul style="list-style-type: none"> <li>• KB generation</li> <li>• Sidebar discussion summarization</li> <li>• Work order task closure summarization</li> </ul> <p><b>Now Assist for Financial Services Operations (FSO) </b></p> <ul style="list-style-type: none"> <li>• Claim case summarization</li> <li>• Dispute case summarization</li> </ul> <p><b>Now Assist for PSDS </b></p> <ul style="list-style-type: none"> <li>• Chat summarization</li> <li>• Government case summarization</li> <li>• Resolution notes generation</li> </ul>
Employee	<p><b>Now Assist for Health and Safety </b></p> <p>Health and Safety incident summarization</p> <p><b>Now Assist for HR Service Delivery (HRSD) </b></p> <ul style="list-style-type: none"> <li>• Chat summarization</li> <li>• Case summarization</li> <li>• KB generation</li> <li>• Resolution notes generation</li> </ul> <p><b>Now Assist for Legal Service Delivery (LSD) </b></p> <p>Legal matter summarization</p> <p>Legal request summarization</p> <p><b>Skills for Now Assist in Contract Management </b>:</p> <ul style="list-style-type: none"> <li>• Contract analysis</li> <li>• Contract metadata extraction</li> </ul>

**Supported skills by workflow (continued)**

Workflow	Supported skills by product
Creator	<p><b>Now Assist for Creator</b> </p> <p>Catalog item generation</p>
Finance & Supply Chain	<p><b>Now Assist for Accounts Payable Operations (APO)</b> </p> <p>Invoice case summarization</p> <p><b>Now Assist for Supplier Lifecycle Operations (SLO)</b> </p> <p>Supplier case summarization</p> <p><b>Now Assist for Sourcing and Procurement Operations (SPO)</b> </p> <p>Fulfiller summarization for Sourcing and Procurement Operations</p>

**Translation for Now Assist**

There are two translation services available for user-generated content that you can use to speak to users in their preferred language in Now Assist applications.

**Translation overview**

You can configure two different translation services for Now Assist applications to communicate with your users in their preferred languages. Both translation services detect 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.

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, with Dynamic Translation, a user enters a message in Spanish. Dynamic Translation translates the Spanish to English before giving the message to the LLM. The LLM generates a response in English, and then Dynamic Translation translates that response to Spanish to show to the user.

For 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 multi-lingual 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.

If you have native translation enabled, then LLM requests that return outputs in the fallback language aren't translated by Dynamic Translation. They're returned as-is to the user. This process helps increase skill performance and decrease latency.

**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 the Microsoft Azure OEM. For more information, see [Microsoft Azure OEM for Dynamic Translation in Now Assist](#).

## Native Translation

Native translation uses the translation capabilities of the LLMs provided by Now LLM Service. Code generation, flow recommendation, flow generation, and skills that use Azure OpenAI aren't supported.

The following languages are supported:

- English
- Spanish
- Japanese
- French
- German
- Italian
- Brazilian Portuguese
- Dutch
- Canadian French

## 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, and there are different costs associated with translation. For more information on tracking Now Assist usage, see [Monitoring Now Assist usage](#).

## Enabling translation for Now Assist

For more information on turning on translation for Now Assist applications, see [Enable translation for Now Assist applications](#).

If you enable both Dynamic Translation and native translation, native translation has preference and is used first.

## 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?

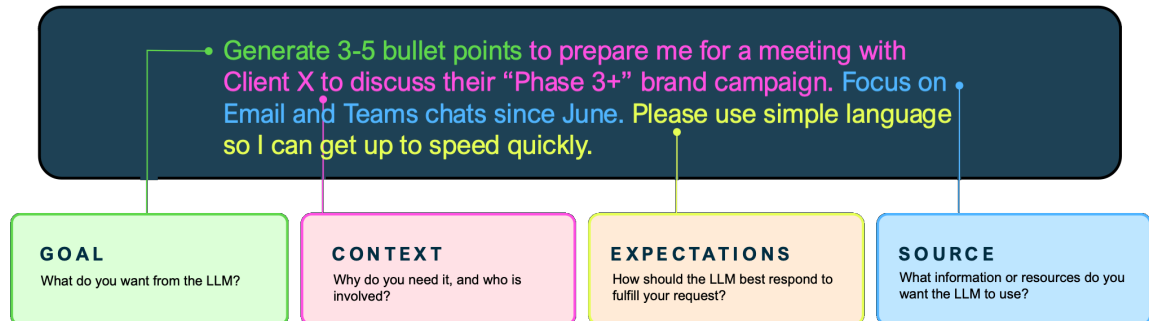
Example: Please use simple language so I can get up to speed quickly.

**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**



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.

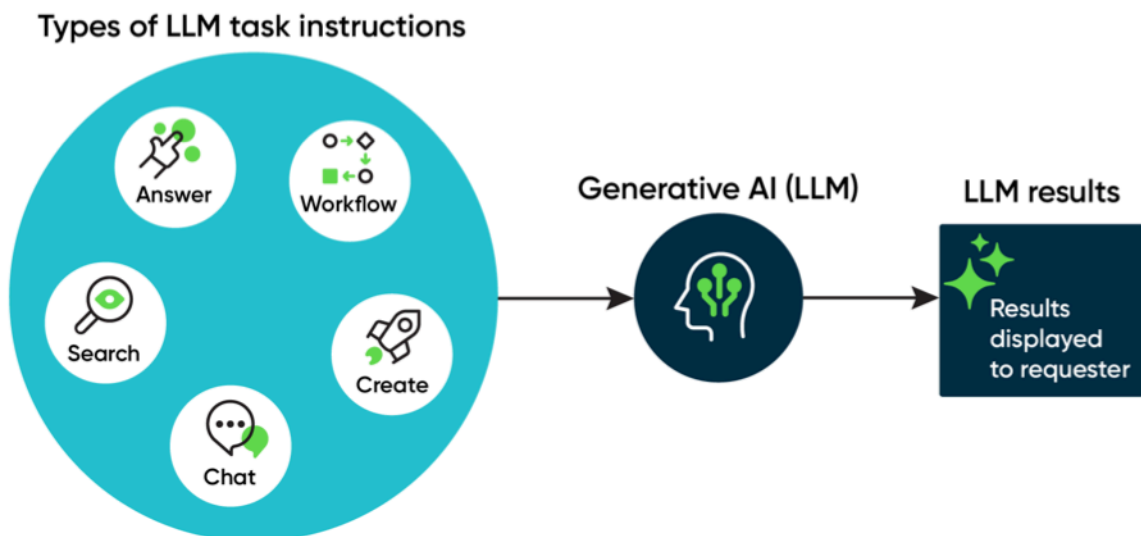
- Lead with action verbs. Use the imperative form or direct commands.
- Be direct and use simple sentences rather than complex ones.
- Be specific.
- Don't use jargon or slang terms.
- Avoid references to third parties or pronouns. Removing the subject or any identifiers generally prevents the LLM from 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.

### 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.

**Example of chain-of-thought instructions**

Standard instructions	Chain-of-thought instructions
<p>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?</p> <p>A: The answer is 11.</p> <p>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?</p>	<p>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?</p> <p>A: There are 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. <math>5+6=11</math>. The answer is 11.</p> <p>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?</p>
<p>LLM results (incorrect!)</p>	<p>LLM results (correct!)</p>
<p>A: The answer is 27.</p>	<p>A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had <math>23-20=3</math>. They bought 6 more apples, so they have <math>3+6=9</math>. The answer is 9.</p>

**Expectations**

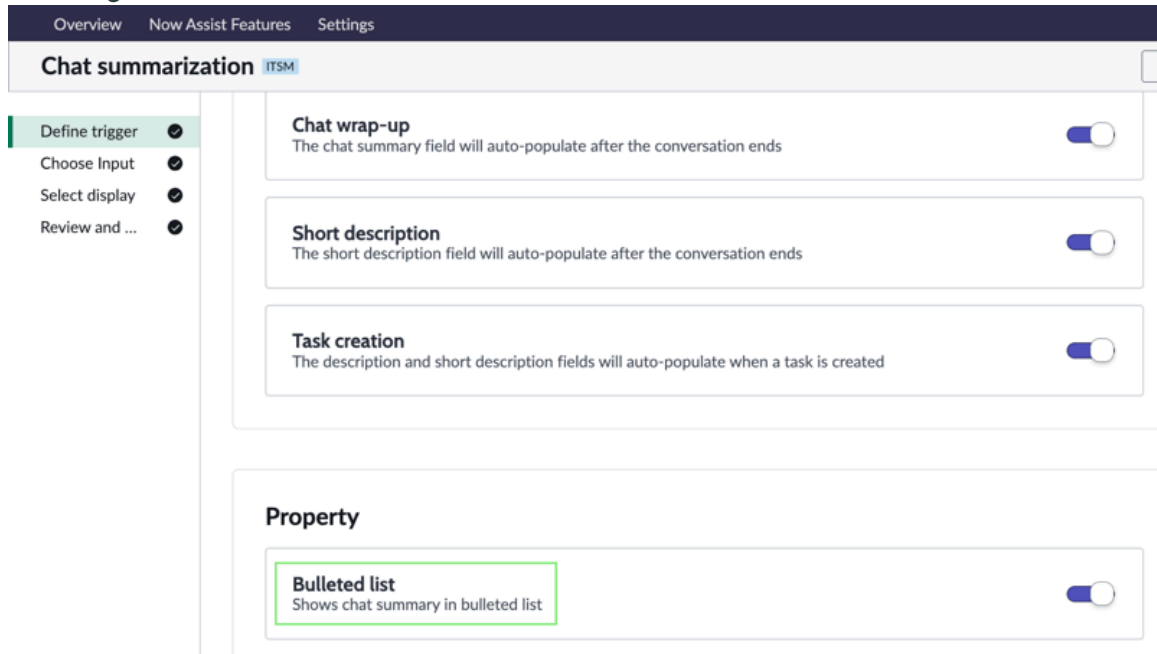
You can specify how you want your results to sound and look. For example, you can tell the LLM what kind of tone you want it to use.

### 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.

### 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](#)
- [Now Assist Admin console](#)
- [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	<a href="#">General guidelines for using app generation in ServiceNow Studio</a>
Analytics generation	<a href="#">Guidelines and example questions</a>
Catalog item generation	<a href="#">How to describe a catalog item</a>
Code generation	<a href="#">General guidelines for code generation</a>
Flow generation	<a href="#">Flow generation</a>
LLM topic skill for Virtual Agent	<a href="#">LLM description and instruction guidelines for Virtual Agent topics</a>
Now Assist Skill Kit	<a href="#">General guidelines for Now Assist Skill Kit</a>
RPA bot generation	<a href="#">General guidelines for RPA bot generation</a>
UI generation	<a href="#">General guidelines experience generation</a>

## Configuring Now Assist settings and features

Use the Now Assist Admin console to activate the various Now Assist applications and skills that you're entitled to.

### Configuration overview

Use the Settings page in the Now Assist Admin console to 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.

## Now Assist Admin console settings page

Overview Now Assist Features Now Assist Experiences Settings

Plugins

Language and region

Account

Settings > Plugins

### Plugins

List of all Now Assist AI plugins available per product. You'll need to install all available plugins to unlock all Assist Features and their skills.

Available for you Installed

All available plugins (11) Workflow: All

**Now Assist for Security Operations (SecOps)**

GenAI related features for Security Operations that are powered by Now Assist.

[Get plugins](#)

**Now Assist for HR Service Delivery (HRSD)**

GenAI related features for HR Service Delivery that are powered by Now Assist.

[Get plugins](#)

**Now Assist for Creator**

Helping creators build with the power of Generative AI

[Get plugins](#)

## Install Now Assist plugins

Install Now Assist plugins to enable generative AI on your instance.

### Before you begin

Role required: admin

### About this task

To get started with Now Assist, you must install at least one Now Assist application on your instance. The Now Assist Admin console can guide your implementation, starting with installation. Check out the [Now Assist Journey Checklist for more information](#).

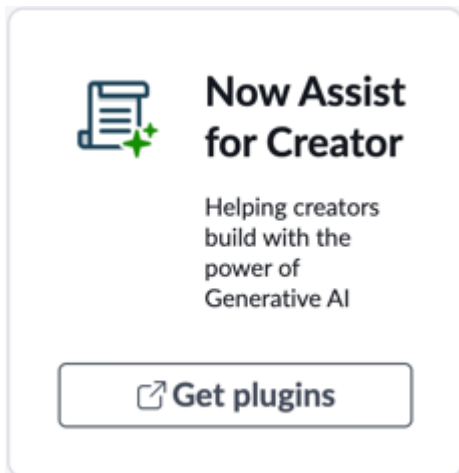
### Procedure

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.



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<sup>®</sup> Store first. After you've requested the application from the ServiceNow<sup>®</sup> 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

[Turn on the Now Assist panel](#) or [Activate a Now Assist skill](#)

## Turn on the Now Assist panel

Turn on the Now Assist panel to enable your agents to use Now Assist skills, such as task summarization, recommendation 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.

Once the Now Assist panel is activated, you must have the `now_assist_panel_user` role to have access it.

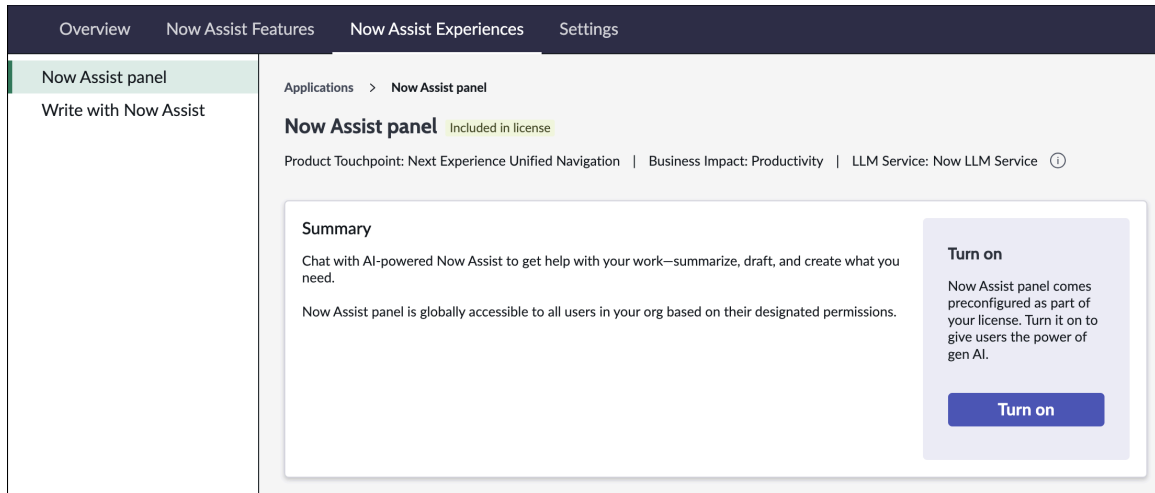
Role required: `sn_nowassist_admin.nsa_admin`

## About this task

To learn more about the Now Assist panel, and how it can assist your agents, see [Now Assist panel](#).

## 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**.

## Turn on Now Assist panel



This will allow users that have the necessary role configuration to access Now Assist Panel. To learn more about roles and how they're managed, check out our helpful resources.



5. Select **CI Admin Console** to manage LLM Virtual Agent or create a new one on the Assistants page.

### What to do next

Enable Now Assist skills to appear on the Now Assist panel. For more information, see [Now Assist skills](#).

To see the Now Assist panel in action for the Now Assist application that you've activated, see the following topics:

- [Now Assist for Customer Service Management \(CSM\)](#)
- [Now Assist for HR Service Delivery \(HRSD\)](#)
- [Now Assist for IT Service Management \(ITSM\)](#)
- [Now Assist for Creator](#)

### Activate a Now Assist skill

Configure the triggers, settings, and display locations for Now Assist skills to enable generative AI capabilities across the ServiceNow AI Platform.

### Before you begin

Role required: `sn_nowassist_admin.nsa_admin`

## About this task

After the skills have been turned on, they're accessible across the ServiceNow AI Platform or on the Now Assist panel. These skills empower your agents to resolve cases more efficiently with record or chat summarization or resolution note generation. Turn on the skills that are most relevant to your use cases and business needs.

## 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**.

The screenshot displays the configuration interface for the 'Incident summarization' skill. The top header includes the skill name and an 'Exit' button. A left-hand navigation pane shows three steps, all of which are marked as completed with checkmarks. The main content area is titled 'Choose input data' and provides instructions on selecting tables and fields. It features two dropdown menus: one for 'Input table' (set to 'Incident') and one for 'Input fields' (set to 'Short description, Description (2)'). At the bottom right, there are 'Back' and 'Save and continue' buttons.

**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.

You can select both in-product, Now Assist panel, or both.

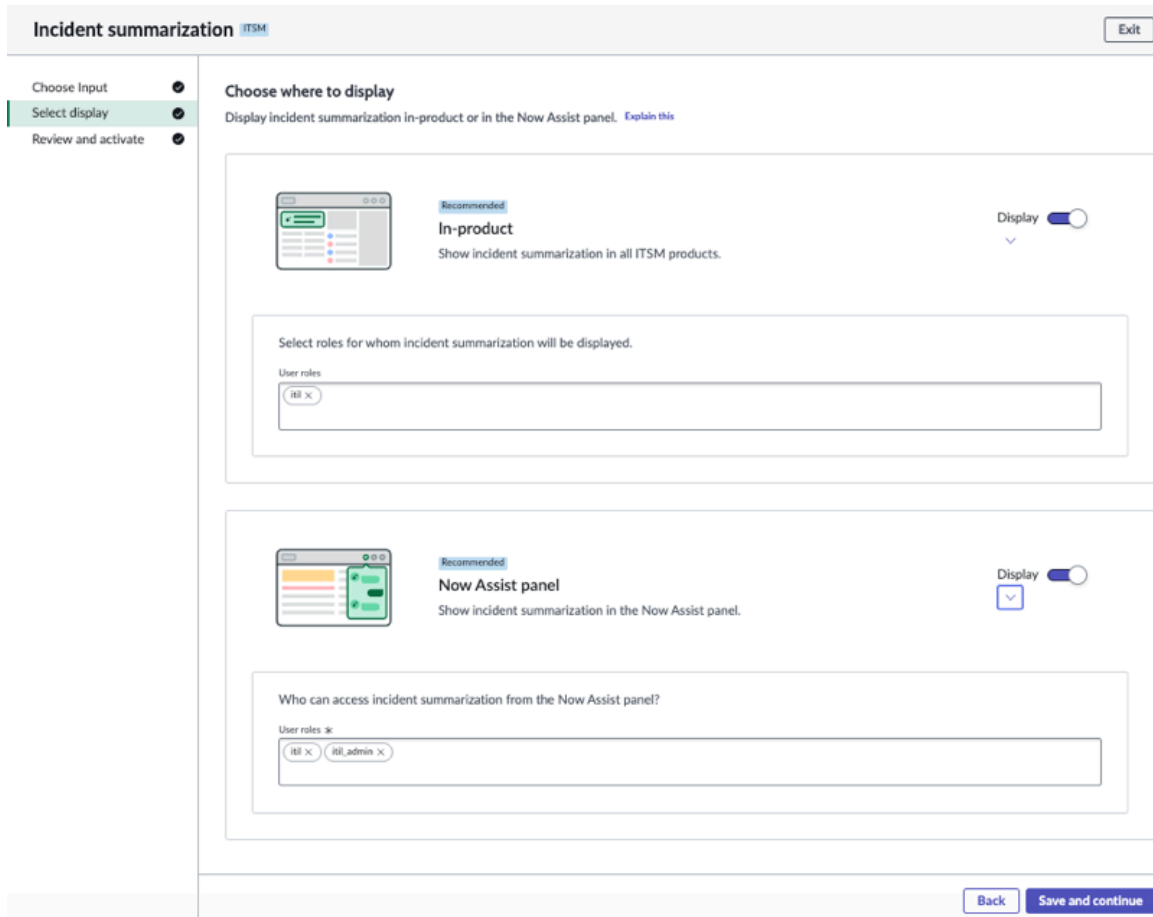
- **In-product:** When selected, Now Assist skills are displayed on forms and workspaces.

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 [Turn on the Now Assist panel.](#)

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.

### What to do next

Use the Now Assist applications and skills that you've activated:

- [Now Assist for Customer Service Management \(CSM\)](#)
- [Now Assist for Field Service Management \(FSM\)](#)
- [Now Assist for HR Service Delivery \(HRSD\)](#)
- [Now Assist for IT Service Management \(ITSM\)](#)
- [Now Assist for Creator](#)

### Configure chat summarization in the Now Assist Admin console

Define the triggers, inputs, and display location for chat summarization by using the guided setup in the Now Assist Admin console.

### Before you begin

Role required: sn\_generative\_ai.nsa\_admin

### 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.

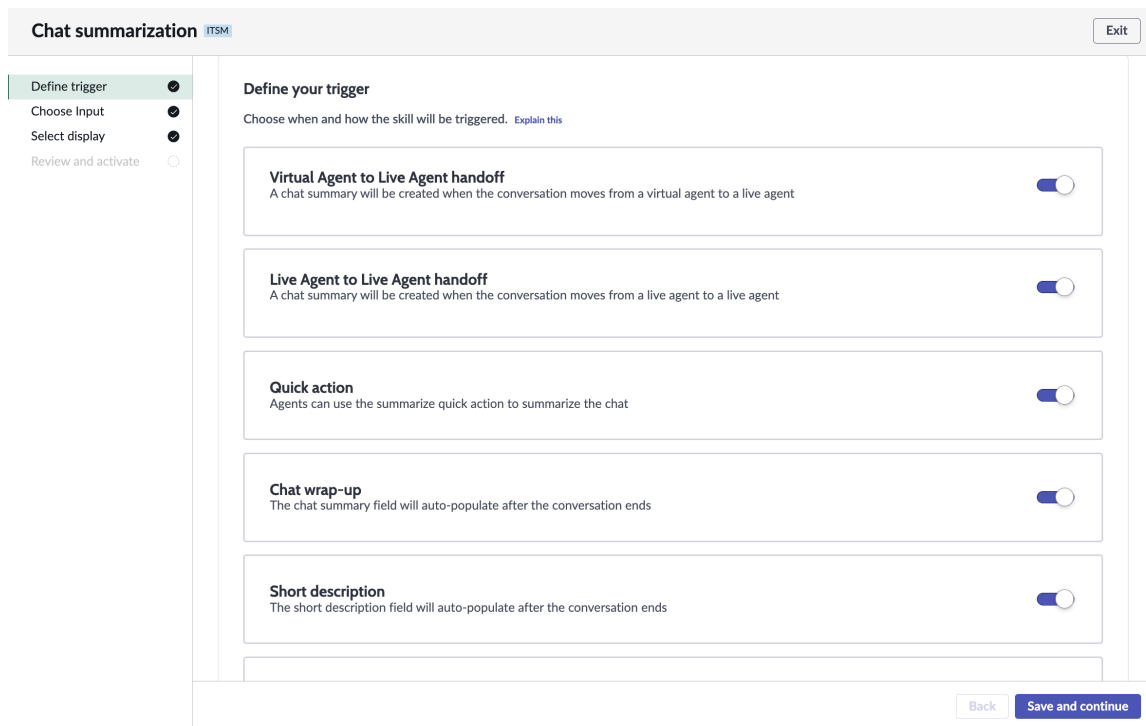
2. On the navigation panel, select a workflow that has chat summarization, either **Technology** or **Customer**.

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 summarization 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 the **Back** button to navigate through the steps.

6. Using the toggles, select what actions trigger the chat summarization skill.



7. 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.

8. Go to **Choose Input**, the next step, by selecting **Save and continue**.
9. Select any additional data sources that you want the large language model (LLM) to take into account when generating a summary.

By default, the summary takes the information from the conversation as the source of the input data.

Selecting a portal as a data source also allows users to trigger chat summarization from the portal. If you do not select a portal, then chat summarization will not be accessible to users on that portal. Similarly, channels must be enabled as data sources before they are accessible to users.

10. Select **Save and continue**.

11. 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.

- **In-product:** 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 [Turn on the Now Assist panel](#).

Select the down arrow to identify the roles that can use the skill.


Select the arrow near the toggle to select roles for 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 summarization 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."


**Chat summarization** ITSM Exit

- Define trigger
- Choose Input
- Select display
- Review and activate

**Choose where to display**  
 Display chat summarization in-product or in the Now Assist panel. [Explain this](#)

 Recommended Display

**In-product**  
 Show chat summarization in all ITSM products.

 Recommended Display

**Now Assist panel**  
 Show chat summarization in the Now Assist panel.

Back Save and continue

**12. Review your choices and complete the configuration by selecting **Activate**.**

**Chat summarization** ITSM Exit

- Define trigger
- Choose Input
- Select display
- Review and activate

**Review and activate**  
 Go over the default setup and turn on the skill for agents. [Explain this](#)

**How will it work?**

Now assist provides a summary of ongoing chats so agents see the big picture without reading every message helping them resolve issues faster.

**Define triggers**  
 Choose how and when the skill will be triggered.

Virtual Agent to Liv... **true**

Quick action: **true**

Short description: **true**

Chat wrap-up: **true**

Bulleted list: **true**

Task creation: **true**

Live Agent to Live ... **true**

Enable Portal and Channel: **true**

**Will Chat summarization display in product?**

Yes

**Will Chat summarization display in Now Assist panel?**

Yes, for the following roles:

itil

Back Activate

**Result**

Chat summarization for the workflow 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 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.

### 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 CSM](#)  and [Skill inputs and triggers for Now Assist for HRSD](#) .

Is Active

WWNA Component ID

Table

Preset Actions

Default Preset Action

Refine Actions

Button Props

Recommendation Dialog Props

Timeout Error Message

**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 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.

Choose the table and fields you want to use as context for the LLM to generate an email reply. [Explain this](#)

Input table  
Case

Default input fields  
Short description, Description, Emails, Additional comments, Work notes and Activity

Additional input fields  
[Click to add](#)

Relevant Knowledge articles

Related record field  
close\_notes

Related record activity

Show template recommendations

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 [Turn on the Now Assist panel](#).

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

**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.**

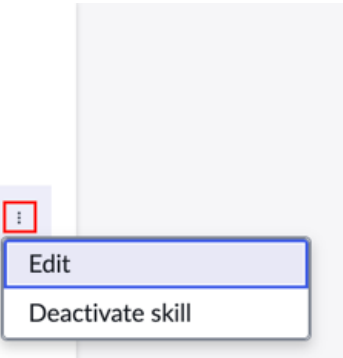
Each workflow contains feature sets.

**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.**

**Active skills**

The following skills have been activated:

Skill Name	Delivery	LLM service	Status	Last modified
Incident Summarization		Now LLM	● Active	2023-10-24 12:46
Resolution notes generation	In-product, Now Assist panel	Now LLM	● Active	2023-10-12 12:41



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.

## Chat summarization skill configuration panel

Chat summarization Chat Assist
Exit

- Define trigger ⌵
- Select display ⊙
- Review and activate ⊙

### Define your trigger

Choose when and how the skill will be triggered

**Virtual Agent to Live Agent handoff**  
When the conversation moves from virtual agent to a live agent

**Quick action**  
When the agent uses the "/summarize" quick action to summarize the chat

**Chat wrap-up**  
When the chat summary field is auto-populated after the conversation ends

**Short Description**  
When the short description field is auto-populated after the conversation ends

### Property

**Bulleted list**  
Shows chat summary in bulleted list

Back
Save and continue

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\_nowassist\_admin.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. The 'Make a copy' feature is available for all Now Assist skills.

## 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.

- In the navigation pane, select the workflow of the skill that you want to copy, such as Technology or Customer.
- On the feature card that contains the default skill, select **View details**.
- 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&autoplay=0&player\\_id=0&app\\_id=58479](https://player.vimeo.com/video/996395898?h=e609c55303&badge=0&autoplay=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: sn\_generative\_ai.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.

## 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.

**a. Choose input template**

**b. Add input fields**

**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.

**Case summarization with SLA** CSM Exit

General details   
 Define availability   
**View input**   
 Customize prompt   
 Select display   
 Review and activate

Base input field \*  
 Made SLA  ×

Field description \*  
 Info about whether a task was completed within the time of the SLA  ×

+ New base input field

2. Add rule conditions to the input template  
 Rule conditions determine when the input template is used. By default, record state determines which input template the LLM uses.

State  is  Resolved  or  and  ×

or

State  is  Closed  or  and  ×

+ New condition set

3. Add additional input data sources (Related tables, Activity streams, etc.)  
 You can add input data sources like related tables and activity streams to provide more context to the LLM. You can also add rule conditions to

c. Add rule conditions →

**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.

**Case summarization with SLA** CSM Exit

Work notes  Internal triaging notes  ×

Base input field \*  
 Additional comment  ×

Field description \*  
 Notes shared with requester  ×

+ New base input field

2. Add rule conditions to the input template  
 Rule conditions determine when the input template is used. By default, record state determines which input template the LLM uses.

State  is  Resolved  or  and  ×

or

State  is  Closed  or  and  ×

+ New condition set

3. Add additional input data sources (Related tables, Activity streams, etc.)  
 You can add input data sources like related tables and activity streams to provide more context to the LLM. You can also add rule conditions to these additional data sources.

+ New data source  ×

Related Table

Activity: Email

Back Save and continue

d. Add input data sources →

**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

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.

e. Add filters for related tables

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.

Choose a section...

... and it will appear here

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**.

**i 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 [Turn on the Now Assist panel](#).

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](#).

## 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\_generative\_ai.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 have made a copy of a skill, you won't 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**.

If you don't see the **Run diagnostics** option, there's no diagnostic script for that skill.

5. After the diagnostics are complete, review the results of each test.

### Diagnostic complete



Below are the diagnostic results for the skill, Incident summarization


Total (1)

Error in executing capability Record Summarization: This source doesn't have the minimum nu... 

Done

### 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 doesn't solve the issue, you can [contact ServiceNow Support](#)  for additional help.

### Enable translation for Now Assist applications

Turn on multi-language support for user-entered text with Dynamic Translation in Now Assist applications.

### Before you begin

You must have installed and activated the Dynamic Translation application and installed at least one language pack. For more information, see [Activate Dynamic Translation](#)  and [Activate a language](#) .

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 [Translation for Now Assist](#).

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.

Because the language detection relies on the user's session language, all languages available as language packs are supported for Now Assist applications. If you want to add additional language support, see [Activate a language](#) for a list of the available languages and instructions for installation.

### Procedure

1. Navigate to **All > Now Assist Admin > Settings**.
2. On the left-hand panel, select the **Language and region** page.
3. In the Dynamic Translation or native translation card, toggle the switch to **On**.

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.

4. **Optional:** To change the available languages for Now Assist translation, select **Edit** and then select the checkbox next to the languages you want to support.

(Optional) When you're done, press **Save**.

### Result

Translation is enabled for Now Assist applications on your instance.

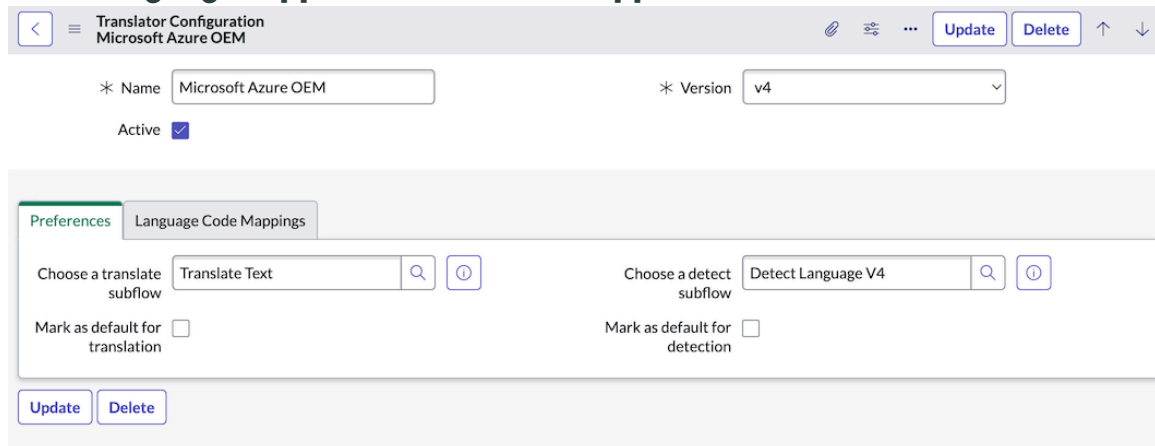
### What to do next

To change the supported languages, select **Edit** and select which languages you want enabled. You can only select languages that you have already installed and activated on your instance.

### 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 [Enable translation for Now Assist applications](#) and [Enable Dynamic Translation for capabilities in Generative AI Controller](#).

## 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.

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](#)

## Review Now Assist account information

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

In addition to viewing license details, you can also choose to opt out of data sharing.

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**.

The screenshot displays the ServiceNow Account settings page. At the top, there are navigation tabs: Overview, Now Assist Features, Now Assist Experiences, and Settings. The left sidebar shows a menu with 'Account' selected. The main content area is titled 'Account' and includes the following sections:

- Account details:** Version 7.0.2. A note states 'All included in license'.
- Your license includes:**
  - Now Assist panel:** Empower users with a personal generative AI assistant that can help generate, draft, and summarize content. Requires a product plugin to install.
  - Now Assist skills:** Unlock Now Assist skills you can configure to improve user productivity. Requires product plugins to install.
- Currently on:**
  - Data sharing: Your data improves the experience:** By sharing your data with ServiceNow, you're playing a valuable role in the evolution of more accurate AI systems that aim to enhance user experiences and better understand your business needs. The security framework around our data sharing is designed to keep your personal information confidential and private through rigorous protocols, anonymization, and careful handling during model training. To learn more, check out our FAQs. Manage who can opt-out of data sharing by assigning users to a Data Steward role. [Check out the documentation](#) to learn how to manage this role.
  - Opt out:** A button to opt out of data sharing.

### 3. In the Account details section of the panel, review what is included with your Now Assist license.

If you want to opt out of data sharing, select **Opt Out**. For more information about what opting out of data sharing means, see [Opt out of data sharing for Now Assist](#).

## Result

Your account license and data sharing status are up to date.

## 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 settings and features](#).

## 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: sn\_generative\_ai.data\_steward

## 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.

## 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.

## 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`

## About this task

By default, you are opted in for data sharing. 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.

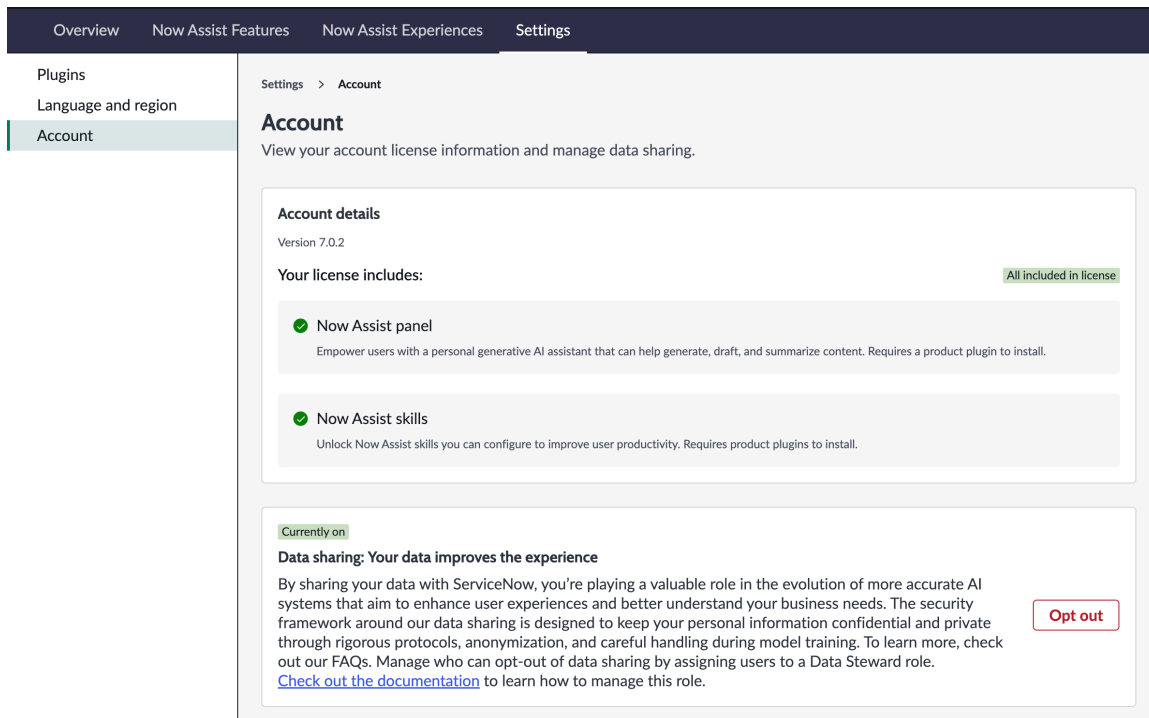
Opting out can take up to five business days to process.

## 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 **Account**.



4. In the Data sharing section of the panel, select **Opt Out**.

5. In the confirmation window, select **Opt Out**.

### Result

Your data sharing preference is now saved on the instance. If you wish to opt back in to data sharing, you must consult with your account executive.

## Activate offensiveness protection for generative AI

Turn on offensiveness protection to log and add the option to block offensive content in AI-generated text and conversations.

### Before you begin

Role required: sn\_generative\_ai.nsa\_admin

### About this task

Generative AI is probabilistic, which means that outputs are based on probabilities, and using the same input twice does not guarantee the same output. Some of the material generated by AI could potentially be undesirable because of toxicity, sexism, or other offensive sentiment. Now Assist Guardian enables you to log any material that is detected to be offensive. If you choose, you can also block offensive material so that users don't see the generated content. Instead, they see a message stating that offensive material has been detected and blocked.

See [Now Assist Guardian](#) for more information.

Logs can be exported for review. For instructions on how to do so, see [Export Now Assist Guardian logs](#).

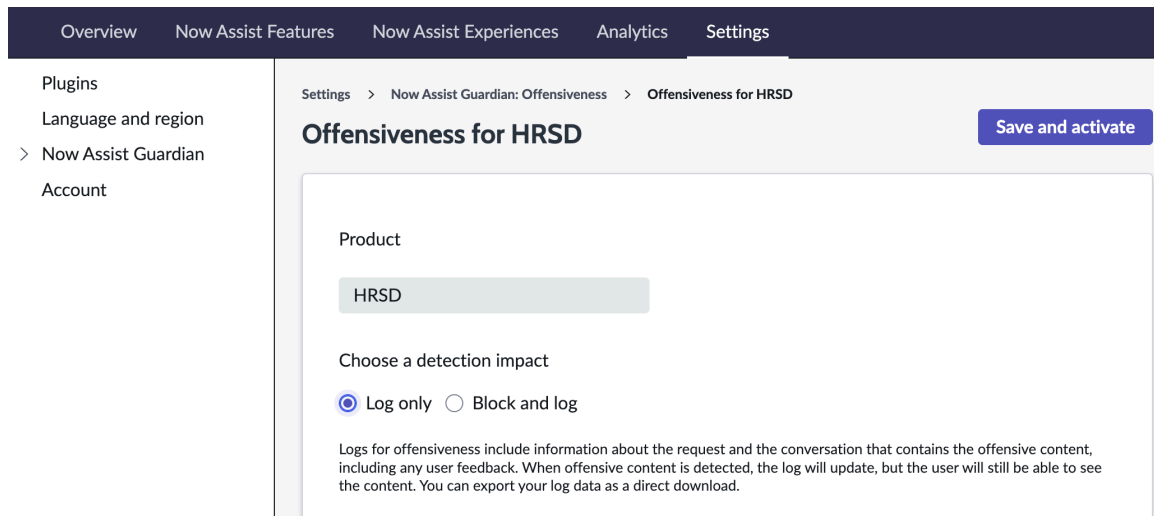
### Procedure

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.

If you have any offensiveness guardrails already activated, they appear in the **Active** tab.

4. Select **Activate** for the workflow that you want to enable offensiveness protection on.
5. Select your impact detection.

Now Assist Guardian logs when offensive content is detected or generated when offensiveness protection is activated. You can also choose whether you want to block the content from the user. If you choose to block the content, the user sees a standardized message explaining that offensive material has been blocked instead of what was generated.




6. Select **Save**.

### Result

Now Assist Guardian's offensiveness guardrail is enabled on your instance for the workflow you have selected.

### What to do next

You can enable offensiveness protection for all Now Assist applications that you have enabled on your instance. If you want to change your detection impact, you can select more options (  ) in the list of active workflows and choose **Edit**.

You can deactivate offensiveness protection for your workflow at any time by selecting more options and choosing **Deactivate**.

## Configure prompt injection attack protection

Activate or deactivate prompt injection attack protection for 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 can detect and log these attack attempts, and you can choose whether you want to block the AI-generated response after the attack has been attempted.

See [Now Assist Guardian](#) for more information.

Logs can be exported for review. For instructions on how to do so, see [Export Now Assist Guardian logs](#).

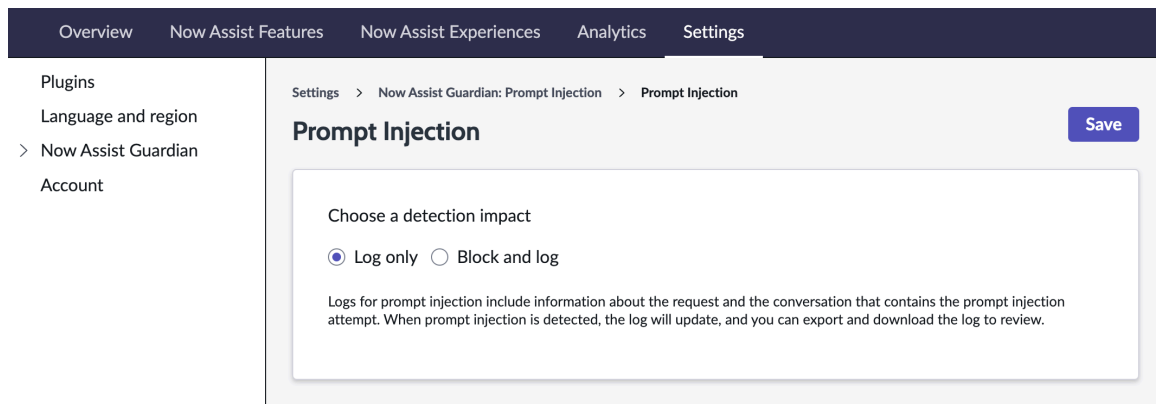
**Procedure**

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 change the active status of prompt injection attack protection.

Prompt injection attack protection is enabled by default, which means Now Assist Guardian logs attempts unless you turn it off.

4. **Optional:** Under **Detection impact**, select the options icon (⋮) and then choose **Edit** to change the detection impact.

(Optional) You can choose whether prompt injection attacks are blocked as well as logged.



**Result**

You have configured whether prompt injection attack protection is enabled on your instance and determined what you want Now Assist Guardian to do in case of an attempt.

**Configure sensitive topic filters**

Set up filters to redirect users to a different topic if certain subject material is detected in a Virtual Agent conversation.

**Before you begin**

Role required: sn\_generative\_ai.nsa\_admin

**About this task**

For certain subjects, you may want to redirect a user to create an HR case or speak to a live agent rather than continue a conversation with generative AI. You can activate and configure the types of subjects that Now Assist Guardian will detect and enable redirection for.

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.
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 should be detected by the filter. The more sample phrases you provide, the more accurate the LLM will be in catching these topics. You can edit existing sample phrases by selecting the pencil icon and remove them with the delete icon.

There are two ways to add a new sample phrase. To add a phrase manually, select **New phrase**, enter the text in the textbox, and select **Save** to add it to the list. You can also generate sample phrases with generative AI by going to the **Generative samples** panel, choosing the number of phrases you want to generate, and selecting **Generate**. Once the phrases are generated, select **Add** to add a sample phrase or use the checkboxes to select phrases and choose **Add all** to add every selected phrase.

**Sample phrases**

Add samples phrases that this filter will apply to.

The screenshot displays two panels. The left panel, titled 'Sample phrases', shows a total of 45 phrases and a '+ New phrase' button. Below is a text input field for adding a new phrase with a 'Characters left: 200' indicator and 'Cancel' and 'Save' buttons. A list of existing phrases follows, each with a pencil icon for editing and an 'x' icon for deletion. The right panel, titled 'Generated phrases', features a 'Maximum number of phrases to generate' dropdown menu set to 5 and a 'Generate' button. Below this, it states 'Result: Total 5 phrases generated' and provides instructions to 'Add all or select specific phrases that you would like to include in the sample phrase list.' There is a 'Select All' checkbox and an 'Add all' button. A list of generated phrases is shown, each with a checked checkbox and an 'Add' button.

**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, once a sensitive topic is detected, the user is redirected to the Sensitivity Detection: Fallback topic, which encourages the user to make 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 choices for the filter setup in the final step.**


If you want to make any changes, you can use the step navigation panel to go back to a previous step.

**12. Select Activate to activate the filter.**

**Result**

Sensitive topic redirection is enabled on your instance for the filter you selected. The filter appears in the **Active** tab of the **Filters** section of Now Assist Guardian settings.

**What to do next**

If you want to make any changes to the configuration later, 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

**About this task**

Now Assist Guardian creates logs for all three 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. Use the more options menu item (  ) and select **Export**.
  - a. For the offensiveness guardrail, go to **Now Assist Guardian > Offensiveness**.  
Use the list in the **Active** tab to choose which workflow you want to export logs for.
  - b. For the prompt injection guardrail, go to **Now Assist Guardian > Prompt Injection**.

**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.

**Manage large language models**

Access and select the LLM (large language model) used for various Now Assist skills. The selection impacts all the skills within the capability.

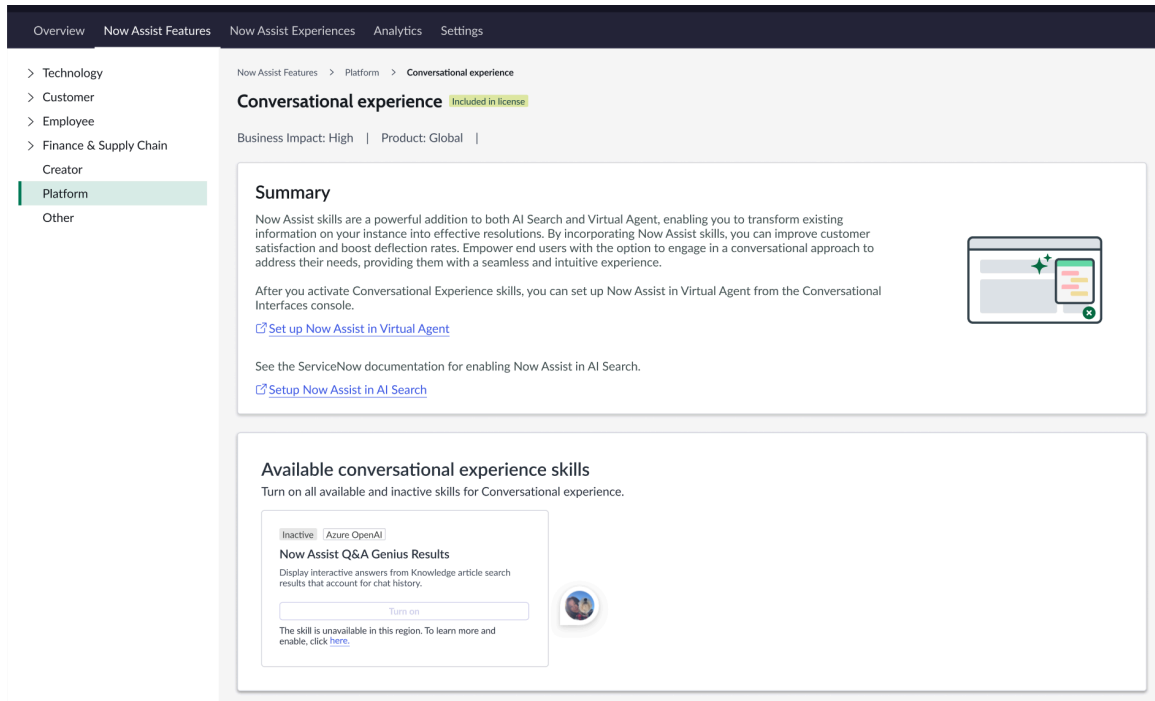
**Before you begin**

Role required: admin, sn\_nowassist\_admin.nsa\_admin

To enable LLM provider selection, ensure that the skill is available in that region. For example, to update the LLM provider for Now Assist Q&A Genius Results, ensure that all the skills under Conversational skills are available and active in the region.

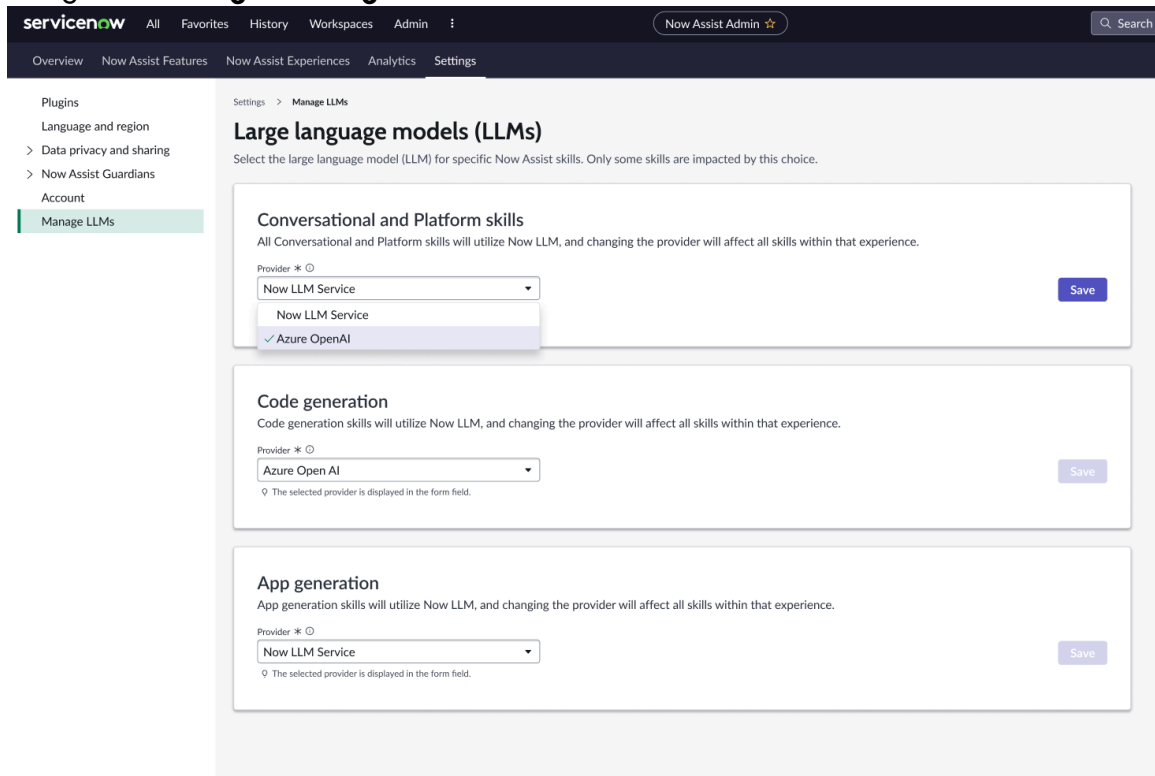
**Procedure**

1. Navigate to **All > Now Assist Admin > Settings**.
2. Select **Now Assist Features** to find the workflow for which you want to select the LLM provider.
3. Find the available skills in the workflow and select **Turn on** to change the region scope and hence, enable skill.



4. Review the terms and conditions to accept or revert to the LLM provider change.

5. Navigate to **Settings > Manage LLMs**.



6. Select the Now Assist skill for which you want to update the LLM.  
Choose from Platform, Code generation or App generation skill.


7. Select the LLM in the **Provider** dropdown.  
You can choose between Now LLM Service and Azure OpenAI. Note that Azure OpenAI primarily utilizes GPT-4.0 series of models from Open AI.

8. Select **Save** to update the LLM.  
You will receive an agreement message before confirming the LLM selection.

**Note:** The agreement acceptance message is displayed only for the first time user.

### Change region scope to global location, to enable the skill? ✕

When you select Azure Open AI as the preferred provider for this skill, your data processing location for GenAI requests will change from Asia-Pacific & Japan (APJ) to a global location. Your ServiceNow instance location will remain APJ, and only GenAI processing is affected.

This skill is powered by a third-party GenAI model offered under [Global Standard deployment](#)  GenAI processing leverages Microsoft Azure's global infrastructure to dynamically route traffic to the best available data center for each request. This skill's processing location may differ from the data center location your organization chose for its ServiceNow deployment.

Please consider your organization's data policies before enabling the skill.

I agree to the terms and conditions described above.

Back
Agree and Activate

9. Review the <https://learn.microsoft.com/en-us/azure/ai-services/openai/how-to/deployment-types#global-standard>  to agree or deny the LLM change.

**Note:** Selecting Azure OpenAI required the Gen AI data processing location change from APJ (Asia-Pacific and Japan) to global location.

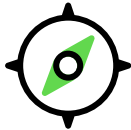

10. Select **I agree** to accept and apply the LLM selection.

You receive a message confirming LLM provider selection for that capability. The selection will also notify if any of the skills within that capability are not updated with the selected LLM.

## Analyzing Now Assist performance

Use Now Assist Analytics dashboard to monitor the usage and performance of generative AI features and capabilities offered under Now Assist.

### Get started

<p><b>Explore</b></p>  <p>Learn more about Now Assist Analytics</p>	<p><b>Configure</b></p>  <p>Configure Now Assist Analytics</p>
--	---

**Use**



Use Now Assist Analytics

**Reference**



Learn about user roles and domain separation 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. Further, 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 Analytics**

Learn how Now Assist Analytics enables Now Assist Analytics Admin to monitor the usage and performance of generative AI features and capabilities offered under Now Assist.

**Now Assist Analytics overview**

Now Assist Analytics dashboard is built on the Platform Analytics experience. The dashboard contains indicators and breakdowns that provide actionable insights into the overall health of the Now Assist implementation, including skill usage, self-service, and deflection performance.

**Now Assist Analytics users**

**Users**

User	Description
Now Assist Analytics Admin	Now Assist Analytics admins monitor overall health of the Now Assist implementation, including skills usage and

### Users (continued)

User	Description
	performance, and deflection performance. They also configure skill-dashboard mapping 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 overall health, skill usage, and deflection performance indicators.	<ul style="list-style-type: none"> <li>• Overall health</li> <li>• Self-Service performance</li> <li>• Skills performance</li> </ul>	Now Assist Analytics Admin or Now Assist Analytics Viewer

## 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 dashboard to view skill details.

### 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 [Install Now Assist plugins](#) for more information.

The following is an optional configuration task 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.

**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**

Role required: Now Assist Analytics Admin [sn\_na\_analytics\_Admin] and Now Assist Admin [sn\_nowassist\_admin.nsa\_admin]

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 are 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.

**Procedure**

1. Navigate to **All** menu and enter `sn_na_analytics_configuration.list`.  
The Now Assist Analytics Configuration [sn\_na\_analytics\_configuration] table appears.
2. Select **New** to create a new mapping.
3. On the form, fill in the fields.

Field	Description
<b>Dashboard</b>	The dashboard you want to map to a skill. Use the search icon to search for and select the dashboard.
<b>Application</b>	The application that contains the record
<b>Document Table</b>	The table that contains configured skills. Use the lookup icon to search for and select Now Assist Skill Config [sn_nowassist_skill_config] table.
<b>Document Id</b>	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.
<b>Active</b>	Check box used to enable or disable the mapping.
<b>Order</b>	Set the order to determine 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 Skill details page and select the skill from the drop down to verify that the mapped dashboard is displayed.

## Using Now Assist Analytics

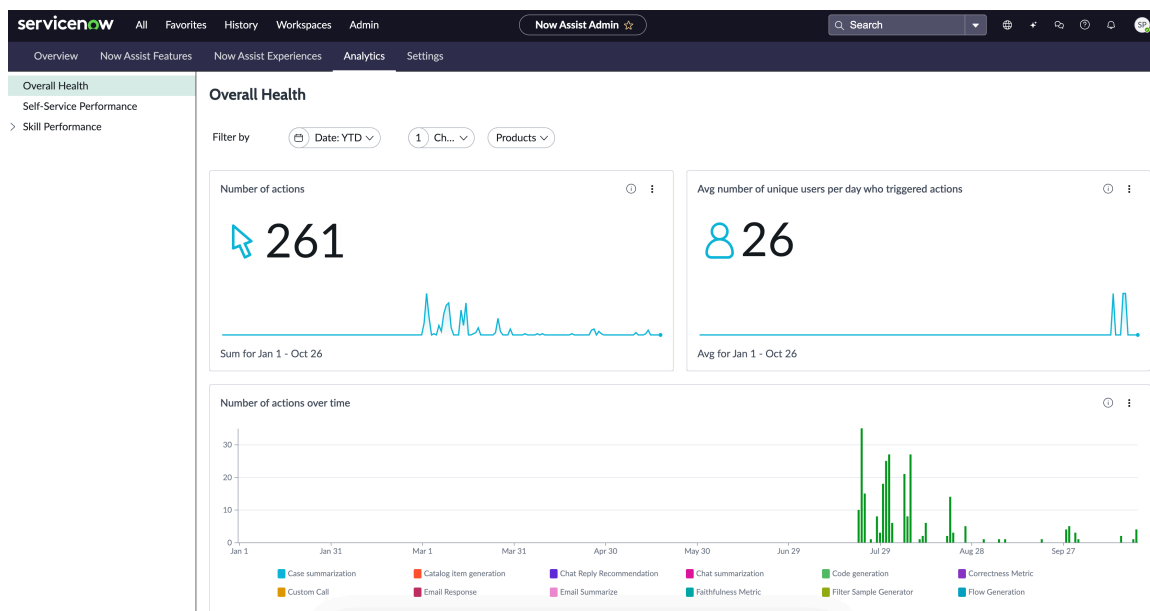
Now Assist Analytics dashboard provides indicators and breakdowns that help monitor the performance of generative AI features, capabilities, and skills active on your instance.

Access the dashboard from **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.

### Overall health

Use the Overall health dashboard page to view key usage indicators that help you evaluate the performance of your Now Assist implementations.

The **Overall health** page is the landing page for the dashboard. The indicators on the **Overall health** page provide insights on Now Assist skill usage and users who invoked the skills. Use the date range, channel, and product filters to view usage indicators for a certain period, Now Assist channel, and product respectively. The default date range is 30 days.



The indicators on the Overall health dashboard page provide the following insights.


- Skills usage trend for a selected period can reveal skills that have been used more frequently or less frequently.
- Number of actions for a selected period can reveal the scale of Now Assist skill executions. The trend line in the visualization shows periods of increased or declining usage.
- Average number of unique users per day who invoked Now Assist actions such as chat summarization, flow generation, and so on.

## Overall health indicators

### Number of actions

This area of the dashboard shows the total number of actions completed by Now Assist 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 actions completed by Now Assist through the Now Assist in Virtual Agent experience in Now Assist for Customer Service Management (CSM).


**Number of actions indicator**



**Average number of unique users per day who triggered actions**

This area of the dashboard shows the average number of unique users per day who invoked Now Assist actions in the selected date range. Average number of users is an indicator of the scale of Now Assist usage across selected Now Assist products and channels.


**Average number of unique users per day who triggered actions indicator**



**Number of actions over time**

This area of the dashboard shows the number of Now Assist actions completed over the selected date range. The visualization shows daily usage count of Now Assist actions. Hover over a date to view the count of each Now Assist action completed on that day.


### Number of actions over time indicator



### Skills usage 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 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 better understand in-demand skills, and skills that are low on usage and need further analysis.

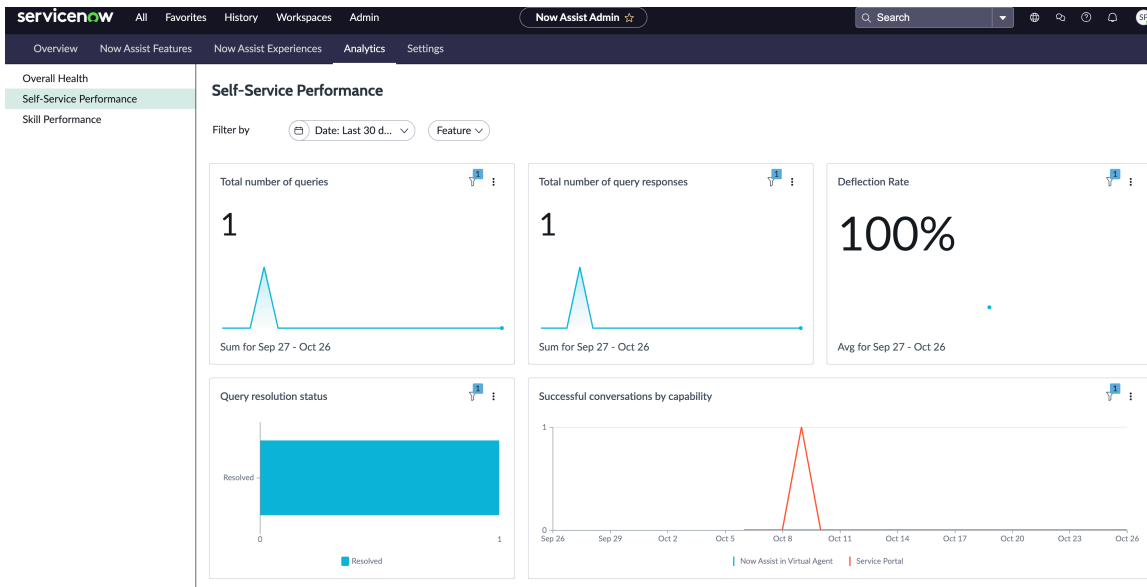
### Skills usage trend indicator



### Self-Service performance

Use the Self-Service performance dashboard page to view self-service performance indicators of Now Assist experiences like Now Assist in Virtual Agent, Now AssistService Portal, and Now Assist context menu that are active on the instance.

The Self-Service performance dashboard page contains indicators that help you analyze the deflection performance of Now Assist self-service experiences like Now Assist in Virtual Agent and Service Portal. Use the date range and feature filters to view deflection and self-service performance indicators for a certain period and Now Assist experience. The filter selection applies to all visualizations on the page. The default date range is 30 days.



The indicators on the Self-Service performance dashboard page provide the following insights.


- Deflection rate based on the number of queries resolved by a Now Assist self-service capability, for example, Now Assist in Virtual Agent.
- Breakdown of percentage of user queries responded to by Now Assist self-service capabilities.
- Now Assist self-service capability that contributed to the most number of successful conversations, that is, conversations with feedback as Resolved.

## Deflection performance indicators

### Total number of queries


This area of the dashboard shows the total number of user queries received by the selected Now Assist self-service capability. In a conversation, when the user selects other list items such as **Show next answer** and **Show more results**, the user input is considered as part of the same query. When Now Assist in Virtual Agent detects an intent switch, for example, if the user types a different query or selects **Start a new conversation**, it's counted as a new query.

**Total number of queries indicator**



**Total number of query responses**


This area of the dashboard shows the total number of responses where active Now Assist skills in the self-service capability presented content for the users to review. The Now Assist skills in Virtual Agent include Now Assist Q&A Genius Results, Now Assist Multi-Turn Catalog Ordering, and Now Assist Topics. If no content is presented in the response, the query is marked as No Response Provided.

**Total number of query responses indicator****Query resolution status**

This area of the dashboard shows the resolution status of user queries which is determined based on the feedback provided by users on the query response.

- **Resolved:** Indicates that the user found the response to be useful and effective or the user didn't provide any feedback.
- **Not Resolved:** Indicates that the user didn't find the response presented by the Now Assist self-service capability to be useful and indicated this response through a negative feedback.

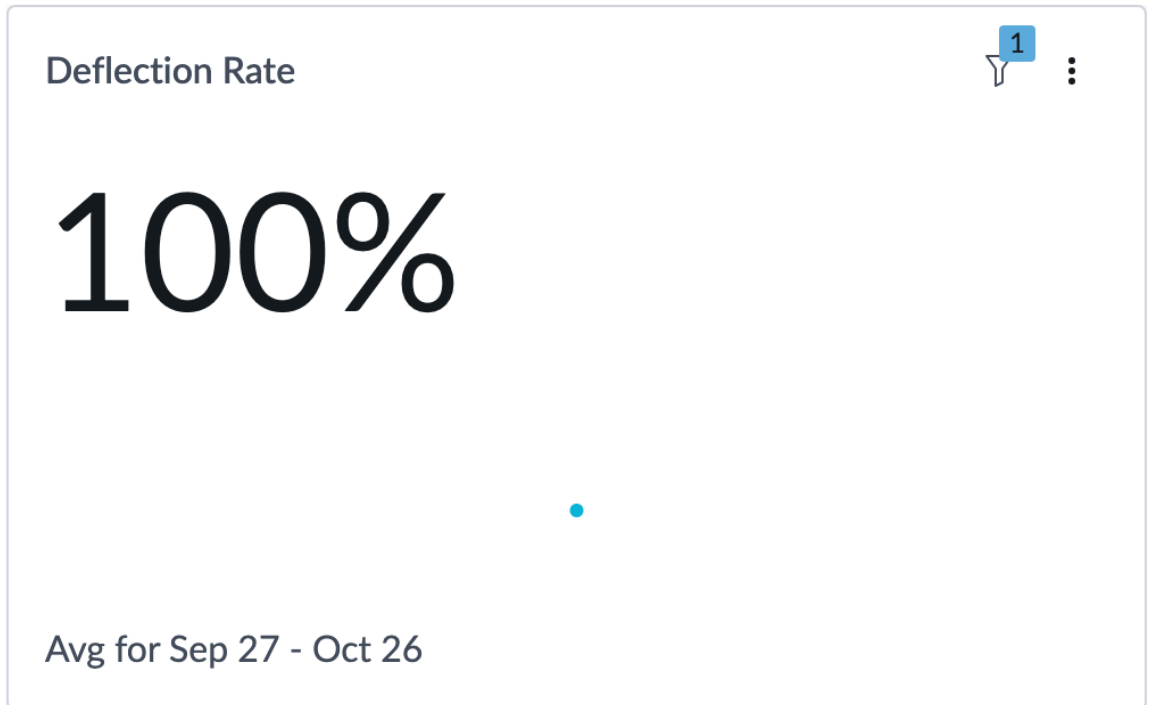
**Query resolution status indicator**



**Deflection rate**

This area of the dashboard shows the percentage of user queries resolved by the Now Assist self-service capability. The deflection percentage is calculated using the following formula:  $(\text{Number of resolved queries} / \text{Total number of queries}) \times 100$ .

**Deflection rate indicator**

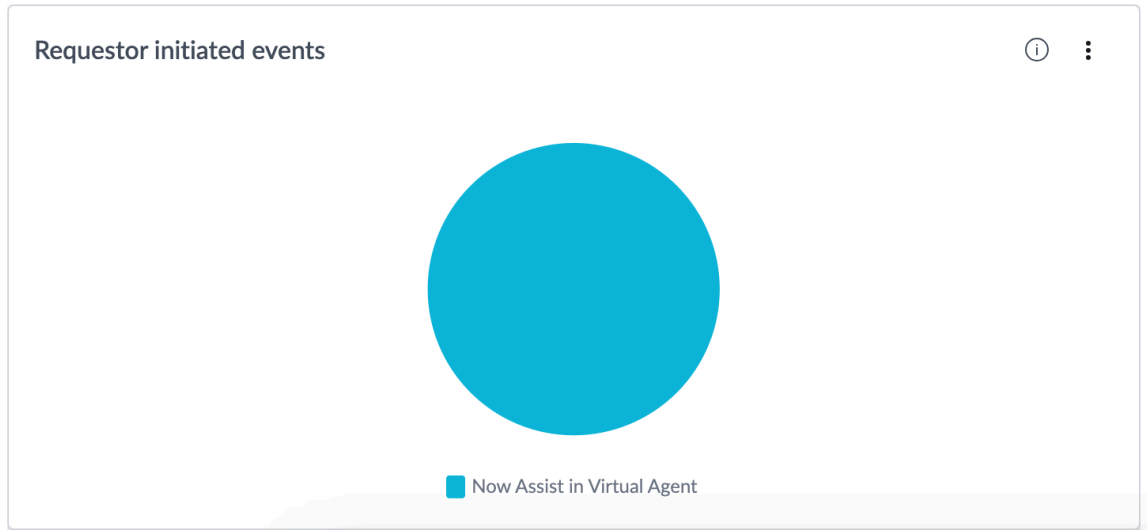


**Self-service indicators**

**Requestor initiated events**

This area of the dashboard shows a breakdown of the queries handled by the selected Now Assist self-service capability. For example, when you select Now Assist in Virtual Agent in the Feature dropdown, the indicator shows the number and percentage of queries handled by Now Assist in Virtual Agent. The indicator helps you understand which self-service capability has effectively handled user queries.


### Requestor initiated events indicator



### Successful conversations by capability

This area of the dashboard shows the Now Assist self-service capabilities by the number of successful conversations they contributed to. Successful conversations are determined based on the feedback provided by the user on the query response.

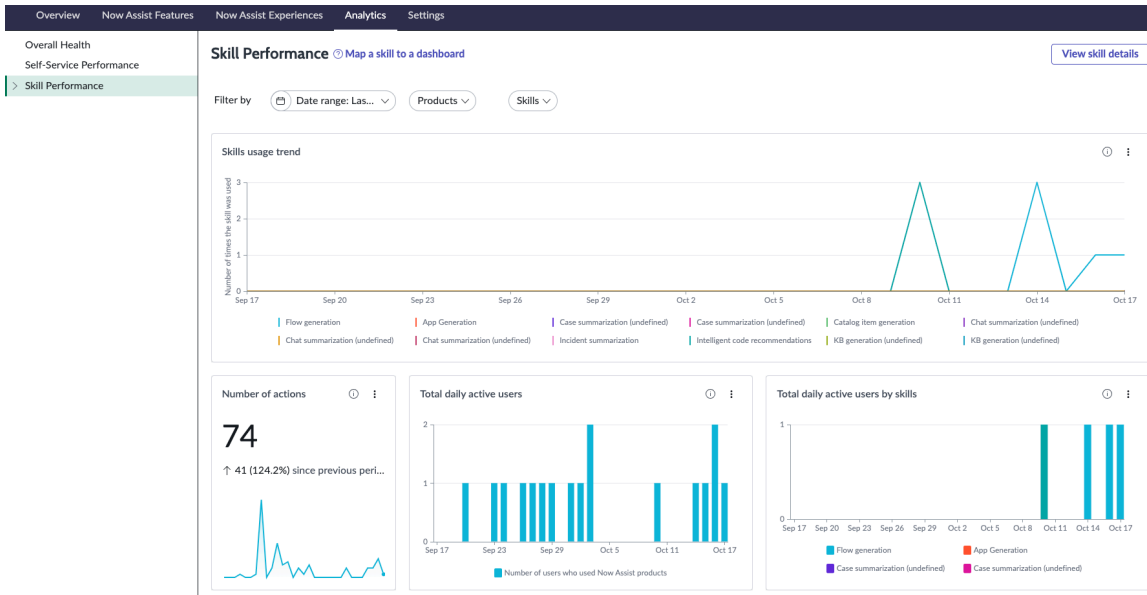
### Successful conversations by capability



### 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 page contains indicators that help you analyze usage and performance of active skills. Use the Date range, Product and Skills filters to breakdown by date range, Now Assist product, and skill respectively. The filter selection applies to all visualizations on the page.



The indicators on the Skills performance dashboard page provide the following insights.

- Skills usage trend visualization for a selected period can reveal skills that have been used more frequently or less frequently.
- 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 visualizations show 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 usage 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 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 better understand in-demand skills, and skills that are low on usage and need further analysis.

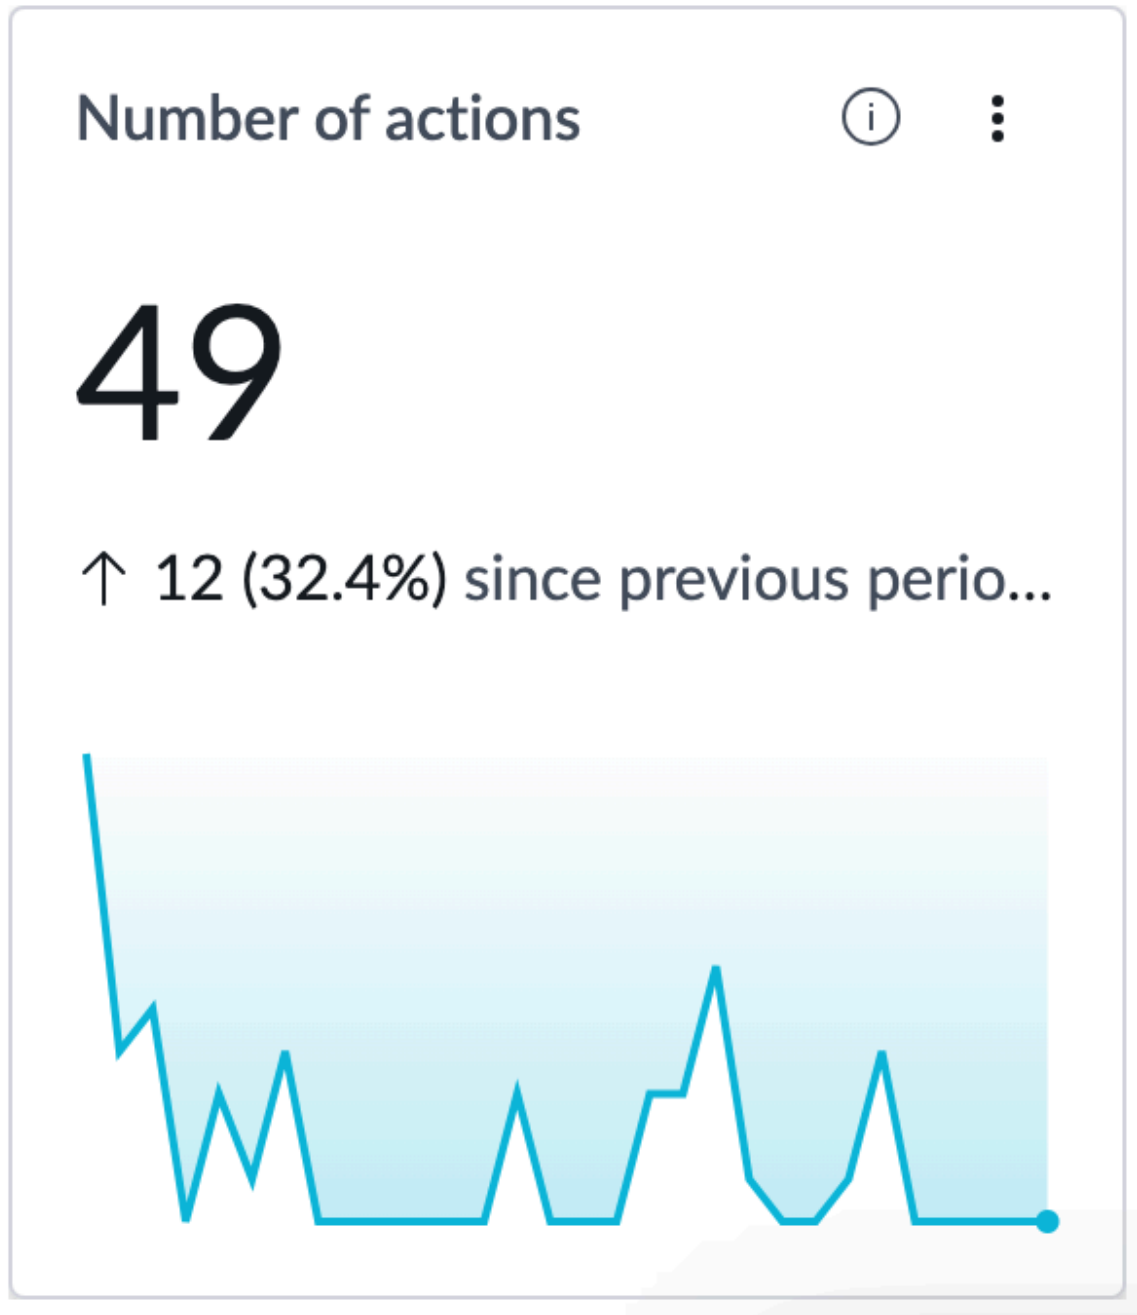
### Skills usage trend indicator



### 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 whereas the trend line comparison shows the increase or decrease in number of actions from the previous period.


**Number of actions indicator**



**Total daily active users**

This area of the dashboard shows cumulative number of daily active users who've used Now Assist products with at least one active skill. Select a filter combination to see skill usage patterns. For example, you can visualize the number of users who've used Chat Summarization skill across all products in the last month.


### Total daily active users indicator



### Total daily active users by skill

This area of the dashboard shows the number of daily active users who've 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.

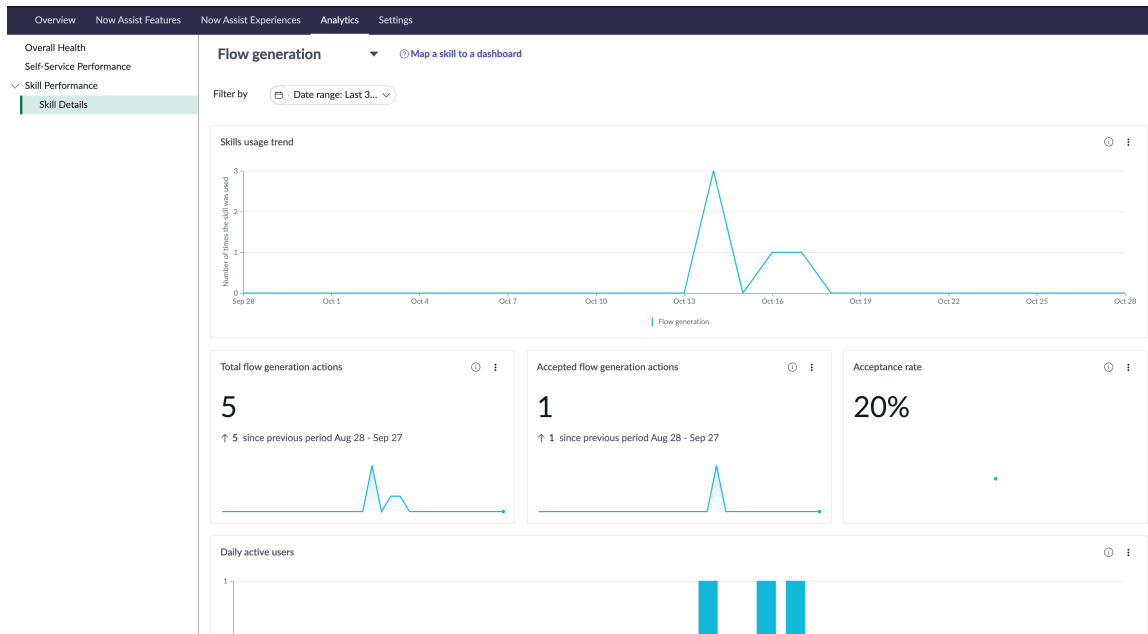
### Total daily active users by skill indicator



### Skill details

Use the Skill details page to view usage and performance indicators of a skill.

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 to view the indicators. The drop-down lists both active and inactive skills. Each skill has a subtitle that identifies the skill family 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.



The indicators on the Skill details dashboard page provide the following insights.

- Skill usage 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. This is because each skill is mapped to its own dashboard that contains a set of indicators related to the skill.


Now Assist Analytics dashboard comes with some default skill-to-dashboard mappings to get you started. The default dashboards are visible to users with 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.

#### 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.

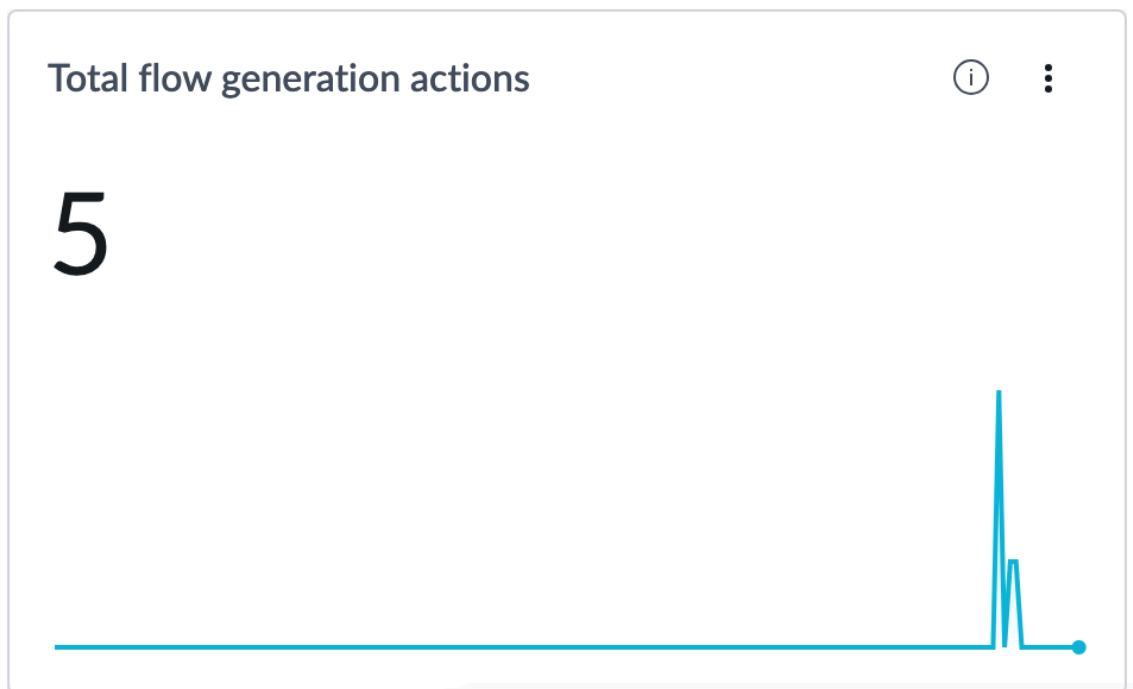
**Skill usage trend indicator**



**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 number of actions from the previous period.

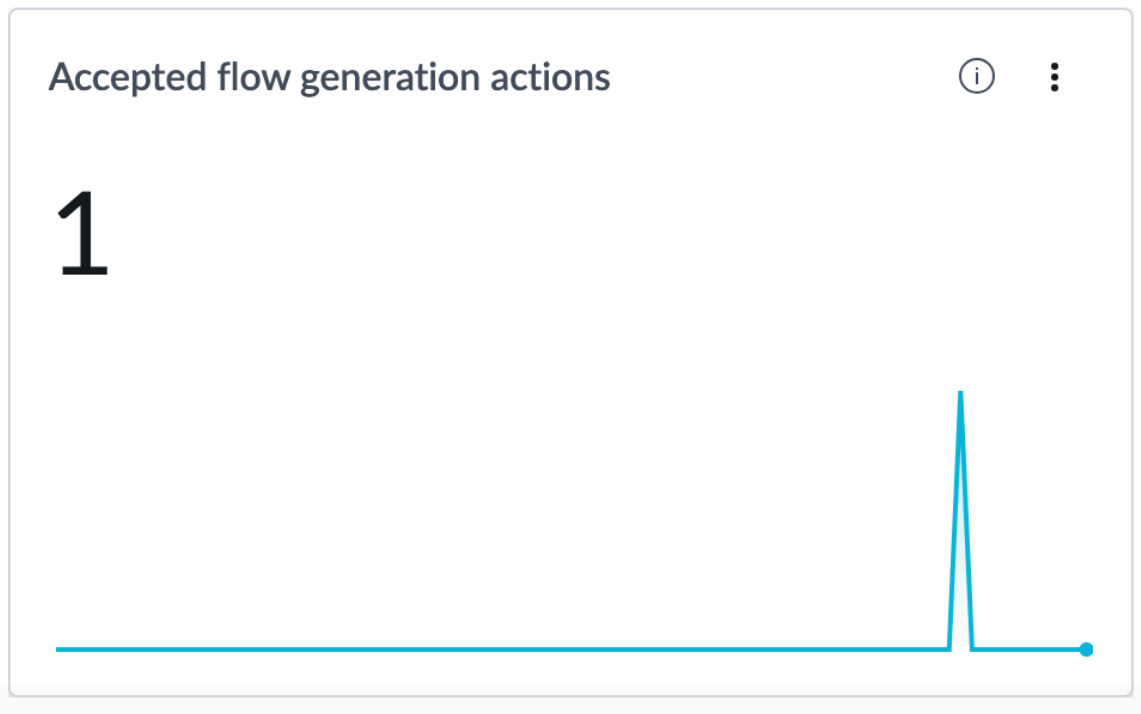
**Total flow generation actions indicator**



**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 which were accepted by the users.

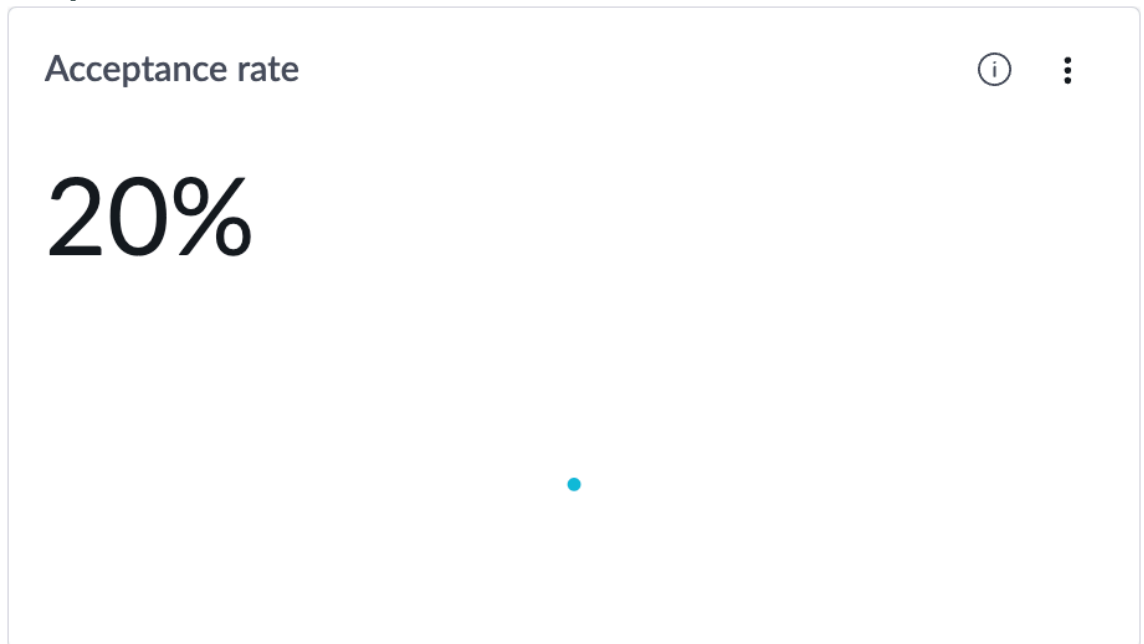
**Accepted flow generation actions indicator**



**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:  $(\text{Total number of accepted flow generation actions} / \text{Total number of flow generation actions}) \times 100$ .

**Acceptance rate indicator**



**Now Assist Analytics reference**

Now Assist Analytics reference topics include information about user roles and indicators on the dashboard.

### Now Assist Analytics roles

Now Assist Analytics requires certain roles to view and manage the dashboard functionality.

#### Now Assist Analytics Admin [sn\_na\_analytics.admin]

This user can view the Now Assist Analytics dashboard in the Now Assist Admin console, and read and write access to some data source tables.

### 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 Viewer [sn\_na\_analytics.viewer]

This user 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 this role is assigned to by default.

None.

### Special considerations

None.

### Now Assist Analytics indicators

Indicators for Now Assist Analytics include information about generative AI log metadata, deflection events, deflection event activities, deflection rate, and more.

To access these indicators, navigate to **Platform Analytics Administration > Indicators**. You must have Now Assist Analytics Admin [sn\_na\_analytics\_admin] role to access the indicators.

These indicators measure data on daily intervals. 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.

## Now Assist Analytics indicators

Now Assist Analytics indicators on your instance provide details about deflection events and activities, deflection rate, skill acceptance rate, and so on. The following table lists the indicators that are included in Now Assist Analytics, along with breakdowns and indicator sources.

### Now Assist Analytics indicators

Indicator	Indicator type	Breakdowns available	Indicator source table	Application
Daily GenAI Deflection Events	Automated	<ul style="list-style-type: none"> <li>• By Deflection Source</li> <li>• By Deflection Event Outcome</li> <li>• By Deflection Event Outcome Type</li> <li>• By Query Resolution Status</li> <li>• By Deflection Log Source</li> </ul>	Now Assist Analytics Deflection Event [sn_na_def_event]	Now Assist Analytics
Daily GenAI Deflection Event Activities	Automated	<ul style="list-style-type: none"> <li>• By Deflection Event Activity Feedback</li> <li>• By Deflection Event Activity Type</li> <li>• By Deflection Event Activity Direction</li> <li>• By Deflection Source</li> <li>• By Deflection Log Source</li> </ul>	Now Assist Analytics Deflection Event Activity [sn_na_def_event_activity]	Now Assist Analytics
Daily GenAI Log Metadata	Automated	<ul style="list-style-type: none"> <li>• By Feedback</li> <li>• By response accepted without edits</li> <li>• By Skills Config</li> <li>• By Skill Family</li> </ul>	Gen AI Log Metadata [sys_gen_ai_log_metadata]	Now Assist Analytics

**Now Assist Analytics indicators (continued)**

Indicator	Indicator type	Breakdowns available	Indicator source table	Application
GenAI Actions per day	Automated	<ul style="list-style-type: none"> <li>• By Skills Config</li> <li>• By GenAI Feature</li> <li>• By Skill Family</li> <li>• By Generative AI Skill Execution Modality</li> </ul>	Generative AI Usage log [sys_gen_ai_usage_log]	Now Assist Admin console
Number of users who used GenAI actions per day	Automated	<ul style="list-style-type: none"> <li>• By Skills Config</li> <li>• By GenAI Feature</li> <li>• By Skill Family</li> <li>• By Generative AI Skill Execution Modality</li> </ul>	Generative AI Usage log [sys_gen_ai_usage_log]	Now Assist Admin console
Acceptance rate	Formula	By Skills Config	Gen AI Log Metadata [sys_gen_ai_log_metadata]	Now Assist Analytics
Acceptance rate without edits	Formula	By Skills Config	Gen AI Log Metadata [sys_gen_ai_log_metadata]	Now Assist Analytics
Deflection Rate (Self-Solved Performance)	Formula	By Deflection Source	Now Assist Analytics Deflection Event [sn_na_def_event]	Now Assist Analytics

**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 text-to-text LLM** [↗](#)

Model used for conversational use cases like Virtual Agent topic execution and conversational catalog and agent assist use cases like alert analysis, AI search, and incident, case, and chat summarization.

#### **Model card for ServiceNow text-to-code LLM** [↗](#)

Model used for code generation.

#### **Model card for ServiceNow flow next-best-action LLM** [↗](#)

Model used for flow recommendations.

#### **Model card for ServiceNow text-to-flow LLM** [↗](#)

Model used for flow generation.

#### **Model card for ServiceNow text-to-text SLM** [↗](#)

Model used for Now Assist Guardian, text-to-cypher and other use cases that demand rapid inference and high throughput.

#### **<https://downloads.docs.servicenow.com/resource/enus/infocard/third-party-llm.pdf>** [↗](#)

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

## **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.

## **Now Assist reference**

Reference topics include information about user roles, data usage, and domain separation for Now Assist.

## 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 [sn\_ace.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.

### 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. 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. You need to enable and configure the Sensitive Data Handling plugin. To configure the plugin, see [Configure sensitive data handling for generative AI](#).

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.

**i 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.

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.

**i 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](#).

## 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](#) .

## Overview of domain separation in the Now Assist Admin console

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.

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](#) 

## Performance Analytics indicators for generative AI

Indicators for generative AI include information about generative AI logs, executions, tokens, latency, number of users who use a Now Assist skill, and more.

These indicators can be found by going to the Analytics Hub. To access the Analytics Hub, navigate to **Performance Analytics > Analytics Hub**. There are two sets of indicators, one for the

Generative AI Log table and one for the Generative AI Usage Log table. Both can be accessed by searching for `gen ai`.

These indicators measure data on daily intervals. 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.

## Generative AI Log table Indicators

Generative AI logs on your instance provide details for generative AI executions on the instance. All indicators can be broken down by definition and skill config so you can keep track of the executions of LLM requests on your instance.

### Total Gen AI Logs

Breakdowns: By Definition, By Skill Config, By Status, By Feedback

### Gen AI Log - Total Internal Now LLM executions

Breakdowns: By Definition, By Skill Config, By Status, By Feedback

### Gen AI Log - Maximum Prompt Token Count

Breakdowns: By Definition, By Skill Config

### Gen AI Log - Total Sum Prompt Token Count

Breakdowns: By Definition, By Skill Config

### Gen AI Log - Minimum latency

Breakdowns: By Definition, By Skill Config

### Gen AI Log - Maximum latency

Breakdowns: By Definition, By Skill Config

### Gen AI Log - Total Sum Response Token Count

Breakdowns: By Definition, By Skill Config

### Gen AI Log - Maximum Response Token Count

Breakdowns: By Definition, By Skill Config

## Generative AI Usage Log table Indicators

Generative AI Usage Log table indicators provide information about how and who is using generative AI features in Now Assist skills and Generative AI Controller capabilities.

### Gen AI Actions per day

Breakdowns: By Generative AI Skill Execution Modality, By Gen AI Feature

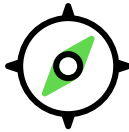



### Number of users who used Gen AI actions per day

Breakdowns: By Gen AI Feature

## 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

<p>Explore</p>  <p>Explore Knowledge Graph</p>	<p>Configure</p>  <p>Configure Knowledge Graph</p>
<p>Use Knowledge Graph</p>  <p>Create and edit Knowledge Graph schemas</p>	<p>Reference</p>  <p>Additional information to configure Knowledge Graph</p>

## Exploring Knowledge Graph

Knowledge Graph provides a customized user experience for enterprises by creating and managing Knowledge Graph schemas that are represented by nodes, node properties, and edges.

### Knowledge Graph overview

ServiceNow Knowledge Graph application provides two main functionalities:

1. Knowledge Graph designer: A user-friendly interface to manage Knowledge Graph schemas and their nodes, node properties, and edges.
2. Prebuilt integrations: Integrations with Now Assist Virtual Agent and AI agents.

### 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, including configuring nodes (tables), properties (columns), and 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.

## 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 and AI agents for User Context: Helps users with personalized responses.
2. Integration with Now Assist Virtual Agent 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 for Employee schema: Helps requesters 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 see [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 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.	Requesters
Simple and easy to use	Creates a complex data model called Knowledge Graph schema with numerous entities and their relation within a few steps.	kg_admin
Easy to manage	Editing Knowledge Graph schemas to add new nodes or edges is simple.	kg_admin

**Knowledge Graph benefits (continued)**

Benefit	Feature	Users
Customizable Knowledge Graph schemas	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

**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 Schemas**

Use Knowledge Graph Designer to create, edit, and manage Knowledge Graph schemas.

**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:

- **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.

The details of the user context passed by default are as below:

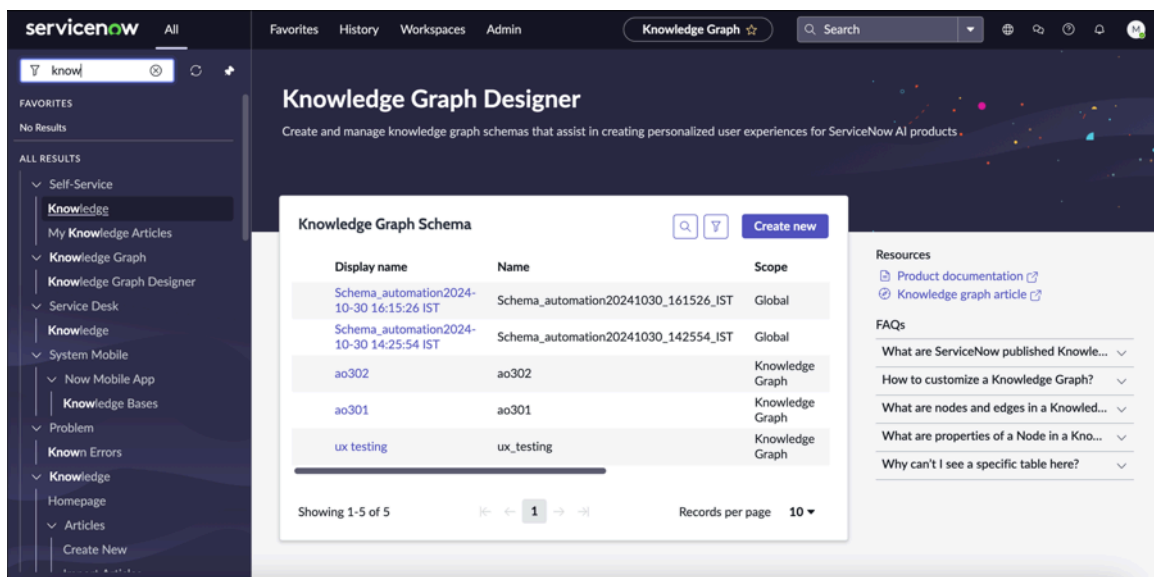
**User Graph Schema**

Entry	Attributes
Grouped as Profile	sys_user table's columns: <ul style="list-style-type: none"> <li>○ name</li> <li>○ first name</li> <li>○ last name</li> <li>○ user name</li> <li>○ employee number</li> </ul>

Entry	Attributes
	<ul style="list-style-type: none"> <li>○ email</li> <li>○ business phone</li> <li>○ mobile phone</li> <li>○ title</li> <li>○ preferred language</li> <li>○ Time format</li> <li>○ date format</li> <li>○ Time zone</li> </ul> <p>cmn_location</p> <ul style="list-style-type: none"> <li>○ city</li> <li>○ state</li> <li>○ country</li> </ul> <p>cmn_department</p> <ul style="list-style-type: none"> <li>○ name</li> <li>○ headcount</li> </ul> <p>core_company: name</p>
Grouped as manager	<p>sys_user table's columns:</p> <ul style="list-style-type: none"> <li>○ name</li> <li>○ first name</li> <li>○ last name</li> <li>○ user name</li> <li>○ employee number</li> <li>○ email</li> <li>○ business phone</li> <li>○ mobile phone</li> <li>○ title</li> <li>○ preferred language</li> <li>○ Time format</li> <li>○ date format</li> <li>○ Time zone</li> <li>○ zip code</li> <li>○ city</li> <li>○ state</li> </ul>
Grouped as reportees	<p>sys_user table's columns:</p>

Entry	Attributes
	<ul style="list-style-type: none"> <li>○ name</li> <li>○ first name</li> <li>○ last name</li> <li>○ user name</li> <li>○ employee number</li> <li>○ email</li> <li>○ business phone</li> <li>○ mobile phone</li> <li>○ title</li> <li>○ preferred language</li> <li>○ Time format</li> <li>○ date format</li> <li>○ Time zone</li> <li>○ zip code</li> <li>○ city</li> <li>○ state</li> </ul>
Grouped as assets	<p>alm_asset table's columns:</p> <ul style="list-style-type: none"> <li>○ display name</li> <li>○ purchase date</li> <li>○ retired date</li> </ul>

- NLQ graph Scheme: Helps requesters 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 see [KB article](#).



You can also create a Knowledge Graph schema from the landing page.

### Leverage Knowledge Graph prebuilt integration with Virtual Agent

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 Virtual Agent are:

1. Integration with Now Assist Virtual Agent for User Context: Helps requesters with personalized responses.
2. Integration with Now Assist Virtual Agent for Slot filling: Helps requesters pre-fill the slots for LLM topics and skills execution using Natural Language Querying of Knowledge Graph.
3. Integration with Now Assist Virtual Agent for NLQ graph: Helps requesters with personalized responses on people queries and Natural Language queries. Also supports people citation card.

**i 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.

### Integration with Now Assist Virtual Agent 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.
- 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 Virtual Agent for Slot filling**

Knowledge Graph enhances the Virtual Agent 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.

**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 Virtual Agent for people queries

Now Assist Virtual Agent 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:

### Example of synthesized response for people on the portal's search results page



## Example of synthesized response with people inline citations in chat

The screenshot shows a chat interface with a header bar containing the 'Now Assist' logo and navigation icons for New chat, Chats, Support, and Settings. A user named Carol Coughlin is in the chat. The AI response includes two paragraphs of text, each with an inline citation for a person named Carol Coughlin. The first citation is for an Account Executive on the Sales team in Hamburg, with email carol.coughlin@example.com. The second citation is for a Product Manager on the AI team in New York City, with email carol.coughlin2@example.com. The response concludes with a request for more information and a feedback bar with thumbs up, thumbs down, and search icons. A disclaimer at the bottom states: 'Some answers generated by AI. Be sure to check for accuracy.'

**Now Assist**

New chat Chats Support Settings

Carol Coughlin

I found 2 people named Carol Coughlin.

Carol Coughlin is an Account Executive on the Sales team. Their location is Hamburg. Their email is [carol.coughlin@example.com](mailto:carol.coughlin@example.com) & Carol Coughlin.

Carol Coughlin is an Product Manager on the AI team. Their location is New York City. Their email is [carol.coughlin2@example.com](mailto:carol.coughlin2@example.com) & Carol Coughlin.

If these aren't who you're looking for, tell me more information, like their job title and location, and I can help find the right person.

Reply...

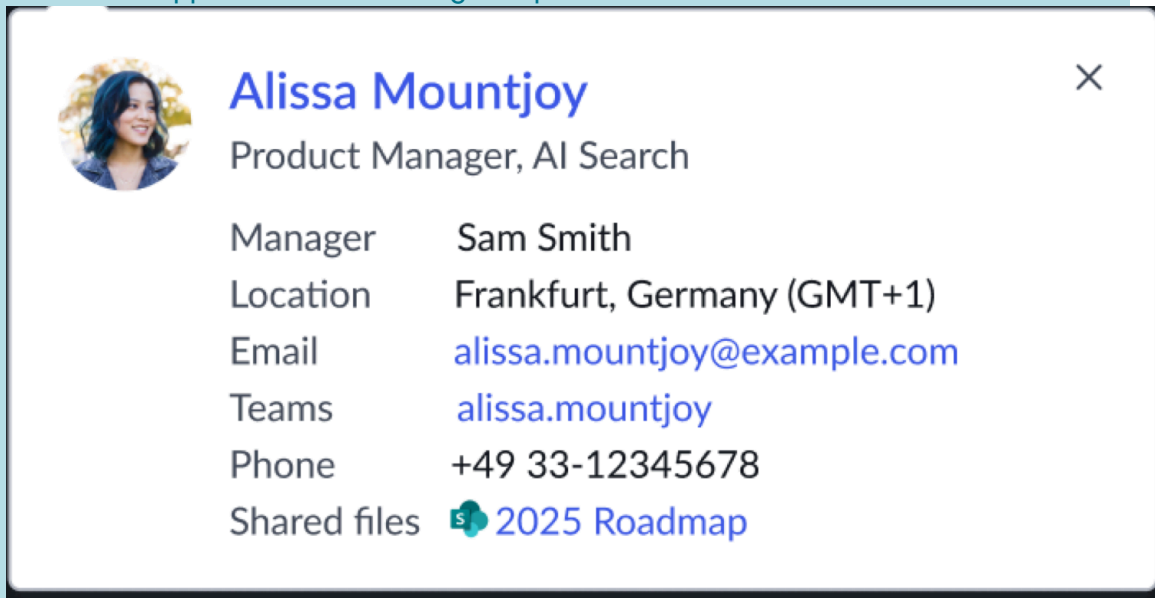
Some answers generated by AI. Be sure to check for accuracy.

Selecting the person's name presents a popover. The information in the popover can include the following information:

- Manager
- Location
- Email
- Teams
- Phone
- Shared files

**i 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](#).



**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.

- 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": [
    {
      "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"`.

## 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 for BUs.

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 OOTB integration with Now Assist Virtual Agent for Slot filling, configure Knowledge Graph in Virtual Agent. Refer [Configuring Now Assist in Virtual Agent](#)  for admin guided setup.

### 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.

- 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.servicenow.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.

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.

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. HRSD schema (available only if you have the Now Assist for HRSD plugin)

### User profile schema

The User profile schema is the default schema that is used for passing the in-session user profile context to the OOTB integration with Virtual Agent for Context. This integration helps with personalized responses.

This schema details the attributes related to a user and comprises the following details:

#### User Profile Schema

Schema	Attributes
sys_user	<ul style="list-style-type: none"> <li>• name</li> <li>• first name</li> <li>• last name</li> <li>• user name</li> <li>• employee number</li> <li>• email</li> <li>• business phone</li> <li>• mobile phone</li> <li>• title</li> <li>• preferred language</li> <li>• time format</li> <li>• date format</li> <li>• timezone</li> <li>• zip code</li> <li>• city</li> <li>• state</li> </ul>

**User Profile Schema (continued)**

Schema	Attributes
cmn_location	<ul style="list-style-type: none"> <li>• city</li> <li>• state</li> <li>• country</li> </ul>
cmn_department	<ul style="list-style-type: none"> <li>• name</li> <li>• headcount</li> </ul>
core_company	name
manager	<ul style="list-style-type: none"> <li>• name</li> <li>• first name</li> <li>• last name</li> <li>• user name</li> <li>• employee number</li> <li>• email</li> <li>• business phone</li> <li>• mobile phone</li> <li>• title</li> <li>• preferred language</li> <li>• time format</li> <li>• date format</li> <li>• timezone</li> <li>• zip code</li> <li>• city</li> <li>• state</li> </ul>
reportees	<ul style="list-style-type: none"> <li>• name</li> <li>• first name</li> <li>• last name</li> <li>• user name</li> <li>• employee number</li> <li>• email</li> <li>• business phone</li> <li>• mobile phone</li> <li>• title</li> <li>• preferred language</li> </ul>

**User Profile Schema (continued)**

Schema	Attributes
	<ul style="list-style-type: none"> <li>• time format</li> <li>• date format</li> <li>• timezone</li> <li>• zip code</li> <li>• city</li> <li>• state</li> </ul>
assets	<ul style="list-style-type: none"> <li>• display name</li> <li>• purchase date</li> <li>• retired date</li> </ul>

**Create a Knowledge Graph schema**

Use the Knowledge Graph to create Knowledge Graph schemas.

**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**.
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.
Mark as Default	Check box used to assign the Knowledge Graph schema as a fallback schema. The default schema is used if the system fails to reach the assigned schema.

4. Use the schema as a fallback option by selecting **Mark as Default**.

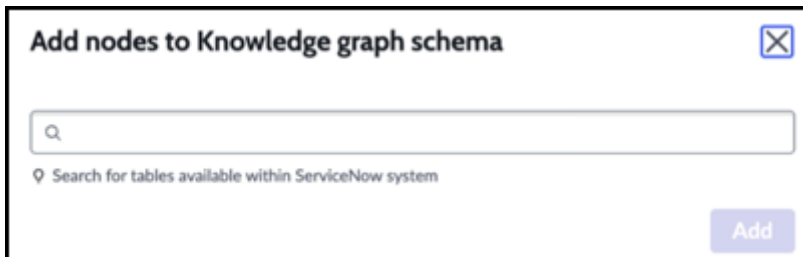
**Note:**

- If the specified Knowledge Graph schema isn't available, the system uses the fallback schema.
- If you select this option, the system overrides the existing selection.

**5. Select Create.**

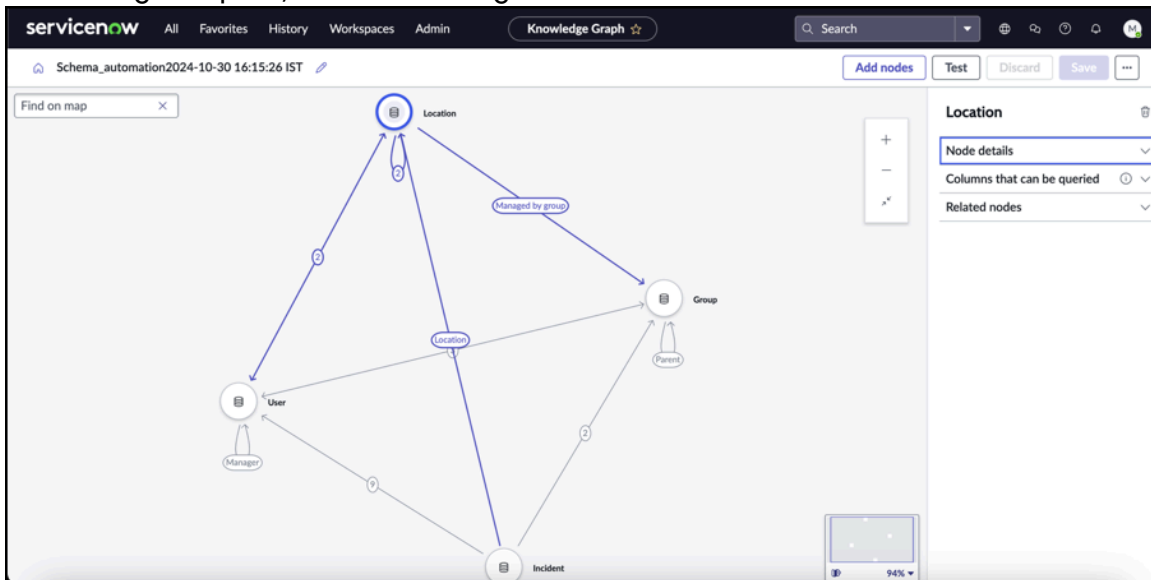
The Add nodes to the Knowledge graph schema window is displayed.

**6. Enter or search for the nodes that you want to add to the Knowledge Graph schema and select Add.**



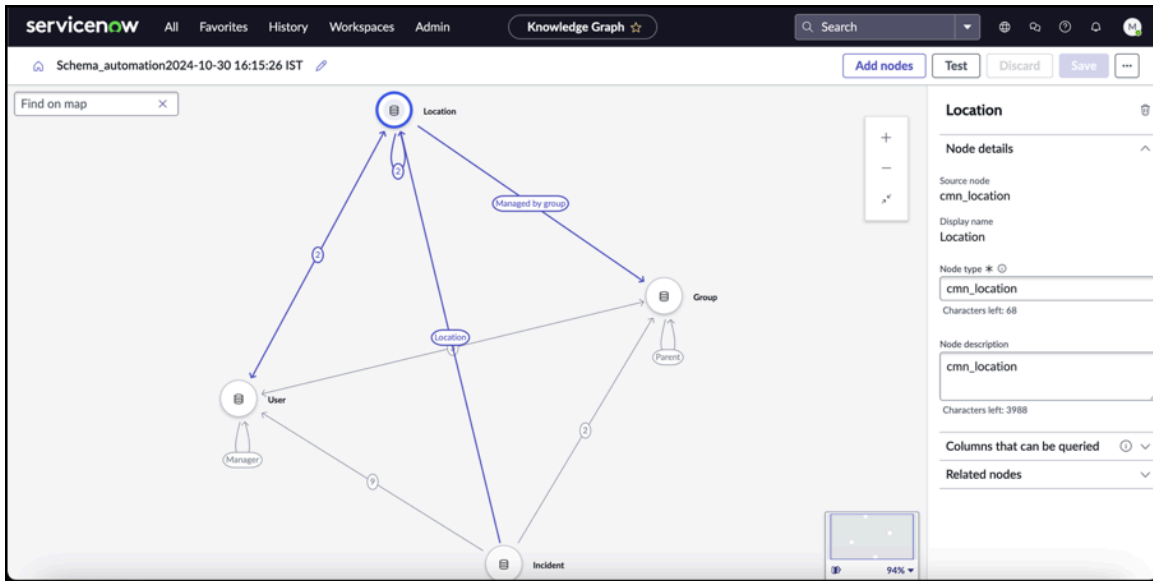
The Knowledge Graph schema is created and displayed on the Knowledge Graph canvas.

**7. In the navigation pane, add the following details:**

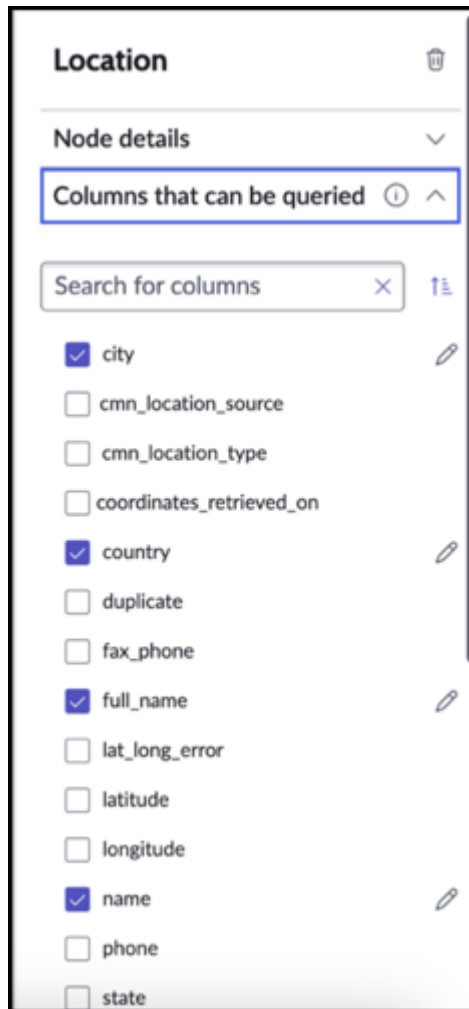


**8. In the Node details section, you can edit the following fields.**

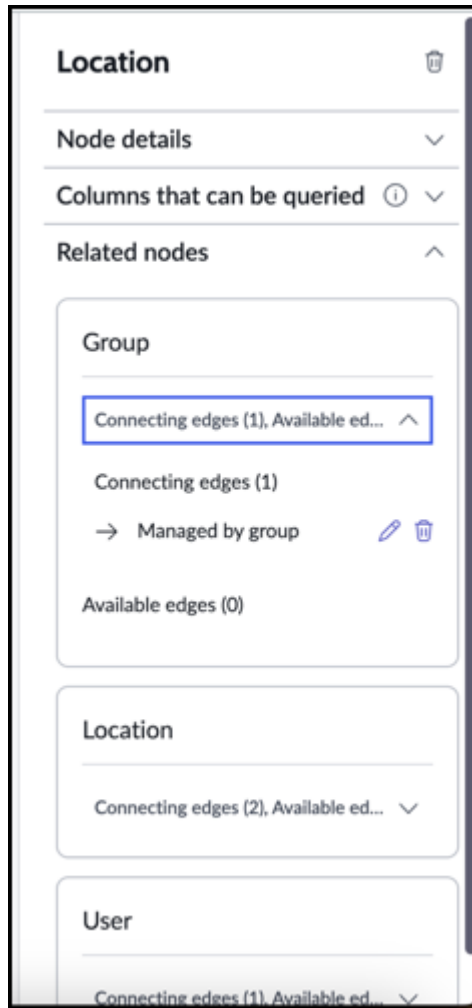
- Node type
- Node Description



9. In the Columns that can be queried section, search for and select the desired columns and select **Save**.



10. Add, delete, or edit edges in the Related nodes section and select **Save**.




### Edit a Knowledge Graph schema

Edit an existing Knowledge Graph schema to customize or update the schema details.

#### Before you begin

Role required: kg\_admin

#### Procedure

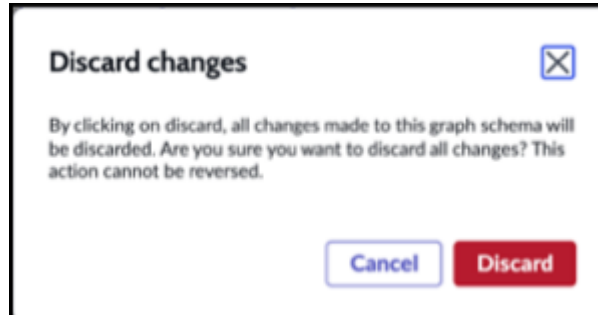
1. Navigate to **All > Knowledge Graph > Knowledge Graph Designer**.  
The UI displays a list of all the Knowledge Graph schemas on the landing page.
2. From the list of existing Knowledge Graph schemas, select a Knowledge Graph schema to edit.  
The Knowledge Graph schema opens in the Knowledge Graph canvas.
3. Select the edit icon (  ).
4. On the form, fill in the fields.

#### Edit Knowledge Graph Schema details form

Field	Description
Display Name	Display name of the Knowledge Graph schema.
Name	Name of the Knowledge Graph schema.

Field	Description
Scope	Scope used to create the Knowledge Graph schema.
Description	Knowledge Graph schema overview for users.
Mark as Default	Check box to select the Knowledge Graph schema as a fallback schema. The default schema is used if the system fails to reach the assigned schema.

- To save the changes, select **Save**.
- Delete all the changes made to the Knowledge Graph schema by selecting **Discard**.



### Manage nodes in a Knowledge Graph schema

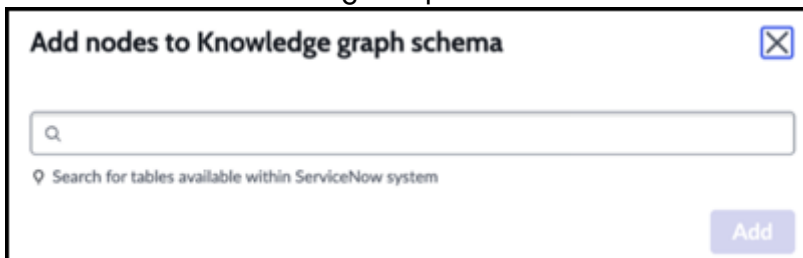
Add or delete a node from an existing Knowledge Graph schema.

#### Before you begin

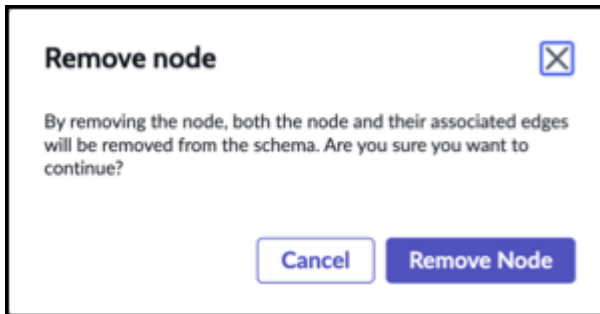
Role required: kg\_admin

#### Procedure

- Navigate to **All > Knowledge Graph > Knowledge Graph Designer**.  
The UI displays a list of all the Knowledge Graph schemas on the landing page.
- Select a Knowledge Graph schema from the list.  
The Knowledge graph schema opens in the Knowledge Graph canvas page.
- From the toolbar, select the **Add nodes** option.  
The Edit Knowledge Graph Schema details form is displayed.
- 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.



- Select **Add**.  
The nodes are added to the Knowledge Graph schema and displayed on Knowledge Graph canvas.
- Select the node to update the node details in the side navigation pane.
- 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**

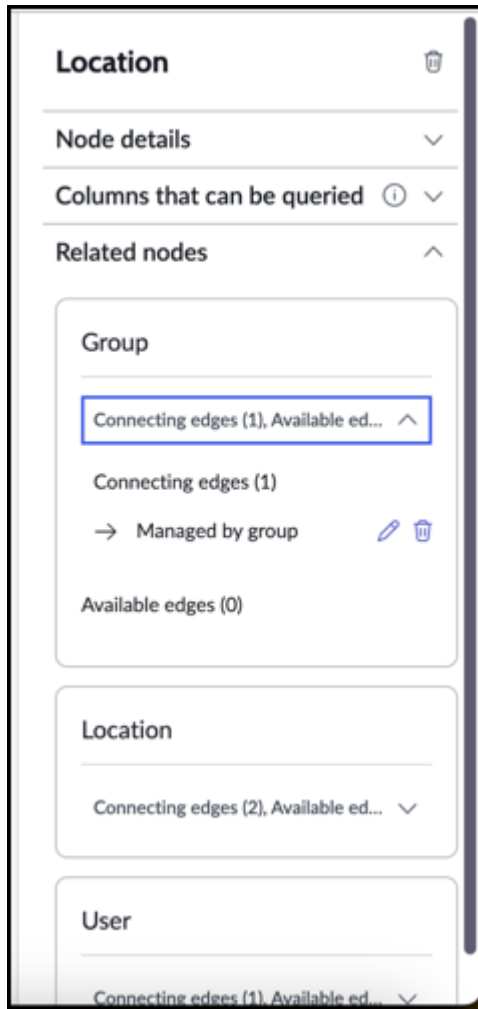
Add, edit, or delete edges that connect the nodes in a Knowledge Graph schema

**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 node to edit.**
- 4. In the side navigation pane, select **Related nodes**.**  
Each related node displays a list of connecting edges and available edges.



5. To establish a connection to an available edge, select the plus icon (+) against the edge. The available edge is moved to **Connecting edges**.
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

### 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 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 and select **Run Query**.

Test Knowledge Graph Schema

Query

Get me the details of Abel Tuter

Run Query

```
"0": [
  {
    "u_sys_id": "62826bf03710200044e0bfc8bcbe5df1",
    "u_first_name": "Abel",
    "u_name": "Abel Tuter",
    "u_phone": null,
    "u_active": "1",
    "u_time_format": null,
    "u_last_name": "Tuter",
    "u_date_format": null,
    "u_time_zone": null,
    "u_preferred_language": "en",
    "u_title": null,
    "u_sys_created_on": "2012-02-18 03:04:52",
    "u_user_name": "abel.tuter",
    "u_email": "abel.tuter@example.com",
    "u_mobile_phone": null
  }
]
```

Response time: 00:00:05

The query response and response time are displayed in the results section.

### 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:
  - 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.
Mark as default	Check box to select the cloned Knowledge Graph schema as a fallback schema. The default schema is used if the system fails to reach the assigned schema.

### Create a copy of this Knowledge Graph Schema ✕

Display Name \*

Name \* ⓘ

Scope \*  
 Global

Description

Characters left: 4000

Mark as default

To be used as fallback schema when a specific graph schema is not available. The selection overrides existing default.

6. Assign the new schema as a fallback schema by selecting the **Mark as Default** option.

**Note:**

- If the specified Knowledge Graph schema isn't available, the system uses the fallback schema.
- If you select this option, the system overrides the existing selection.

7. Select **Create**.

## 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.

**Terminologies (continued)**





Term	Description
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.

**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. You can also create your own prompts from scratch.

<p style="text-align: center;">Explore</p>  <p style="text-align: center;">Explore Generative AI concepts and terminology.</p>	<p style="text-align: center;">Configure</p>  <p style="text-align: center;">Configure Generative AI Controller provider capabilities.</p>
<p style="text-align: center;">Use</p>  <p style="text-align: center;">Use Generative AI Controller provider capabilities.</p>	<p style="text-align: center;">Reference</p>  <p style="text-align: center;">Reference for Generative AI Controller.</p>

## 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. Further, 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.

### 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, and IBM watsonx. 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 <a href="#">Generate Content</a> and <a href="#">Summarize</a>
Analyze user sentiment to identify and alleviate end-user concerns.	<a href="#">Sentiment Analysis Workflow Studio Action</a>
Query a large language model (LLM) directly.	<a href="#">Generic Prompt Workflow Studio Action</a>
Write scripts for AI model capabilities that are designed to increase the accuracy and scalability of your custom content.	<a href="#">Script with generative AI</a>

Benefit	Feature
Integrate with third-party AI service providers to customize your AI experience	OpenAI, Azure OpenAI, Google AI, Aleph Alpha, IBM watsonx

## 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](#).
- [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.
- Add and configure [sensitive data handling](#) and [semantic filtering](#) to tailor security needs to your use cases.
- [Use Generative AI Controller](#) with Virtual Agent Designer, Flow Designer, and scripts.

## 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](#).

## Overview of domain separation and Generative AI Controller

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](#).

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.

### Configure sensitive data handling for generative AI

Set up and configure how personally identifiable information and other sensitive data are removed from generative AI prompts.

### Enable recursive summarization for large inputs

Enable recursive summarization to retain context for large inputs in LLM calls.

### 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

### Install product plugins

Choose product workflows to empower your users with Now Assist and start unlocking Features and Skills.



#### Technology

Use Now Assist to elevate day-to-day operations and increase employee agility.

[Browse plugins](#)



#### Customer

Use Now Assist to improve customer experiences and deepen brand loyalty.

[Browse plugins](#)



#### Employee

Use Now Assist to boost employee engagement and strengthen morale.

[Browse plugins](#)

The Sensitive Data Handler plugin is required to successfully use Generative AI Controller. It is installed with the latest version of the application, but if you have an earlier version of Generative AI Controller on your instance you may need to add the Sensitive Data handler manually.

### 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 Aleph Alpha

Configure your API credentials to use Aleph Alpha in custom workflows and Virtual Agent Designer topics.

#### Before you begin

You must have an Aleph Alpha account to set up your API key.

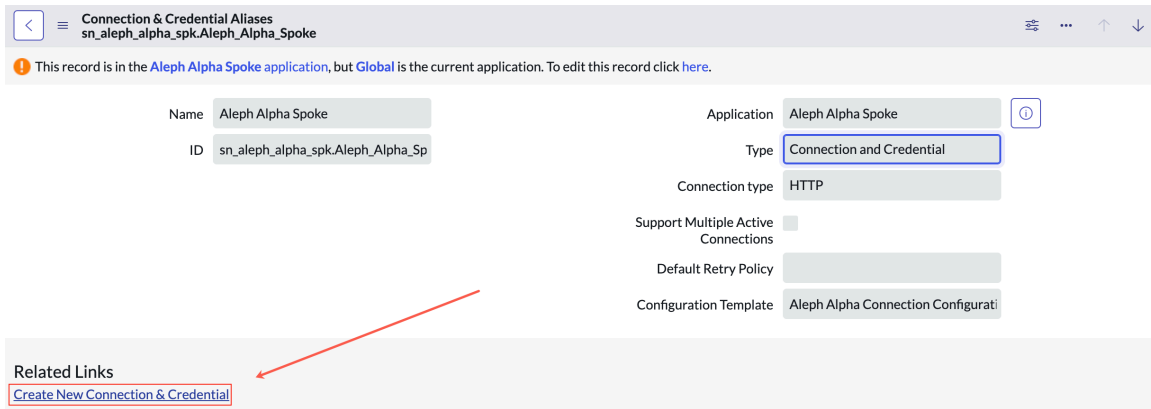
Role required: admin

#### About this task

In order to use Aleph Alpha 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 Aleph Alpha Spoke.



3. Select the **Create New Connection & Credential** related link.

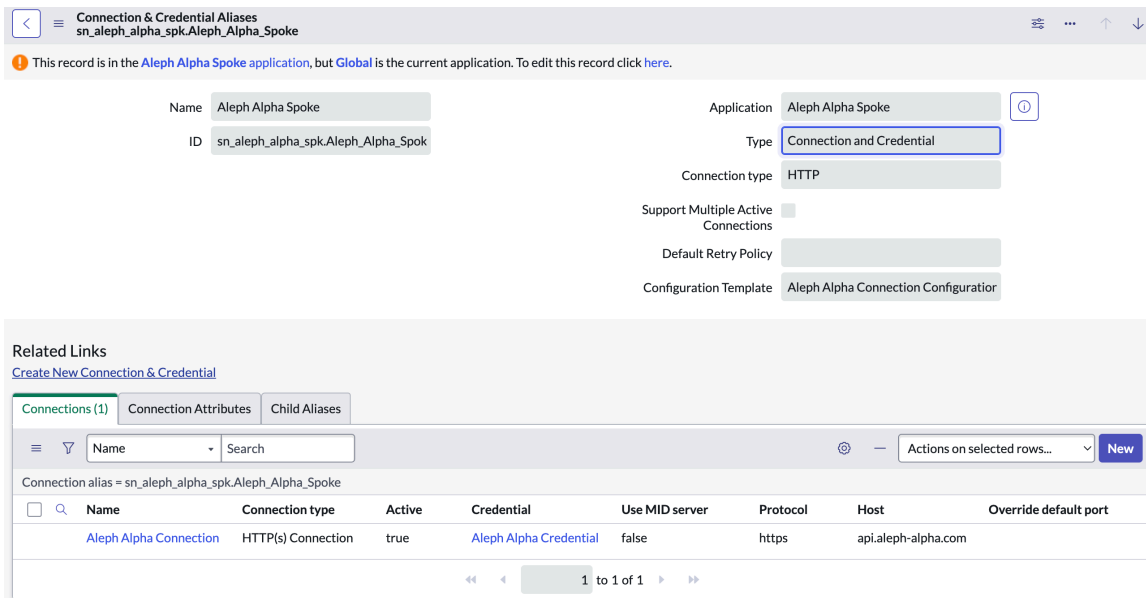
4. Enter your API key.

You can create an API key by logging into Aleph Alpha and selecting **Get Token** on the User Profile page.

5. Select **Create**.

### Result

You can now use Aleph Alpha Completion OneExtend Capabilities in Flow Designer, Virtual Agent Designer, and scripts to create custom experiences with generative AI.



### 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.

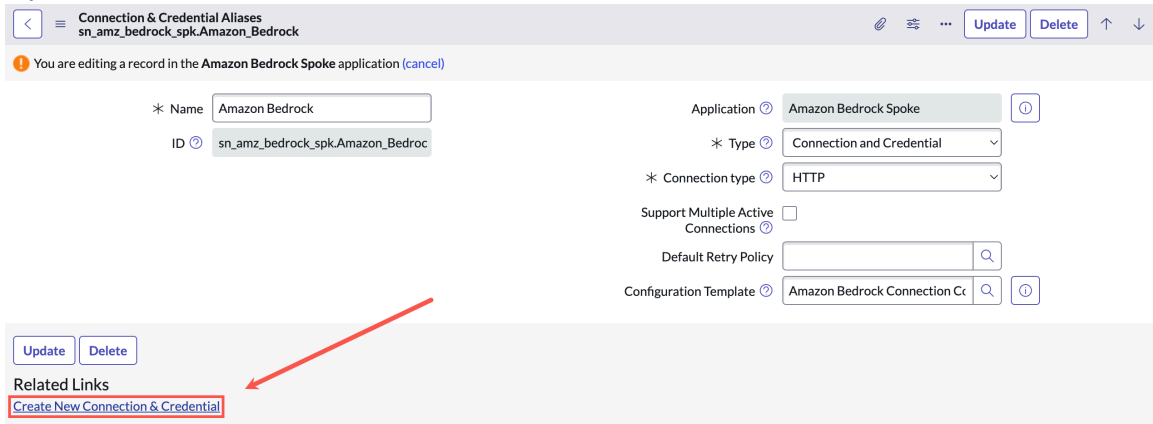
Role required: admin

#### About this task

To use Amazon Bedrock as your LLM provider for Generative AI Controller capabilities, you must have an active connection configured.

## Procedure

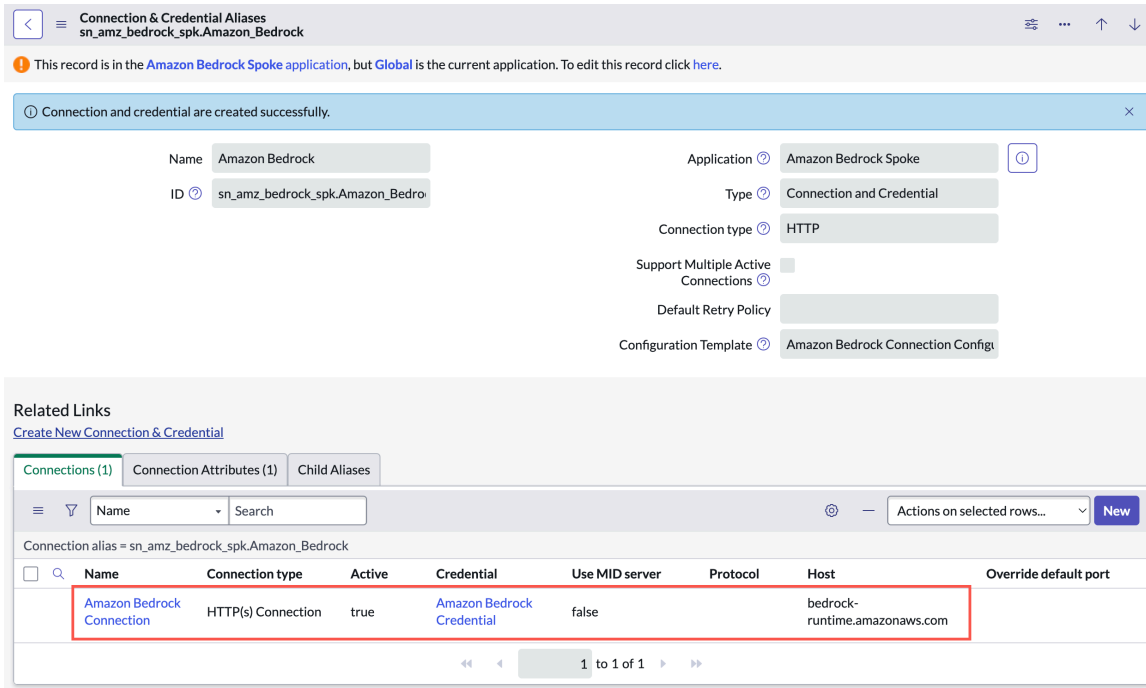
1. Navigate to **All > Connections & Credentials > Connections & Credential Aliases**.
2. Open the record for Amazon Bedrock.



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.



## 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.

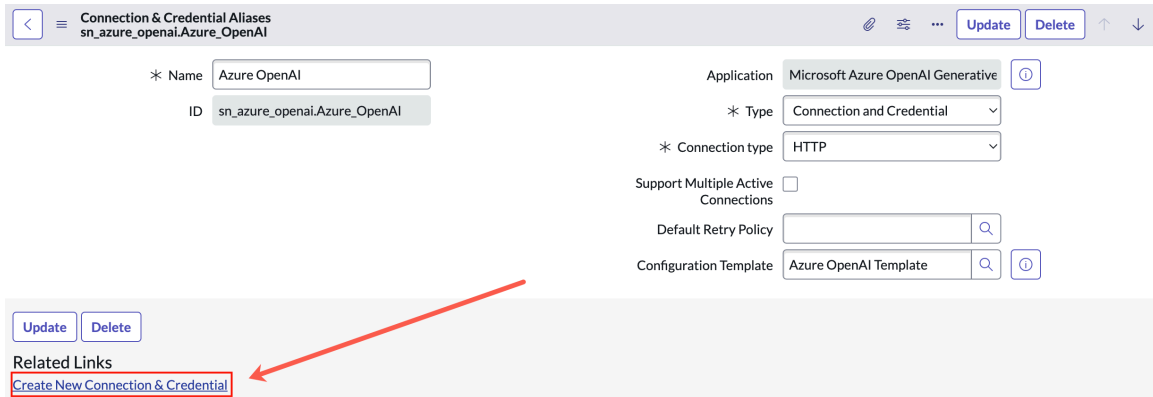
Role required: admin

### 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.

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.

The screenshot shows the configuration page for a connection alias named 'Azure OpenAI'. The configuration includes:

- Name: Azure OpenAI
- ID: sn\_azure\_openai.Azure\_OpenAI
- Application: Microsoft Azure OpenAI Generative
- Type: Connection and Credential
- Connection type: HTTP
- Support Multiple Active Connections:
- Default Retry Policy: (empty)
- Configuration Template: Azure.OpenAI Template

Below the configuration is a table of related links and a table of connections. The 'Connections' table has one entry:

Name	Connection type	Active	Credential	Use MID server	Protocol	Host	Override default port
Azure OpenAI Connection	HTTP(s) Connection	true	Azure OpenAI Credential	false		openai-platform.openai.azure.com	

### 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 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.

### 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.

The screenshot shows the configuration page for a connection alias named 'Google Bard MakerSuite'. The configuration includes:

- Name: Google Bard MakerSuite
- ID: sn\_google\_bard\_spk.Google\_Bard\_Mi
- Application: Google Bard Spoke
- Type: Connection and Credential
- Connection type: HTTP
- Support Multiple Active Connections:
- Default Retry Policy: Google Bard Retry Policy
- Configuration Template: Google Bard Connection Confi

Below the configuration is a 'Related Links' section with a link 'Create New Connection & Credential' highlighted by a red box and a red arrow pointing to it.

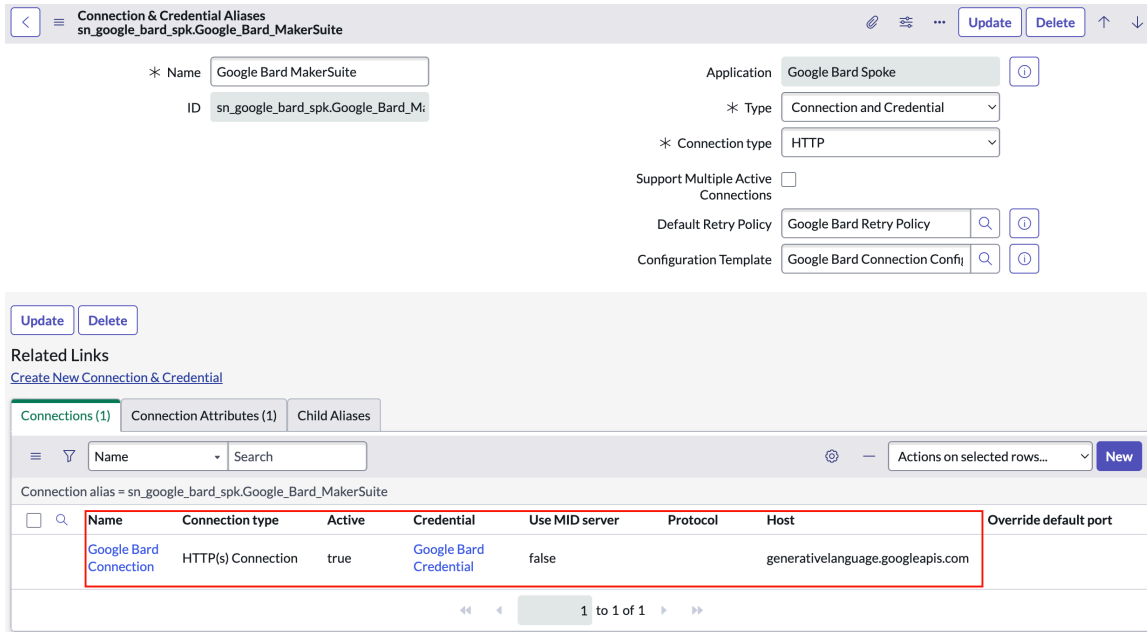
4. Enter your API key.

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.

**Before you begin**

You must have a Google Cloud project and the permissions to generate new OAuth credentials.

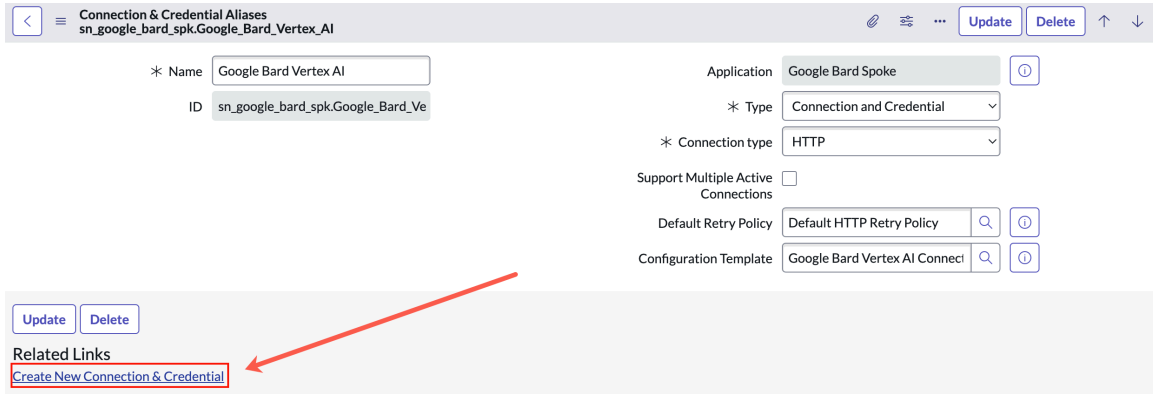
Role required: admin

**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: <b>a. Application type: Web application</b> <b>b. Authorized redirect URI:</b> URL in the OAuth Redirect URL field, usually <code>&lt;instance&gt;.service-now.com/oauth_redirect.do</code> For more information, see the <a href="#">Google documentation for creating OAuth client IDs</a> . 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**.

**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.

Update Delete

Related Links  
[Create New Connection & Credential](#)

Connections (1) Connection Attributes (3) Child Aliases

Name Search Actions on selected rows... New

Connection alias = sn\_google\_bard\_spk.Google\_Bard\_Vertex\_AI

Name	Connection type	Active	Credential	Use MID server	Protocol	Host	Override default port
Google Bard OAuth Connection	HTTP(s) Connection	true	Bard Credentials.1	false		us-central1-alphaplatform.googleapis.com	

1 to 1 of 1

### 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.

Update Delete

Related Links  
[Create New Connection & Credential](#)

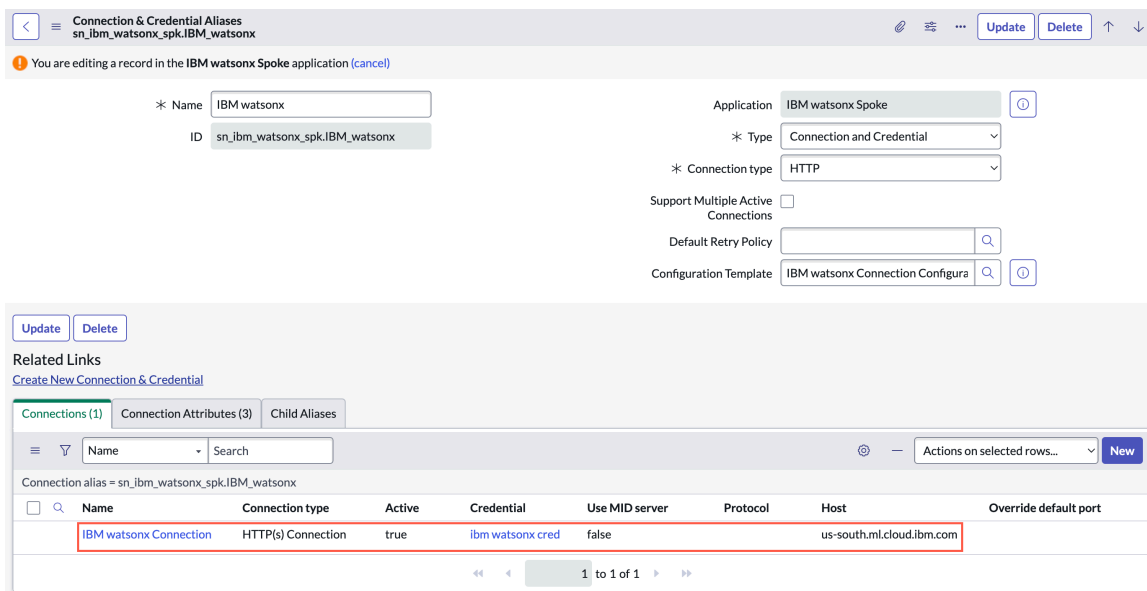
4. In the API key field, enter your API key.

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.

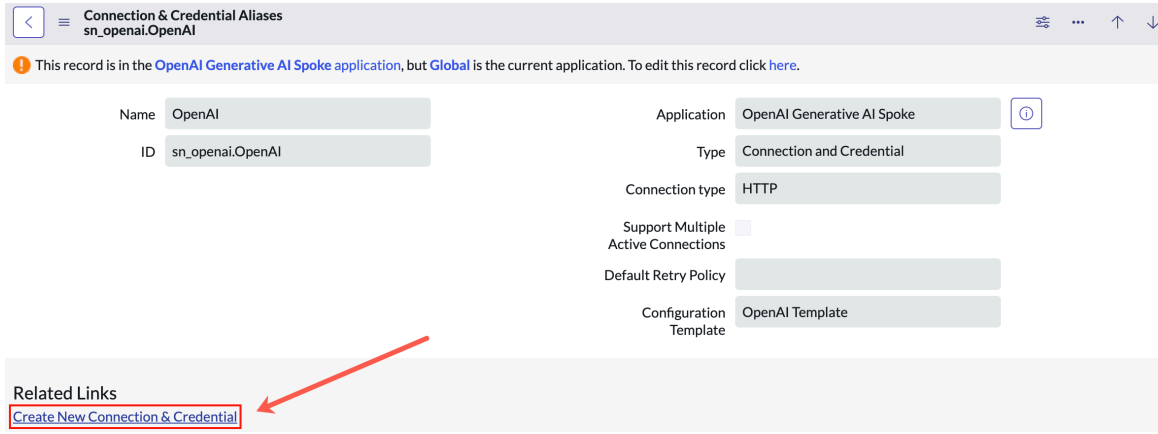
Role required: admin

### 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.



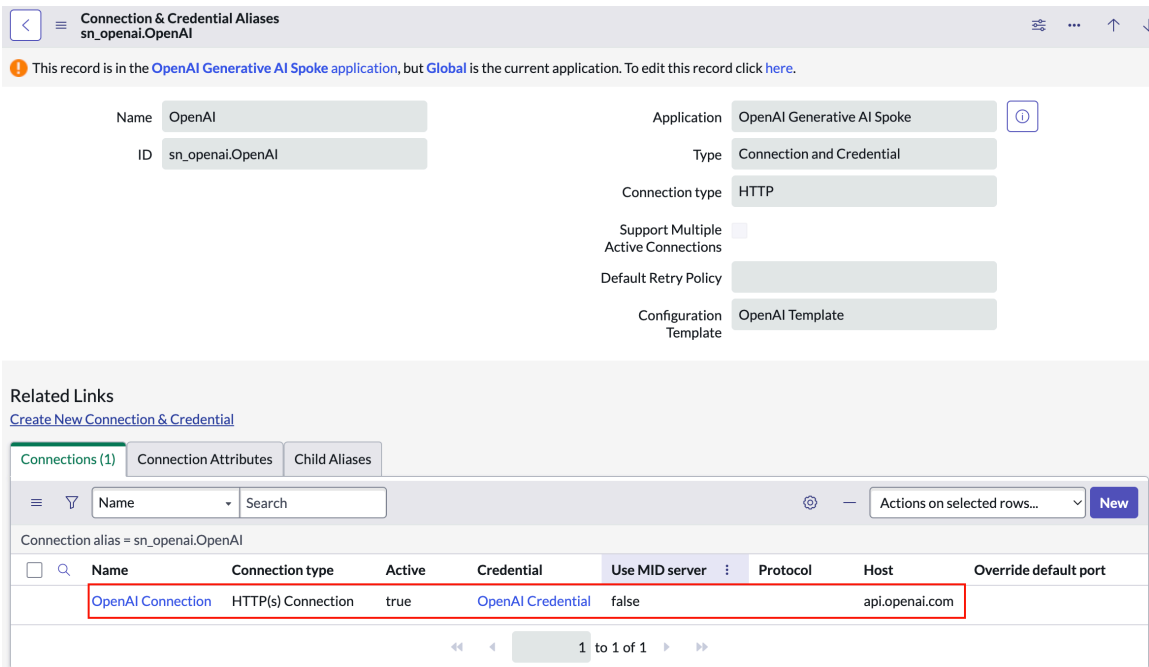
4. In the API key field, enter the API key.

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.



### 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.
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.

Update Delete

Related Links  
[Create New Connection & Credential](#)

Connections (1) | Connection Attributes | Child Aliases

Search: [Name] [Search] [Actions on selected rows...] [New]

Connection alias = My\_Model\_Alias

Name	Connection type	Active	Credential	Use MID server	Protocol	Host	Override default port
My Model Connection	HTTP(s) Connection	true	My Model Credentials	false	https	api-inference.huggingface.co	

1 to 1 of 1

**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.

**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 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.

OneExtend Capability Attributes (6)

Name	Active	Apply Filter	Attribute Prompt Template	Contains Large Input	Data Type	Definition	Description	Hidden Attribute
error	true	false		false	String	Summarize (OpenAI Completion)	Error Message	false
errorCode	true	false		false	String	(empty)	Error Code	false
provider	true	false		false	String	Summarize (OpenAI Completion)	Provider name	false
response	true	false		false	String	Summarize (OpenAI Completion)	Summary	false
status	true	false		false	String	Summarize (OpenAI Completion)	Status	false
textToSummarize	true	false		false	String	Summarize (OpenAI Completion)	Text to Summarize	false

Attributes you can use in your prompts

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:

```
{"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
  }
}

```

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](#).

### 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.

### 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).

OneExtend Capability Sentiment Analysis

Name:  Application: Generative AI Controller

Active:  Domain: global

Description: This capability can be used to calculate the sentiment of an input utterance.

Update Delete

OneExtend Capability Attributes (6) OneExtend Capability Definitions (10) **OneExtend Definition Configs (10)**

for text Search

Capability = Sentiment Analysis

Active	Default	Definition	Domain
true	false	Sentiment Maker Suite (Google Bard Chat ...)	global
true	<b>true</b>	Sentiment (OpenAI Chat Completion)	global
true	false	Sentiment GPT4 (Azure OpenAI Chat Compl)	global
true	false	Sentiment Vertex AI (Google Bard Chat Co...)	global
true	false	Sentiment (Azure OpenAI Completion)	global
true	false	Sentiment GPT4 (OpenAI Chat Compl)	global
true	false	Sentiment Maker Suite (Google Bard Compl...)	global
true	false	Sentiment (Azure OpenAI Chat Completion)	global
true	false	Sentiment (OpenAI Completion)	global
true	false	Sentiment Vertex AI (Google Bard Complet...)	global

1 to 10 of 10

### Model Options

You have the option to choose between several different models for each capability.

Capability Definition	Model
<ul style="list-style-type: none"> <li>○ OpenAI Completion</li> <li>○ Azure OpenAI Completion</li> </ul>	GPT-3
<ul style="list-style-type: none"> <li>○ OpenAI Chat Completion</li> <li>○ Azure OpenAI Chat Completion</li> </ul>	GPT-3.5
<ul style="list-style-type: none"> <li>○ GPT4 (OpenAI Chat Compl)</li> <li>○ GPT4 (Azure OpenAI Chat Compl)</li> </ul>	GPT-4
<ul style="list-style-type: none"> <li>○ AI Studio (Google Cloud Completion)</li> <li>○ AI Studio (Google Cloud Chat Completion)</li> <li>○ Vertex AI (Google Cloud Completion)</li> <li>○ Vertex AI (Google Cloud Chat Completion)</li> </ul>	Google Gemini
Aleph Alpha Completion	Luminous
IBM watsonx	Granite
Amazon Bedrock	Titan

**Result**

Flows, topics, and scripts that use the generative AI capability use the provider you've specified.

**What to do next**

Create flows, build Virtual Agent topics, and write scripts with generative AI to improve efficiency when creating content for agents and users.

**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.

**Enable a generative AI capability in Virtual Agent Designer**

Add generative AI functionality to Virtual Agent topics to generate text, summarize information, analyze user sentiment, and interact with large language models (LLMs).

**Before you begin**

You must be on Vancouver patch 2 or a later release.

Role required: admin

**Procedure**

1. Go to the OneExtend Builder Configs table by entering `sys_one_extend_builder_config.list` in the navigation filter.
2. Open the Virtual Agent Designer record.
3. In the OneExtend Builder Capabilities related list, select **New**.
4. On the form, fill in the fields.

**OneExtend Builder Capability form**

Field	Description
Capability	Capability that is available in the Virtual Agent Designer with your preferred provider. You must have at least one provider for the capabilities, but you don't need to use the same provider for all the capabilities.
Execution Mode	Mode that determines when a capability is executed in relation to other processes. For Virtual Agent Designer, select Async. The Sync and Fire and Forget options aren't supported.

### Generic Prompt builder definition form

OneExtend Builder Capability  
Virtual Agent Designer

Active

\* Capability

\* Execution Mode

Callback Script

Description

Application

Domain

\* Builder

Icon [Click to add...](#)

5. Select **Submit**.

6. Repeat steps 3 through 5 for each capability that you want to enable.

### Result

Generative AI capabilities are available for Virtual Agent Designer.

### What to do next

Create topics and topic blocks and enhance your Virtual Agent with generative AI capabilities.

You can also access the Workflow Studio actions directly with the Action topic block in Virtual Agent Designer.

### Configure sensitive data handling for generative AI

Set up and configure how personally identifiable information and other sensitive data is removed from generative AI prompts.

### Before you begin

You can create a new Sensitive Data Regex record with a pattern or you can modify an existing record. The Generative AI category will apply the pattern to your generative AI prompts.

Role required: admin


### About this task

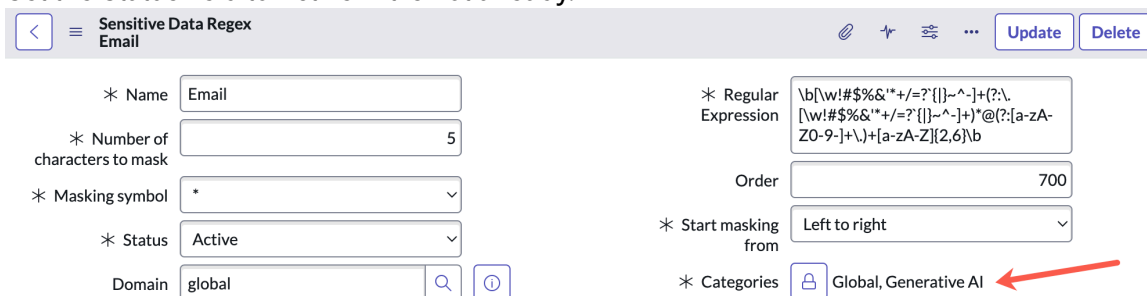
Personally identifiable information and other sensitive data can be masked so that it does not appear in generative AI prompts. Placeholder text is sent with the prompt instead, and that placeholder text is replaced with the original text after the response has been received. This two-way masking ensures your users see the correct values, but the LLM is not exposed to any sensitive information.

There are some considerations for using the Sensitive Data Handler with generative AI.

- The sensitive data detection toggle in the Conversational Interfaces Settings is unrelated to the Generative AI Controller, which uses sensitive data handling capabilities separately from this utility.
- The prompt, response, edited\_response and additional\_data fields in a conversational topic are masked during data extraction.
- The prompt is masked before sending it to the large language model (LLM) by using the sensitive data handling plugin. Only regular expressions (regexes) in the Generative AI category are used to mask the prompt.

### Procedure

1. Navigate to **All > Sensitive Data Handling > All**.
2. Select the Sensitive Data Regex record you would like to apply for generative AI prompts.
3. Unlock the Categories field by selecting the lock icon .
4. Enter **Generative AI** in the target field and select the record.
5. Select the lock icon again to lock the Categories field.
6. Set the Status field to Active if it is not already.



7. Select **Update** to save the record.

### 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.

### 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.

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.

## 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

### 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 [Using language detection and dynamic machine translation in Virtual Agent enhanced chat conversations](#) .


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.

## 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: <ul style="list-style-type: none"> <li>○ <b>LLM Provider:</b> For the LLM Provider, select the generative AI provider mapping.</li> <li>○ <b>Capability Definition:</b> For the Capability Definition, select the One Extend Capability Definition.</li> <li>○ <b>All Capabilities:</b> Applies to all providers and all capability definitions.</li> </ul>
Rule Type	Rule type for the rate limit rule.  Define the rule type using the following options: <ul style="list-style-type: none"> <li>○ <b>Instance:</b> Accumulates data from all the users of that instance.</li> <li>○ <b>User:</b> Keeps the count of the requests per user.</li> </ul>
Request limit per hour	Maximum number of requests to be received per hour.
Active	Status of the rate limit rule.

3. Select **Submit**.

## Using Generative AI Controller

Use Generative AI Controller to integrate generative AI with Workflow Studio, Virtual Agent Designer, and scripting.

### Use Generative AI Controller with Workflow Studio

Use the Generative AI Controller capabilities in flows with Workflow Studio.

### Use Generative AI Controller with Virtual Agent Designer

Use the Generative AI Controller capabilities in Virtual Agent topics with Virtual Agent Designer.

### Use Sentiment Analysis with Workflow Studio

Use the Analyze Sentiment capability to determine the user attitude and the use that results in a flow with Workflow Studio.

### Use a script with Generative AI Controller

Write scripts to harness the power of generative AI directly.

### Use Generative AI Controller with Workflow Studio

Summarize the information automatically and add it to a form with the Generative AI Controller Summarize action in Workflow Studio. You can build a flow to summarize any field.

### Before you begin

Role required: admin

### About this task

In this example, you can build a flow to summarize the comments in the work notes of an incident and use that summary in the resolution notes.

For more information on tracking Generative AI Controller usage, see [Monitoring Now Assist usage](#).

### Procedure

1. Navigate to **All > Process Automation > Flow Designer**.
2. Select **New > Flow**.
3. Enter the name of the flow, such as **Summarize Comments and Work Notes**.
4. Select **Submit**.
5. Select **Add a trigger**.
6. Select when you would like the flow to run.

To follow this example, select **Record > Updated**.

**Work Notes Summary** Inactive View:

TRIGGER

Trigger

RECORD	Updated
Created	
Updated	Trigger initiates from a ServiceNow record update that meets the condition filter. Configure the trigger to initiate the flow (Once) triggers the flow once for the life of the record, (Only if not currently running) triggers the flow for every unique change if the flow is not currently running, (For each unique change) triggers the flow for every unique change, even if the flow is currently running, or (For every update) triggers the flow for every update to the record, including if the same update previously occurred.
Created or Updated	
SCHEDULED	
Daily	Note: Flows that have a record trigger that runs (For each unique change) can produce recursions when run in a non-interactive session. When such flows make a change to the trigger record, the change meets the flow trigger conditions and causes a recursion.
Weekly	
Monthly	
Run Once	Note: For a flow that uses Approvals, it is recommended that you configure flows to run (Once).
Repeat	

APPLICATION

7. Select the table that the flow should run on.

For this example, select the Incident table.

8. Select the **Add filters** button to add a condition for the flow.

The Condition is **State changes to Resolved**.

TRIGGER

Incident Updated where (State changes to 6) + 🗑️

Trigger

\* Table  ✕

Condition All of these conditions must be met


or

Run Trigger

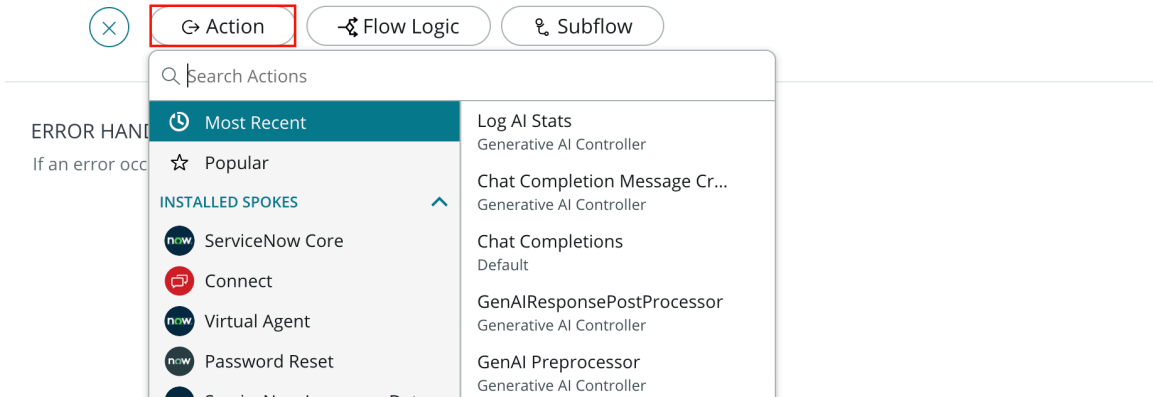
9. Select **Done**.

10. Under Actions, select **Add an Action, Flow Logic**, or **Subflow** and then select **Action** to open the Actions panel.

TRIGGER

 Incident Updated where (State changes to 6)

ACTIONS Select multiple

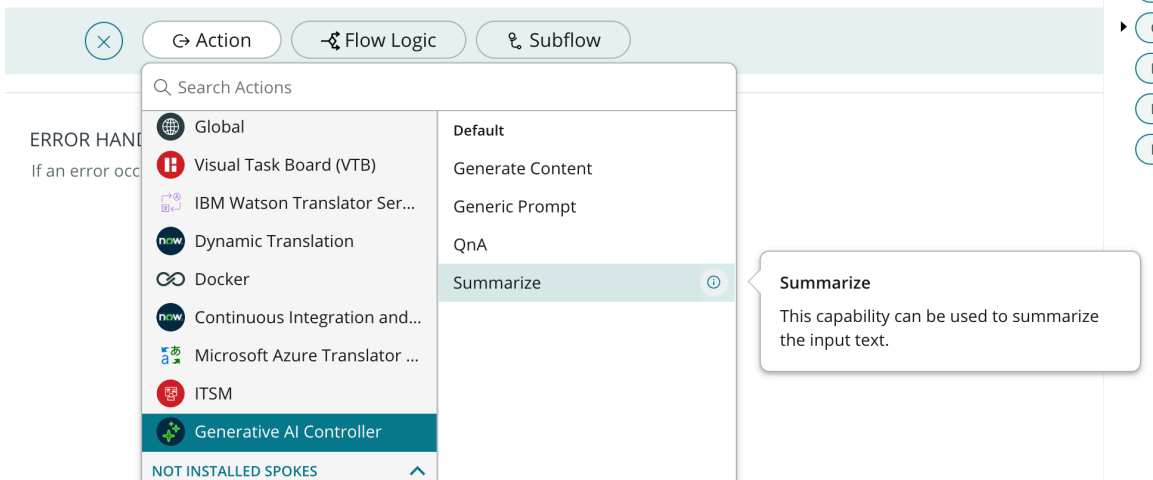



**11. Select the **Generative AI Controller** > **Summarize** action and select it.**

TRIGGER

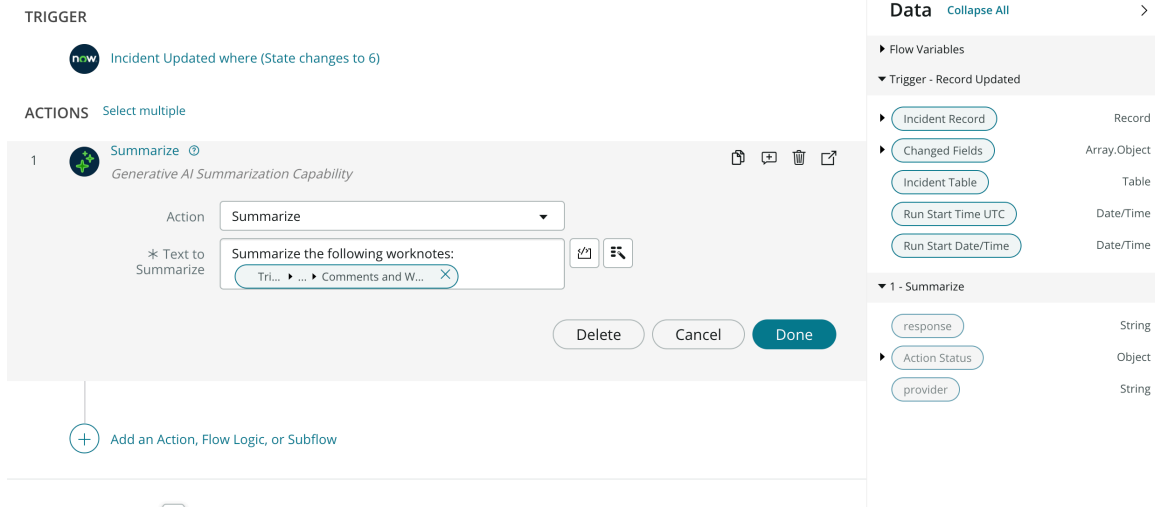
 Incident Updated where (State changes to 6)

ACTIONS Select multiple



**12. Drag a data pill from the Data panel or select the data pill picker icon (  ) to search. In this case, drag the **Trigger** > **Incident Record** > **Comments and Work notes** pill.**

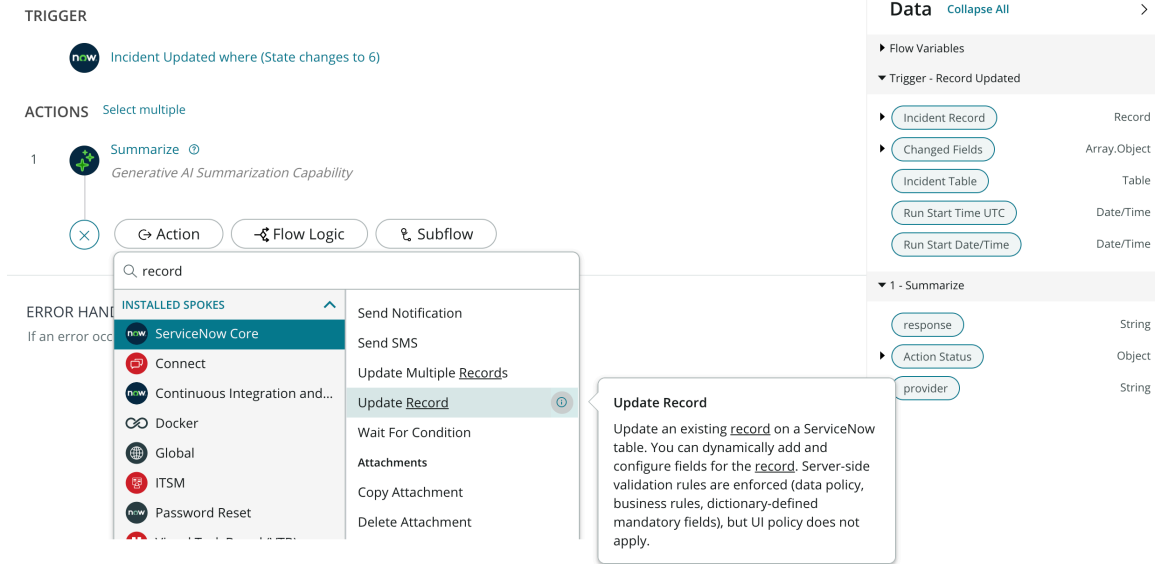
**Note:** If the field that you want to summarize isn't a String field, you must cast the field as a String. You can cast the field as a String with scripting or by adding text to the **textToSummarize** field in addition to the data pill, like in the following example.



13. Select **Done**.

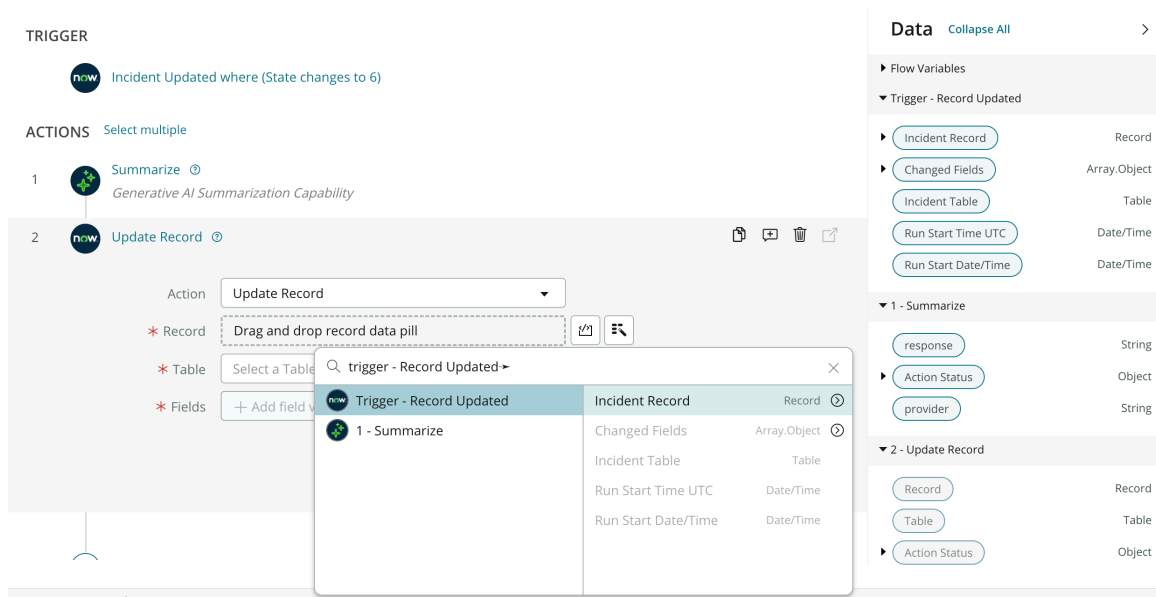
14. Add an action by selecting **Add an Action, Flow Logic, or Subflow**.

15. In the action selector, search for **Update Record** and select the action.



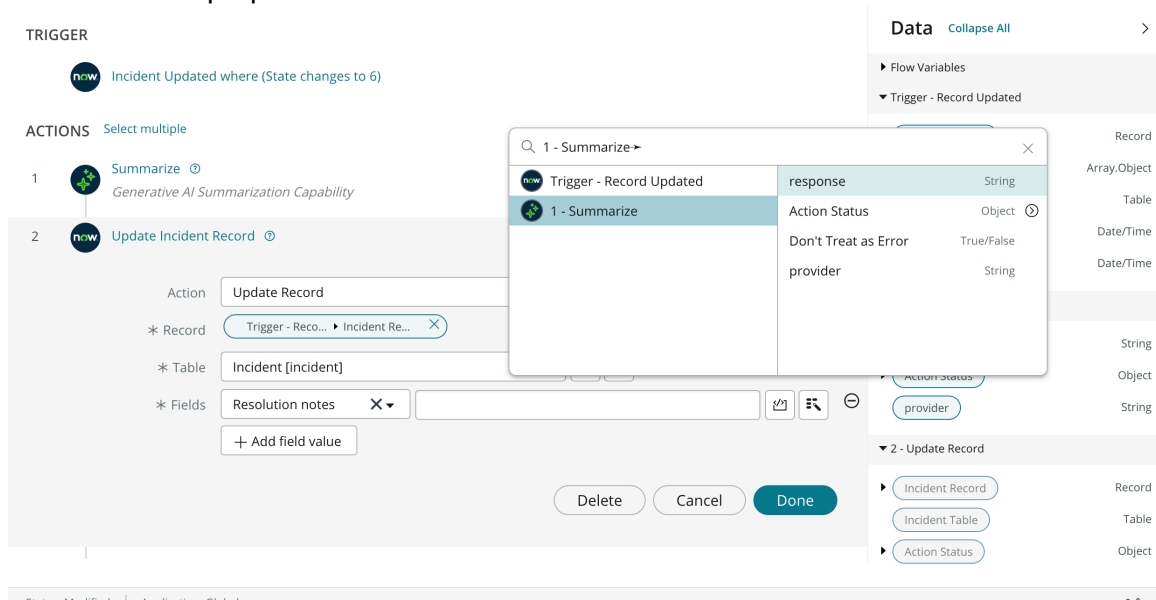
16. Drag the **Trigger > Incident Record** data pill into the **Record** field or select the icon to search.

The Table field is filled in automatically based on the Trigger record.

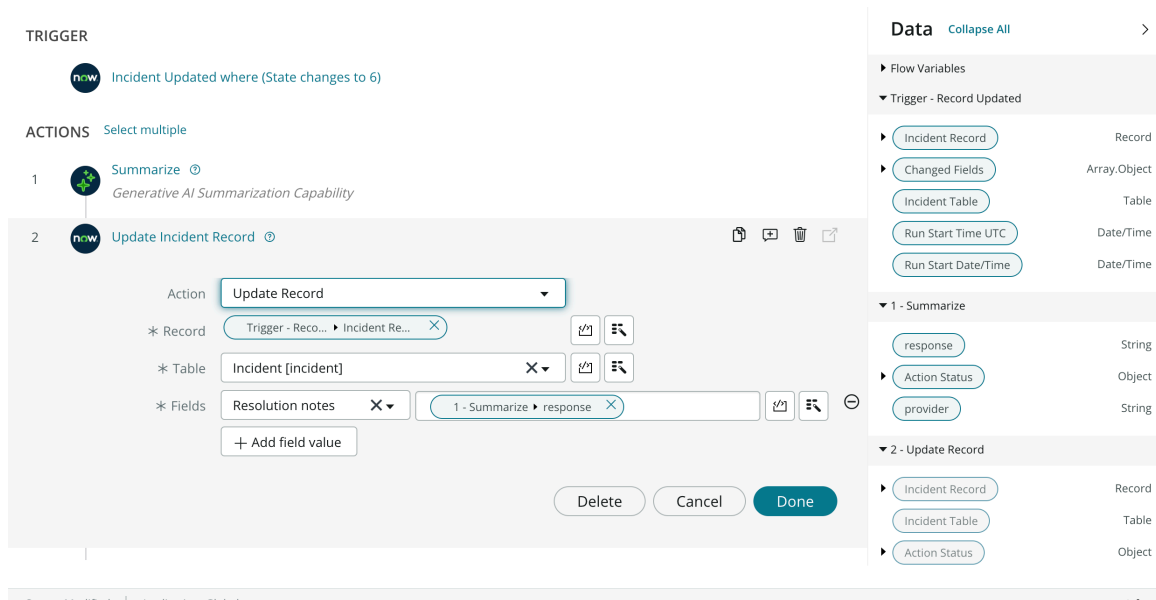


17. Select **Add field value** and search for the fields that you want to update. In the following example, select **Resolution notes**.

18. Add the **Summarize > Response** data pill into the data field by dragging it from the Data panel or with the data pill picker.



19. Select **Done**.



20. Save the flow by selecting **Save**.

### What to do next

After you've saved the flow, activate it to migrate it to production or test it further.

### Use Generative AI Controller with Virtual Agent Designer

Use Generative AI Controller capabilities inside Virtual Agent topics to draft knowledge articles.

### Before you begin

You must enable the capabilities for Virtual Agent Designer first. For more information, see [Enable a generative AI capability in Virtual Agent Designer](#).

Role required: admin

### About this task

This example uses the Generate Content capability of the Generative AI Controller to draft knowledge article content from a Virtual Agent conversation.

For more information on tracking Generative AI Controller usage, see [Monitoring Now Assist usage](#).

### Procedure

1. Navigate to **All > Conversational Interfaces > Virtual Agent > Designer**.
2. Create a topic by selecting the Create card.
3. Enter a unique name for the topic.

More topic options can be configured, such as Natural Language Understanding (NLU) model or access restrictions.

This example uses `KB Article Generation` as the topic name.

4. Select **Create**.
5. Create a user input node between Start and End by dragging the **User Input > Text** component from the Components panel.
6. Enter a node name on the info panel on the canvas.

In this example, the node name is KB Article Content Request.

The screenshot displays the ServiceNow flow designer interface. On the left, a 'Components' panel lists various input and response nodes. The 'Text' node under 'User Input' is highlighted with a red box. The main workspace shows a flow starting at a 'Start' node, leading to a 'KB Article Content Request' node, which then leads to an 'End' node. On the right, the configuration panel for the 'Text input' node is visible, showing fields for Node name, Variable name, Prompt, Input format, and NLU entity.

7. In the prompt field, enter a message to display to the user.

Because the user supplies the topic for the knowledge base (KB) article content, the prompt in this example is What topic would you like to see a Knowledge Base article about?.

8. Add User Input nodes to fill in all required fields.

**Note:** To create a record successfully, all required fields must have a value. The next steps describe how to fill in the required fields for a KB article. If your form has no required fields, you don't need to add any User Input nodes.

- a. Add another **User Input > Text** node and name it **Article Title Request**.
- b. Add a prompt for the required **Article title** field such as **What should the title of the Knowledge Article be?**.

The screenshot shows the ServiceNow Flow Designer interface. The flow consists of three nodes: a 'Start' node, a 'KB Article Content Request' node, and an 'Article Title Request' node, all connected in a vertical sequence. The 'Article Title Request' node is highlighted. On the right, the configuration panel for the 'Text input' node is visible, showing the following settings:

- Node name: Article Title Request
- Variable name: article\_title\_request
- Prompt: What should the title of the Knowledge Article be?
- Input format: Text
- NLU entity: Define an NLU entity from a model

9. On the Components panel, scroll down until you see the Capabilities section and drag the **Generate Content** action next in the Topic flow.

10. In the Info panel, name the node.

The node name that is used in the example is **AI KB Generator**.

The screenshot shows the ServiceNow Flow Designer interface with the 'Generate Content' node added to the flow. The flow now consists of four nodes: 'Start', 'KB Article Content Request', 'Article Title Request', and 'AI KB Generator', connected in a vertical sequence. The 'AI KB Generator' node is highlighted. On the right, the configuration panel for the 'Generate Content' node is visible, showing the following settings:

- Node name: AI KB Generator
- Wait for response? (checked)
- topic (string): [empty field]
- Output mapping table:

Enable	Name	Variable Name
<input checked="" type="checkbox"/>	provider (string)	Provider
<input checked="" type="checkbox"/>	response (string)	Response
<input checked="" type="checkbox"/>	error (string)	Error
<input checked="" type="checkbox"/>	status (string)	Status

11. Open up the data pill selector by selecting the data pill icon (🍯) next to the topic field.

The screenshot shows the ServiceNow Flow Designer interface. The flow consists of the following steps: Start, KB Article Content Request, Article Title Request, AI KB Generator, and End. The 'AI KB Generator' node is highlighted with a red '1' and a data pill icon. The configuration panel for the 'Generate Content' node is open on the right. In the 'topic (string)' field, a data pill icon (🍯) is selected, opening a dropdown menu. The dropdown menu shows 'Input Variables' and 'Script Variables'. Under 'Input Variables', 'Article Title Request' is selected, and its value is 'String'. The 'status (string)' field is also visible with a checked checkbox and a 'Status' variable name.

12. Use the text that was created in the interaction by selecting the **Input Variables > Response** variable.

This screenshot shows the same flow as in step 11, but the configuration panel for the 'Generate Content' node is updated. The 'topic (string)' field now contains the text 'Input Variables > KB Article...' with a data pill icon. The 'Output mapping' section is expanded, showing a table with the following entries:

Enable	Name	Variable Name
<input checked="" type="checkbox"/>	provider (string)	Provider
<input checked="" type="checkbox"/>	response (string)	Response
<input checked="" type="checkbox"/>	error (string)	Error
<input checked="" type="checkbox"/>	status (string)	Status

13. Drag the **Utilities > Record Action** component from the Components panel after the Generate Content node.

14. Enter a name for the node.

For this example, the node name is **Create KB Article**.

15. Fill in the other required fields.

a. For the Action type, select whether you want to create a record or update an existing one.

In this example, the Action type is **Create a record**.

b. Select the table for the new or updated record.

This example uses the **Knowledge (kb\_knowledge)** table.

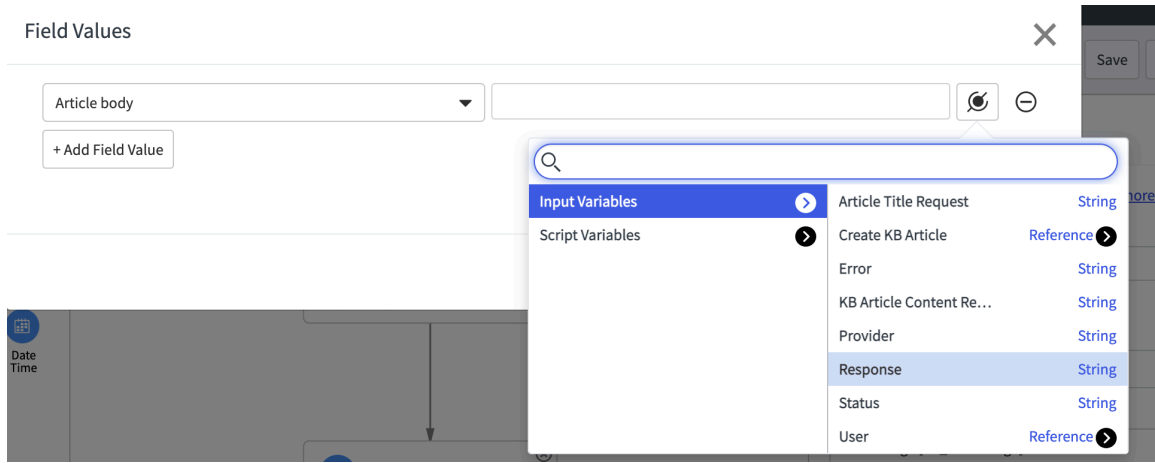
16. Under Field, select the **Add Field** button to open the Field Values window.

17. Set a field value by selecting **Add Field Value**.

18. Select the field that you would like to update from the list.

You can add as many fields as you like. In this example, the fields are **Article body**, **Knowledge base**, and **Short description**.

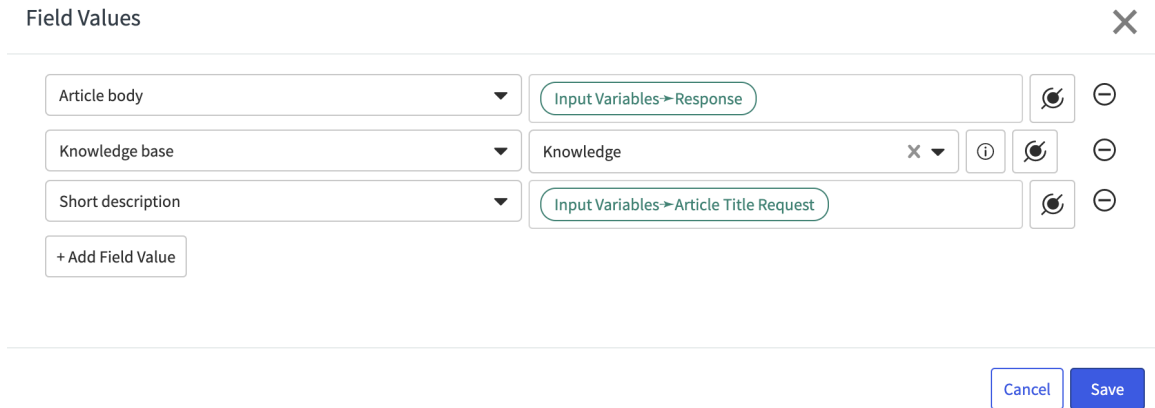
19. Select the data pill icon to open the data pill selection menu.



20. For the **Article body** field, select **Input Variables > Response** to use the response that is generated by the Generate Content node.

21. For the **Knowledge base** field, select **Knowledge** or your preferred category.

22. For the **Short description** field, select the data pill icon and add **Input Variables > Article Title Request**, or the name of the user input node that is prompting for a title.



23. Save your field selections and close the Field Values window by selecting **Save**.

24. Save the topic by selecting **Save**.

### What to do next

This example topic creates simple knowledge articles. Use additional user input, such as context and background information, for more complex articles.

**Note:** Review any generated content before you publish. See [AI Limitations](#) for more details.

### Use Sentiment Analysis with Workflow Studio

Create a flow to upgrade the assignment group of a case if the user has a negative sentiment in the short description.

### Before you begin

Role required: admin

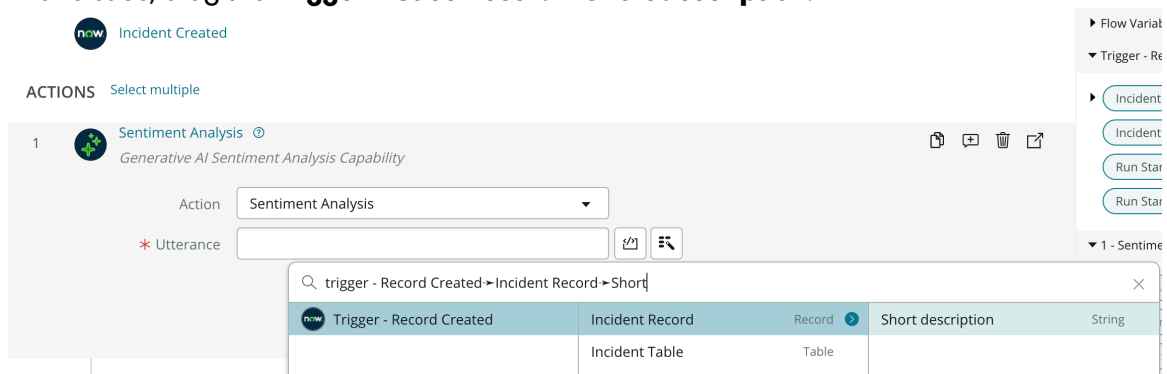
### About this task

The Sentiment Analysis capability helps determine whether the user-generated text has a positive or negative sentiment. If a user is expressing negative statements such as "Level 1 support couldn't help me," you could upgrade the assignment group automatically for the case to help the user have a better experience.

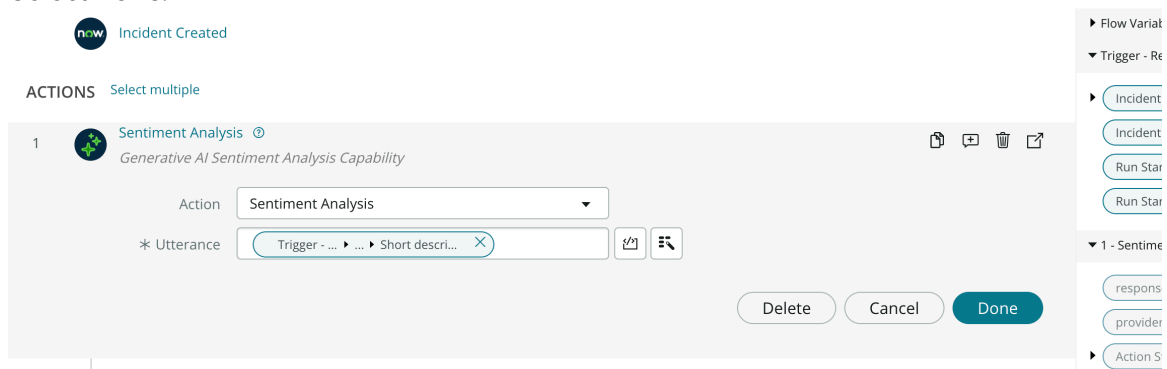
For more information on tracking Generative AI Controller usage, see [Monitoring Now Assist usage](#).

### Procedure

1. Navigate to **All > Process Automation > Flow Designer**.
2. Select **New > Flow**.
3. Enter the name of the flow, such as `Change assignment group if negative sentiment`.
4. Select **Submit**.
5. Select **Add a trigger**.
6. Select **Record > Created** to update a new record.
7. Select the table that the flow should run on.  
For this example, select the Case (sn\_customerservice\_case) table.
8. Select **Done**.
9. Under Actions, select **Add an Action, Flow Logic, or Subflow** and then select **Action** to open the Actions panel.
10. In the action selector, search for the **Generative AI Controller > Sentiment Analysis** action and select it.
11. Drag a data pill from the Data panel or select the data pill picker icon.  
In this case, drag the **Trigger > Case Record > Short description**.

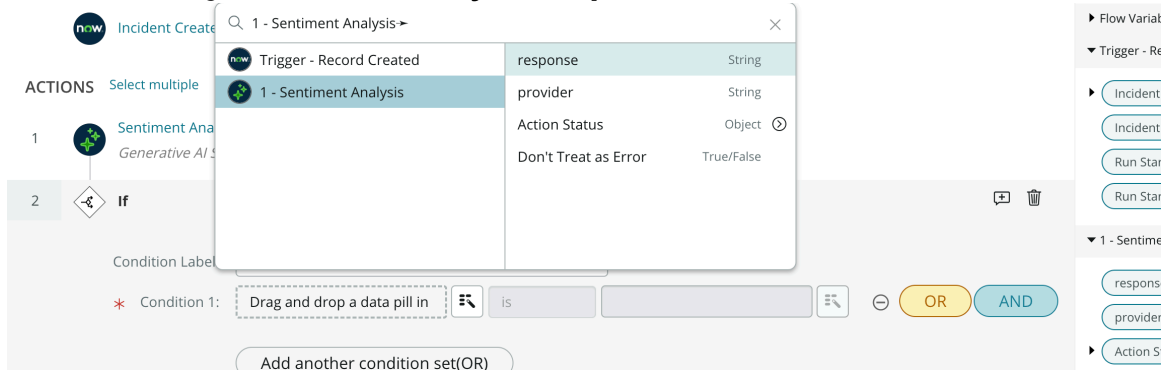


12. Select **Done**.



13. Add an action by selecting **Add an Action, Flow Logic, or Subflow**.
14. Select **Flow Logic > If**.
15. Give the **Condition label** a name, such as `Negative sentiment`.
16. For the **Condition** field, drag a data pill or select the data pill picker icon to search for the correct condition.

In this case, drag the **Sentiment Analysis > response**.



17. Select **contains the value Negative**.

**Note:** Capitalization matters. This flow won't work if the value is set to **negative**.



18. Search for **Update Record** and select the action.

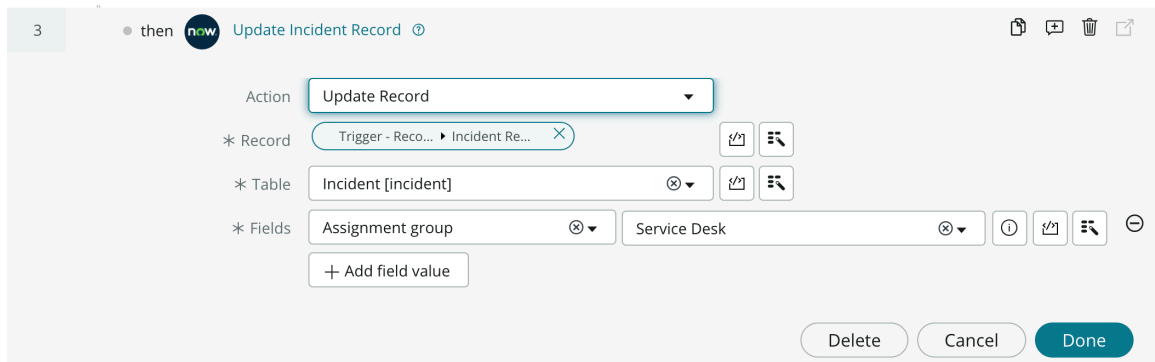
19. Drag the **Trigger > Case Record** data pill into the **Record** field or select the icon to search.

The **Table** field is filled in automatically based on the Trigger record.

20. Select **Add field value** and search for the fields that you want to update. In this example, select **Assignment group**.

21. Select which value you would like to set the field to.

The field value depends on your use case. In this example, the value is set to **Service Desk**.



22. Select **Done**.

23. Save the flow by selecting **Save**.

### What to do next

After you've saved the flow, you must activate the flow to migrate it to production or test it further.

### Use a script with Generative AI Controller

Use background scripts and Generative AI Controller Generic Prompt to interact directly with the generative AI model API.

#### Before you begin

You must be on Vancouver patch 2 or above to script with `sn_one_extend.OneExtendUtil` objects.

Role required: admin

#### About this task

With scripting, you have greater flexibility in how you use Generative AI Controller capabilities across the platform. By accessing the capability programmatically, you can capture the LLM's response and use that to determine next actions.

For more information on tracking Generative AI Controller usage, see [Monitoring Now Assist usage](#).

#### Procedure

1. Navigate to **All > System Definition > Scripts - Background**.
2. Create the text content for a Generic Prompt by creating a `request` array with a singular `executionRequests` object.
3. Inside the `executionRequests`, add the `payload` and `capabilityID` attributes. The `capabilityID` is the `sys_id` of the capability on the OneExtend Capabilities (`sys_one_extend_capability`) table, such as Generic Prompt.
4. Add the required inputs as attributes in the `payload` object.

Capability	Required attributes
Generic Prompt	<code>prompt</code> : Prompt for the generative AI to respond to.
Generate Content	<code>topic</code> : Subject for content generation.
Sentiment Analysis	<code>utterance</code> : Text to analyze for user sentiment.
Summarize	<code>textToSummarize</code> : Topic for the AI to generate a summary.

The following is an example `request` array for the Generic Prompt capability.

```
var request = {
  "executionRequests": [
    {
      "payload": {
        "prompt": "Can you act like my business partner and
        give me some advice on a pitch?"
      },
      "capabilityId": "0c90ca79533121106b38ddeeff7b12d7"
    }
  ]
};
```

5. Call the `execute` method on `sn_one_extend.OneExtendUtil`.

**6. Optional:** To display the response after running the background script, print the response in readable JSON by using the following code:

```
gs.info(JSON.stringify(sn_one_extend.OneExtendUtil.execute(request)));
```

## Reference for Generative AI Controller

Reference topics provide information about Generative AI Controller flow actions, tables, and properties.

### Flow actions

Action	Description
Generate Content	Generate content with an artificial intelligence model to create bodies of text quickly, such as generating knowledge base articles from interactions, cases, or incidents.
Summarize	Summarize large bodies of text or complicated exchanges to make the transfer of information easier, such as during handoffs between virtual and live agents.
Sentiment Analysis	Analyze user sentiment based on text that has been input in interactions and forms to improve customer experiences.
Generic Prompt	Use Generative AI to generate ideas and brainstorm on any topic, such as ideas for increasing engagement.

### 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.

Name	Table	Description
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 Metadatas	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, and guardrail activity.
Gen AI Model Version Mapping	sys_gen_ai_model_version_mapping	Maintains mappings between AI models, their providers, and configuration details. It shows how source models map to target models along with metadata like skill type, model type, resource associations, and provider information.

## 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

### Generative AI Controller properties (continued)

Property	Description
	Default value: true
com.sn.generative.ai.moderation.message	<p>Message that is displayed if the OpenAI or Azure OpenAI moderation tools identify the content that goes against their terms of service.</p> <p>Type: string</p> <p>Default value: The response cannot be displayed as it's deemed inappropriate by OpenAI.</p>
com.glide.one.extend.token.buffer	<p>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.</p> <p>Type: integer</p> <p>Default value: 250</p>

### External links

Provider	Data policy	Usage policy
Amazon Bedrock	<a href="#">Data protection</a>	<a href="#">AWS Service Terms</a>
Aleph Alpha	<a href="#">Data privacy</a>	<a href="#">Terms and Conditions</a>
Google Cloud	<a href="#">Google Cloud Platform Terms of Service</a>	<a href="#">Google Cloud Platform Terms of Service</a>
IBM watsonx	<a href="#">Keeping your data secure and compliant</a>	<a href="#">Foundation model terms of use in watsonx.ai</a>
Microsoft Azure OpenAI	<a href="#">Data, privacy, and security for Azure OpenAI Service</a>	<a href="#">Code of conduct for Azure OpenAI Service</a>
OpenAI	<a href="#">API data usage policies</a>	<a href="#">Usage policies</a>

### Generate Content action

Use the Generate Content action with Workflow Studio and Virtual Agent to generate bodies of text.

## Fields

Field	Description
Topic	Subject for the AI model to generate text about. You can enter text directly or use pills from the Data panel.

## Inputs

Provide a value for each input that your flow needs. To add dynamic values, you can also drag pills from the Data panel or select them from the pill picker.

### Topic

Data type: *String*

Subject for the AI model to generate text about. You can enter text directly or use pills from the Data panel.

## Outputs

These outputs appear in the Data panel. You can use them as inputs elsewhere in your flow.

### Response

Data type: *String*

The content generated by the AI model.

### Provider

Data type: *String*

Name of the AI model that generated the response.

### Action Status

Data type: *String*

Details about the success of the action.

## Generic Prompt action

Use the Generic Prompt action for Workflow Studio to respond to a general prompt with an AI model.

## Fields

Field	Description
Prompt	Statement or question for the AI model to generate a response to. You can enter text directly or select pills from the Data panel.

## Inputs

Provide a value for each input that your flow needs. To add dynamic values, you can also drag pills from the Data panel or select them from the pill picker.

### Prompt

Data type: *String*

You can enter text directly or select pills from the Data panel.

## Outputs

These outputs appear in the Data panel. You can use them as inputs elsewhere in your flow.

### Response

Data type: *String*

Content generated by the AI model.

### Provider

Data type: *String*

Name of the AI model that generated the response.

### Action Status

Data type: *String*

Details about the success of the action.

## Sentiment Analysis action

Use the Sentiment Analysis action with Workflow Studio and Virtual Agent to detect the sentiment of your users from their input.

## Fields

Field	Description
Utterance	Text to analyze for user sentiment.

## Inputs

Provide a value for each input that your flow needs. To add dynamic values, you can also drag pills from the Data panel or select them from the pill picker.

### Utterance

Data type: *String*

User-generated text to analyze the sentiment. You can enter text directly or use pills from the Data panel.

## Outputs

These outputs appear in the Data panel. You can use them as inputs elsewhere in your flow.

### Response

Data type: *String*

Content generated by the AI model.

### Provider

Data type: *String*

Name of the AI model that generated the response.

### Action Status

Data type: *String*

Details about the success of the action.

## Summarize action

Use the Summarize action with Workflow Studio to summarize information with the AI model.

## Fields

Field	Description
Text to Summarize	Text to summarize by the AI model. You can enter text directly or drag a pill from the data panel.

## Inputs

Provide a value for each input that your flow needs. To add dynamic values, you can also drag pills from the Data panel or select them from the pill picker.

### Input

Data type: *String*

Text to summarize by the AI model.

**Note:** You must add some amount of text directly as well as any data pills from the pill picker.

## Outputs

These outputs appear in the Data panel. You can use them as inputs elsewhere in your flow.

### Response

Data type: *String*

Content generated by the AI model.

### Provider

Data type: *String*

Name of the AI model that generated the response.

### Action Status

Data type: *String*

Details about the success of the action.






## Now Assist AI agents

The ServiceNow<sup>®</sup> 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.

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.

<p>Explore</p>  <p>Learn more about Now Assist AI agents</p>	<p>Configure</p>  <p>Configure the agentic AI experience in AI Agent Studio</p>	<p>Create an AI agent</p>  <p>Create your own custom AI agents with advanced multi-agent reasoning frameworks</p>
<p>Create an agentic workflow</p>  <p>Create your agentic workflows with planned instructions and triggers</p>	<p>Reference</p>  <p>Learn more about the user roles, tables, and properties in Now Assist AI agents</p>	

Looking for an AI agent?

- 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 on your instance, see [Find AI agents](#).
- To find agents that might not be installed on your instance, visit the [AI Agent Marketplace](#) on the ServiceNow Store.

**i 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. Further, 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](#).

## Exploring Now Assist AI agents

Learn how the Now Assist AI agents application can help to improve live agent productivity by using AI agents to perform tasks. AI agents simulate human-like intelligence to handle various tasks that range from automated responses to complex problem solving.

### Now Assist AI agents terminology

- **Agentic system:** An [agentic system](#) is a type of software or AI that perceives its environment, decides 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.
- **AI agent:** On the ServiceNow AI Platform agentic system, an AI agent contains a set of large language model (LLM) instructions with the tools to perform these specific tasks.
- **Agentic workflow:** In the agentic system on your instance, an agentic workflow contains a set of LLM instructions with one or more AI agents that can execute an objective.

### AI agents overview

The Now Assist AI agents application is designed to securely leverage your data, workflows, and integrations natively on the ServiceNow AI Platform. 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.

Use AI agents to do the following tasks for your organization:

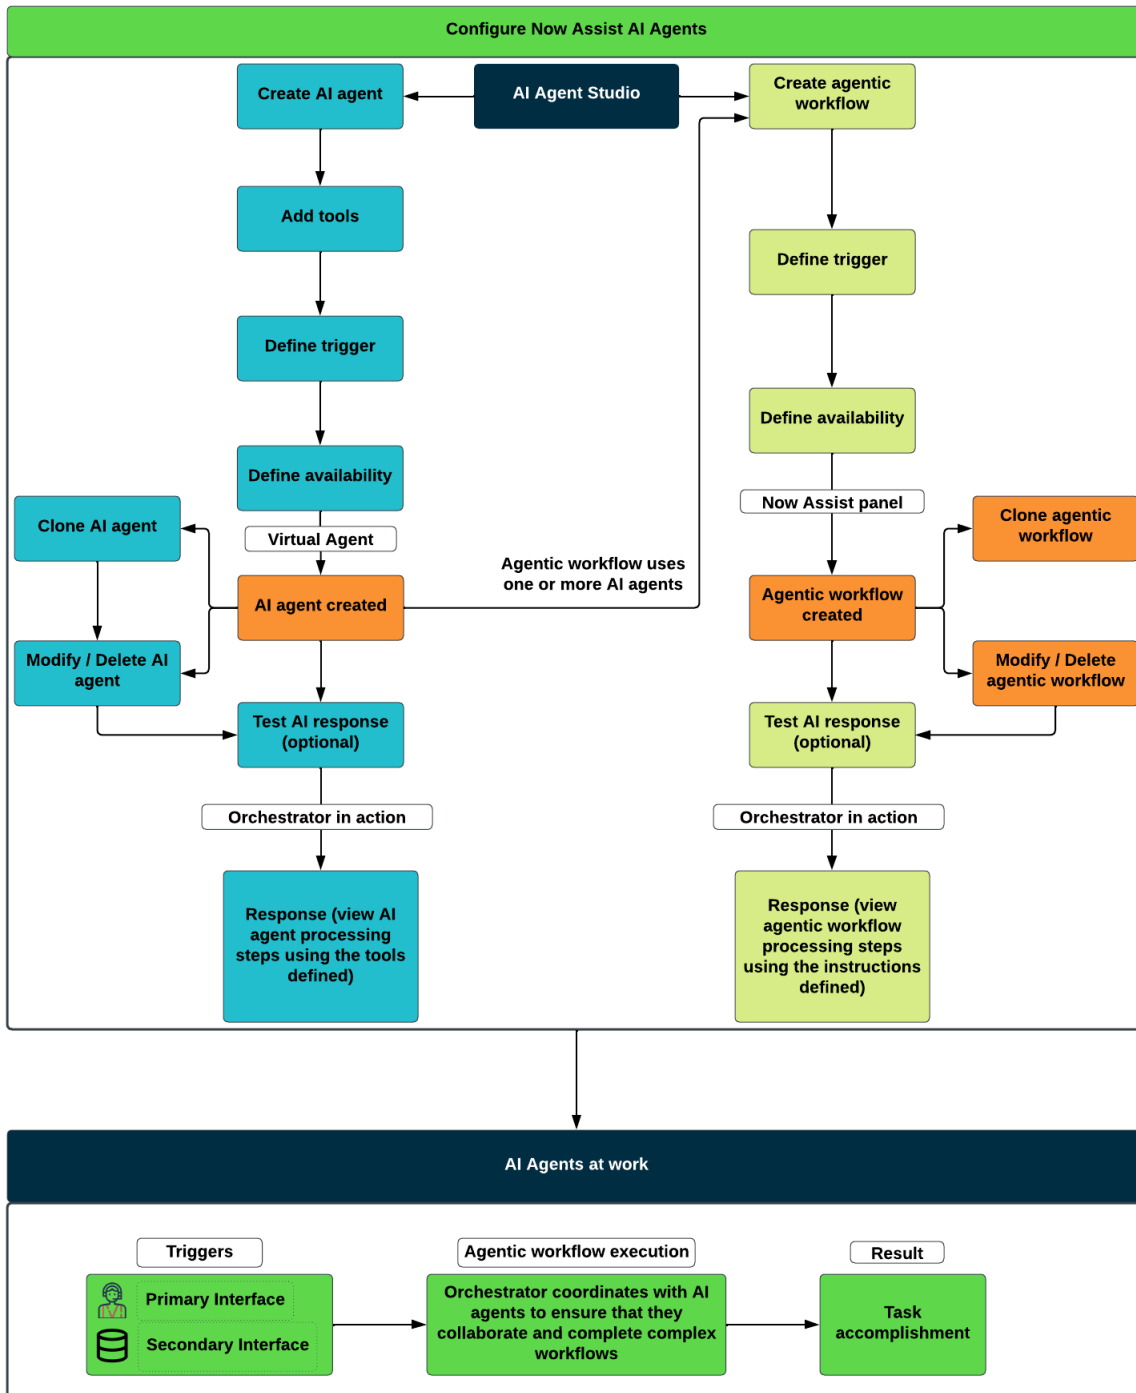
- Dynamically figure out a plan to resolve an incident or case.
- Collaborate with other AI agents for subtasks as needed.
- Take feedback from humans as needed.
- Work on knowledge-intensive tasks that need information, understanding, and skills to accomplish that task.
- Find a solution by using multiple sources of knowledge or similar incidents.

- Plan to solve an issue with human collaboration.
- Avoid iterations and discovery by defining a specific flow.

You can create, duplicate, or modify agentic workflows. You must activate the base system agentic workflows before they can be used. For more information about activating an agentic workflow, see [agentic workflow template](#). For more information about cloning an agentic workflow, see [Clone an agentic workflow](#).

The base system agentic workflows remain read only. You must duplicate them to use them. For more information about cloning an agentic workflow, see [Clone an agentic workflow](#). You can modify the duplicated base system agentic workflow to meet your requirements. For more information, see [Modify an agentic workflow](#). You must activate the base system agentic workflows before they can be used. For more information about activating an agentic workflow, see [agentic workflow template](#).

## 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

panel. To learn more about what an AI Agent Orchestrator is, see the [Now Assist AI Agents Orchestrator](#) section in this topic.

## 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.

## Now Assist AI agents Orchestrator

You can use the AI Agent Orchestrator to verify that your AI agent teams work together and track agentic workflows across the enterprise.

The AI Agent Orchestrator in the agentic framework is a coordinator entity that handles collaboration across multiple agents by routing requests from one agent to other agents to help accomplish a complex task. The AI Agent Orchestrator can also reach out to one of the agents involved in the task accomplishment to get the missing context or information from the user for task completion.

AI agents don't work in isolation. If they get stuck or fail completing a task, they can go back to the AI Agent Orchestrator and ask for help, using the iterative aspect of the AI agents.

## Citations in Now Assist AI agents

When you execute AI agents and agentic workflows in Now Assist AI agents, you can see citations on the Now Assist panel by default 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.

What to test: Steps for Issue Resolution
Task: help me resolve INC0010003 Test
Restart

Agent Studio Chat

There are similar incidents related to a boot-up error in the Cisco Webex Meetings server. The affected user in these incidents is Abel Tuter. Incident INC0010017 was created from a similar incident INC0010016. Incident INC0010016 was created from incident INC0010003. Incident INC0010017

Show less

**Sources**

- 1 Incident INC0010017
- 2 Incident INC0010016

just now

There are several knowledge articles available for resolving boot-up errors in Cisco WebEx Meetings Server. One article, KB0010003, provides a detailed guide for when the server does not boot up. Another article, KB0010002, offers instructions for issues occurring after...

Show more

**Sources**

- 1 Detailed guide for KB0010003
- 2 Instructions for KB0010002
- 3 Troubleshooting steps for KB0000052

▶  
Task Start

↓  
Orchestrator

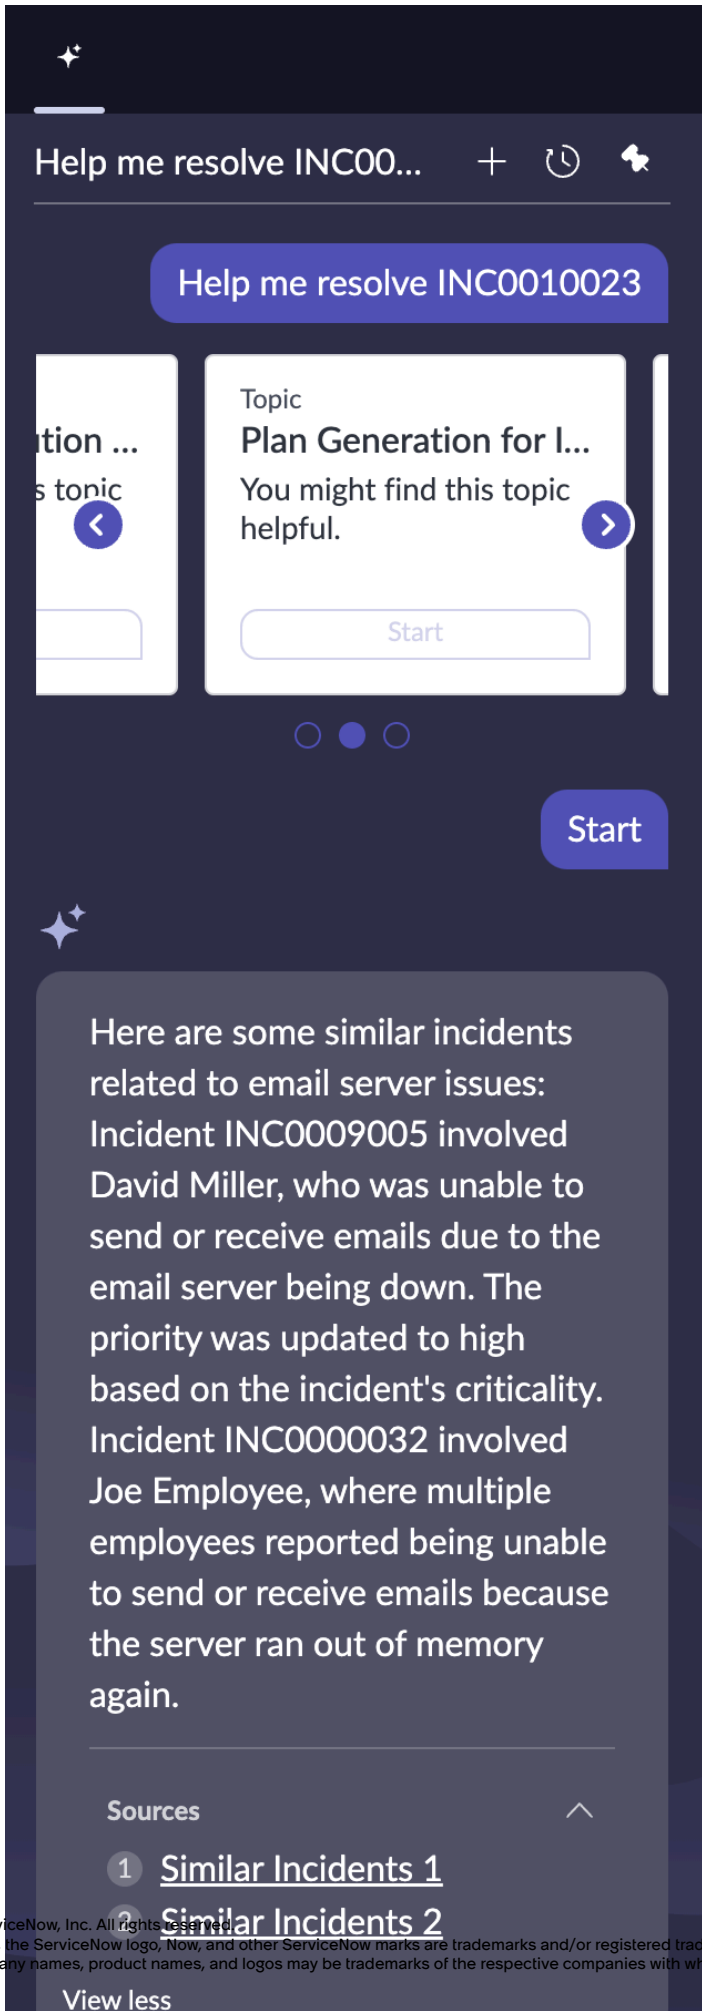
↙  
⦿  
Communicator  
AI Agent

↘  
⦿  
Next Best Action A...  
AI Agent

### AI agent decision logs

Observe the AI agents as they work to solve the task. Watch their interactions, decisions, and thought processes as they happen in real time.

Orchestrator	Success
Next Best Action Agent	Success
Communicator	Ongoing
Next Best Action Agent	Queued



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](#).

## 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.

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.

## 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&autoplay=0&player\\_id=0&app\\_id=58479](https://player.vimeo.com/video/1096896405?h=81ff7b3861&badge=0&autoplay=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.

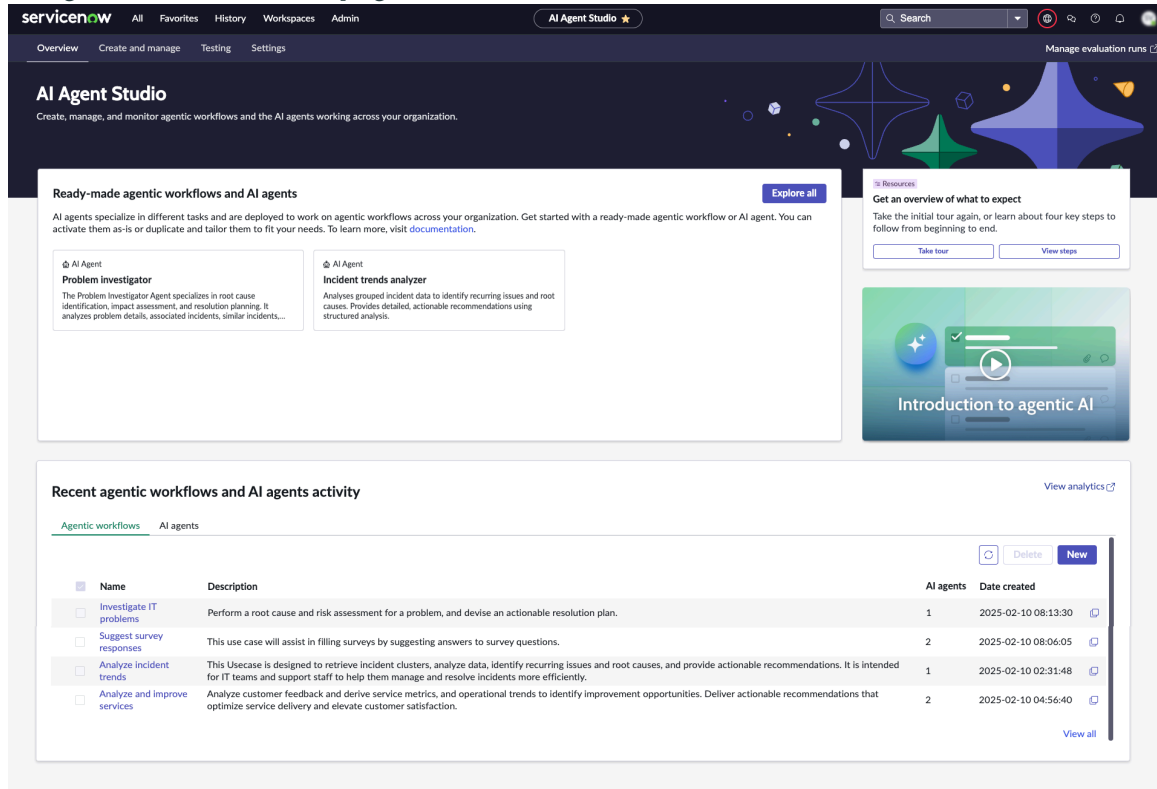
In the Ready-made agentic workflows and AI agents section, you can find the templates and ready-made AI agents and agentic workflows that you can incorporate as-is or in custom workflows that you create. You can even explore more templates to find the ones that best fit your needs.

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.

## AI Agent Studio overview page

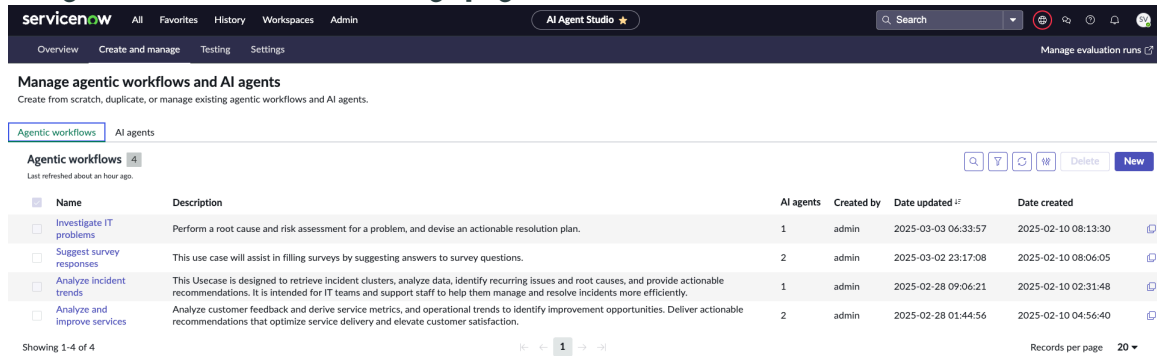


## Managing agentic workflows and AI agents

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



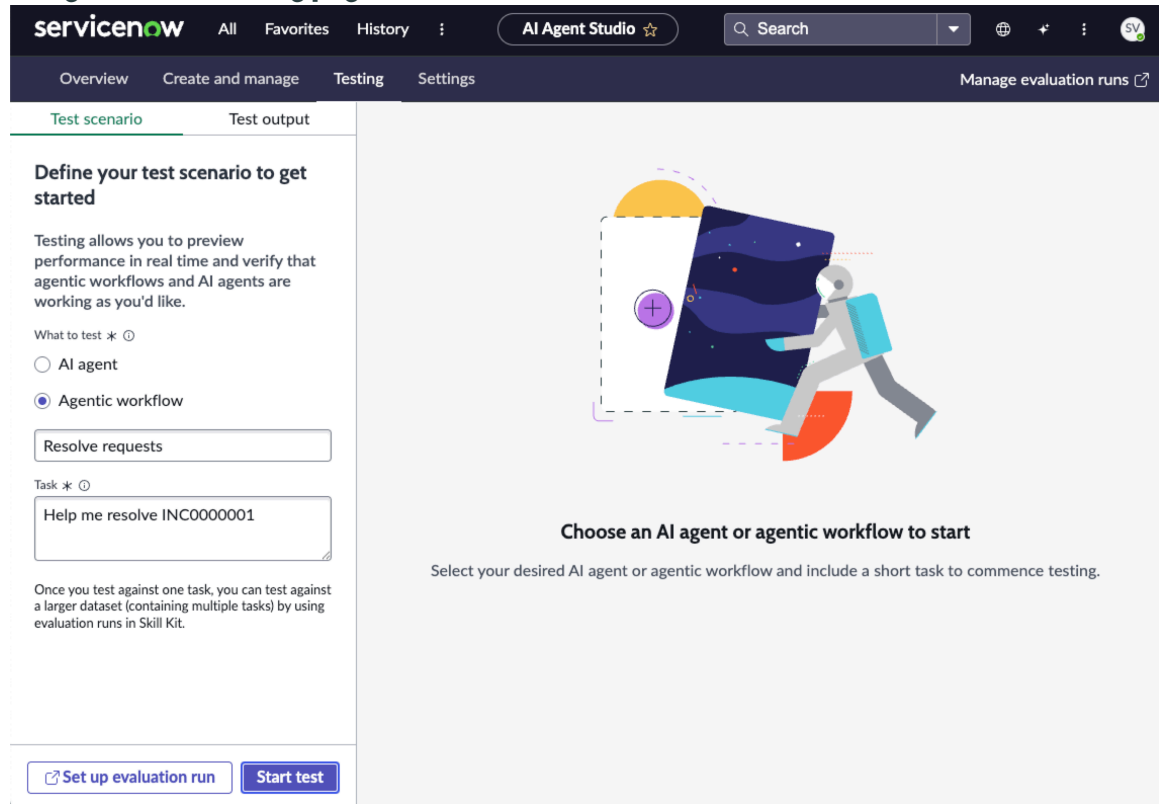
## AI agent testing

From the AI Agent Studio testing page, you can test your agentic workflows and AI agents by asking them to complete short tasks that would be part of their more advanced workflows. For example, you could task the Resolve an incident agentic workflow to resolve a specific incident.

After the task begins, you can see the AI agent conversation, the AI Agent Orchestrator, and the chain of logic that is used to achieve the task.

The following example shows the inputs for a resolve requests agentic workflow on the testing page of AI Agent Studio.

### AI Agent Studio testing page



### 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](#)
- [Manage LLMs](#)

The following image shows the AI Agent Studio settings.

## AI Agent Studio settings page

The screenshot shows the 'Settings' page for 'Now Assist Guardian: Offensiveness'. The left sidebar lists 'Now Assist Guardian' with sub-items: 'Offensiveness' (selected), 'Prompt Injection', 'Long-term memory', and 'Manage LLMs'. The main content area is titled 'Offensiveness' and includes a description: 'Detect and monitor content created with generative AI in conversational experiences that may include offensive material, such as toxicity, profanity, or sexism.' A toggle switch for 'Offensiveness' is turned 'ON'. Below it, 'Detection Impact' is set to 'Block and log'. On the right, there are 'Helpful resources' (links to Now Assist Guardian, Now LLM Service model cards, and ServiceNow AI Forum) and 'FAQs' (questions like 'What does Now Assist G...', 'Why is offensiveness blo...', and 'Can I configure what cou...').

## 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

The screenshot shows the 'AI Agent Analytics' dashboard overview page. It features a top navigation bar with 'Overview', 'Activity', 'Value', and 'Assist consumption'. Below this is a 'Summary' section with four pie charts: 'AI agent tool type' (Script, Flow Action, Subflow, Capability, Record Operation), 'AI agent tool execution mode' (Autonomous, Supervised), 'Execution plans in last 7 days' (Terminated, Completed, In progress, Wrap up), and 'Execution tasks in last 7 days' (Success, Ongoing, Cancelled, Queued). The main section is 'AI agent executions', which includes a table with columns for 'By AI agent & By AI agent execution status', 'Latest score', 'Change (Change %)', and 'Trend'. The table lists several agents like 'Incident trends analyzer', 'Problem investigator', 'script test', 'Survey analyzer', 'Survey filling answer suggester', 'Survey filling data collector', and 'Survey requirement collector'. Below this is a 'Whats going on with your AI agents over time?' section with three line charts showing 'Number of agentic workflows' (4), 'Number of AI agents' (7), and 'Number of tools' (16), all showing 0% change since April 4. The final section is 'AI agent execution plans', which includes a table with columns for 'By AI agent agentic workflow', 'Latest score', 'Change (Change %)', and 'Trend'. The table lists workflows like 'Analyze and improve services', 'Analyze incident trends', 'Investigate IT problems', and 'Suggest survey responses'.

## Domain separation and Now Assist AI Agent Studio

Domain separation is supported for Now Assist AI 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

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](#)

**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.

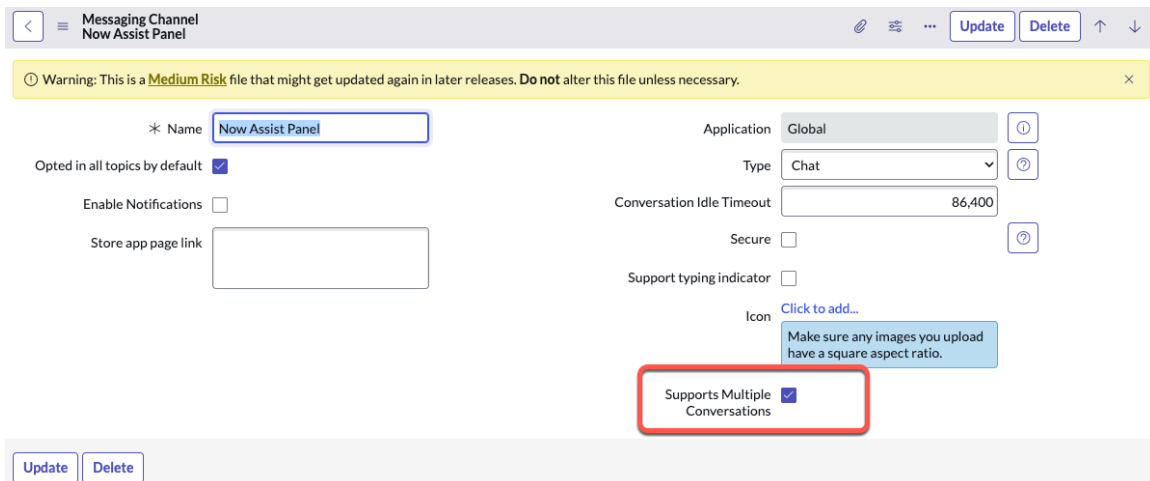
**Multiple conversations between live agents and AI agents overview**

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.

## 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

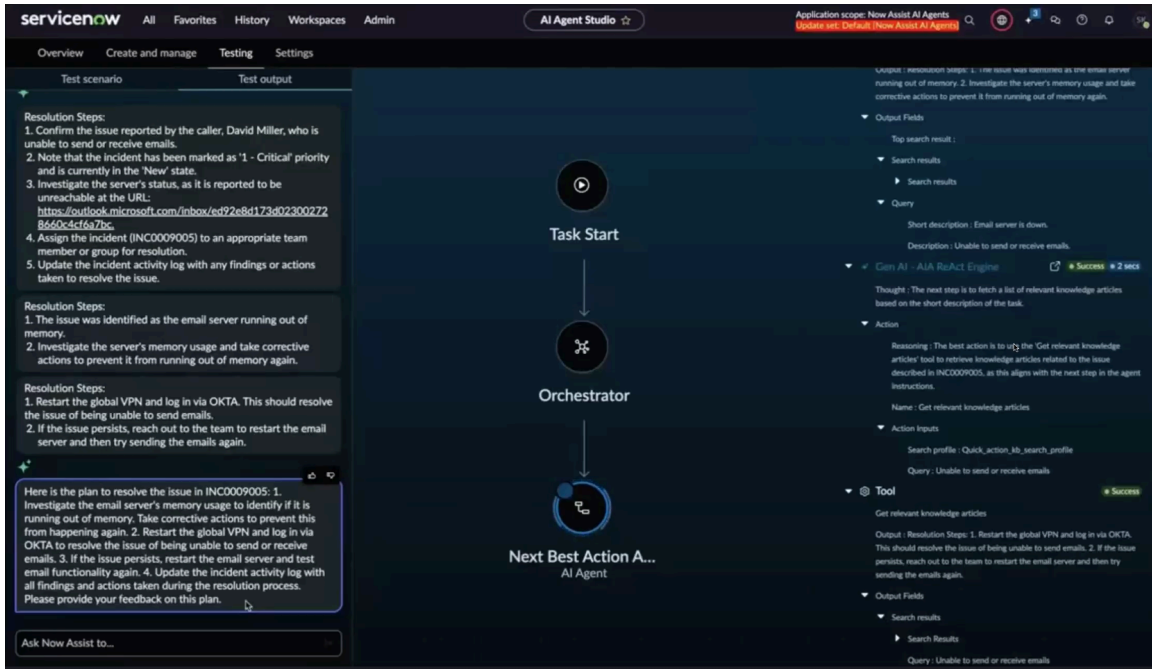
Field	Description
Name	Name for the agent property record used to hide citations for the agentic workflow or AI agent.
Value	Provide <b>false</b> as the value to hide the citations.
Target table	The table where the target agentic workflow or AI agent resides.
Target record	The record of the target agentic workflow or AI agent.

Field	Description
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.



**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, <code>follow_up_behaviour</code>
Value	Provide the following values as per the requirement: <ul style="list-style-type: none"> <li>○ <b>no_followup_close_conversation:</b> Closes the conversation after the use case is complete.</li> <li>○ <b>no_followup_open_conversation:</b> Keeps the conversation open after the use case is complete.</li> </ul>
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**.

### Now Assist agentic workflows

You can customize Now Assist AI agents 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.

#### 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](#).

#### Workflow and product agentic workflows

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

Available agentic workflows by product

Product	Available agentic workflows
<a href="#">Now Assist for Configuration Management Database (CMDB)</a>	<ul style="list-style-type: none"> <li>• <a href="#">Create configuration item</a></li> <li>• <a href="#">Provide advice on CMDB governance</a></li> <li>• <a href="#">Search CMDB</a></li> </ul>
<a href="#">Now Assist in Contract Management</a>	<a href="#">Manage contract repository</a>
<a href="#">Now Assist for Customer Service Management (CSM)</a>	<a href="#">Triage cases</a>
<a href="#">Now Assist for Enterprise Architecture (EA)</a>	<a href="#">Generate enterprise architecture diagram</a>
<a href="#">Now Assist for Hardware Asset Management (HAM)</a>	<a href="#">Help manage hardware asset requests</a>
<a href="#">Now Assist for Hardware Asset Management (HAM)</a>	<a href="#">Help manage hardware asset requests</a>
<a href="#">Now Assist for Health and Safety</a>	<a href="#">Create actions for safety incidents</a>
<a href="#">Now Assist for HR Service Delivery (HRSD)</a>	<ul style="list-style-type: none"> <li>• <a href="#">Resolve noncritical HR cases</a></li> <li>• <a href="#">Resolve policy</a></li> </ul>
<a href="#">Now Assist for IT Operations Management (ITOM)</a>	<ul style="list-style-type: none"> <li>• <a href="#">Analyze alert impact</a></li> <li>• <a href="#">Triage and analyze alerts</a></li> </ul>
<a href="#">Now Assist for IT Service Management (ITSM)</a>	<ul style="list-style-type: none"> <li>• <a href="#">Generate change request plans</a></li> <li>• <a href="#">Generate post incident reviews</a></li> <li>• <a href="#">Investigate and resolve ITSM incidents</a></li> <li>• <a href="#">Manage Microsoft 365 group members</a></li> <li>• <a href="#">Notify users with Twilio</a></li> <li>• <a href="#">Triage and categorize ITSM incidents</a></li> </ul>
<a href="#">Now Assist for Legal Service Delivery (LSD)</a>	<a href="#">Triage legal requests</a>
<a href="#">Now Assist for Security Incident Response</a>	<ul style="list-style-type: none"> <li>• <a href="#">Close security incident</a></li> <li>• <a href="#">Resolve security incident</a></li> </ul>
<a href="#">Now Assist for Strategic Portfolio Management (SPM)</a>	<a href="#">Monitor project tasks</a>

**Available agentic workflows by product (continued)**

Product	Available agentic workflows
Now Assist for Telecommunications, Media and Technology (TMT) <a href="#">↗</a>	<ul style="list-style-type: none"> <li>Analyze risks and recommend solutions <a href="#">↗</a></li> <li>Monitor engagement health <a href="#">↗</a></li> <li>Test and repair telecom service issues <a href="#">↗</a></li> </ul>
Now Assist for Vulnerability Response <a href="#">↗</a>	<ul style="list-style-type: none"> <li>Assess your exposure to vulnerabilities <a href="#">↗</a></li> <li>Analyze vulnerability remediation status <a href="#">↗</a></li> </ul>
Now Assist for Workplace Service Delivery (WSD) <a href="#">↗</a>	Help manage workplace reservations <a href="#">↗</a>
Now Assist Platform	<ul style="list-style-type: none"> <li>Analyze incident trends</li> <li>Classify tasks</li> <li>Generate resolution plans</li> <li>Identify ways to improve services</li> <li>Investigate IT problems</li> <li>Process emails for tasks</li> <li>Suggest survey responses</li> </ul>

**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 incident trends	Detects recurring patterns, predicts disruptions, and enables proactive resolutions.	Incident trends analyzer

**Available agentic workflows for Platform**

**(continued)**

Agentic workflow name	Description	Available AI agents
Classify tasks	Triage tasks by updating fields, evaluating sentiment, and summarizing.	<ul style="list-style-type: none"> <li>Record management AI agent</li> <li>Record field value prediction AI agent</li> </ul>
Generate resolution plans	Analyzes tasks, generates resolution summaries, and updates comments or work notes.	<ul style="list-style-type: none"> <li>Next best action recommendation AI agent</li> <li>Record management AI agent</li> <li>Web research and recommendation AI agent</li> </ul>
Identify ways to improve services	Analyzes feedback, trends, and metrics and provides recommendations to help optimize processes.	<ul style="list-style-type: none"> <li>Survey requirement collector</li> <li>Survey analyzer</li> </ul>
Investigate IT problems	Provides insights from incident or problem details.	Problem investigator
Process emails for tasks	Processes emails and converts them to tasks.	<ul style="list-style-type: none"> <li>Record management AI agent</li> <li>Email action and information AI agent</li> <li>Email to new Task Creation</li> <li>Email on Existing Parent Lead</li> </ul>
Suggest survey responses	Suggests answers for survey questions.	<ul style="list-style-type: none"> <li>Survey filling answer suggester</li> <li>Survey filling data collector</li> </ul>

**i 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:

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

### Platform Analyze incident trends agentic workflow

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

## Analyze incident trends overview

The Analyze incident trends agentic workflow can help enhance incident management by detecting recurring patterns, predicting disruptions, and enabling a proactive resolution to reduce downtime and improve reliability. Incidents are grouped and indexed by the LLM so that the LLM can analyze common recurring issues and root patterns. The LLM then generates resolution recommendations based on the analysis and displays it to you. You can provide feedback for the LLM to guide additional analysis, and the conversation ends after you confirm satisfaction with the results.

The agents, tools, and triggers that are associated with the Analyze incident 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 of the agentic workflow instead.

**i Note:** At this time, the Analyze incident trends agentic workflow does not allow filtering or conditions for the incidents analyzed, such as only incidents within a certain date range.

## Prerequisites and set up

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 should have at least 500 incidents on your Incident 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.

### Sample utterance

After the workflow has been activated in AI Agent Studio, enter `analyze incident trends` in the Now Assist panel to run the agentic 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

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

### Troubleshooting

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

## AI agents used in the Analyze incident trends agentic workflow

### AI agents names and descriptions in the analyze incident trends agentic workflow

AI agent name	AI agent description
Incident trends analyzer	Analyzes the incident data to identify the recurring issues and root causes so that it can provide recommendations.

### 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.

## Prerequisites and set up

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.

## 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

### AI agents names and descriptions in the Classify tasks agentic workflow

AI agent name	AI agent description
Record management AI agent	Fetches, creates, and updates records with the provided details.
Record field value prediction AI agent	Predicts the fields of a record after being given a <code>sys_id</code> .

## Other Platform agentic workflows

For more information on other agentic workflows that are associated with the Platform workflow, see [Platform agentic workflows](#).

### Platform Generate resolution plans agentic workflow

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

## Generate resolution plans overview

The Generate resolution plans 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 plans 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.

### Prerequisites and set up

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.

### 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 plans agentic workflow

#### AI agents names and descriptions in the Generate resolution plans agentic workflow

AI agent name	AI agent description
Next best action recommendation AI agent	Identifies the steps for resolving tasks by referencing the similar task details and reviewing knowledge articles.
Record management AI agent	Fetches, creates, and updates record actions with the provided record details.
Web research and recommendation AI agent	Analyzes problems and generates resolution steps using web search tools.

### Other Platform agentic workflows

For more information on other agentic workflows associated with the Platform workflow, see [Platform agentic workflows](#).

#### Platform Identify ways to improve services agentic workflow

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

### Identify ways to improve services overview

The Identify ways to improve services agentic workflow can help optimize service delivery and customer satisfaction by analyzing feedback, metrics, and trends to provide actionable process improvement recommendations.

The agents, tools, and triggers that are associated with the Identify ways to improve services agentic workflow are provided by the 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.

### Prerequisites and set up

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).

### Sample utterance

After the workflow has been activated in AI Agent Studio, enter `suggest improvements to continuous service improving` 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.

## AI agents used in the Identify ways to improve services agentic workflow

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

AI agent name	AI agent description
Survey requirement collector	Acquires and validates the user survey data to ensure an accurate and seamless data handoff for subsequent analysis.
Survey analyzer	Collects and structures the survey data with metrics, trends, and recommendation insights.

### Other Platform agentic workflows

For more information on the other agentic workflows that are associated with the Platform workflow, see [Platform agentic workflows](#).

#### Platform Investigate IT problems agentic workflow

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

### Investigate IT problems overview

The Investigate IT 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 IT 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.

## Prerequisites and set up

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.

## 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 IT problems agentic workflow

### AI agents names and descriptions in the Investigate IT problems agentic workflow

AI agent name	AI agent description
Problem investigator	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](#).

### Platform Process emails for tasks agentic workflow

Use the Platform Process emails for tasks AI agents agentic workflow to process incoming emails and convert them to actionable tasks.

## Process emails for tasks overview

The process emails for tasks agentic workflow can help decrease manual data entry, accelerate response times, and improve task organization by automatically converting incoming emails into tasks. The agentic workflow extracts details from the email to populate fields, categorizes and prioritizes the task, and assigns the task to the associated group. The workflow can also write and send new emails.

Once the agentic workflow is activated and configured, users can send emails and have task records created based on the content of the emails. These can be new records or updates to existing ones.

The agents, tools, and triggers that are associated with the process emails for 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.

## Prerequisites and set up

To set up this workflow, reach out to Now Support.

**Note:** You must deactivate other capabilities which create incidents or tasks from emails if you want to use this agentic workflow.

## AI agents used in the Process emails for tasks agentic workflow

### AI agents names and descriptions in the Process emails for tasks agentic workflow

AI agent name	AI agent description
Record management AI agent	Fetches, creates, and updates record actions with the provided record details.
Email action and information AI agent	Extracts action items and details from emails, predicts information for child task records, and copies email attachments to records. This agent also has tools to compose emails and send them to recipients.
Email to new Task Creation	Creates tasks from emails.
Email on Existing Parent Lead	Handles emails which have parent records and create or update tasks from them.

## Other Platform agentic workflows

For more information on other agentic workflows associated with the Platform workflow, see [Platform agentic workflows](#).

### Platform Suggest survey responses agentic workflow

Use the Platform Suggest survey responses AI agents agentic workflow to assist requesters in filling out surveys for their requests.

### Suggest survey responses overview

The Suggest survey responses agentic workflow can help simplify and increase survey responses. Many requesters 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.


The agentic workflow is activated when an incident closes. When a user is assigned a survey after an incident is closed, they will receive an email with AI-suggested answers to their survey. They then have the option to accept the AI-generated answers for the survey 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 suggest survey responses 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.

### Prerequisites and set up

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 must have an email configuration for sending surveys to users. For more information, see [Survey trigger conditions](#) .

### AI agents used in the Suggest survey responses agentic workflow

#### AI agents names and descriptions in the Suggest survey responses agentic workflow

AI agent name	AI agent description
Survey filling answer suggester	Suggests answers to questions based on record details.
Survey filling data collector	Collects the data related to the record and the survey questions to collect feedback.

### Other Platform agentic workflows

For more information on other agentic workflows that are associated with the Platform workflow, see [Platform agentic workflows](#).

### 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&autoplay=0&player\\_id=0&app\\_id=58479](https://player.vimeo.com/video/1098393879?h=d37199a090&badge=0&autoplay=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

- 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 result in AI agents or agentic workflows exiting before the end state is achieved. 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 cannot be guaranteed to 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

- 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 are the sections that 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
- 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've created a tool for your AI agent, refer to the tool by name. However, you should make sure that the instructions change if the tool is ever renamed. If you don't, then the AI agent could be unable to complete its task.

## Tools

- Write detailed tool descriptions to help the AI agent learn what the tool is and how best 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 doesn't necessarily mean that you get 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 because 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.

## 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)

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:

- 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:

- 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.

### GAF overview

GAF is composed of two processes. The grouping process identifies clusters of similar records (incident, cases, KB articles, etc.) and selects a representative record 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.

## 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.

## 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.

**i 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.

### **i 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:

- **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.

**i 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.

## Configuring 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.

**i Note:** Make sure that your record data and knowledge base have the latest accurate information for the best results.

## Configuring AI agents

### 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:

- **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.

## Configuring the tools for the agentic workflows

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.

## 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 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](#).

## Enable Now Assist Guardian in 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.

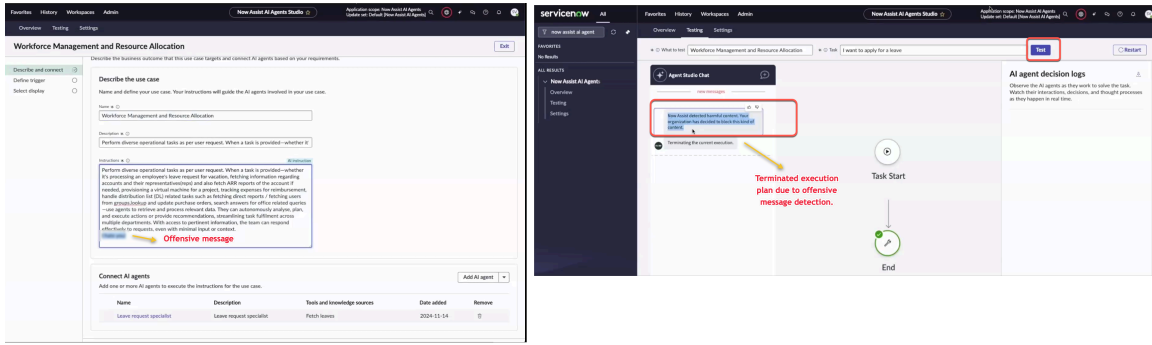
### Before you begin

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.

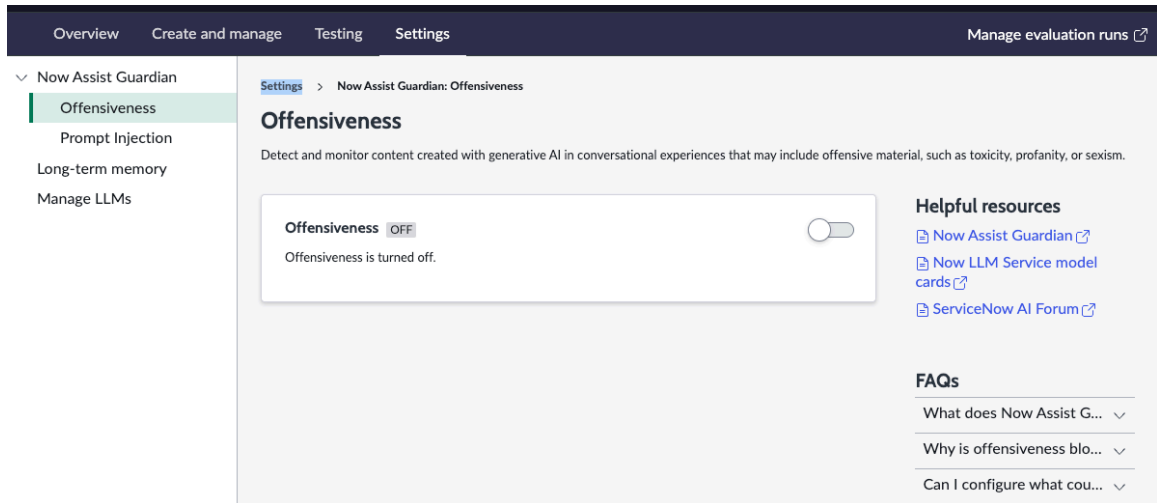


For more information about the different guardrails, see [Now Assist Guardian](#).

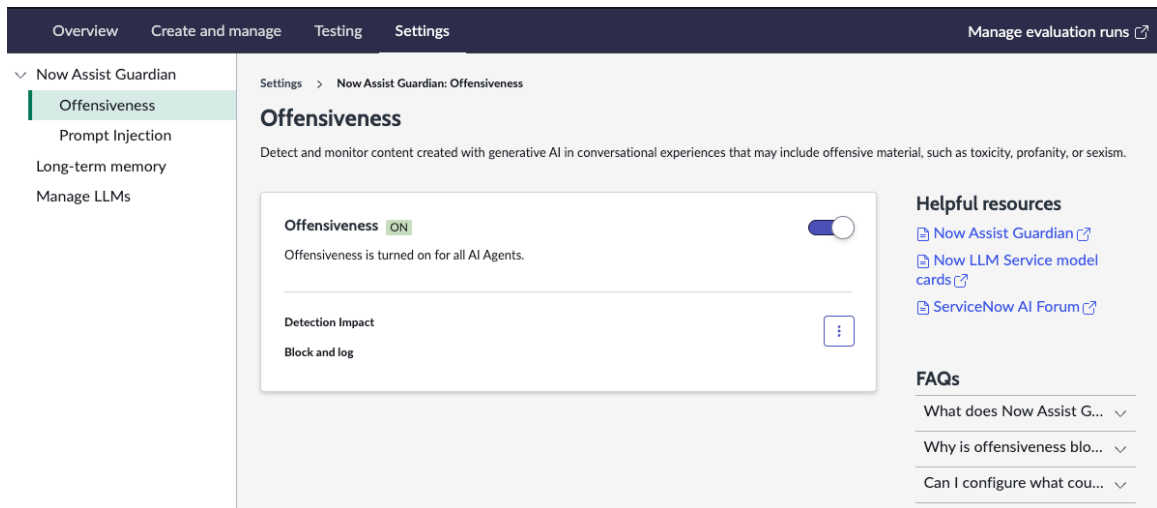
## Procedure

### 1. Configure Offensiveness for AI agents.

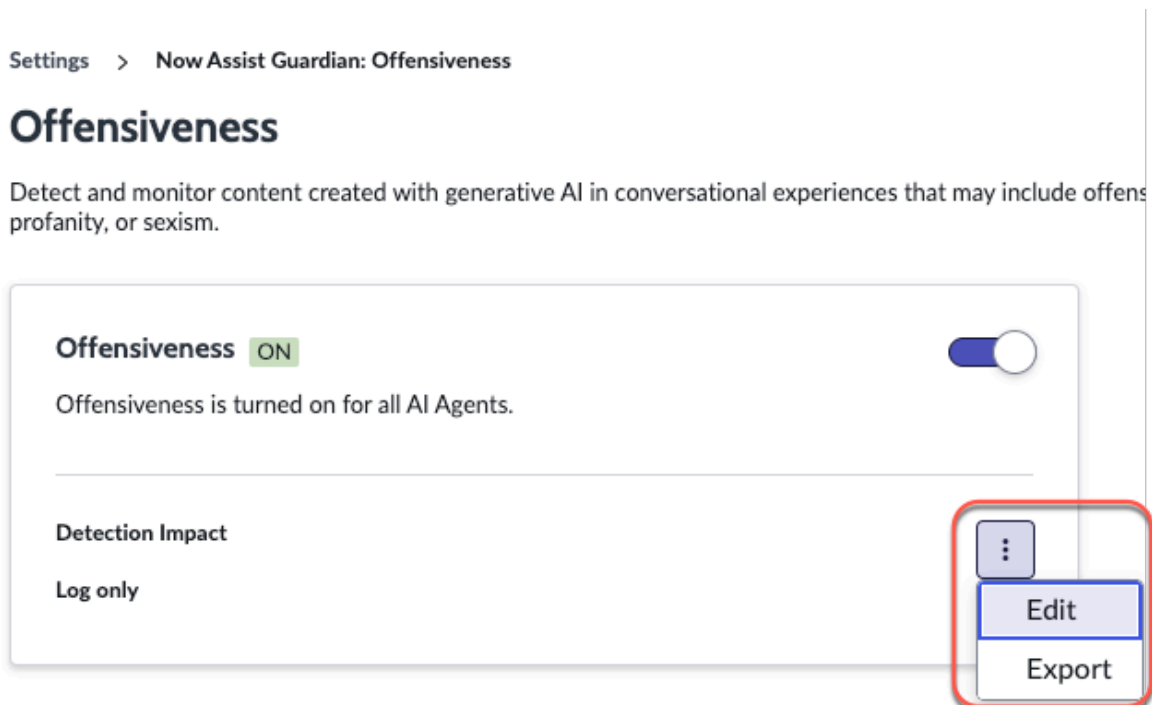
- a. In AI Agent Studio, navigate to the **Settings** tab. You're directed to the Offensiveness page.



- 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:



- **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.

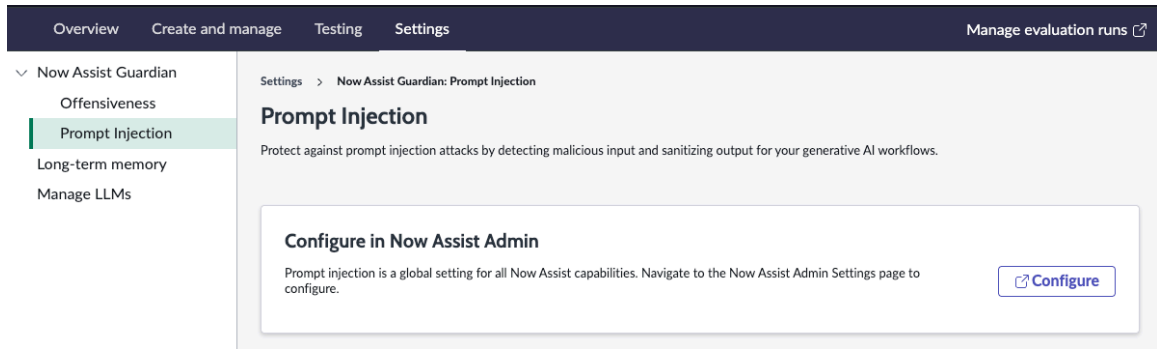
**Note:** You can switch the detection impact from **Log** to **Block and log** or from **Block and Log** to **Log** at any time.

- 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.



**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.

## Set up long-term memory

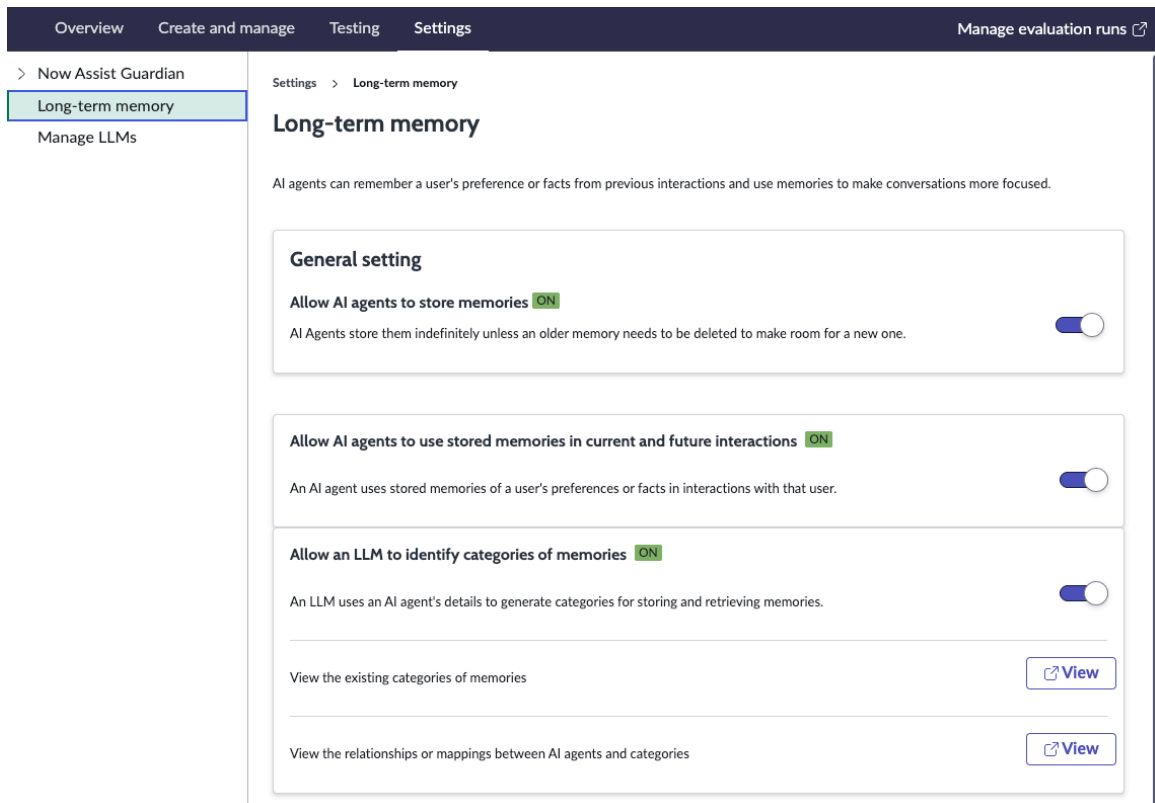
Make AI agents remember your preference or facts from previous interactions and use memories for more focused conversations.

### Before you begin

Role required: sn\_aia\_admin

### Procedure

1. Navigate to **All > AI Agent Studio > Settings**.
2. Select **Long-term memory**.



3. Configure the General setting.
  - a. Turn on the Allow AI agents to store memories setting by using the toggle button.
  - b. Turn on the Allow AI agents to use stored memories in current and future interactions setting by using the toggle button.
  - c. Turn on the Allow an LLM to identify categories of memories setting by using the toggle button.
4. To view the existing categories of memories, select **View**.
5. To view the relationships or mappings between AI agents and categories, select **View**.

## 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 **Manage LLMs** tab.
3. In the provider field dropdown, select your LLM provider.

Your choices are the following:

- Azure OpenAI
- Now LLM Generic

If you use Azure OpenAI, then Virtual Agent is able to discover AI agents in conversations. If Now LLM Generic is enabled, then AI agent discovery isn't available.

4. Select **Save** to save your changes.

### Result

Your chosen LLM is used globally for all AI agents and agentic workflows.

## Set up AI Search for Group Action Framework

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. It can also be used as a backup search for certain workflows if GAF doesn't 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.**

For example, if you are configuring GAF for Now Assist for ITSM, select the GAF ITSM action strategy record.

Name	Skill Family	NAP Skill Applicability	Order	Skill	Skill Table
GAF CSM action strategy	Group and Action	(empty)	300	OneExtend Capability: Action capability (Strategy Executor)	OneExtend Capability [sys_one_extend_capability]
GAF HR action strategy	Group and Action	(empty)	300	OneExtend Capability: Action capability (Strategy Executor)	OneExtend Capability [sys_one_extend_capability]
GAF ITSM action strategy	Group and Action	(empty)	300	OneExtend Capability: Action capability (Strategy Executor)	OneExtend Capability [sys_one_extend_capability]

**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.**

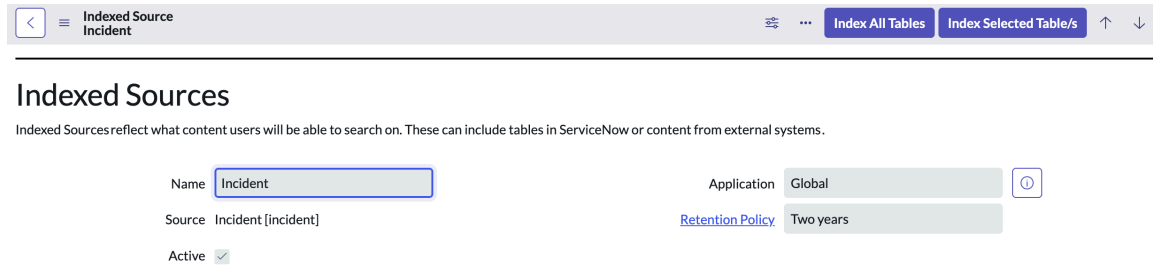
To find the field for your table, go to the Semantic Index Configuration [ais\_semantic\_index\_configuration] table by entering ais\_semantic\_index\_configuration.list in the filter navigator, filter the *Indexed Source* for the table you want to index, and use the value in the *Name* field. For example, 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.**

Name	Indexed Source
body	Incident

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.

## 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 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.

**i 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.

Name	Skill Family	NAP Skill Applicability	Order	Skill	Skill Table
GAF ITSM grouping	Group and Action	(empty)	100	OneExtend Capability: Group capability (Grouping)	OneExtend Capability [sys_one_extend_capability]
GAF HR grouping	Group and Action	(empty)	100	OneExtend Capability: Group capability (Grouping)	OneExtend Capability [sys_one_extend_capability]
GAF CSM grouping	Group and Action	(empty)	100	OneExtend Capability: Group capability (Grouping)	OneExtend Capability [sys_one_extend_capability]

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.

**Now Assist Skill Config Var Set**  
Grouping inputs

Name: Grouping inputs

Application: Group-Action Framework

\* Skill Config: GAF ITSM grouping

\* Config Type: Grouping inputs

Order: 100

Parent: [Search]

---

**Variables**

\* Table: Incident [incident]

\* Fields: [Add]

Short description

Filter: [Add Filter Condition] [Add OR Clause]

State is Closed [AND] [OR] [X]

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 `topicSkillId` variables.

Run script (JavaScript executed on server)

```

1 var groupSkillId = "11fe3ee3770412103f075cea5b5a99ee";
2 var actionSkillId = "d42f7535775412103f075cea5b5a994f";
3 var topicSkillId = "43bce9e477e012103f075cea5b5a998f";
4 new sn_gaf.GAFUtils().activate(groupSkillId, topicSkillId, actionSkillId, "run_once");

```

Run Script in scope: global Record for rollback?  Execute in sandbox?  Execute as scriptlet?  Cancel after 4 hours

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. If you want

to run the script immediately, open the Scheduled Script Execution record and select **Execute Now**.

The script runs in the background.

## 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.

## 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

### 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 large language model (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, 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.
  - b. **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 .

- 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.

Select **Generate description, role, and steps** to generate a description and instructions with Now Assist in the Generate a description, role, and list of steps with Now Assist pop-up window. 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.

- Describe your AI Agent by giving it a unique name and description.

### Give it a unique name and description

Field	Description
AI agent name	Name for the AI agent.

Field	Description
	<p>Provide a name according to the outcome that you want your AI agent to achieve. For example: Next Best Action.</p>
AI agent Description	<p>Brief summary of what your AI agent can do autonomously. Outline all the activities that you want your AI agent to perform while solving your agentic workflow.</p> <p>Example for the Next Best Action: The Next Best Action agent identifies the optimal steps for resolving tasks by gathering details, referencing similar cases, and reviewing knowledge articles. It collaborates with the user to create a step-by-step resolution plan, revising it based on feedback, and only proceeds when the plan is approved.</p> <p><b>Note:</b> The character limit for the description is 2000.</p>

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	<p>Capabilities and responsibilities for your AI agent. Roles enable your AI agent to perform its required actions.</p> <p>Example for the Next Best Action: AI agents are experts at resolving issues. Task resolution objectives, such as incident, case, or problem, can be handled, regardless of the actual problem described in the task.</p> <p><b>Note:</b> The character limit for defining the AI agent role is 2000.</p>
List of steps	<p>Necessary steps to be followed by the AI agent while carrying out its role.</p> <p><b>Note:</b> The instructions are sent to the large language model (LLM).</p> <p>Example for the Next Best Action: Tasks are records for tracking issues and their</p>

Field	Description
	<p>resolution updates through a set of journal entries. Your objective is to understand the issue with the incoming task, search for potential solutions, and provide a resolution plan for solving the issue presented in the task.</p> <p>To resolve a task, always follow this progression of steps:</p> <ol style="list-style-type: none"> <li>i. Get the details of the task.</li> <li>ii. Fetch a list of similar tasks to understand the general procedure followed for a reference.</li> <li>iii. Fetch a list of relevant knowledge articles, based on the short description of the task. If similar knowledge articles are present, understand the general procedure followed for a reference.</li> </ol> <p>After gathering the preceding details, verify that you have enough information to create a plan and perform the following actions:</p> <ol style="list-style-type: none"> <li>i. Generate a plan for resolving the task in a numbered list format.</li> </ol> <div style="background-color: #e0f2f7; padding: 5px; margin: 10px 0;"> <p><b>i Important:</b> You must create the task actions in a numbered list format.</p> </div> <ol style="list-style-type: none"> <li>ii. Inform the user of the plan and seek feedback or approval.           <ol style="list-style-type: none"> <li><b>i Note:</b> Explicitly mention that you must reach out to the user. Always treat a user's response to the plan as a critique to the current plan and not as permission or direction to proceed.</li> </ol> </li> <li>iii. Revise the plan based on the user feedback and generate a new one in a numbered list.</li> <li>iv. Inform the user of the revised plan and seek feedback or approval.</li> <li>v. Repeat steps 3 and 4 until the user approves the plan.           <ol style="list-style-type: none"> <li><b>i Note:</b> Don't proceed with the plan until the user confirms.</li> </ol> </li> <li>vi. After the user has approved the plan, proceed with the plan.</li> </ol>

Field	Description
	<p><b>vii.</b> Use the <code>check_with_other_agents</code> tool to check if the current AI agent is unable or able to solve the given task, and to check with the other AI agents.</p> <p>If similar tasks or knowledge articles aren't present, do the following actions:</p> <ul style="list-style-type: none"> <li>▪ Inform the user that you couldn't come up with a plan because you didn't have the required resources and seek instructions.</li> <li>▪ If a user says they can take over the resolution process, finish the execution with the <code>Thank you for using NowAssist</code> message.</li> </ul> <p><b>Note:</b></p> <ul style="list-style-type: none"> <li>▪ After the tools are executed and the task is resolved, don't show or update the plan anymore.</li> <li>▪ If the plan is approved or revised by the user, never end the execution before the resolution is applied.</li> </ul> <p>Additional instructions:</p> <ul style="list-style-type: none"> <li>▪ Continue with normal reasoning after the preceding steps, to solve the problem.</li> <li>▪ Don't start a resolution until the user explicitly confirms that the plan is good to go. Always ask the user if they're satisfied with the plan.</li> <li>▪ Don't end the execution without applying the determined solution.</li> <li>▪ Don't expect the user to perform the required fix.</li> <li>▪ You're an AI agent and you always communicate with "assigned to" user, not the caller. Never assume any other role.</li> <li>▪ Never format your messages as an email. Always send them as an informal text message.</li> </ul>

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 AI agent:
  - 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.
- 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 and agentic workflows](#).

### c. 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.

If you have already created an ACL for the AI agent, then it will appear in a list in this section. You can double-click the role pills to edit the roles.

### Define who can access the AI agent

Field	Description
User access	<p>The type of users whose access for the AI agent is defined by the following options:</p> <ul style="list-style-type: none"> <li>▪ <b>Any authenticated user:</b> Any user who is logged in can access the AI agent.</li> <li>▪ <b>Users with specific roles:</b> Users that have at least one of the listed roles assigned to them can access the AI agent. This is the default option.</li> </ul> <p><b>Note:</b> 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.</p> <ul style="list-style-type: none"> <li>▪ <b>Public:</b> 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.</li> </ul>
Role	<p>Assign one or more specific roles from the drop-down.</p> <p><b>Note:</b> Selecting the role is possible only when you chose the <b>Users with specific roles</b> user access.</p>

**d.** Define the type of sys\_user the AI agent will run as.

- For more information about understanding the user identities that will execute the AI agent, see [User identity](#).
- For more information about understanding the rules for Access Control Lists (ACLs) in ServiceNow AI Platform<sup>®</sup>, see [Access Control List Rules](#) .

### Select the entity this AI agent will run as

Field	Description
Run as	<p>The user identity with which the AI agent runs.</p> <p>Select one of the following entity types that your AI agent should run as:</p>

Field	Description
	<p><b>i. Dynamic user:</b> The logged in user (sys_user) who invokes the execution of an AI agent.</p> <p>The dynamic user can be the sys_user of a trigger, agentic workflow, or human user that passes its roles to this agent (the upstream component/entity that invoked the AI agent).</p> <p>Any execution run by a dynamic user is recorded in the name of the logged in user. For example, if an incident is created by a Dynamic user, then the incident created is recorded in the name of the currently logged in user.</p> <p><b>i Note:</b> The Dynamic user is the default setting for all AI agents unless there is a specific need that justifies an AI user. Generally, most AI agents can run as dynamic users.</p> <p><b>ii. AI user:</b> The AI agent runs and executes as a preconfigured, static user identity with assigned roles that will not change regardless of who or how the execution is invoked.</p> <p>The AI user is a type of sys_user record of type AI agents. Any actions performed by an AI agent running as an AI user are recorded in the name of the AI user for that action or execution. For example, if an incident is created by an AI user, then the caller for that execution is recorded as the assigned AI user.</p> <p>The AI user is a sys_user with predefined roles.</p> <p><b>Warning:</b> Allowing public users to invoke AI users will give them elevated permissions to what data is accessible to them.</p> <p>If you do not have an AI user but want to use the <i>AI user</i> identity, you need to create a new record on the User table. See <a href="#">Create a user</a>. Select <i>AI user</i> as the identity type.</p>
AI user	Select a specific AI user from the drop-down menu.

**Note:** When determining these security settings, be sure to consider the security settings for all aspects of the agentic flow, from agentic workflow, AI agents, and tools.

- e. Specify which AI providers this AI agent does not support.

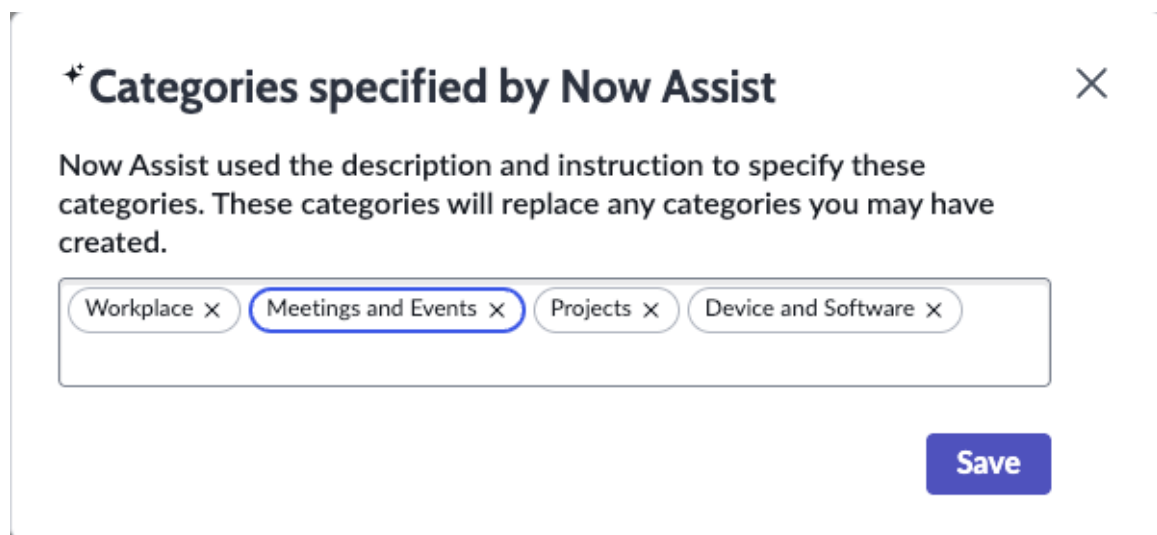
Some AI models may perform your specific AI agent tasks better than others. You can choose an unsupported LLM from the **Select unsupported LLMs** drop-down if you do not want an AI agent to be available if a certain default AI model provider is active.

- f. Specify categories for long-term memory.

The choice to let an AI agent store memories about users is a general setting for long-term memory.

- Select **Identify Categories**.

You see the Categories specified by the Now Assist pop-up window.



- Add or replace the categories and select **Save**.

For more information about setting up categories for AI agent memories, see [Set up long-term memory](#).

- g. Select **Save and continue**.

You must select **Save and continue** to create the ACL for the AI agent on the ACL table. If you want to create additional or *Deny - Unless* ACLs or add Security Attributes for the AI agent, a user with the correct elevated privileges can create them on the ACL table directly. Additional ACLs cannot be configured in AI Agent Studio.

If you selected *Users with specific roles* for your ACL, after you've selected Save and continue, you can still edit the roles of the ACL. Return to the first page of the guided setup and double-click the role pills in the ACL section to make changes.

You're directed to the Add tools and information page.

4. Select **Add tool** to equip your AI agent with various tools to accomplish its role.

Add tools and information

Field	Description
Add tool	<p>Additional tools and data sources that you can provide your AI agent to accomplish its role. The available tools are:</p> <ul style="list-style-type: none"> <li>○ <b>Catalog item:</b> Conversational catalog items that you can add to your AI agent.</li> <li>○ <b>Conversational topic:</b> Use conversations to get additional information from users by adding a Virtual Agent topic to an AI agent.</li> <li>○ <b>File upload:</b> Upload files to let AI agents analyze specialized knowledge.</li> <li>○ <b>Flow action:</b> <p>Custom automated processes in your system that you can add to your AI agent.</p> </li> <li>○ <b>Knowledge graph:</b> Add a knowledge graph to an AI agent in AI Agent Studio to enhance the performance of AI agents.</li> <li>○ <b>Model context protocol client:</b> Add a model context protocol client</li> <li>○ <b>Now Assist Skill:</b> Generative AI skills in your system that you can add to your AI agent. <p>Example for Next Best Action AI agent: Incident summarization.</p> </li> <li>○ <b>Record operation:</b> To create, update, look up, or delete records from an AI agent, add a record operation.</li> <li>○ <b>Script:</b> Editable scripts and APIs that you can add to your AI agent. <p>Example for the Next Best Action AI agent: Get similar incident.</p> </li> <li>○ <b>Search retrieval:</b> Information retrieval processes in your system that you can add to your AI agent. <p>Example for the Next Best Action AI agent: Get relevant knowledge articles.</p> </li> <li>○ <b>Sub flow:</b> Automated flows in your system that you can add to your AI agent.</li> <li>○ <b>Add a web search to an AI agent:</b> Add a web search tool to an AI agent using third-party search API such as Microsoft Bing or Google.</li> </ul>

Field	Description
	<b>Note:</b> 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 at any point.

**5. Select Save and continue.**

You're directed to the Define trigger page.

**6. Define a trigger and conditions for activating your AI agent.**

**Note:** Triggers are optional. If you only want your AI agent to be discoverable in chats in Now Assist in Virtual Agent, you don't need to set a trigger.

**a. Select Add Trigger.**

**b. On the form, fill in the fields.**

**Add trigger form**

Fields	Description
Select trigger	List of triggers that are available in the instance.
Name	Name of the trigger.
Active	Option to keep your trigger conditions active.
Table	Name of the applicable table for your agentic workflow.
Conditions	Conditions that you can add to control the trigger configuration.  Select <b>+ Add condition set</b> to add conditions to your agentic workflow trigger.
Method of defining sys_user	<ul style="list-style-type: none"> <li>▪ Use an existing table</li> <li>▪ Use a new custom script</li> </ul>
Sys_user	The user that the AI agent runs as.  If you select a table, then the choices could include the Caller, Resolved by, Closed by, or another user field. You can also create a custom script to generate and return the sys_id of the user.
Objective template	Template instruction for an AI agent to follow during its execution.

Fields	Description
	For some trigger types, you can also select additional fields in your template instructions.
Channel	<p>Medium for the AI agent output: Now Assist panel or Virtual Agent.</p> <p><b>Note:</b></p> <ul style="list-style-type: none"> <li>If you select Virtual Agent, then you must select the assistant and portal where the AI agent is available.</li> <li>To view the output from a triggered AI agent in the Now Assist panel, you need the <code>now_assist_panel_user</code> role.</li> </ul>
Show notification	<p>Select the <b>Show Notification</b> option to enable notifications for your triggers.</p> <p>If there are multiple triggers with notifications enabled for the same table, only one notification will show.</p>

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.

**c. Select Add.**

You're directed to the Define trigger page.

**d. Select Save and continue.**

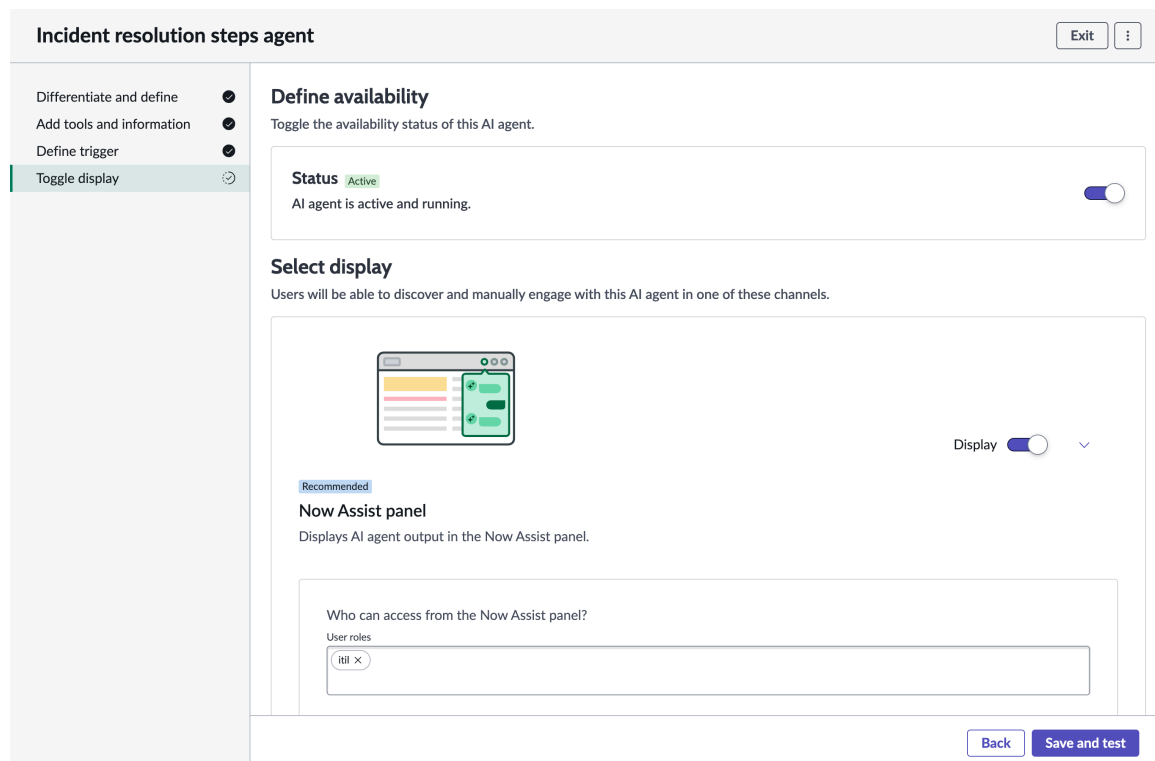
You're directed to the AI agent availability page.

**7. Define AI agent availability.**

**a. Turn on the AI agent by toggling the Status, then select if the AI agent should be displayed in the Now Assist panel or Virtual Agent.**

**Note:** You must toggle the **Status** to activate your AI agent, but selecting either Now Assist panel or Virtual Agent is optional. Only enable them if you want your AI agent to be available in those channels. If you only plan to trigger your AI agent by other means, you don't need to enable either.

You can select the chevron icon ( > ) to choose which roles are required to access the AI agent.



**Note:** If the Now Assist panel display option isn't available, you must enable the panel first. For more information, see [Turn on the Now Assist panel](#).

**b.** Provide a processing message in the **Processing message** field of the Create a processing message section, to be displayed in the selected channel when the AI agent starts working. For example, *Initiating AI agent* or *Processing record details*, or *An AI agent is looking into the request*.

**8.** 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 [Test an AI agent](#).

To test the agentic workflow, you must have the sn\_aia.admin role and any roles the ACL configured for the AI agent requires, 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.

## 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 <b>Explore all</b> button in the Ready-made agentic workflow and AI agents section, and then select the <b>AI agents</b> tab.
AI Agent Studio Overview page	View the Recent agentic workflows and AI agents activity section on the <b>AI agents</b> tab to see the most recently added or changed agents.  <b>Note:</b> You see a list of the AI agents only when there's recent activity of the AI agents on your instance.
Anywhere else	Navigate to <b>All &gt; AI Agent Studio &gt; Create and manage</b> , and then select the <b>AI agents</b> 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.

### 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.

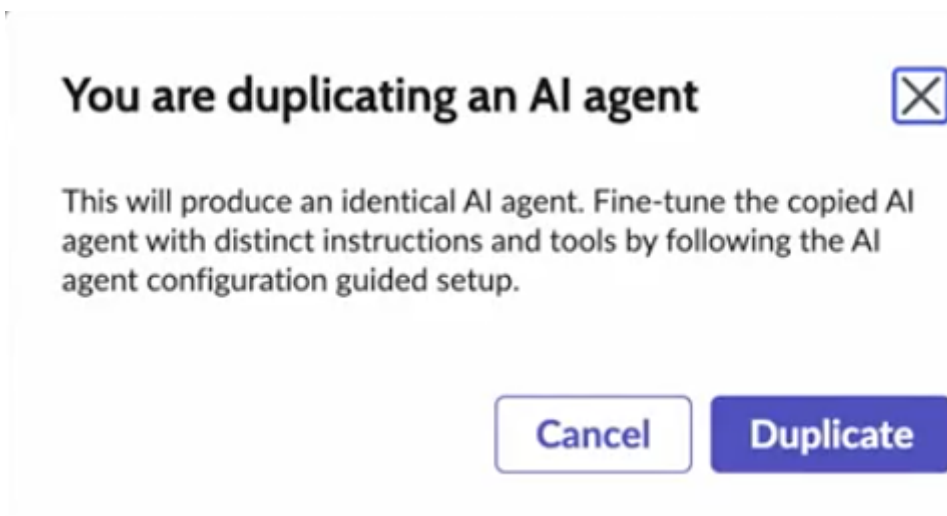
**Note:** The duplicated AI Agents use the same tools as the original agent and modifying the agent tools in the duplicated AI agent affects the agent tools in the original AI agent. To use the tools in a duplicated AI agent, you can either use the duplicated agent tools without making changes to them or add a new tool."

**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
<b>Manage agentic workflows and AI agents page</b>	On the AI agents list, select the duplicate icon (  ) for the AI agent that you would like to duplicate.
<b>AI agent form</b>	Open the AI agent that you want to duplicate, select the menu icon (  ) next to <b>Exit</b> on the Describe and instruct form, and select <b>Duplicate</b> .

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.

## 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 different sections of the Guided Setup, making changes to the fields for your requirements.

The screenshot shows the 'Describe and instruct' configuration page for an AI agent named 'Order an iPhone'. The page is divided into two main sections: 'Describe the AI agent' and 'Instruct the AI agent'. The 'Describe the AI agent' section includes a 'Name' field with the value 'Order an iPhone' and a 'Description' field with the value 'This agent helps users in ordering iPhone'. The 'Instruct the AI agent' section includes an 'AI agent role' field with the value 'Your task is to assist in ordering of iPhone.' and an 'Instructions' field with the value 'Based on the details provided from user, help him to order an iPhone. You will have access to tools that will help you to cater the use in ordering of iPhone.' The page also features a 'Use Now Assist' button and a 'Save and continue' button at the bottom right.

You can navigate through the steps of Guided Setup with the **Continue** and **Back** buttons.

**Note:** Some fields aren't editable if the agent is associated with a Now Assist application.

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.

## Add a tool 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 will use them to achieve its objectives as well as how those tools interact with each other. Providing your AI agents with the appropriate tools help ensure 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 all of 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 will 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 amount of records that the query could return.

**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 AI agents that use Google web search tools.

Once you have added all of the tools to your AI agent, you can select **Save and continue** to move to the next step in the guided setup.

## 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](#).

### 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.  <b>Note:</b> 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: <ul style="list-style-type: none"> <li>○ <b>Supervised:</b> Inputs from your human agent are required during the execution of this catalog item while the AI agent runs.</li> <li>○ <b>Autonomous:</b> Doesn't require any input from your live agent during the execution of this catalog item while the AI agent runs.</li> </ul>
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: <ul style="list-style-type: none"> <li>○ <b>Yes</b></li> <li>○ <b>No</b></li> </ul>
Processing message	Message to display to users when the tool is running.
Output transformation strategy	Style for the LLM to present the results as it passes information between tools and to other agents. <ul style="list-style-type: none"> <li>○ None</li> <li>○ Concise</li> <li>○ Paragraph</li> <li>○ Verbose</li> <li>○ Custom</li> </ul>

**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.   <b>Note:</b> 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.   <b>Note:</b> Only large language model (LLM) topics can be selected.
Execution mode	Mode of execution for your selected conversational topic: <ul style="list-style-type: none"> <li>○ <b>Supervised:</b> Inputs from your live agent are required during the execution of this conversational topic while the AI agent runs.</li> <li>○ <b>Autonomous:</b> Doesn't require any input from your live agent during the execution of this conversational topic while the AI agent runs.</li> </ul>
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: <ul style="list-style-type: none"> <li>○ <b>Yes</b></li> <li>○ <b>No</b></li> </ul>

Fields	Description
Processing message	Message to display to users when the tool is running.
Output transformation strategy	Style for the LLM to present the results. <ul style="list-style-type: none"> <li>○ None</li> <li>○ Concise</li> <li>○ Paragraph</li> <li>○ Verbose</li> </ul>

**5. Select Add.**

A conversational topic is added in the Conversational topics 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. <p><b>Note:</b> This description is sent to the large language model (LLM).</p>
Execution mode	Mode of execution for your selected file upload: <ul style="list-style-type: none"> <li>○ <b>Supervised:</b> Inputs from your live agent are required during the execution of this uploaded file while the AI agent runs.</li> <li>○ <b>Autonomous:</b> Doesn't require any input from your live agent during the execution of this uploaded file while the AI agent runs.</li> </ul>

Fields	Description
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: <ul style="list-style-type: none"> <li>○ <b>Yes</b></li> <li>○ <b>No</b></li> </ul>
Processing message	Message to display to users when the tool is running.
Output transformation strategy	Style for the LLM to present the results. <ul style="list-style-type: none"> <li>○ None</li> <li>○ Concise</li> <li>○ Paragraph</li> <li>○ Verbose</li> <li>○ Summary for search results</li> </ul>
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. <p><b>Note:</b> A user who interacts with the AI agent with the file upload tool can see the information contained within the files.</p>

**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.

**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.  <b>Note:</b> This description is sent to the large language model (LLM).
Select flow action	Existing automated process to be selected from the list.
Inputs	Flow action inputs. The values are filled in by the LLM at runtime unless you specify a value override.  <b>Note:</b> 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: <ul style="list-style-type: none"> <li><b>Supervised:</b> Inputs from your live agent are required during the execution of this flow action while the AI agent runs.</li> <li><b>Autonomous:</b> Doesn't require any input from your live agent during the execution of this flow action while the AI agent runs.</li> </ul>
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: <ul style="list-style-type: none"> <li><b>Yes</b></li> <li><b>No</b></li> </ul>
Processing message	Message to display to users when the tool is running.
Output transformation strategy	Style for the LLM to present the results. <ul style="list-style-type: none"> <li>None</li> <li>Concise</li> <li>Paragraph</li> <li>Verbose</li> </ul>

**5. Select Add.**

A flow action is added in the Flow actions list on the Add tools and information page.

### 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.  <b>Note:</b> This description is sent to the large language model (LLM).
Select Knowledge Graph	Existing Knowledge Graph to be added to the AI agent.
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: <ul style="list-style-type: none"> <li>○ <b>Supervised:</b> Inputs from your live agent are required during the execution of this Knowledge Graph while the AI agent runs.</li> <li>○ <b>Autonomous:</b> Doesn't require any input from your live agent during the execution of this Knowledge Graph while the AI agent runs.</li> </ul>
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent:

Fields	Description
	<ul style="list-style-type: none"> <li>○ <b>Yes</b></li> <li>○ <b>No</b></li> </ul>
Processing message	Message to display to users when the tool is running.
Output transformation strategy	Style for the LLM to present the results. <ul style="list-style-type: none"> <li>○ None</li> <li>○ Concise</li> <li>○ Paragraph</li> <li>○ Verbose</li> </ul>

**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.


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.   <b>Note:</b> This description is sent to the large language model (LLM).
Select skill	Existing skill to be selected from the list.

Fields	Description
Inputs	<p>Skill inputs. The values are filled in by the LLM at runtime unless you specify a value override.</p> <p><b>Note:</b> 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.</p>
Execution mode	<p>Mode of execution for your skill:</p> <ul style="list-style-type: none"> <li>○ <b>Supervised:</b> Inputs from your live agent are required during the execution of this tool while the AI agent runs.</li> <li>○ <b>Autonomous:</b> Doesn't require any input from your live agent during the execution of this tool while the AI agent runs.</li> </ul>
Display output	<p>Permission to display the output of the tool execution in the Now Assist panel or in Virtual Agent:</p> <ul style="list-style-type: none"> <li>○ <b>Yes</b></li> <li>○ <b>No</b></li> </ul>
Processing message	<p>Message to display to users when the tool is running.</p>
Output transformation strategy	<p>Style for the LLM to present the results.</p> <ul style="list-style-type: none"> <li>○ None</li> <li>○ Concise</li> <li>○ Paragraph</li> <li>○ Verbose</li> </ul>

**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

**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
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.  <b>Note:</b> This description is sent to the large language model (LLM).
Inputs	<ul style="list-style-type: none"> <li><b>Input name:</b> Name of the input for the LLM to use.</li> <li><b>Description:</b> Description of the input to give the LLM context.</li> <li><b>Value override:</b> Value for the input. If you leave it blank, generative AI fills in the value for you.</li> </ul> <b>Note:</b> 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	<ul style="list-style-type: none"> <li><b>Create record:</b> Requires values for fields in the new record.</li> <li><b>Delete records:</b> Requires the conditions to determine which records to delete.</li> <li><b>Look up records:</b> Requires the conditions to determine which records to look up, number of results specified, which fields to return, the result order, and the result sort type.</li> <li><b>Update records:</b> Requires the values for fields in the updated record and conditions to determine which records to update.</li> </ul>
Execution mode	Mode of execution for your selected record operation: <ul style="list-style-type: none"> <li><b>Supervised:</b> Inputs from your live agent are required during the execution of this tool while the AI agent runs.</li> <li><b>Autonomous:</b> Doesn't require any input from your live agent during the execution of this tool while the AI agent runs.</li> </ul>

Fields	Description
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: <ul style="list-style-type: none"> <li>○ <b>Yes</b></li> <li>○ <b>No</b></li> </ul>
Processing message	Message to display to users when the tool is running.
Output transformation strategy	Style for the LLM to present the results. <ul style="list-style-type: none"> <li>○ None</li> <li>○ Concise</li> <li>○ Paragraph</li> <li>○ Verbose</li> </ul>

**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.

**Add a script form**

Fields	Description
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.   <b>Note:</b> 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.

Fields	Description
	Example: Input name is <i>task_record_sys_id</i> and description is <i>sys_id of the Task Record, this is mandatory.</i>
Script	JavaScript code that includes an object class or a function for your script.  <b>Note:</b> For improved security, use <code>GlideRecordSecure</code> instead of <code>GlideRecord</code> and <code>addUserEncodedQuery()</code> instead of <code>addEncodedQuery()</code> .
Execution mode	Mode of execution for your script: <ul style="list-style-type: none"> <li><b>Supervised:</b> Inputs from your live agent are required during the execution of this tool while the AI agent runs.</li> <li><b>Autonomous:</b> Doesn't require any input from your live agent during the execution of this tool while the AI agent runs.</li> </ul>
Display output	Permission to display the output of the tool execution in the Now Assist panel or in Virtual Agent: <ul style="list-style-type: none"> <li><b>Yes</b></li> <li><b>No</b></li> </ul>
Processing message	Message to display to users when the tool is running.
Output transformation strategy	Style for the LLM to present the results. <ul style="list-style-type: none"> <li>None</li> <li>Concise</li> <li>Paragraph</li> <li>Verbose</li> </ul>

**5. Select Add.**

A script tool is added in the Scripts section on the Add tools and information page.

**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
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.   <b>Note:</b> 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 <a href="#">Search profiles in AI Search</a>  .
Search sources	One or multiple sources that you want to add to your search profile. To learn more about the search sources, see <a href="#">Search profiles in AI Search</a>  .
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: <ul style="list-style-type: none"> <li>○ <b>Semantic:</b> This search is carried out with the logical meaning of the search query.</li> <li>○ <b>Keyword:</b> This search is carried out with the defined keywords.</li> <li>○ <b>Hybrid:</b> This search is a combination of both keyword and semantic searches.</li> </ul>
Semantic indexes	Fields on the source table that are indexed for a semantic search.   <b>Note:</b> Semantic indexed fields are required if the search criteria are semantic or hybrid.

Fields	Description
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: <ul style="list-style-type: none"> <li>○ <b>Supervised:</b> Inputs from your live agent are required during the execution of this tool while the AI agent runs.</li> <li>○ <b>Autonomous:</b> Doesn't require any input from your live agent during the execution of this tool while the AI agent runs.</li> </ul>
Display output	Permission to display the output of the tool execution in the Now Assist panel or in Virtual Agent: <ul style="list-style-type: none"> <li>○ <b>Yes</b></li> <li>○ <b>No</b></li> </ul>
Processing message	Message to display to users when the tool is running.
Output transformation strategy	Style for the LLM to present the results. <ul style="list-style-type: none"> <li>○ None</li> <li>○ Concise</li> <li>○ Paragraph</li> <li>○ Verbose</li> </ul>

**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

**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.  <b>Note:</b> This description is sent to the large language model (LLM).
Select subflow	Existing process automation capability you want to add.
Inputs	Subflow inputs. The values are filled in by the LLM at runtime unless you specify a value override.  <b>Note:</b> 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: <ul style="list-style-type: none"> <li>○ <b>Supervised:</b> Inputs from your live agent are required during the execution of this tool while the AI agent runs.</li> <li>○ <b>Autonomous:</b> Doesn't require any input from your live agent during the execution of this tool while the AI agent runs.</li> </ul>
Display output	Permission to display the output of the execution in the Now Assist panel or in Virtual Agent: <ul style="list-style-type: none"> <li>○ <b>Yes</b></li> <li>○ <b>No</b></li> </ul>
Processing message	Message to display to users when the tool is running.
Output transformation strategy	Style for the LLM to present the results. <ul style="list-style-type: none"> <li>○ None</li> <li>○ Concise</li> <li>○ Paragraph</li> <li>○ Verbose</li> </ul>

**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

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.  <b>Note:</b> 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 overrides or leave them blank for generative AI to fill in the details for you. <ul style="list-style-type: none"> <li>○ <b>Country:</b> Country where results come from.</li> <li>○ <b>Max tokens:</b> Maximum number of tokens to include in the response.</li> <li>○ <b>Number of results:</b> Total number of results acquired.</li> <li>○ <b>Search query:</b> Value to search for</li> <li>○ <b>Sites or domains:</b> Websites where you want to search.</li> </ul> <b>Note:</b> 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.

Fields	Description
Execution mode	<p>Mode of execution for your selected web search:</p> <ul style="list-style-type: none"> <li>○ <b>Supervised:</b> Inputs from your live agent are required during the execution of this tool while the AI agent runs.</li> <li>○ <b>Autonomous:</b> Doesn't require any input from your live agent during the execution of this tool while the AI agent runs.</li> </ul>
Display output	<p>Permission to display the output of the execution in the Now Assist panel or in Virtual Agent:</p> <ul style="list-style-type: none"> <li>○ <b>Yes</b></li> <li>○ <b>No</b></li> </ul>
Processing message	<p>Message to display to users when the tool is running.</p>
Output transformation strategy	<p>Style for the LLM to present the results.</p> <ul style="list-style-type: none"> <li>○ None</li> <li>○ Concise</li> <li>○ Paragraph</li> <li>○ Verbose</li> </ul>

**5. Select Add.**

A web search is added in the Web search list on the Add tools and information page.

**Test an AI agent**

Analyze the performance of an AI agent when it provides a resolution plan for your AI agentso you can see that it functions the way that you defined it.

**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 must select the **AI agent** radio button to test an AI agent. Select the AI agent that you want to test and provide a task with a concise summary and a reference number in the **Task** field to begin testing the AI agent.

When testing, you can see how the AI agent interacts with the AI Agent Orchestrator, an agent that directs different AI agents, and the Communicator, an agent that focuses on facilitating communication between agents and the Orchestrator.

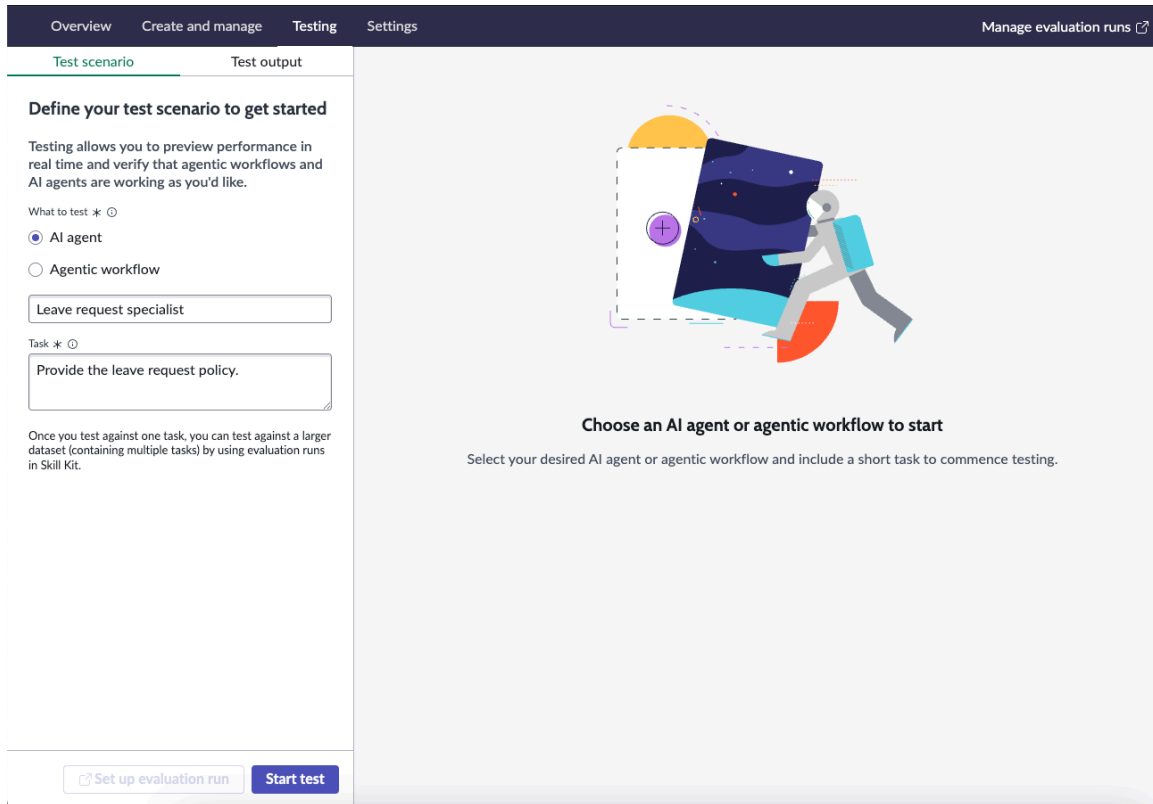
**Procedure**

1. Navigate to **All > AI Agent Studio > Testing.**
2. In the **Test scenario** tab, under the What to test section, select the **AI agent** button.
3. In the Select one drop-down list, select an AI agent that you want to test.

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 **Output** tab, where you can see Now Assist executing operations to test the AI agent.

## 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 or all of its tools, you will be notified that you don't have the necessary access and the test won't execute.

## Testing an AI agent

The screenshot displays the 'Testing' tab in the AI Agent Studio. On the left, a chat window shows a simulated conversation: 'You have 240 hours of vacation leave available' followed by 'Do you have any additional queries or questions regarding your leave request?'. Below the chat is a text input field with a 'Please type your response here' placeholder and a send button. In the center, a task flow diagram starts with 'Task Start', leading to an 'Orchestrator' node, which then branches into three parallel AI Agent nodes: 'Leave request spec...', 'Leave request spec...', and 'Communicator'. On the right, the 'AI agent decision logs' panel shows a table of activity:

AI agent decision logs		<a href="#">Download logs</a>
Orchestrator	Success	▼
Leave request specialist AI Agent	Success	▼
Orchestrator	Success	▼
Communicator AI Agent	Ongoing	▼
Leave request specialist AI Agent	Ongoing	▼

- A simulated chat experience begins on the Now Assist panel between your fulfiller agent and AI agent.
- 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.

**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**.

## 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

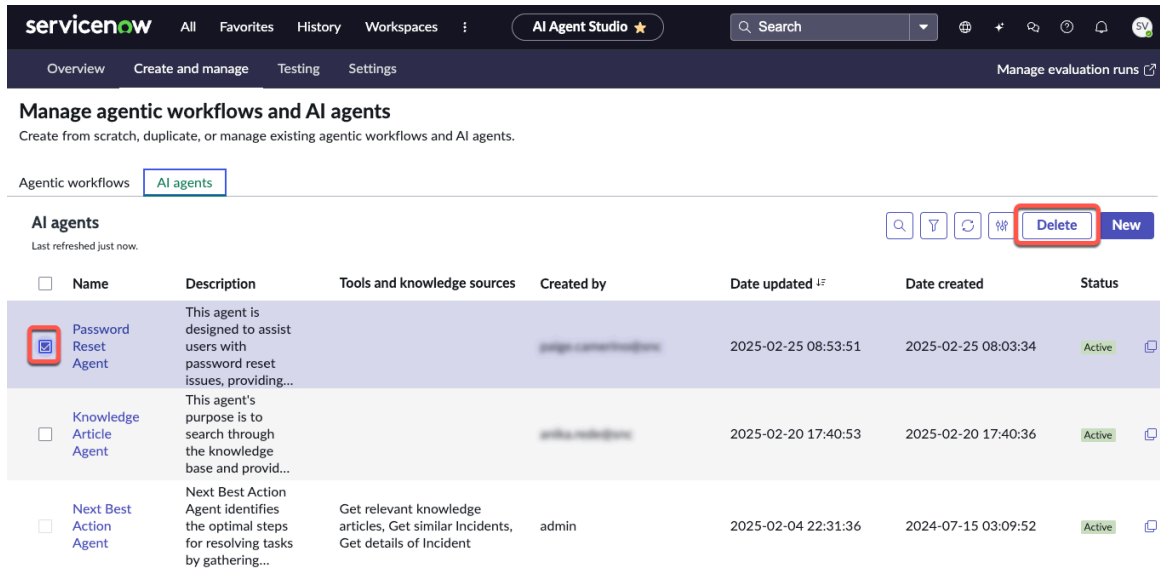
### About this task

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**.



**Note:** You can also delete multiple AI agent cases by selecting multiple AI agent records at a time if you define the ACLs.

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 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 with agentic AI. In AI Agent Studio, you must define an agentic workflow and connect it with an AI agent to execute complex goals involving 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**. You're directed to the New agentic workflow page on the Describe and connect step.

Overview Create and manage Testing Activity Settings Manage evaluation runs ⓘ

New agentic workflow Exit

- Define key requirements
- Add a preferred trigger
- Select a UI display

### Define key requirements

Each agentic workflow is directed by a large language model (LLM) that interprets instructions describing:

- when to use this workflow
- the steps required to complete the workflow
- the AI agents needed to complete the steps.

+ Generate description and steps

#### Write a clear description and list of steps

Using precise language, describe the purpose of this workflow and create a list of steps. For each step, define a specific action and any criteria. View versions

Workflow Name   
For example, "Customer request"

Workflow description  Description for LLM  
For example, "This agentic workflow finds and updates customer incidents or requests."

List of steps  Instruction for LLM  
Define one or more steps the agentic workflow should follow

#### Add AI agents that can perform these steps

The LLM compares the terminology in an AI agent's name and description to a workflow step, looking for a match. That's how it assigns AI agents. Add AI agent

#### Define who can access the agentic workflow

Users with these roles can discover and use this agentic workflow in UI displays. Once you save these roles, an access control will be automatically generated.

User access   
Users with specific roles

Role

#### Select the entity this agentic workflow will run as

This agentic workflow can run as the sys\_user of whichever entity (trigger, agentic workflow, or user) invokes it, which determines data access. Or it can run as a type of sys\_user called an AI user that has fixed roles. [Learn more](#)

Run as   
Dynamic user

+ Ask Now Assist to suggest AI agents Recommend AI Agents

Once you feel good about your agentic workflow description and instructions Now Assist can recommend AI Agents to add.

#### Unsupported model providers

List any model providers that this agentic workflow does not support.

Select unsupported LLMs

Save and continue

- If you want to use AI to help you with the agentic workflow setup, select **Generate description and steps** to create a draft of an agentic workflow description and instructions with Now Assist. On the Generate a description and list of steps with Now Assist pop-up window, you can select **Generate** to write the name, description, and instructions fields for you. You can make changes to those fields after the text has been generated or try again with new instructions.
- Describe the business outcome that this agentic workflow targets and connect the AI agents to the agentic workflow based on your requirements.

**Note:** If you have generated an agentic workflow by using Now Assist, the name, description, and list of steps fields are automatically filled in for you. You can still make changes!

- Give the agentic workflow a name, description, and list of steps to complete.

For examples on writing successful descriptions and instructions for the large language model (LLM), see [General guidelines for creating AI agents and agentic workflows](#). The guidelines explain some factors to consider when writing your description and list of steps and offer some examples.

**Write a clear description and list of steps form**

Fields	Description
Workflow Name	Name of your agentic workflow. Provide a name for the business challenge that you want to solve.  Example: Issue Resolution
Workflow description	Brief summary of what business problem that your agentic workflow is going to solve. This description is given to the LLM to understand the overall objective.
List of steps	Guided action plan to be followed by your AI agents in the agentic workflow. These steps should be a detailed numbered list that includes specific instructions for as many cases as possible, including potential errors.

- b.** Assign one or more AI agents to execute the instructions for the agentic workflow.
  - i.** In the Add AI agents that can perform these steps section, select **Add AI agent** and fill in the name of the AI agent in the **AI agent** field.
  - ii.** Select **Add**.
  - iii.** If you want to create an AI agent for your agentic workflow, then in the Add AI Agent drop-down list, select **Create new AI agent**.

For more information, see [Create an AI agent](#).

**Note:** If no AI agents are available for selection, the **Add AI agent** button is inactive.

The selected AI agent is added to the agentic workflow.

- c.** Configure Access Control Lists (ACLs) for the agentic workflow.

**Note:** The ACLs determine who has access to discover and execute the agentic workflow. 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 agentic workflow without creating an ACL, you must generate an ACL before you can make other modifications.


If you have already created an ACL for the agentic workflow, then it will appear in a list in this section. You can double-click the role pills to edit the roles.

**Define who can access the agentic workflow**

Field	Description
User access	The type of users whose access for the agentic workflow is defined by the following options:

Field	Description
	<ul style="list-style-type: none"> <li>▪ <b>Any authenticated user:</b> Any user who is logged in can access the agentic workflow.</li> <li>▪ <b>Users with specific roles:</b> Users that have at least one of the assigned roles to them can access the agentic workflow. This is the default option.</li> </ul> <p><b>Note:</b> If a user doesn't have access to an agentic workflow 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.</p> <ul style="list-style-type: none"> <li>▪ <b>Public:</b> Any user can access the agentic workflow even without logging in. Use this option only when you want guests to be able to access the agentic workflow.</li> </ul>
Role	<p>Assign one or more specific roles from the drop-down.</p> <p><b>Note:</b> Selecting the role is possible only when you chose the <b>Users with specific roles</b> user access.</p>

**d.** Define the type of sys\_user the agentic workflow will run as.

- For more information about understanding the user identities that will execute the agentic workflow, see [User identity](#).
- For more information about understanding the rules for Access Control Lists (ACLs) in ServiceNow AI Platform<sup>®</sup>, see [Access Control List Rules](#) .

**Select the entity the agentic workflow will run as**

Field	Description
Run as	<p>The user identity with which the agentic workflow runs.</p> <p>Select one of the following entity types that your agentic workflow should run as:</p> <ul style="list-style-type: none"> <li><b>i. Dynamic user:</b> The logged in user (sys_user) who invokes the execution of an agentic workflow.</li> </ul> <p>The dynamic user can be the sys_user of a trigger, agentic workflow, or human user that passes its roles to this agentic workflow (the upstream component/entity that invoked the agentic workflow).</p>

Field	Description
	<p>Any execution run by a dynamic user is recorded in the name of the logged in user. For example, if an incident is created by a Dynamic user, then the incident created is recorded in the name of the currently logged in user.</p> <p><b>i Note:</b> The Dynamic user is the default setting for all agentic workflows unless there is a specific need that justifies an AI user. Generally, most agentic workflow can run as dynamic users.</p> <p><b>ii. AI user:</b> The agentic workflow runs and executes as a preconfigured, static user identity with assigned roles that will not change regardless of who or how the execution is invoked.</p> <p>The AI user is a type of sys_user record of type AI agents. Any actions performed by an agentic workflow running as an AI user are logged in the name of the AI user for that action or execution. For example, if an incident is created by an AI user, then the caller for that execution is recorded as the assigned AI user.</p> <p>The AI user is a sys_user with predefined roles.</p> <p><b>Warning:</b> Allowing public users to invoke AI users will give them elevated permissions to what data is accessible to them. Agentic workflows very rarely need to run as AI users.</p> <p>If you do not have an AI user but want to use the <i>AI user</i> identity, you need to create a new record on the User table. See <a href="#">Create a user</a>. Select <i>AI user</i> as the identity type.</p>
AI user	Select a specific AI user from the drop-down menu.

**i Note:** When determining these security settings, be sure to consider the security settings for all aspects of the agentic flow, from agentic workflow, AI agents, and tools.

**e.** In the Ask Now Assist to suggest AI agents section, select **Recommend AI Agents** to explore the suggested recommendations for the AI agents.

**f.** Specify which AI providers this agentic workflow does not support.

Some AI models may perform your specific agentic workflow tasks better than others. You can choose an unsupported LLM from the **Select unsupported LLMs** drop-down if you do not want an agentic workflow to be available if a certain default AI model provider is active.

**g. Select Save and continue.**

Once you select **Save and continue**, Now Assist runs a check to try to identify if the agentic workflow is a potential duplicate of an existing one. If you see the Possible duplications found pop-up window, you can review the flagged agentic workflows and verify that the one you're creating is unique.

If there's an existing similar agentic workflow, it's possible for the wrong one to run when a user attempts to trigger it. Try to make your description and list of steps as unique as possible to avoid duplicates.

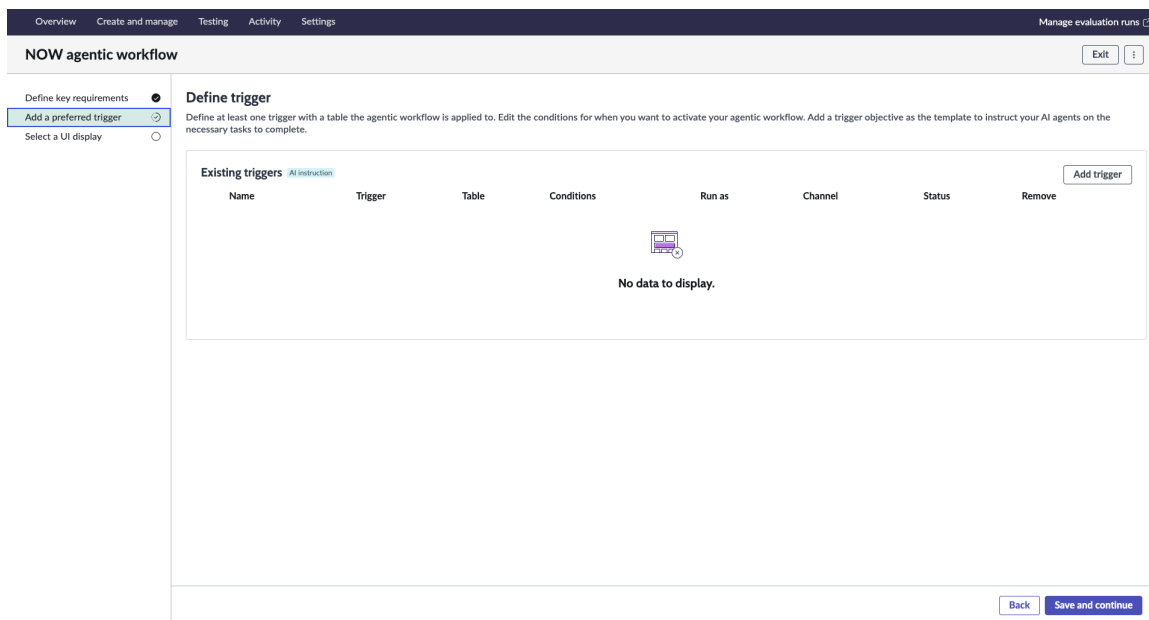
You must select **Save and continue** to create the ACL for the AI agent on the ACL table. If you want to create additional or *Deny - Unless* ACLs or add Security Attributes for the AI agent, a user with the correct elevated privileges can create them on the ACL table directly. Additional ACLs cannot be configured in AI Agent Studio.

If you selected *Users with specific roles* for your ACL, after you've selected Save and continue, you can still edit the roles of the ACL. Return to the first page of the guided setup and double-click the role pills in the ACL section to make changes.

You're directed to the Add a preferred trigger page.

**4. Define a preferred trigger and the conditions for activating your agentic workflow.**

**Note:** Triggers are optional. If you only want your AI agent to be discoverable in chats in the Now Assist panel, you don't need to set a trigger.



**a. Select Add Trigger.**

**b. On the form, fill in the fields.**

Trigger conditions for the agentic workflows set the conditions required for a user to execute an agentic workflow.

## Add trigger



Select trigger \*

Created

Name \*

NOW trigger for agentic workflow

Active

Table \*

.NET Application

Conditions\* [+ Add condition set](#) ⚠ At least one condition is required.

Method of defining sys\_user \* ⓘ

Use an existing table  Use a new custom script

Sys\_user \* ⓘ

Objective template \* ⓘ

Help me resolve \${number}

Channel \* ⓘ

Now Assist panel  Virtual Agent

Show Notification

No

[Cancel](#) [Add](#)

### Add trigger form

Fields	Description
Select trigger	List of triggers that are available in the instance.
Name	Name of the trigger.
Active	Option to keep your trigger conditions active.
Table	Name of the applicable table for your agentic workflow.  Example for the Issue Auto Resolution agentic workflow: Incident
Conditions	Conditions that you can add to control the trigger configuration.

Fields	Description
	Select <b>+ Add condition set</b> to add conditions to your agentic workflow trigger.
Method of defining sys_user	<ul style="list-style-type: none"> <li>Use an existing table</li> <li>Use a new custom script</li> </ul>
Sys_user	<p>The user that the AI agent runs as.</p> <p>If you select a table, then the choices could include the Caller, Resolved by, Closed by, or another user field. You can also create a custom script to generate and return the sys_id of the user.</p>
Objective template	<p>Template instruction for an AI agent to follow during its execution.</p> <p>For some trigger types, you can also select additional fields in your template instructions.</p>
Channel	<p>Medium for the agentic workflow output: Now Assist panel or Virtual Agent.</p> <p><b>i Note:</b> To view the output from a triggered agentic workflow in the Now Assist panel, you need the now_assist_panel_user role.</p>
Show notification	<p>Select the <b>Show Notification</b> option to enable notifications for your triggers.</p> <p>If there are multiple triggers with notifications enabled for the same table, only one notification is shown.</p>

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.

**i 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.

**c. Select Add.**

You're directed to the Define trigger page.

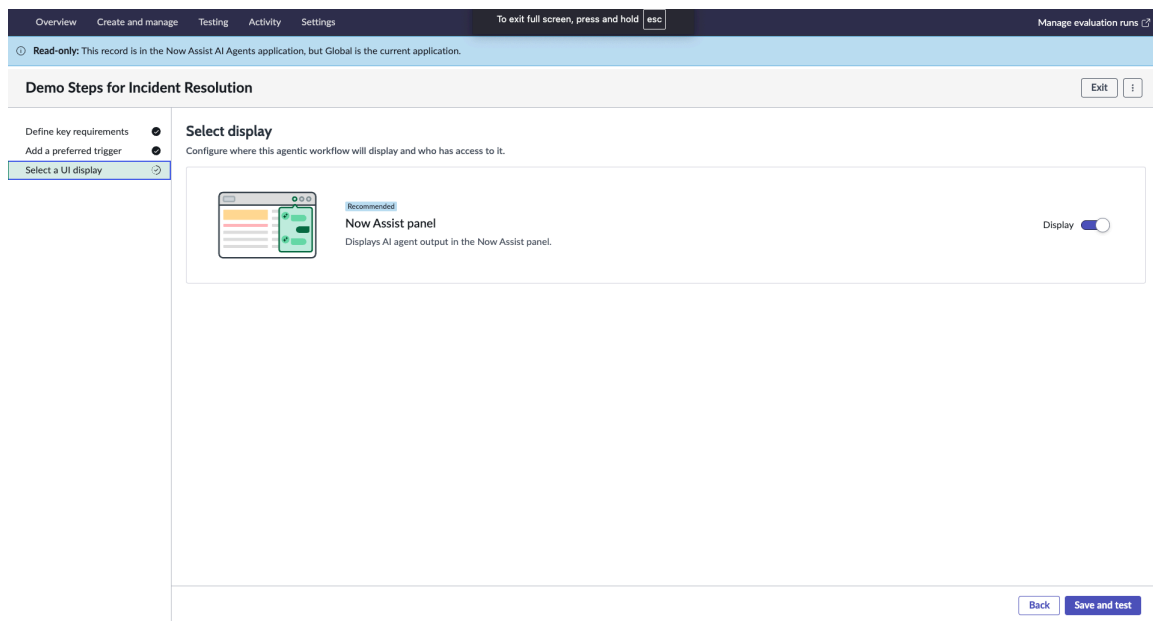
**d. Select Save and continue.**

You're directed to the Select a UI display page.

**5. Configure the display and access for your agentic workflow from the Select a UI display page.**

**a. Turn on the Display toggle to display your agentic workflow output on the selected channel.**

**Note:** You must toggle the **Status** to activate your agentic workflow, but selecting Now Assist panel is optional. Only enable it if you want your agentic workflow to be discoverable in that channel. If you only plan to trigger your agentic workflow by other means, you don't need to enable it.



**Note:** If the Now Assist panel display option isn't available, you must enable the panel first. For more information, see [Turn on the Now Assist panel](#). You must have the `now_assist_panel_role` to view the output from a triggered agentic workflow in the Now Assist panel.

**b. You can select the next icon ( > ) to access the User roles drop-down list and select a user role for who should be accessing the agentic workflow.**

**c. Select Save and test.**

You're directed to the agentic workflow testing page. For more information about testing an agentic workflow, see [Test an agentic workflow](#).

To test the agentic workflow, you must have the `sn_aia.admin` role and any roles the ACL configured for the agentic workflow requires, if applicable.

**Result**

An agentic workflow is created and can be seen in the agentic workflow list in the Create and manage 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.

## 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
<b>Manage agentic workflows and AI agents page</b>	On the agentic workflows list, select the duplicate icon (  ) for the agentic workflow that you would like to duplicate.
<b>AI agent form</b>	Open the agentic workflow that you want to duplicate, select the menu icon (  ) next to <b>Exit</b> on the Describe and instruct form, and select <b>Duplicate</b> .

You see a confirmation message in a pop-up window.

## You are duplicating a use case



This will produce an identical use case, but the triggers will be inactive. Fine-tune the copied use case with distinct instructions and triggers by following the use case configuration guided setup.

Cancel

Duplicate

3. Create a copy of the agentic workflow with the same information from the original agentic workflow's record by selecting **Duplicate**.

### **i** 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**.

- i** 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 triggers and change display settings for agentic workflows in AI Agent Studio that come installed with Now Assist applications.

### Before you begin

Role required: sn\_aia.admin

## 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. Some may come with predefined triggers that must be set to activate, but others don't have triggers at all and you must add your own.

These predefined agentic workflows can also be used as templates for your own customized ones. You can [duplicate a 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 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.

## 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 AI agents list on your AI Agent Studio.
2. Select the agentic workflow that you want to modify.
3. Modify the different sections of Guided Setup, making changes to the fields for your own requirements.

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 Guided Setup with the **Continue** and **Back** buttons.

**i Note:** Some fields aren't editable if the agent is associated with a Now Assist application.

4. Navigate to the last step and select **Test** to save your changes and begin testing your modified agentic workflow.

## Result

Your agentic workflow is modified and ready to test.

## Test an agentic workflow

Test your agentic workflow in AI Agent Studio to analyze its performance 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 the agentic workflow to confirm that it functions the way that you defined it. You must select the **Agentic workflow** radio button to test your agentic workflow. Select the agentic workflow that you want to test and provide a task with a concise summary and a reference number in the **Task** field to begin testing the agentic workflow.

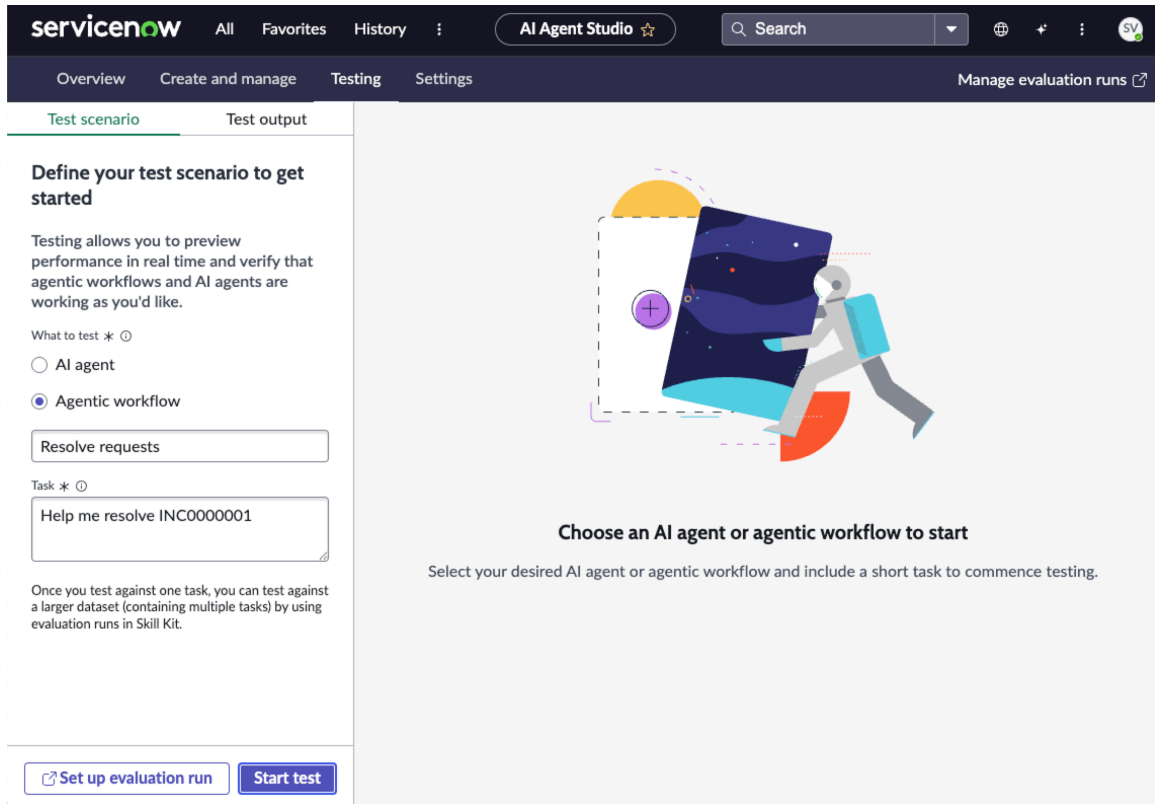
During the agentic workflow testing, you see a simulated chat experience in the AI Agent Studio chat on the **Output** tab. The chat experience is a result of the agentic workflow execution done by the fulfiller and AI agent.

## Procedure

1. Navigate to **All > AI Agent Studio > Testing**.
2. On the **Test scenario** tab, in the What to test section, select the **Agentic workflow** radio button.
3. In the Select one drop-down list, select an agentic workflow that you want to test.
4. In the Task drop-down list, provide a concise summary of the task to be achieved.

**i Note:** In the task summary, include a reference number for a specific record for better results during your testing.

5. Select **Start test**.



You're directed to the **Output** tab, where you see Now Assist executing the operations to test your agentic workflow.

## 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.

Your AI agent starts to execute the test autonomously to resolve the agentic workflow by taking the input from the live agent as required.

## Testing an agentic workflow

The screenshot displays the 'Testing' tab in the AI Agent Studio. On the left, a 'Test scenario' panel shows a 'Test output' section with a proposed plan for resolving an email access issue. The central area features a workflow diagram starting with 'Task Start', followed by an 'Orchestrator', and then an 'Issue Resolver Agent AI Agent'. This agent branches into three parallel tool calls: 'Get details of Incid...', 'Get similar Incidents', and 'Get relevant knowl...'. On the right, the 'AI agent decision logs' panel provides a detailed view of the 'Issue Resolver Agent' process, including 'Action Reasoning', 'Action Inputs', and 'Output Fields'.

- A simulated chat experience begins on the Now Assist panel between your fulfiller agent and AI agent.
- 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**.

## Delete an agentic workflow

Delete an agentic workflow from AI Agent Studio.

### Before you begin

Role required: sn\_aia\_admin

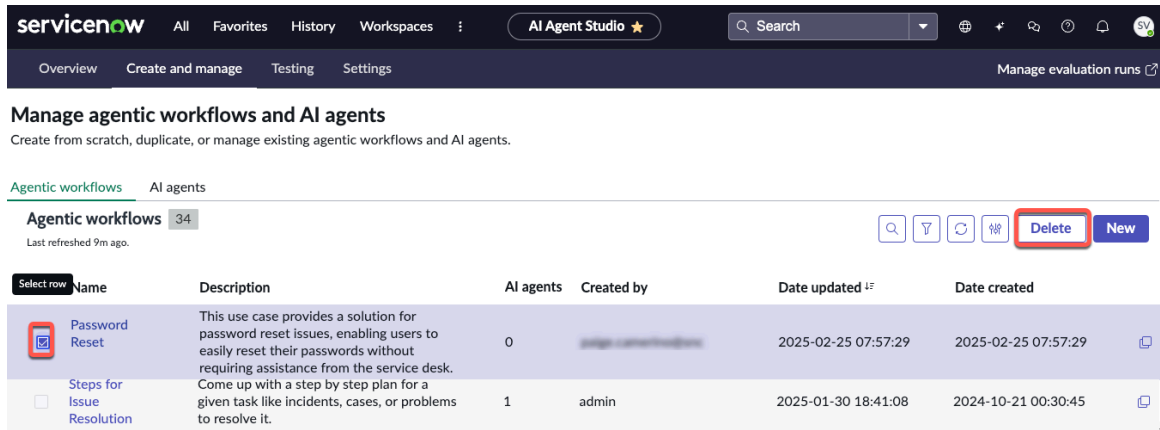
### About this task

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**.



**Note:** You can also delete multiple agentic workflows by selecting multiple agentic workflow records at a time if you define the ACLs.

3. In the confirmation pop-up window, select **Delete**.

### Result

The selected agentic workflows are deleted from the Agentic workflows list AI Agent Studio.

## Run an agentic evaluation

Evaluate an agentic workflow against a dataset of your choice to monitor performance and evaluate it against different benchmarks.

[https://player.vimeo.com/video/1076149291?h=648f0dfc45&badge=0&autoplay=0&player\\_id=0&app\\_id=58479](https://player.vimeo.com/video/1076149291?h=648f0dfc45&badge=0&autoplay=0&player_id=0&app_id=58479)

### Before you begin

Evaluation runs require execution log data of the agentic workflow you want to evaluate. For a new agentic workflow, you can create execution logs by testing in AI Agent Studio. For more information about testing agentic workflows, see [Test an agentic workflow](#).

For more information about getting started with agentic evaluations, see [General guidelines for agentic evaluation runs](#).

Role required: sn\_aia.admin

### 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 an agentic workflow and then select **Set up evaluation run**. A modal appears to ask if you want to be redirected to Now Assist Skill Kit. Select **Open Skill Kit**. You'll be 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 workflow that you want to evaluate.

< Categorize Incident Agentic Workflow Evaluation Save as draft

---

Set up an evaluation run. ^

- Add general info
- Select evaluation methods
- Choose a dataset
- Review summary

### General info

Tell us a bit more about the evaluation run you want to set up.

Name \*

Description

Agentic Workflow \*  [Go to workflow](#)

Agentic workflow description  
Categorize a new incident by assigning category, subcategory and CI to the incident

Exit setup Continue

**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 in the list 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 method.**

Overall task completeness evaluation is selected by default. Running multiple evaluation methods at a time can help provide a more comprehensive overview of the agentic workflow's performance.

To see more information about each plan, you can expand the card for each evaluation plan by selecting the chevron icon (∨).

< Categorize Incident Agentic Workflow Evaluation Save as draft

---

Set up an evaluation run. ^

- Add general info
- Select evaluation methods
- Choose a dataset
- Review summary

### Select evaluation methods

Default evaluation methods  
These are the primary evaluation methods and are pre-selected because they provide an overall assessment of the agentic workflow.

**Overall task completeness evaluation** ∨  
The Overall Task Completeness Evaluation metric assesses whether an AI agent successfully completes its assigned task.

Select additional evaluation methods

**Tool performance evaluation** ∨  
The Tool Performance Evaluation metric assesses an AI agent's ability to select the most appropriate tool for each step while completing a task.

**Tool calling evaluation** ∨  
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.

**6. Choose your dataset.**

- a. Select an existing dataset or create your own.
- b. To create a new dataset, fill out the form.

**Choose a dataset form**

Field name	Description
Name	Name of the dataset.
Description	General description of the dataset and its intended purpose.
Max records (optional)	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 will be ignored for that evaluation run.
Filters	Conditions for narrowing down the AI execution log records you want to include in the dataset. By default, the agentic workflow that you're evaluating is selected as a filter condition.

- c. 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 notice any place where you want to make changes, you can select **Back** to go to a previous step, or you can select the step in the sidebar.

< Categorize Incident Agentic Workflow Evaluation Save as draft

Set up an evaluation run. ^

- ⌚ Add general info
- |
- ⌚ Select evaluation methods
- |
- ⌚ Choose a dataset
- |
- ⌚ Review summary

**General info**

Name  
Categorize Incident Agentic Workflow Evaluation

Description  
This evaluation run aims to assess the deployment readiness of the Categorize Incident agentic workflow. It will help identify areas for improvement, if any, and outline steps to enhance the workflow's effectiveness

Agentic workflow  
Categorize incident

Agentic workflow description  
Categorize a new incident by assigning category, subcategory and CI to the incident

**Dataset**

Categorize Incident Agentic Workflow Eva To be generated

Description  
This dataset contains execution logs to evaluate Categorize Incident agentic workflow.

Records for evaluation  
225

**Evaluation methods** 4

Plan evaluation

Tool performance evaluation

Tool calling evaluation

Exit setup
Back
Start evaluation

## 8. Select **Start evaluation**.

### Result

Your evaluation run executes. The time it takes for an evaluation run to complete varies, but once it has been complete 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](#).

### General guidelines for agentic evaluation runs

Learn about agentic evaluation runs and different recommendations for evaluating your agentic workflows against datasets to check for completion, performance, and tool execution.

### Overview of agentic evaluation runs

Evaluation runs for agentic workflows evaluate agentic workflow executions for different metrics, such as task completion, performance, and tool execution. You can create datasets using logs for agentic workflows.

### When to run agentic evaluations

#### Run after you have collected enough data.

Evaluation runs are measured against logs of agentic workflow activity on your instance.

#### Run agentic evaluations when you make significant changes.

After making updates to the agentic workflow, you can execute an agentic evaluation run to track the efficacy of the new version.

### Choosing an evaluation method

#### 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.

#### Use multiple evaluation methods at a time.

Choosing multiple evaluation methods can provide a better overall picture of the agentic workflow's performance.

## Creating a dataset

### Use filters to target the right data.

Add filters to the execution logs to control exactly what you're measuring your agentic workflow against. Filter different time frames to verify that you're measuring the latest version of a workflow. 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.

## Examples of using AI agents

Review different ways that you can leverage the Now Assist AI agents application in agentic workflows across the platform.

### 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&autoplay=0&player\\_id=0&app\\_id=58479](https://player.vimeo.com/video/1098087911?h=d89ede16f9&badge=0&autoplay=0&player_id=0&app_id=58479)

### 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.

## Now Assist AI agents reference

Find more information about user roles, tables, and the different properties that are installed in Now Assist AI agents.

### 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.

### 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: <b>azure_openai</b> .
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: <b>3</b>
sn_aia.continuous_tool_execution_limit	Defines the maximum consecutive number of uses for the same tool.  Default value: <b>7</b>
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: <b>false</b>
sn_aia.ltm.category.auto_create	Enables AI-generated categories if no matching categories exist.  Default value: <b>false</b>

**System properties to use for configuring AI agents (continued)**

Property	Description
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: <b>false</b>
sn_aia.maximum_agent_tools	Defines the maximum number of tools that can be assigned to an AI agent.  Default value: <b>20</b>
sn_aia.max_scheduled_trigger_query	Defines how many records are processed when a scheduled trigger is detected.  Default value: <b>10</b>  Write_role: admin
sn_aia.mid_skill_switch_enabled	Enables mid-skill switching.  Default value: <b>false</b>
sn_aia.react_failure_retry_max_limit	Defines the maximum number of retries in case of an execution failure.  Default value: <b>3</b>
sn_nowassist_va.router_redirect_va_agentic	Determines AI agent discovery in Virtual Agent. If set to <i>NEVER</i> , Virtual Agent continues Q&A without any agentic AI.  Default: <i>ROUTER_DECISION</i>
com.glide.cs.dynamic.capability.timeout	Defines the Timeout for AIA Proficiency Descriptor.  Default value: <b>180</b> .  Write_role: admin.
sn_aia.enable_follow_up	Enables users to continue the conversation with follow-ups after the agentic workflow execution is complete.  Default value: <b>true</b> .
sn_aia.follow_up_message	Defines a follow-up message sent after execution is completed.  Default Value: <b>How else can I help you?</b>
sn_aia.allow_context_sharing	Enables the sharing of short-term memory, allowing context to persist across execution within the same conversation.

**System properties to use for configuring AI agents (continued)**

Property	Description
	Default Value: <b>true</b>
sn_aia.agent_strategy_choice_enabled	Enables to show the LLM reasoning strategy in the agent setup screen.  Default Value: <b>false</b>
sn_aia.context_sharing_strategy	This property defines the strategy to use for storing short-term memory for an execution.  Default Value: <b>summarise</b>
sn_aia.enable_agent_tool_input_value_overrides	Enables you to override the agent tool input value.  Default Value: <b>true.</b>
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: <b>1.</b>
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: <b>true.</b>
sn_aia.quick_mode_failure_retry_max_limit	Defines maximum limit for retries in case of a failure in Quick Mode execution.  Default value: <b>3.</b>
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: <code>getUserContext</code> .  List of available user information: <ul style="list-style-type: none"> <li>• profile</li> <li>• manager</li> <li>• reportees</li> <li>• assets</li> </ul> <p>The user information can also be customized by overriding the method <code>getUserContext</code> via <code>UserContextUtil</code>.</p>

### System properties to use for configuring AI agents (continued)

Property	Description
	<p>If customized, the property must define the comma separated list of keys generated by the customized <code>getUserContext</code> method.</p> <p>Default value: <b>profile</b>.</p>

### 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]	<p>List of tools executed by the plan ID.</p> <p><b>Note:</b> The records in the Tools Executions table expire and become unavailable after a period of 13 months.</p>
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]	List of the records created for hiding citations specific to AI agents and agentic workflows.
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.

**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.

**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.

**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.
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 95% of LLM requests in a system to be processed.
Number of closed tasks	Total number of tasks closed.

**Automated indicators for AI agents**

(continued)

Name	Description
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.
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.

### 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
Activity	Time created

**Filters on the AI Agent Analytics dashboard pages (continued)**

Page	Available filters
Performance	<ul style="list-style-type: none"> <li>• Time created</li> <li>• Agentic workflow</li> <li>• AI agent</li> <li>• Tool</li> </ul>
Troubleshooting	Time created

**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
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.
Execution tasks in last 7 days	Pie chart	Execution tasks divided by status (Success, Ongoing, Cancelled, Queued).
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.

**Data visualizations on the AI Agent Analytics dashboard overview page (continued)**

Title	Type	Description
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.
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 activity 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 value page**

Title	Type	Description
Tasks closed using AI agents	Trend line	Percentage of tasks closed using AI agents. Trend line is shown.
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.

**Data visualizations on the AI Agent Analytics dashboard value page (continued)**

Title	Type	Description
Efficiency gain supporting indicators	List	List of indicators. The list is broken down by indicator. Columns span across a time span.
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.

**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 workflow latency	Score card	Time in seconds of 95% of agentic workflows in a system to be processed.
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.

Data visualizations on the AI Agent Analytics dashboard performance page (continued)

Title	Type	Description
Agentic workflow latency	Line chart	Agentic workflow latency over time. Uses three 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 three 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 three 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.
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 three different indicators: P90 tracks the fastest 90% of executions, P95 tracks the fastest 95% of executions, and P99 tracks the fastest 99% of executions.

Data visualizations on the AI Agent Analytics dashboard performance page (continued)

Title	Type	Description
AI agent latency (weekly average)	Line chart	AI agent latency over time by week. Uses three 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 three 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.
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 three 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 three 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 three different indicators: P90 tracks the fastest 90% of executions, P95 tracks the fastest 95% of executions, and P99 tracks the fastest 99% of executions.

**Data visualizations on the AI Agent Analytics dashboard performance page (continued)**

Title	Type	Description
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 three 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 three 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.

**Data visualizations on the AI Agent Analytics dashboard troubleshooting page**

Name	Type	Description
Agentic workflow errors	List	List of agentic workflow errors.
LLM 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 <a href="#">role masking</a> and data access.
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.

**Data visualizations on the AI Agent Analytics dashboard security page (continued)**

Name	Type	Description
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 agentic workflows are accomplishing their objectives. A Now LLM Service model judges the agentic workflow based on the execution logs of that agentic workflow. 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 workflow. Recommended actions give you suggestions for deployment or improvement to help ensure that the agentic workflows that you deploy are performing up to your standards.

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**.

### Overall task completeness evaluation run results

Label	Description	Recommended action	Default threshold
Excellent	Tasks were consistently performed at a high standard. The agentic workflow 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.

#### 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.

**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.*

(continued)

Number	Score	Description
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 completeness, and input format completeness were successful.
0	False	One or more of input key completeness, input value completeness, or input format completeness wasn't successful.

**Now Assist glossary**

Learn about the terms and concepts used in Now Assist AI agents.

**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

In the agentic system on your instance, an agentic workflow contains a set of LLM instructions with one or more AI agents that can execute an objective.

### AI agent

An AI agent contains a set of large language model (LLM) instructions with the tools to perform these specific tasks.

### Citation

Small, interactable number next to AI-generated content in the Now Assist panel that cites the source of the information.

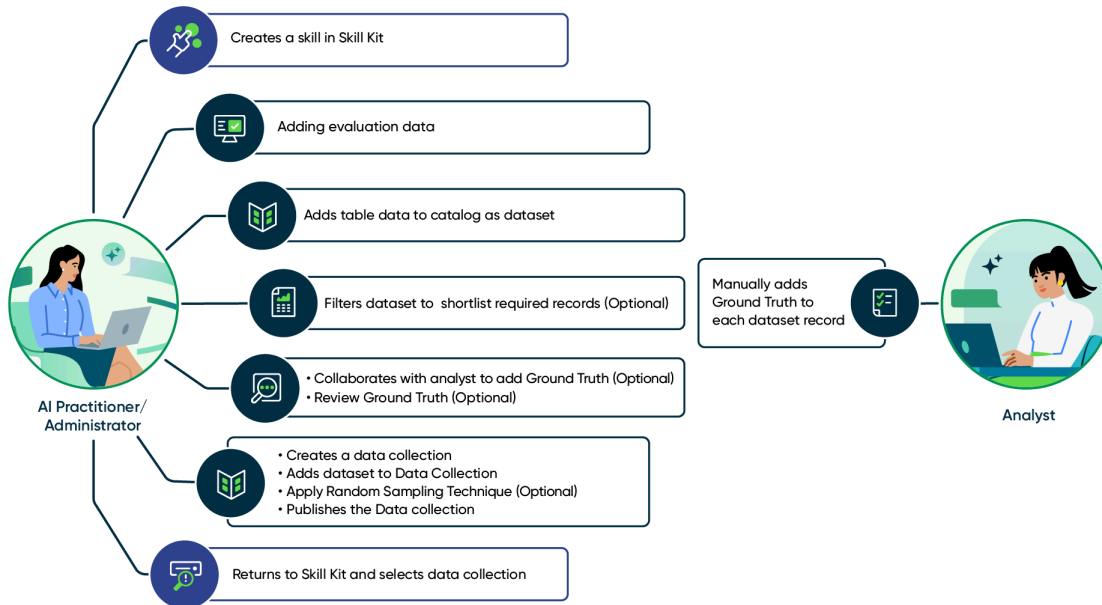
### Orchestrator

AI Agent Orchestrator provides a structured process to ensure seamless collaboration between AI agents so that the teams of AI agents can work together in harmony to manage and govern AI agents as well as track AI use cases and add value across the enterprise.

## Now Assist Data Kit



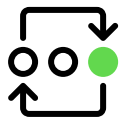

Use ServiceNow® Now Assist Data Kit to add datasets to a data catalog. The curated data works with ServiceNow SDK to enable the AI skill development and evaluation.

Now Assist Data Kit workflow overview.



### Get started

Now Assist Data Kit

<p>Explore</p>  <p>Learn how Now Assist Data Kit can help create custom datasets</p>	<p>Configure</p>  <p>Configure Now Assist Data Kit for prompts and skills</p>	<p>Use</p>  <p>Use Now Assist Data Kit to create datasets to add to data collections</p>
		<p>Reference</p>  <p>Reference information for Now Assist Data Kit</p>

**i 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.

**i 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. Further, 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 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.

**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.	<a href="#">Add a dataset</a>	AI developer
Create smaller datasets from existing datasets.	<a href="#">Create a derived dataset</a>	AI developer
Merge multiple datasets and create a collection that can be used by Now Assist Skill Kit.	<a href="#">Activate a skill</a>	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](#)

## 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

#### 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. On the **Data catalog** tab, select **Add New**.
3. On the Choose data form, select the import type, table, and column.  
There's an option to add columns, such as group notes and comments. These columns aren't stored in the incident table but in a separate table. You can add these columns by using a script.
4. Select *Add a filter if needed*.
5. Review the records and select **Save & next**.
6. On the form, fill in the fields.

#### Add dataset info form

Field	Description
Dataset name	Name of the dataset.
Category	Category that the dataset is in.
Dataset description	Description of the dataset.
Dataset source type	Dataset source that you can select from the drop-down menu.

7. Add tags to identify the dataset.
8. Navigate to the Data governance section and select each check box.
9. Select **Add dataset**.  
The dataset is added to the catalog.

**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.

**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. From the Dataset view page, select **Create derived dataset**.
2. Add information for the new derived dataset.

**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.

3. 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.
4. Select **Run**.
5. 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.

**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

**About this task**

After you successfully add a derived dataset to the data catalog, a pop-up window opens and then you can select **Add ground truth**.

**Procedure**

1. Select how you want to add the ground truth for the records in your dataset.
2. On the form, fill in the fields.

**Add ground truth info form**

Field	Description
Type	Currently, Add manually is supported.
Ground truth type	Type.
Ground truth guideline	Guidelines for labelers and linguists who manually add the ground truth.
Column name	Name.

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.

12. Select **Confirm** to make your collection available.

The data collection is published and available in Mobile SDK for evaluation.

## Now Assist Data Kit reference

Reference topics for Now Assist Data Kit.

### 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.

### 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 Skill Kit

Use ServiceNow<sup>®</sup> 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&autoplay=0&app\\_id=58479](https://player.vimeo.com/video/985874233?h=2ccc18794b&badge=0&autoplay=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.

### **i** 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

- [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. Further, 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

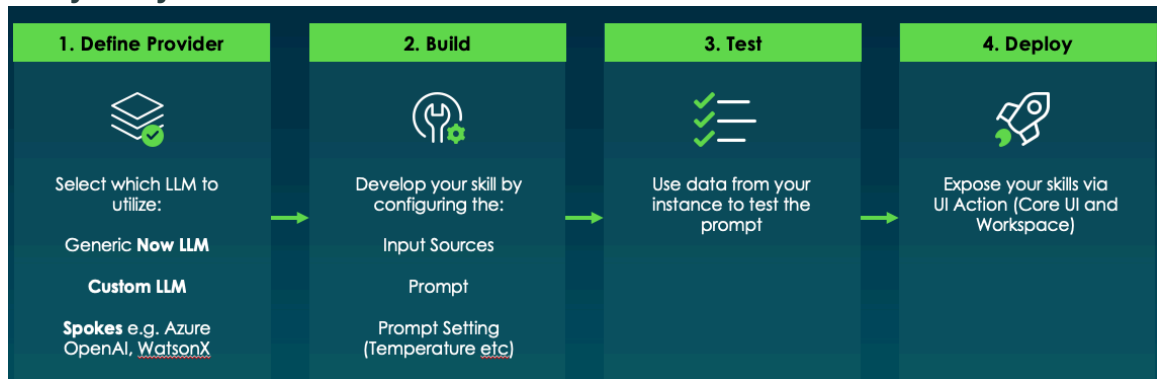
**Now Assist Skill Kit Users**

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.

### 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 directly to a UI action.

### Now Assist Skill Kit benefits

#### Now Assist Skill Kit benefits

Benefit	Feature	Users
Create custom solutions by building a custom skill or workflow.	<a href="#">Create a skill</a>	AI developer
Create and edit prompts for skills and configure where you want to bring in data from to augment your prompt.	<a href="#">Create a prompt</a>	AI developer
Test and iterate on your skill before activating it.	<a href="#">Test a prompt</a>	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](#)

### 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.

### 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:

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.

### 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

		"True" Underlying Accuracy Gap			
		5%	10%	15%	20%
Data used in evaluation	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).

### 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?

**i 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.

### 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

**New skill** ✕

Skill name \*

Description

Use the description above to recommend an AI generated prompt (powered by Now LLM Generic)

Default provider \*

Provider API \*

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.

### Add skill input

Datatype \* ⓘ

Table name \* ⓘ

Name \* ⓘ

Description

Mandatory  
 Truncate

Choose test record ⓘ

Select insert inputs for the prompt.

### 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}}
```

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.

## Add tool



Type

Flow Action

Name \* ⓘ

ChildIncidents

Resource \* ⓘ

IncidentDetailsFetcher

### Inputs

Name ⓘ	Datatype ⓘ	Value ⓘ
incidentSysId	string	{{incident.sys_id}}
Name ⓘ	Datatype ⓘ	Value ⓘ
numberOfComments	string	2

### Output

Name ⓘ	Type ⓘ
--------	--------

## 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**.

## 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](#).

## Security for Now Assist Skill Kit

Enable security controls for Now Assist skills and custom skills through access control lists (ACLs) and user identities.

### Access control lists

The access control lists in Now Assist Skill Kit enhance the security of Now Assist skills and are set to determine users who can invoke a skill.

### Configure ACLs in Now Assist Skill Kit

You can configure ACLs for Now Assist Skill Kit when you create or edit a skill or when you activate a skill from Now Assist Admin console.

ACLs configured in Now Assist Skill Kit for are *Allow-If* and role-based.

The *Allow-If* logic grants access to data or resources if any of the conditions in the ACL are met. The other type of ACL is *Deny-Unless*. *Deny-Unless* ACLs block access to data or resources unless a condition is met, even if there are other conditions like *Allow-If* ACLs that would normally grant someone access.

There are two possible options for ACLs created in Now Assist Skill Kit:


- *Any authenticated user*: Grants access to any user authenticated on the instance, regardless of role
- *Select roles*: Requires you to select the roles that grant access

Each skill must have its own unique ACL. You can't create a skill or save changes to a skill without an ACL. To configure an ACL for a skill, see .

## 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.**

>

## Configurations

Adjust prompt settings to change output quality.

Provider  
Now LLM Generic  
Provider API  
Now LLM Generic

Model ⓘ

llm\_generic
▼

Temperature \* ⓘ

0.2

Maximum response tokens \* ⓘ

1000

Maximum request tokens \*

6942

ⓘ Maximum number of tokens allowed in request.

---

**Usage conditions** +

Conditions determine when to use this prompt based on user inputs.

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.

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.

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](#).

## 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**.

### Skill Settings

Create or manage contents of your skill

General information

- > Providers
- > Evaluation Methods

+

Deployment Settings

### Deployment Settings Save

#### Now Assist Features

Workflow \* ⓘ Product \* ⓘ Feature \* ⓘ

Technology

ITSM

Child Incident Summarization

#### Feature Details

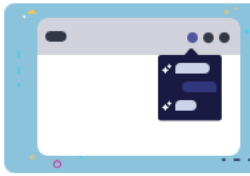
Name \* ⓘ


Child Incident Summarization

Description

Use generative AI to summarize child incidents.

#### Allow the Admin to enable this skill from the following

Now Assist Panel  
  
Allow this skill to be triggered through the Now Assist Panel

UI Action  
  
Adds a UI action button to the core UI, to trigger the skill

6. On the form, fill in the fields.

### Deployment settings form

Field	Description
Workflow	<p>The high-level category that this skill pertains to, such as <b>Technology</b>, <b>Employee</b>, <b>Creator</b>, or <b>Platform</b>. You can also select <b>Other</b> if none of the categories fit.</p> <p>The workflow that you choose is where the skill appears in the Now Assist Admin console.</p>
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

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](#)

## Using Now Assist Skill Kit

Use Now Assist Skill Kit to create and publish prompts and custom skills for Now Assist.

Building a custom skill in Now Assist Skill Kit is a sequential process. The following steps take you from an empty skill through to activation, where end users can trigger it from the platform.

### Create a skill

Create a skill or clone an existing base system skill. Clone a Now Assist skill if it's close to what you need but requires modification.

- [Create a skill](#)
- 

### Build the prompt

A prompt is an instruction template sent to the AI model when the skill runs. Define what information the skill receives as inputs, write the prompt text, and optionally add tools that gather additional context before the prompt executes.

- [Create a prompt](#)
- 
- [Add a retriever](#)
- 
- [Use prompt assistance](#)

### Test and evaluate

Before publishing, test your prompt against real records on your instance to verify the output. You can also run a formal evaluation, which measures output quality against expected results using correctness and faithfulness scores.

- [Test a prompt](#)
- [Evaluate a prompt](#)

### Publish and activate

When your skill is ready, finalize the prompt and publish it. Publishing makes the skill visible to a Now Assist admin, who then activates it in Now Assist Admin to make it available to your users.

- [Finalize and publish a skill](#)
- [Activate a skill](#)

## Call a skill from a script

After a skill is published, you can call it from a server-side script, allowing you to integrate the skill into automated workflows or business rules.

- [Call a custom skill from a script](#)

## 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

### Procedure

1. Navigate to **All > Now Assist Skill Kit > Home**.
2. Select **Create new skill**.
3. On the form, fill in the fields.

**New skill** ✕

Skill name \*

Description

Use the description above to recommend an AI generated prompt (powered by Now LLM Generic)

Default provider \*

Provider API \*

### New skill form

Field	Description
Skill name	A name for the skill.
Description	A description of the skill.
Default provider	Available providers: <ul style="list-style-type: none"> <li>○ Now LLM Service</li> <li>○ External LLM <ul style="list-style-type: none"> <li>▪ Spokes</li> <li>▪ Custom LLM</li> </ul> </li> </ul>

Field	Description
	<p>For more information on setting up a custom large language model (LLM), see <a href="#">Configure a generic large language model (LLM) connector</a></p> <p>Available prebuilt spokes that enable you to connect with an external LLM:</p> <ul style="list-style-type: none"> <li>○ Microsoft Azure OpenAI Generative AI Spoke</li> <li>○ OpenAI Generative AI Spoke</li> <li>○ Aleph Alpha</li> <li>○ WatsonX</li> <li>○ Google Gemini (MakerSuite and Vertex AI)</li> </ul> <p><b>Note:</b> The spokes don't consume Integration Hub transactions. The spokes consume assists.</p>
Provider API	The provider of the API for your chosen LLM.

**4. Select Create skill.**

**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. To learn more about creating a prompt, see [Create a prompt](#).

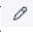
**Create a prompt**

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

**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**.

## Add skill input

Datatype \* ⓘ

Table name \* ⓘ

Name \* ⓘ

Description

Mandatory

Truncate

Choose test record ⓘ

Cancel

Add skill input

### Skill input options

Input type	Description
Datatype	<ul style="list-style-type: none"> <li>○ Record</li> <li>○ String</li> <li>○ Numeric</li> <li>○ Boolean</li> <li>○ Simple Array</li> <li>○ JSON Object</li> <li>○ JSON Array</li> </ul>
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.
Truncate	The Truncate option means that, if your prompt is too big, the prompt context will be 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	

Input type	Description
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. Add a tool.



10. On the form, fill in the fields.

**Tools form**

Field	Description
Type	Flow Action.
	Retriever.
	Subflow.
	Script.
Name	The name of the tool.
Conditions	<ul style="list-style-type: none"> <li><input type="radio"/> Always</li> <li><input type="radio"/> Script</li> </ul>

11. Select **Add**.

12. If you're not ready to finalize the prompt and publish the skill, select **Save** or **Save as**.

**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 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 **Add Tool**.
4. Select **Retriever** from the drop-down menu.
5. Select **Configure retriever**.
6. 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 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	<ul style="list-style-type: none"> <li><input type="radio"/> Hybrid</li> <li><input type="radio"/> Semantic</li> <li><input type="radio"/> Keyword</li> </ul> <p><b>Note:</b> If you choose Hybrid or Semantic, you can make selections for chunking and reranking. To learn more about chunking and reranking, see <a href="#">Retriever chunking and reranking</a>.</p>

7. Select **Next**.
8. 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.
9. **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.

**10. Select Next.**

**Note:** If you selected **Hybrid** or **Semantic** search criteria, see [Retriever chunking and reranking](#) to complete setting up your retriever.

**11. Review the retriever tool information.**

**12. 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	<ul style="list-style-type: none"> <li>○ 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.</li> <li>○ 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</li> </ul>

Field	Description
	<p>linked together into a large text and broken into smaller passages.</p> <p>Don't use the small to big chunking strategy if you're using full text or truncate index chunking configuration.</p>
Chunking unit	<ul style="list-style-type: none"> <li>○ Words</li> <li>○ Sentences</li> </ul>
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.

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**.


### 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. 
4. Select the **Library** tab to search for a prompt from the prompt library.

## 💡 Prompt assistance

Use a prompt template to get started on your prompt engineering journey

**Library** AI generated

Templates prompts are those provided by your organization.

Providers **Now LLM Generic** ▼

5.

6. Select the **AI generated** tab to have AI help you create a prompt based on your needs.

## 💡 Prompt assistance

Use a prompt template to get started on your prompt engineering journey

Library **AI generated**

AI generated prompts are those that can be created from a text description that you can provide.

Describe what you need ✖ ⓘ

[Generate AI Prompt](#)

Disclaimer: The prompt will be generated by a Generative AI model. Be sure to review it for accuracy before using.

### 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


### 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.fdi.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 **Evaluation** tab.
4. Create a dataset from a table or data collection.

**Create a data set**

Method	Steps
Create a dataset from a table	<ol style="list-style-type: none"> <li>a. Give the dataset a name and description.</li> <li>b. Select <b>Table</b>.</li> <li>c. Find the table that you want to use.</li> <li>d. Select the maximum number of records that you want to use.</li> <li>e. Add conditions.</li> <li>f. Select <b>Generate Preview</b>.</li> <li>g. Select the mappings.</li> <li>h. Select <b>Create</b>.</li> </ol>
Create a dataset from a data collection	<ol style="list-style-type: none"> <li>a. Give the dataset a name and description.</li> <li>b. Select Data Collection.</li> <li>c. Select a data collection that you created in Now Assist Data Kit.</li> <li>d. Select <b>Generate Preview</b>.</li> <li>e. Select the mappings.</li> <li>f. Select <b>Create</b>.</li> </ol>

5. Select the add icon <sup>+</sup> for **Evaluation Runs**.
6. Give the evaluation run a name and description.
7. Select one or more prompts that you want to evaluate.
8. Select **Save & Next**.
9. Select a dataset.
10. Select **Save & Next**.
11. Expand the **Quality** tab.
12. 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, and writing quality relative to the given instruction. This metric helps to check that the text accurately


Evaluation method	Metric	Description
		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.

13. Select **Save & Next**.

14. Review the evaluation choices that you made.

15. Select **Save & Evaluate**.

16. **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**.

**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**

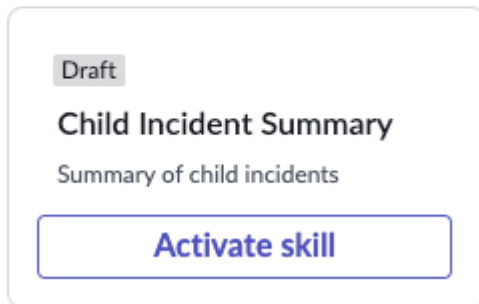
Role required: admin

**Procedure**

1. Navigate to **All > Now Assist Admin > Features**.
2. Find the custom skill that you created.
3. Select **Activate skill**.

**All available Child Incident Summarization skills**

Set up and configure each skill. A skill delivers generative AI functionality.


**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.

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);

```

## 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 <a href="#">Now Assist Overview</a> 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.
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.

## 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
 <p>Learn about the Now Assist Readiness Evaluation app.</p>	 <p>Run scheduled jobs and configure the app through the Now Assist Readiness Evaluation guided setup.</p>	 <p>Learn about how to use the Now Assist Readiness Evaluation dashboards and assessments to prepare for implementing agentic AI, Now Assist, or both.</p>

**i 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. Further, 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 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.

**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.	<a href="#">Run the GenAI/AgenticAI Assessment scheduled job</a>	Admins
Review an overall dashboard of your issues before implementing agentic and generative AI Now Assist.	<a href="#">Using Now Assist Readiness Evaluation dashboard</a>	Admins
View your overall go or no-go status to implement Now Assist generative AI skills or	<a href="#">Assessing go or no-go</a>	Admins

**Now Assist Readiness Evaluation benefits (continued)**

Benefit	Feature	Users
agentic AI agents, and view the go or no-go status' issues categorized by percentage.		
Review summary information of overall implementation readiness.	<a href="#">Reviewing your Now Assist assessment</a> and <a href="#">Reviewing your agentic AI assessment</a>	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.	<a href="#">Reviewing your Now Assist assessment</a> and <a href="#">Reviewing your agentic AI assessment</a>	Admins
Receive an estimated remediation effort of development work days to fix issues before implementation.	<a href="#">Reviewing your Now Assist assessment</a> and <a href="#">Reviewing your agentic AI assessment</a>	Admins

**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.
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.

**i 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.

**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.

**i 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 go or no-go](#)

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.

- [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.

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.

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.

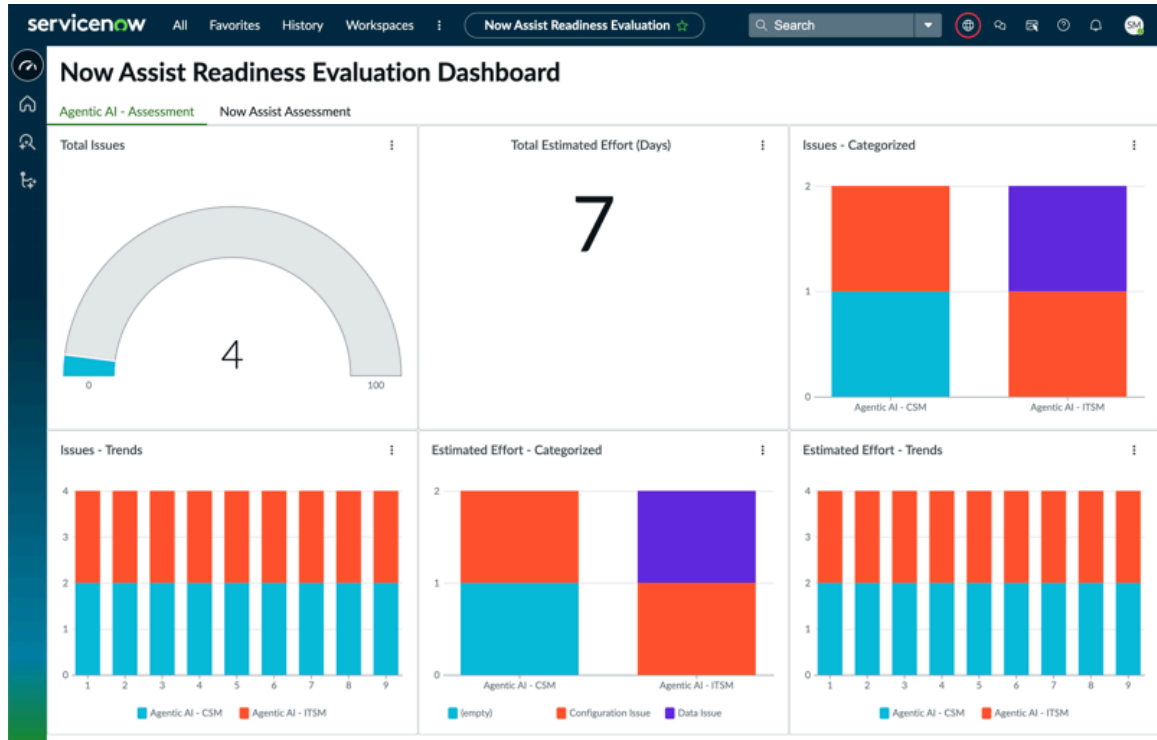
**Issues - Categorized**

This area of the dashboard shows the percentage of each issue category for each product in a bar chart.

**Estimated Effort - Trends**

This area of the dashboard shows the trends over time in a bar chart for each products' estimated work effort.

**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.

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.

**Total Issues**

This area of the dashboard shows the total issues identified.

**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.

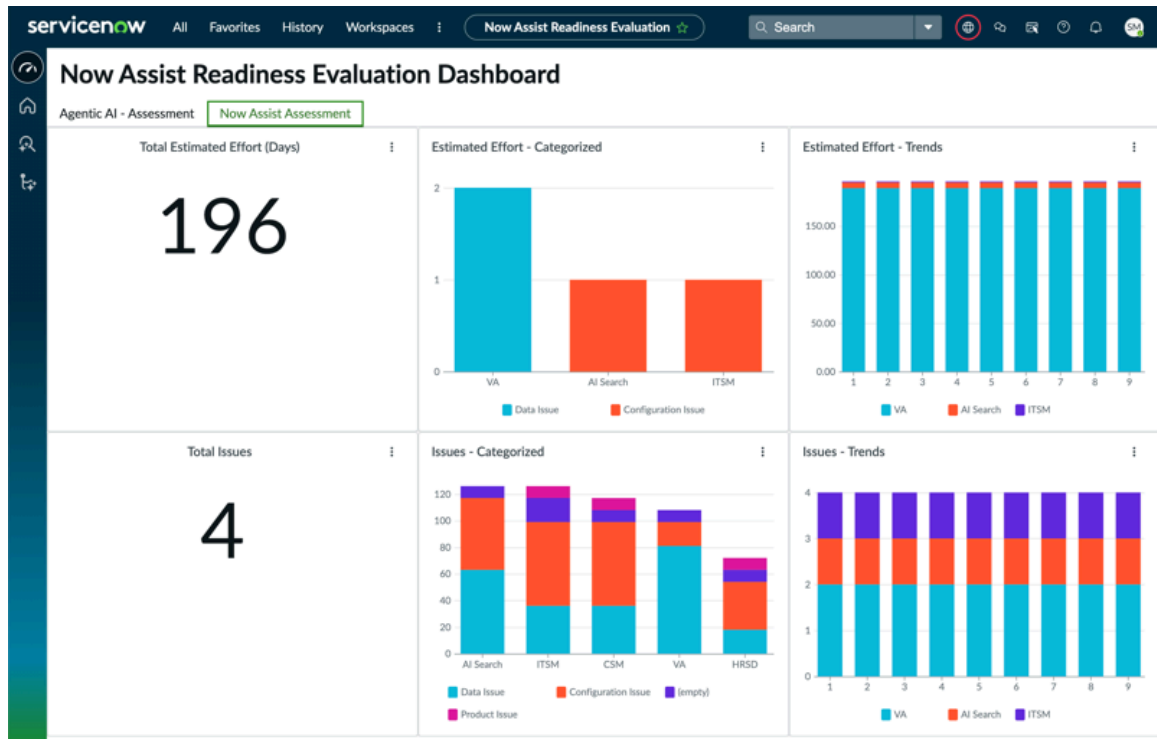
**Estimated Effort - Trends**

This area of the dashboard shows the trends over time in a bar chart for each products' estimated work efforts.

### 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.

### Example of the Now Assist Readiness Evaluation Dashboard's Now Assist Assessment



### Assessing go or no-go

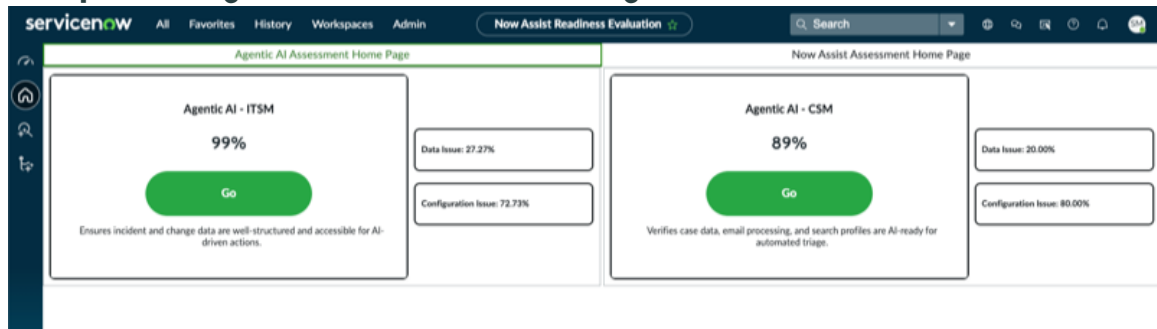
The Now Assist Readiness Evaluation app shows you a high-level overview of your go or no-go status for your organization's readiness to implement agentic AI, generative AI, or both in Now Assist.

Navigate to **Workspaces > Now Assist Readiness Evaluation** and select the Home icon  to review your assessments' home page.

An overall go or no-go status is determined by percentage rules. A green **Go** button appears if over 75% of your instance is ready to implement agentic AI, generative AI, or both in Now Assist. A yellow **No-go** 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 go or no-go status for Now Assist agentic AI in ITSM and CSM.

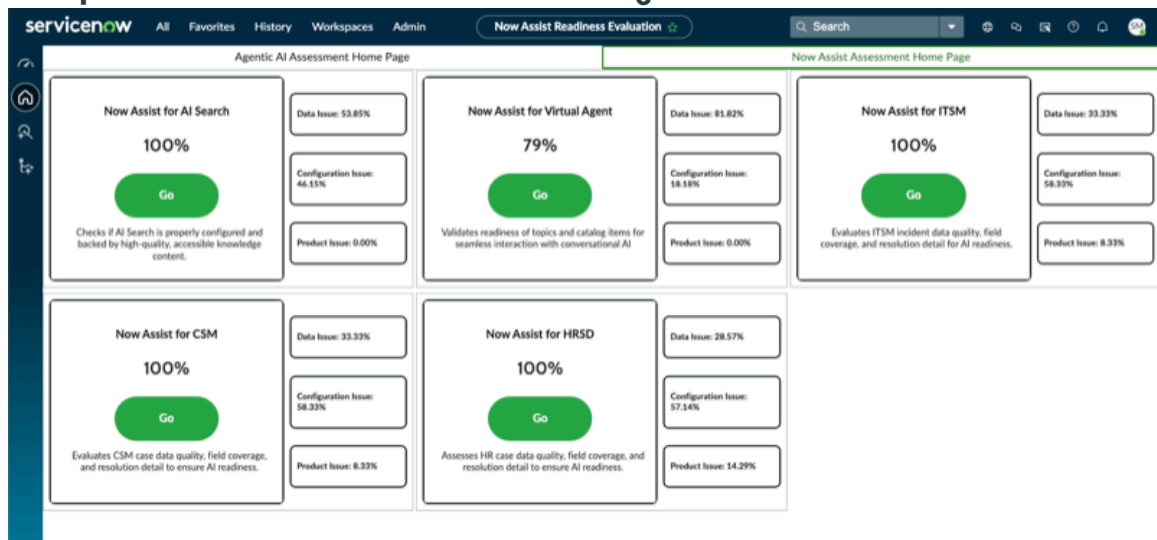
### Example of the Agentic AI Assessment Home Page



The Now Assist Assessment Home Page tab details the go or no-go status for Now Assist products and features, including:

- 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.

**Example of Legends estimated remediation effort sizing**

Estimated Remediation Effort	Time (in days)
Small	1-5
Medium	6-10
Large	11-30
XL	31-90

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](#).

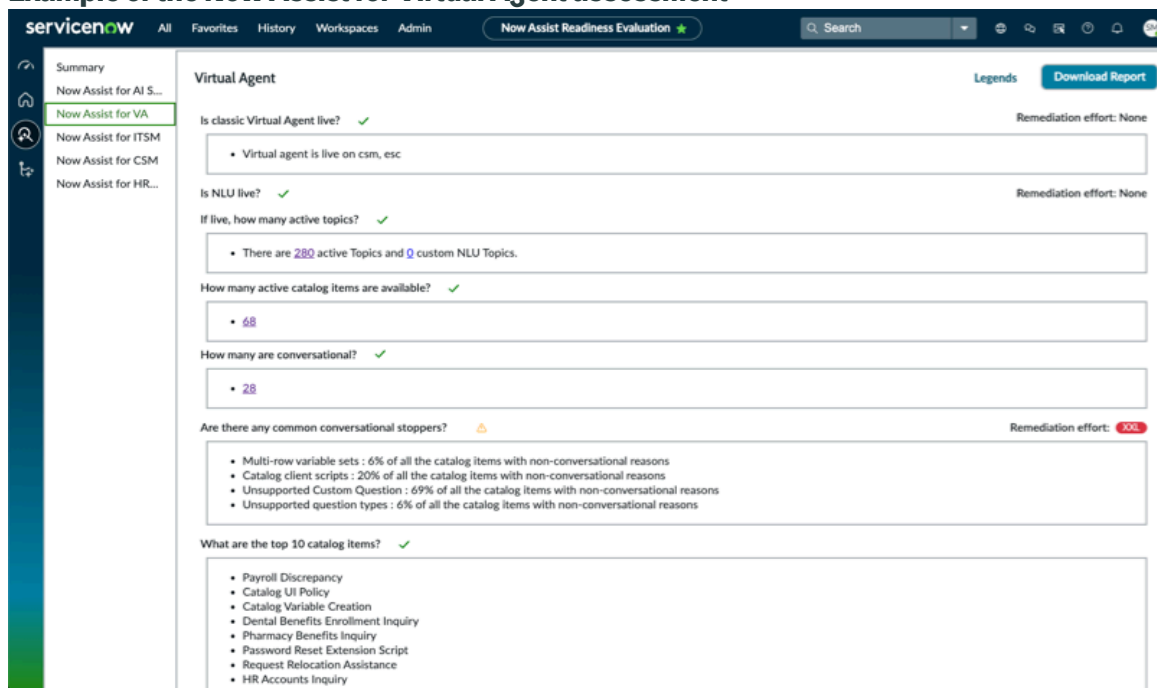
The assessments include visual indicators to guide you on non-blocker and blocker results. The following list details the descriptions of the assessment's visual indicators:

- ✓ : Non-blocker area and no further action needed.
- ✗ : Blocker area and further action is needed to resolve issue before implementation.
- ⚠ : Non-blocker area warning, but the issue should be evaluated before implementation.

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).

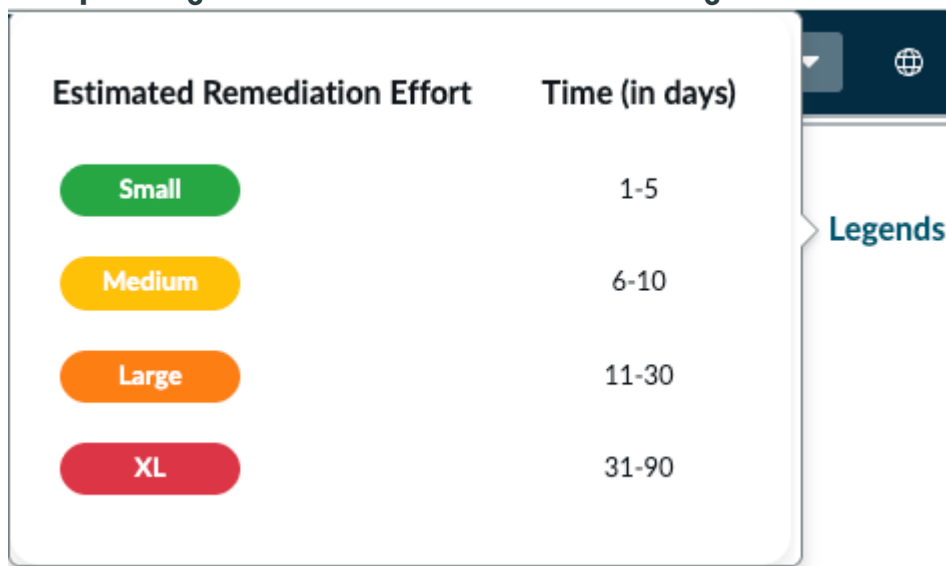
Navigate to **Workspaces > Now Assist Readiness Evaluation** and select the agentic\_AI assessment icon  to review your agentic AI assessment.

**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.

#### Example of Legends estimated remediation effort sizing



Estimated Remediation Effort	Time (in days)
Small	1-5
Medium	6-10
Large	11-30
XL	31-90

After reviewing the summary, select the individual assessment tabs to review those specific assessments. For the Agentic AI - ITSM tab, you must first choose a use case from the **Select Use Case** drop-down menu: **Categorize Incidents** or **Generate change request plans**. The CSM page doesn't have any specific use cases to select. The questions for each assessment differ. For example, the questions and answers on the **Agentic AI - ITSM > Categorize Incidents** assessment are different than the questions and answers on the **Agentic AI - CSM** 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](#).

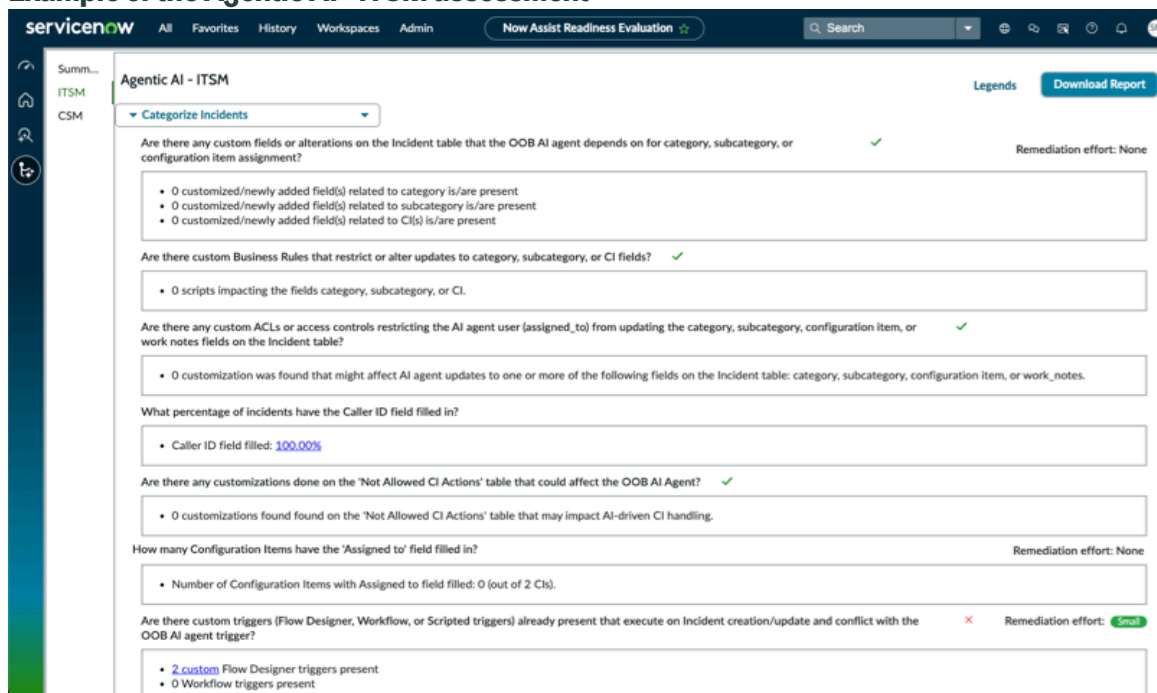
The assessments include visual indicators to guide you on non-blocker and blocker results. The following list details the descriptions of the assessment's visual indicators:

- ✓ : Non-blocker area and no further action needed.
- ✗ : Blocker area and further action is needed to resolve issue before implementation.
- ⚠ : Non-blocker area warning, but the issue should be evaluated before implementation.

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.

**i 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



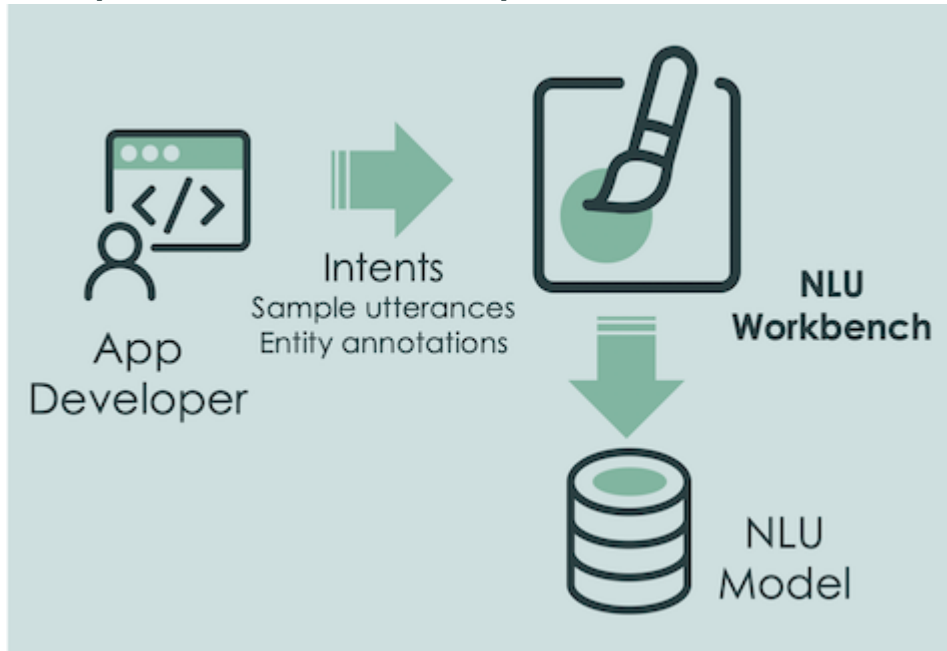
After reviewing the assessment and making the suggested assessment changes, rerun your scheduled jobs to see how your results improve. For more information on rerunning a specific job or all jobs, see [Run the GenAI/AgenticAI Assessment scheduled job](#).

## 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.

## Overview of NLU

### 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

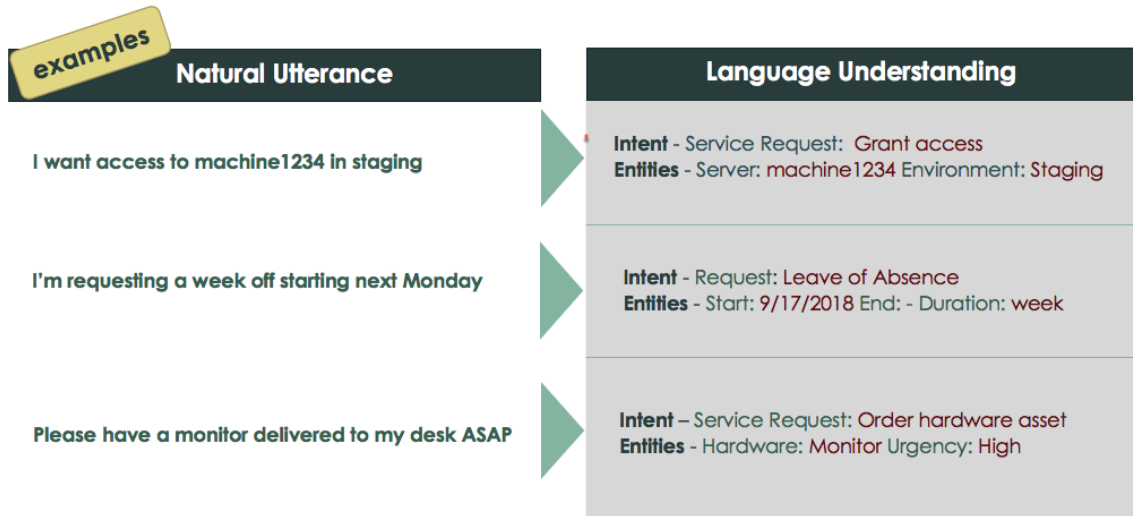
Vocabulary is used to define or overwrite word meanings. For example, you can assign the synonym "Microsoft" to the acronym "MS".

### 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

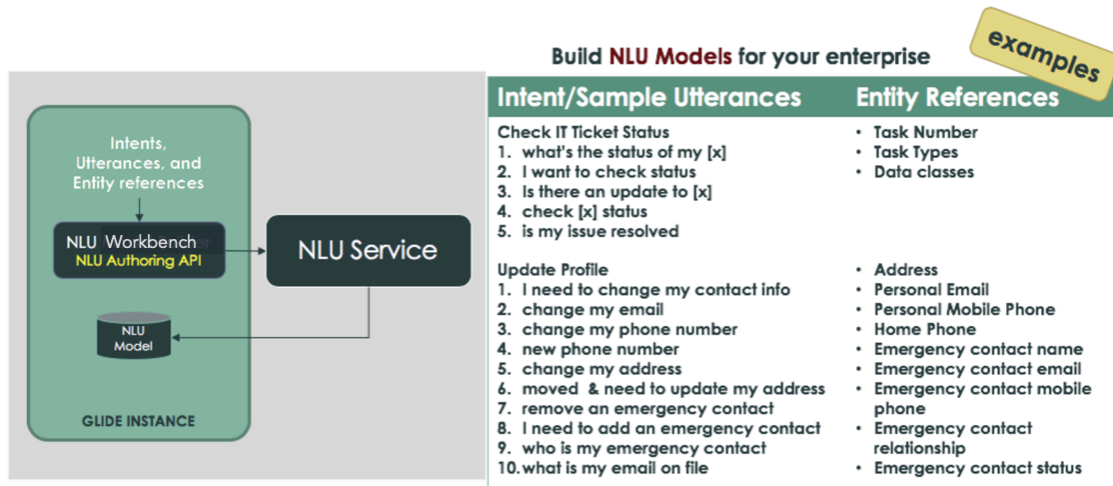


### NLU Workbench

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

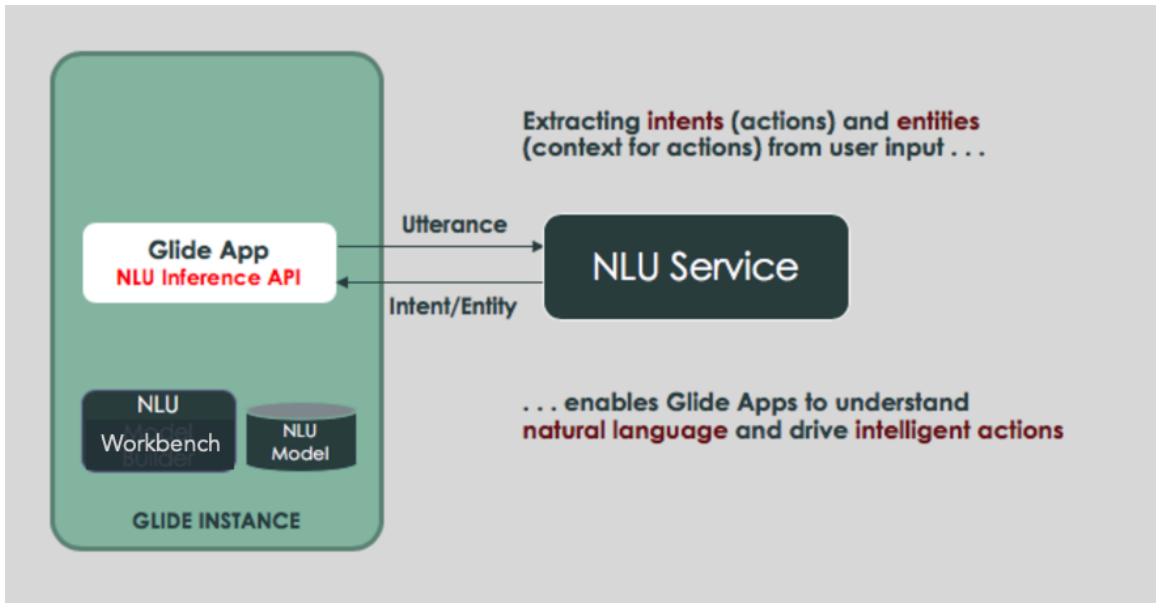


For information on how to build and use an NLU model, see: [Create an NLU model](#).

### 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

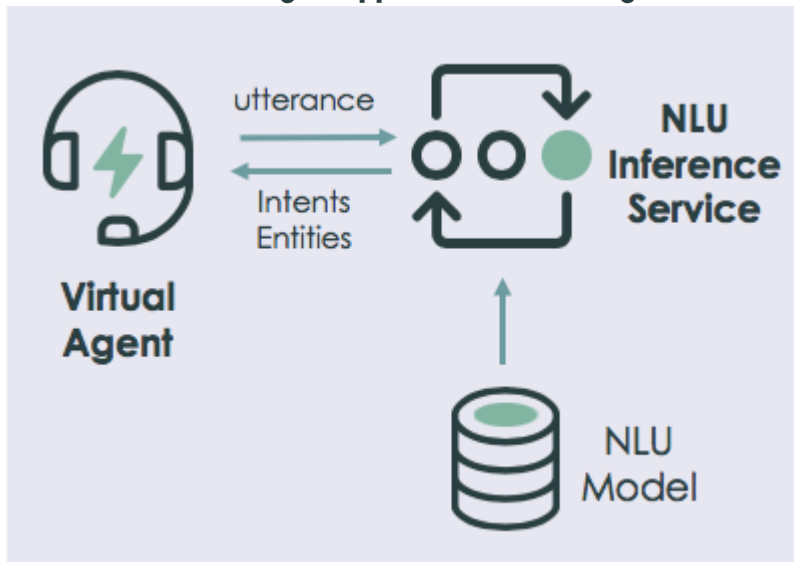


Here, the system uses an inference API to train NLU algorithms by using sample record data to identify intents and entities that are strong candidates for accurate prediction.

### NLU model consumption



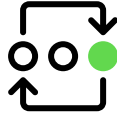

Other ServiceNow<sup>®</sup> applications consume NLU model output, such as Virtual Agent.

#### Overview of Virtual Agent application consuming NLU



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

<p>Explore</p>  <p>Learn about NLU concepts and features.</p>	<p>Use</p>  <p>Create, test, translate, and publish your NLU models.</p>	<p>Integrate</p>  <p>Learn how Virtual Agent administrators can update NLU models from within Virtual Agent Designer.</p>
	<p>Reference</p>  <p>Learn about using models in different languages for use in other applications</p>	

## 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

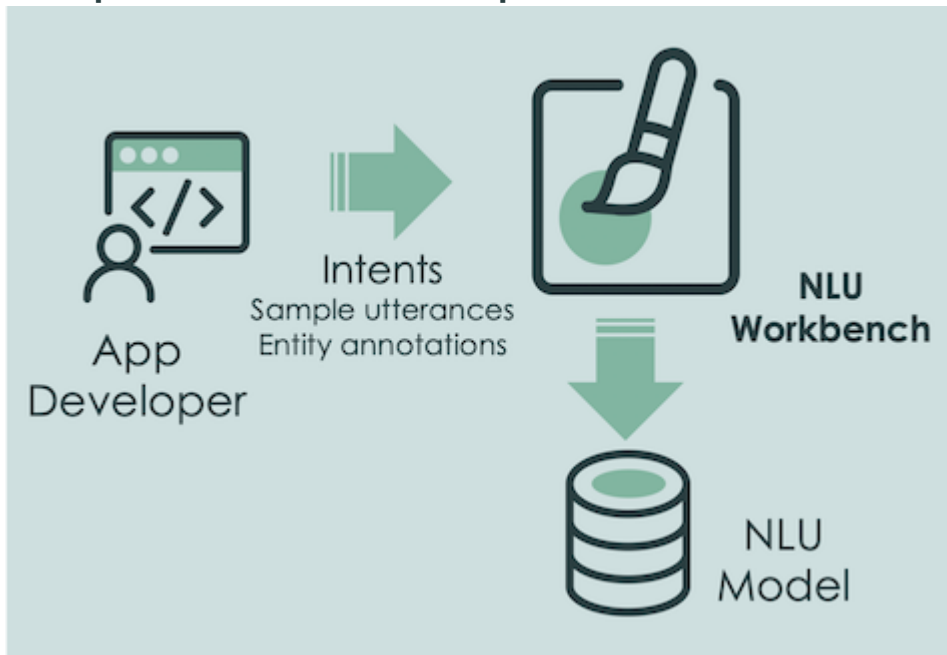
ServiceNow® Natural Language Understanding (NLU) provides an NLU Workbench and an NLU inference service that you can use to enable the system to learn and respond to human-expressed intent. By entering natural language examples into the system, you help it understand word meanings and contexts so it 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

**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.

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 <a href="#">NLU Expert Feedback Loop</a> 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
<a href="#">Intent Discovery</a> 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 <a href="#">Install Intent Discovery</a> .

Plugin	Description
<a href="#">NLU Workbench - Advanced Features</a> sn_nlu_workbench	Delivers the Model Performance, Multi-model Batch Testing, Cross-model Conflict Review, and Expert Feedback Loop features.  For more information, see <a href="#">Install NLU Workbench - Advanced Features</a> .

### 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.

**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**

**i 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**

**i 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.

## 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

**i 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	Recommendation usage
Maximum number of utterances per intent <i>glide.nlu.utterances_per_intent.value_limit</i>	200	NLU Workbench	Use fewer than 200 utterances per intent to keep your model well balanced in terms of intent size.

**Model Settings (continued)**

Label and Name	Default value	Plugin	Recommended usage
			<p><b>i Note</b> Value must be greater than 5 and less than or equal to 300.</p>
Maximum number of records in a Table vocabulary source <i>glide.platform.ml.api.max_nlu_lookupsource_records</i>	100,000	NLU Workbench	Keep the value under 100,000.
Maximum number of values in a List vocabulary source <i>glide.nlu.static_lookup.value_limit</i>	1,000	NLU Workbench	Keep the value under 1,000.
Enable pre-built vocabulary for software names <i>glide.mlpredictor.option.nlu.@LookupSources:software</i>	enabled	NLU Workbench	Enable pre-built vocabulary so the system can recognize software names.
Enable pre-built vocabulary for hardware names <i>glide.mlpredictor.option.nlu.@LookupSources:hardware</i>	enabled	NLU Workbench	Enable pre-built vocabulary so the system can recognize hardware names.

**Advanced Settings**

Label and Name	Default value	Plugin	Recommended usage
Maximum number of records for Intent Discovery classification <i>sn_nlu_discovery.intent_discovery_max_classification_limit</i>	300,000	Intent Discovery	Keep the number of records less than 500,000.
Minimum number of records for Intent Discovery classification <i>sn_nlu_discovery.intent_discovery_min_classification_limit</i>	10,000	Intent Discovery	Use at least 10,000 records to get high quality results.

Advanced Settings (continued)

Label and Name	Default value	Plugin	Recommended usage
Minimum number of records for NLU performance analysis <i>sn_nlu_workbench.glide.nlu.performance_min_clustering_records</i>	5,000	NLU Workbench - Advanced Features	Use at least 5,000 records to get high quality results.
NLU Conflict Detection - Moderate Threshold <i>sn_nlu_workbench.glide.nlu.conflict.moderate_threshold</i>	.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 <i>sn_nlu_workbench.glide.nlu.conflict.critical_threshold</i>	.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 <i>sn_nlu_workbench.glide.nlu.batch_test_max_import_rows</i>	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 <i>glide.mlpredictor.option.nlu.activeLearning.label_candidate_table_max_responses</i>	300	NLU Workbench - Advanced Features	Pull no more than 300 utterances from your users' Virtual Agent chat logs to display in the Expert Feedback Loop application. The minimum number of utterances a user should review before tuning the model
The minimum number of utterances a user should review before tuning the model <i>sn_nlu_workbench.glide.nlu.optimize.min_labeled_data</i>	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 <i>glide.mlpredictor.option.nlu.activeLearning.va_chat_log_row_limit - 3000</i>	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.

**Advanced Settings (continued)**

Label and Name	Default value	Plugin	Recommended usage
Size limit on Label Candidate Table (used for pruning the table) <i>glide.mlpredictor.option.nlu.activeLearning_label_candidate_table.max_data_size - 10000</i>	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) <i>glide.mlpredictor.option.nlu.activeLearning_label_table.max_data_size - 10000</i>	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. <i>glide.mlpredictor.scheduled.nlu.model.training</i>	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)
<ul style="list-style-type: none"> <li>• Brazilian Portuguese</li> <li>• Chinese (simplified)</li> <li>• Danish</li> <li>• Dutch</li> <li>• English</li> <li>• Finnish</li> <li>• French</li> <li>• French Canadian</li> </ul>	<ul style="list-style-type: none"> <li>• Brazilian Portuguese</li> <li>• Dutch</li> <li>• English</li> <li>• French</li> <li>• French Canadian</li> <li>• German</li> <li>• Italian</li> <li>• Japanese</li> </ul>	<ul style="list-style-type: none"> <li>• Chinese (simplified)</li> <li>• Danish</li> <li>• Finnish</li> <li>• Korean</li> <li>• Norwegian</li> <li>• Polish</li> <li>• Swedish</li> </ul>

Available in Tokyo, Utah, and Vancouver	Intent and entity support (in Utah and Vancouver)	Intent-only support (in Utah and Vancouver)
<ul style="list-style-type: none"> <li>• German</li> <li>• Italian</li> <li>• Japanese</li> <li>• Korean</li> <li>• Norwegian</li> <li>• Polish</li> <li>• Portuguese</li> <li>• Spanish</li> <li>• Swedish</li> </ul>	<ul style="list-style-type: none"> <li>• Portuguese</li> <li>• Spanish</li> </ul>	

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.

## 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.

### 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.

## NLU Workbench

Create, manage, and tune your NLU models to better understand what your users are saying. [Learn more](#)

Virtual Agent (VA)   Issue Auto Resolution (IAR)   AI Search

---

### Virtual Agent models

+ Create new model

Show Prebuilt Models
All Languages ▾

Model	Status	Used In VA	Model Type	Enabled Intents	Mapped Intents
ITSM model for Virtual Agent (English) <span>↗</span>	Draft Saved	No	Custom	6	0

⏪ ⏩ 1 ⏪ ⏩

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

displayed. For more information on model languages and grouping, see [Multilingual model](#)

## Virtual Agent models

[+ Create new model](#)

Search  Show Prebuilt Models  All languages ▼


>	Model	Status	Used In VA	Model Type	Enabled Intents	Mapped Intents	Last Modified
▼	ITSM model for Virtual Agent 2						
	English(Primary)	Draft Saved	No		6	0	2023-06-22 12:09:54
	Brazilian Portuguese	Draft Saved	No		6	0	2023-06-24 14:22:21

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


### Boost your model performance



**Tune your model**

Provide and gather feedback on real end-user chats from your org's subject matter experts.

[Go to Expert Feedback Loop](#)



**Discover new intents**

Analyze your users chat and task data to identify which intents you should create next.

[Go to Intent Discovery](#)

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.

## 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.

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.



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.

## NLU Workbench

Create, manage, and tune your NLU models to better understand what your users are saying. [Learn more](#)

[Virtual Agent \(VA\)](#)   [Issue Auto Resolution \(IAR\)](#)   [AI Search](#)

### Virtual Agent models

+ Create new model

Search   Show Prebuilt Models All Languages ▾

Model	Status	Used In VA	Model Type	Enabled Intents	Mapped Intents
<a href="#">ITSM model for Virtual Agent (English)</a>	Draft Saved	No	Custom	6	0

|< < 1 > >|

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.

## Model management phases

Use these iterative phases to refine and improve your model over time.




### Build and train your model

View phase

Improve intents by listing the various ways users might express the same need.

---

- Add or remove utterances
- Configure entities to capture important context
- Build and maintain your model's default test set
- Provide synonyms for company specific words and terms



### Test and publish your model


View phase

Test your model's performance before publishing.

---

- Assess the model quality using default test set
- Identify which intents need improvement
- Review model predictions and publish

---



### Tune your model

Tune Model

Provide and gather feedback to improve your model's performance.

---

- You have utterances that need feedback
- Give feedback on predictions made on end user utterances in Virtual Agent
- Improve model quality by adding feedback to the model and default test set

## 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](#).

## 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.

## NLU Workbench

Create, manage, and tune your NLU models to better understand what your users are saying. [Learn more](#)

Virtual Agent (VA)   Issue Auto Resolution (IAR)   AI Search

---

### Virtual Agent models

+ Create new model

Show Prebuilt Models
 All Languages ▾

Model	Status	Used In VA	Model Type	Enabled Intents	Mapped Intents
<a href="#">ITSM model for Virtual Agent (English)</a>	Draft Saved	No	Custom	6	0

⏪ ⏩ **1** ⏪ ⏩

## 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

How do you want to create your model? ✕



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.

## 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**.

## START FROM BLANK

### First, let's fill out some model details

Add a unique name and description to describe the model.

Name \*

HR Model for Virtual Agent

---

Short description

Natural language for Human Resources user requests

---

What is the primary language? ⓘ

Select... ▼

What are you creating this for? ⓘ

Virtual Agent ▼

What business area is this for? (optional) ⓘ

-- Not specified -- ▼

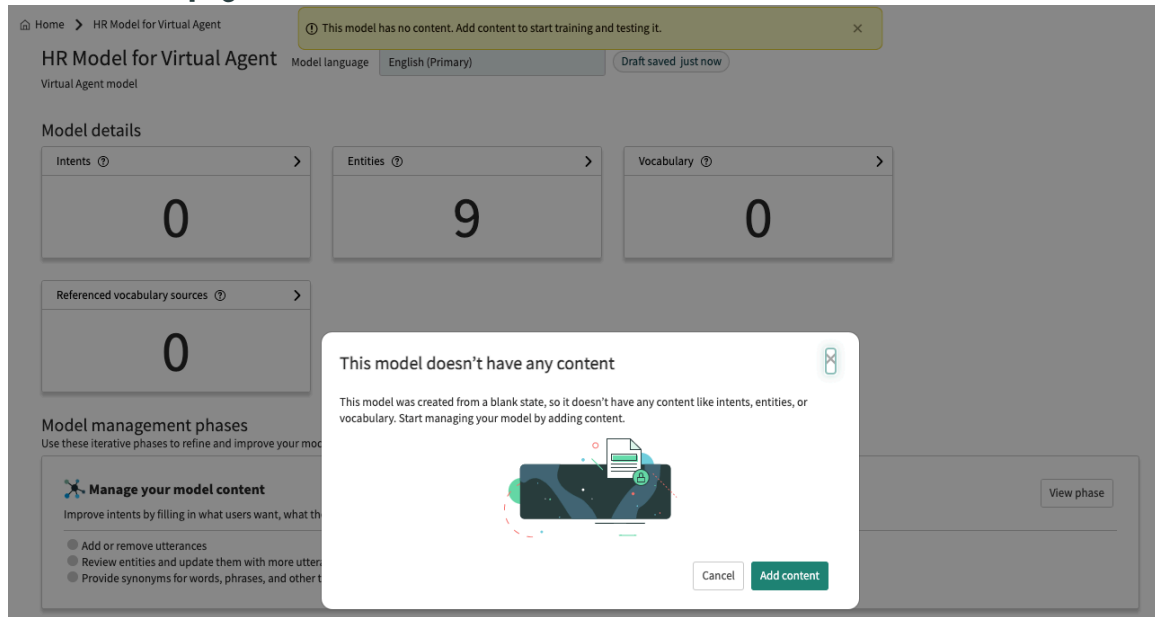
Cancel
Next

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.

## 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.

### 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.



### IMPORT DATA FROM CSV

## First, let's fill out some model details

Add a unique name and description to describe the model.

Name \*

Calendar Model

---

Short description

Model for answering and performing calendar requests

---

What is the primary language? ⓘ

English - en

What are you creating this for? ⓘ

Virtual Agent

What business area is this for? (optional) ⓘ

-- Not specified --

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**.




### IMPORT DATA FROM CSV

## Now, select which file you want to use

We'll take the utterances and intents in this CSV file and upload them in the next step. Make sure the file is in the right format before importing it. Otherwise, the data won't transfer correctly.

Intent	Utterance
██████████ ████	██████████ █████
██████████ ████	██████████ ██████████



10. Choose the CSV or XLSX (Excel Workbook) file from the pop-up.

**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.



## CREATE USING PREBUILT MODEL

### First, let's fill out some model details

Add a unique name and description to describe the model.

Name \*

Human Resources VA Model

---

Short description

Virtual Agent Model for responding to HR requests

---

What is the primary language? ⓘ

English - en

What are you creating this for? ⓘ

Virtual Agent

What business area is this for? (optional) ⓘ

HR x

Cancel Next

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.



## CREATE USING PREBUILT MODEL

### First, let's decide which prebuilt model to use

Prebuilt models come with NLU Workbench. They contain intents and utterances you can use to create the foundation of your new model.

Prebuilt model

HR NLU for VA

Cancel Next

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.



## CREATE USING PREBUILT MODEL

### Pick which intents to use

These are the intents that are included in the prebuilt model you selected. Select the ones that are most relevant to the model you're creating, and we'll use them (plus their associated data) to quickly create the foundation for you new model.

Search intents

<input type="checkbox"/> Intent ↑	Utterances	Entities
<input checked="" type="checkbox"/> AddEmergencyContact	28	4
<input type="checkbox"/> AskaQuestion	3	0
<input checked="" type="checkbox"/> Benefits	81	0
<input checked="" type="checkbox"/> CreateHRGeneralInquiryCase	92	0
<input checked="" type="checkbox"/> DeleteEmergencyContact	31	1
<input checked="" type="checkbox"/> GeneralHRInquiry	92	0
<input checked="" type="checkbox"/> NewHireOrientation	60	0
<input checked="" type="checkbox"/> PayDiscrepancy	59	0
<input checked="" type="checkbox"/> RequestForLeave	61	0
<input checked="" type="checkbox"/> UpdateAddress	29	0
<input checked="" type="checkbox"/> UpdateEmail	34	1

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.

### Before you begin

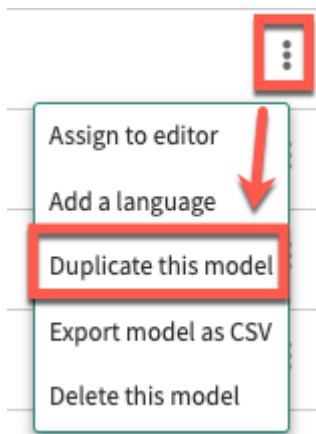
- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Identify an existing NLU model to copy from. The source model should have the language and purpose you want for the new model. If a suitable model does not exist, create and train one.
- Role required: nlu\_admin or admin.

### About this task

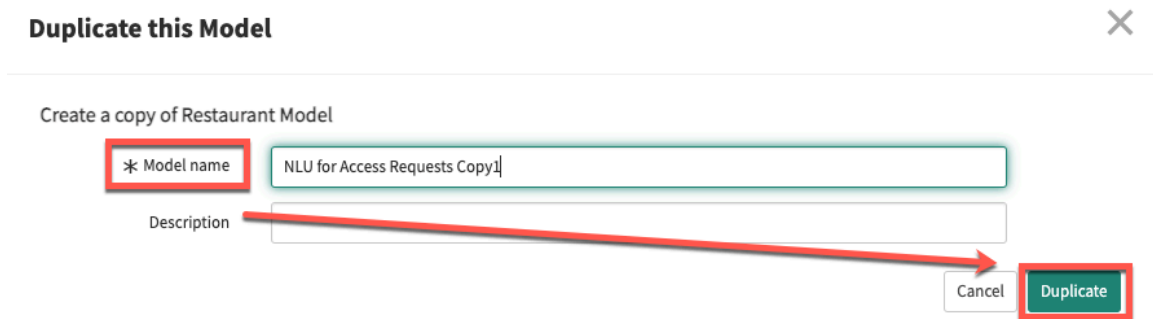
In this example procedure, you've already created a model that's titled NLU for Access Requests. You want to customize it for another department. You duplicate the model to create a new, separate model. To customize the new model, you then add content related to the other department.

### 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.

## 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

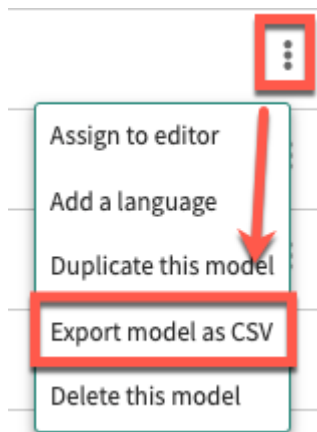
## About this task

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](#).

## 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 [Update set batching](#). The following video provides a demonstration of this process.

### Procedure

1. Select **All** and enter `sys_nlu_model.list` into the navigator.



2. Select the **Model Name** of the model.

Model Name	Display name	Language	Description	Application	State	Created	Created by	Updated	Updated by
global_of3b8e580b203300a1b16c2367673a75_...	ITSM NLU for Virtual Agent FR	fr	French model used by ITSM VA Conversations	ITSM NLU Model for Virtual Agent Convers...	Published	2021-03-30 11:11:56	admin	2022-01-16 09:35:10	admin
global_of3b8e580b203300a1b16c2367673a75_...	ITSM NLU for Virtual Agent DE	de	German model used by ITSM VA Conversations	ITSM NLU Model for Virtual Agent Convers...	Published	2021-03-30 07:06:37	admin	2022-01-16 09:35:10	admin
global_ofb737eac7f70010c59db91703c26027_...	NLU Common Entities FR	fr		NLU Common Model	Draft	2020-03-30 02:26:24	admin	2021-03-19 05:40:43	admin
global_771250608732101001c8de05b7cb0bd1_...	ISC NLU	en		Global	Published	2020-08-19 07:44:01	admin	2021-04-07 07:25:38	admin

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) This record is in the [ITSM NLU Model for Virtual Agent Conversations application](#), but **Global** is the current application [To edit this record click here.](#)

4. In the **Related Links** section, click **Add model to current update set.**

You are editing a record in the ITSM NLU Model for Virtual Agent Conversations application (cancel)

This item is read-only based on its protection policy.

Model Name	global_of3b8e580b203300a1b16c2367673a75_...	Application	ITSM NLU Model for Virtual Agent Ci
Display name	ITSM NLU for Virtual Agent FR	Language	fr
State	Published	Confidence Threshold	0.55
Description	French model used by ITSM VA Conversations		
Created For	virtual_agent		

[Train](#)

Related Links

[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.

Adding model to update set

Failed at 100%

You are attempting to add a record to the system default update set, please create a [New Local Update Set](#) and set that as your current update set.

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.

Name	Application	State	Installed from	Created	Created by	Parent	Batch Base
ITSM model migration - Global Records 7	Global	Complete		2023-06-26 16:23:25	admin	ITSM model migration	ITSM model migration
ITSM model migration - Global Records 9	Global	Complete		2023-06-26 16:23:25	admin	ITSM model migration	ITSM model migration
ITSM model migration - Global Records 8	Global	Complete		2023-06-26 16:23:25	admin	ITSM model migration	ITSM model migration
ITSM model migration - Global Records 4	Global	Complete		2023-06-26 16:23:24	admin	ITSM model migration	ITSM model migration
ITSM model migration - Global Records 5	Global	Complete		2023-06-26 16:23:24	admin	ITSM model migration	ITSM model migration
ITSM model migration - Global Records 6	Global	Complete		2023-06-26 16:23:24	admin	ITSM model migration	ITSM model migration
ITSM model migration - Global Records 2	Global	Complete		2023-06-26 16:23:23	admin	ITSM model migration	ITSM model migration
ITSM model migration - Global Records 3	Global	Complete		2023-06-26 16:23:23	admin	ITSM model migration	ITSM model migration
ITSM model migration - Global Records	Global	Complete		2023-06-26 16:23:20	admin	ITSM model migration	ITSM model migration
ITSM model migration	ITSM NLU Model for Virtual Agent Convers...	Complete		2023-06-26 16:22:19	admin	(empty)	ITSM model migration

**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

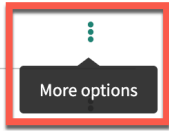
Deleting an NLU model removes the model and its contents, including its default test set, from the system. You cannot delete a model if it meets any of the following criteria:

- 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.

Model	Status	Last Published	Created for	Enabled intents	Mapped intents ⓘ	Total Entities	Editors
<a href="#">HR NLU for VA EFL3 (English)</a>	Trained 3 days ago	2022-12-05 15:40:53	Virtual Agent	22	0	6	
<a href="#">HR NLU for VA EFL2 (English)</a>	updated 25 days ago	2022-11-08 01:17:25	Virtual Agent	22	0	6	

4. Select the option to **Delete this model**.
5. Select the check box to acknowledge that all references to this model will be deleted.

### Delete HR NLU for VA EFL1

This cannot be undone. The model's default test set will be deleted. Please acknowledge the following to delete this model.

All references to this model will be deleted

Cancel

Delete model

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.

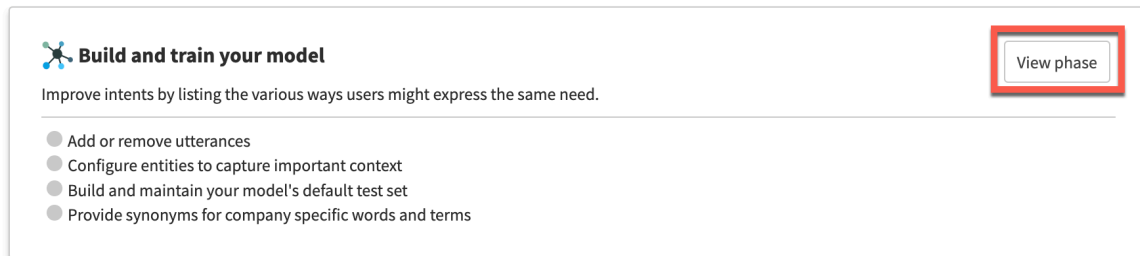
Models are made up of the following content:

- **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**.

### Model management phases

Use these iterative phases to refine and improve your model over time.



**Build and train your model**

Improve intents by listing the various ways users might express the same need.

[View phase](#)

- Add or remove utterances
- Configure entities to capture important context
- Build and maintain your model's default test set
- Provide synonyms for company specific words and terms

## 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](#).

## 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.

Model details

Settings

* Name	<input type="text" value="IT Model"/>
Short description	<input type="text" value="Describe the purpose of this model"/>
Language	<input type="text" value="English - en"/>
Created for <sup>①</sup>	<input type="text" value="Virtual Agent"/>
Business area (optional) <sup>①</sup>	<input type="text" value="-- Not specified --"/>
Confidence Threshold(%)	<input type="text" value="60"/>
Ignore punctuation <sup>①</sup>	<input checked="" type="checkbox"/>

Save

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.

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.

## Build and train your model

Intents (22) Entities (16) Vocabulary (28) Test set (266)

**1**

intent has low utterance count

**15**

intents have low test utterance count

**2**

intents have critical conflicts

Resolve any issues with existing intents, or create new intents to improve the model content. A minimum of 2 intents are required to train a model.

[+ New intent](#) [Import intents](#)

Update Showing: Select...

Name	Utterances	Test utterances	Associated entities	Critical conflicts	Mapped objects	Last updated	Enabled
<a href="#">#UpdateAddress</a>	29	24	0	0	--	2022-11-13 22:21:27	<input checked="" type="checkbox"/>
<a href="#">#UpdateEmail</a>	52	32	1	0	--	2022-11-13 22:21:27	<input checked="" type="checkbox"/>
<a href="#">#UpdateEmergencyContact</a>	61	21	4	0	--	2022-11-13 22:21:27	<input checked="" type="checkbox"/>

[Train model](#) [Try model](#)

**TRAIN MODEL**

✔ Last trained 2 months ago

Content changes

English - en  
0 intent(s)  
0 entitie(s)  
0 vocabulary

[Train](#)

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.

Select an intent to access the intent details page. The **Utterances** tab lists all the utterances currently in the intent.

## #UpdateEmail

Utterances (52)
Associated entities (1)
Conflicts (12)
Settings

Use @ to refer to values (words) in your existing vocabulary source, which you can always add to. To extract specific words from your utterances, select the word to annotate them as entities.

	Utterance	Source <span style="font-size: 0.8em;">?</span>	Last modified <span style="font-size: 0.8em;">↑</span>	
<input type="checkbox"/>	<span style="font-size: 0.8em;">“</span> new@email.com is my new email	Manual	2022-11-13 22:21:57	<span style="font-size: 0.8em;">✎</span> <span style="font-size: 0.8em;">📄</span> <span style="font-size: 0.8em;">⋮</span>
<input type="checkbox"/>	<span style="font-size: 0.8em;">“</span> my email address is wrong	Manual	2022-11-13 22:21:57	<span style="font-size: 0.8em;">✎</span> <span style="font-size: 0.8em;">📄</span> <span style="font-size: 0.8em;">⋮</span>

Perform action on selected rows... ▼

Train model
Try model

**TRAIN MODEL**

✔ Last trained 2 months ago

**Content changes**

English - en  
0 intent(s)  
0 entitie(s)  
0 vocabulary

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.

## #AddEmergencyContact

Utterances (48) **Associated entities (4)** Conflicts (16) Settings

These entities are pieces of info that users provide to help identify the specifics on what they need.

Name ↓	Type	Associated Intents	Utterances	Intent match
<a href="#">relationship</a>	Simple	1	8	Select...▼
<a href="#">phonenumber</a>	Pattern	3	14	Select...▼
<a href="#">name</a>	Simple	3	16	Select...▼
<a href="#">email</a>	Pattern	3	9	Select...▼

**Train model** Try model

---

TRAIN MODEL

● Last trained 2 months ago

Content changes

English - en  
0 intent(s)  
0 entity(s)  
0 vocabulary

**Train**

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.

## Build and train your model

[Intents \(22\)](#) [Entities \(16\)](#) [Vocabulary \(28\)](#) [Test set \(266\)](#)

**1**

intent has low utterance count

**15**

intents have low test utterance count

**2**

intents have critical conflicts

Resolve any issues with existing intents, or create new intents to improve the model content.

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.

- 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 uxinfo dev list"
- Utterance B: "I'm mistakenly removed from the arlo-drury-directreports group"
- Result: The system doesn't recognize uxinfo dev or arlo-drury-directreports and can't use these words to predict the intent.
- Solution: Add uxinfo dev 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.

## Creating an intent

### Create an intent



Fill in the properties for this intent

* Intent name	PayDiscrepancy
Description	

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**

There is something wrong with my pay check		
Something is wrong with my paycheck		
salary discrepancy		
report paycheck issues		
Report a problem with my payslip		
Report a problem with paycheck		
problem with compensation		
please investigate pay discrepancy		
please help me with my paycheck		
payroll help		
paycheck wrong		
need to change bank information		

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](#).

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
<ul style="list-style-type: none"> <li>Date</li> <li>Time</li> <li>Duration</li> <li>Location</li> <li>Phone number</li> <li>Phone extension</li> <li>Email</li> <li>Email domain</li> <li>Website address</li> <li>Computer drive</li> <li>Executable file</li> <li>IP address</li> <li>Bytes</li> <li>Proper Names</li> </ul>	<ul style="list-style-type: none"> <li>Task number</li> <li>Incident number</li> <li>Change number</li> <li>Case number</li> <li>KB article number</li> <li>Problem number</li> </ul>	<ul style="list-style-type: none"> <li>Standard phrases</li> <li>List – ServiceNow Table: Field Examples                             <ul style="list-style-type: none"> <li>Change request: Type</li> <li>Task: Impact</li> <li>Task: Urgency</li> </ul> </li> <li>List – user defined</li> <li>Regular expressions</li> </ul>

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.

EMPLOYEE EXPERIENCE PORTAL		CUSTOMER SERVICE PORTAL
ITSM	HR	CSM
<ol style="list-style-type: none"> <li>Check IT ticket status</li> <li>Open IT ticket</li> <li>Order an item</li> <li>Search Knowledge Base</li> <li>Process approval</li> <li>My Assigned Equipment</li> <li>Reset password</li> <li>Update assigned task</li> <li>Create change request</li> <li>Escalate Ticket</li> <li>Resolve Incident</li> <li>Create Problem</li> <li>Update Change Request</li> <li>Identify Next Change Window</li> <li>Identify Scheduled Changes</li> </ol>	<ol style="list-style-type: none"> <li>Guest Wifi Access</li> <li>Reset RSA Token</li> <li>Local Admin Access</li> <li>Repository Access</li> <li>VPN connectivity</li> <li>Slow Computer</li> <li>Walk-up Support</li> <li>Manage Distribution List</li> <li>Meeting Room Issues</li> <li>Printer Issues</li> <li>Email Setup</li> <li>Email Issues</li> <li>Get Password Reset link</li> <li>Service Disruptions</li> </ol>	<ol style="list-style-type: none"> <li>General HR Inquiry</li> <li>Request for Leave</li> <li>Pay Discrepancy</li> <li>Update Address</li> <li>Update Email</li> <li>Update Phone Number</li> <li>Add Emergency Contact</li> <li>Delete Emergency Contact</li> <li>Update Emergency Contact</li> </ol>
<p><b>VA Common Topics</b></p> <ul style="list-style-type: none"> <li>Greetings</li> <li>End Conversations</li> <li>What can you do</li> <li>Live agent</li> <li>Feedback</li> </ul>		

For more NLU and Virtual Agent context, refer to the following product documentation:

- 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
<ul style="list-style-type: none"> <li>• CSM Virtual Agent conversations [com.sn_csm.virtualagent]</li> <li>• Human Resources Scoped App: Virtual Agent Conversations [com.sn_hr_virtual_agent]</li> <li>• ITSM Virtual Agent conversations [com.snc.itsm.virtualagent]</li> </ul>	<p>Enables <a href="#">Pre-built Virtual Agent topics, topic blocks, and ServiceNow NLU models</a> for the Customer Service Management, HR Service Delivery, and IT Service Management applications.</p> <ul style="list-style-type: none"> <li>• Requires the Glide Virtual Agent (com.glide.cs.chatbot) plugin to be activated.</li> <li>• Automatically activates the NLU Workbench (com.snc.nlu_studio) plugin.</li> </ul>

## 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:

- 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**.

Home > HR NLU for VA EFL3 - English (Primary) > Build and train your model updated 2 months ago

### Build and train your model

Intents (22) Entities (16) Vocabulary (28) Test set (266)

**1**  
intent has low utterance count

**15**  
intents have low test utterance count

**2**  
intents have critical conflicts

Resolve any issues with existing intents, or create new intents to improve the model content. A minimum of 2 intents are required to train a model.

+ New intent
Import intents

Showing: Select...

Name	Utterances	Test utterances	Associated entities	Critical conflicts	Mapped objects	Last updated	Enabled
#UpdateAddress	29	24	0	0	--	2022-11-13 22:21:27	<input checked="" type="checkbox"/>

Train model Try model

**TRAIN MODEL**

Last trained 2 months ago

Content changes

English - en  
0 intent(s)  
0 entitie(s)  
0 vocabulary

Train

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.

## Import Intents



We recommend importing in small batches for optimal performance. The underlying entities and utterances will also be imported.

Intent	Model Name ↑	#Utterances	#Entities	Created by	Last updated
<input checked="" type="checkbox"/> Application: ITSM NLU Model for Virtual Agent Conversations (43)					
<input type="checkbox"/> ManageDistributionList	ITSM NLU for Virtual Agent	38	0	admin	2021-08-05 12:07:47
<input type="checkbox"/> IntelligentOpenITTicket (Deprecated)	ITSM NLU for Virtual Agent	29	1	admin	2021-08-05 12:07:47
<input type="checkbox"/> MeetingRoomIssues	ITSM NLU for Virtual Agent	35	1	admin	2021-08-09 08:47:05
<input type="checkbox"/> RemoveOwnerFromGroupDistributionList	ITSM NLU for Virtual Agent	15	0	admin	2021-08-05 12:07:48
<input checked="" type="checkbox"/> OpenITTicket	ITSM NLU for Virtual Agent	29	1	admin	2021-08-05 12:07:48
<input type="checkbox"/> GuestWifiAccess	ITSM NLU for Virtual Agent	42	0	admin	2021-08-05 12:07:47

1 selected
Cancel
Import

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.

## 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 screenshot shows the 'Build and train your model' page for 'HR NLU for VA EFL3 - English (Primary)'. The page has a breadcrumb trail: Home > HR NLU for VA EFL3 - English (Primary) > Build and train your model. Below the breadcrumb is the title 'Build and train your model' with a sub-header 'Intents (22) Entities (16) Vocabulary (28) Test set (266)'. A red box highlights three issue cards:

- 1** intent has low utterance count
- 15** intents have low test utterance count
- 2** intents have critical conflicts

Below the cards, there is a note: 'Resolve any issues with existing intents, or create new intents to improve the model content.'

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:

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 <a href="#">Enable or disable a secondary model intent</a> .

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.

## #GeneralHRInquiry

Train model Try model ⋮

Utterances (92) Associated Entities (0) **Conflicts (3)** Settings

We found overlapping utterances in these intents.  
Remove or edit any utterances so that each utterance only applies to one intent. Any changes you make will automatically update the model and improve its performance.

Hide moderate conflicts

Critical	This intent: GeneralHRInquiry	Conflicting intent: CreateHRGeneralInquiryCase
1	“ help me with this Human Resources question ”	“ help me with this hr question ”
2	“ Human Resources related query ”	“ hr related query ”
3	“ speak to Human Resources ”	“ speak to hr ”

### 7. Point to the utterance that you want to edit or remove.

Critical	This intent: GeneralHRInquiry	Conflicting intent: CreateHRGeneralInquiryCase
1	“ help me with this Human Resources question ”	“ help me with this hr question ”
		✎ 🗑️
		⊖ Ignore

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.

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)

Home > HR NLU for VA TIP - English (Primary) > Build and train your model

## Build and train your model

Intents (22) **Entities (16)** Vocabulary (28) Test set (150)

Choose which system entities to use and review the user defined entities being used in the model. If you want to add new entities, you can annotate them in an individual utterance. We recommend using at least 5 examples for user defined entities.

Import entity

Name	Type	Model Availability	Associated Intents	Created by	Last updated ↓	Enable
System Entities (10)						
User Defined Entities (6)						
<a href="#">whattoedit</a>	Simple		1		2022-10-18 12:01:22	
<a href="#">phonetype</a>	Simple		1		2022-10-18 12:01:22	
<a href="#">name</a>	Simple	✓	3		2022-10-18 12:01:21	

Train model Try model

### TRAIN MODEL

Last trained 4 months ago

### Content changes

English - en  
1 intent(s)  
0 entitie(s)  
0 vocabulary

Train

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](#).

## 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

#### conferenceRoom

Entity (Mapped)

Train model

Try model



Utterances (1)

Settings

Define the properties of the entity

\* Entity Name

Type

Use this when you want to take what the user said and look to a source that it's been mapped to and extract a value for it.

Model Availability

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 <sup>?</sup>  
 Create a list of values and learn alternatives through labeled utterances <sup>?</sup>

\* Vocabulary source

See @conferenceRoom  
 Source fields: Actual value, Alternate values

Mapped \* value for the entity

Save

## 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.

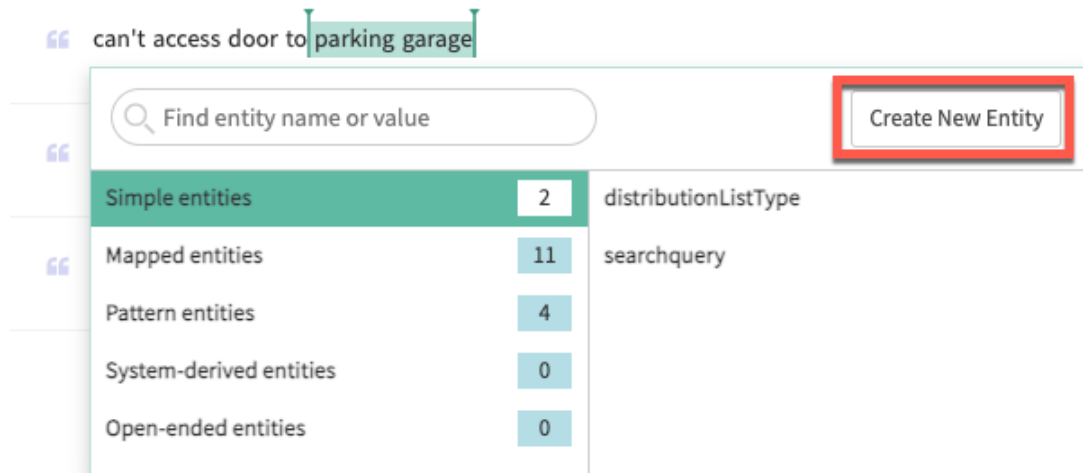
### About this task

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.

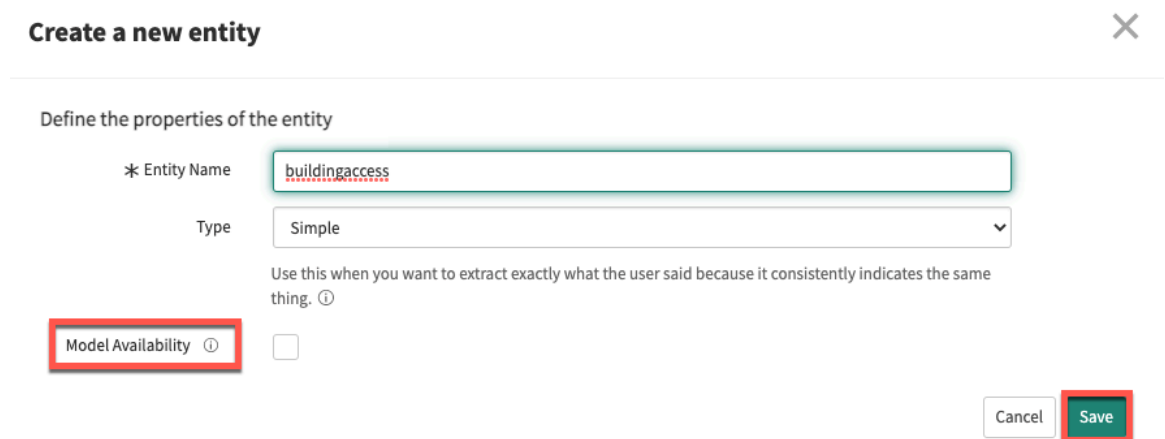
For the following example procedure, you've already created an intent that's titled SubmitAccessRequest and you're creating a simple entity for the type of access the user is requesting.

**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.

## 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.

### Example utterance using a vocabulary source

Use @ to refer to values (words) in your existing vocabulary source, which you can always add to. To extract specific words from your utterances, select the word to annotate them as entities. We recommend adding at least 15 utterances per intent to account for variations in what users say.



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.

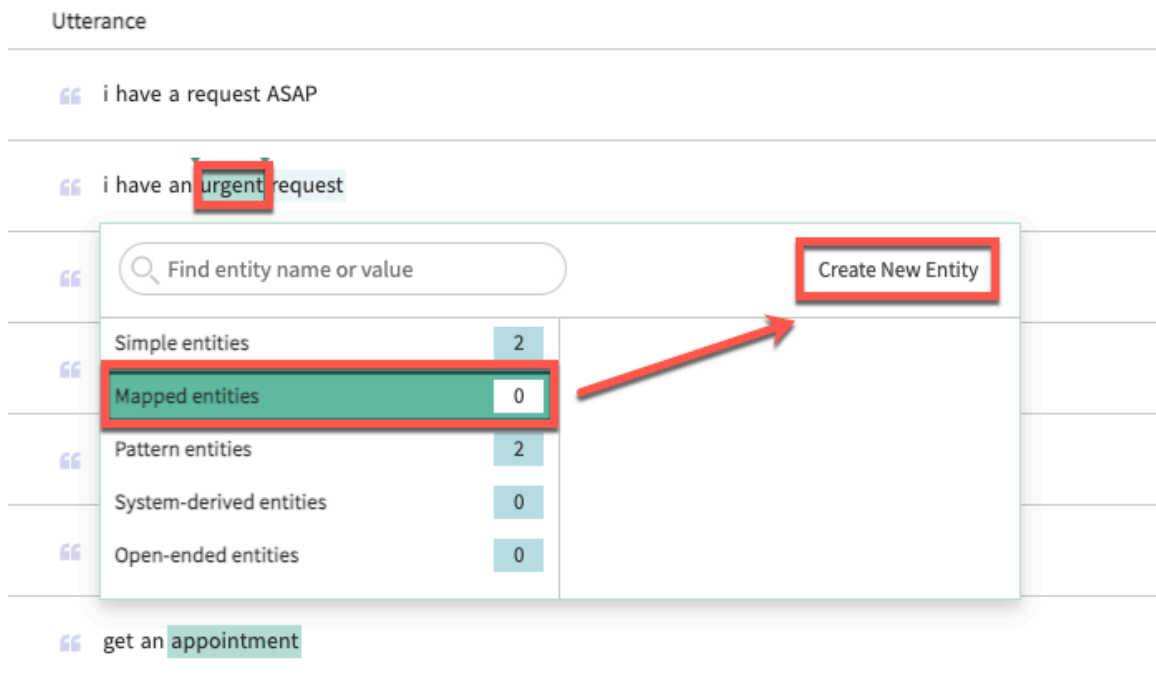
In this example procedure, you create a mapped entity for urgency.

**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.

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.

**Create a new entity** ✕

---

Define the properties of the entity

\* Entity Name

Type

Use this when you want to take what the user said and look to a source that it's been mapped to and extract a value for it.

Model Availability   ⓘ

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 ⓘ

Create a list of values and learn alternatives through labeled utterances ⓘ

Provide values for this entity

**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.

“ i have an urgent request

🔍 Find entity name or value Create New Entity

Simple entities	2	priority	low
Mapped entities	1		medium
Pattern entities	2		high
System-derived entities	0		
Open-ended entities	0		

## 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**.

# #CheckITTicketStatus

Utterances (67) Associated Entities (2) Conflicts (0) Settings

Use @ to refer to values (words) in your existing vocabulary source, which you can always add to. To extract specific words from your utterances, select the word to annotate them as entities.

Add

Utterance	Source	Last modified	
<input type="checkbox"/> how is INC1234567 doing ?	Manual	2022-08-17 10:53:41	<span>✎</span> <span>📄</span> <span>⋮</span>
<input type="checkbox"/> <span style="border: 1px solid #ccc; padding: 2px;">Find entity name or value</span>		22-08-17 53:41	<span>✎</span> <span>📄</span> <span>⋮</span>
Simple entities <span style="float: right;">1</span>	changenumber		
Mapped entities <span style="float: right;">11</span>	email	22-08-17 53:41	<span>✎</span> <span>📄</span> <span>⋮</span>
<b>Pattern entities <span style="float: right;">4</span></b>	meetingID		
System-derived entities <span style="float: right;">0</span>	ticketnumber	22-08-17 53:41	<span>✎</span> <span>📄</span> <span>⋮</span>
Open-ended entities <span style="float: right;">0</span>			
<input type="checkbox"/> I need to check the status of my incident	Manual	2022-08-17 10:53:41	<span>✎</span> <span>📄</span> <span>⋮</span>

**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}`

**Create a new entity** ✕

Define the properties of the entity

\* Entity Name

Type

A Pattern Entity is an entity that matches based on a regular expression that you define. Pattern Entities don't require annotated utterances

Model Availability

\* Regex

[See documentation for Regular Expressions](#)

Cancel
Save

## 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.

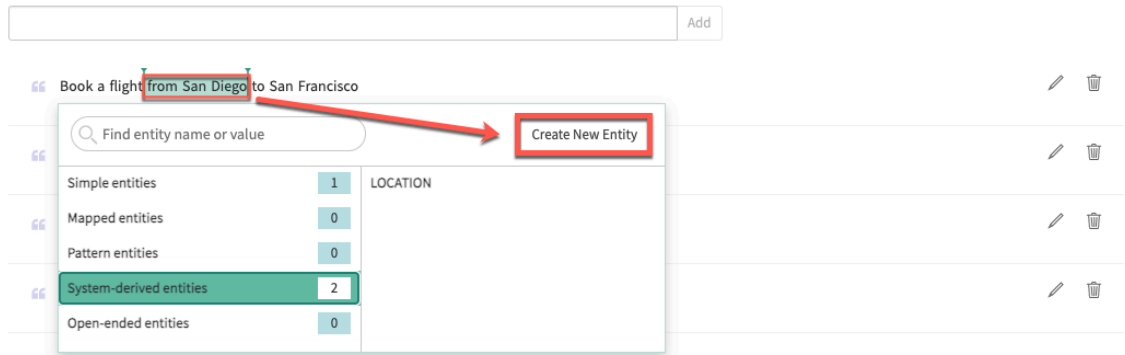
In this example scenario, you have an example utterances book a flight from San Diego to San Francisco. For this step, click **from San Diego**.

# #FlightBooking

Train model Try model

Utterances (4) Associated Entities (2) Conflicts (0) Settings

Use @ to refer to values (words) in your existing vocabulary source, which you can always add to. To extract specific words from your utterances, select the word to annotate them as entities. We recommend adding at least 15 utterances per intent to account for variations in what users say.



**6. Select Create New Entity.**

**7. In the Create a new entity screen, configure the fields.**

For this example, use these configurations.

- o **Entity Name:** Enter `FromLocation`
- o **Type:** Select **System-derived**
- o **Model Availability:** Select this check box if you want this entity to be included in all intents in your model
- o **Parent Entity:** Select **LOCATION**

### Create a new entity ✕

Define the properties of the entity

\* Entity Name

Type

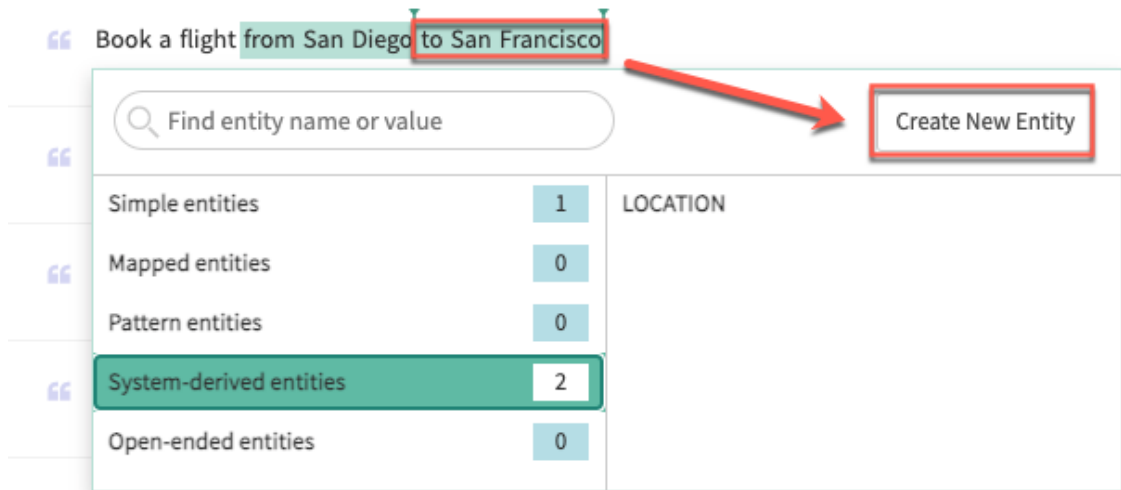
A System-Derived Entity is useful when you want to annotate the same system entity more than once in an utterance. For example, you want to take time off from Monday to Friday so you annotate them as two different entities (Start Date and End Date), where both are derived from Date as their parent system entity

Model Availability

\* Parent Entity

**8. Select Save.**

**9. On the same utterance, select another word or phrase.**  
 In this example scenario, you click **to San Francisco**.



**10. Select Create New Entity.**

For this example, use these configurations.

- o **Entity Name:** Enter `ToLocation`
- o **Type:** Select **System-derived**
- o **Model Availability:** Select this check box if you want this entity to be included in all intents in your model
- o **Parent Entity:** Select **LOCATION**

**Create a new entity** ✕

---

Define the properties of the entity

\* Entity Name

Type

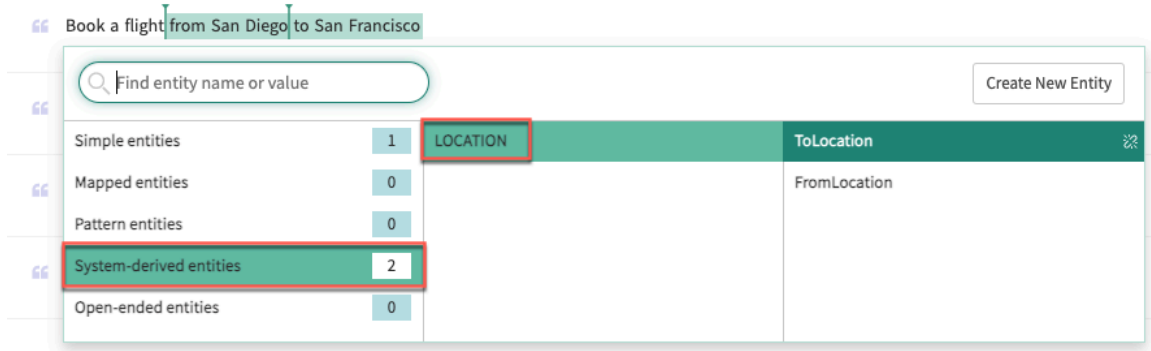
A System-Derived Entity is useful when you want to annotate the same system entity more than once in an utterance. For example, you want to take time off from Monday to Friday so you annotate them as two different entities (Start Date and End Date), where both are derived from Date as their parent system entity

Model Availability ⓘ

\* Parent Entity

**11. Select Save.**

**Result:** Your two system-derived entities appear in the entity window.



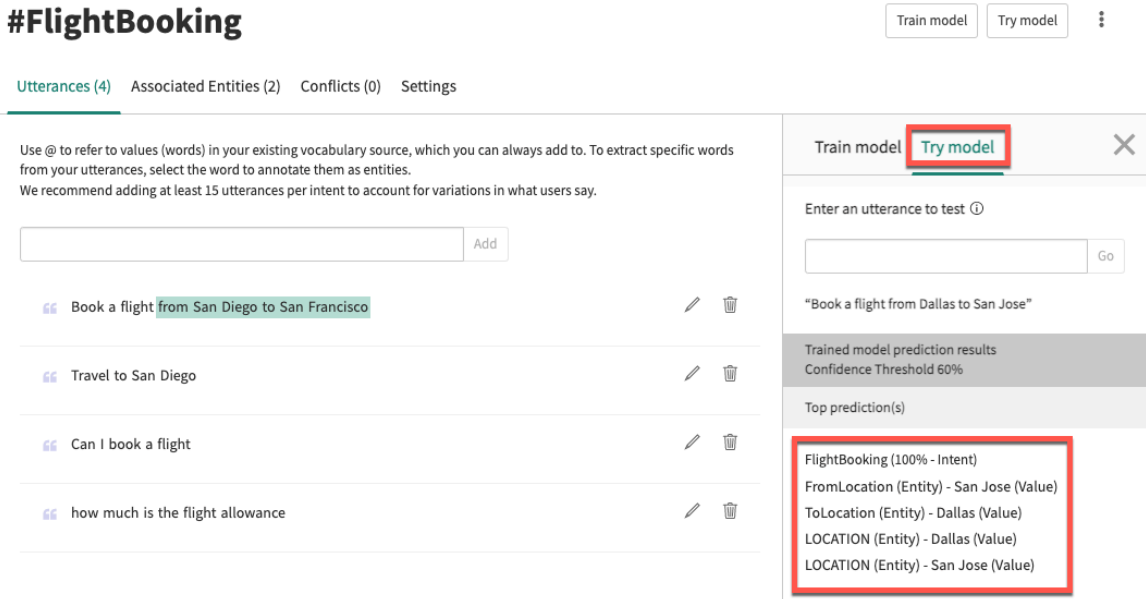
**What to do next**

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**  
**#FlightBooking**



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.

- 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.

## Entity window #OrderMerch

Train model Try model

Utterances (3) Associated Entities (1) Conflicts (0) Settings

Use @ to refer to values (words) in your existing vocabulary source, which you can always add to. To extract specific words from your utterances, select the word to annotate them as entities. We recommend adding at least 15 utterances per intent to account for variations in what users say.

Input field:  Add

Utterances:

- Can I buy some mugs
- How much is the polo shirt
- I want to order a hoodie

Entity selection modal:

Find entity name or value

Create New Entity

Simple entities	0
Mapped entities	0
Pattern entities	0
System-derived entities	0
Open-ended entities	0

### 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**

### Create entity

Create a new entity

Define the properties of the entity

\* Entity Name: merch

Type: Open-ended

When you mark a word as an open-ended entity, the system ignores the entity value and focuses on the context to better infer the intent.

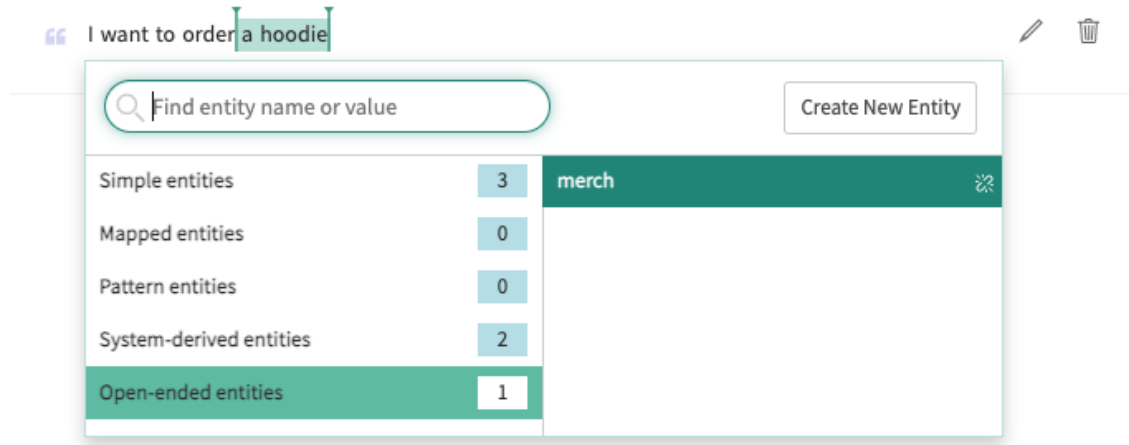
"I want to order an iPhone" → #OrderCatalogItem, CatalogItem

Context Entity value Intent Entity

Cancel Save

### 8. Select **Save**.

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.



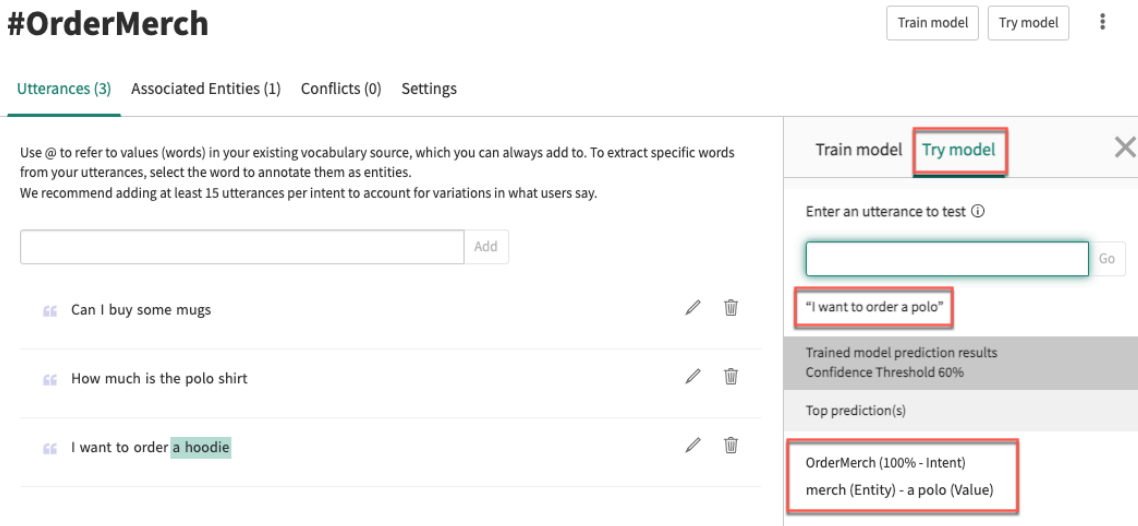
### What to do next

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**.

### #OrderMerch



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.

### 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**.

## Build and train your model

Intents (38) **Entities (38)** Vocabulary (36) Test set (12)

Choose which system entities to use and review the user defined entities being used in the model. If you want to add new entities, you can annotate them in an individual utterance.  
We recommend using at least 5 examples for user defined entities.



Name	Type	Model Availability	Associated Intents	Created by	Last updated ↓	Enable
▶ System Entities (10)						
▼ User Defined Entities (28)						

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.

## Import Entity



Please select one or more entities to import. The sample utterances where the entities are annotated are not imported.

Entity	Model Name ↑	Type	Created by	Last updated
▶ Application: Universal Request: NLU Model for Virtual Agent Conversations (1)				
▶ Application: NLU Common Model (7)				
▶ Application: Global (76)				
▼ Application: HR Service Delivery NLU Model for Virtual Agent Conversations (4)				
<input checked="" type="checkbox"/>	email	HR NLU for VA	Pattern	admin 2019-04-08 00:38:46
<input checked="" type="checkbox"/>	relationship	HR NLU for VA	Simple	admin 2019-04-08 00:28:37
<input checked="" type="checkbox"/>	phonenumber	HR NLU for VA	Pattern	admin 2021-06-16 02:37:55
<input checked="" type="checkbox"/>	name	HR NLU for VA	Simple	admin 2019-04-07 22:49:46
▶ Application: ITSM NLU Model for Virtual Agent Conversations (4)				
▶ Application: Customer Service NLU Model for Virtual Agent Conversations (1)				
▶ Application: Admin Experience Framework (1)				

4 selected
Cancel
Import

**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.

Name	Type	Model Availability	Associated Intents	Created by	Last updated ↓ Enable
▶ System Entities (9)					
▼ User Defined Entities (4)					
<a href="#">relationship</a>	Simple	✓	0	admin	2022-01-25 11:33:50
<a href="#">name</a>	Simple	✓	0	admin	2022-01-25 11:33:50
<a href="#">email</a>	Pattern	✓	0	admin	2022-01-25 11:33:50
<a href="#">phonenumber</a>	Pattern	✓	0	admin	2022-01-25 11:33:50

### 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](#)

## 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.

✕

**Create a new entity**

---

Define the properties of the entity

\* Entity Name

Type

A Pattern Entity is an entity that matches based on a regular expression that you define. Pattern Entities don't require annotated utterances

Model Availability

\* Regex

[See documentation for Regular Expressions](#)

### Case number

Case tasks use the format CS1234567. The regex code is `CS\d{7}` where CS = case record and `d{7}` = 7 digits.

### Create a new entity



Define the properties of the entity

\* Entity Name

Type

A Pattern Entity is an entity that matches based on a regular expression that you define. Pattern Entities don't require annotated utterances

Model Availability

\* Regex

[See documentation for Regular Expressions](#)

Cancel Save

### 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.

### Create a new entity

Define the properties of the entity

\* Entity Name

Type

A Pattern Entity is an entity that matches based on a regular expression that you define. Pattern Entities don't require annotated utterances

Model Availability

\* Regex

[See documentation for Regular Expressions](#)

Cancel Save

### 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.

## Create a new entity



Define the properties of the entity

* Entity Name	<input type="text" value="usphonenumber"/>
Type	<input type="text" value="Pattern"/>
A Pattern Entity is an entity that matches based on a regular expression that you define. Pattern Entities don't require annotated utterances	
Model Availability ⓘ	<input type="checkbox"/>
* Regex	<input type="text" value="\d{10} (?:\d{3}-)\d{4} (\d{3})\d{3}-?\d{4}"/>

[See documentation for Regular Expressions](#)

Cancel

Save

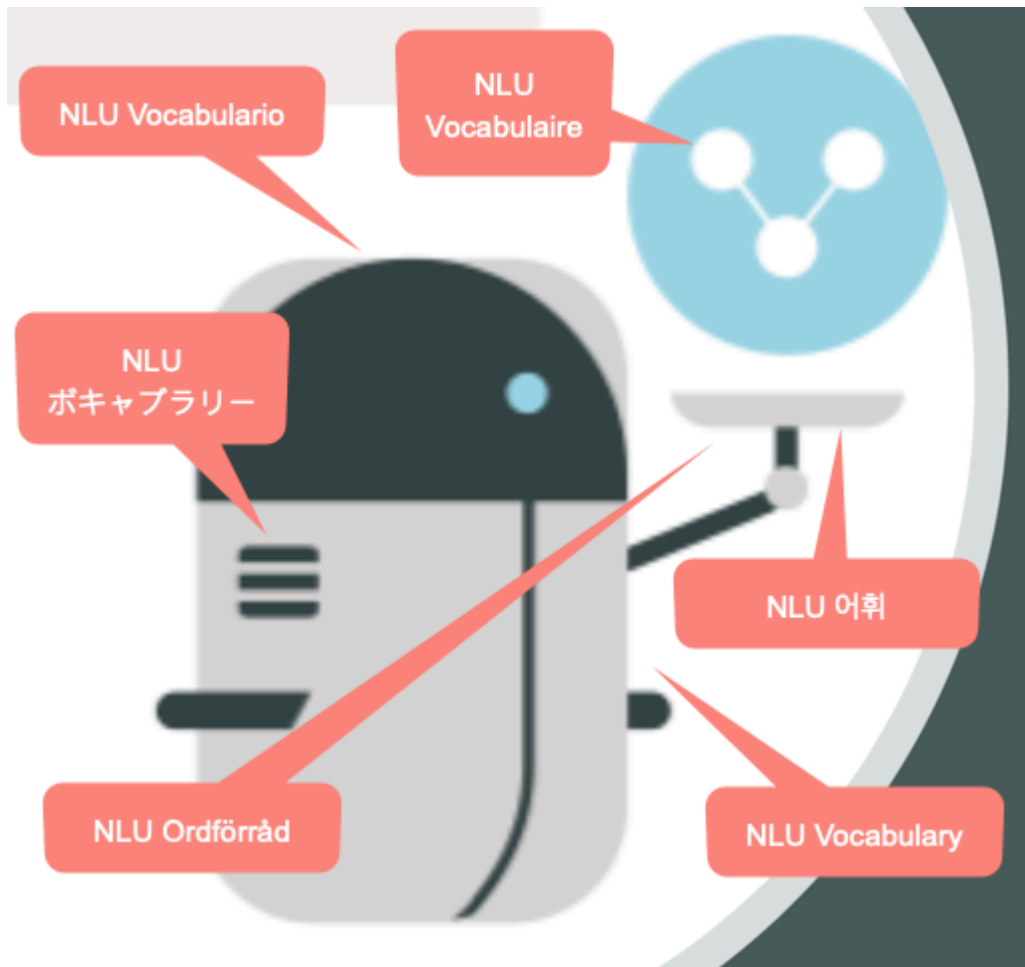
## Regex resources

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.




## 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.

## Build and train your model

Intents (2) Entities (10) **Vocabulary (3)** Test set


Vocabulary provides your model with a synonym for each word or phrase that is unique to your company or business area. If an utterance contains the vocabulary, the model will instead use the synonym when predicting an intent. Use a single word or short phrase as the synonym. For providing a synonym for a list of values, [go to vocabulary sources](#).

 Add a vocabulary

Vocabulary  	Synonym 	Type 		
\b(?:)pwd\b	password	Pattern		
\b(?:)AD\b	active directory	Pattern		
shareholder	stockholder	Regular		

**Train model** Try model

TRAIN MODEL

 Last trained 6 days ago

Content changes

English - en  
0 intent(s)  
0 entitie(s)  
3 vocabulary

Train

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 <a href="#">Create a regular vocabulary item</a> .
Pattern	A regular expression (regex) that can capture specific formats such as email addresses and phone numbers.  See <a href="#">Create a pattern vocabulary item</a> .

### 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.

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"	<ul style="list-style-type: none"> <li>• Issue: the system doesn't recognize the acronym sfcrm, so can't accurately predict the intent.</li> <li>• Solution: Add sfcrm as a regular vocabulary item, and provide a synonym such as CRM software.</li> </ul>
"I need to install Word"	<ul style="list-style-type: none"> <li>• 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.</li> <li>• 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.</li> </ul> <p><b>Note:</b> To extract the specific software name to use in a Virtual Agent conversation topic, annotate it as an entity in the utterance.</p>

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.

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**.


## Build and train your model

Intents (22) Entities (16) **Vocabulary (28)** Test set (610)

Vocabulary provides your model with a synonym for each word or phrase that is unique to your company or business area.

If an utterance contains the vocabulary, the model will instead use the synonym when predicting an intent.

Use a single word or short phrase as the synonym. For providing a synonym for a list of values, [go to vocabulary sources](#).

 Add a vocabulary

**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.

### Add a vocabulary

Enter a word, short phrase, or regular expression as a vocabulary, and provide a synonym.

Note: Regular vocabulary is case insensitive by default. Use pattern vocabulary for case sensitive scenarios.

Type ⓘ	<input type="text" value="Regular"/>
* Vocabulary	<input type="text" value="credentials"/>
* Synonym	<input type="text" value="password"/>

**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.

### 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**.


## Build and train your model

Intents (22)   Entities (16)   **Vocabulary (28)**   Test set (610)

Vocabulary provides your model with a synonym for each word or phrase that is unique to your company or business area.

If an utterance contains the vocabulary, the model will instead use the synonym when predicting an intent.

Use a single word or short phrase as the synonym. For providing a synonym for a list of values, [go to vocabulary sources](#).

 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".

## Add a vocabulary



Enter a word, short phrase, or regular expression as a vocabulary, and provide a synonym.

Note: Regular vocabulary is case insensitive by default. Use pattern vocabulary for case sensitive scenarios.

Type ⓘ	<input type="text" value="Pattern"/>
* Vocabulary (Regex)	<input type="text" value="\b(?:)(mfa( authentication)?)\b"/> <a href="#">See documentation for Regular Expressions</a>
* Synonym	<input type="text" value="multi-factor authentication"/>

## 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**.

## Vocabulary sources

ServiceNow Tables (6) **My lists (7)**

All languages ▾

You can manage your lists of words, names or phrases that can be referenced from NLU models.



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.

### Create a new list to refer to

✕

---

Give the list a handle and provide a word or short phrase synonym for the list item to use when you refer to it. The synonym needs to be in the same language as the vocabulary source.

[Show me an example](#)

\* Handle ⓘ

Language

\* Synonym

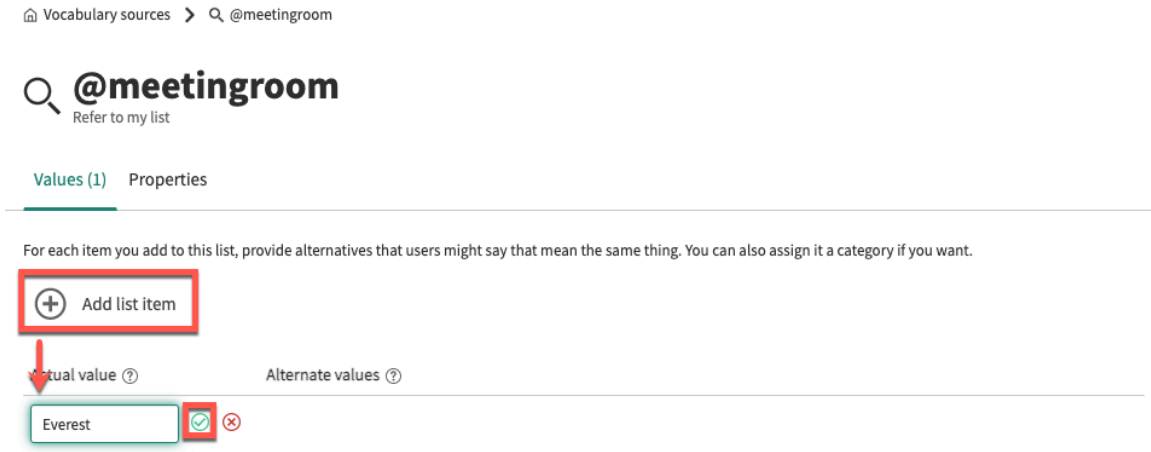
Enable Fuzzy matching (Allow slight misspellings, partial matches etc.)

Make case sensitive

5. Click **Create**.

Your list vocabulary source draft appears in the My lists section of the Vocabulary sources screen.

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.

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.

## Add another table to refer to

Refer to a table

Select the ServiceNow table you want to refer to and specify the fields you want to look up and use with your models. Provide a word or short phrase synonym for items that are in this table. The synonym needs to be in the same language as the vocabulary source.

[Show me an example](#)

* Table	Location (cmn_location) 
* Handle ⓘ	@ Location
* Synonym	Location

- 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:

- **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-separated 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.

▼ Fields

Select fields you want to refer to from this table

\* Field name

Country (country) Options +

City (city) Options + -

Fields can appear together e.g. [First name] [Last name]

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

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](#).

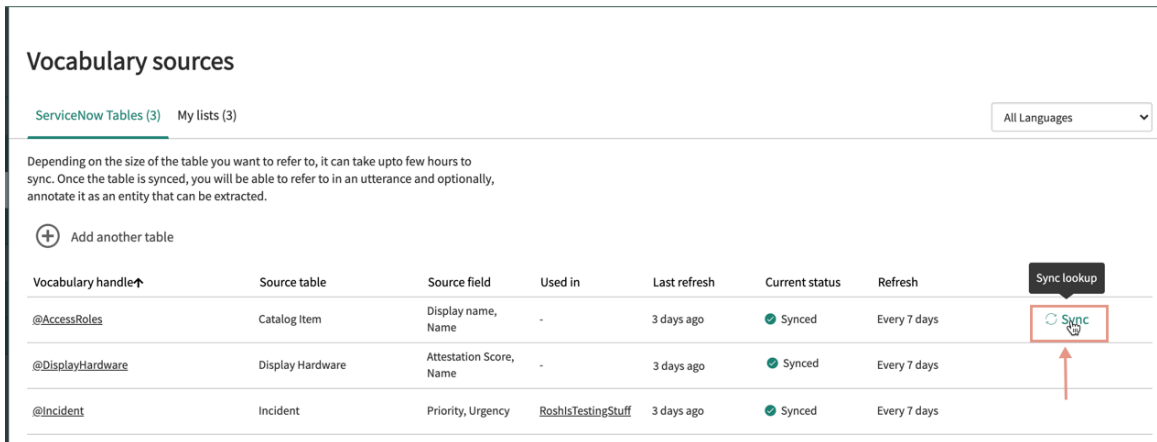
You can also manually sync a table vocabulary source, such as before model training.

**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.

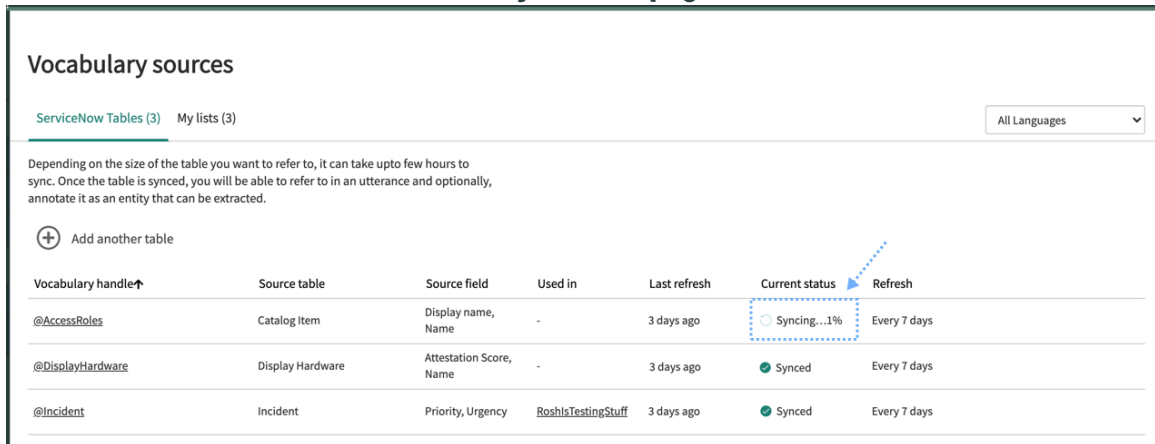


3. Select **Sync**.  
The vocabulary source starts synchronization.

**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.

## ServiceNow Tables tab of the Vocabulary Sources page

**Vocabulary sources**

ServiceNow Tables (3) My lists (3) All Languages

Depending on the size of the table you want to refer to, it can take upto few hours to sync. Once the table is synced, you will be able to refer to in an utterance and optionally, annotate it as an entity that can be extracted.

+ Add another table

Vocabulary handle↑	Source table	Source field	Used in	Last refresh	Current status	Refresh
@AccessRoles	Catalog Item	Display name, Name	-	3 minutes ago	Synced	Every 7 days
@DisplayHardware	Display Hardware	Attestation Score, Name	-	3 days ago	Synced	Every 7 days
@Incident	Incident	Priority, Urgency	RoshisTestingStuff	3 days ago	Synced	Every 7 days

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

#ResetPassword  
Intent

Utterances (12) Associated Entities (1) Settings

Use @ to refer to values (words) in your existing vocabulary source, which you can always add to. To extract specific words from your utterances, select the word to annotate them as entities.

 Add

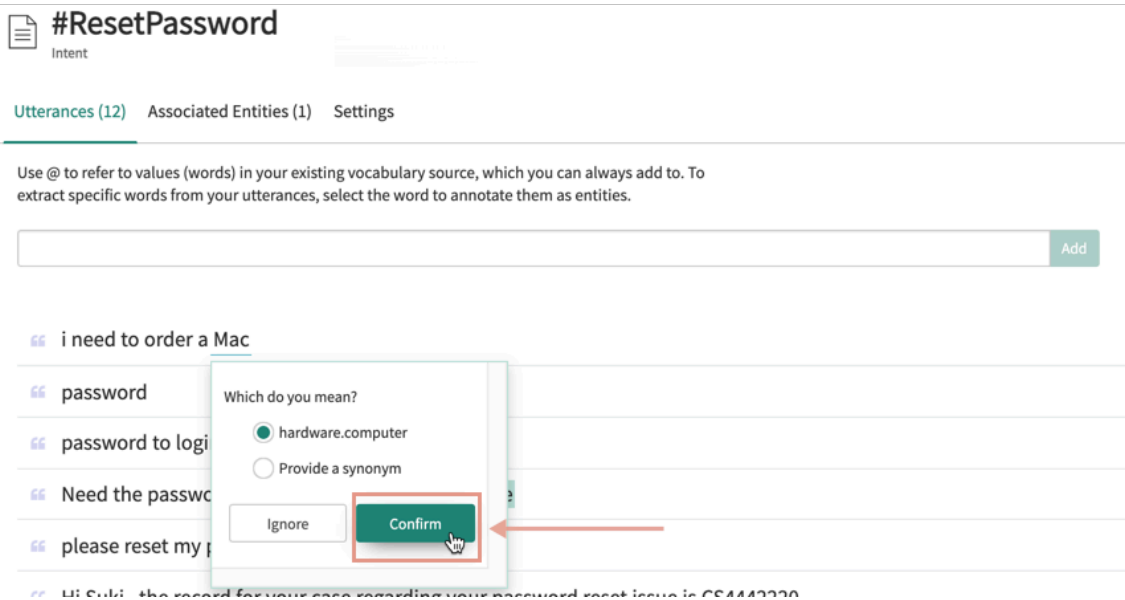
“ i need to order a Mac ▲

“ password

“ password to login

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



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.

**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.

- 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.

Home > HR NLU for VA EFL1 - English (Primary) > Build and train your model

## Build and train your model

Intents (21) Entities (16) Vocabulary (28) **Test set (103)**

This is your model's default test set. You can add utterances to test your model's performance in the next phase.

For an overview of test sets, [check out this video](#).

Test Coverage ⓘ

# 100%

Excellent

[View details](#)

Test Utterance	Expected Intent	Source ⓘ	Last modified ↓
<input type="checkbox"/> I want to submit a request for time off	RequestForLeave	Manual	2023-01-12 20:23:13

- 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.

HR NLU for VA EFL1 Model language English (Primary) Trained a day ago

Model details Settings

Intents ⓘ >  

# 21

Entities ⓘ >  

# 16

Vocabulary ⓘ >  

# 28

Vocabulary sources ⓘ >  

# 1

**Test Coverage ⓘ >**  

# 100%

- Navigate to **All > Multi-model Batch Testing > Test sets** tab. Find the name of your model. Default test sets are labeled as

**Default.** [Test set - HR NLU for VA test](#) Default HR NLU for VA test 151 71.33% 4.67% 2022-10-31 01:54:42 ● Done ○ Run test 🗑️ Delete

### 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**.

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.

- 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**.

## Model management phases

Use these iterative phases to refine and improve your model over time.

**Build and train your model**

View phase

Improve intents by listing the various ways users might express the same need.

---

- Add or remove utterances
- Configure entities to capture important context
- Build and maintain your model's default test set
- Provide synonyms for company specific words and terms

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

**Build and train your model**

Intents (39)
Entities (35)
Vocabulary (36)
Test set (11)

**39**  
intents have low test utterance count

Resolve any issues with existing intents, or create new intents to improve the model content.  
A minimum of 2 intents are required to train a model.

New intent
 Import intents

Showing: Select...

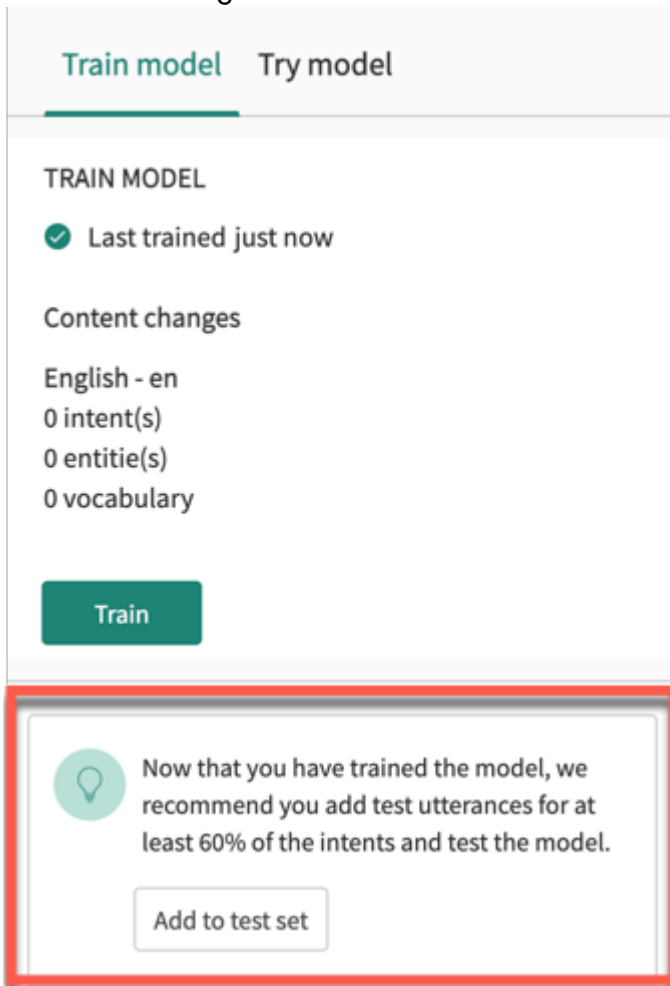
Name	Utterances	Test utterances	Associated entities	Mapped objects	Last updated	Enabled <span style="font-size: 0.8em;">?</span>
<div style="display: flex; justify-content: space-between; align-items: center;"> <div style="border-bottom: 1px solid #ccc; padding-bottom: 5px;"> <p><b>Train model</b>   Try model</p> <p><b>TRAIN MODEL</b></p> <p><span style="color: #0070c0;">ⓘ</span> The model was never trained</p> <p>Content changes</p> <p>English - en 39 intent(s) 25 entitie(s) 36 vocabulary</p> </div> <div style="border: 2px solid red; padding: 5px; font-weight: bold; color: #0070c0; background-color: #0070c0; color: white;">Train</div> </div>						

any.

6. Select the **Train** button.

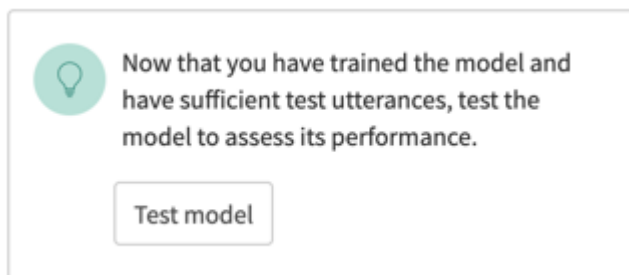
Result: The system displays a progress bar during training. When finished, the system displays one of two recommendations:

- 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](#).



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](#).

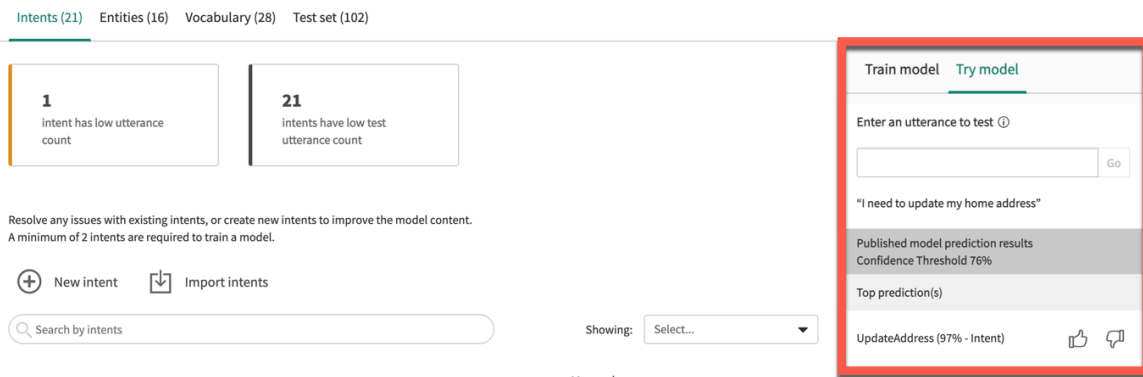


model.

7. To manually try individual utterances, select the **Try model** tab.

8. In the text field under **Enter an utterance to test**, type an utterance and select **Go**.

### Build and train your model





Intents (21) Entities (16) Vocabulary (28) Test set (102)

1 intent has low utterance count

21 intents have low test utterance count


Resolve any issues with existing intents, or create new intents to improve the model content. A minimum of 2 intents are required to train a model.

 New intent  Import intents

Search by intents

Showing: Select...



Train model **Try model**

Enter an utterance to test 

"I need to update my home address"

Published model prediction results  
Confidence Threshold 76%

Top prediction(s)

UpdateAddress (97% - Intent)  

## 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.

## 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.

# Build and train your model

Showing:

Train model
Try model

Enter an utterance to test ⓘ

Go

“help with hr”

Published model prediction results  
Confidence Threshold 76%

Top prediction(s)

CreateHRGeneralInquiryCase (96% - Intent)

👍

👎

**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.

Train model
Try model



Enter an utterance to test ⓘ

“help with hr”

Published model prediction results  
Confidence Threshold 76%

Top prediction(s)

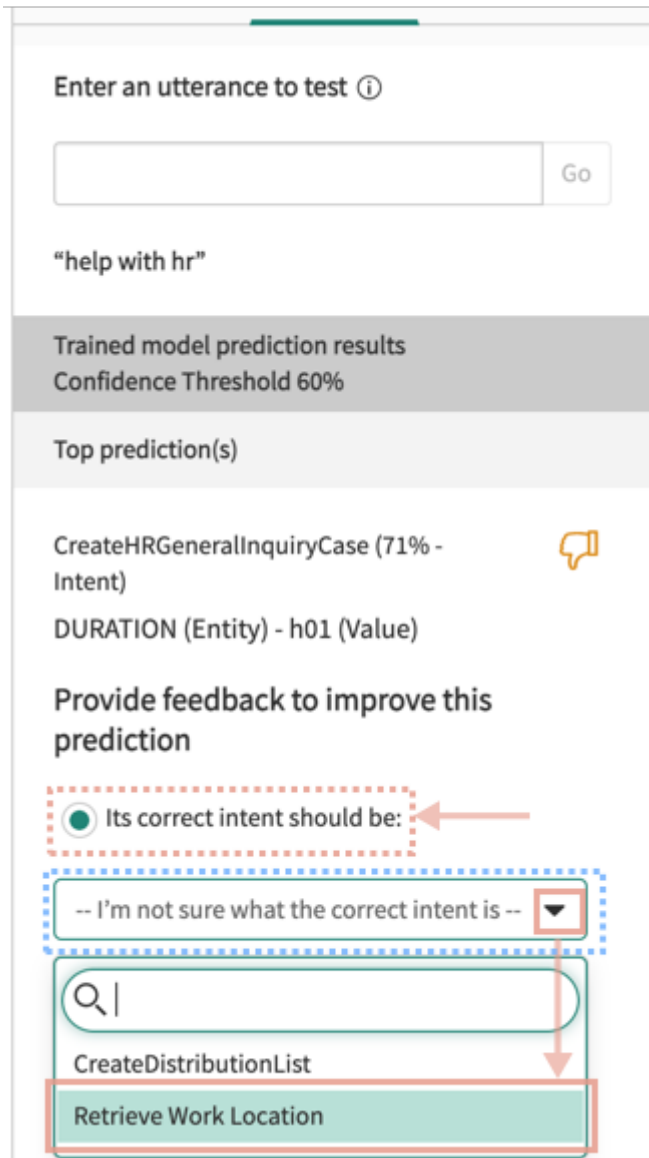
CreateHRGeneralInquiryCase (96% - Intent)

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:

- 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.

Train model Try model

Enter an utterance to test ⓘ

Go

“ਮੇਰੀ ਵਰਚੁਅਲ ਮਸ਼ੀਨ broke down”

Trained model prediction results  
Confidence Threshold 60%

No intent was predicted for this utterance

If this isn't what you expected, you can give feedback on the prediction.

[Give feedback](#)

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.

Train model **Try model** ✕

Enter an utterance to test ⓘ

“ਮੇਰੀ ਵਰਚੁਅਲ ਮਸ਼ੀਨ broke down”

Trained model prediction results  
Confidence Threshold 60%

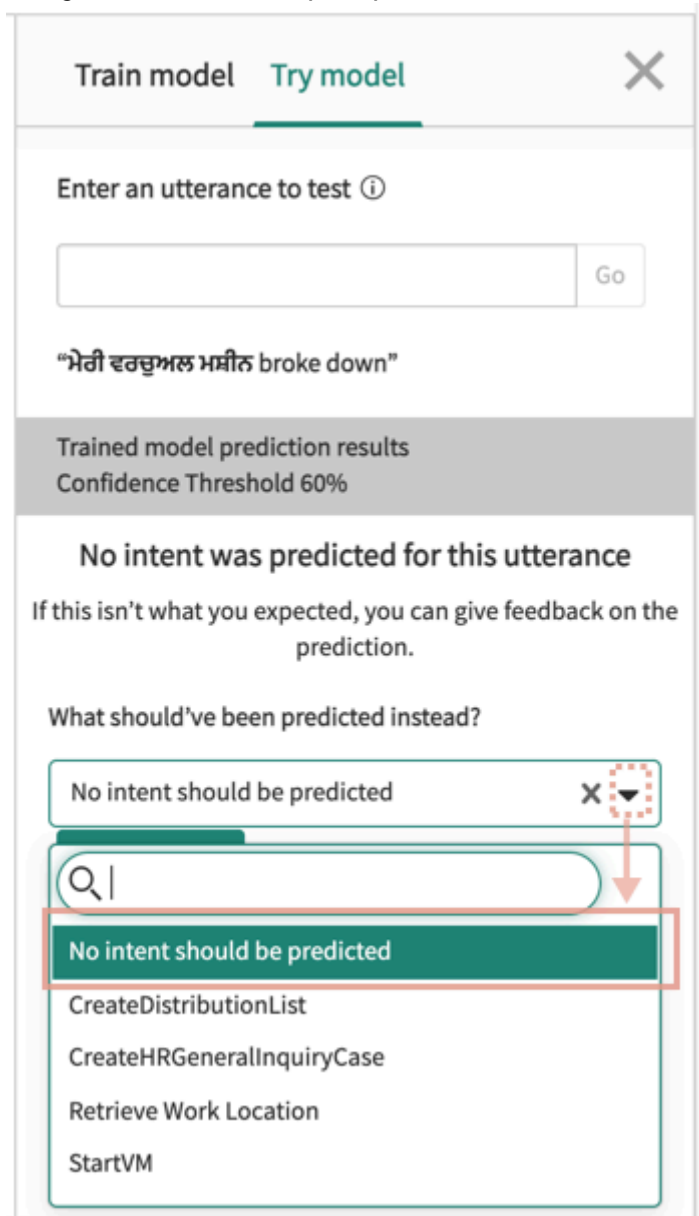
**No intent was predicted for this utterance**

If this isn't what you expected, you can give feedback on the prediction.

What should've been predicted instead?

 ✕ ▼

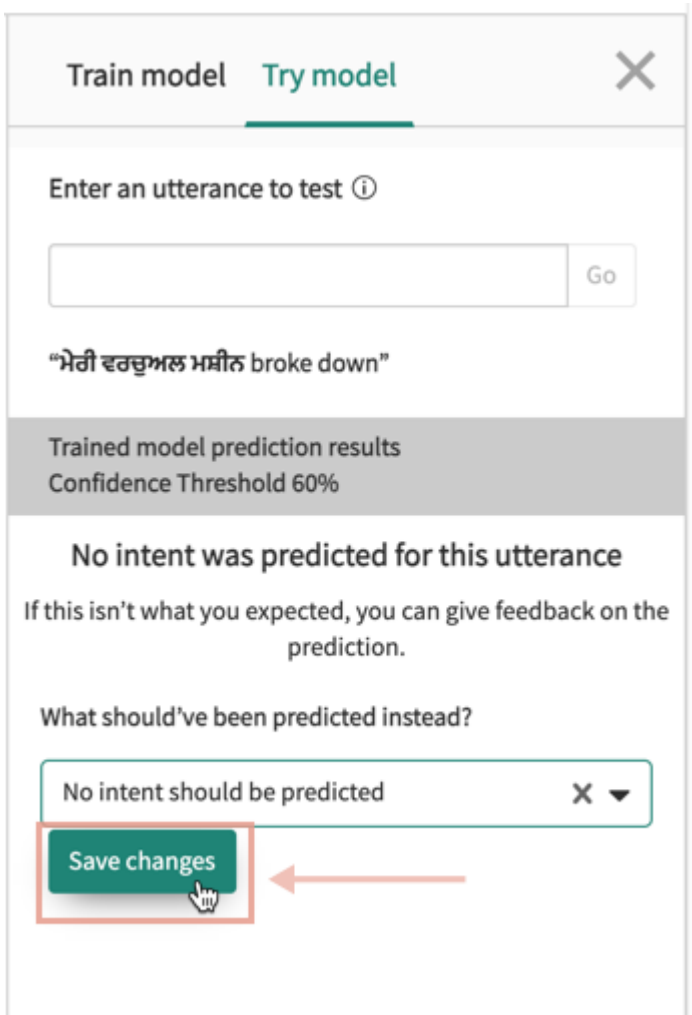
So instead of choosing an intent from the prompt, this user selects the **No intent should be**



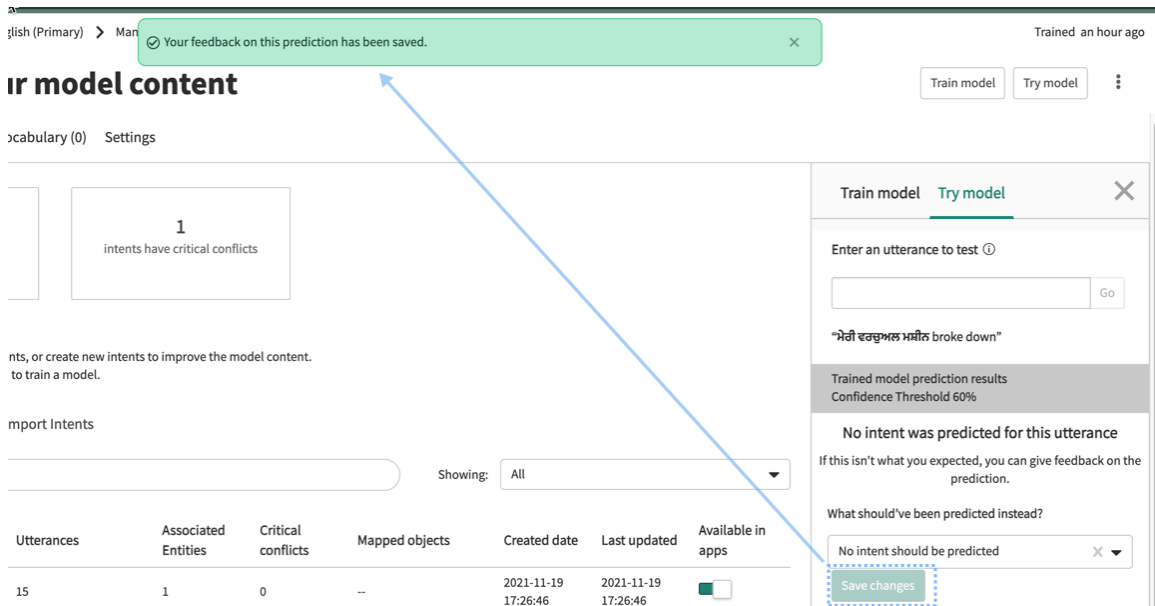
predicted option.

**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**.



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.



## 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.

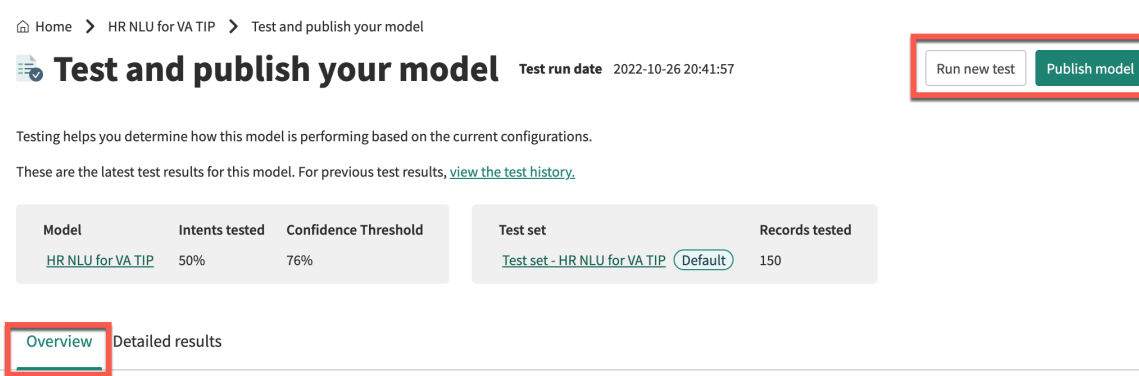
**Note:** Testing your model requires the Multi-model Batch Testing feature, available with the NLU Workbench - Advanced Features application from ServiceNow® Store. For more information, see [Install NLU Workbench - Advanced Features](#).

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.

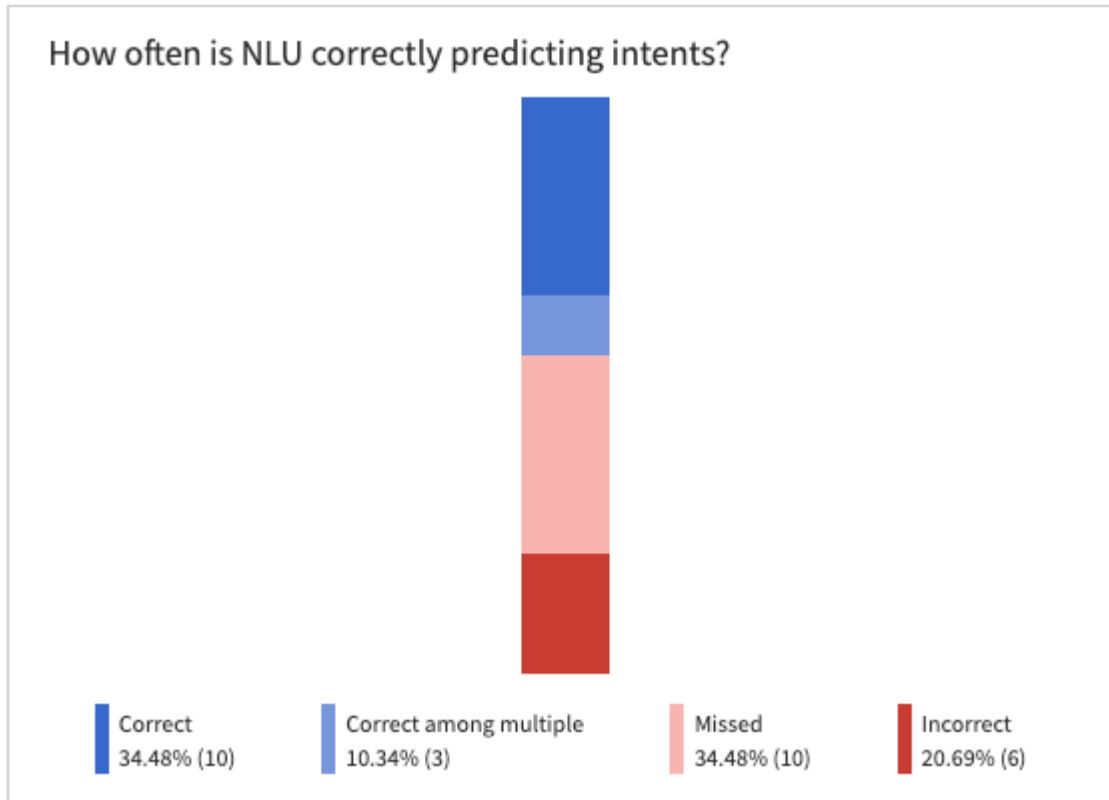
**Test run date**

2023-01-09 17:36:31

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 <b>Not relevant</b> , 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](#).

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.

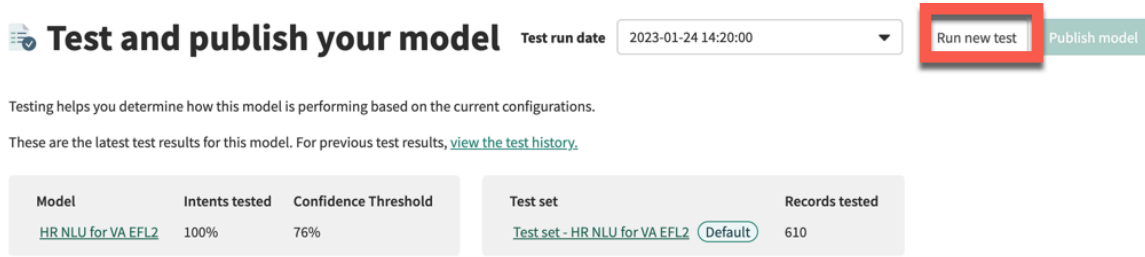
- If an expected intent in your test set does not correspond to any intent in the model, the utterances with those intents are not used for testing. They are not included in the test results.
- The mid-conversation responses of Dialog Acts can't be tried or tested in NLU Workbench.
- When the model returns 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.

In this example scenario, you've trained your model and want to assess the performance.

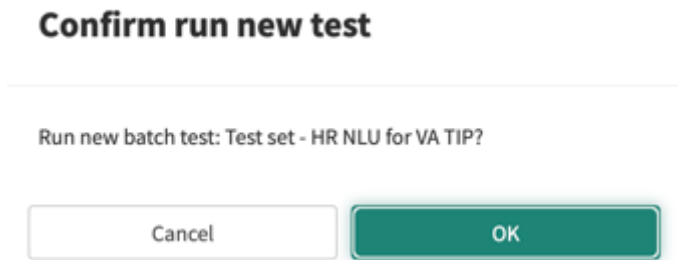
**Procedure**

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.

## 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**.

Home > HR NLU for VA TIP > Test and publish your 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.

## 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

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.

**#UpdateEmail**

Utterances (37) Associated Entities (1) Conflicts (0) Settings

Use @ to refer to values (words) in your existing vocabulary source, which you can always add to. To extract specific words from your utterances, select the word to annotate them as entities.

Enter an utterance to test

Go

"wrong email address"

Trained model prediction results  
Confidence Threshold 76%

Top prediction(s)

UpdateEmail (87% - Intent)  
SOFTWARE (Entity) - email (Value)

Published model prediction results  
Confidence Threshold 76%

Top prediction(s)

UpdateEmail (85% - Intent)  
SOFTWARE (Entity) - email (Value)

### 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.

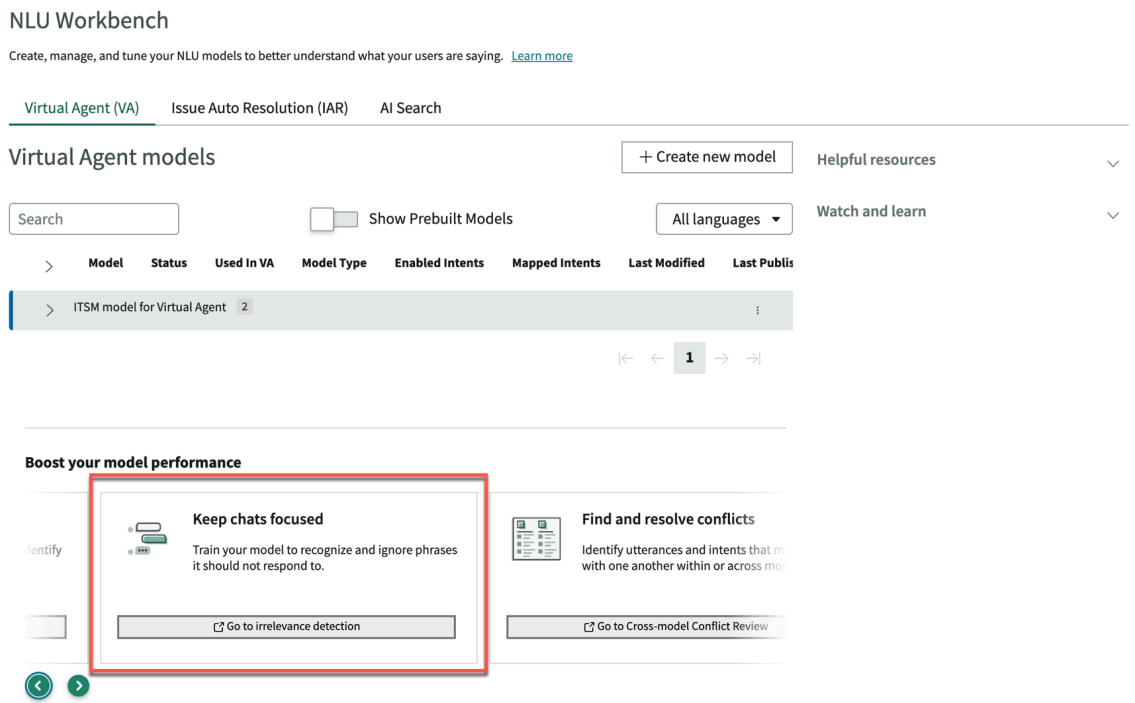
The Irrelevance detection table is where you can manage utterances marked as **Not relevant**.

### 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.

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.

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.

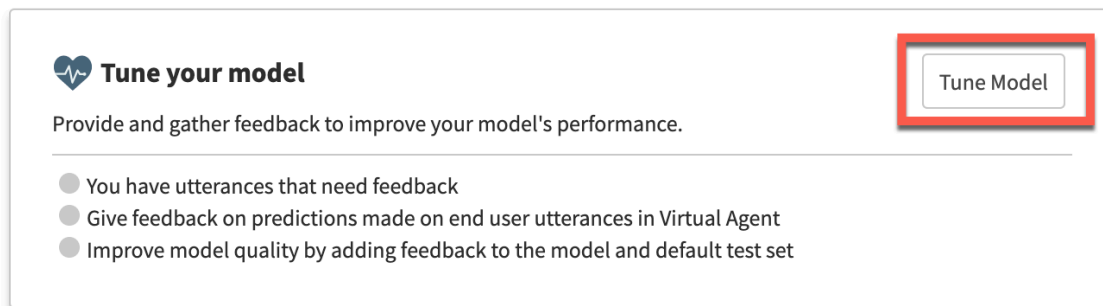
- When testing models against test sets, results are considered Correct if no intent is predicted for an irrelevant utterance.
- Utterances that were marked as **Not relevant** can be re-assigned later. For example, if a new intent is created, existing irrelevant utterances can be manually assigned to 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<sup>®</sup> 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.

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.

Home > HR NLU for VA EFL1

HR NLU for VA EFL1 Model language English (Primary) Published 5 days ago

Virtual Agent model for HR

Model details **Settings**

\* Name

Short description

Language

Created for

Business area (optional)

Ignore punctuation

### Model threshold settings

The confidence threshold determines what intents from the model can get predicted. Allow the system to automatically set the optimal threshold for your model, or manually set the threshold.

Threshold type  Automatic (Optimal)  Manual

Confidence Threshold(%)

Save

## 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:


- **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.

Breakdown of current results



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.

## Virtual Agent models

+ Create new model

Search  Show Prebuilt Models  All languages ▾

>	Model	Status	Used In VA	Model Type	Enabled Intents	Mapped Intents	Last Modified
▾	ITSM model for Virtual Agent 2						⋮
	English(Primary)	Draft Saved	No		6	0	2023-06-22 12:09:54
	Brazilian Portuguese	Draft Saved	No		6	0	2023-06-24 14:22:21

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.

### 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

▼	NLU for Access Requests	3 languages	Virtual Agent	⋮		
	<a href="#">English (Primary)</a>	Draft saved 23 days ago	12	0	37	⋮
	<a href="#">French</a>	Needs review	12	0	37	⋮
	<a href="#">German</a>	Needs review	12	0	37	⋮

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.

## Viewing model languages list

Home > Booking

Booking Model language English (Primary) updated 14 days ago

Virtual Agent model for Customer Support

Model details

Intents ?

## 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.
```

The page titled **Review how we've grouped your models** opens.

Home > Review how we've grouped your models

6 accepted

## Review how we've grouped your models

Do this later **Finish**

To make managing your NLU easier, we've identified these primary models and grouped secondary models under each one. Review these and edit as needed.

VM Management NLU for Virtual Agent-CAI English - en (Primary) Accepted **Edit**

Secondary languages	Intent mappings	Entity mappings
VM Management NLU for Virtual Agent-CAI DE (German - de) 4 intents, 1 entities	4 mapped	1 mapped

## 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.

Home > Review how we've grouped your models > Review what's in this grouping

## Review what's in this grouping

Cancel **Save**

### Review model grouping

If you'd like to delete this, remove all secondary models.

\* Primary model name

\* Secondary models

**AI Search custom model ES (es)** 2 intent(s), 5 entities

### Review the intents and entities in this 1 language

To simplify grouping, all intent and entity names in other languages must be the same as those in the primary model. Any variations will be renamed to match the primary model.

AI Search custom mo...

**AI Search custom model ES (Spanish - es)** 2 intent(s), 5 entities **Intents** 2 mapped, 0 unmapped **Entities** 5 mapped, 0 unmapped

Language	English (Primary)	Spanish - es
Intents	catalogFind	catalogFind
	personFind	personFind

- 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.

Here are your other models.

You can designate them as primary or secondary models. You can also group them under an existing primary model.

Model ↑	Language	Enabled intents	Entities	
AI Search custom model	English - en	2	5	<b>Group</b>

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.

[Home](#) > [Review how we've grouped your models](#) > [Review what's in this grouping](#)

## Group what's in these models

Cancel Save

Create model grouping

\* Primary model name

\* Secondary models

**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.

Translation mode	Description
<b>Use AI</b>	<p>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 <a href="#">Dynamic Translation</a>.</p> <p>To activate the Use AI mode, <a href="#">Localization Framework</a> must be configured and activated.</p>
<b>Use a third party</b>	<p>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 <a href="#">Localization Framework</a> and <a href="#">Translation modes</a>.</p>
<b>No translation</b>	<p>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.</p>

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.

## NLU Workbench

Create, manage, and tune your NLU models to better understand what your users are saying. [Learn more](#)

Virtual Agent (VA) Issue Auto Resolution (IAR) AI Search

### AI Search models

+ Create new model

Helpful resources

Search



Show Prebuilt Models

All languages

Watch and learn

Model Type	Enabled Intents	Mapped Intents	Last Modified	Last Published
Custom	0	0	2023-07-06 08:39:38	⋮
Custom	2	0	2023-07-05 09:44:54	⋮
Pre-built	0	0	2023-06-11 21:58:05	

- View localization request
- Manage editors
- Add a language**
- Duplicate this model
- Export model as CSV
- Delete this model

4. In the **Add a language to this model** window, choose the desired language for the secondary model and the translation method.

### Add a language to this model ✕

Choose an additional language and a translation method.

Primary language

English

Additional language \*

French

How to translate \* ⓘ

**Use AI**

Automatically translate with AI

**Use a third party**

Request a translation from a professional translator

**No translation**

Do not translate the utterances

Cancel

**Add**

5. Select the **Add** button.

The secondary model is created and the **Build and train your model** phase of the secondary model loads.

### 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.

#### Intents list

Name	Utterances	Associated Entities	Critical conflicts	Mapped objects	Created date ↓	Last updated	Enabled ⓘ
<a href="#">#401kBenefitsInquiry</a>	19	0	0	--	2022-05-10 01:27:07	2022-05-10 01:27:07	<input checked="" type="checkbox"/>
<a href="#">#AddEmergencyContact</a>	28	4	0	--	2022-05-10 01:27:07	2022-05-10 01:27:07	<input checked="" type="checkbox"/>

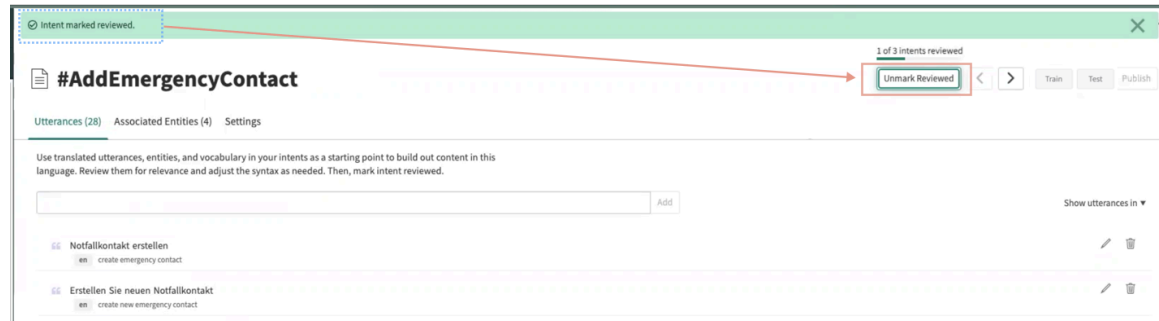
Enabling of intents works in the following ways:

- 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.

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.

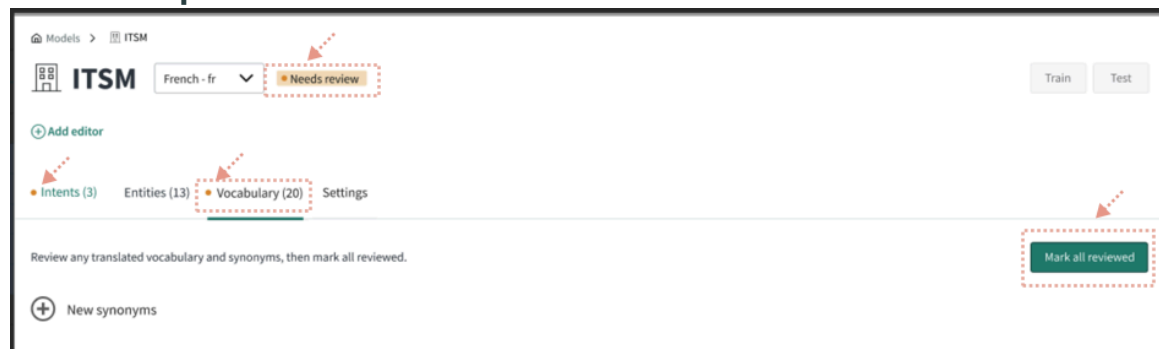
### Switch between the Reviewed and Unmark Reviewed states for an intent



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.

## 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

The screenshot shows the 'Virtual Agent Conversation Setup JA' model in the 'Needs review' state. At the top, there is a summary box indicating '0 of 5 intents reviewed'. Below this, a table lists five intents, all with a 'Needs review' status. The first intent, '#End Conversations', is highlighted with a red box and an arrow. To the right of the table, an 'Enable' button is highlighted with a blue dashed box and an arrow. The table columns are: Status, Name, Utterances, Associated Entities, Mapped objects, Created date, and Last updated.

Status	Name	Utterances	Associated Entities	Mapped objects	Created date	Last updated
Needs review	#End Conversations	4	0	--	2021-04-20 14:38:13	2021-04-20 14:38:19
Needs review	#Feedback	6	0	--	2021-04-20 14:38:13	2021-04-20 14:38:19
Needs review	#Greetings	15	0	--	2021-04-20 14:38:13	2021-04-20 14:38:20
Needs review	#Live Agent	25	0	--	2021-04-20 14:38:13	2021-04-20 14:38:21
Needs review	#What Can You Do	6	0	--	2021-04-20 14:38:13	2021-04-20 14:38:22

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.

## 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.

**Build and train your model** Needs review

Intents (22) Entities (11) Vocabulary (28) Test set (0)

Review the translations of the utterances and entities within each intent before marking the intent as "Reviewed". After you've marked all intents as "Reviewed", you can test or train the model.

Search by intents Showing: Needs review

Reviewed status	Name	Utterances	Test utterances	Associated entities	Mapped objects	Last updated	Enabled
Needs review	<a href="#">#RequestForLeave</a>	61	0 <span>Low test utterances</span>	0	--	2022-05-13 03:04:54	<input type="checkbox"/>
Needs review	<a href="#">#TuitionReimbursementRequest</a>	19	0 <span>Low test utterances</span>	0	--	2022-05-13 03:04:55	<input type="checkbox"/>
Needs review	<a href="#">#UpdateAddress</a>	28	0 <span>Low test utterances</span>	0	--	2022-05-13 03:04:56	<input type="checkbox"/>

**TRAIN MODEL**

The model was never trained

Content changes

German - de  
22 intent(s)  
6 entity(s)  
28 vocabulary

You have 6 entities that are missing annotations. Annotate utterances with entities for entity prediction  
[See entities](#)

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.**

## NLU Workbench

Create, manage, and tune your NLU models to better understand what your users are saying. [Learn more](#)

Virtual Agent (VA) Issue Auto Resolution (IAR) **AI Search**

### AI Search models

+ Create new model

Helpful resources

Show Prebuilt Models

All languages ▾

Watch and learn

Status	Used In AI Search	Model Type	Enabled Intents	Mapped Intents	Last Modified	Last Published
Draft Saved	No	Custom	0	0	2023-07-06 08:39:38	⋮
Trained	No	Custom	2	0	2023-07-05 09:44:54	⋮
Draft Saved	No	Custom	2	0	2023-07-16 17:47:08	⋮
Draft Saved	No	Pre-built	0	0	2023-06-11 21:58:05	⋮

- View localization request
- Manage editors**
- Add a language
- Duplicate this model
- Export model as CSV
- Delete this model

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.

### Add editors to the AI Search custom model ES model ✕

Language

Spanish

Editors \*

Search editors

Cancel

Add

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.

✔ Model "NLU Common Entities ES Copy1" assigned to an editor
✕

### 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.

## 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.

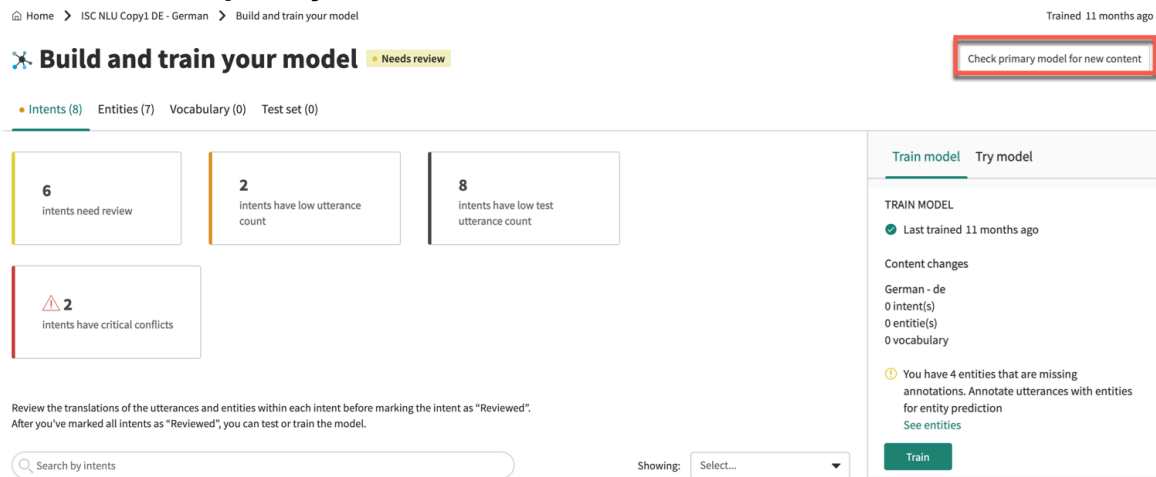
### About this task

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.

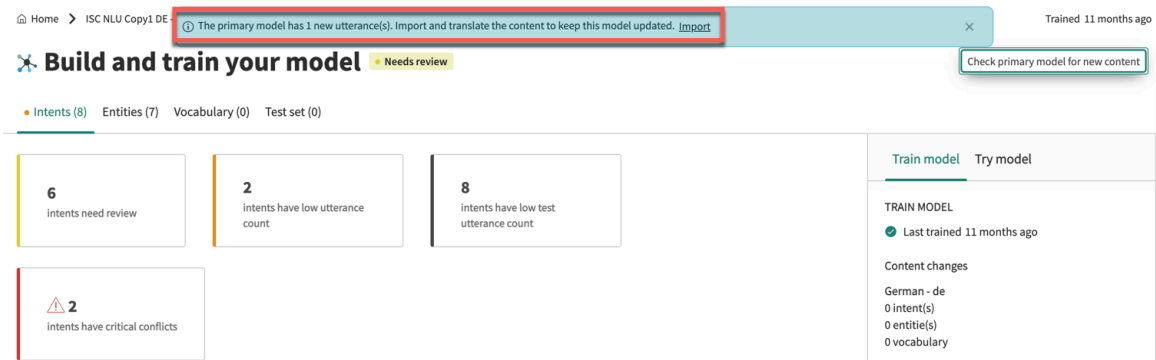
You can import primary model content to only one secondary model at a time.

### 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.



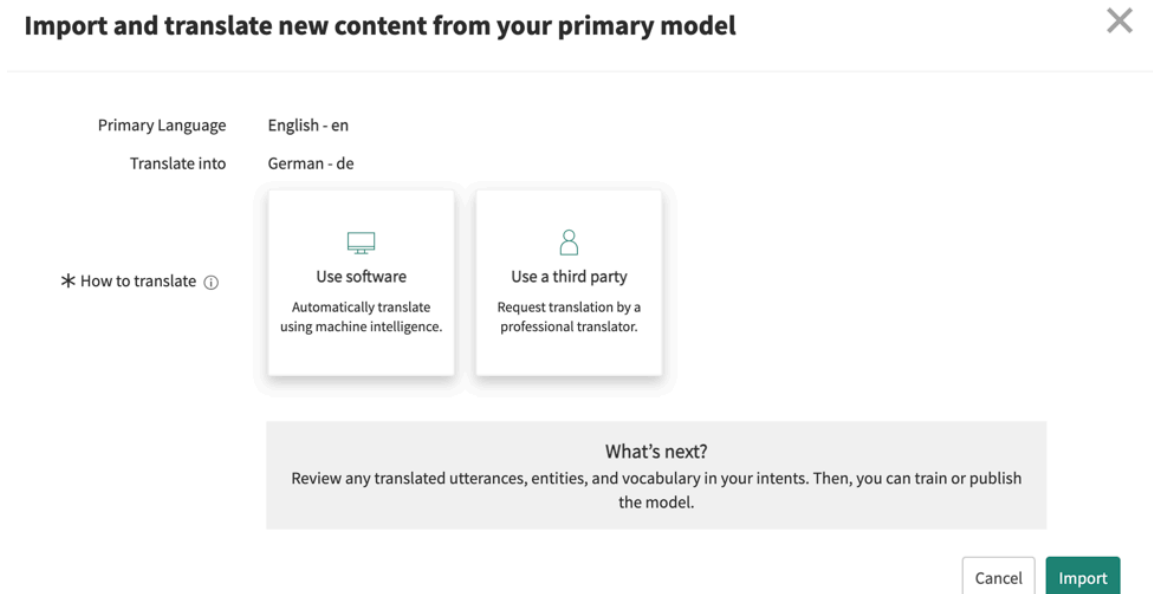
The banner Import and translate the content to keep this model updated appears at the top of the screen, along with a summary of the updates in the primary model.



**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.



**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](#).

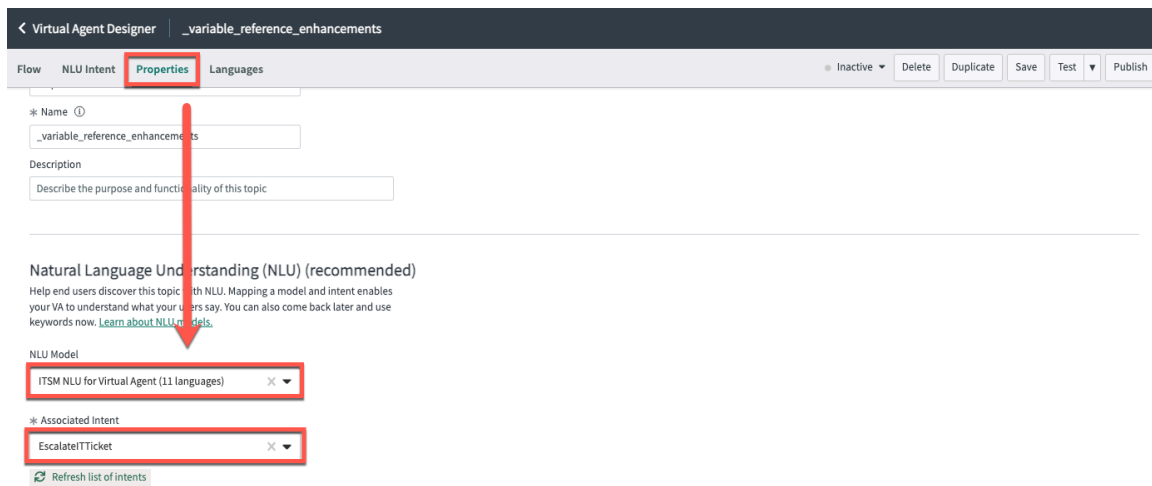
After the NLU model is complete and associated with a Virtual Agent conversation topic, administrators with the virtual\_agent\_admin or admin role can do the following from within the Virtual Agent Designer user interface.

- Update NLU intent utterances.
- Train, test, and publish the NLU model.

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.

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.

Visit the [ServiceNow Store](#) website to view all the available apps and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the [ServiceNow Store version history release notes](#).

**3.** 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 <b>Ignore</b> .
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.

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.

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.

### Batch Test

Run an analysis

Test results Test sets

All languages ▼

Use large data sets with labeled utterances to test how your models are performing.

Test set	Models	Type	Test utterances	Correct	Incorrect	Last run	Status
<a href="#">Test set for HR</a>	HR for VA, HR for Search	Test	154	71.52%	4.64%	2022-05-15 22:02:20	✔ Done
<a href="#">Optimization Test Set</a>	HR for VA	Optimize	154	71.52%	4.64%	2022-05-12 19:41:11	✔ Done
<a href="#">Test set for HR</a>	HR Ger, HR VA Ger	Test	154	96.43%	2.38%	2022-05-12 04:04:50	✔ Done
<a href="#">Test set for CSM</a>	NLU CSM 1, NLU CSM 2, NLU CSM 3	Test	154	71.52%	4.64%	2022-05-12 04:05:09	✔ Done

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.

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**

3. Click **Create test set**.



### Create new test set

\* Test set name

Language English - en ▾

Test set file

Please Note: Ensure that your file has "Utterance" and "Expected intent" columns:

Utterance	Expected intent
-----	-----
-----	-----

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**

Batch Testing > HR VA Test

**HR VA Test**

Utterances	Language	Last updated	Status
924	English - en	2021-04-11 17:19:49	Uploaded a month ago



Test utterance	Expected intent
I need to update my phone number for HR.	UpdatePhoneNumber
I have a new phone number	UpdatePhoneNumber

Select a CSV or XLSX (Excel Workbook) file with your additional utterances to import.

### Import utterances window

#### Import a list of utterances paired with correct intents



\* Test set file

Select file

Please Note: Ensure that your file has "Utterance" and "Expected Intent" columns:

Utterance	Expected intent
-----	-----
-----	-----

Cancel

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](#).

## 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.

### Run a new batch test



Test the model to validate its quality and identify areas to improve the model before publishing it.

\* Select model(s)

\* Select a test set

Cancel Run

4. Select a test set from the list.
5. Click **Run**.

### Batch Test

✔ Your batch test run is being processed. ✕

Test results Test sets

A

Use large data sets with labeled utterances to test how your models are performing.

Search

Test set	Models	Test utterances	Correct	Incorrect	Last run	Status
<a href="#">Test set - HR NLU for VA test</a> <span>Default</span>	HR NLU for VA test , HR NLU Translation DEF	150				<span>Testing...</span> <span>✕ Cancel run</span>

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](#).

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.

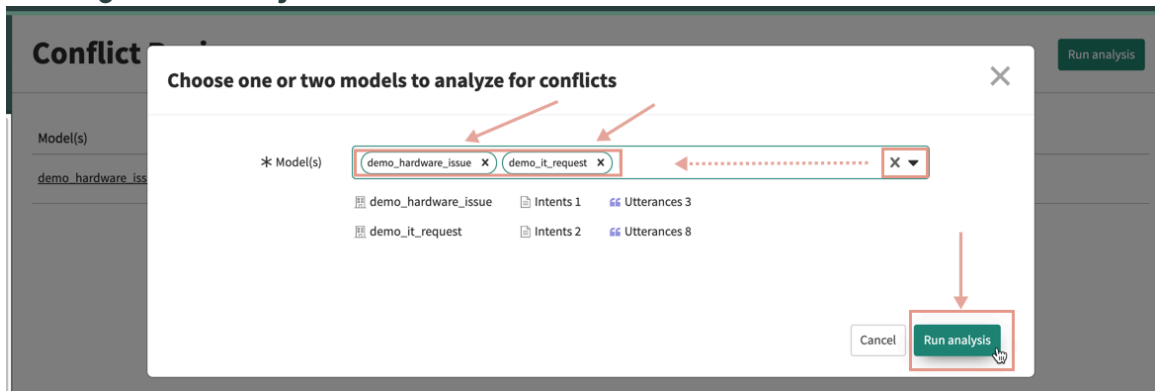
When a conflict is detected, you can use one of following actions to resolve the conflict:

- 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**.

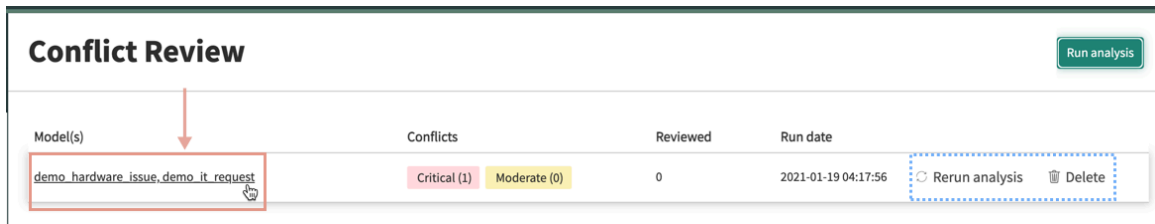
## Running conflict analysis



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.

## Reviewing conflict details

Conflict Review > demoHardware\_issue, demo\_it\_request

### Conflict Review

Model(s)	Trained on	Conflicts	Version
demoHardware_issue	2021-01-19 04:17:24	Critical (1)	2021-01-19 04:17:56
demo_it_request	2021-01-19 04:17:33	Moderate (0)	

[Rerun analysis](#)

**0 of 1**  
Critical conflicts reviewed

**Filter conflicts**

Show Moderate (0)

Hide reviewed (0)

**We found overlapping utterances in these two intents**

Remove or edit any utterances that are under the wrong intent, so each utterance is only under one. Any changes you make will automatically update the model to improve NLU performance.

# laptop\_issue
# laptop\_not\_working

1
“ laptop is really slow

“ laptop is really slow

Ignore

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**.

## Reviewing overlapping utterances

Conflict Review > demoHardwareIssue, demoItRequest

### Conflict Review

Model(s)	Trained on	Conflicts ⓘ	Version
demoHardwareIssue	2021-01-19 04:17:24	Critical (1)	2021-01-19 04:17:56
demoItRequest	2021-01-19 04:17:33	Moderate (0)	

Rerun analysis

**0 of 1**  
Critical conflicts reviewed

**We found overlapping utterances in these two intents**

Remove or edit any utterances that are under the wrong intent, so each utterance is only under one. Any changes you make will automatically update the model to improve NLU performance.

#laptop\_issue    #laptop\_not\_working

① “laptop is really slow”    “laptop is really slow”

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

Home > demoHardwareIssue > #laptop\_not\_working Trained 11 hours ago

### #laptop\_not\_working

Intent

Delete Train Test Publish

Utterances (3) Associated Entities (1) Settings

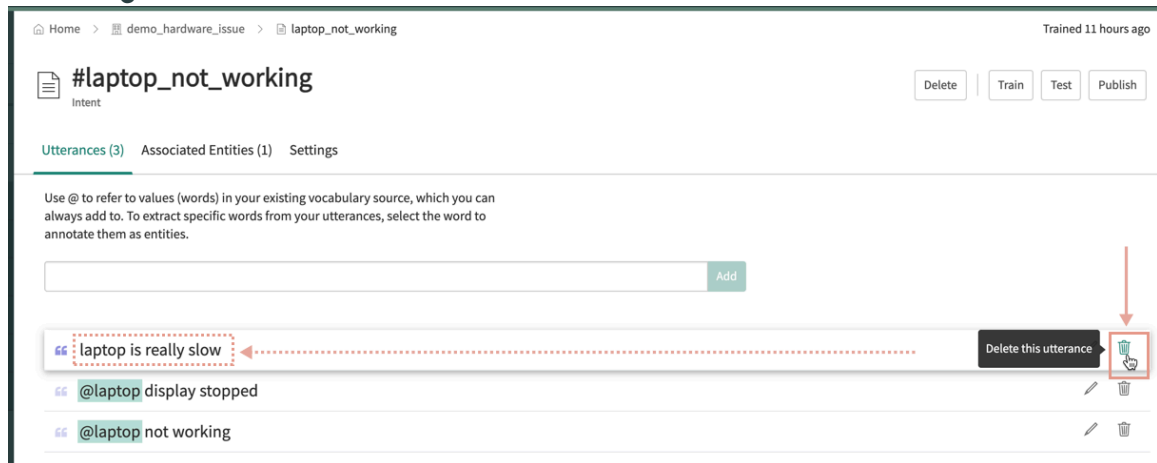
Use @ to refer to values (words) in your existing vocabulary source, which you can always add to. To extract specific words from your utterances, select the word to annotate them as entities.

Add

- “laptop is really slow” ✎ 🗑
- “@laptop display stopped” ✎ 🗑
- “@laptop not working” ✎ 🗑

7. Click the **Delete this utterance** trash can icon.

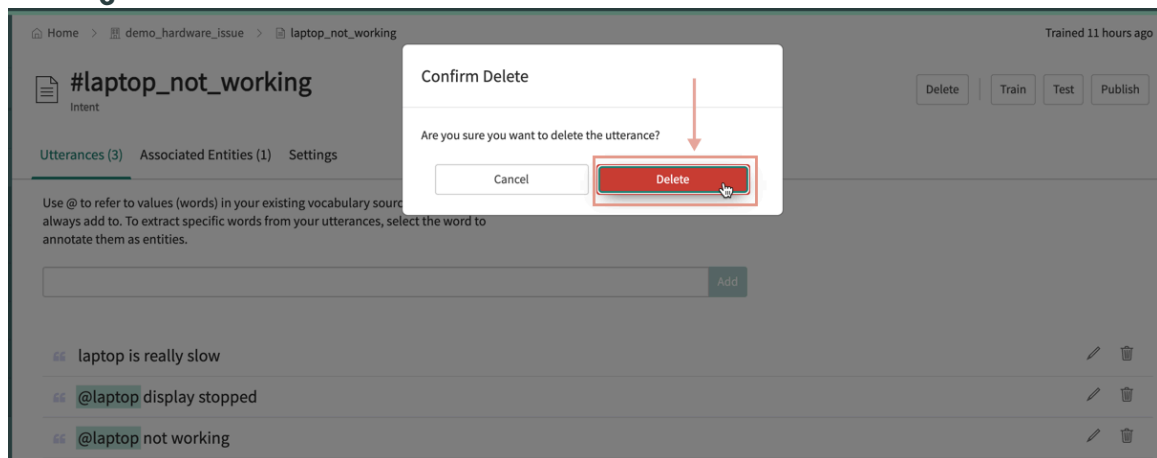
## Resolving conflicts



The Confirm Delete screen appears.

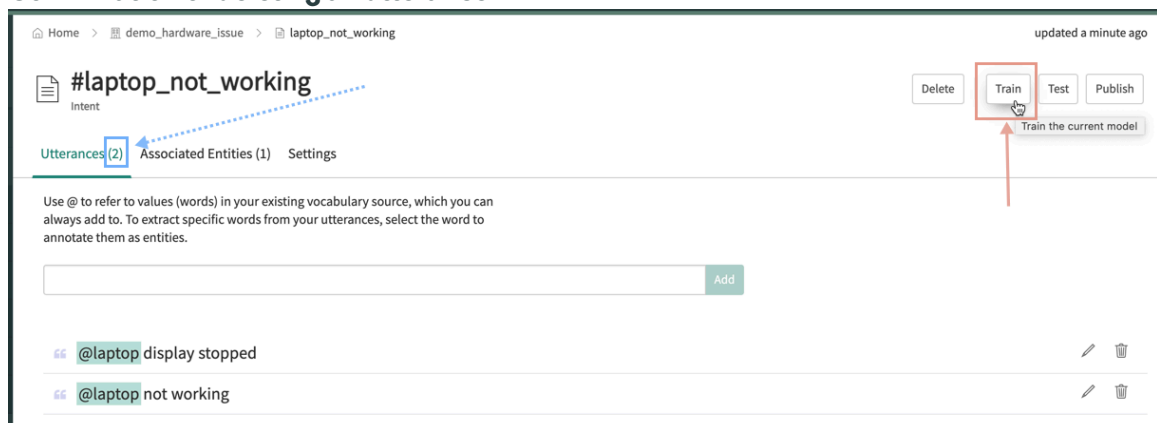
### 8. Select **Delete**.

## 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

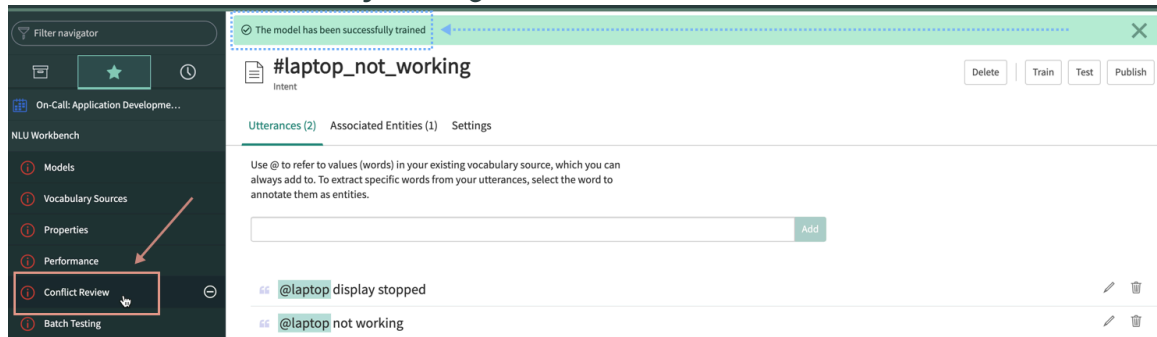


### 9. Select **Train**.

A banner appears on the Intent screen, confirming the model is successfully trained.

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.

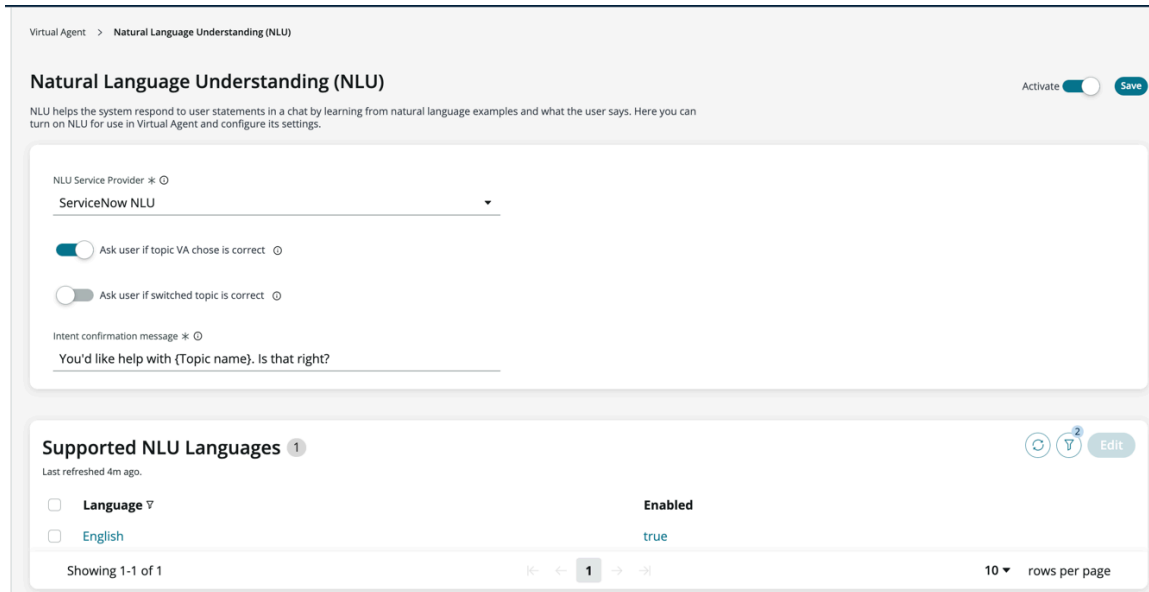
### Installation

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

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.

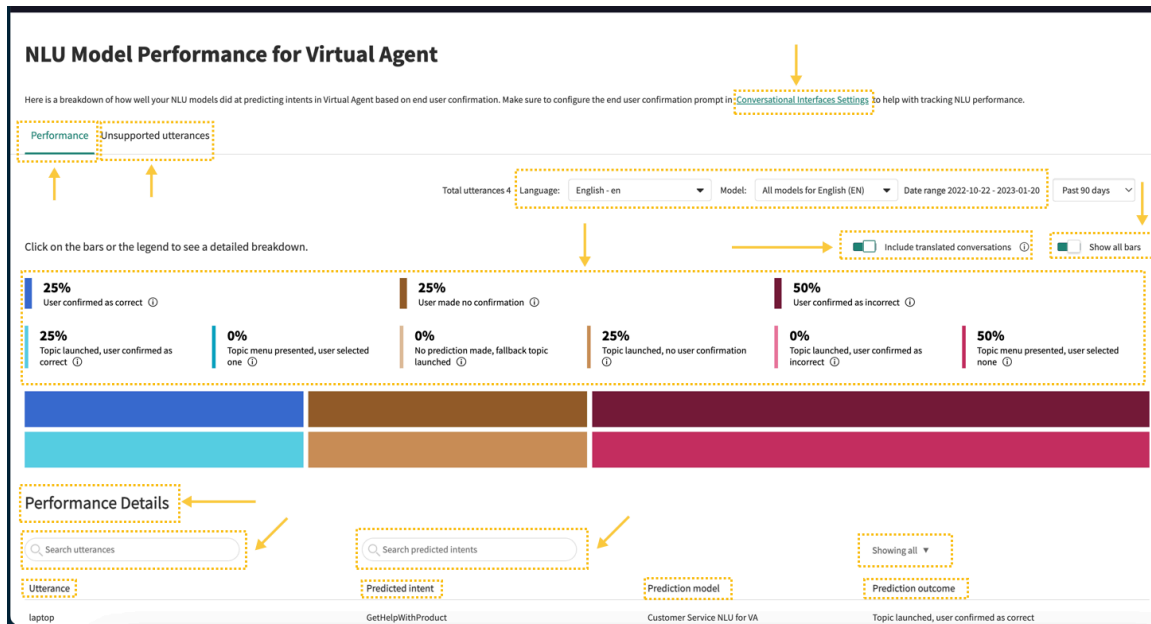
- The **Performance** tab and its colored chart which shows data describing scenarios that occur when users interact with VA. This chart is segmented in shades of three colors, each representing a user scenario. The rows in the chart that have text and percentages are referred to as the legend. The three 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

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
<b>User confirmed as correct</b>	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.
<b>User made no confirmation</b>	The topic was launched but the user did not confirm whether or not it met their need.

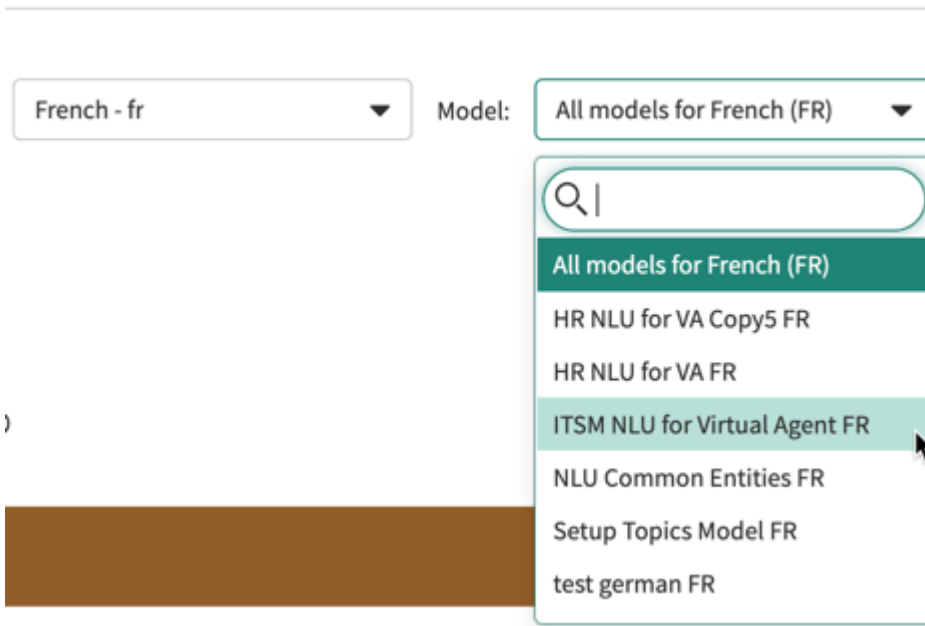
### Scenario Definitions (continued)

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.
<b>User confirmed as incorrect</b>	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.



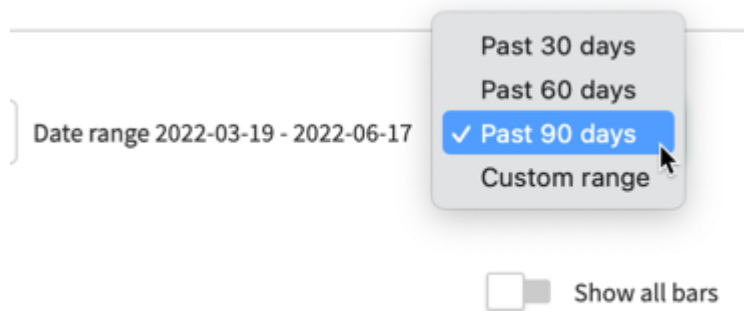
There are other interactions. For example, from the default view of the NLU Model Performance landing page:

- 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

#### NLU Model Performance for Virtual Agent

Here is a breakdown of how well your NLU models did at predicting intents in Virtual Agent based on end user confirmation. Make sure to configure the end user confirmation prompt in [Conversational Interfaces Settings](#) to help with tracking NLU performance.

Performance **Unsupported utterances** ←

We have grouped utterances where NLU did not make a prediction or end user confirmed the prediction is incorrect to help you quickly identify areas of improvement. To make model improvements, visit [Expert Feedback Loop](#) to provide feedback on end user utterances. Show additional ▾ Clusters 2 Rerun analysis

% utterances	# utterance	Top keywords in this cluster
66.67%	6	hello, salut, morning, good, ciao
33.33%	3	help, issue, need, question

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.

## Reviewing the top representative utterances in a cluster

### NLU Model Performance for Virtual Agent

Here is a breakdown of how well your NLU models did at predicting intents in Virtual Agent based on end user confirmation. Make sure to configure the end user confirmation prompt in [Conversational Interfaces Settings](#) to help with tracking NLU performance.

% utterances	# utterance	Top keywords in this cluster
66.67%	6	hello, salut, morning, good, ciao
33.33%	3	help, issue, need, question

Top utterances

- I have a question
- what can you help me with
- need help with an issue

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.

### 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.

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:

1. At the top of the Expert 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.

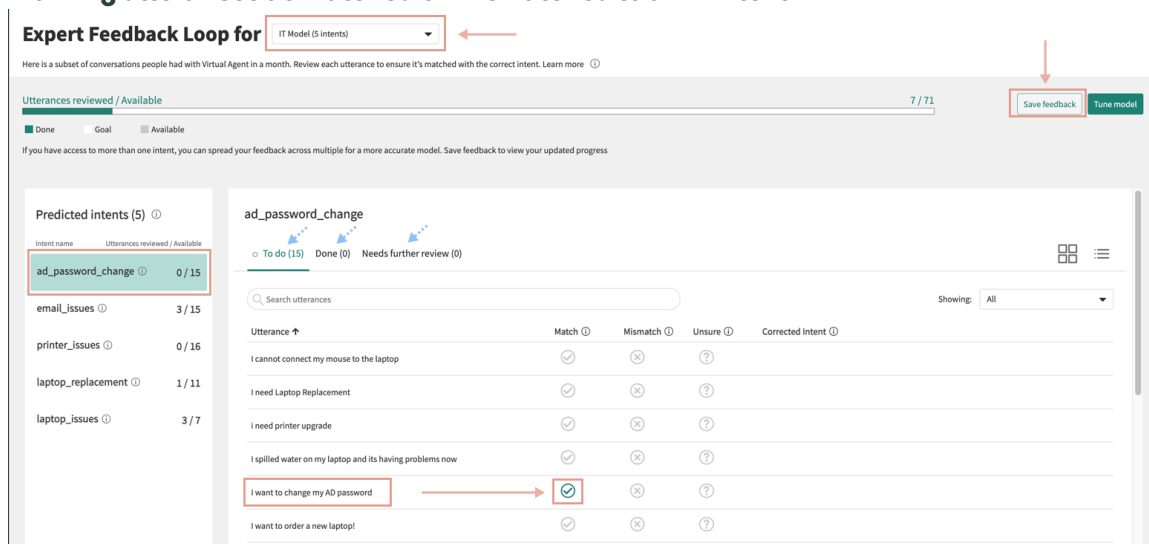
2. Click the **ad\_password\_change** intent so it loads its utterances into the **To do** section for your review.
3. Review the 15 utterances in the intent. In this scenario, you've started by correctly marking the I want to change my AD password utterance as a match to the **ad\_password\_change** 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



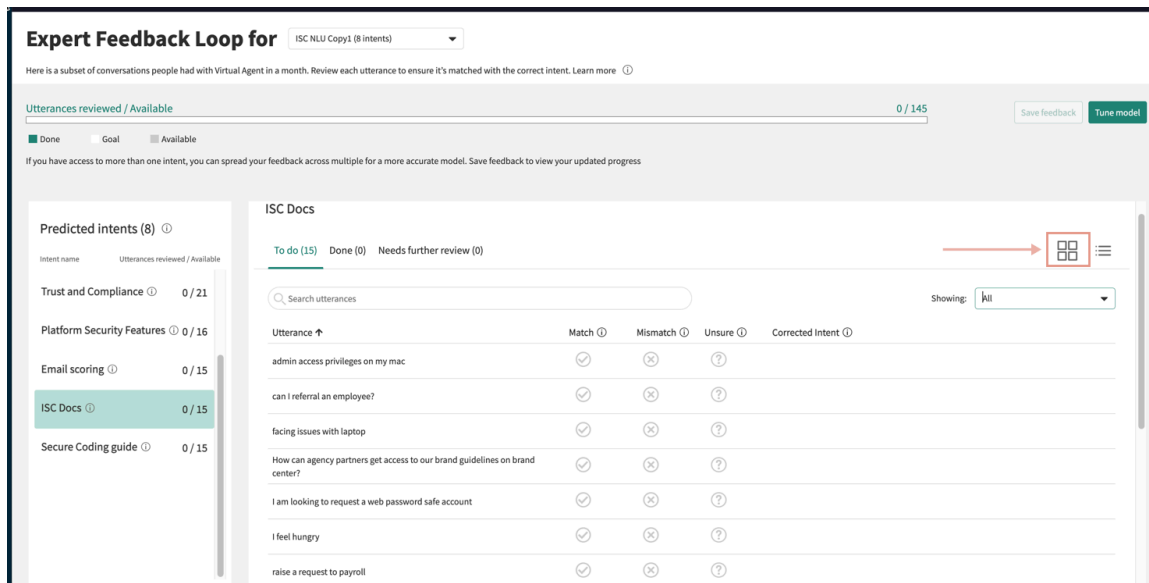
The utterances in the **ad\_password\_change** intent are easy to mark as a match to its predicted intent because there's at least one utterance within the intent that has the word password in it.

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.

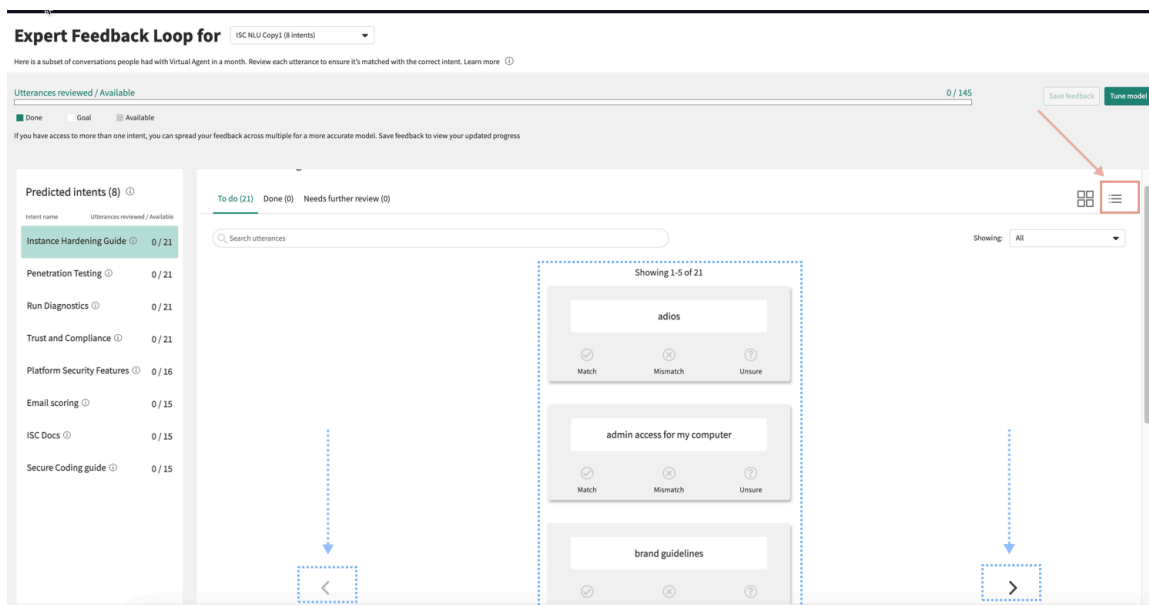
If you want to see your utterances grouped by their value, click the **Showing** filter and select one of the options in its prompt. The values are **All**, **NLU\_Match**, and **Unsure**.

## 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.



## 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.

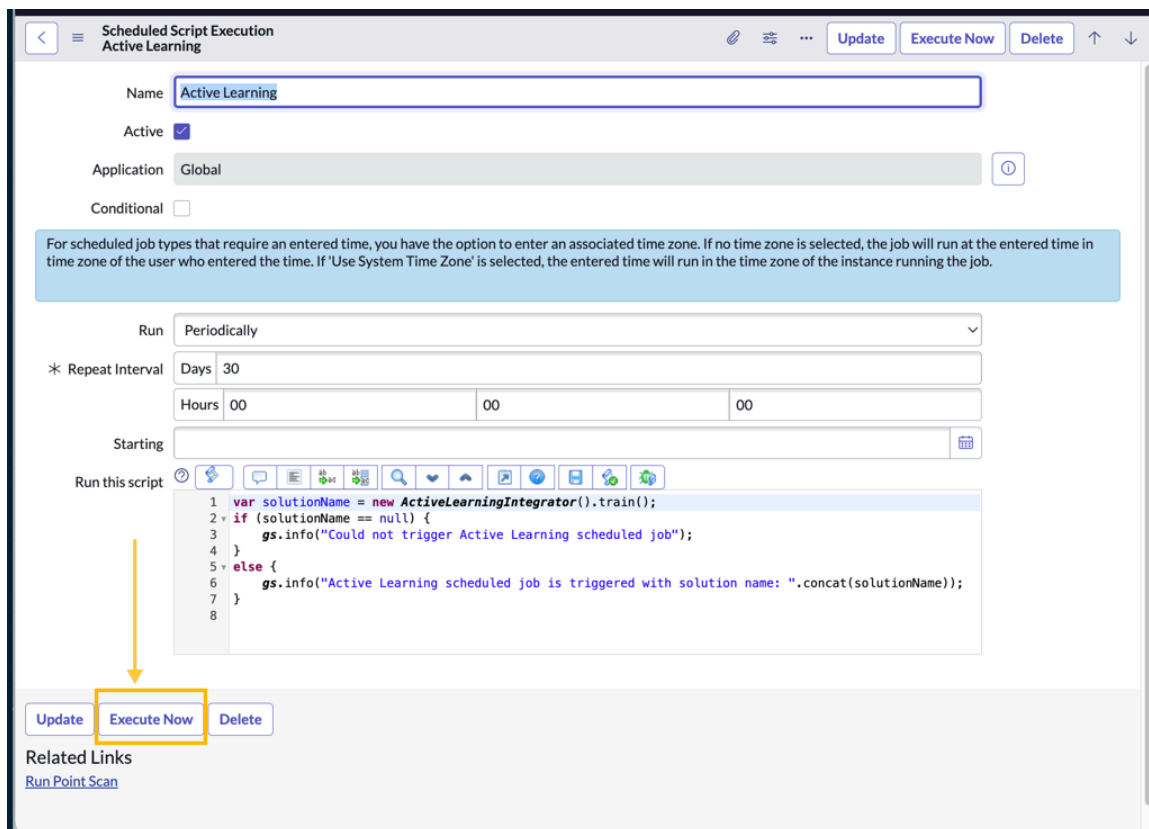
## Procuring additional VA feedback data on demand

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**.

Name	Active	Class	Updated
0f3e357a1b9b8110c353a710604bcb3	true	Training Request Schedule	2022-05-10 00:21:59
506cf12787321300f0187c736cb0b80	true	Scheduled Script Execution	2018-10-05 10:07:58
76f49c491b83c110c353a710604bcb39	true	Training Request Schedule	2022-04-21 19:57:02
Activate report view ACLs for Employee Profile	true	Scheduled Script Execution	2022-04-19 23:14:31
Activate report view ACLs for Templated Snippets	true	Scheduled Script Execution	2022-04-19 23:14:23
Active Learning	true	Scheduled Script Execution	2021-10-19 05:57:18
Activity Stream Reaper	true	Scheduled Script Execution	2020-05-14 04:41:59
ActSub - clean up activity stream	true	Scheduled Script Execution	2021-04-15 04:24:03
Add checkpoint descriptions	false	Scheduled Script Execution	2015-08-11 15:18:01
Add Identifier Fields In Recommended Rules	false	Scheduled Script Execution	2019-04-17 17:30:37

3. Click **Execute Now**.



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.

**i 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.

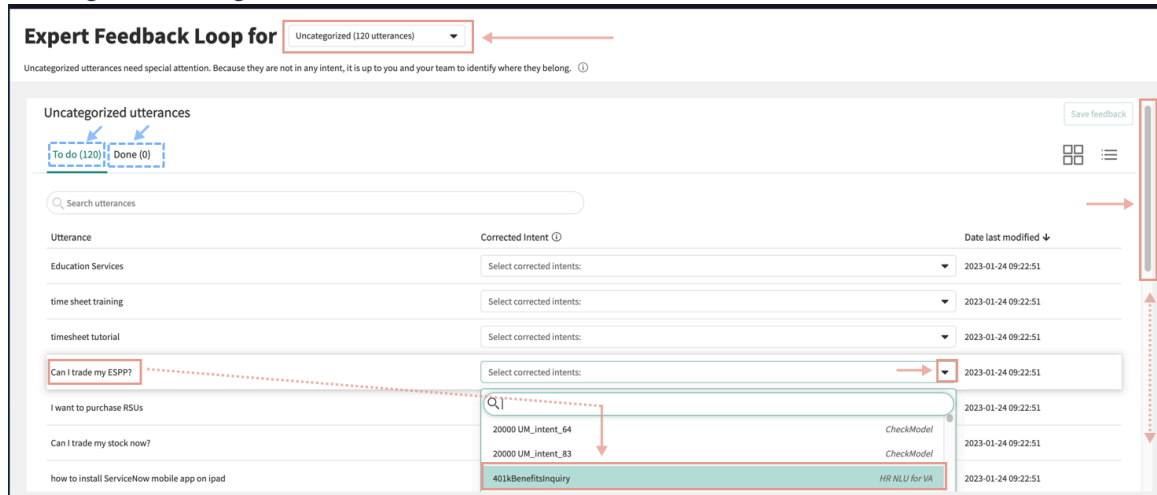
Here are the steps you use to connect an uncategorized utterance to an intent:

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.

**i 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



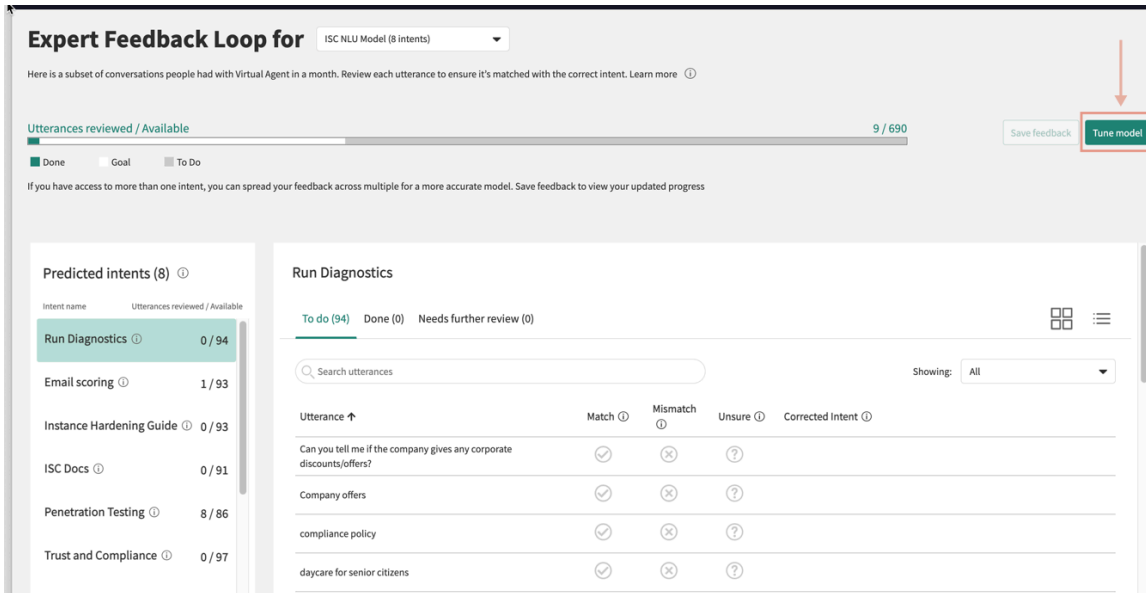
Regarding the image above, let's cover a few items that help describe how the uncategorized utterances page works, as follows. The **To do** section collects utterances that must be connected to an intent. On this particular day we see there are 120 such utterances. Note that the **Done** section on 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 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.

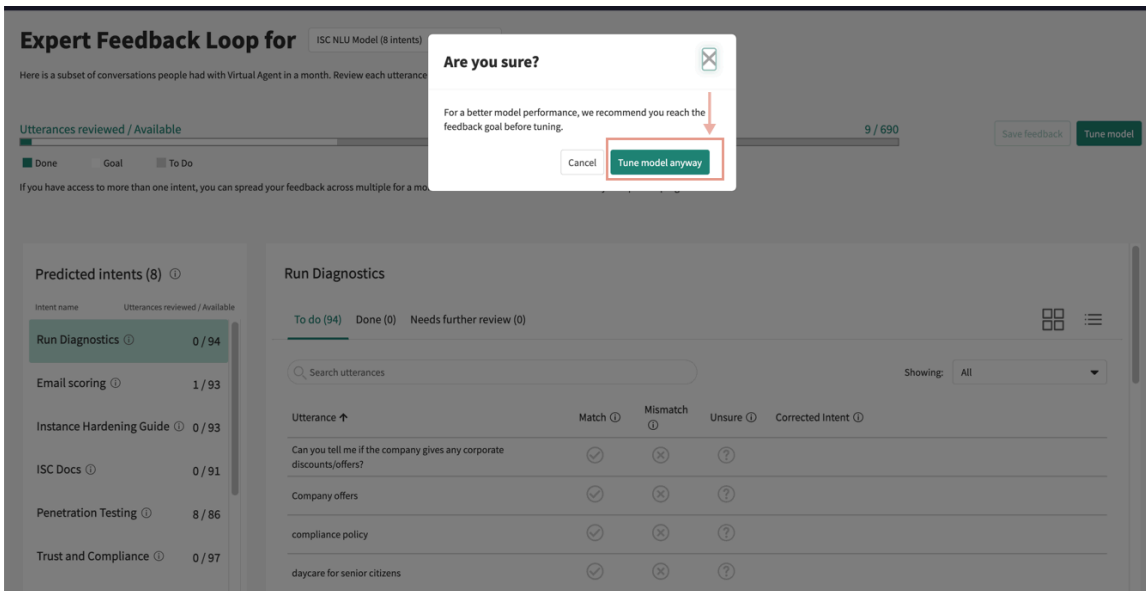
### Expert Feedback Loop data in the Tune Model phase

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.

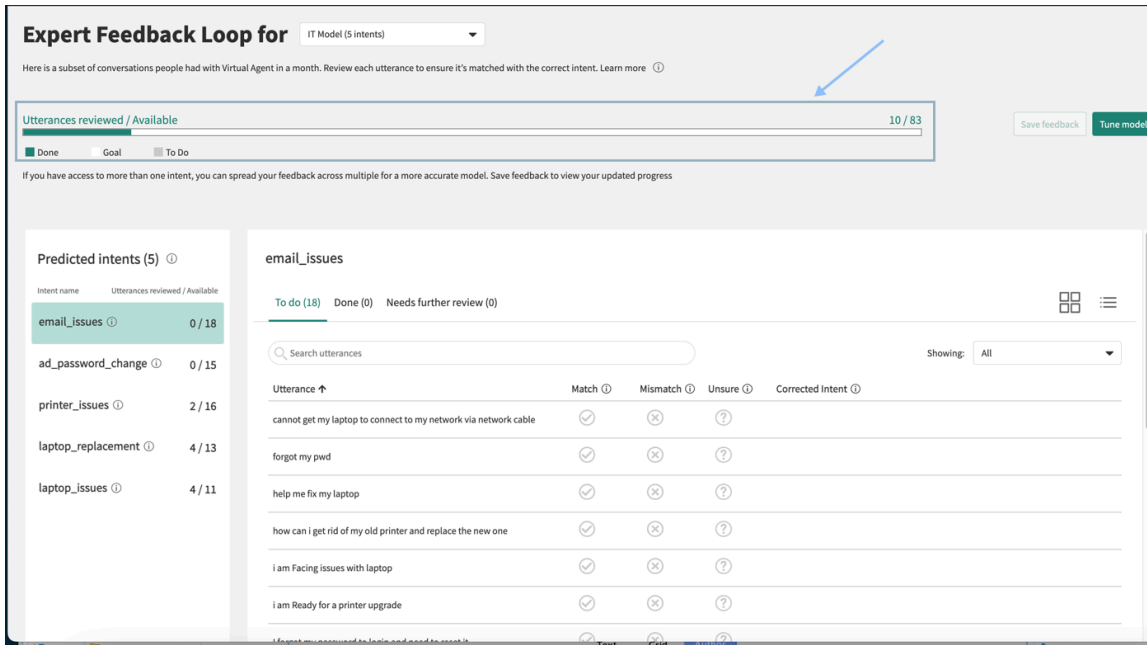


After reviewing utterances in the Expert Feedback Loop, you can push a portion of the feedback data to the default test set of your model. These utterances are then directly added to the **Test utterances** tab of your model. This helps you to continuously maintain and update your test set with real 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.

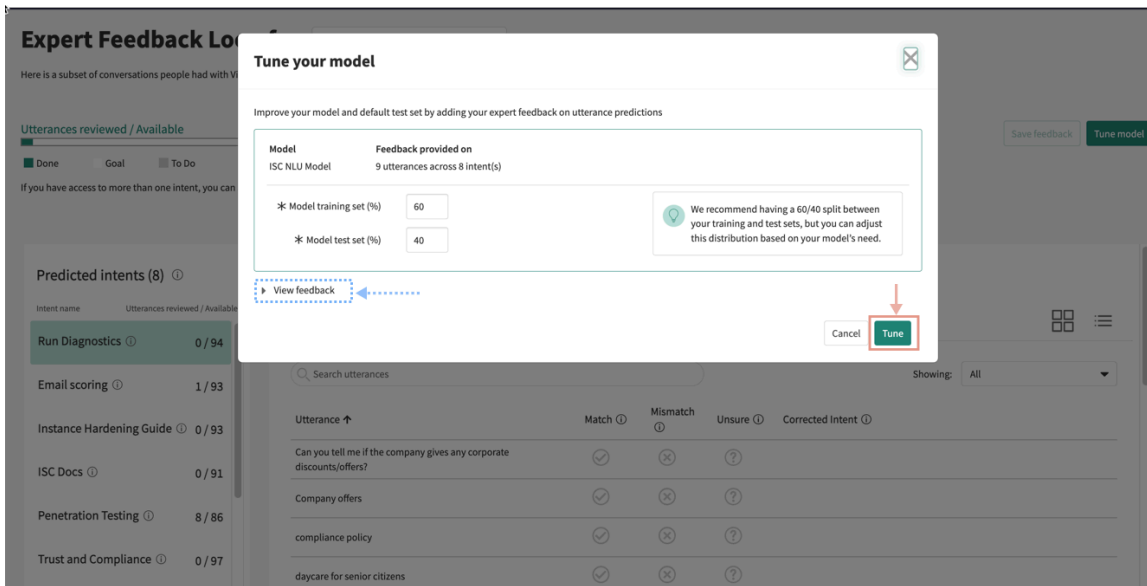


As you keep marking and saving your feedback data, you can see the ratio of reviewed utterances (in green), and those that are still available for review (in white) on the progress bar image below.



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.

- 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.

Name	Next action	Trigger type	Job ID	State
506ef12787321300f019f7c736cb0b80	2023-01-31 19:52:58	Repeat	(empty)	Ready
<b>Active Learning</b>	2023-02-09 22:29:28	Repeat	(empty)	Ready
Activity Stream Reaper	2023-01-31 20:39:28	Repeat	(empty)	Ready
ActSub - clean up activity stream	2023-02-01 00:00:00	Daily	(empty)	Ready
ActSub - clean up activity stream	2023-02-01 00:00:00	Daily	(empty)	Ready
Aggregate AgentZero Metrics	2023-02-01 13:15:00	Interval	RunScriptJob	Ready
AI Search Auto Provisioning	2023-01-31 19:54:23	Interval	RunScriptJob	Ready
AIS Clean Up State Configurations	2023-02-01 00:00:00	Daily	(empty)	Ready
AIS Index Purge State Document	2023-02-01 16:00:00	Daily	RunScriptJob	Ready
AIS Signals Migration Job	2023-01-31 19:44:29	Repeat	AisSignalsMigrationJob	Ready
Analytics - Generate User Hashes	2023-02-01 10:30:00	Daily	(empty)	Ready
Analytics - Register Analytics Buckets w...	2023-01-31 19:56:22	Repeat	(empty)	Ready
Analytics - UX Metrics Apps Registration	2023-02-01 10:33:00	Daily	(empty)	Ready
Analyze PA scores and PA snapshot tables	2023-02-01 15:00:00	Daily	(empty)	Ready
API Monthly Requestor Stats	2023-02-01 02:00:00	Daily	(empty)	Ready
API Monthly Stats	2023-02-02 02:27:00	Day in Month	(empty)	Ready
ApnsTokenRefresh	2023-01-31 19:43:12	Interval	ApnsTokenRefreshJob	Ready
App Operation Queue Health Monitor	2023-01-31 19:42:00	Interval	RunScriptJob	Ready
App Status Count	2023-02-01 00:48:49	Repeat	(empty)	Ready
Application Service Manual Ep Cleanup	2023-01-31 20:21:50	Repeat	(empty)	Ready

2. On the Schedule page, click **Active Learning**. A record for the Active Learning job appears, as shown in the image below.

Name: Active Learning

Active:

Application: Global

Conditional:

For scheduled job types that require an entered time, you have the option to enter an associated time zone. If no time zone is selected, the job will run at the entered time in time zone of the user who entered the time. If 'Use System Time' is selected, the entered time will run in the time zone of the instance running the job.

Run: Periodically

\* Repeat Interval: Days 30, Hours 00, 00, 00

Starting: [Calendar Widget]

```

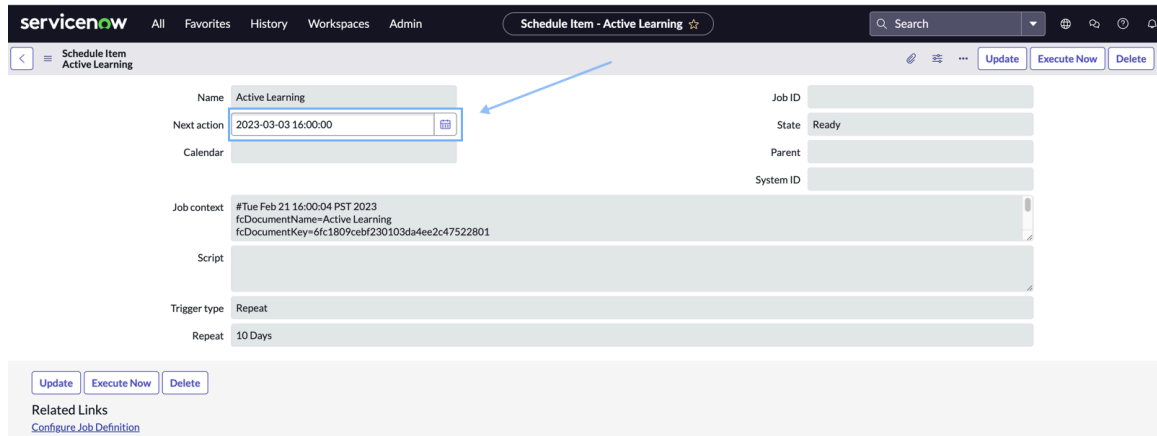
1 var solutionName = new ActiveLearningIntegrator().train();
2 if (solutionName == null) {
3   gs.info("Could not trigger Active Learning scheduled job");
4 }
5 else {
6   gs.info("Active Learning scheduled job is triggered with solution name: ".concat(solutionName));
7 }
8 
```


Update Execute Now Delete

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.

- 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.

### Summary usage and roles

Use the `nlu_admin` or `admin` role to access IAR Tuning in the NLU Workbench. Accessing IAR Tuning in the NLU Workbench requires at least the `nlu_feedback_admin` role (the `nlu_admin` role contains the `nlu_feedback_admin` role). Also, the `virtual_agent_admin` role contains the `nlu_admin` role.

Selecting a IAR model name in the list takes you to tuning. For Issue Auto Resolution, you are not 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.

## Exploring the NLU Workbench

### NLU Workbench

Create, manage, and tune your NLU models to better understand what your users are saying. [Learn more](#)

Virtual Agent (VA) Issue Auto Resolution (IAR) AI Search

#### Virtual Agent models

Show Prebuilt Models

All languages ▾

[+ Create new model](#)

	Model	Status	Used In VA	Model Type	Enabled Intents	Mapped Intents	Last Modified	Last Published
>	ITSM model for Virtual Agent <sup>2</sup>							
	English(Primary) <a href="#">↗</a>	Draft Saved	No	Custom	6	0	2023-06-22 12:09:54	
	Brazilian Portuguese <a href="#">↗</a>	Draft Saved	No	Custom	6	0	2023-06-24 14:22:21	
>	Setup Topics Model <sup>17</sup>							
>	ITSM NLU for Virtual Agent <sup>11</sup>							
>	[Read Only] Password Reset NLU Model <sup>6</sup>							
>	VM Management NLU for Virtual Agent CAI <sup>11</sup>							
>	NLU Common Entities <sup>5</sup>							

< 1 >

- Helpful resources
- Need more context? We've got a variety of resources to help.
- [Conversational Interfaces Academy video series](#)
  - [NLU fundamentals course](#)
  - [VA activation quick start guide](#)
  - [VA and NLU quick start guide](#)
  - [NLU FAQ and troubleshooting guide](#)
  - [Best practices on building NLU models](#)
  - [Best practices on using vocabulary](#)
  - [Best practices on testing NLU models](#)
  - [Best practices on single vs. multiple models](#)
  - [NLU product documentation](#)
  - [Virtual Agent & NLU community forum](#)
- Watch and learn

#### Boost your model performance

**Tune your model**

Provide and gather feedback on real end-user chats from your org's subject matter experts.

[Go to Expert Feedback Loop](#)

**Discover new intents**

Analyze your users chat and task data to identify which intents you should create next.

[Go to Intent Discovery](#)

**Keep chats focused**

Train your model to recognize and ignore phrases it should not respond to.

[Go to Irrelevance detection](#)

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 gray. In the Model column of the list, when you select the caret to the left of the model name, the model changes color from gray 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.

What do you want to tune for? [Explain this](#)



Please choose from one of the options below

<p><b>Precision</b> <span>Recommended</span></p> <p>A model tuned for precision will only match issues with intents when it's very confident. This will result in fewer errors, but less automation.</p>	<p>Match rate</p> <p><b>71%</b></p> <p>Change: +4%</p>	<p>IAR coverage</p> <p><b>58%</b></p> <p>Change: -7%</p>
<p><b>Automation</b></p> <p>A model tuned for automation will match issues with intents more frequently. This will result in more automation, but also more errors.</p>	<p>Match rate</p> <p><b>62%</b></p> <p>Change: -5%</p>	<p>IAR coverage</p> <p><b>76%</b></p> <p>Change: +11%</p>
<p><b>Balance</b></p> <p>A model tuned for balance will attempt to equalize matching issues with intents and limiting errors.</p>	<p>Match rate</p> <p><b>67%</b></p> <p>Change: +0%</p>	<p>IAR coverage</p> <p><b>65%</b></p> <p>Change: +0%</p>

Cancel Done

### 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**.

## 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.

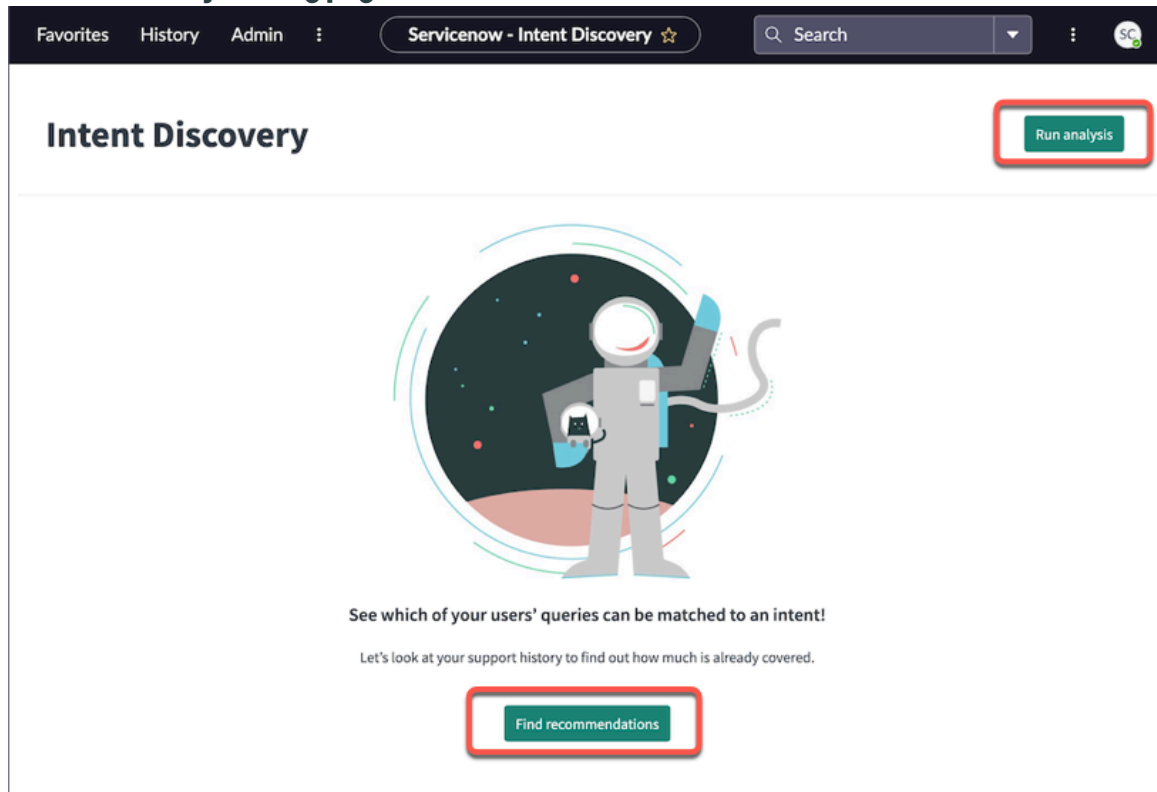
## Intent Discovery report details

- 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



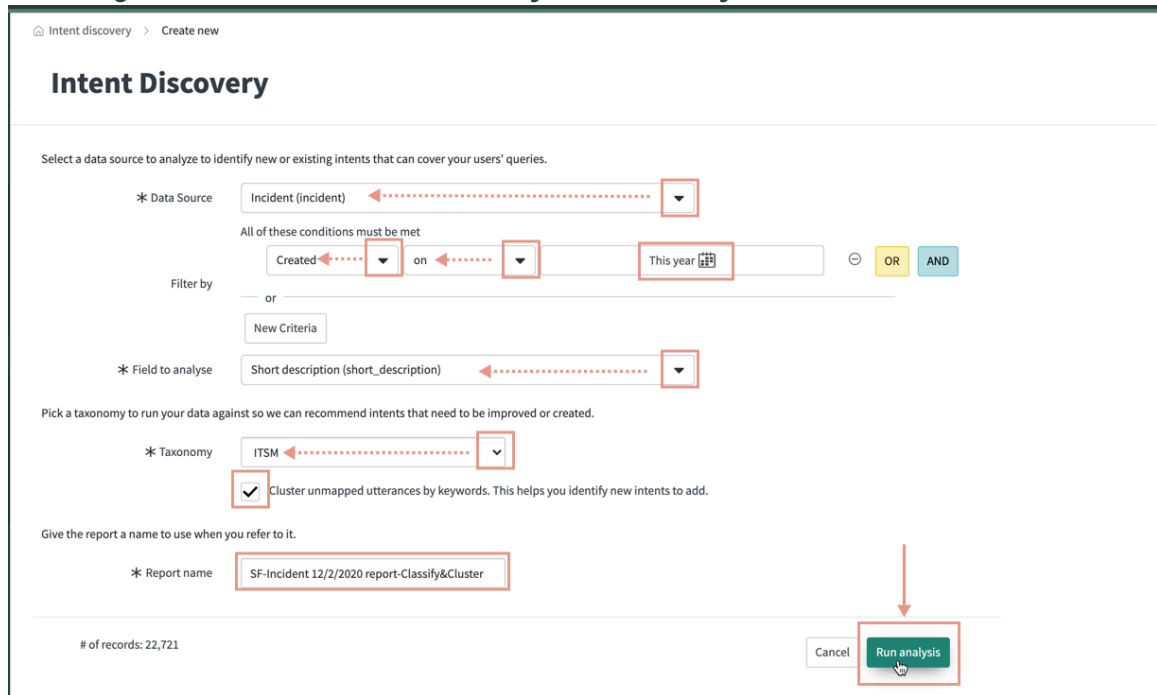
## Running an analysis on the report

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]**

- 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.

## An ongoing run analysis

### Intent Discovery Run analysis

Check on your latest analyses or run a new one to get recommendations on intents that cover queries you're getting from your users. Each of these reports is an analysis of a dataset you identified against a prebuilt model.

Name	Data source	% covered	Recommended intents	Not covered (clustered) ⓘ	Total records	Run date	Status
SF-Incident 12/2/2020 report-Classify&Cluster	Incident	--	--	--	--		Clustering...
No test messages	Incident	38%	122	1	26,737	2020-11-23 13:20:18	Done
Incident 11/20/2020 Cluster Only	Incident	38%	122	224	27,106	2020-11-20 10:26:41	Done
Incident 11/20/2020 Classification Only	Incident	38%	122	1	27,106	2020-11-20 10:20:34	Done
Brad Incident 11/19/2020	Incident	38%	122	224	27,106	2020-11-20 04:08:32	Done
Mimi test	Incident	38%	126	224	27,106	2020-11-12 16:37:07	Done
Brad 3	Incident	--	0	0	0		
Brad Incident 11/9/2020 - last 6 months	Incident	40%	65	170	17,081	2020-11-10 02:41:48	Done

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

### Intent Discovery Run analysis

Check on your latest analyses or run a new one to get recommendations on intents that cover queries you're getting from your users. Each of these reports is an analysis of a dataset you identified against a prebuilt model.

Name	Data source	% covered	Recommended intents	Not covered (clustered) ⓘ	Total records	Run date	Status
SF-Incident 12/2/2020 report-Classify&Cluster	Incident	42%	107	208	22,721	2020-12-02 15:58:34	Done <span style="float: right; font-size: 0.7em; border: 1px solid #ccc; padding: 2px;">Delete report</span>
No test messages	Incident	38%	122	1	26,737	2020-11-23 13:20:18	Done
Incident 11/20/2020 Cluster Only	Incident	38%	122	224	27,106	2020-11-20 10:26:41	Done
Incident 11/20/2020 Classification Only	Incident	38%	122	1	27,106	2020-11-20 10:20:34	Done
Brad Incident 11/19/2020	Incident	38%	122	224	27,106	2020-11-20 04:08:32	Done
Mimi test	Incident	38%	126	224	27,106	2020-11-12 16:37:07	Done
Brad 3	Incident	--	0	0	0		
Brad Incident 11/9/2020 - last 6 months	Incident	40%	65	170	17,081	2020-11-10 02:41:48	Done

**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.

## Reviewing a recommended intent

Intent discovery > Incident 12/16/2020 - SF Test

### Incident 12/16/2020 - SF Test

Data source	Filter	Field	Taxonomy	Total records	Run date	
Incident	Created on This quarter	Short description	ITSM	10,009	2020-12-16 14:30:00	<a href="#">Run again</a>

**27%** Records covered by recommendations      **73%** Remaining records

Start by selecting intents and adding them to an existing or new model. Show Additional ▾

Records matched ↓	Intent	Content work needed	Sample utterances	Added to
521 (5.2%)	ITSM - Issue - Email *	Review topic flow and utterances	54	
375 (3.7%)	ITSM - Request - Networking - MFA - Token *	Review topic flow and utterances	28	
266 (2.7%)	ITSM - Issue - Collaboration Software Issues *	Review topic flow and utterances	113	
242 (2.4%)	ITSM - Issue - Hardware - Troubleshoot *	Review topic flow and utterances	33	
216 (2.2%)	ITSM - Request - Data - User - Account - Password *	Review topic flow and utterances	43	
162 (1.6%)	ITSM - Issue - VPN *	Review topic flow and utterances	21	

**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

Intent discovery > Incident 12/16/2020 - SF Test

### Incident 12/16/2020 - SF Test

Data source	Filter	Field	Taxonomy	Total records	Run date	
Incident	Created on This quarter	Short description	ITSM	10,009	2020-12-16 14:30:00	<a href="#">Run again</a>

**27%** Records covered by recommendations      **73%** Remaining records

Start by selecting intents and adding them to an existing or new model. Show Additional ▾

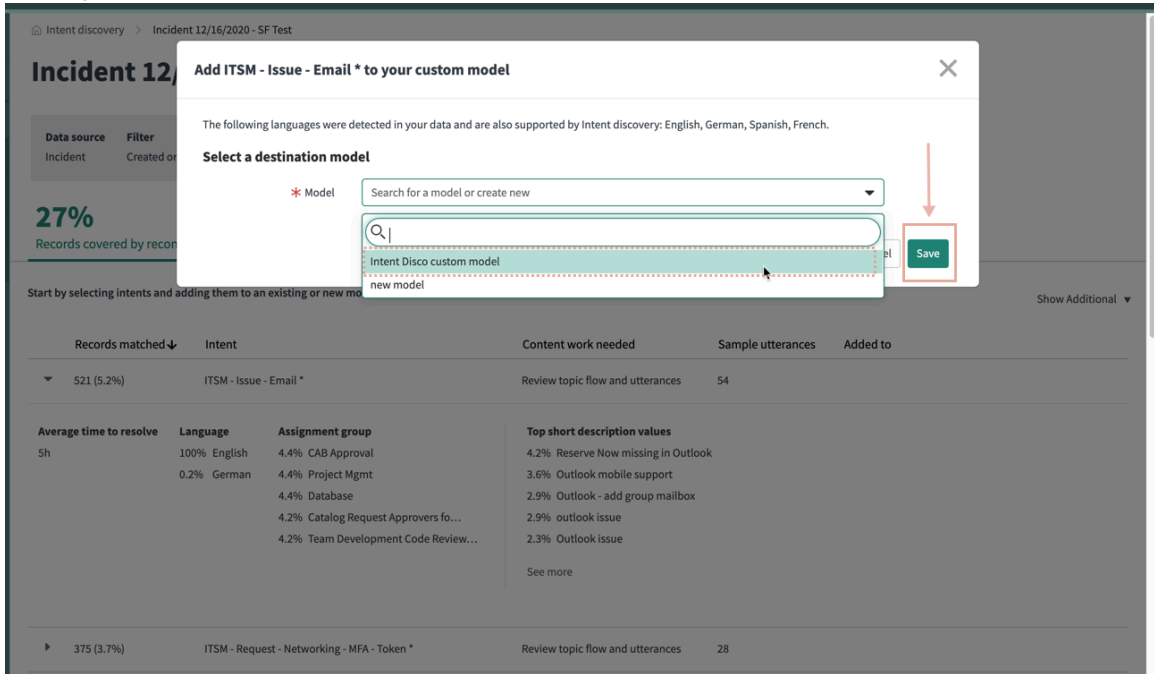
Records matched ↓	Intent	Content work needed	Sample utterances	Added to								
521 (5.2%)	ITSM - Issue - Email *	Review topic flow and utterances	54	<input checked="" type="checkbox"/> Add to model <input type="checkbox"/> Ignore								
<table border="0"> <tr> <td><b>Average time to resolve</b></td> <td><b>Language</b></td> <td><b>Assignment group</b></td> <td><b>Top short description values</b></td> </tr> <tr> <td>5h</td> <td>100% English 0.2% German</td> <td>4.4% CAB Approval 4.4% Project Mgmt 4.4% Database 4.2% Catalog Request Approvers fo... 4.2% Team Development Code Review...</td> <td>4.2% Reserve Now missing in Outlook 3.6% Outlook mobile support 2.9% Outlook - add group mailbox 2.9% outlook issue 2.3% Outlook issue See more</td> </tr> </table>					<b>Average time to resolve</b>	<b>Language</b>	<b>Assignment group</b>	<b>Top short description values</b>	5h	100% English 0.2% German	4.4% CAB Approval 4.4% Project Mgmt 4.4% Database 4.2% Catalog Request Approvers fo... 4.2% Team Development Code Review...	4.2% Reserve Now missing in Outlook 3.6% Outlook mobile support 2.9% Outlook - add group mailbox 2.9% outlook issue 2.3% Outlook issue See more
<b>Average time to resolve</b>	<b>Language</b>	<b>Assignment group</b>	<b>Top short description values</b>									
5h	100% English 0.2% German	4.4% CAB Approval 4.4% Project Mgmt 4.4% Database 4.2% Catalog Request Approvers fo... 4.2% Team Development Code Review...	4.2% Reserve Now missing in Outlook 3.6% Outlook mobile support 2.9% Outlook - add group mailbox 2.9% outlook issue 2.3% Outlook issue See more									
375 (3.7%)	ITSM - Request - Networking - MFA - Token *	Review topic flow and utterances	28									
266 (2.7%)	ITSM - Issue - Collaboration Software Issues *	Review topic flow and utterances	113									

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.

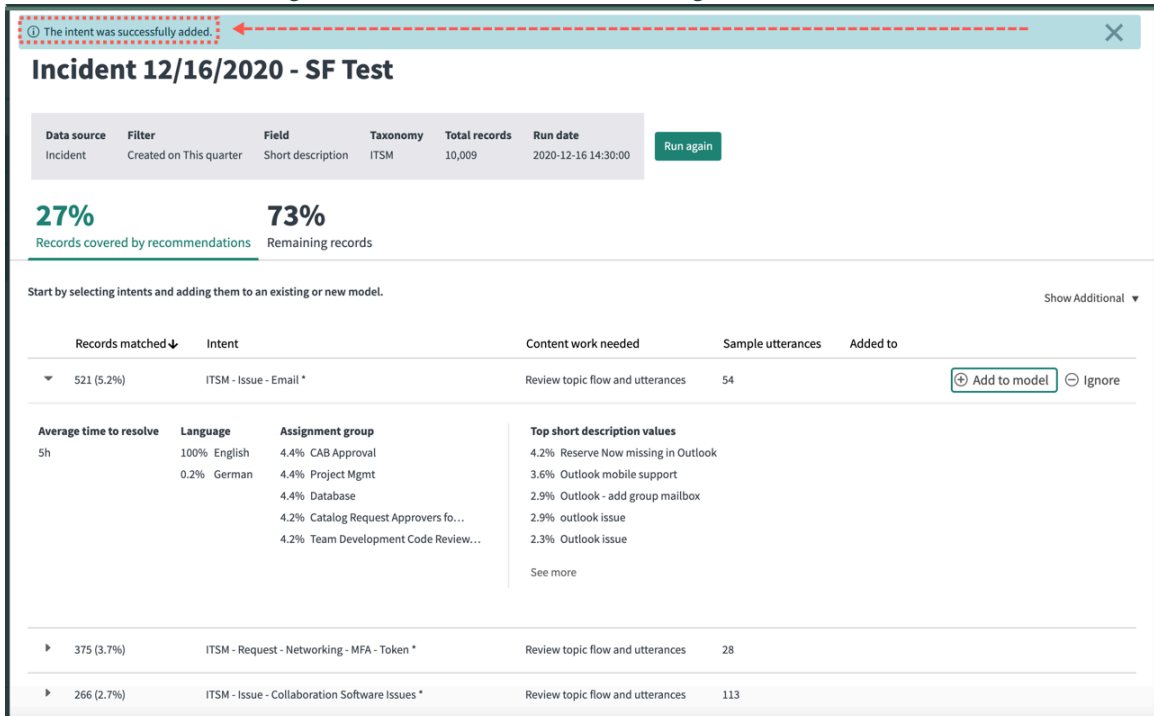
4. Select **Save**.

**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.

## View a recommended intent in the target model

Home > Intent Disco custom model Published 24 minutes ago

**Intent Disco custom model** Delete Clone Train Test Publish

Intents (2) Entities (8) Vocabulary (0) Settings

Review your existing intents or create new ones.

+ New Intent ↓ Import Intents

Name	Utterances	Associated Entities	Mapped objects	Created by	Last updated
#EmailSetup	55	1	--	admin	2020-12-16 15:15:24
#EmailIssues	55	0	--	admin	2020-12-16 15:12:09

## 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

Intent discovery > Incident 12/16/2020 - SF Test

**Incident 12/16/2020 - SF Test**

Data source	Filter	Field	Taxonomy	Total records	Run date	Run again
Incident	Created on This quarter	Short description	ITSM	10,009	2020-12-16 14:30:00	Run again

27% Records covered by recommendations 73% Remaining records

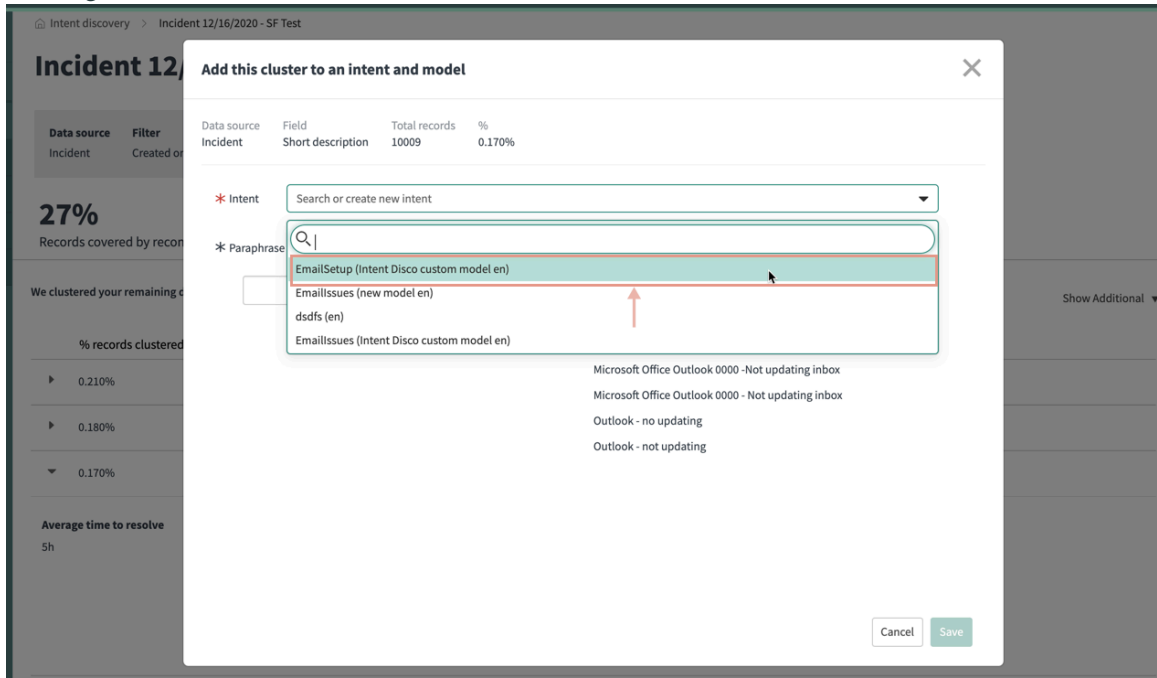
We clustered your remaining data to show you common records. Select to add to new or existing intents. Show Additional

% records clustered	# of records	Example	Added to
▶ 0.210%	21	Preferred Name Change	
▶ 0.180%	18	!! From AP Invoice queue to HelpDesk Acme	
<input checked="" type="checkbox"/> 0.170%	17	Outlook not updating	<input checked="" type="checkbox"/> Add to intent <input type="checkbox"/> Ignore
▶ 0.150%	15	C2 Imaging Invoice SI0000143556	

Average time to resolve	Language	Assignment group	Top short description values
5h	100% English	12% US Presidents Group 1 5.9% ATF Service Level Management... 5.9% Project Mgmt 5.9% Catalog Request Approvers > ... 5.9% ATF_TestGroup_ServiceDesk	41% Outlook not updating 24% Microsoft Outlook-Inbox Not Updating 12% Microsoft Office Outlook 0000 -Not updating inbox 12% Microsoft Office Outlook 0000 - Not updating in... 5.9% Outlook - no updating

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.

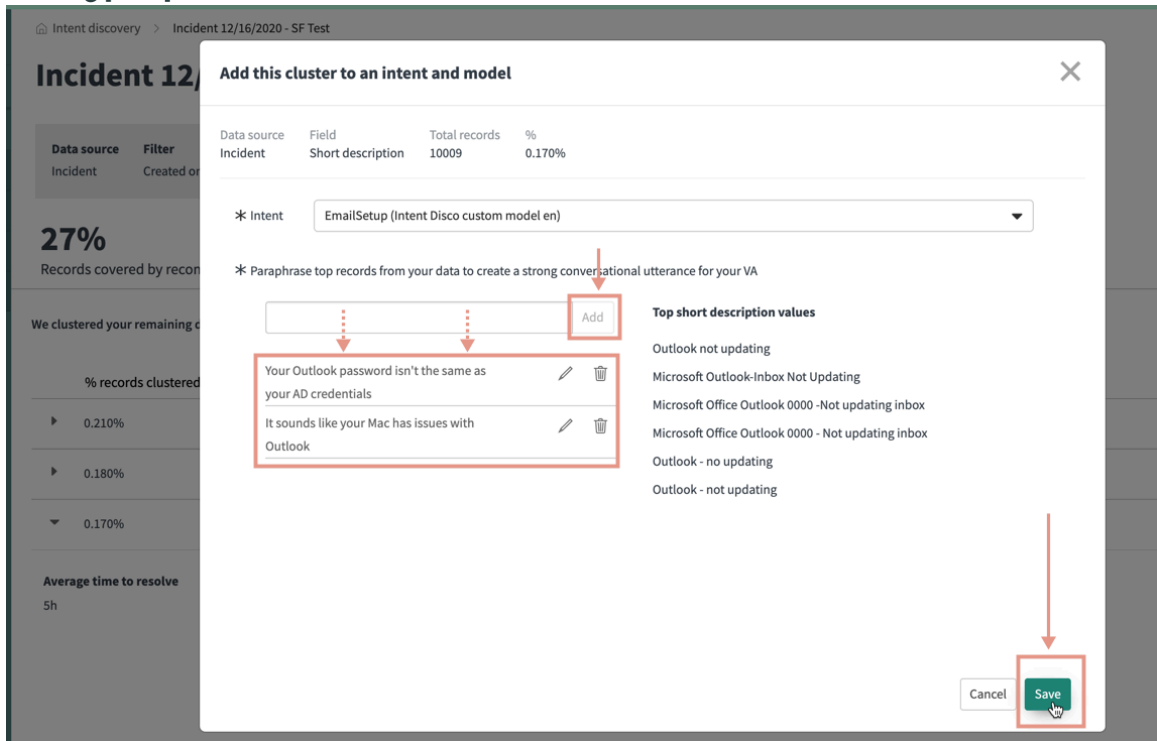
## Adding a cluster to an intent and model



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.

## Confirmation of adding paraphrased utterances to an intent

Added 2 utterances to "EmailSetup (Intent Disco custom model en)". View and edit these in the Models tab on the left or through the link in the row below.

### Incident 12/16/2020 - SF Test

Data source	Filter	Field	Taxonomy	Total records	Run date	Run again
Incident	Created on This quarter	Short description	ITSM	10,009	2020-12-16 14:30:00	Run again

**27%** Records covered by recommendations    **73%** Remaining records

We clustered your remaining data to show you common records. Select to add to new or existing intents. Show Additional ▾

% records clustered	# of records ↓	Example	Added to
▶ 0.210%	21	Preferred Name Change	
▶ 0.180%	18	!! From AP Invoice queue to HelpDesk Acme	
☑ 0.170%	17	Outlook not updating	Intent Disco custom model > EmailSetup

Average time to resolve	Language	Assignment group	Top short description values
5h	100% English	12% US Presidents Group 1 5.9% ATF Service Level Management... 5.9% Project Mgmt 5.9% Catalog Request Approvers > ... 5.9% ATF_TestGroup_ServiceDesk	41% Outlook not updating 24% Microsoft Outlook-Inbox Not Updating 12% Microsoft Office Outlook 0000 - Not updating inbox 12% Microsoft Office Outlook 0000 - Not updating in... 5.9% Outlook - no updating

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

We clustered your remaining data to show you common records. Select to add to new or existing intents. Show Additional

% records clustered	# of records ↓	Example	Added to
▼ 12.12%	8	Seem to have an issue with my hard drive...	

Average time to resolve	Language	Assignment group	Top short description values
Not available	100% English	None available	13% Seem to have an issue with my hard drive... 13% Having problems with Sales Tools performance 13% Network file shares access issue 13% CPU load high for over 10 minutes 13% Wireless access is down in my area

▼ 7.576%	5	EMAIL Server Down Again	<span>⊕ Add to intent</span> <span>⊖ Ignore</span>
----------	---	-------------------------	--

Show Additional

Clusters with added intents

Ignored clusters

## Running another analysis on your Intent Discovery report

1. Select **Run Again**.

## Selecting the version of the analysis to run

Added 2 utterances to "EmailSetup (Intent Disco custom model en)". View and edit these in the Models tab on the left or through the link in the row below. ✕

### Incident 12/16/2020 - SF Test

Data source	Filter	Field	Taxonomy	Total records	Run date	
Incident	Created on This quarter	Short description	ITSM	10,009	2020-12-16 14:30:00	<b>Run again</b>

**27%** Records covered by recommendations      **73%** Remaining records

We clustered your remaining data to show you common records. Select to add to new or existing intents. Show Additional ▾

% records clustered	# of records ↓	Example	Added to
▶ 0.210%	21	Preferred Name Change	
▶ 0.180%	18	!! From AP Invoice queue to HelpDesk Acme	
▼ 0.170%	17	Outlook not updating	<a href="#">Intent Disco custom model &gt; EmailSetup</a>

Average time to resolve	Language	Assignment group	Top short description values
5h	100% English	12% US Presidents Group 1	41% Outlook not updating
		5.9% ATF Service Level Management...	24% Microsoft Outlook-Inbox Not Updating
		5.9% Project Mgmt	12% Microsoft Office Outlook 0000 -Not updating inbox
		5.9% Catalog Request Approvers > ...	12% Microsoft Office Outlook 0000 - Not updating in...
		5.9% ATF_TestGroup_ServiceDesk	5.9% Outlook - no updating

**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

Added 2 utterances to "EmailSetup (Intent Disco custom model en)". View and edit these in the Models tab on the left or through the link in the row below. ✕

### Incident 12/16/2020 - SF Test

Data source	Filter	Field	Taxonomy	Total records	Run date	
Incident	Created on This quarter	Short description	ITSM	10,009	2020-12-16 14:30:00	<b>In progress...</b> <b>Cancel run</b>

**27%** Records covered by recommendations      **73%** Remaining records

We clustered your remaining data to show you common records. Select to add to new or existing intents. Show Additional ▾

% records clustered	# of records ↓	Example	Added to
▶ 0.210%	21	Preferred Name Change	
▶ 0.180%	18	!! From AP Invoice queue to HelpDesk Acme	
▼ 0.170%	17	Outlook not updating	<a href="#">Intent Disco custom model &gt; EmailSetup</a>

Average time to resolve	Language	Assignment group	Top short description values
5h	100% English	12% US Presidents Group 1	41% Outlook not updating
		5.9% ATF Service Level Management...	24% Microsoft Outlook-Inbox Not Updating
		5.9% Project Mgmt	12% Microsoft Office Outlook 0000 -Not updating inbox
		5.9% Catalog Request Approvers > ...	12% Microsoft Office Outlook 0000 - Not updating in...
		5.9% ATF_TestGroup_ServiceDesk	5.9% Outlook - no updating

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**.

## Selecting the new version of the report

Added 2 utterances to "EmailSetup (Intent Disco custom model en)". View and edit these in the Models tab on the left or through the link in the row below.

You have a new report available. Select the latest from dropdown.

Data source	Filter	Field	Taxonomy	Total records	Run date
Incident	Created on This quarter	Short description	ITSM	10,009	2020-12-16 14:30:00

27% Records covered by recommendations | 73% Remaining records

We clustered your remaining data to show you common records. Select to add to new or existing intents.

% records clustered	# of records	Example	Added to
0.210%	21	Preferred Name Change	
0.180%	18	!! From AP Invoice queue to HelpDesk Acme	
0.170%	17	Outlook not updating	Intent Disco custom model > EmailSetup

**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

### Intent Discovery

Check on your latest analyses or run a new one to get recommendations on intents that cover queries you're getting from your users. Each of these reports is an analysis of a dataset you identified against a prebuilt taxonomy.

Name	Data source	Covered by recommended intents	Not covered (clustered)	Total records	Run date	Status
Incident 12/16/2020 - SF Test	Incident	62	170	10,009	2020-12-16 19:41:02	Done
Incident 12/14/2020 C	Incident	64	169	10,066	2020-12-15 09:03:21	Done
Incident 12/14/2020	Incident	0	0	0		

## 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](#).

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](#) website to view all the available apps and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the [ServiceNow Store version history release notes](#).

3. 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.





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

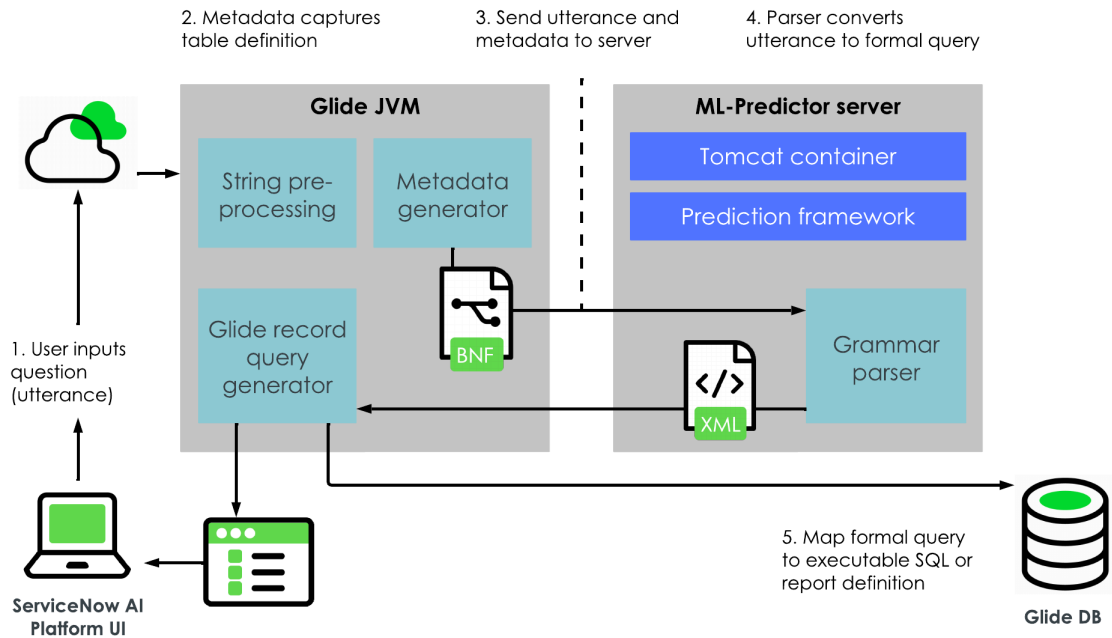
Choose one of these tiles to get started.

<p style="text-align: center;"><b>Explore</b></p>  <p style="text-align: center;">Learn about NLQ concepts and features.</p>	<p style="text-align: center;"><b>Use</b></p>  <p style="text-align: center;">Use NLQ to query your data with requests in natural language.</p>
<p style="text-align: center;"><b>Configure</b></p>  <p style="text-align: center;">Set up NLQ for your environment.</p>	<p style="text-align: center;"><b>References</b></p>  <p style="text-align: center;">Get details about properties and roles installed with NLQ.</p>

## 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.

## NLQ data flow for user input and processing



ServiceNow<sup>®</sup> 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](#).

## 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	<a href="#">Create a report with Analytics Q&amp;A</a> ↗
Configuration Management Database (CMDB)	<p>Query your CMDB data without needing to know table relationships or data structures.</p> <ul style="list-style-type: none"> <li>• <a href="#">Querying the CMDB</a> ↗</li> <li>• <a href="#">Intelligent Search for CMDB</a> ↗</li> </ul> <p>English is the only supported language for CMDB.</p>
NLQ with AI Search in global search	<a href="#">NLQ Genius Results</a> ↗ 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.

## 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

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.

**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	<p>Type of synonym.</p> <ul style="list-style-type: none"> <li>○ <b>Table or Column:</b> Synonym for a value on a table or a specific column in that table.</li> <li>○ <b>Record:</b> Synonym for a value on the CMDB tables [cmdb_rel_type_table]. For more information, see <a href="#">Querying the CMDB</a> .</li> </ul> <p><b>Note:</b> Supports cmdb_rel_type,cmdb_group, cmdb_ci_service_technical, cmdb_ci_service_discovered, cmdb_ci_query_based_service tables.</p>
Application	<p>[Read only] Application scope that can use the synonym in a query. The default is <b>Global</b>.</p> <p>When creating a synonym, ensure that you are in the scope you want for the synonym.</p>
Table	<p>Source table.</p> <p>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.</p>
Column name	<p>Specific column on the source table.</p>
Language	<p>Language of the synonym. Should match the language of the source value.</p>
Synonyms	<p>Words or short phrases that the system should map to column names or tables, when converting the utterance to a formal query.</p> <p>Separate multiple synonyms with a comma.</p>
Extend to All Child Tables	<p><b>Table or Column</b> type only.</p> <p>Select this option to make any child tables of the source table inherit the synonyms.</p>

4. Select **Submit** if new, or **Update** if you are modifying an existing record.

## 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.

The screenshot shows the 'NLQ Synonym New record' configuration page. The fields are as follows:

- Synonym Type:** Table or Column
- Application:** Global
- Table:** Change Request [change\_request]
- Column name:** Active
- Language:** en
- Synonyms:** open, in progress
- Extend To All Child Tables:**

A 'Submit' button is located at the bottom left of the form area.

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.

- 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 <b>Global</b> . 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_to ISEMPTY`.

Update Delete

\* Table

\* Conditions Add Sort

All of these conditions must be met
   
 
OR AND

New Criteria

\* Synonyms

Application

Update Delete

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.

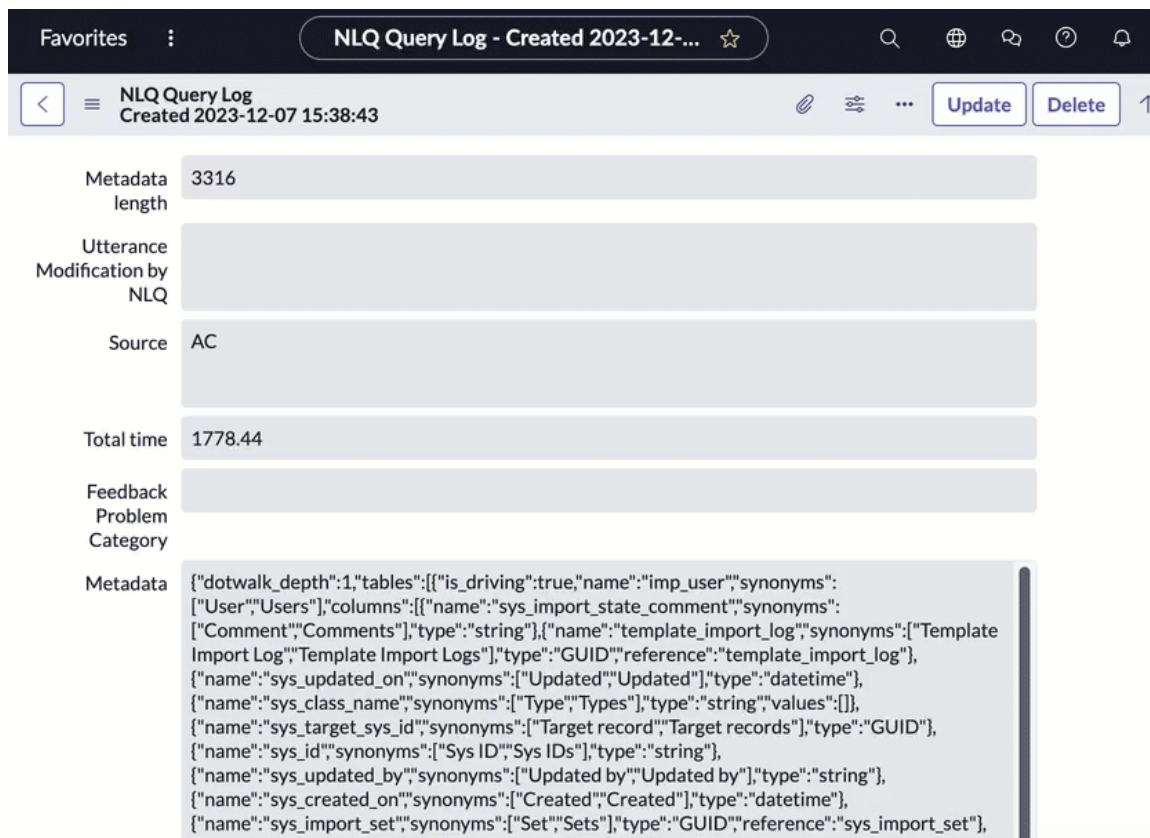
- **Output Source**: how the results were generated. The value **BNF** indicates a rules-based method. The value **GAI** indicates the Now LLM Service (fallback) method.
- **Source**: the location from which the query was initiated. The value **AC** indicates Analytics Center. The value **CMDB\_WS** indicates CMDB Workspace.
- **Utterance**: the original natural-language request which triggered NLQ.

## Procedure

1. Navigate to **All > NLQ > Logs**.

Utterance	Successful	Table	Output	Metadata generation time	Prediction service time	Total time	Error message	User	Metadata	Updated
show me servers running databases	false	Server [cmdb_ci_server]	NO MATCH	360.193	255.975	616.273	We weren't able to understand your reque...	(empty)	["dotwalk_depth":1,"tables":["is_drivin...	2022-07-07 16:48:27
cmdb_ci name show me servers running dat...	true	Configuration Item [cmdb_ci]	["type": "list","sumfield": ""aggreat...	684.908	286.27	971.224		(empty)	["dotwalk_depth":1,"tables":["is_drivin...	2022-07-07 16:48:10
show me linux servers	true	Linux Server [cmdb_ci_linux_server]	["type": "list","sumfield": ""aggreat...	1053.294	411.103	1464.51		(empty)	["dotwalk_depth":1,"tables":["is_drivin...	2022-07-07 16:30:41
show me all servers with tag key abc and...	true		["type":"cmdb","nodes":{"label":"Server...	735.245	343.898	1079.191		(empty)	["dotwalk_depth":1,"tables":["is_drivin...	2022-07-07 16:43:40
show me databases running on servers	true		["type":"cmdb","nodes":{"label":"Databa...	551.146	353.673	904.853		(empty)	["dotwalk_depth":1,"tables":["is_drivin...	2022-07-07 15:28:51
configuration items connected by servers...	true		["type":"cmdb","nodes":{"label":"Config...	1424.053	1070.169	2494.27		(empty)	["dotwalk_depth":1,"tables":["is_drivin...	2022-07-07 17:18:39
Find business applications that consume...	true		["type":"cmdb","nodes":{"label":"Busine...	489.925	449.925	939.898		(empty)	["dotwalk_depth":1,"tables":["is_drivin...	2022-07-07 16:32:53

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.

## 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**

Utterance	Output	Matched words to table names	Confidence scores	Schema info generation time	Computing best matches time	Source
servers in san diego with app se...	cmdb_ci_service_auto.cmdb_ci_service.cmdb...	["servers": ["cmdb_ci_server":"cmdb_ci_vm..."/>	["cmdb_ci_server":"1.0";cmdb_ci_servic...	0.011	42.769	CMDB_WS
servers in san diego with app se...	cmdb_ci_service_auto.cmdb_ci_service.cmdb...	["servers": ["cmdb_ci_server":"cmdb_ci_vm..."/>	["cmdb_ci_server":"1.0";cmdb_ci_servic...	0.004	40.3	CMDB_WS
show all app services name SAP c...	cmdb_ci_server.cmdb_ci_spkg.cmdb_ci_serv...	["server": ["cmdb_ci_server":"cmdb_ci_vm..."/>	["cmdb_ci_bj_ksa":"0.9";cmdb_ci_server"...	133.245	44.837	CMDB_WS
Show all soqtvr installed on ser...	cmdb_ci_service.cmdb_ip_service_ci.cmdb_...	["installed": ["cmdb_software_instance"]...]	["cmdb_ci_server":"0.90909090909091;...]	141.418	48.248	CMDB_WS
Show all soqtvr installed on ser...	cmdb_ci_service.cmdb_ip_service_ci.cmdb_...	["installed": ["cmdb_software_instance"]...]	["cmdb_ci_server":"0.90909090909091;...]	95.738	42.258	CMDB_WS
Show me all Tecical service offe...	cmdb_ci_service.task_service_offering.cm...	["application": ["cmdb_ci_appil"];service...	["cmdb_ci_service":"1.0";cmdb_ci_servic...	172.383	39.617	CMDB_WS

### 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.

Name	Value	Type	Application	Description	Updated	Updated by
com.snc.listv2.nlq.lists.append_query	true	true   false	Global	When true (default), NLQ inputs will alw...	2020-07-24 14:35:37	admin
com.snc.listv2.nlq.lists.enabled	true	true   false	Global	Enable natural language query search opt...	2020-07-14 14:32:37	admin
com.snc.par.nlq.report_designer.enabled	true	true   false	Global	Enable Natural Language Queries (NLQ) in...	2020-06-11 03:29:14	admin
glide.cmdb.query.nlq.activated	false	true   false	Global	Activate query builder NLQ search feature	2021-10-13 14:23:44	admin

Property	Description
com.snc.listv2.nlq.lists.append_query	<ul style="list-style-type: none"> <li>• True: NLQ inputs add onto existing queries via an "and" operator</li> <li>• False: New NLQ input replaces any existing queries</li> </ul>

Property	Description
	<p>Example: You run two queries.</p> <ul style="list-style-type: none"> <li>• Query 1: Incidents with critical priority</li> <li>• Query 2: assigned to John Smith</li> </ul> <p>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.</p>
com.snc.listv2.nlq.lists.enabled	<ul style="list-style-type: none"> <li>• True: Enables NLQ search option for List v2</li> <li>• False: Removes NLQ search option for List v2</li> </ul>
com.snc.nlq.gai_enabled	<ul style="list-style-type: none"> <li>• True: The Now LLM Service fallback is available</li> <li>• False: The Now LLM Service fallback is not available</li> </ul> <p>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.</p>
com.snc.par.nlq.report_designer.enabled	<ul style="list-style-type: none"> <li>• True: Enables NLQ in Report Designer</li> <li>• False: Removes NLQ in Report Designer</li> </ul>
glide.cmdb.query.nlq.activated	<ul style="list-style-type: none"> <li>• True: NLQ search feature is active in CMDB Query Builder</li> <li>• False: NLQ search feature is inactive in CMDB Query Builder</li> </ul>
glide.service_portal.ais_nlq_enabled	<ul style="list-style-type: none"> <li>• True: Enables NLQ in global search</li> <li>• False: NLQ is not available in global search</li> </ul>

## 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:

- 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.







# 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.

<p><b>Explore</b></p>  <p>Learn about Predictive Intelligence and machine learning</p>	<p><b>Install</b></p>  <p>Install Predictive Intelligence and its associated apps</p>	<p><b>Configure</b></p>  <p>Configure Predictive Intelligence on the platform</p>
<p><b>Train</b></p>  <p>Create and train machine learning solutions.</p>	<p><b>Use</b></p>  <p>Use Predictive Intelligence for machine learning.</p>	<p><b>Reference</b></p>  <p>Get details about Predictive Intelligence components such as properties and language support.</p>

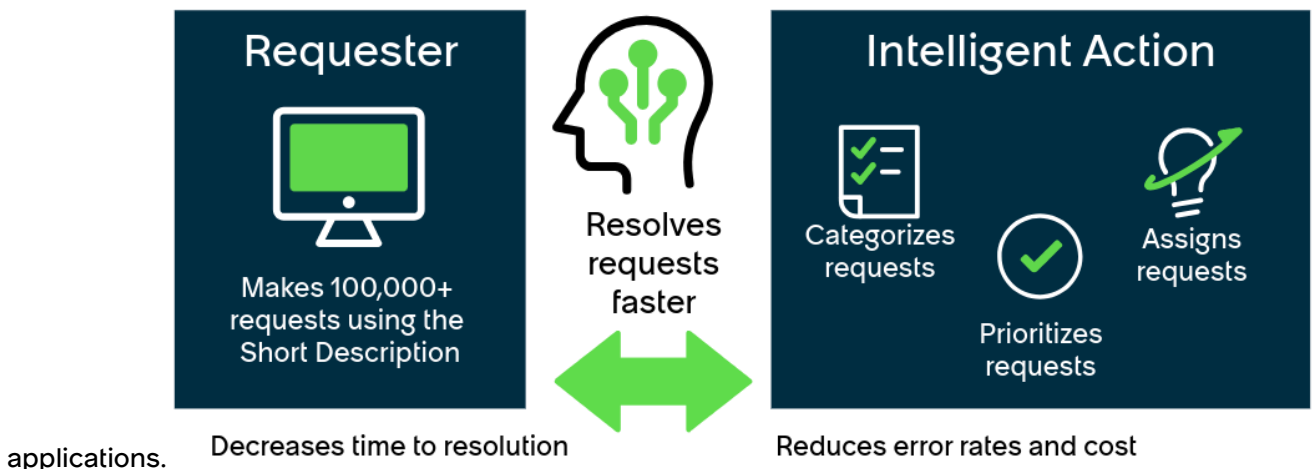
## Explore Predictive Intelligence

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 to use artificial intelligence and machine learning to improve the work experience. You can create and train models on the platform and integrate with other ServiceNow products and


## Predictive Intelligence



The following introduces the underlying concepts behind Predictive Intelligence and the different frameworks available.

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 combine the existing model with new training data.

### 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.

## 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 frameworks in the Xanadu release. Each framework has different solution types to train the system to predict, recommend, and organize data outcomes. A trained solution can be invoked by any application through an API to make a prediction. More information can be found in [Predictive Intelligence frameworks](#).

## Predictive Intelligence frameworks

Predictive Intelligence provides three different model frameworks in the Xanadu 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 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

**i 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](#).

## 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](#).

## 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.

- [Configuration tips for Predictive Intelligence](#)
- [Create a custom stopwords list](#)
- [Create a word corpus](#)
- [Quick start tests for Predictive Intelligence](#)
- [Activate a solution version](#)
- [Export trained solutions to production](#)
- [Configuring advanced settings for your ML solutions](#)

### 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
<p>The solution training remains in Waiting for Training status for too long, as the scheduler job is using an incorrect Glide callback instance URL.</p>	<p>Ensure the <code>glide.servlet.uri</code> property in the Glide instance is set to the correct instance URL. This issue can occur when:</p> <ul style="list-style-type: none"> <li>• An instance is cloned from production, yet it still refers to the production URL for the <code>glide.servlet.uri</code> property.</li> <li>• The Glide instance is provisioned and runs the training for the first time.</li> </ul>
<p>New categories have been added and aren't yet having an impact on training.</p>	<p>This is expected behavior, as the new categories may not yet have sufficient data until the solution is retrained.</p>
<p>The solution training fails.</p>	<p>When the training fails, click the <b>Show Training Progress</b> related link on the solution screen to determine where the potential problem resides.</p>
<p>The solution training fails due to user authentication.</p>	<p>Navigate to <b>System Security &gt; Users</b> and ensure the <code>sharedservice.worker</code> user is set to Active.</p>
<p>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."</p>	<p>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:</p> <ol style="list-style-type: none"> <li>1. Ensure the distribution of the output field isn't skewed.</li> <li>2. Retrain the model by changing the date filters to use a larger amount of data.</li> <li>3. If the input fields are not fully populated, add a filter to remove null records.</li> </ol>
<p>The solution has data in multiple languages but the coverage and precision results are poor.</p>	<p>Use the following options to help improve your metrics.</p> <p>Option 1: Update the processing language of the solution to the most prominent non-English language.</p> <p><b>Note:</b> English is applied by default for all datasets.</p>

Issue	Resolution or suggested action
	<p>Option 2: If there's sufficient data for each language/region:</p> <ol style="list-style-type: none"> <li>1. Add a filter criteria for a specific language/region where the primary language can be identified (Dutch, English, French, German, Japanese, or Spanish).</li> <li>2. Generate a solution for each language/region and apply the proper processing language to each solution.</li> </ol>

## Solution prediction

Issue	Resolution or suggested action
<p>The prediction fails and returns a Java exception where the cause is unknown.</p>	<ol style="list-style-type: none"> <li>1. Search for the exception in the Predictive Intelligence Glide logs.</li> <li>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.</li> </ol>
<p>There is no prediction applied to the incident/case record but the prediction returns a value when tested in the Rest API Explorer.</p>	<p>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.</p> <ol style="list-style-type: none"> <li>1. Navigate to <b>System Web Services &gt; REST &gt; REST API Explorer</b> to find the confidence level for the prediction. See <a href="#">Test a classification solution prediction</a>.</li> <li>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 <b>Class</b> page appears.</li> <li>3. Check the <b>Estimated Precision</b> and <b>Estimated Coverage</b> 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.</li> <li>4. Adjust the class precision and coverage values to increase coverage or precision. See <a href="#">Tune a trained classification solution</a>.</li> </ol>

## 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 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.

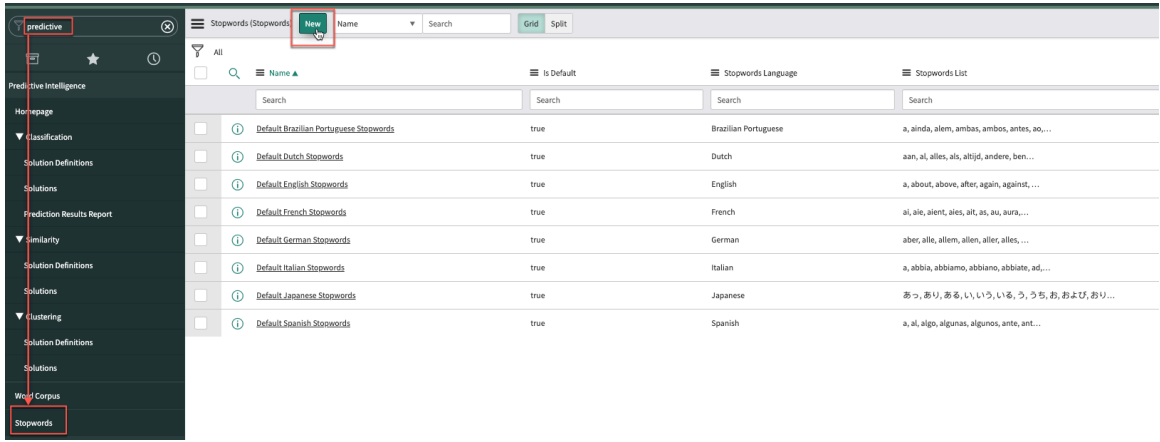
Predictive Intelligence provides you with default stopwords lists for each language the system supports. Examples of stopwords include words such as in, the, and the names of people and companies. You can also define your own stopwords list that's comprised of words specific to your organization and 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 Portuguese language.


#### Procedure

1. Navigate to **All > Predictive Intelligence > Stopwords**.
2. On the Stopwords list, click **New**.



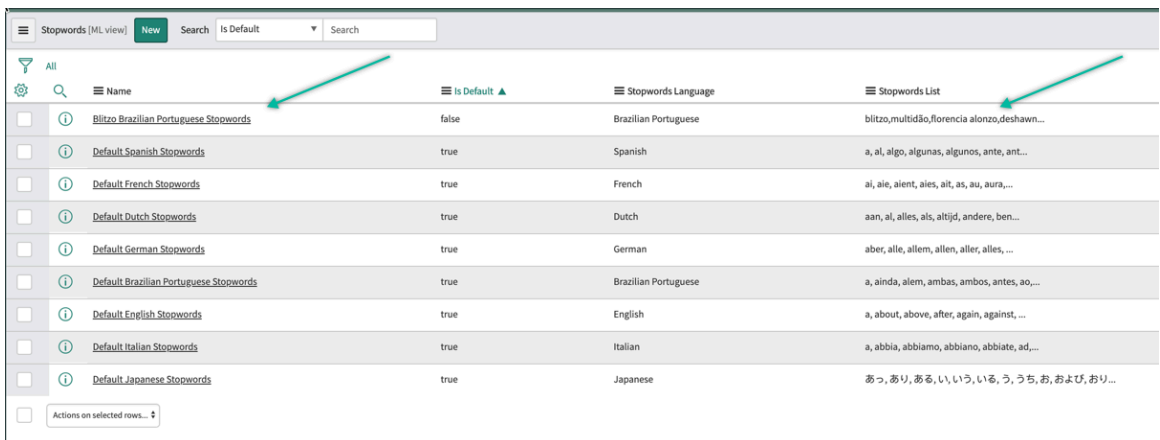
3. In the Stopwords form, configure these fields.

Field	Value
Name	Enter a unique name for the list, such as the name of your company and the processing language. For example, <b>Blitzo Brazilian Portuguese Stopwords</b> .
Stopwords Language	Select <b>Brazilian Portuguese</b>
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**.

Name	Is Default	Stopwords Language	Stopwords List
Blitzo Brazilian Portuguese Stopwords	false	Brazilian Portuguese	blitzo,multidido,florencia alonzo,deshawn...
Default Spanish Stopwords	true	Spanish	a, al, algo, algunas, algunos, ante, ant...
Default French Stopwords	true	French	ai, aie, aient, aies, ait, as, au, aura...
Default Dutch Stopwords	true	Dutch	aan, al, alles, als, altijd, andere, ben...
Default German Stopwords	true	German	aber, alle, allem, allen, aller, alles, ...
Default Brazilian Portuguese Stopwords	true	Brazilian Portuguese	a, ainda, alem, ambas, ambos, antes, ao...
Default English Stopwords	true	English	a, about, above, after, again, against, ...
Default Italian Stopwords	true	Italian	a, abbia, abbiamo, abbiamo, abbiate, ad...
Default Japanese Stopwords	true	Japanese	あっ, あり, ある, い, いう, いる, う, うち, お, および, おり...

### 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 Xanadu 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 <b>Active Incidents</b> and <b>Published KBs</b> , 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.

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 <b>Incidents closed in last 6 months</b> .
Table	Select the table that contains the data you want to include in your word corpus. For this use case, select <b>Incident [incident]</b> .   <b>Note:</b> 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: <b>[Closed] [is not empty] and [Created in last 6 Months]</b> .
Field List	For this use case, select <b>Short description, Description, and Resolution notes</b> .
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.

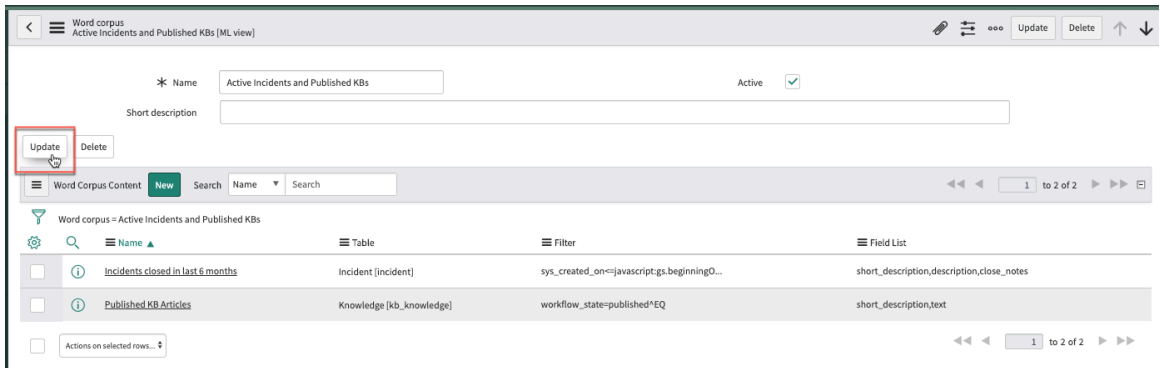
8. Select **Submit**.
9. In the Word Corpus Details section, select **New**.

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 <b>Published KB Articles</b> .
Table	Select the table that contains the data you want to compare to your first content component. For this use case, select <b>Knowledge [kb_knowledge]</b> .  <b>Note:</b> 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: <b>[Workflow] [is] [Published]</b> .
Field List	Select <b>Short description</b> and <b>Article body</b> .

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 apply an upgrade or develop 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.

**Predictive Intelligence: Classification and Similarity Solution Prediction test suite**

Test	Description	Release version
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.	New York
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.	New York
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.	New York
PI: Glide upgrade test for Similarity solution	Validate that the similarity model prediction API calls on active models are successful after a glide upgrade.	New York

**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.

## 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.

**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 a data source](#)

**Advanced Solution Settings for Classification**

From the Xanadu Patch 9 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% accuracy for all records you train in the Email class.

#### 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.


- 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 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.

#### 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.



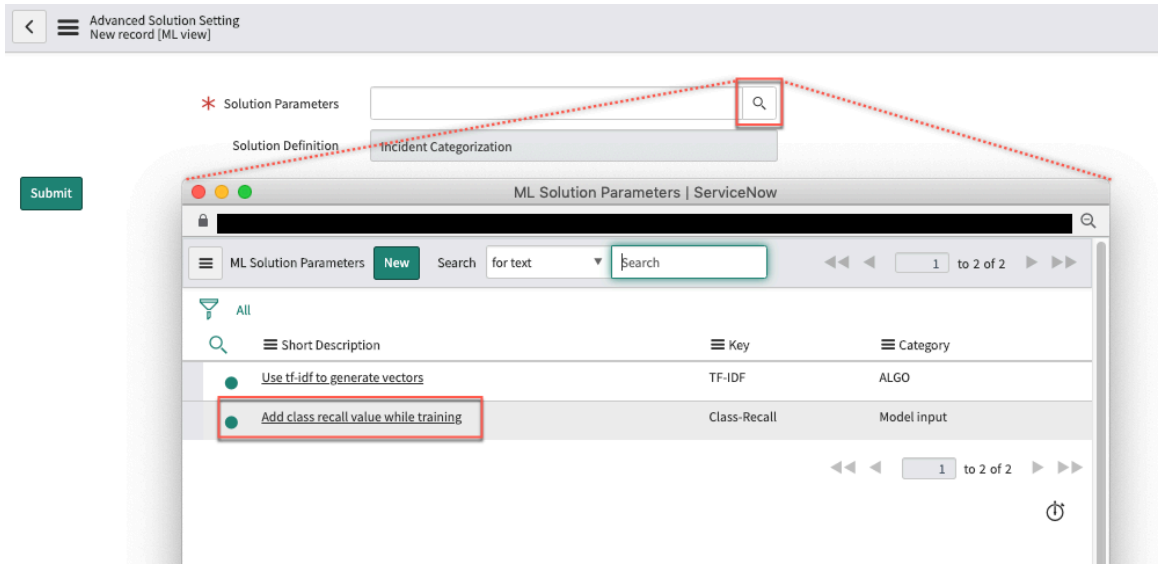
3. On the Advanced Solution Settings tab in the Related Links section of the form, click **New**.



4. Create a parameter record.

a. In the **Solution Parameters** field, click the search icon.

b. In the ML Solution Parameters screen, select **Add class recall value while training**.



5. Click **Submit**.

The Advanced Solution Setting record appears.




6. Configure the **User inputs** field.

a. Enter the class name on which you want to track the Recall value.  
In this scenario, you enter **Phish** for the **ClassName**.

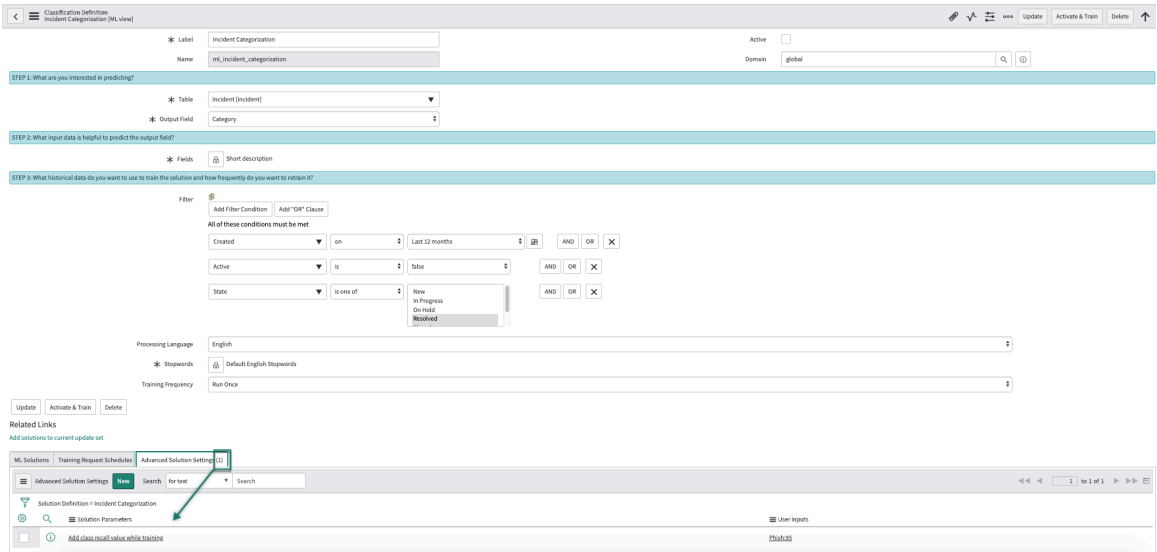
b. Enter the Recall value.  
Enter **90** for the **RecallValue**.

Here you're specifying Phish as the target class, and 95 is the recall percentage you're requesting the system to deliver during solution training.



7. Click **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.

## Related topics

- [Create and train a classification solution](#)
- [Configuring target metrics for a trained classification solution](#)
- [Configure TF-IDF for solutions](#)

## 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.

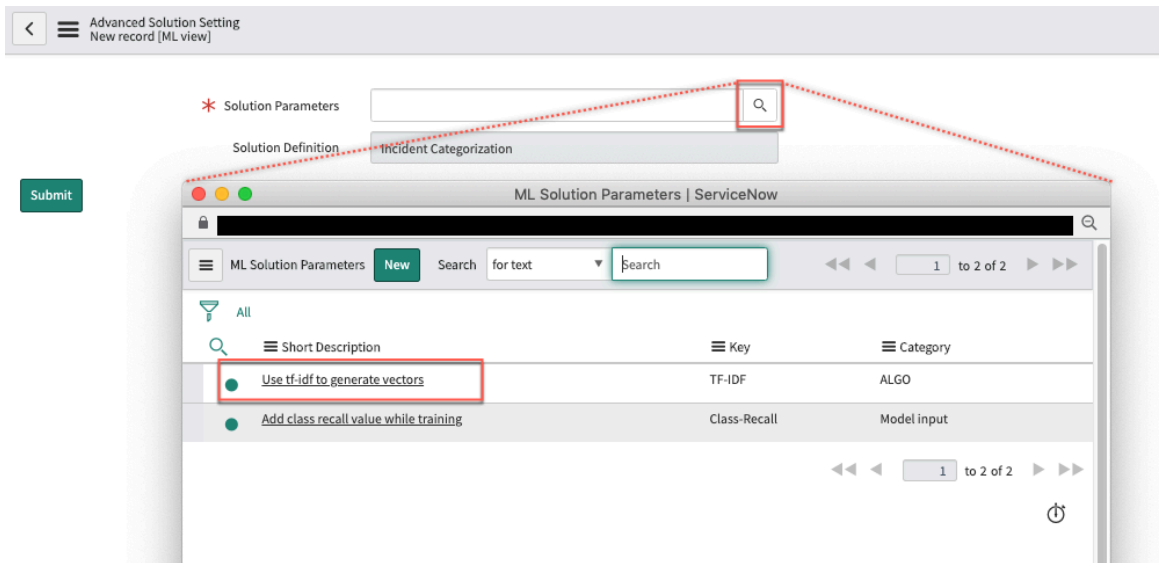


3. On the Advanced Solution Settings tab in the Related Links section of the form, click **New**.




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**.



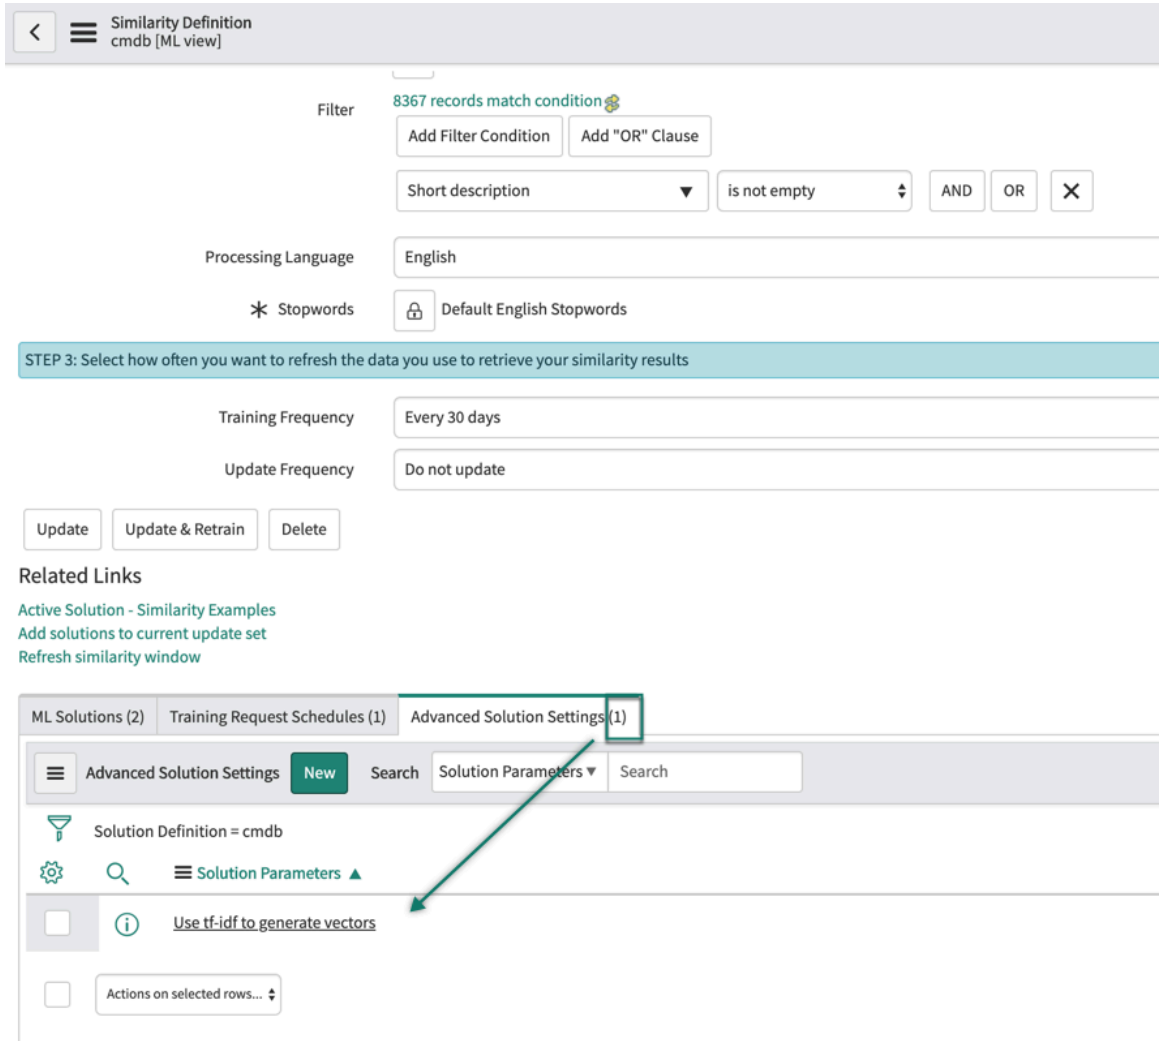
5. Click **Submit**.

The Advanced Solution Setting record screen refreshes.



**6. Click Submit.**

**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

**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.

- 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


## 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.

In this example scenario, you apply XGBoost to both a classification solution and a regression solution.

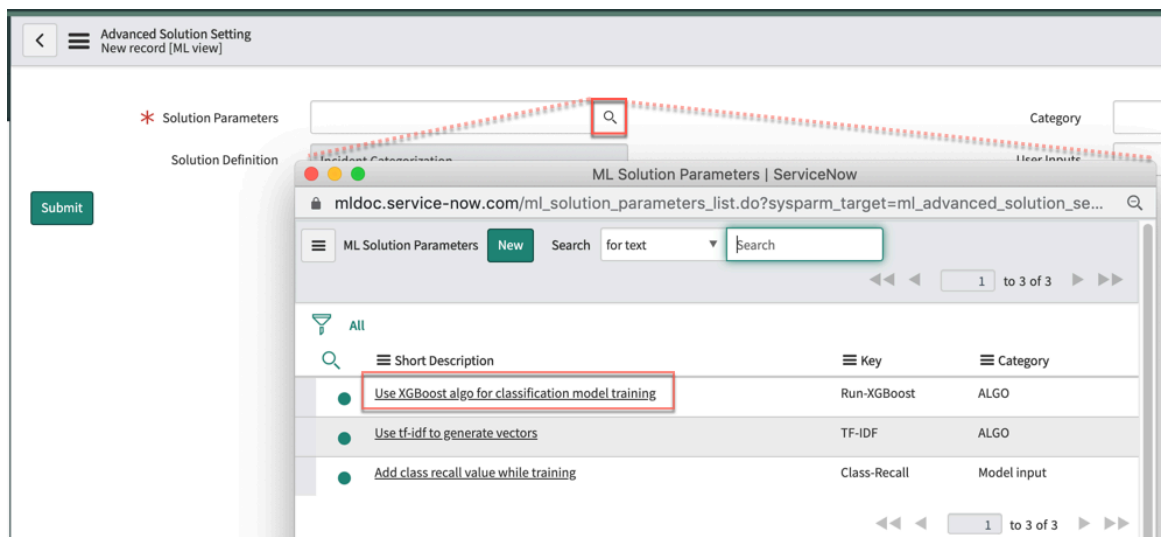
## Procedure

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, 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 XGBoost algo for classification model training**.



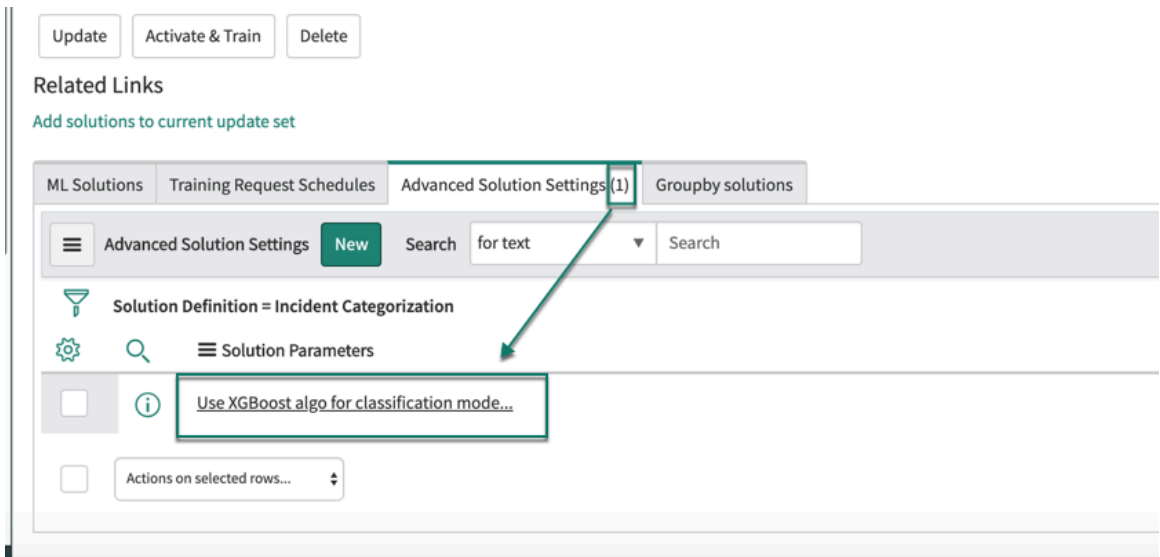
**5. Click Submit.**

The Advanced Solution Setting record screen refreshes.



**6. Click 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.**

**8.** In this second scenario, you open a regression solution definition form.

**9.** Repeat steps 1-5 from the previous classification solution example, except this time you're using a regression solution.

**10. Click 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.

## 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.

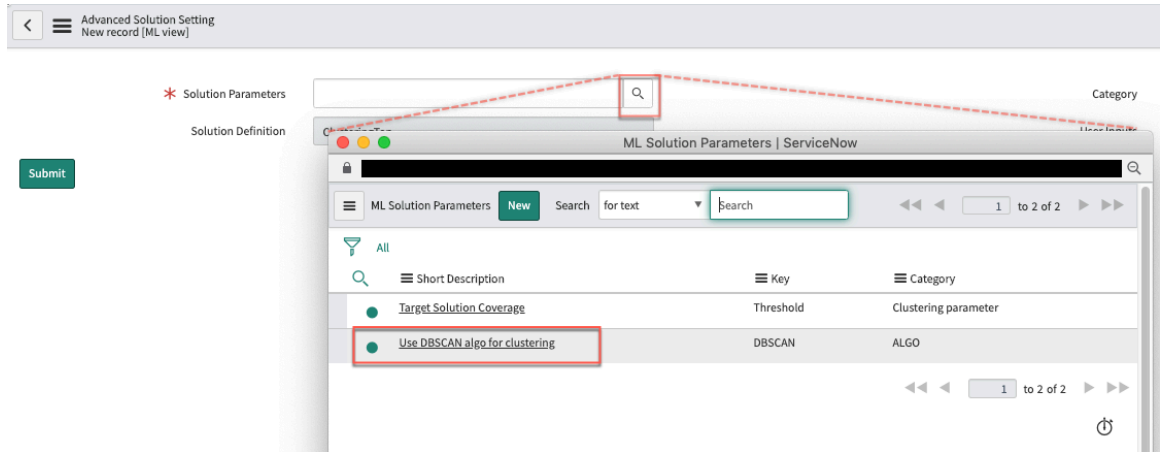
In this example scenario, you use an Incidents clustering solution definition form.

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.


a. In the **Solution Parameters** field, select the search icon.

b. In the ML Solution Parameters screen, select **Use DBSCAN algo for clustering**.



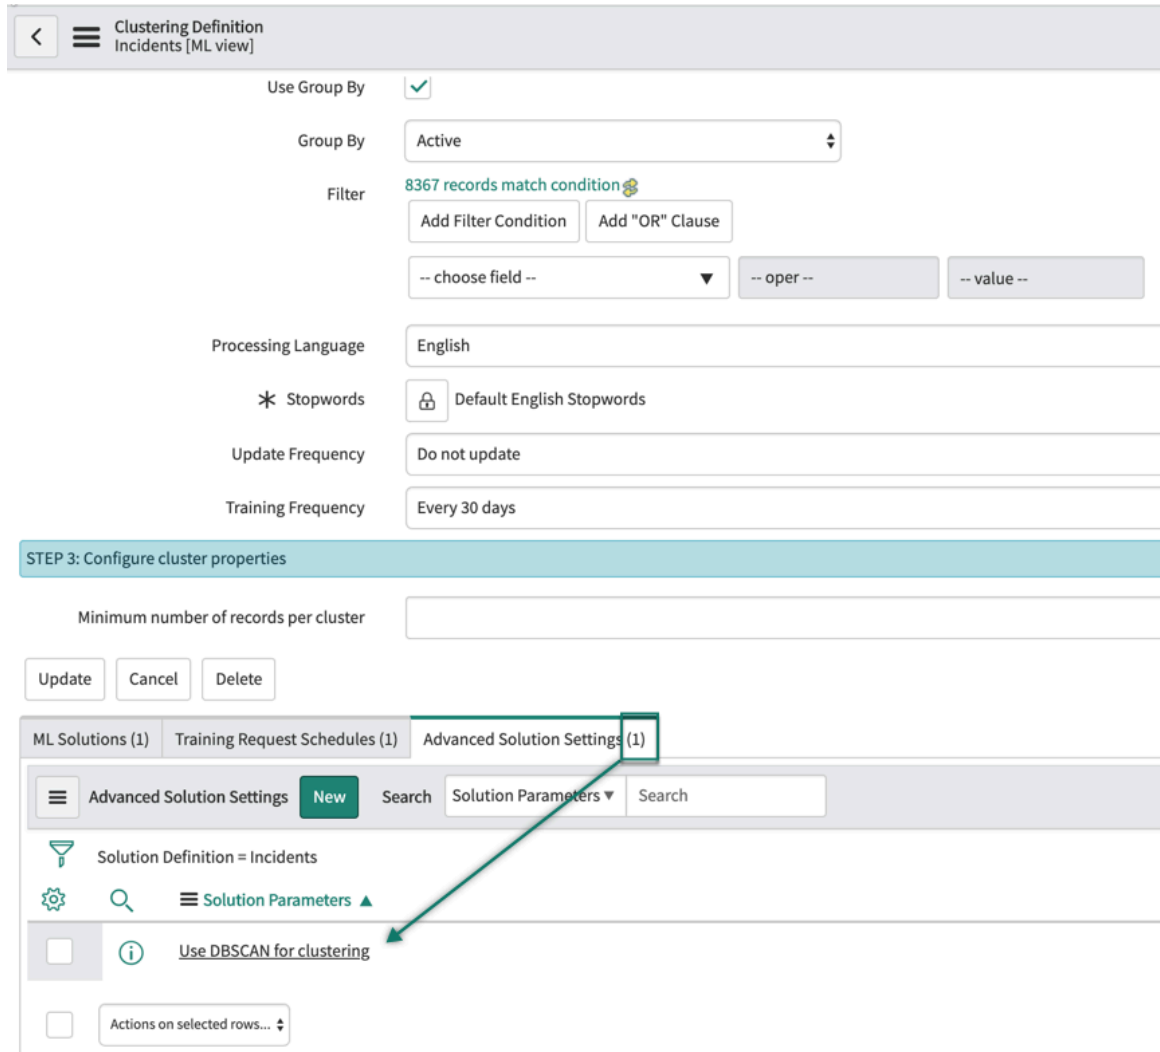
5. Select **Submit**.

The Advanced Solution Setting record appears. The field **User Inputs** is grayed out because it does not apply to this algorithm.



6. Select **Submit**.

**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

### About this task

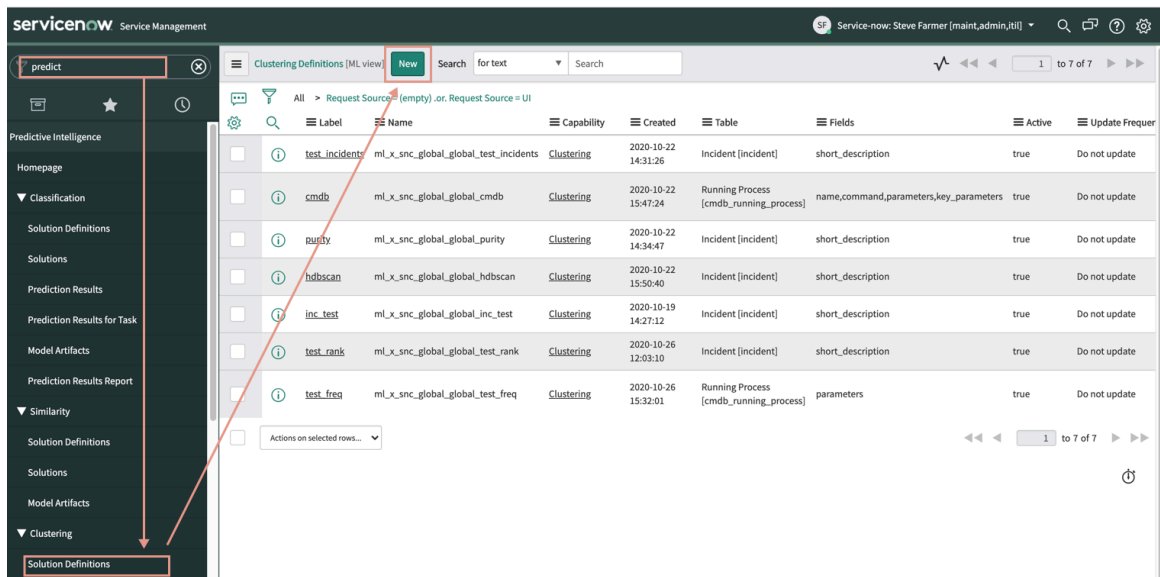
You can apply the HDBSCAN algorithm to help the system identify data samples that aren't assigned to any cluster. For example, you can apply HDBSCAN to support Topic Discovery.

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

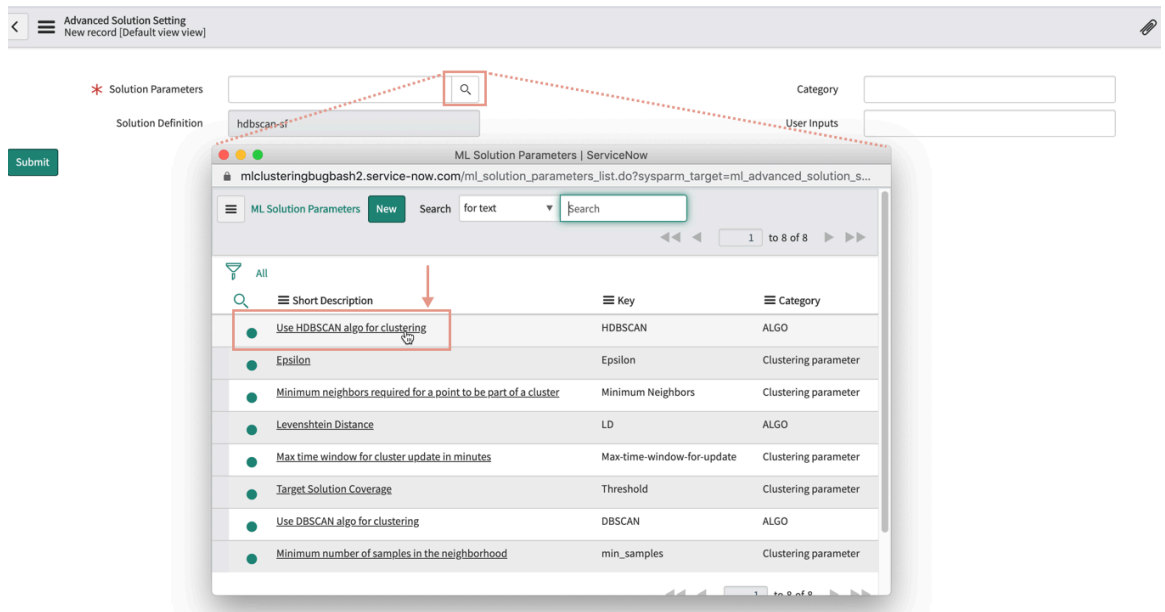
4. Select **Submit & Train**.

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.


a. In the **Solution Parameters** field, click the search icon.

b. In the ML Solution Parameters screen, select **Use HDBSCAN algo for clustering**.



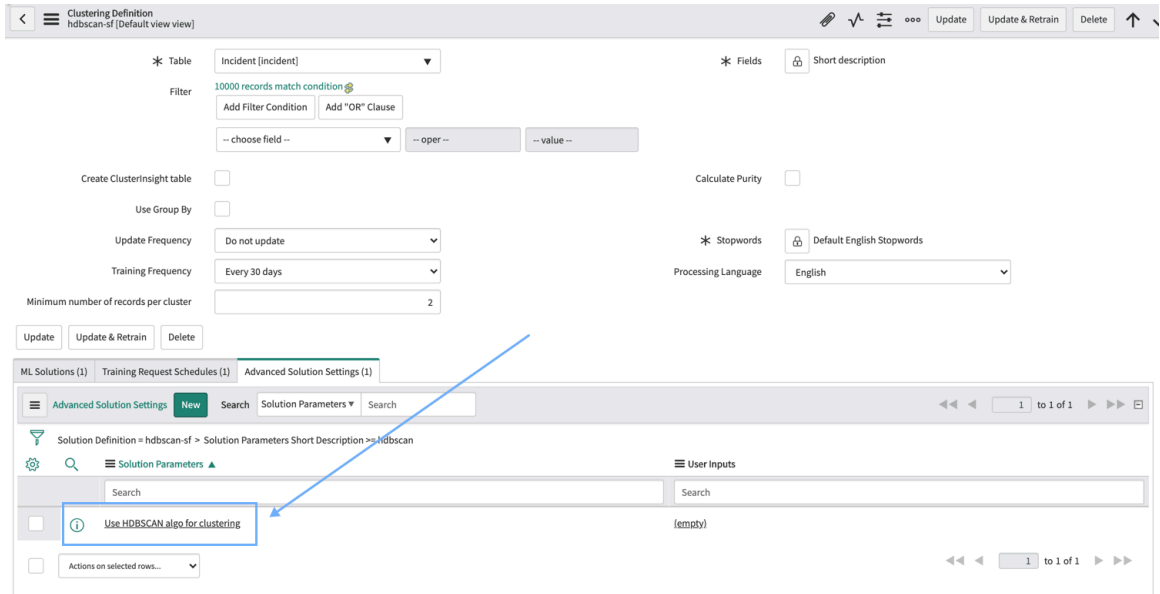
7. Select **Submit**.

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**.

**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.

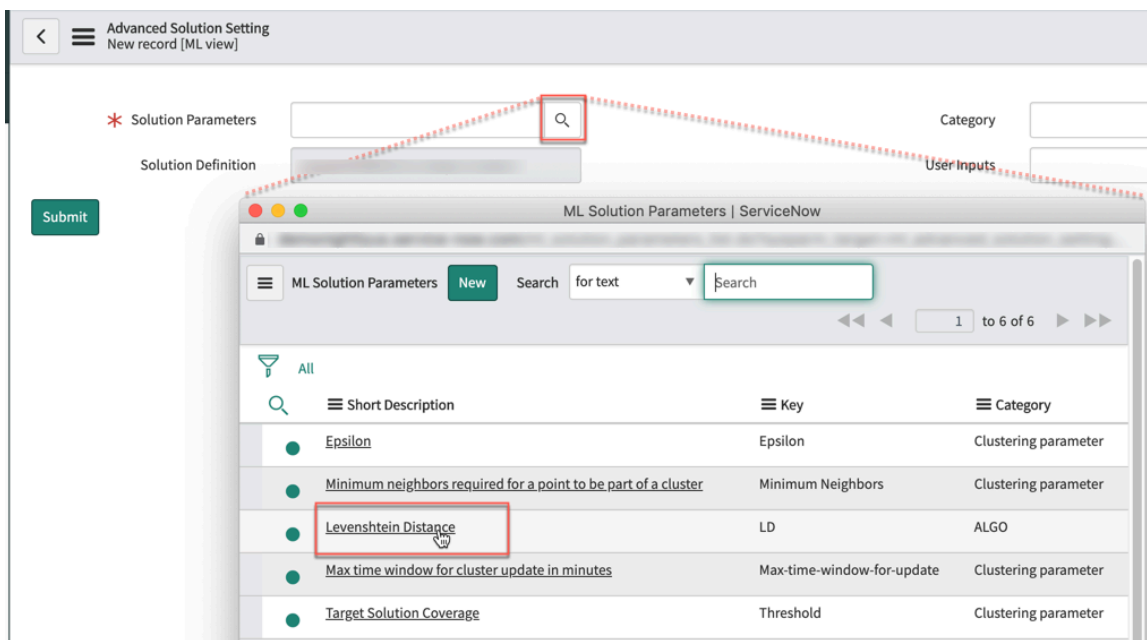
In this example scenario, you train your solution definition by using the third option referenced above.

**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**.



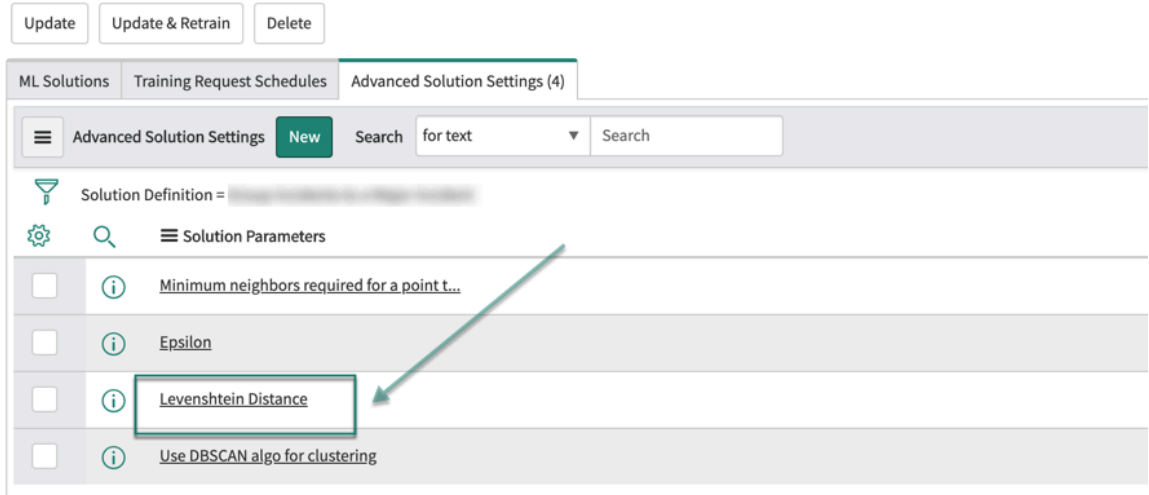
5. Select **Submit**.

The Advanced Solution Setting record screen refreshes.

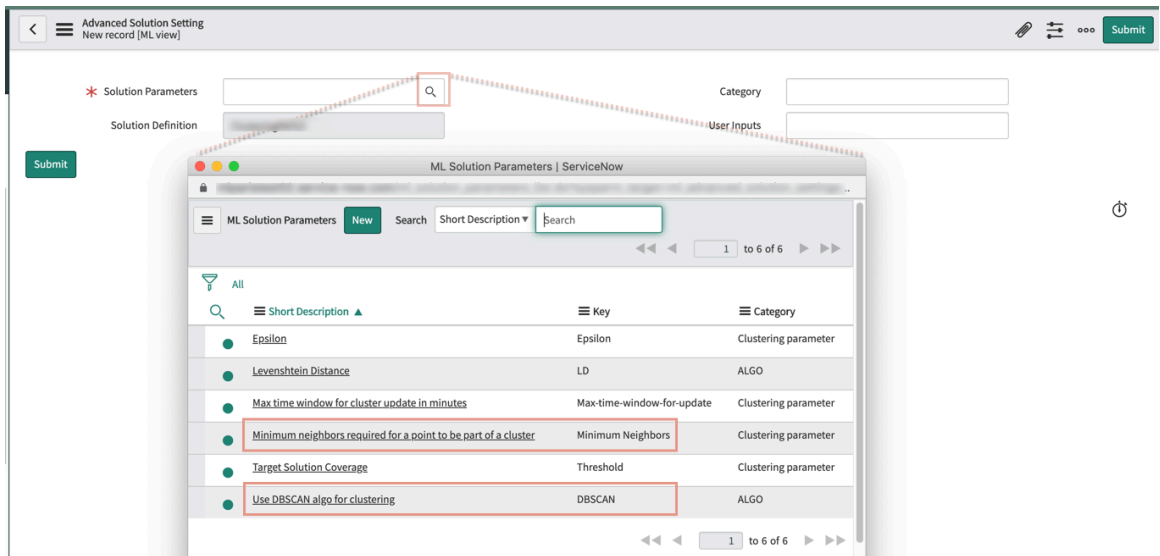


6. Select **Submit**.


**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.

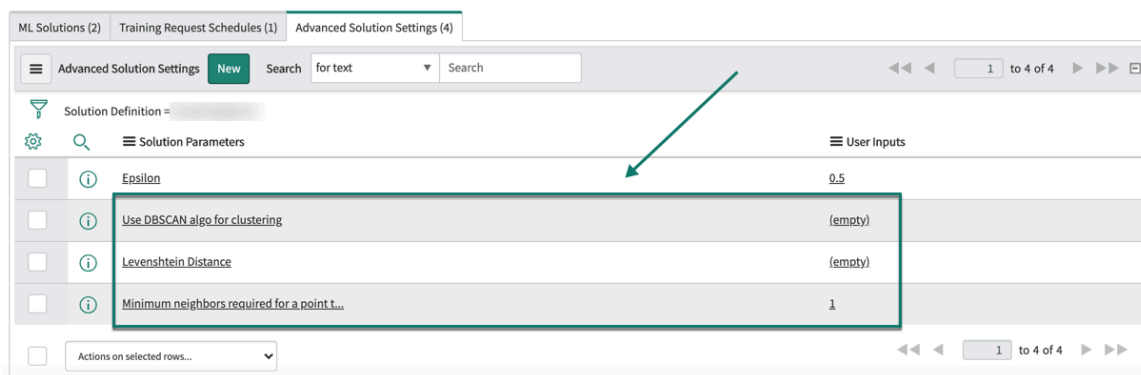


When you select, configure, and submit the **Minimum Neighbors** solution parameter, be sure to set the **User Inputs** field with a value of 1. Only some parameters have a User Inputs field.



**Result:**

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.

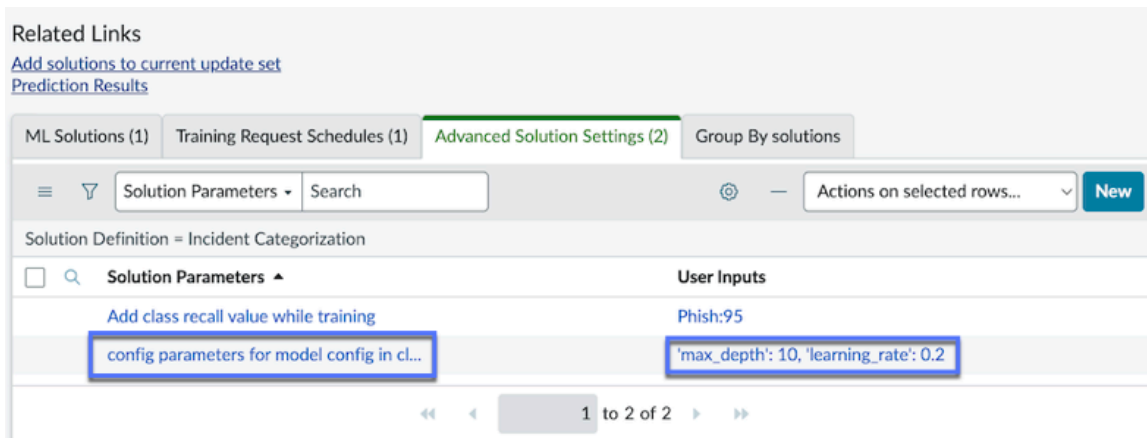
**Note:** When pasting the parameters into the solution definition form, don't include JSON's enclosing label, braces, or parentheses. Enter the key-value pairs in the format shown here, after editing the values according to your requirements.

```
'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,
```

```
'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:

```
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.

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);
```

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:

```
{
  "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.

The screenshot shows the 'Advanced Solution Setting' form for a new record in the ML view. The 'Solution Parameters' field is active, and the search term 'Remove others label' is entered. The 'Category' is set to 'preprocessing parameter' and 'User Inputs' is set to 'False'. A 'Submit' button is visible in the top right corner.

## 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](#).

- Clustering solutions:

Groups similar records into clusters to identify patterns and major incidents. See [Create and train a clustering solution](#).

- Regression solutions:

**i Important:** With the Xanadu 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 create new ones.

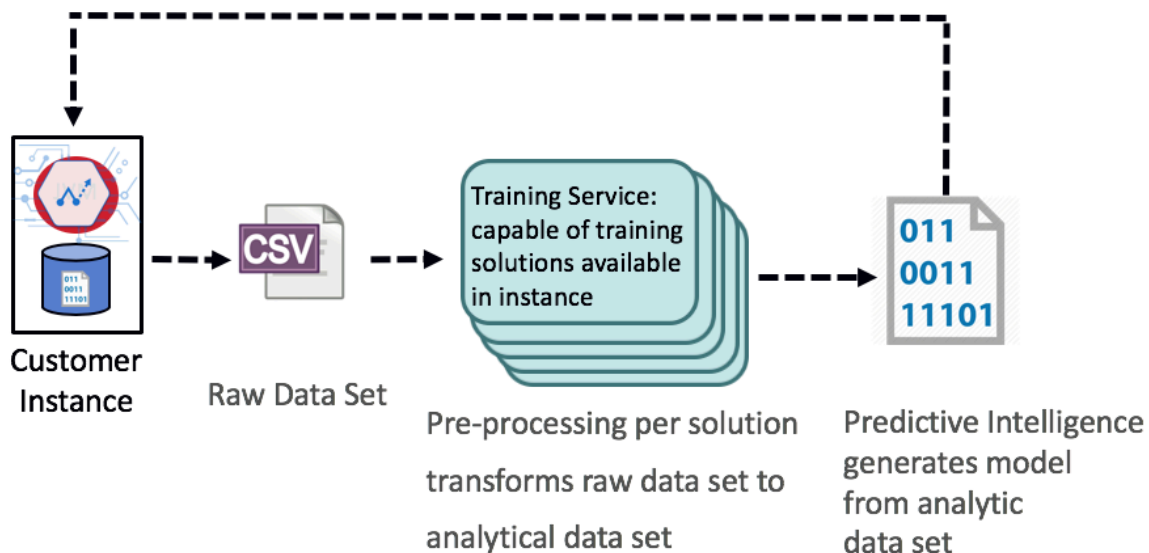
Uses historic data to predict numeric outputs, such as estimating the time it takes to resolve an incident or case. See [Create and train a regression solution](#).

## 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. As 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.

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.

## 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

**Note:** In the Xanadu 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.

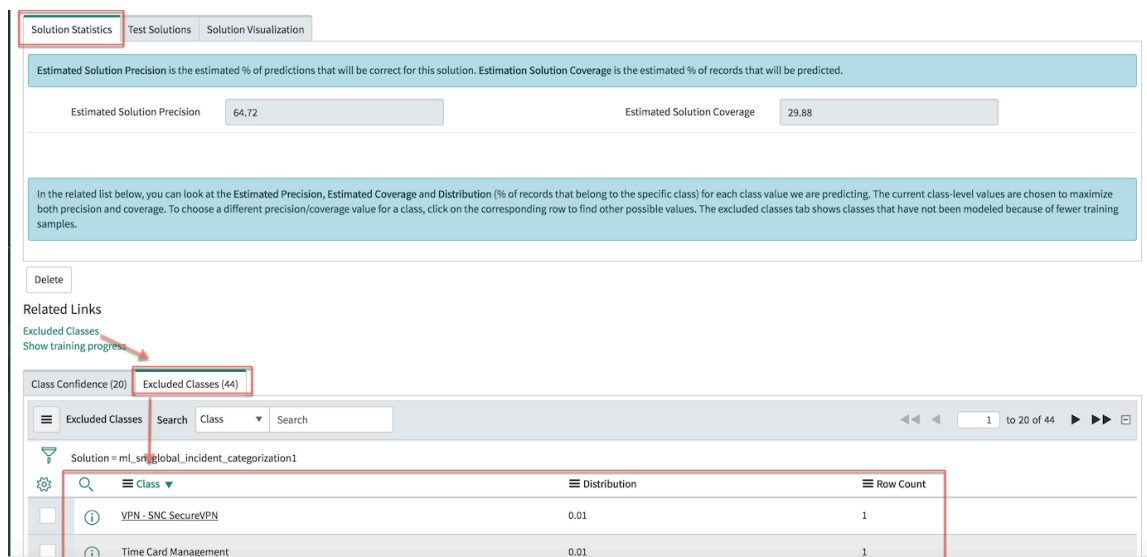
### About this task

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.

**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 <a href="#">Create a word corpus</a> .  <b>Note:</b> 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.

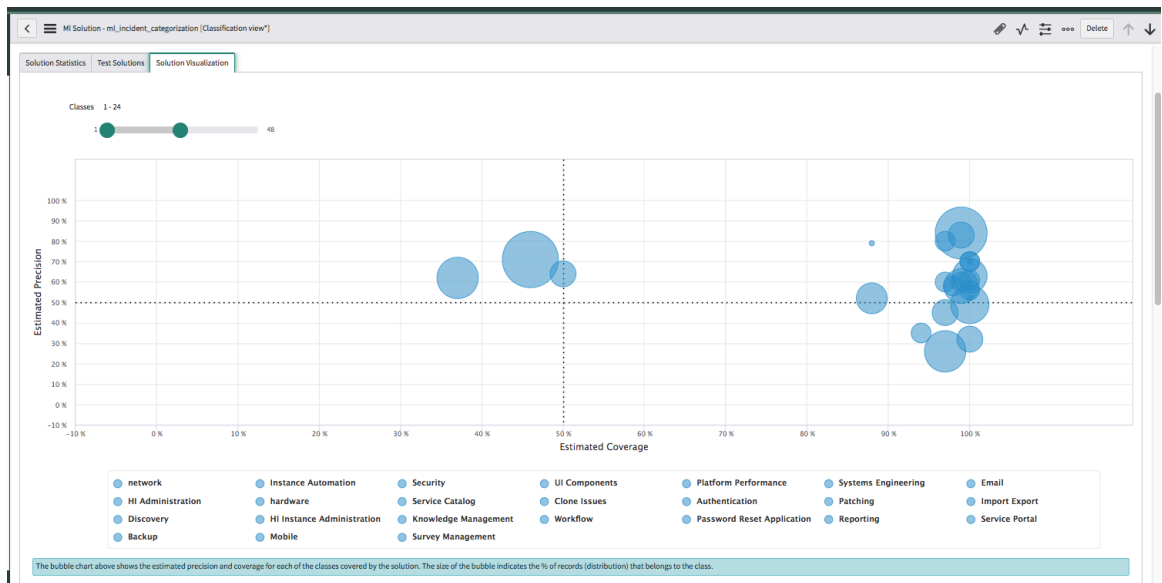
Field	Value
	<p>In general, a good output field has these characteristics.</p> <ul style="list-style-type: none"> <li>○ It's a choice field or a string field with a finite set of possible values.</li> <li>○ It has some causal connection to the input fields.</li> </ul> <p>For example, on the default <b>Incident Categorization</b> solution definition, the output field is set to <b>Category</b>.</p>
Fields	<p>Select the input fields that you want the solution to use to generate a prediction.</p> <p>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 <b>Short description</b> field, so it's a great input field to use in general. You could also use <b>Resolution notes</b> as an input field, as it too might reference the incident class in the detailed notes for the incident.</p> <p>In general, good input fields have these characteristics.</p> <ul style="list-style-type: none"> <li>○ The fields are available to users when creating records.</li> <li>○ 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.</li> <li>○ The field has a default value and should not be blank.</li> </ul> <p>All default solution definitions use the <b>Short description</b> field.</p>
Filter	<p>Click <b>Add Filter Condition</b> to apply conditions to the records you're training.</p> <p>For example, the <b>Incident Categorization</b> solution definition uses a filter with these conditions: <b>[Created][on][Last 12 months] AND [Active][is][false] AND [State][is one of][Resolved   Closed]</b></p> <p>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.</p> <p><b>Note:</b> 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.</p> <p>In general, a good filter has these characteristics.</p> <ul style="list-style-type: none"> <li>○ The training records are inactive and their states indicate work completed within your standard process, such as resolved or closed.</li> <li>○ 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.</li> </ul>

Field	Value
	<ul style="list-style-type: none"> <li>○ The training records contain multiple examples of each target field value you want the solution to predict.</li> <li>○ The training records include common variations of the input fields.</li> </ul> <p>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.</p>
Processing Language	<p>Select the dominant language of the dataset that you're training on the solution definition. If the dataset language is Italian, choose <b>Italian</b>. 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.</p> <p><b>Note:</b> 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.</p>
Stopwords	<p>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 <b>Default Italian Stopwords</b> list appears. The <b>Default English Stopwords</b> list is also included. If you create a custom stopwords list, you can select it from the Stopwords field to add to your solution.</p>
Training Frequency	<p>Select how often the system regenerates the solution. The available options range from <b>Run once</b> up to <b>Every 180 days</b>.</p> <p><b>Note:</b> The minimum number of records required for classification solution training is set at 10,000.</p> <p>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.</p> <p>When your solution definition is fairly stable, consider scheduled trainings, as data can age over time, thus degrading the accuracy of your prediction model.</p> <p><b>Note:</b> 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.</p>

4. Click the appropriate context menu option or button for your solution definition.

Option	Description
<b>Save or Save &amp; Train</b>	Save your solution definition record so you can return to it later, or save and submit it for training.
<b>Submit or Submit &amp; Train</b>	Create your solution definition record and submit it, or submit and train it.

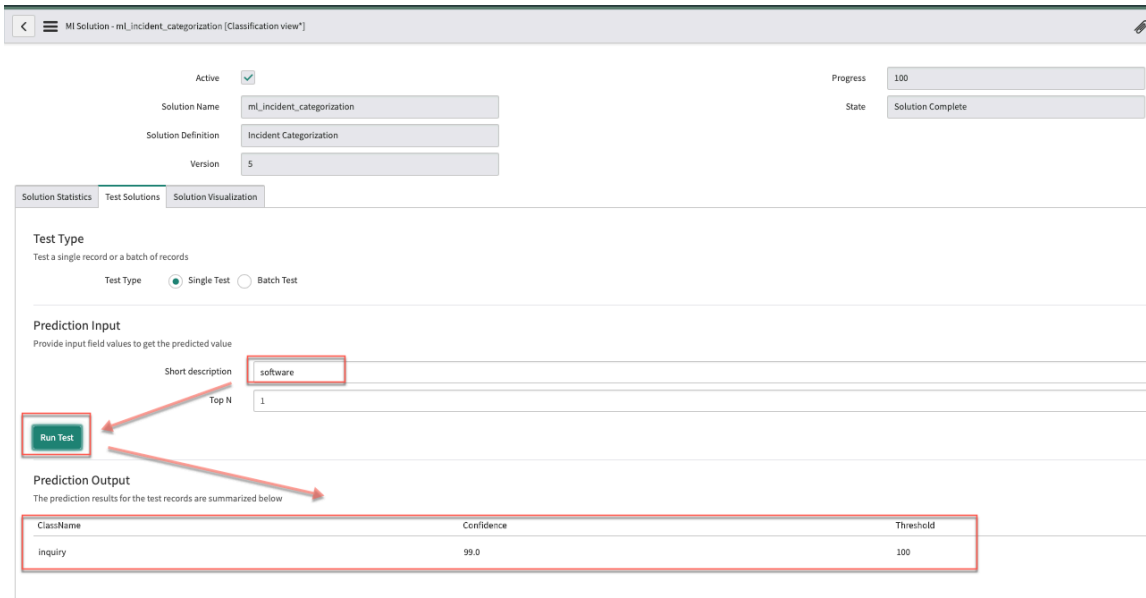
5. If you submitted the solution for training, click **OK** on the **Training Activation** window to confirm.
  - The system schedules the solution for training with the nearest training service. The system sends you a notification when the training completes, including any errors that may have occurred 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 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 statistics don't meet your threshold of usefulness. For example, exclude it if you don't want the model to predict 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

### About this task

If your classification solution doesn't provide sufficient precision or coverage, you can exclude a particular incident category from prediction.

Excluding a class from prediction only lasts until the next time you train your solution. If a class still doesn't provide sufficient precision or coverage values, you may want to consider deactivating the solution until it provides better results.

Typically, you exclude a class from prediction if you only want a person to manually set the excluded class value. For example, exclude the class if 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 only select 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 this class will never produce 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.

- 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.

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.

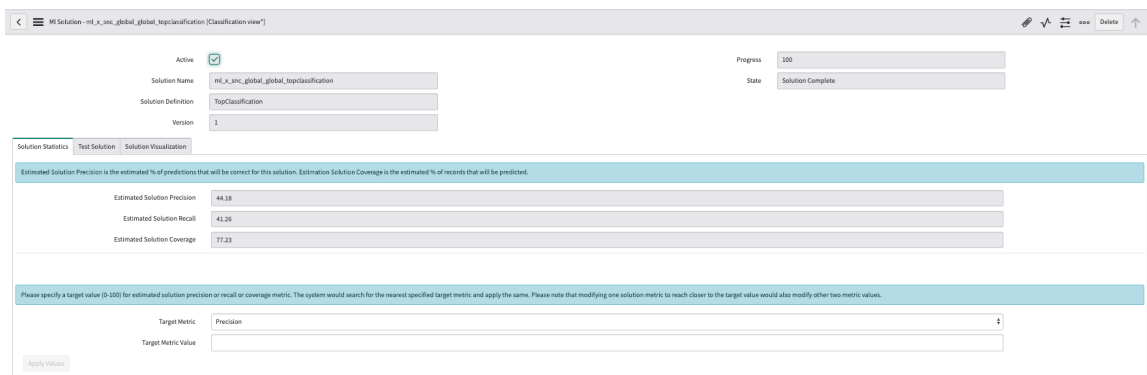
Applying a value to a metric asks the system to train its predictions to favor the metric you set based on the highest percentage value, and at a cost to the other metrics. The system tries to meet these values but may not set them exactly as you request due to how the data you're training is distributed.

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 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**.

This screenshot shows the 'Solution Visualization' interface. At the top, there are three tabs: 'Solution Statistics', 'Test Solution', and 'Solution Visualization'. Below the tabs, a light blue banner reads: 'Estimated Solution Precision is the estimated % of predictions that will be correct for this solution. Estimation Solution Coverage is the estimated % of records that will be predicted.'

The main area displays three horizontal bars representing metrics:

- Estimated Solution Precision: 100
- Estimated Solution Recall: 0
- Estimated Solution Coverage: 0

Below these bars, a text box instructs: 'Please specify a target value (0-100) for estimated solution precision or recall or coverage metric. The system would search for the nearest specified target metric and apply the same. Please note that modifying one solution metric to reach closer to the target value would also modify other two metric values.'

The configuration fields are:

- Target Metric: Precision (selected in a dropdown menu)
- Target Metric Value: 80

An 'Apply Values' button is located at the bottom left of the configuration area.

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.

This screenshot shows the 'Solution Visualization' interface after the target has been applied. The banner text remains the same. The metric values are now:

- Estimated Solution Precision: 80.04
- Estimated Solution Recall: 24.91
- Estimated Solution Coverage: 30.52

The configuration fields below are identical to the previous screenshot:

- Target Metric: Precision
- Target Metric Value: 80

The 'Apply Values' button is still present.

### 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**.

This screenshot shows the 'Solution Visualization' interface with the 'Coverage' target selected. The banner text is the same. The metric values are:

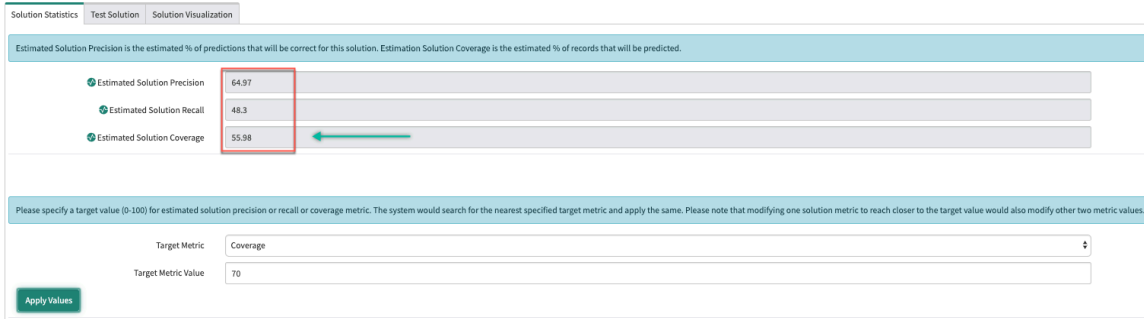
- Estimated Solution Precision: 80.18
- Estimated Solution Recall: 32.49
- Estimated Solution Coverage: 35.99

The configuration fields are:

- Target Metric: Coverage (selected in a dropdown menu)
- Target Metric Value: 70

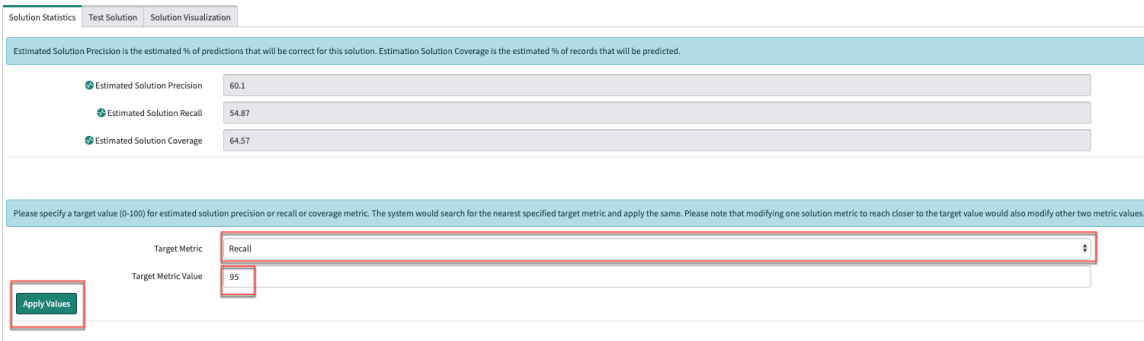
The 'Apply Values' button is at the bottom left.

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.

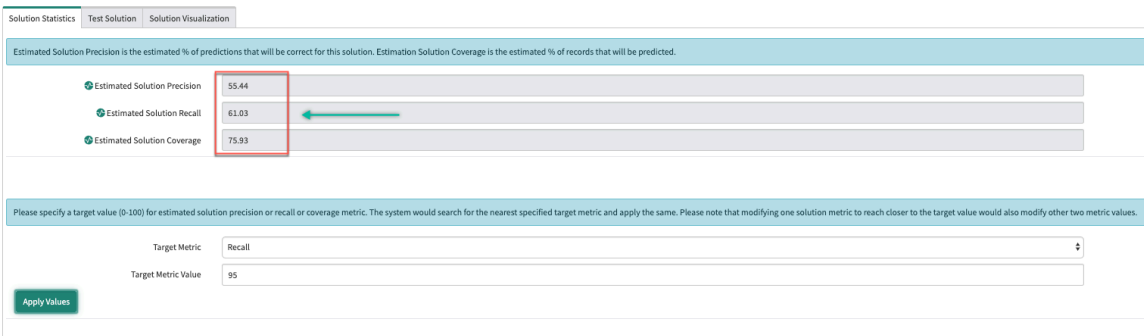


### Recall configuration example

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**.



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.

Name	Estimated Precision	Estimated Coverage	Estimated Recall	Distribution
Content Management System CMS	100	0	0	0.18
Platform Outage	100	0	0	0.19
HR Service Management	100	0	0	0.16
Internet ISP	100	0	0	0.17
Aspenzala	100	0	0	0.24
Edge Encryption	100	0	0	0.18
Domain Support	100	0	0	0.29
Security	100	0	0	0.99
Database Indexes	100	0	0	0.12
Scheduled Job Processing	100	0	0	0.13
Platform Licensing	100	0	0	0.12
Automation	100	0	0	0.23
UI Policy Client Script	100	0	0	0.17
Instance Automation	91	93	81	1.66
Event Management	83	100	45	0.17
Close Issues	83	93	86	2.61
Language and Translations	83	100	36	0.21

### 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.

### 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
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](#) 

## 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

**i Important:** In the Xanadu 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.

< ☰ Similarity Definition
New record View: ML
📎 ⚙️ ⋮
Submit
Submit & Train

\* Label

Name

STEP 1: Please select the table & field(s) you want to use to retrieve your similarity results, and the table and fields you want compared

\* Table

\* Test Table

\* Fields

\* Test Fields

Filter Add Filter Condition Add "OR" Clause

Processing Language

Stopwords  Default English Stopwords

STEP 2: Select how often you want to refresh the data you use to retrieve your similarity results

Training Frequency

Update Frequency


Submit
Submit & Train

### 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**.

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 <code>Match Knowledge Articles to Incidents</code> .
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	<p>If you have a legacy similarity solution, you can select a relevant word corpus from the <b>Word Corpus</b> field in the definition form.</p> <p><b>Note:</b> Starting from the Washington DC release, a word corpus is not required because a pre-trained model is used instead. The <b>Word Corpus</b> field is not visible in the definition form for pre-trained models.</p> <p>For more information, see <a href="#">Create a word corpus</a>.</p>
Table	<p>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 <b>Knowledge [kb_knowledge]</b> table, because its KB Article records might provide information relevant to the Incidents that you're trying to resolve.</p> <p>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.</p>
Test Table	<p>In the Test Table field, select the table containing the records that you want to target. In this use case, you select the <b>Incident [incident]</b> table, as it contains the Incident records that you're trying to resolve.</p> <p><b>Note:</b> 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.</p>
Fields	<p>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 <b>Short description</b> and <b>Article body</b>. Including the article body increases your chances of capturing informative details regarding the subject.</p> <p><b>Note:</b> <b>Journal Type</b> is not a supported data type.</p>

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 <b>Short description</b> of the Incident records you're trying to resolve.
Filter	Select <b>Add Filter Condition</b> to apply conditions to the Fields records you're using as an information source. For example, in this use case you could set a <b>workflow_state=published</b> condition to retrieve published KB articles only.  <b>Note:</b> 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.  <b>Note:</b> 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 <b>Default Italian Stopwords</b> list appears. The <b>Default English Stopwords</b> list is also included.  To use a custom stopwords list, select the lock icon(  ) and then search in the <b>Select target record</b> field.
Training Frequency	Select a retraining frequency. The available options range from <b>Run Once</b> up to <b>Every 180 days</b> .
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 <b>Every 15 minutes</b> , 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 <b>Every 1 day</b> .

Field	Value
	<p><b>Note:</b> 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.</p>

4. Select the appropriate button for the solution definition.

Option	Description
<b>Save</b>	Save your solution definition record so you can return to it later.
<b>Submit &amp; Train</b>	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 updates your Solution Definition form, where it delivers paired solution examples ranked by their degree of similarity.
- When training completes, the system uploads the solution as an Attachment record.

### What to do next

Review the trained similarity solution examples on the Related Links section of your Solution Definition 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 the Similarity Score Threshold values for your similarity examples.
- Role required: admin or ml\_admin

### Procedure

1. Navigate to **All > Predictive Intelligence > Similarity > Solutions**.
2. In the ML Solutions list, locate your solution and click the Reference Lookup icon (i).
3. Click **Open Record**.
4. In the Solution Statistics section, enter a new numerical value that represents a percentage in the **Similarity Score Threshold** field.

For example, imagine that the current score 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 score of 90.

**5.** In the Context Menu, click **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 score to 90, the degree of similarity in your word corpus is accurate at least 91% of the time.

## 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

**i Important:** In the Xanadu 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 <code>Group Incidents to a Major Incident</code> .
Name	As you enter your solution Label, this field automatically populates with a system-assigned name based on your Label value.
Word Corpus	<p>If you have a legacy clustering solution, you can select a relevant word corpus from the <b>Word Corpus</b> field in the definition form.</p> <p><b>i Note:</b> With the Xanadu release, a word corpus is not required, because a pre-trained model is used instead. The <b>Word Corpus</b> field is not visible in the definition form for pre-trained models.</p> <p>For more information, see <a href="#">Create a word corpus</a>.</p>

Field	Value
Table	<p>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 <b>Incident [incident]</b> table as it contains incident records you want to group together for a major incident analysis.</p> <p>When you assign a table value, a link appears in the form that shows the number of records that match your current conditions.</p>
Fields	<p>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 <b>Short description</b>.</p> <p><b>i Note:</b> When selecting a reference type field, you must dot-walk to the field's property name. For example, instead of <code>short_description</code>, enter <code>short_description.name</code>.</p>
Use Group By	<p>Select this check box only if you want to group input records by a field before creating clusters.</p> <p><b>i Note:</b> Selecting this check box activates the <b>Group By</b> list. If you don't select the check box, all the table records are grouped into clusters.</p>
Group By	<p>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.</p>
Purity Fields	<p>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 <b>Category</b> and <b>Assignment group.Name</b>.</p>
Filter	<p>Add filter conditions to apply to the input field records that you want to include in your clusters.</p> <ul style="list-style-type: none"> <li>○ The maximum number of records for clustering is limited to 300,000.</li> <li>○ For best results, aim for at least 2000 records as a minimum.</li> </ul> <p><b>i Note:</b> Script includes can't be referenced from the Filter. Use database views as an alternative.</p>
Processing Language	<p>Select the dominant language of the dataset you're training on the solution definition. If the dataset language is Italian, choose <b>Italian</b>. 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.</p>

Field	Value
	<p><b>i Note:</b> 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.</p>
Stopwords	<p>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 <b>Default Italian Stopwords</b> list appears. The <b>Default English Stopwords</b> list is also included. If you create a custom stopwords list, you can select it from the Stopwords field to add to your solution.</p>
Update Frequency	<p>Select how often you want the system to update your clusters with new and updated records.</p> <p><b>i Note:</b> The system pulls records based on the Group By filter conditions that you set on your clustering solution, if any.</p> <p>For example, if you select <b>Every 15 minutes</b>, 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.</p> <p>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.</p> <p>You can also choose not to update your clusters at all.</p>
Training Frequency	<p>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.</p> <p><b>i Note:</b> 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.</p>
Minimum number of records per cluster	<p>Enter the minimum number of records you want a cluster to contain. The value you enter must be 2 or higher.</p>

**4. Select the appropriate context menu option or button for your solution definition.**

Option	Description
<b>Save or Save &amp; Train</b>	Save your solution definition record so you can return to it later, or save and submit it for training.
<b>Submit or Submit &amp; Train</b>	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

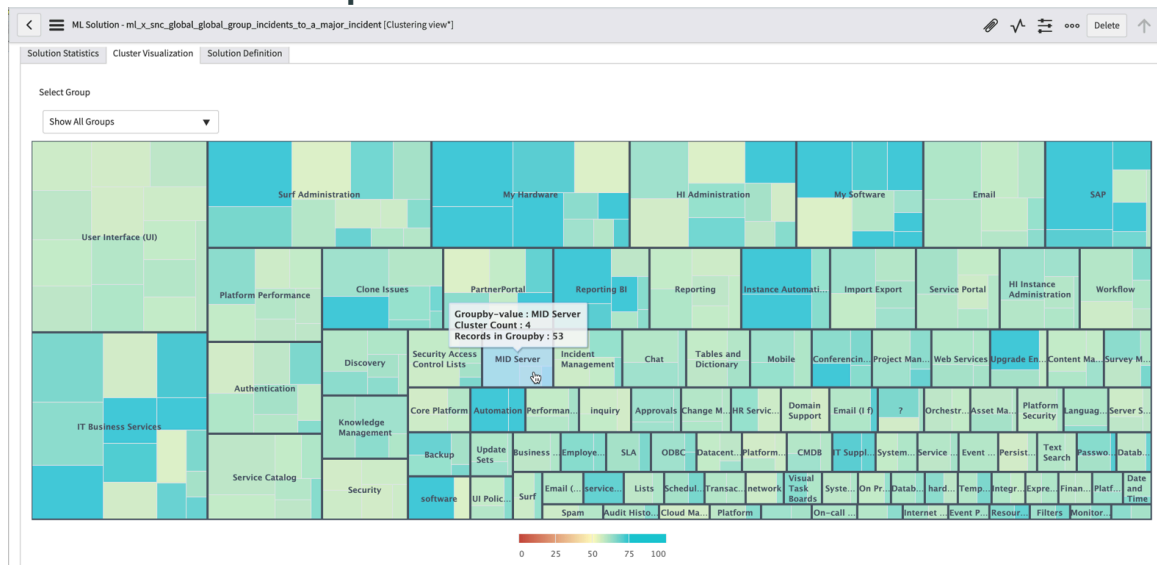
The system trains the solution and notifies you in real time when the training completes.

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.

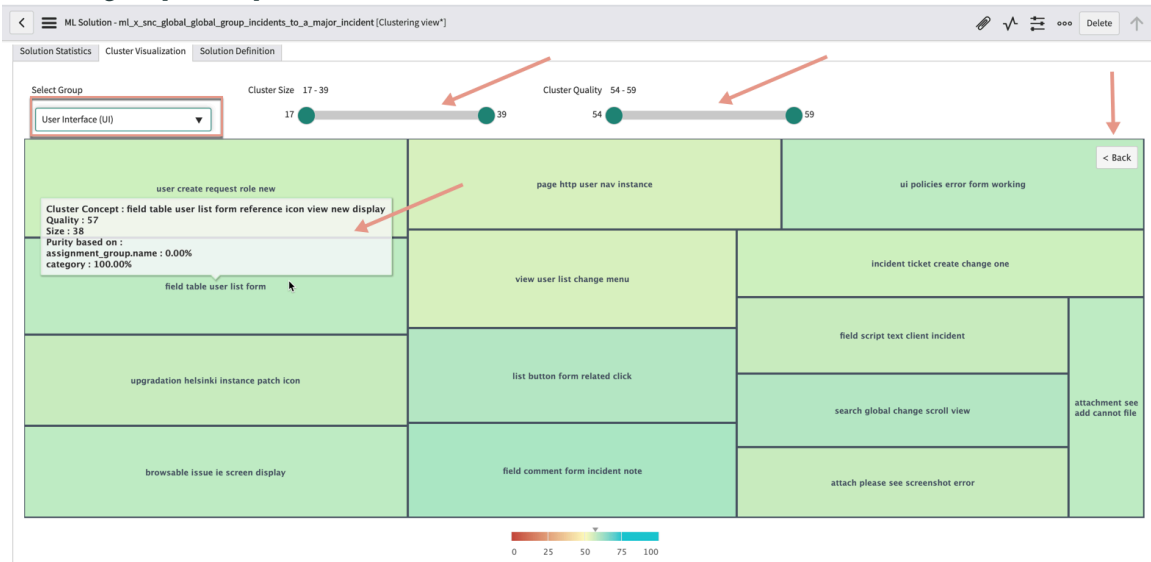
### Cluster visualization example



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.

## Cluster details page

Record ID	Groupby	Cluster	Solution	Rank
Incident:INC0014640	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.47
Incident:INC0011932	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.46
Incident:INC0018153	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.4
Incident:INC0015158	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.42
Incident:INC0013794	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.51
Incident:INC0011551	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.41
Incident:INC0016071	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.45
Incident:INC0012325	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.46
Incident:INC0017948	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.42
Incident:INC0017329	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.43
Incident:INC0012471	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.42
Incident:INC0017640	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.41
Incident:INC0013100	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.5
Incident:INC0011627	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.45
Incident:INC0012185	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.38
Incident:INC0017541	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.4
Incident:INC0012492	User Interface (UI)	116	ml_x_snc_global_group_incidents_to_a...	0.42

## 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

Cluster Id	Cluster Quality	Cluster size	Groupby Value
330	59.77	10	Email (1 f)
51	72.8	23	My Software
93	62.82	7	Email
111	56.58	23	Platform Performance
38	53.09	44	Surf Administration
234	58.71	9	Transaction and Session Management
97	58.54	20	Email
5	57.87	5	SAP
193	64.74	3	Incident Management
236	61.44	6	Approvals
354	57.88	6	Text Search
318	59.53	5	Lists
299	59.38	9	Date and Time
191	56.57	14	Incident Management
346	58.82	20	Event Management
35	60.72	29	Surf Administration
54	100	41	My Software
258	58.17	9	Performance Analytics
12	89.32	50	IT Business Services

- 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

Number of clusters	New clusters created	Records assigned to new clusters	Records clustered
0	370	5,469	5,469

### 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.

Cluster Summary  
55 [Default view view]

Cluster Id: 55

Cluster Name:

Groupby Value:

**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.

**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**

Cluster Summary  
13 [Default view view]

Cluster Id: 13

Cluster Name:

Groupby Value:

Solution: ml\_x\_snc\_global\_global\_cmdb

Cluster size: 115

Cluster Quality: 100

Update Delete

ML Cluster Detail **Generate Cluster Sample** New Search Record ID Search

Cluster = 13

Sample. Record ID Groupby Solution Rank

**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.
3. Select the **Calculate Purity** check

STEP 2: Select the input data for clustering, the frequency at which existing clusters are updated, and how frequently the data is reclustered

\* Table: Incident [incident]

Filter: 10000 records match condition @

Add Filter Condition Add "OR" Clause

-- choose field -- -- oper -- -- value --

Create Clusterinsight table

Use Group By

Update Frequency: Do not update

Training Frequency: Every 30 days

Minimum number of records per cluster: 2

box. Update Update & Retrain Delete

\* Fields: Short description

Calculate Purity

Purity Fields

Please select upto 10 fields to calculate purity for clusters. Empty field triggers auto Purity calculation.

\* Stopwords: Default English Stopwords

Processing Language: English

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.

Purity Fields

Available

Selected

Number

Opened

Short description

Caller (+)

Priority

State

Category

Assignment group (+)

Assigned to (+)

>

<

🔒

↑

↓

Please select upto 10 fields to calculate purity for clusters. Empty field triggers auto Purity calculation.

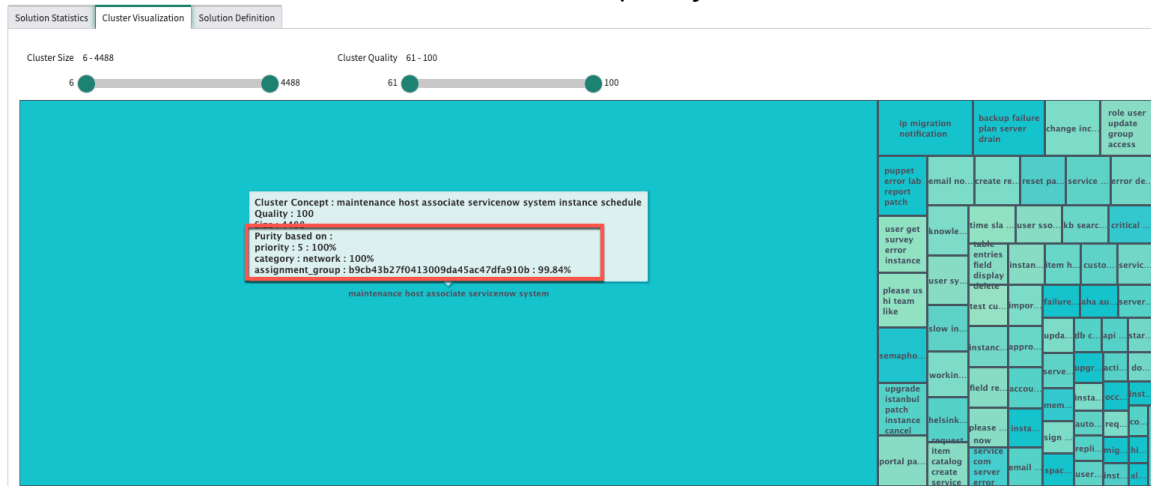
**5. Select Update or Update & Retrain.**

**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.

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 a data source**

Analyze a cluster by any of the dimensions available on the source table.

**Before you begin**

- Create a clustering solution definition or use an existing one.
- Role required: admin or ml\_admin

**About this task**

Analyze a cluster with a data source by accessing the cluster insight table. You can add a cluster insight table when creating a solution or editing an existing solution. For more information, see <https://www.servicenow.com/community/>

STEP 2: Select the input data for clustering, the frequency at which existing clusters are updated, and how frequently the data is reclustered

\* Table Incident [incident] ▼

Filter 10000 records match condition 🗑️

Add Filter Condition Add "OR" Clause

-- choose field -- ▼ -- oper -- -- value --

Create ClusterInsight table

p/2301006 [🔗](#)

## Procedure

1. Navigate to **All > Predictive Intelligence > Clustering > Solutions**.
2. Select a solution.
3. Under **Cluster Visualization**, select a cluster.
4. On the **Cluster Insight** table, select the list menu icon (☰) next to a column header, and select **Bar Chart**.

## Result

The **Create a Report** page shows your selected insight data in a bar chart. You can use this method on different insights and tables and with different visualizations.

## Create and train a regression solution

Train your solution by using historical data to predict numeric outputs, such as a temperature or a stock price. For example, you can use regression to estimate the time it takes to resolve an incident or a case.

## Before you begin

**Important:** Support for new regression solutions is deprecated in the Xanadu release. You will still be able to edit and train any existing solutions, but you won't be able to create new ones.

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 actuals 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, click **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 <b>Regression Test for DB Restore</b> .
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	<p>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 <b>Incidents in the last 3 months</b>.</p> <p>If you don't have a relevant word corpus, follow the steps <a href="#">to create a word corpus</a> first. When the word corpus is complete, you can select it from the Word Corpus field in your Regression Definition form.</p> <p>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.</p> <p><b>Note:</b> A pre-trained model is used instead of the Word Corpus for users who activated Predictive Intelligence starting in Utah.</p>
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	<p>Select the field whose value you want the predictive model to set.</p> <p>In general, a good output field is a numeric, integer, or floating point field.</p>

Field	Value
	<p>In this example scenario, you use the <b>Duration</b> field to measure a length of time. The output field should generate a numeric value.</p>
Fields	<p>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 <b>Short description</b>, <b>Source datacenter</b>, <b>Target datacenter</b>, and <b>Database size</b>. (short_description, Sourcedc, Targetdc, and Dbsize.) Input field types can be string, nominal, or numeric.</p>
Filter	<p>(Optional) Add filter conditions to the output field records that you want to train using regression.</p> <p><b>Note:</b></p> <ul style="list-style-type: none"> <li>○ The minimum number of records for regression training is 10,000 records.</li> <li>○ The maximum number of records for regression training is limited to 300,000.</li> </ul>
Processing Language	<p>Select the primary language of the dataset that you're training on the solution definition. If the dataset language is Italian, choose <b>Italian</b>. 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.</p> <p><b>Note:</b> 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.</p>
Stopwords	<p>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 <b>Default Italian Stopwords</b> list is displayed. The <b>Default English Stopwords</b> 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 <b>Default English Stopwords</b> list.</p>
Training Frequency	<p>Select how often the system regenerates the solution based on records matching the <b>Filter</b>. Your options include:</p> <ul style="list-style-type: none"> <li>○ Run Once</li> <li>○ Every 30 days</li> <li>○ Every 60 days</li> <li>○ Every 90 days</li> </ul>

Field	Value
	<ul style="list-style-type: none"> <li>○ Every 120 days</li> <li>○ Every 180 days</li> </ul> <p>In this scenario, you select Every 30 days</p> <p>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.</p> <p><b>Note:</b></p> <ul style="list-style-type: none"> <li>○ The minimum number of records required for regression solution training is set at 10,000.</li> <li>○ 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.</li> </ul>

4. Click the appropriate context menu option or button for your solution definition.

Option	Description
<b>Save or Save &amp; Train</b>	Save your solution definition record so you can return to it later, or save and submit it for training.
<b>Submit or Submit &amp; Train</b>	Create your solution definition record and submit it, or submit and train it.

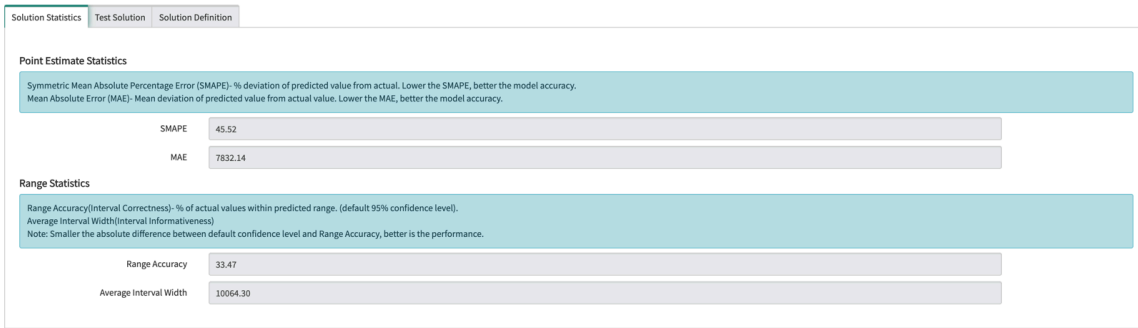
5. If you submitted the solution for training, click **OK** on the **Training Activation** window to confirm.

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. 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 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. Click the **Run Test** button to find the prediction output.

**Test Type**  
 Test a single record or a batch of records  
 Test Type:  Single Test  Batch Test

**Prediction Input**  
 Provide input field values to get the predicted value

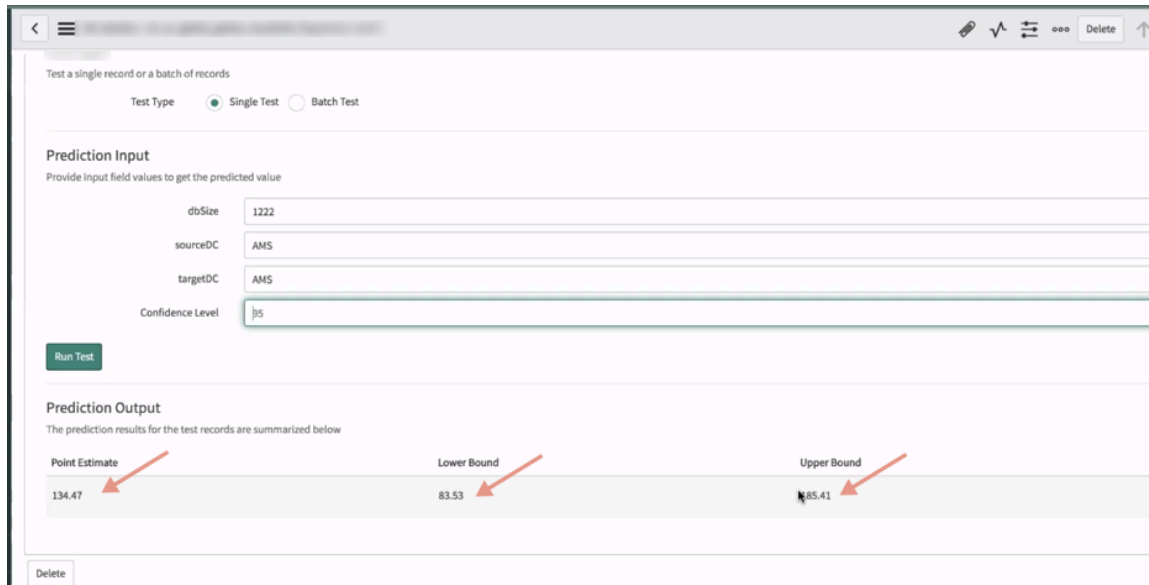
dbSize	1222
sourceDC	AMS
targetDC	AMS
Confidence Level	95

**Run Test** button highlighted with a red box.

**Prediction Output**  
 The prediction results for the test records are summarized below

Delete button and Related Links (Show training progress) are also visible.

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.

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

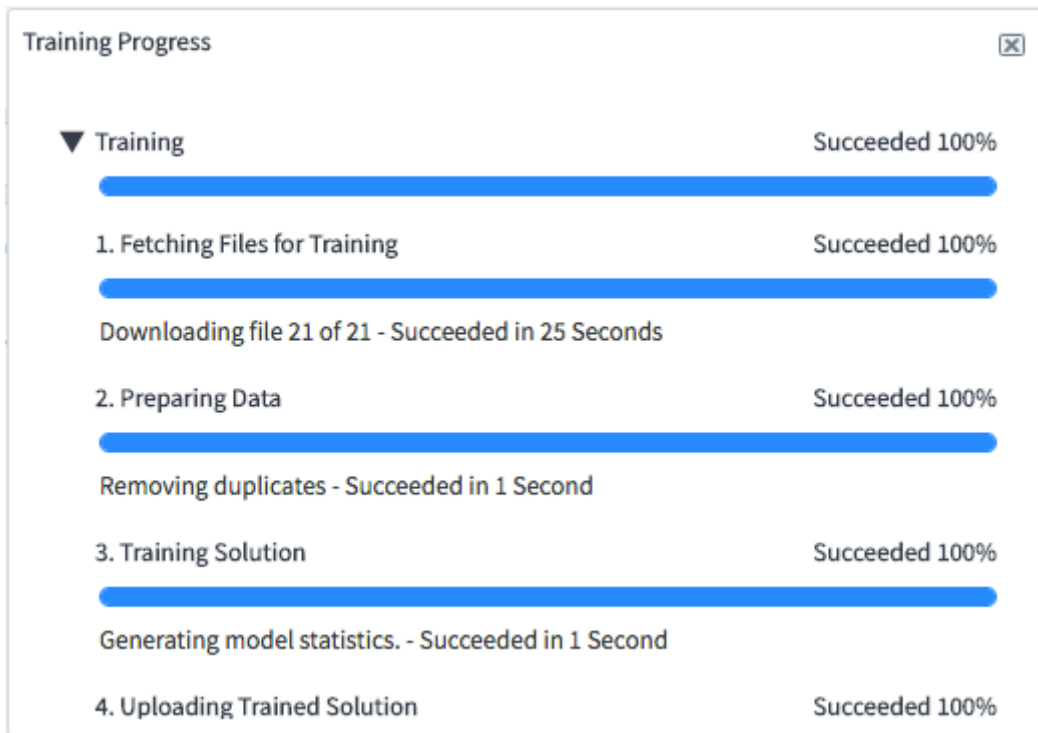
1. Navigate to **All > Predictive Intelligence > Classification > Solutions** or **Predictive Intelligence > Similarity > Solutions**.
2. In the ML Solutions list, select the solution whose progress or statistics you want to view.

### 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.

### Before you begin

- Role required: admin, ml\_admin, or ml\_report\_user

### About this task


**Important:** With the Xanadu 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 Xanadu 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.

## 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 and scores generated during solution training to determine if the similarity score threshold is sufficient.

## Before you begin

- Train a similarity solution
- Role required: admin or ml\_admin


## About this task

Solution training generates paired data record examples with a percentile score that represents the degree of similarity between the two records. The higher the score, the higher the similarity. A score of 100 indicates identical records and a score of 0 indicates dissimilar records.

The solution only returns similarity results that have a score that's higher than the threshold.

- Note:** The similarity filters specified in the solution definition aren't applied for similarity examples and are only applied during prediction.

## Procedure

1. Navigate to **All > Predictive Intelligence > Similarity > Solutions**.
2. In the ML Solutions list, locate your solution and click the Reference Lookup icon.
 
3. Click **Open Record**.
4. In the Related Links section, click **Similarity Examples**.
5. Review the similarity examples and their threshold scores to determine the accuracy and coverage levels you want applied to your solution similarity results.
 

The higher the similarity score, the more precise it is and the less coverage it offers. The lower the similarity score, the more coverage it has and the less precision it offers.
6. Based on your review, determine whether to increase or decrease the similarity score threshold value for your similarity solution.

## What to do next

If you decide to adjust the score for your similarity solution, [update its similarity solution threshold](#).

## 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.
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.

**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](#) ↗

## 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:

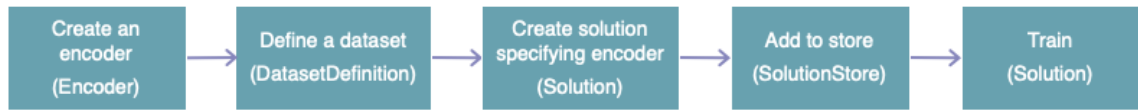
### ML API flow – Solution, encoder, and estimate training



**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.

### ML API flow – Solution API training with encoder



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 breakdown to learn how to configure and train a solution.

#### Configure and train a solution



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();
```

- [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\(\)](#)
- [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"}
```

<code>getLatestVersion()</code>	<code>getStatus()</code>
<a href="#">ClassificationSolution - getLatestVersion()</a>	<a href="#">ClassificationSolutionVersion - getStatus()</a>

getLatestVersion()	getStatus()
<i>ClusteringSolution</i> - <a href="#">getLatestVersion()</a>	<i>ClusteringSolutionVersion</i> - <a href="#">getStatus()</a>
<i>Encoder</i> - <a href="#">getLatestVersion()</a>	<i>EncoderVersion</i> - <a href="#">getStatus()</a>
<i>PredictabilityEstimate</i> - <a href="#">getLatestVersion()</a>	<i>PredictabilityEstimateVersion</i> - <a href="#">getStatus()</a>
<i>RegressionSolution</i> - <a href="#">getLatestVersion()</a>	<i>RegressionSolutionVersion</i> - <a href="#">getStatus()</a>
<i>SimilaritySolution</i> - <a href="#">getLatestVersion()</a>	<i>SimilaritySolutionVersion</i> - <a href="#">getStatus()</a>

### 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');
```

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",
      "predictedSysId": ""
    }
  ]
}
```

## Testing and monitoring predictions

Evaluate the coverage and precision of your machine-learning (ML) solutions by testing them. Once deployed, track the performance over time to see how your solutions perform. Use the information to make changes to your solutions to improve the predictions.

### 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

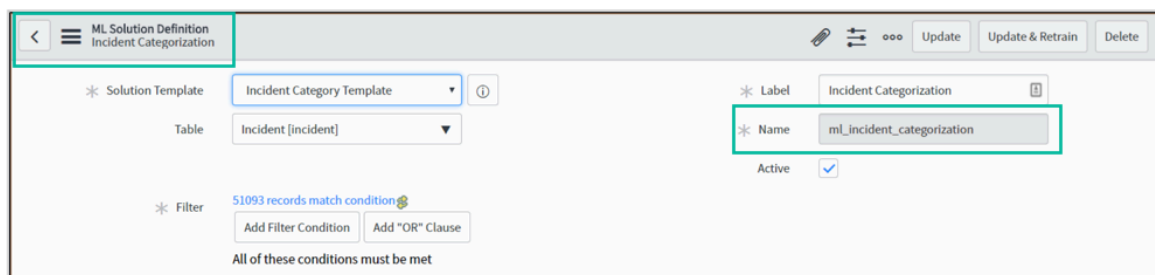
### About this task

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.

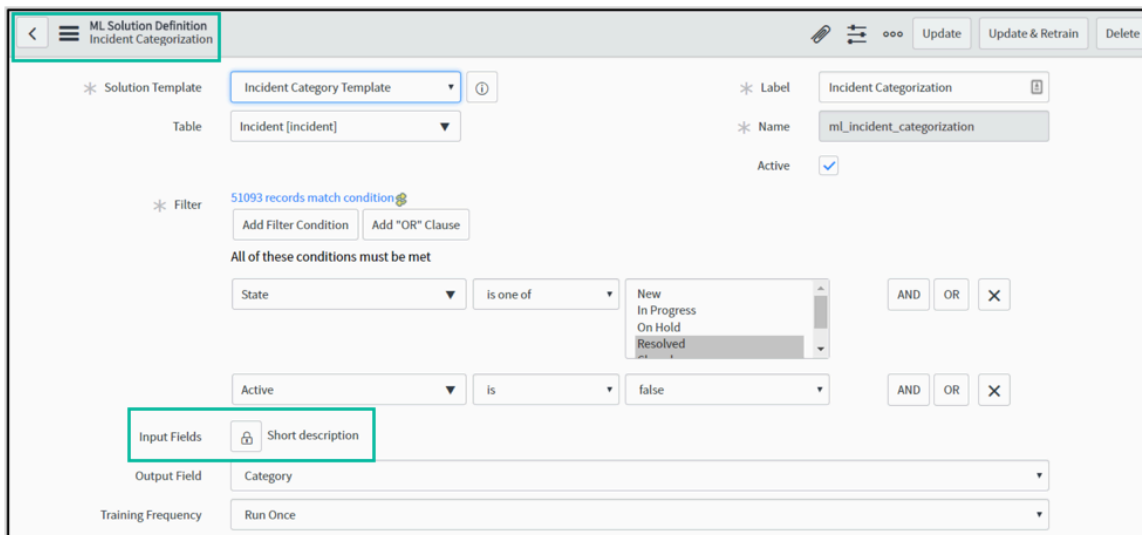
This scenario illustrates a classification solution prediction for a hypothetical ML solution that you have previously created and trained. You can also use the REST API Explorer to test a similarity solution prediction.

## 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
<b>Namespace</b>	now (leave as default)
<b>API Name</b>	Predictive Intelligence
<b>API Version</b>	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`.

**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 example, `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, `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](#) 

**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.

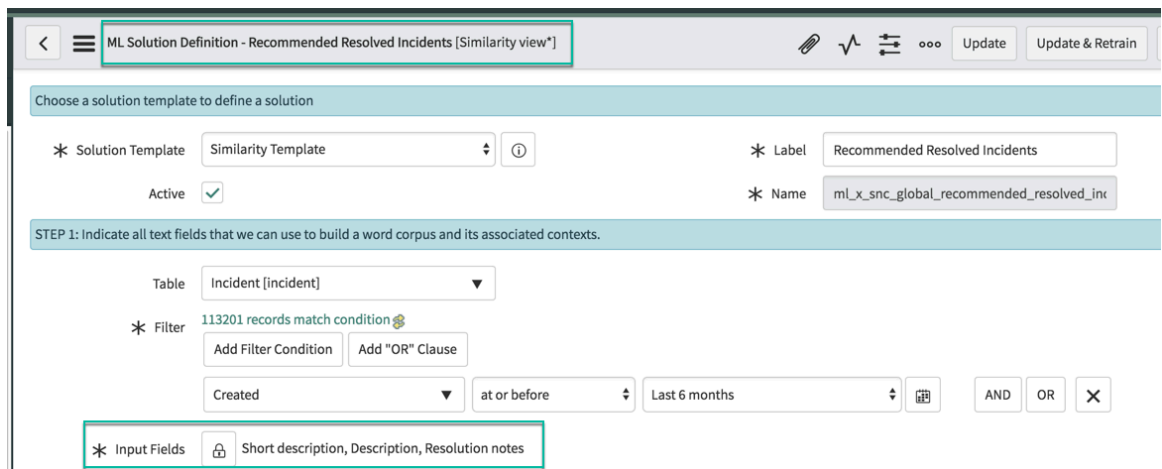
This scenario illustrates a similarity solution prediction for a hypothetical ML solution that you have previously created and trained. You can also use the REST API Explorer to test a classification solution prediction.

## 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**.
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
<b>Namespace</b>	now (leave as default)
<b>API Name</b>	Predictive Intelligence
<b>API Version</b>	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.

#### Response Body

```

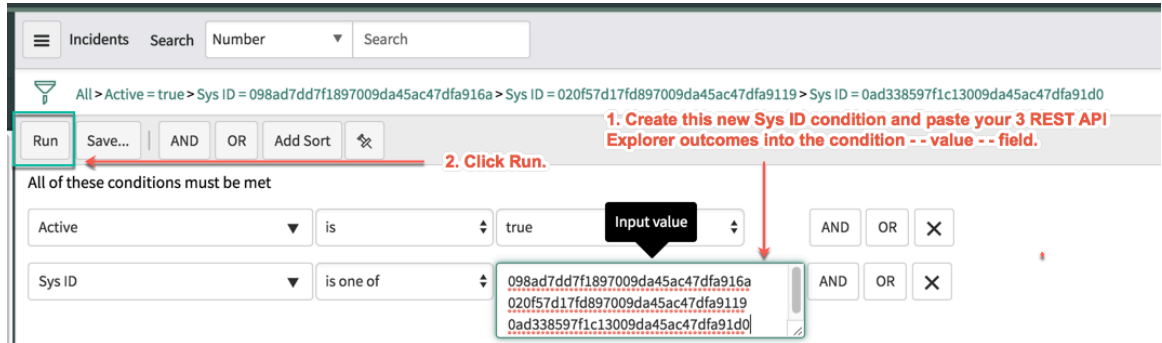
},
"output": [
  {
    "outcome": "098ad7dd7f1897009da45ac47dfa916a",
    "confidence": "83.38670134544373",
    "threshold": 80
  },
  {
    "outcome": "020f57d17fd897009da45ac47dfa9119",
    "confidence": "62.12549805641174",
    "threshold": 80
  },
  {
    "outcome": "0ad338597f1c13009da45ac47dfa91d0",
    "confidence": "61.85050010681152",
    "threshold": 80
  }
]

```

14. In your original browser tab, navigate to **ServiceDesk > Incidents**.

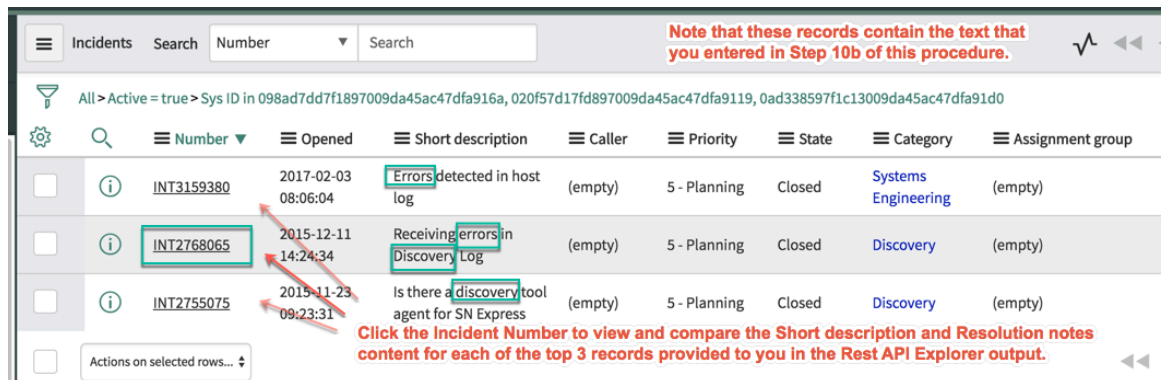
15. As shown in the image below, set filter conditions for the three REST API outcomes to the Incidents table list view.

- 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**.

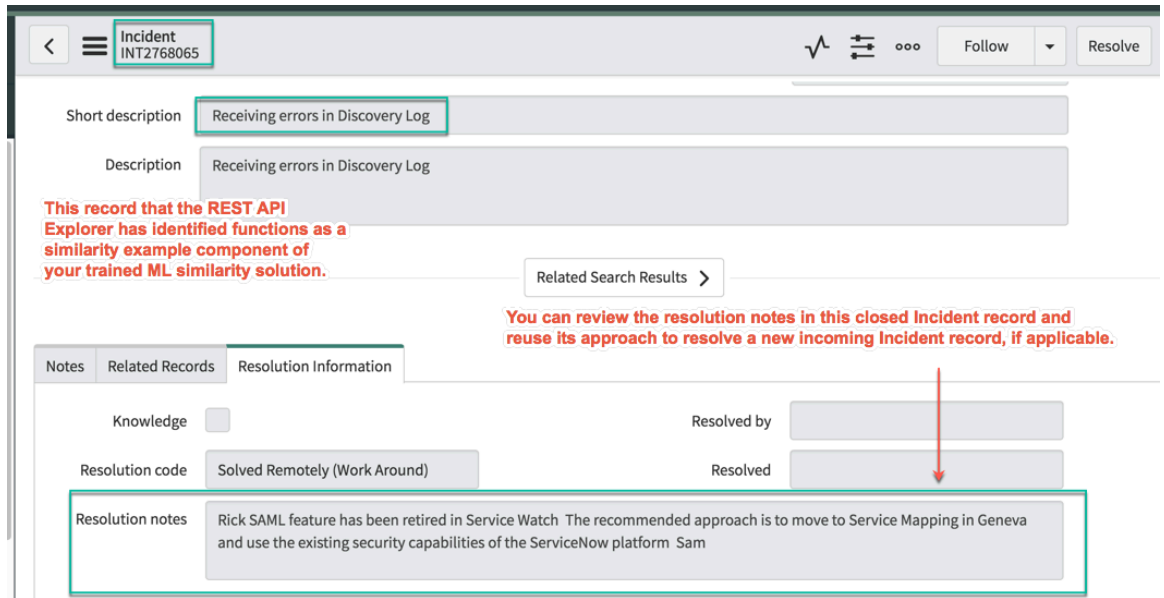


16. Per the image below, compare the returned list of incidents with the input for the prediction output in the REST API Explorer.

- 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.

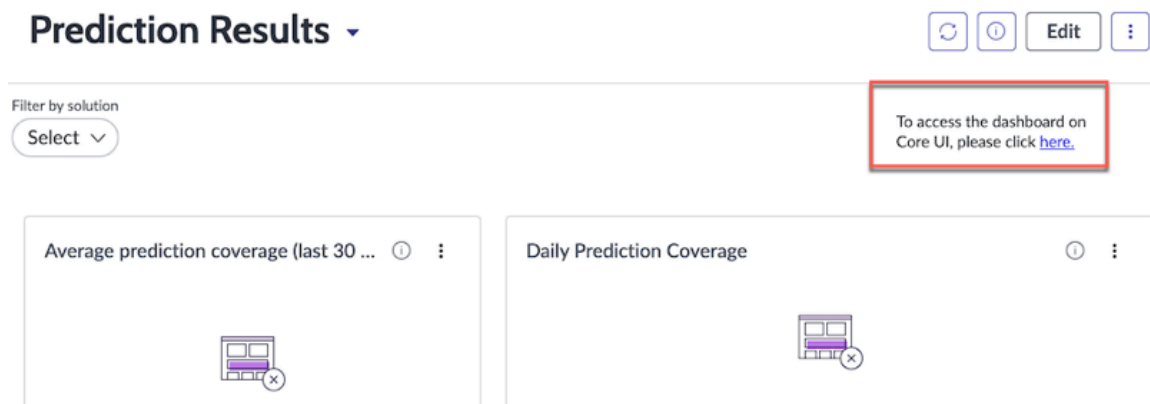
**Before you begin**

- 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 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 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**.

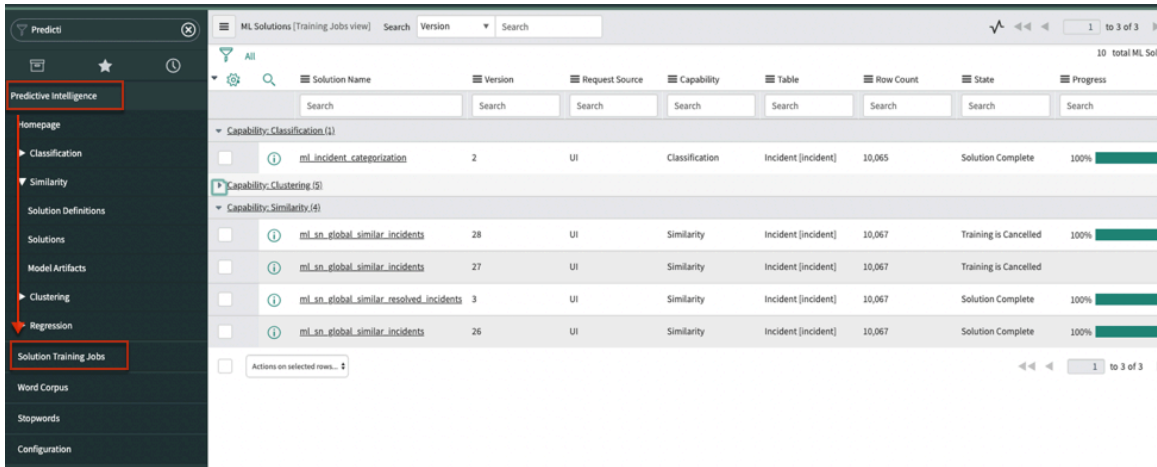
**i Note:** The ML scheduler limits the number of trainings an instance can commit to 50 new ML training requests per instance within a 24 hour window. This excludes scheduled re-training requests. In addition, clustering and similarity updates are also excluded from this limit, even if the new training requests exceed 50 within a 24 hour window.

## ML Training Jobs view summary

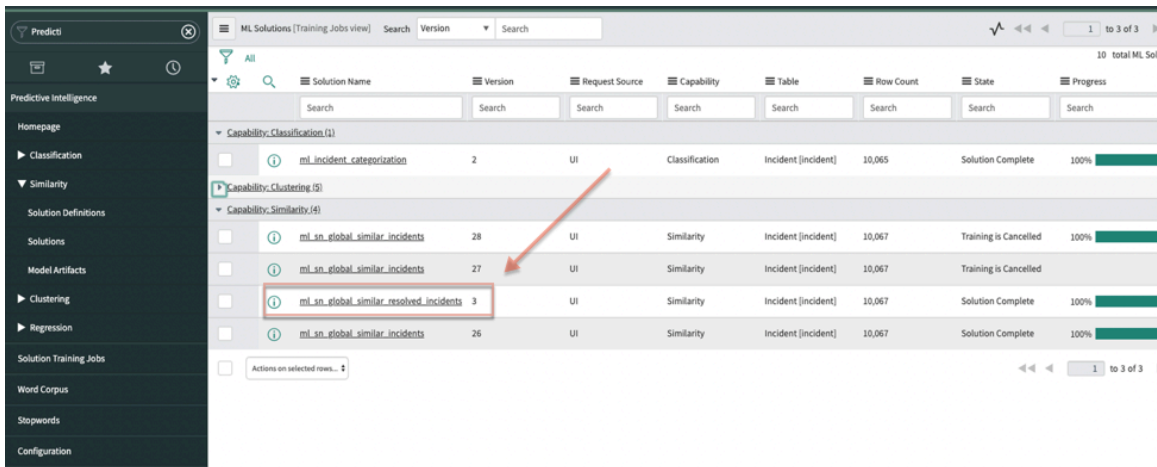
The view shows all ML training jobs grouped by the four Predictive Intelligence capability frameworks: classification, similarity, clustering, and regression.

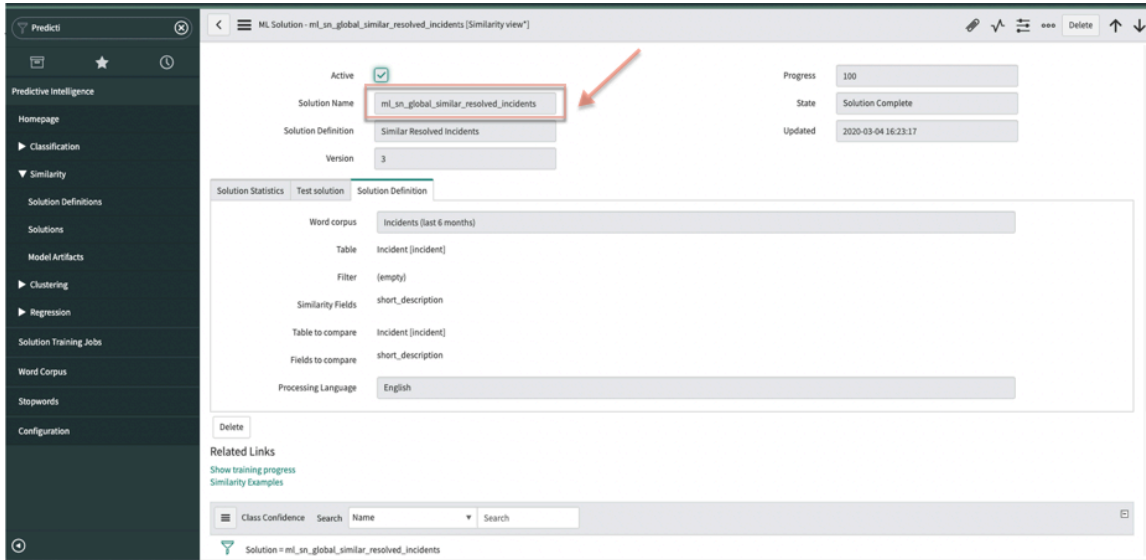
Each record displays values such as solution name, version, training state, and training completion percentage.

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.





## 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](#).

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**.

Name	Description	Label	Plural	Updated
sn_customerservice_knowledge_view	Join kb_knowledge, kb_template_what_is, ...	Knowledge View	Knowledge View All	2020-03-25 23:51:02
sn_fsm_mobile_native_questionnaires_db_view	View for native questionnaires associate...	FSM Questionnaires		2020-03-12 16:07:47
sn_fsm_mobile_use_part_view	View joining "alm_asset" and "sm_asset_u...	sn_fsm_mobile_use_part_view		2020-04-23 02:12:36
sp_page_widgets	DB view to retrieve all widgets across a...	Service Portal - Page Widget Mapping	Service Portal - Page Widget Mapping	2020-02-21 04:41:34
subscribed_users_with_role	All subscribed users with role	Subscribed Users with Role	Subscribed Users with Role	2019-04-23 05:25:22
task_assessment_detail	Detailed task assessment responses view ...	Task Assessment Details	Task Assessment Details	2020-02-10 22:42:58
upgrade_history_update_xml		Upgrade History Sys Update XML	Upgrade History Sys Update XMLs	2020-02-10 22:42:58
users_by_skill_category		Users By Skill Category	Users By Skill Category	2019-02-12 18:01:08
user_has_license_role	Joins sys_user_has_role with license_rol...			2020-04-23 16:28:17
user_role_license_role	Join user_role to license_role to easily...	User Role to License Role	User Roles to License Roles	2020-04-16 11:48:17
user_skills_by_group		User Skills By Group	User Skills By Group	2019-02-12 16:57:05

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.

**Database View**  
sn\_customerservice\_knowledge\_view

This record is in the Customer Service application, but Global is the current application. To edit this record click here.

Name: sn\_customerservice\_knowledge\_view  
Application: Customer Service  
Label: Knowledge View  
Plural: Knowledge View All

Description: Join kb\_knowledge, kb\_template\_what\_is, kb\_template\_how\_to, kb\_template\_faq, and kb\_template\_kcs tables for viewing all fields in knowledge.

Related Links  
Try It

View Tables	Search	Order	Search
View = sn_customerservice_knowledge_view			
Table	Order	Variable prefix	Where clause
kb_knowledge	100	knowledge	knowledge.active=1 && knowledge.workflow...
kb_template_kcs_article	200	kcs	knowledge_sys_id=kcs_sys_id
kb_template_how_to	200	howto	knowledge_sys_id=howto_sys_id
kb_template_what_is	200	whatis	knowledge_sys_id=whatis_sys_id
kb_template_faq	200	faq	knowledge_sys_id=faq_sys_id

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.

Label	Name	Capability	Created	Table	Fields	Active
Similar Knowledge Articles All	ml_sn_sn_customerservice_global_similar...	Similarity	2020-03-19 07:07:11	Knowledge View [sn_customerservice_knowledge_view]	knowledge_short_description,knowledge_te...	true
Similar Open Problems	ml_sn_global_global_similar_open_problems	Similarity	2020-02-24 12:03:46	Problem [problem]	short_description	true
User profile based recommendation	ml_sn_sn_hr_core_global_user_profile_bas...	Similarity	2020-02-21 07:05:33	HR Profile [sn_hr_core_profile]	user.department.name,position.position,l...	true
Similar Knowledge Articles for HR Task	ml_x_sn_hr_core_global_similar_knowledge...	Similarity	2020-01-29 02:14:37	Knowledge [kb_knowledge]	short_description,description,text	true
Similar Knowledge Articles for HR Case	ml_sn_sn_hr_core_global_case_to_articles	Similarity	2020-01-09 00:33:13	Knowledge [kb_knowledge]	short_description,text,description	true
Similar Open Change Requests	ml_sn_global_global_similar_open_change_...	Similarity	2020-01-03 06:36:54	Change Request [change_request]	short_description	false
Similar Knowledge Articles	ml_sn_global_articles_similar_to_cases	Similarity	2019-10-03 21:38:11	Knowledge [kb_knowledge]	text,short_description,description	true
Similar Open Incidents	ml_sn_global_similar_open_incidents	Similarity	2019-09-18 15:43:19	Incident [incident]	short_description	true

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.

Similarity Definition  
Similar Knowledge Articles All [ML view]

This record is in the Customer Service application, but Global is the current application. To edit this record click here.

Label: Similar Knowledge Articles All    Active:

Name: ml\_sn\_sn\_customerservice\_global\_simila

STEP 1: Please select a word corpus that is relevant to your solution, or create a new word corpus so you can select it from this form

Word Corpus: All Articles and Cases

STEP 2: Please select the table & field(s) you want to use to retrieve your similarity results, and the table and fields you want compared

Table: Knowledge View [sn\_customerservice\_knowledge\_view]    Test Table: Case [sn\_customerservice\_case]

Fields: knowledge\_short\_description,knowledge\_text,knowledge\_description    Test Fields: short\_description

Filter: 14991 records match condition  
knowledge\_workflow\_state=published\*knowledge\_kb\_knowledge\_base=05ff44289f011200550bf7b6077fca3\*ORknowledge\_kb\_knowledge\_base=2b292926c3302200e7c7d44d81d3ae10^EQ

Processing Language: English

Stopwords: Default English Stopwords

STEP 3: Select how often you want to refresh the data you use to retrieve your similarity results

Training Frequency: Every 30 days

Update Frequency: Every 1 day

Related Links  
Active Solution - Similarity Examples  
Add solutions to current update set  
Refresh similarity window

## ServiceNow® apps and features that use Predictive Intelligence

Review this topic to see a sample of ServiceNow applications and features that leverage Predictive Intelligence functionality.

When applicable, ServiceNow teams work together to create products that apply the artificial intelligence and machine learning capabilities that Predictive Intelligence provides.

These teams have customers in different business units (BUs) and industries. For example, the Customer Service Management (CSM) marketing, product, design, development, and documentation teams partner with the platform Predictive Intelligence team to deliver BU-specific products to ServiceNow CSM customers.

Here is a listing of some of the ServiceNow products that use Predictive Intelligence functionality.

<a href="#">Machine Learning solutions for Strategic Portfolio Management</a>	<a href="#">Machine Learning solutions for Knowledge Management</a>	<a href="#">Machine Learning solutions for IT Service Management</a>
<a href="#">Machine Learning solutions for HR Service Delivery</a>	<a href="#">Machine Learning solutions for HR Service Delivery Agent Workspace</a>	<a href="#">Issue assignment using the Governance, Risk, and Compliance Predictive Intelligence plugin</a>
<a href="#">Machine Learning solutions for Vulnerability Response</a>	<a href="#">Machine Learning solutions for Customer Service Management</a>	<a href="#">Machine Learning solutions for Search administration</a>
<a href="#">Machine Learning solutions for Field Service Management</a>	<a href="#">Machine Learning solutions for Event Management</a>	<a href="#">Machine Learning solutions for Flow Designer</a>

### 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 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.

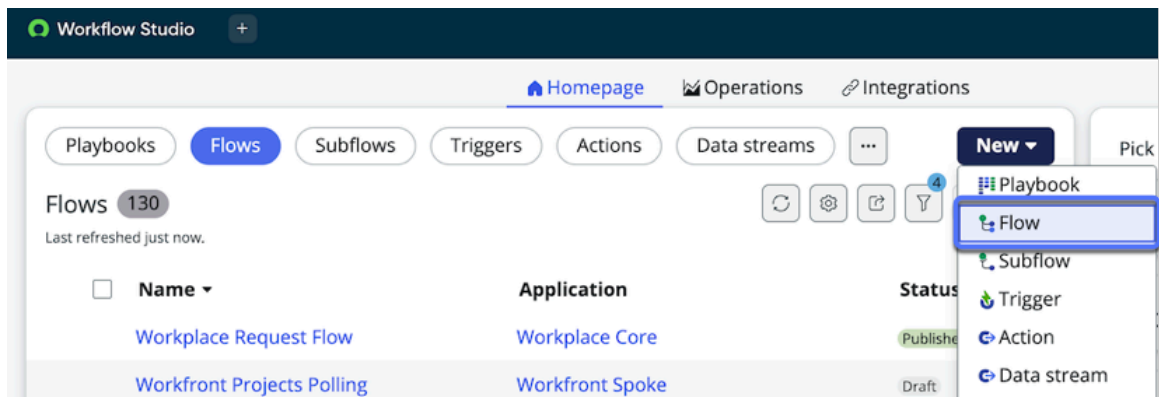
In this example procedure, you create a flow that implements the ml\_incident\_categorization solution in a Workflow Studio action. You can find this solution by searching on the ML Solutions [ml\_solution] table, as shown in the image below. Confirm that the solution you use has been trained and its **Active** value is set to **true**.

Search	Search	Search	Search	Search	Search	Search
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

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 <code>Auto-assign Category to Incident</code> .
Description	Enter a brief summary description of what the flow delivers. For example, you enter the following: <code>When an incident is created, it automatically triggers this flow, which uses ML Solutions to predict the correct Category for the incident.</code>
Application	Select <b>Global</b> .
Protection	Select <b>--None--</b> or <b>Read-only</b> . In this scenario, you select <b>--None--</b> .
Run As	Select <b>User who initiates session</b> .

Field	Description
Run with role(s)	Leave blank.
Flow priority default	Medium (default).

4. Select **Build flow**.

### Let's get the details for your flow

Flow name \* ⓘ

Application \* ⓘ

Description ⓘ

[Hide additional properties](#)

Protection ⓘ

Run as ⓘ

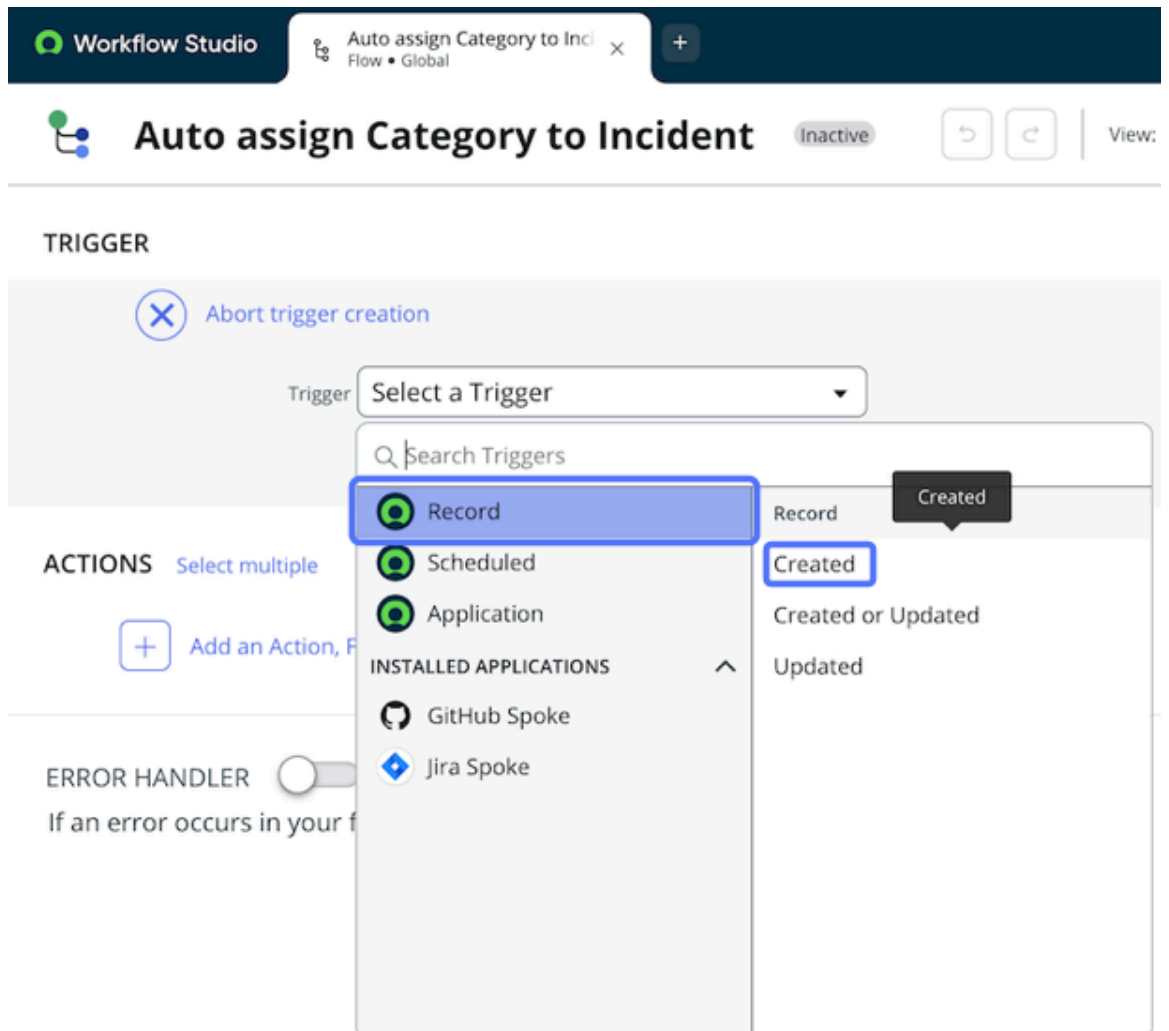
Run with roles ⓘ

Flow priority default ⓘ

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**.

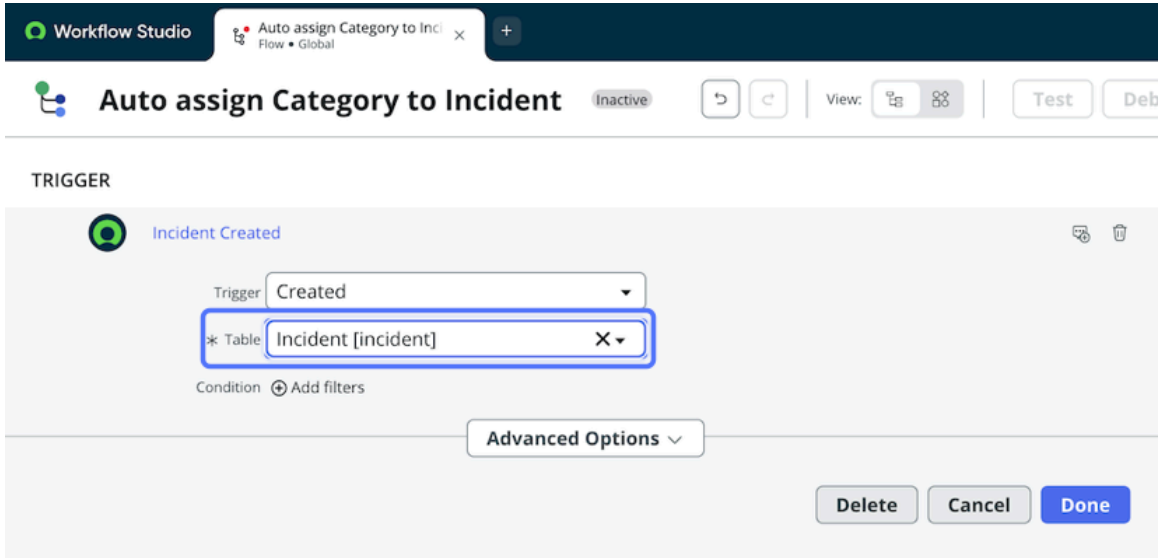
5. In the TRIGGER section of the Flow screen, configure the following fields to create a trigger for the flow.

- 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.



c. **Condition:** Select **Add filters** if you want to add any conditions to the flow.

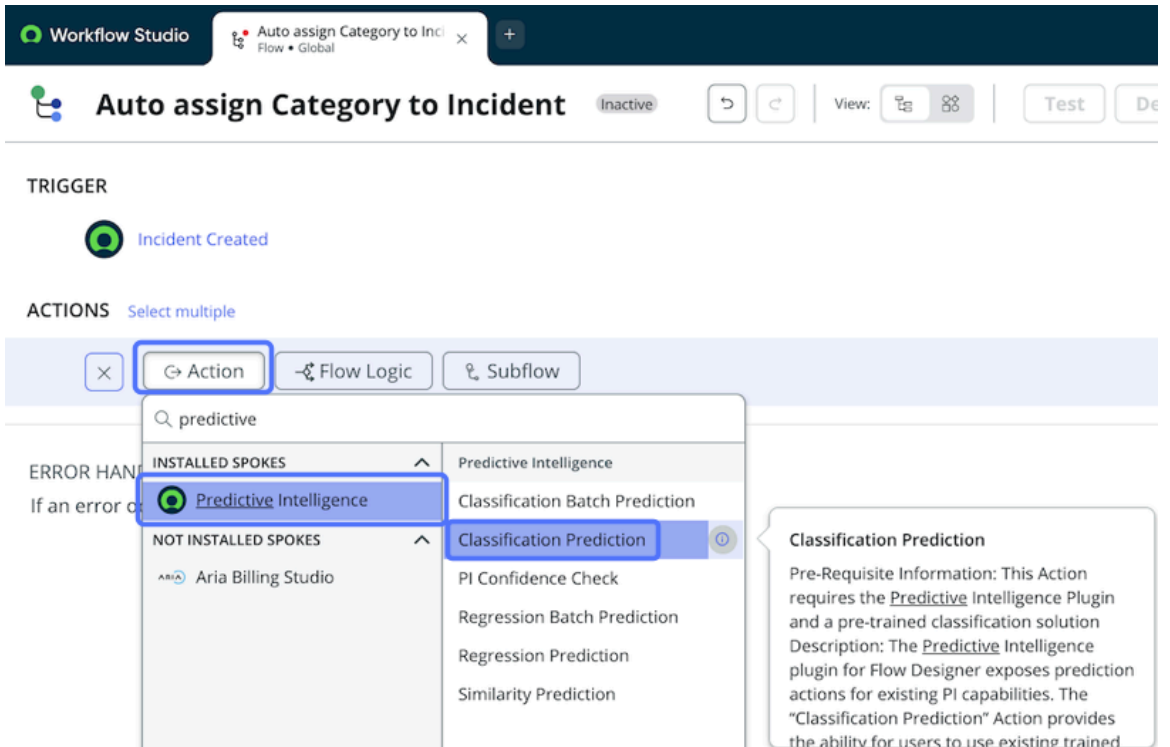
d. **Optional:** Open the **Advanced Options** panel to view additional conditions you can apply to the flow.

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**.



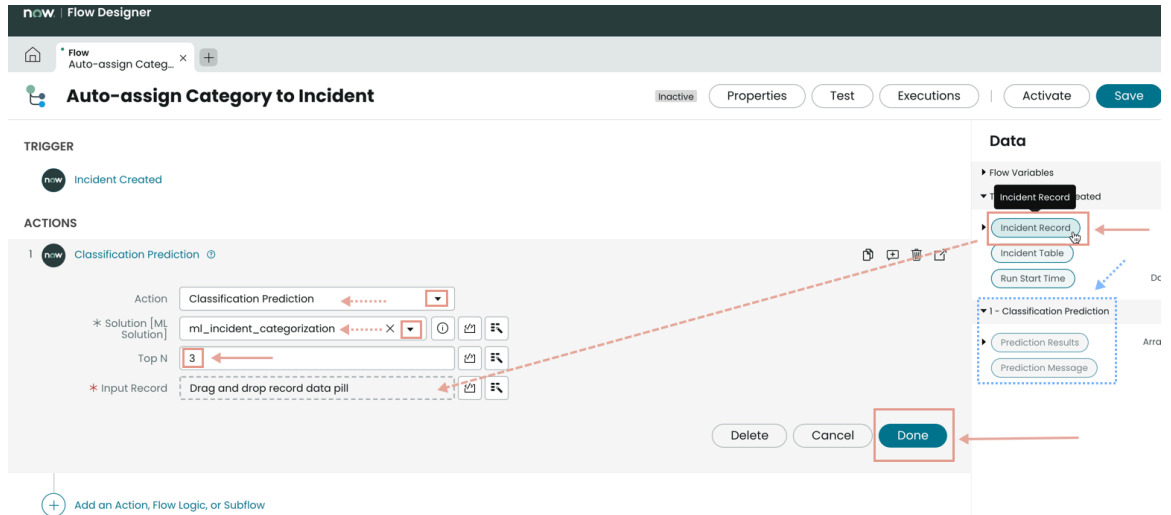
Select the information icon (i) to see the description of a Classification Prediction.

**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.

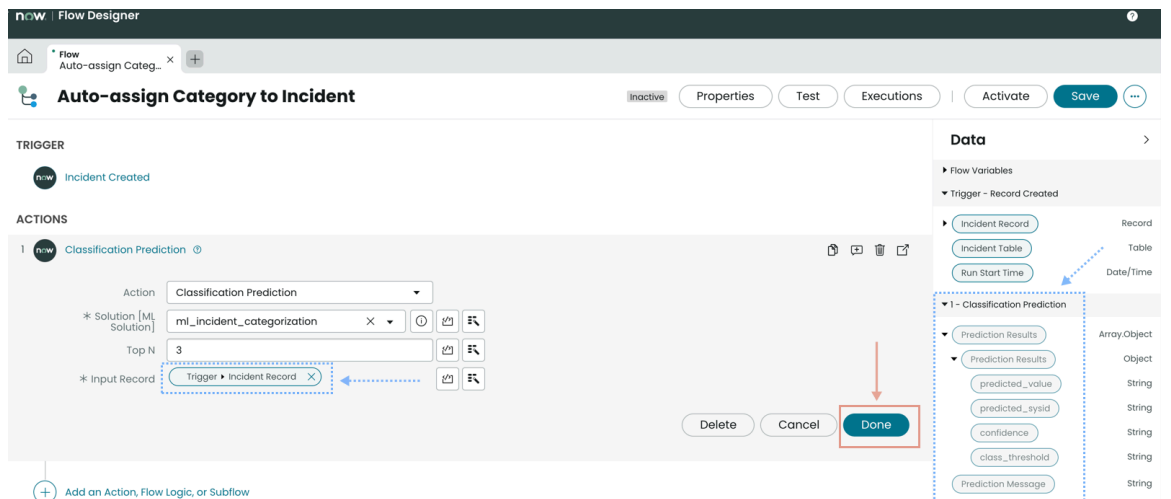


The Action, Solution Name, Top N, and Input Record values provide a base for the Category prediction.

**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.**

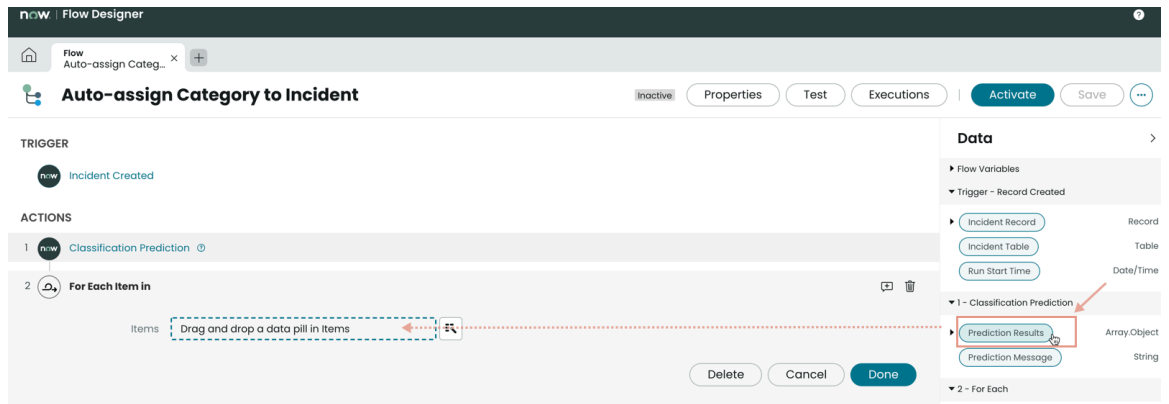
**Result:** The Classification Prediction action is completed in the flow and its data pills appear in the Data section of the screen.



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.



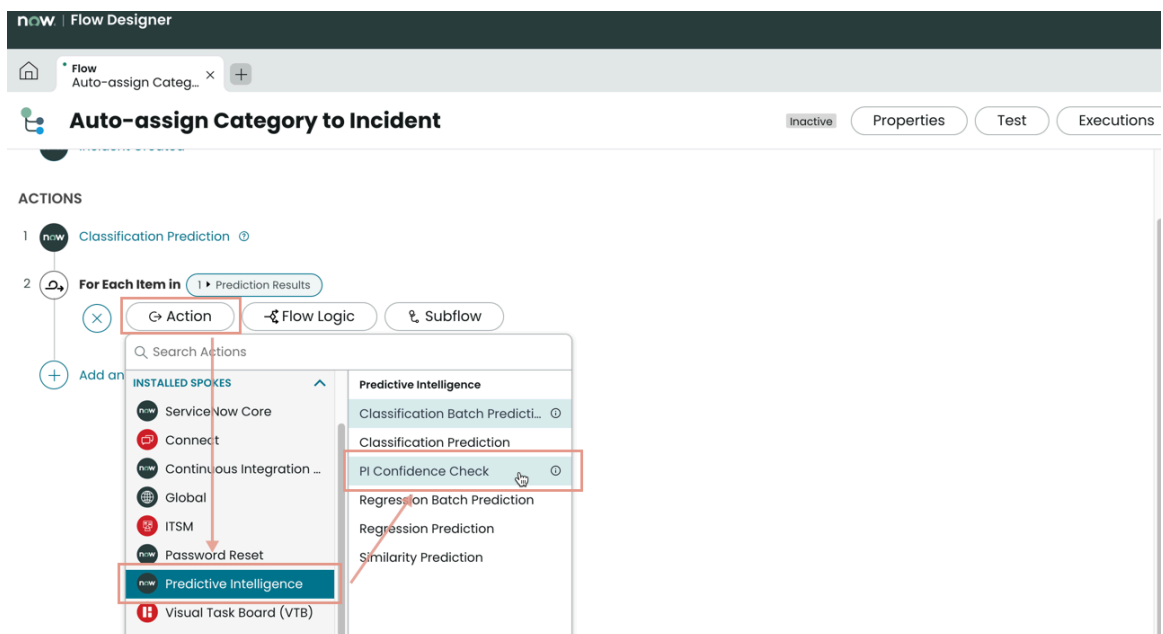
**Note:** In order to access the list of items in the Regression Prediction action, you don't need the **For Each Item in** flow logic.

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.

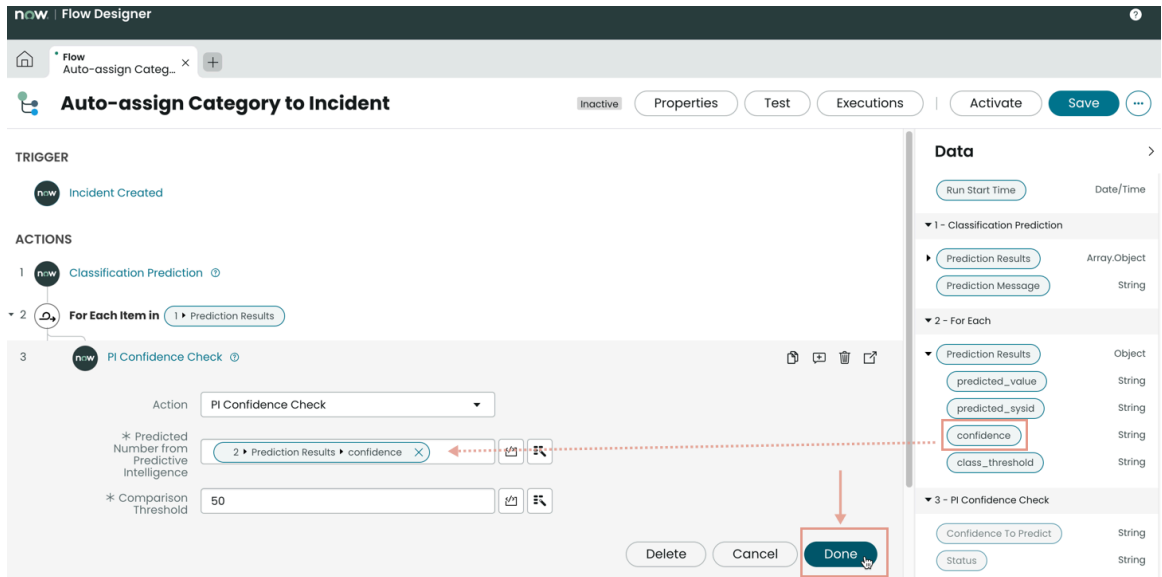


9. Drag and drop the **confidence** data pill into the **Predicted Number from Predictive Intelligence** field.

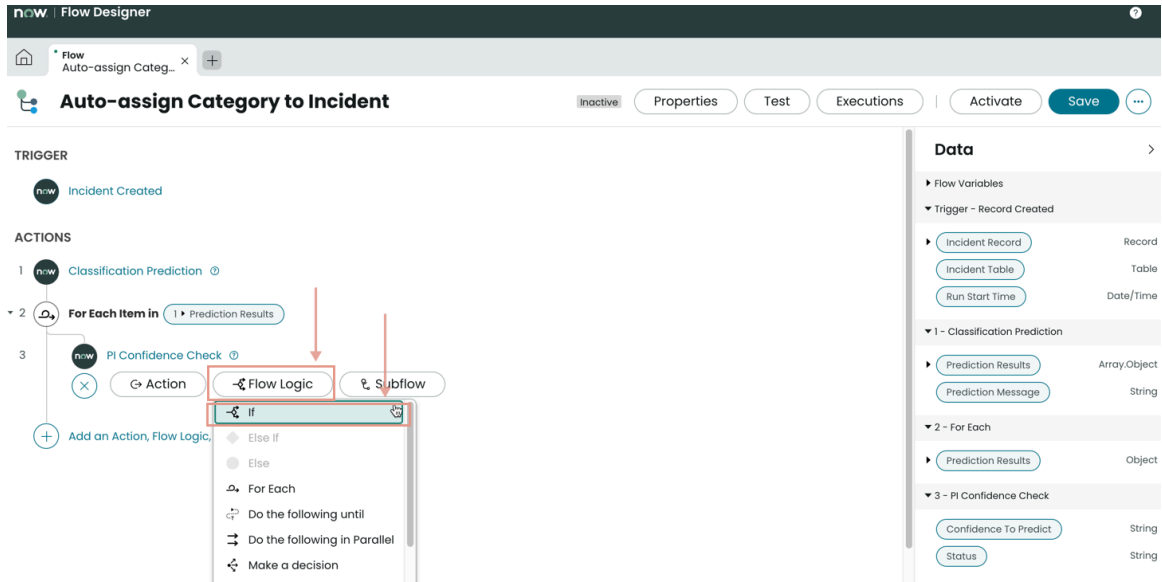
10. Enter 50 in the **Comparison Threshold** field.

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**.

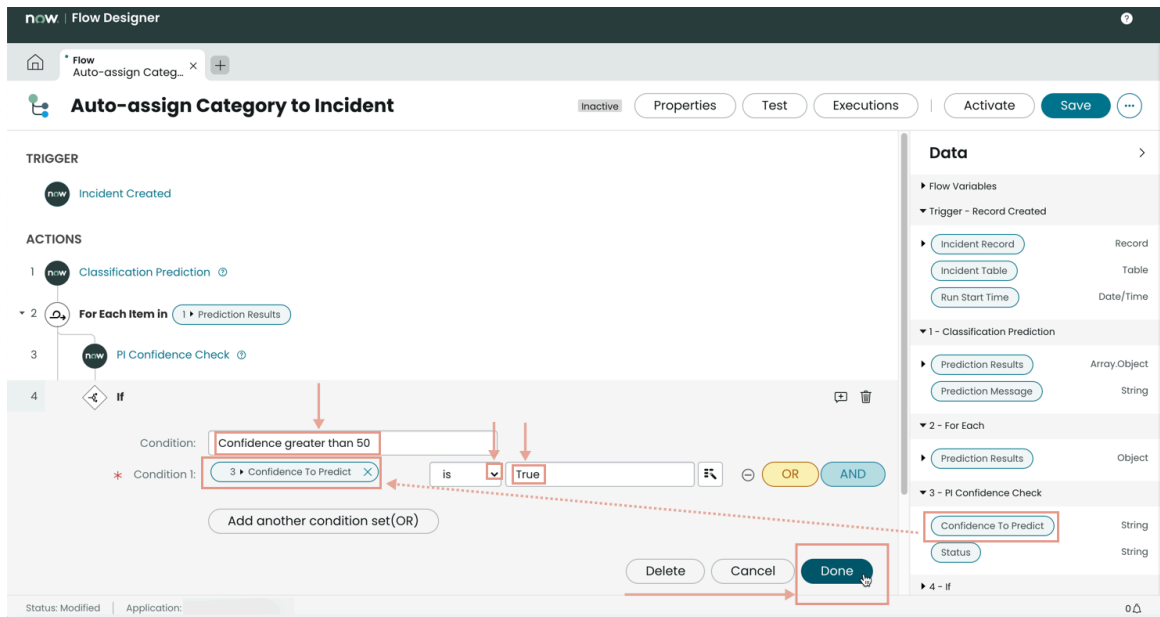


12. Select **Flow Logic > If** to add a condition to the flow.



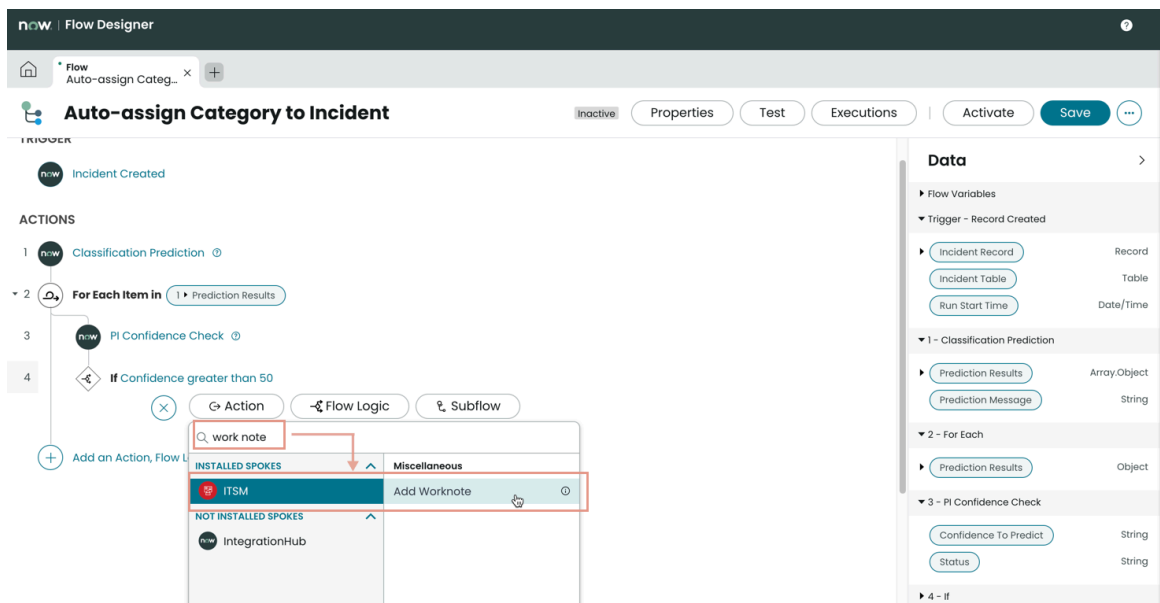
13. Configure the following fields to define the first part of the condition flow logic.

- **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**.



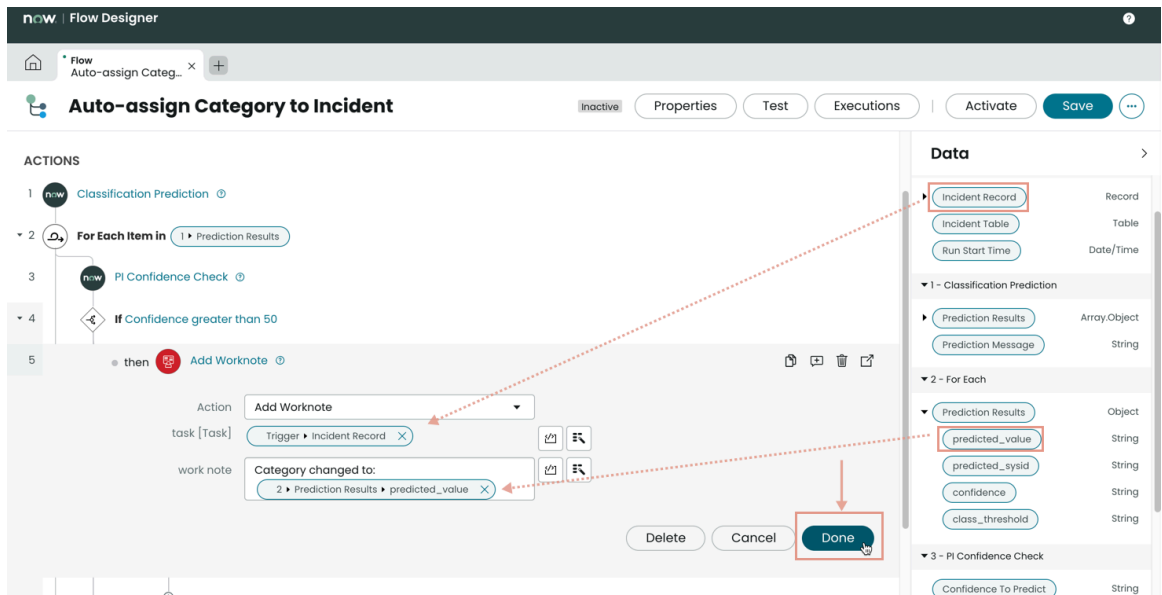
14. Select **Action** and enter **worknote** into the search field.

15. Select **ITSM > Add Worknote** to add a work note as the second part (the conclusion) of the condition.



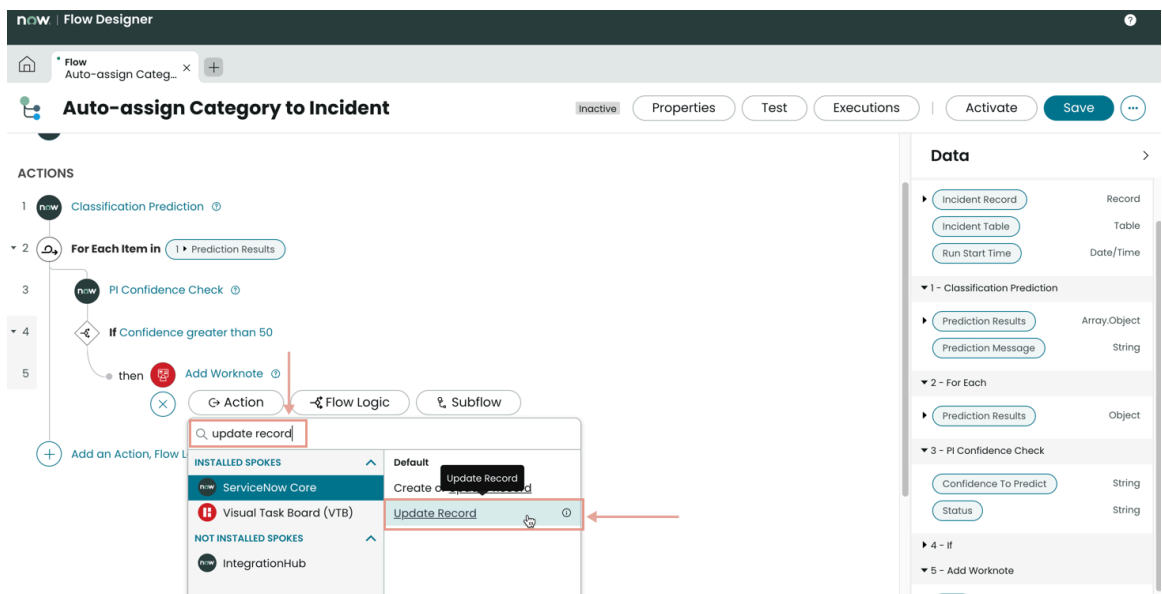
16. Configure the following fields to define the second and final part of the condition flow logic.

- **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.

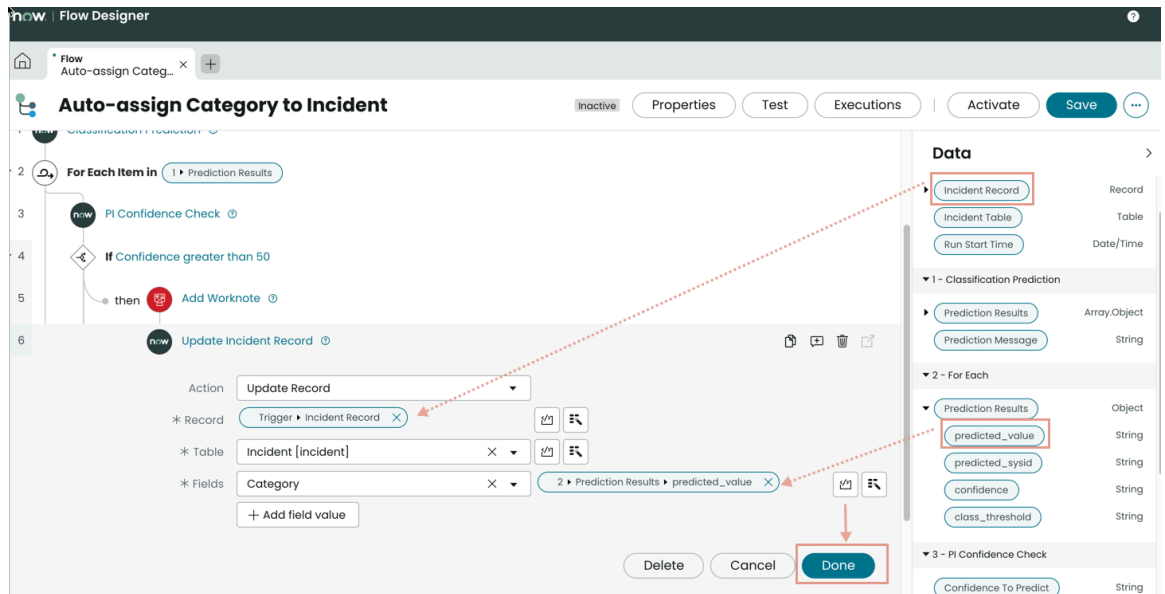
18. Select **Update Record**.



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**.

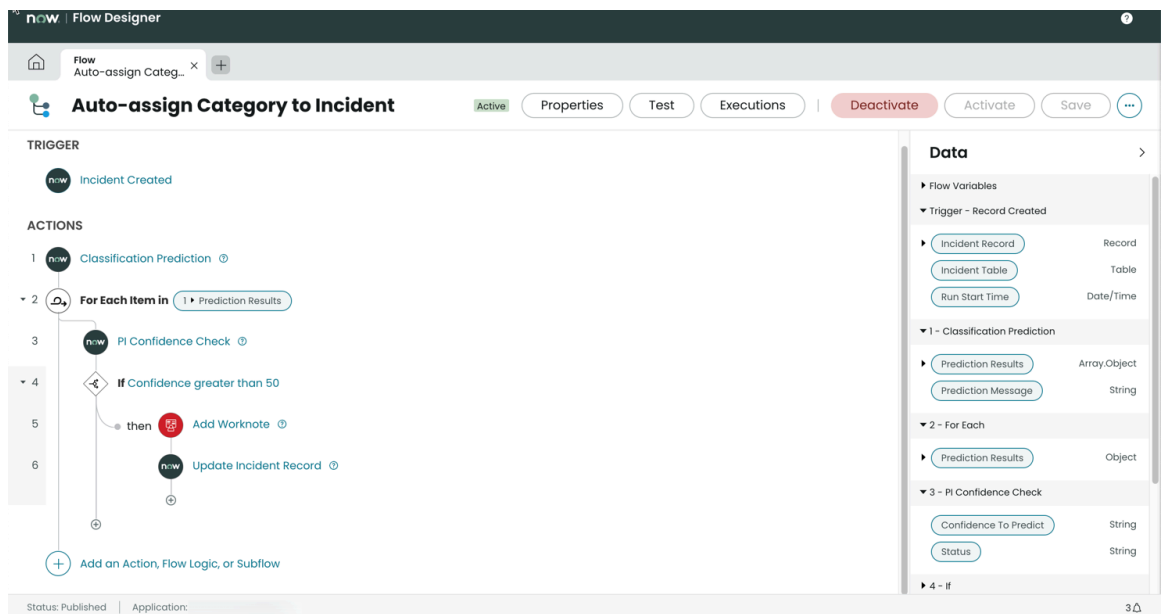


**20. Select Save.**

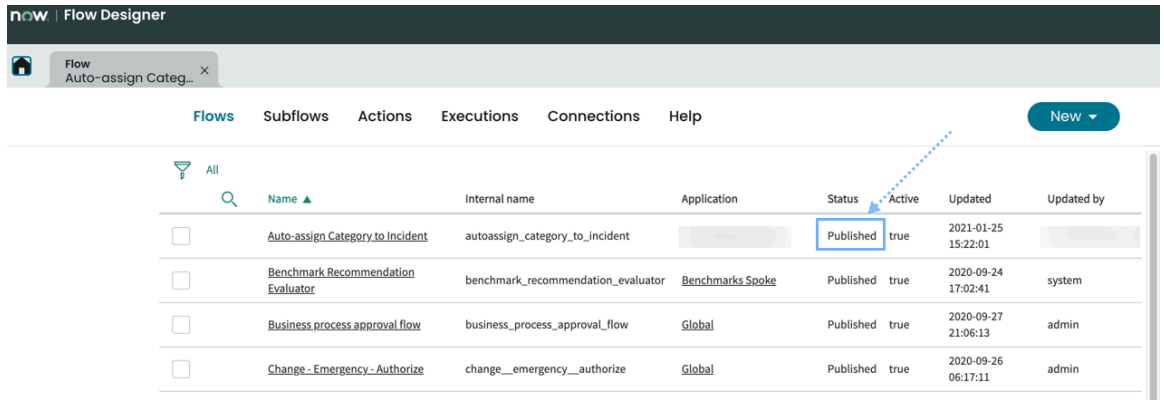
**21. Select Activate.**

**Result**

- Your Auto-assign Category to Incident flow is activated and complete.

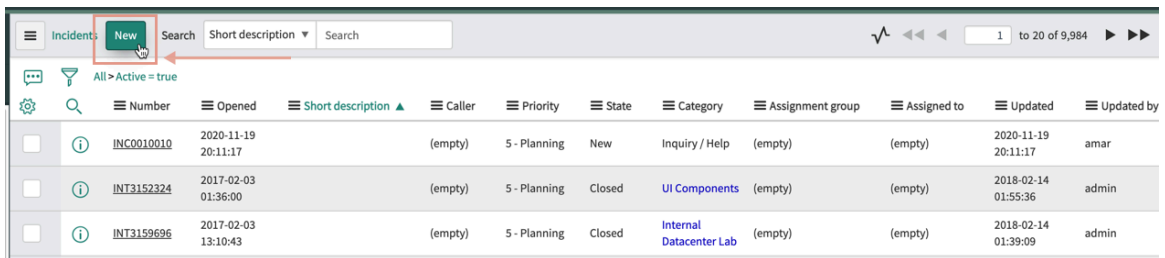


- It also appears as published in the Flows column on the Workflow Studio home screen.



**22. Navigate to Incidents.**

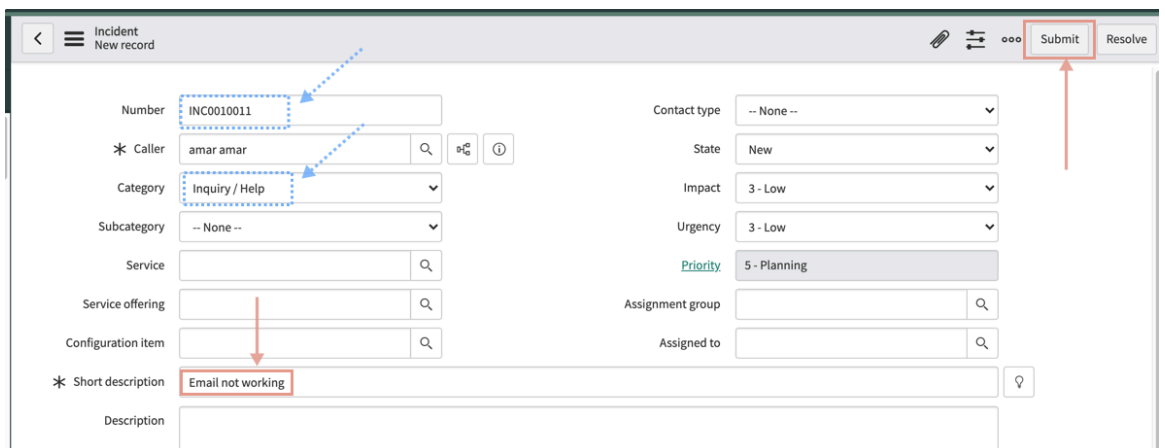
- 23. Select New to create a test incident record in the Incidents table.**  
 In this example scenario, you create record INC0010011.



- 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.**



**Result**

The system updates the incident record to show that its Category value has changed from **Inquiry / Help** to **Email**.

Incident INC0010011

Number: INC0010011

Caller: amar amar

Category: **Email**

Subcategory: -- None --

Service:

Service offering:

Configuration item:

Short description: Email not working

Description:

Contact type: -- None --

State: New

Impact: 3 - Low

Urgency: 3 - Low

Priority: 5 - Planning

Assignment group:

Assigned to:

Incident INC0010011

Short description: Email not working

Description:

Related Search Results >

Notes | Related Records | Resolution Information

Watch list: [lock] [share]

Work notes list: [lock] [share]

Work notes: [empty field]

Additional comments (Customer visible) [checkbox] Post

Activities: 2

SM: Service-now: [maint\_adm...]

Category changed to: Email

Work notes • 2021-01-25 15:41:59

## Related topics

[Machine Learning solutions for Flow Designer](#)

## Using MLSolutionFactory scriptable objects

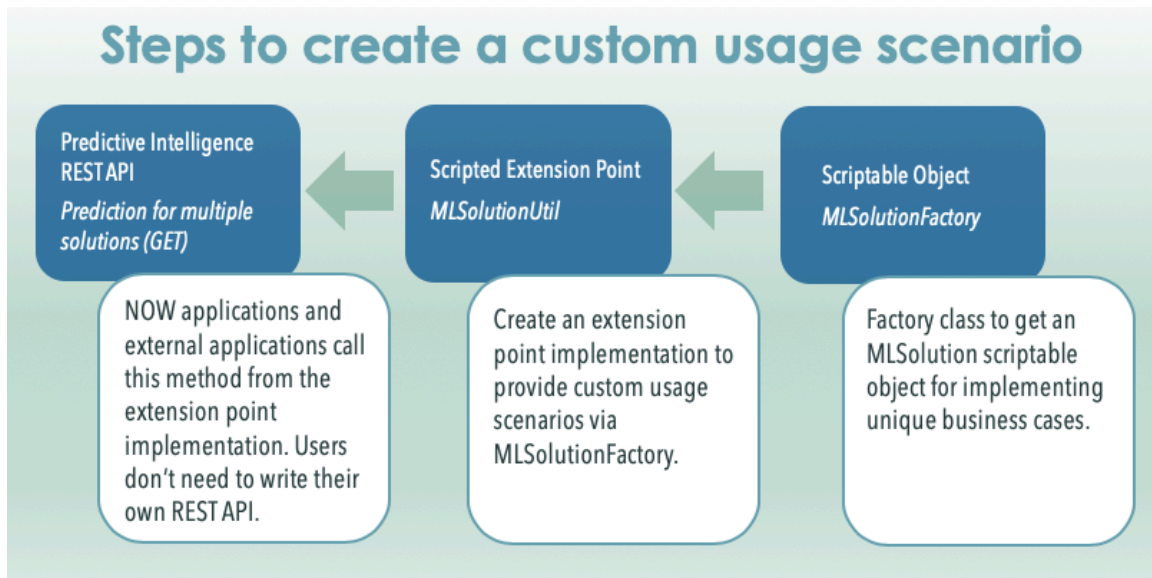
MLSolutionFactory scriptable objects enable defining ML functionality. You can use the APIs to compose data-driven functionality, such as subclustering large clusters or clusters with multiple PRBs attached.

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 ServiceNow AI 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.

## 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.

## 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 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 definition can be configured in Predictive Intelligence per domain, which ensures that the data used by the solution is domain-specific data. After a solution is trained within a domain, the solution's predictions can be applied within that domain.

### 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
Column Level Encryption	Yes	Ensure the sharedservice.worker user has the same encryption module role that's been used for encryption.
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 Predictive Intelligence properties control certain aspects of its machine-learning solution frameworks and solution training process, as well as other areas of the application.

Using the maint role, navigate to **Predictive Intelligence > Configuration** to review or edit these properties.

 **Note:** Some properties require other roles to review or edit.

**Predictive Intelligence properties**

Property	Property Name	Description
Override ml_solution_definition record's readonly state during training	glide.platform_ml.override_training_lock	<ul style="list-style-type: none"> <li>• True: Overrides the <b>readonly</b> state during training, enabling you to edit ml_solution_definition.</li> <li>• False: ml_solution_definition remains in the <b>readonly</b> state.</li> </ul>
Minimum number of records for Regression	glide.platform_ml.api.min_regression_records	<p>Sets the minimum number of records required for Regression solution training.</p> <ul style="list-style-type: none"> <li>• Type: Integer</li> <li>• Default value: 10000</li> </ul> <p><b>i Note:</b> Support for new regression solutions is deprecated with the Washington DC release. You can edit and train any existing solutions, but you can't create new ones.</p>
Maximum number of records for Regression	glide.platform_ml.api.max_regression_records	<p>Sets the maximum number of records that can be used in Regression solution training.</p> <ul style="list-style-type: none"> <li>• Type: Integer</li> <li>• Default value: 300000</li> </ul>

**Predictive Intelligence properties (continued)**

Property	Property Name	Description
		<p><b>i Note:</b> Support for new regression solutions is deprecated with the Washington DC release. You can edit and train any existing solutions, but you can't create new ones.</p>
The time (in ms) an in-training solution will wait without updates before timing out	glide.platform_ml.training_timeout	<p>Sets the time-out time in milliseconds for an in-training solution.</p> <ul style="list-style-type: none"> <li>• Type: Integer</li> <li>• Default value: 21600000 (ms)</li> </ul>
Maximum model size	glide.platform_ml.api.model_size	<p>Sets the maximum number of records you can use to train a model.</p> <ul style="list-style-type: none"> <li>• Type: Integer</li> <li>• Default value: 524288000</li> </ul>
Maximum number of records used in training	glide.platform_ml.api.csv_max_line	<p>Sets the maximum number of records that can be used in Classification solution training.</p> <ul style="list-style-type: none"> <li>• Type: Integer</li> <li>• Default value: 300000</li> </ul>
Minimum number of records used in training	glide.platform_ml.api.csv_min_line	<p>Sets the minimum number of records required for Classification solution training.</p> <ul style="list-style-type: none"> <li>• Type: Integer</li> <li>• Default value: 10000</li> </ul>
Maximum number of days worth of records on request can retrieve	glide.platform_ml.api.csv_split_days	<p>Sets the maximum number of days a single request can retrieve from your records.</p>

**Predictive Intelligence properties (continued)**

Property	Property Name	Description
		<ul style="list-style-type: none"> <li>• Type: Integer</li> <li>• Default value: 30</li> </ul>
Maximum number of records per table for word corpus	glide.platform_ml.api.max_wordcorpus_records	<p>Sets the maximum number of records per table for Word Corpus creation for Similarity and Clustering solutions.</p> <ul style="list-style-type: none"> <li>• Type: Integer</li> <li>• Default value: 300000</li> </ul> <p><b>i Note:</b> With the Washington DC release, 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.</p>
Maximum number of records for similarity window (to lookup results)	glide.platform_ml.api.max_similarity_window_records	<p>Sets the maximum number of records that the Similarity Window can retrieve to look up results.</p> <ul style="list-style-type: none"> <li>• Type: Integer</li> <li>• Default value: 100000</li> </ul>
Maximum number of records for Clustering	glide.platform_ml.api.max_clustering_records	<p>Sets the maximum number of records you can include in a cluster.</p> <ul style="list-style-type: none"> <li>• Type: Integer</li> <li>• Default value: 100000</li> </ul>

### 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.

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.

None.

### Special considerations

None.




## 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.

### Request apps on the Store






Visit the [ServiceNow Store](#) website to view all the available apps and for information about submitting requests to the store. For cumulative release notes information for all released apps, see the [ServiceNow Store version history release notes](#).

### Get started

<p style="text-align: center;"><b>Explore</b></p>  <p style="text-align: center;">Learn about DocIntel concepts and features</p>	<p style="text-align: center;"><b>Configure</b></p>  <p style="text-align: center;">Set up DocIntel to process different types of documents</p>	<p style="text-align: center;"><b>Integrate</b></p>  <p style="text-align: center;">Integrate your DocIntel application with other applications</p>
---	--	--

<p style="text-align: center;"><b>Use</b></p>  <p style="text-align: center;">Use the DocIntel application to process documents by detecting and analyzing text in documents</p>	<p style="text-align: center;"><b>Monitor</b></p>  <p style="text-align: center;">Track DocIntel usage and effectiveness</p>	<p style="text-align: center;"><b>References</b></p>  <p style="text-align: center;">Get details about the properties that are installed with DocIntel</p>
<p style="text-align: center;"><b>Now Assist in Document Intelligence</b></p>  <p style="text-align: center;">Use generative AI to analyze and extract information from documents</p>		

## 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).

### Document Intelligence overview

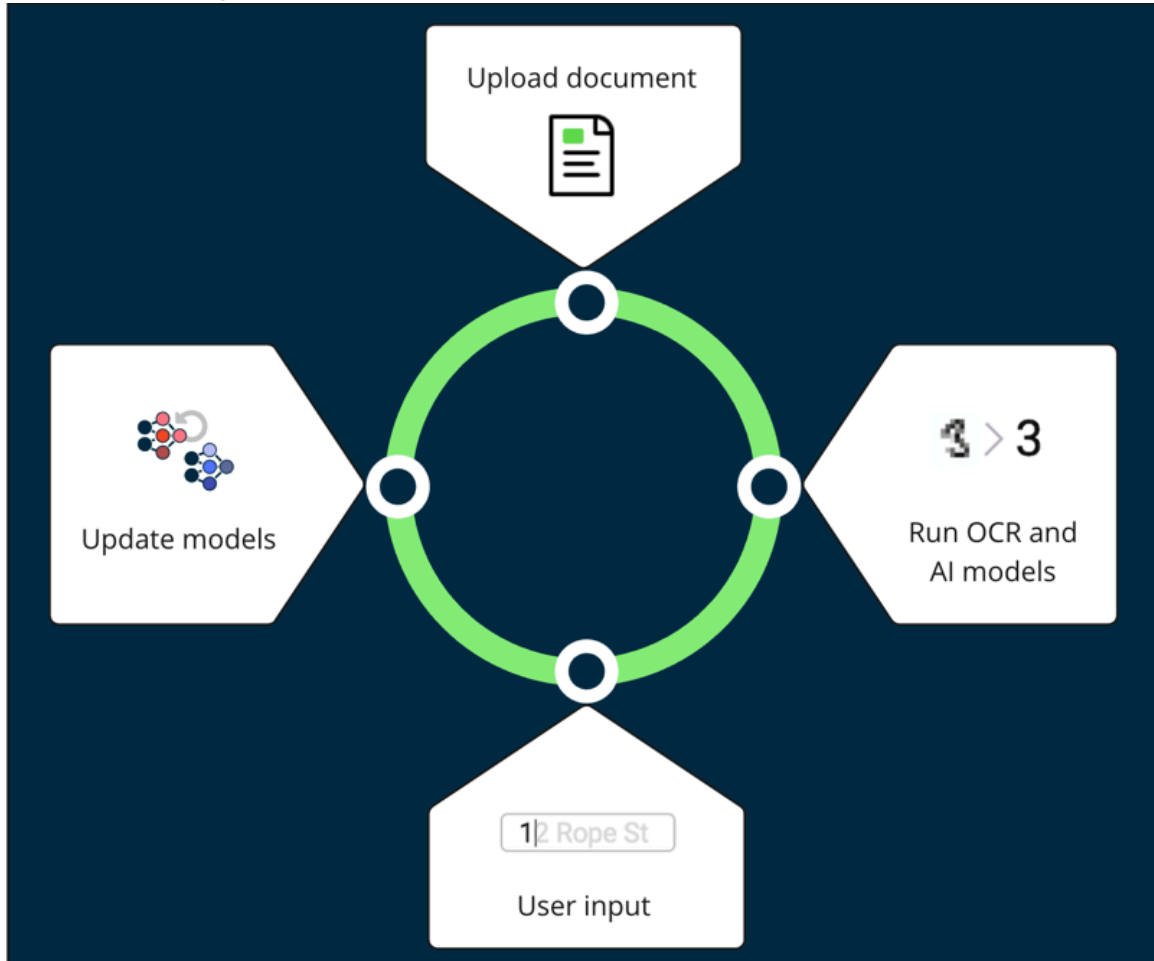
Document Intelligence uses optical character recognition (OCR) together with AI models to detect and extract information from documents with various structures and formatting. This enables you to accurately extract information and automate document processing.

## 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.

### 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.

## Document Intelligence benefits

### Benefits of Document Intelligence

We take a phased, responsible approach to increase automation **with minimal disruption to business**



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]

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.

### 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.](#)

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.

- o [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.

### Checklist

Setup task	Description
<p>Verify that the ServiceNow core applications or plugins that are required to support Document Intelligence are installed and activated.</p>	<p>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.</p> <ul style="list-style-type: none"> <li>• Predictive Intelligence (com.glide.platform_ml)</li> <li>• Platform Document Intelligence (com.glide.platform_ml_di)</li> <li>• Document Intelligence UIB Component (sn_docintel_iframe)</li> </ul> <p>For more information, see <a href="#">Install Document Intelligence</a>.</p>
<p>(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.</p>	<p>Install the <a href="#">Document Intelligence Admin</a> application from the ServiceNow® Store. Review the listing for information on dependencies, licensing or subscription requirements, and release compatibility.</p> <p>Key features of the application include:</p> <ul style="list-style-type: none"> <li>• Easily configure your document extraction solutions by defining what you want to extract from your documents.</li> <li>• Monitor the performance of your document extraction solutions and adjust the level of automation.</li> </ul>
<p>Verify that you've assigned the required ServiceNow AI Platform roles.</p>	<p>The following roles are used across the Document Intelligence features:</p> <ul style="list-style-type: none"> <li>• The administrator (admin) installs the application from the ServiceNow Store and assigns the roles.</li> <li>• sn_docintel.admin</li> </ul>

**Checklist (continued)**

Setup task	Description
	<ul style="list-style-type: none"> <li>• sn_docintel.creation_agent</li> <li>• sn_docintel.extraction_agent</li> <li>• sn_docintel.manager</li> <li>• sn_docintel.viewer</li> </ul>
Domain separation	<p>Review the <a href="#">domain separation topic</a> information if you intend to separate data, processes, and administrative tasks.</p> <p>See <a href="#">Domain separation and Document Intelligence</a>.</p>
Verify that file extensions required to support Document Intelligence are included in any glide.attachment.extensions system property customizations.	<p>The glide.attachment.extensions system property is empty by default. When customizing it, ensure that the following file extensions are included:</p> <ul style="list-style-type: none"> <li>• jpeg</li> <li>• jpg</li> <li>• png</li> <li>• pdf</li> </ul> <p>For more information, see <a href="#">Restrict attachment file extensions</a>.</p>
Copy any use cases needed to support your Document Intelligence implementation.	<p>Follow the steps in <a href="#">Duplicate a document extraction use case</a> to copy a use case along with its fields, field groups, integrations, flows, and all the related machine learning (ML) models.</p>
Import any use cases needed to support your Document Intelligence implementation.	<p>Follow the steps in <a href="#">Import a document extraction use case</a> to import a use case along with its fields, field groups, integrations, flows, and all the related machine learning (ML) models.</p>

**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.

## 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**.

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.

## 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
<p>Choose a time to schedule the upgrade</p>	<p>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.</p>
<p>Update any custom code that points directly to DocIntel database tables</p>	<p>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.</p> <p><b>Note:</b> Document extraction use cases may not be available for use until the migration is complete.</p> <p>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.</p> <p>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.</p> <p>For the list of tables, see <a href="#">Components installed with Document Intelligence</a>.</p>
<p>Add cross-scope records for integrations</p>	<p>For any integrations with Document Intelligence, add cross-scope records for the new DocIntel database tables.</p> <p><b>Note:</b> Ensure that the existing cross-scope records for the old tables are not removed.</p> <p>See the Cross-scope records section for the list of records to be added.</p> <p>For more detail on cross-scope records, see <a href="#">Cross-scope privilege record</a>.</p>

## 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	<p>Some cases where you can lose data include:</p> <ul style="list-style-type: none"> <li>• 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.</li> <li>• 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.</li> <li>• 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.</li> </ul>
Post-upgrade migration of legacy use cases imported through update sets	<p>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.</p> <p>You may want to run this script if you must migrate any imported use cases requiring migration to the platform tables (sys_di_).</p> <ol style="list-style-type: none"> <li>1. Navigate to <b>All &gt; System Definition &gt; Scheduled jobs</b>.</li> <li>2. Open the <i>DocIntel migrate remaining data</i> scheduled job.</li> </ol>

### Document Intelligence 3.0 or later post-upgrade checklist (continued)

Post-upgrade task	Description
	<p><b>3.</b> Select the <b>Active</b> check box.</p> <p><b>4.</b> Select <b>Execute Now</b>.</p> <p><b>⚠ Warning:</b> Never use the <i>DocIntel migrate remaining databatch</i> job to rerun migration after deleting all records from <i>sys_di</i> tables. You should only use it to migrate missing data from <i>di_tables</i> to <i>sys_di</i> tables.</p>

## 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
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.

### 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 <b>Auto-sync</b> to enable the exact match option. This lets the user select an icon to

Setting	Description
	view only the recommendations that match exactly what they type.
Candidate score	Select <b>Auto-sync</b> to make confidence scores visible next to each recommendation.  For more information, see <a href="#">Confidence scores</a> .
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.
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.

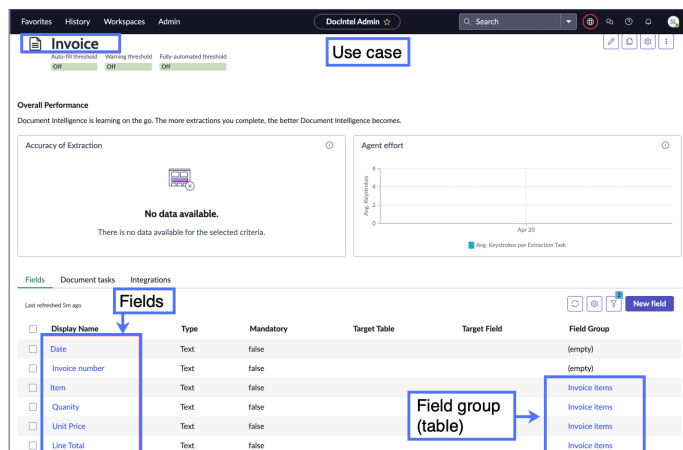
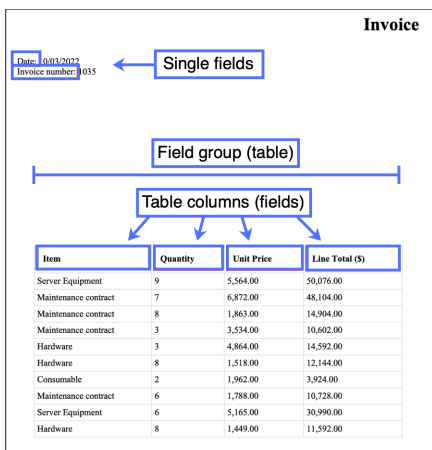
## Set up document extraction use cases

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.

## 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.

## Workflow

Set up a document extraction use case in the following steps.

### 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.

### 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](#).

## 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.

## 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**.

## 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.

## 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.

**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.

Extraction mode	Description
Recommendation	<p>DocIntel provides recommendations for the fields in the Document Intelligence workspace. Choose the recommendation or manually enter the value. All fields must be reviewed.</p> <p>Recommendations are ordered based on how confident the AI is in the prediction. As DocIntel continues processing your documents, recommendations improve over time.</p>
Auto-fill	<p>DocIntel auto-fills the fields in the Document Intelligence workspace. All fields must be reviewed.</p> <p>Auto-fill works only if the AI has enough confidence to make the prediction. You can change the confidence threshold by updating the <b>Auto-fill threshold</b> field in the use case.</p>
Fully automated	<p>DocIntel automatically extracts the data for all fields and processes the document task if the confidence scores for all required fields</p>

Extraction mode	Description
(Straight-through processing)	<p>are above the defined confidence threshold. Fields don't need to be reviewed.</p> <p>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.</p>

## 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.

- 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.

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.

### 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.

#### Before you begin

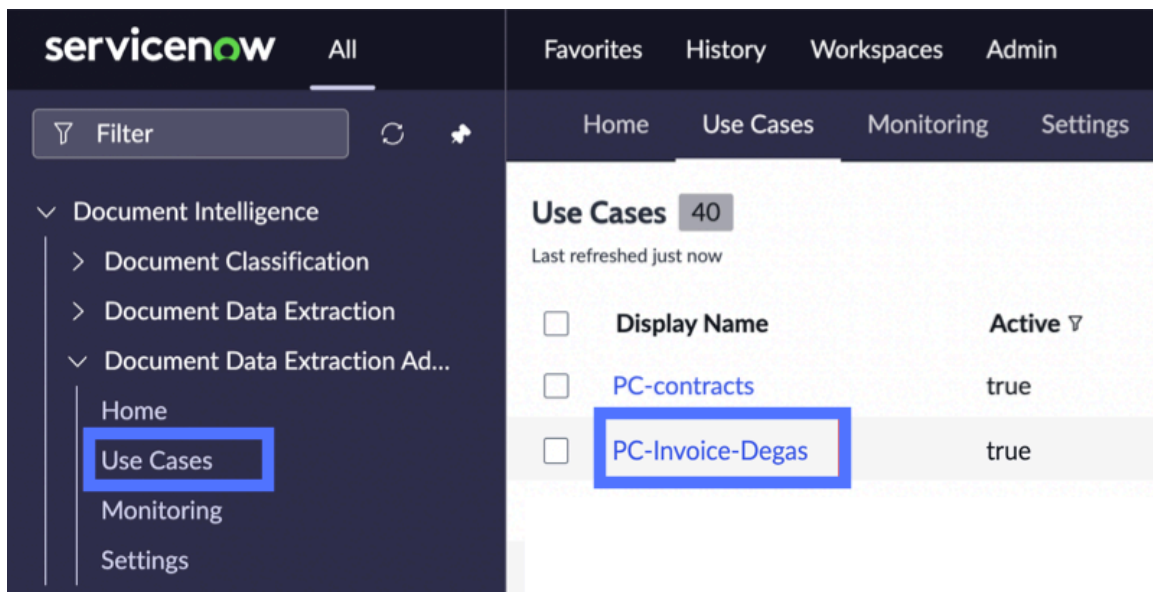
Role required: sn\_docintel.manager

#### 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.

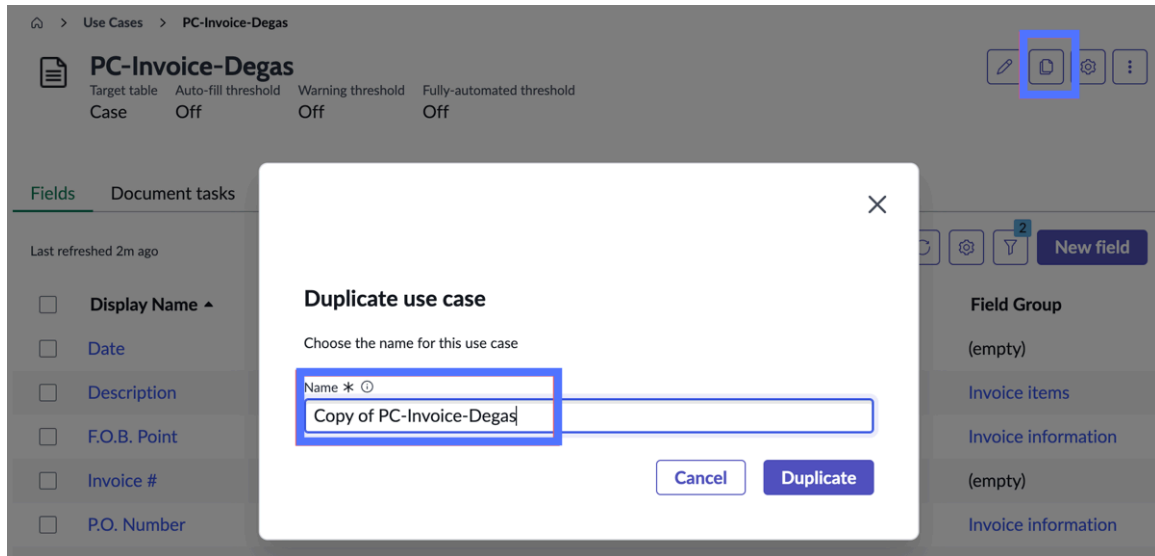
### Procedure

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.



**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.

**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](#).

**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.

#### 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**.
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.

#### 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.

## 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.

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.

## 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.
Autofill Threshold	<p>DocIntel only auto-fills the classes (fields) if the confidence score of the top recommendation is at or above the percentage you define.</p> <p>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.</p> <p>This field is available only if the Auto-fill mode is enabled.</p>
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	<p>DocIntel shows a warning for empty fields and auto-filled fields with a confidence score at or below the percentage that you define.</p> <p>This field is available only if the Auto-fill mode is enabled.</p>

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.

**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.
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

The screenshot displays the configuration page for a use case named 'PC-DC-Invoice'. The configuration includes several threshold and processing options. Below the configuration, the 'Fields (3)' tab is active, showing a table of fields associated with this use case.

Display Name	Name	Type	Order	Active
Invoice	invoice	Text	0	true
Survey	survey	Text	0	true
Receipt	receipt	Text	0	true

### Train a use case

Train the document classification use case with user input from completed document tasks to improve Document Intelligence recommendations over time.

### 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

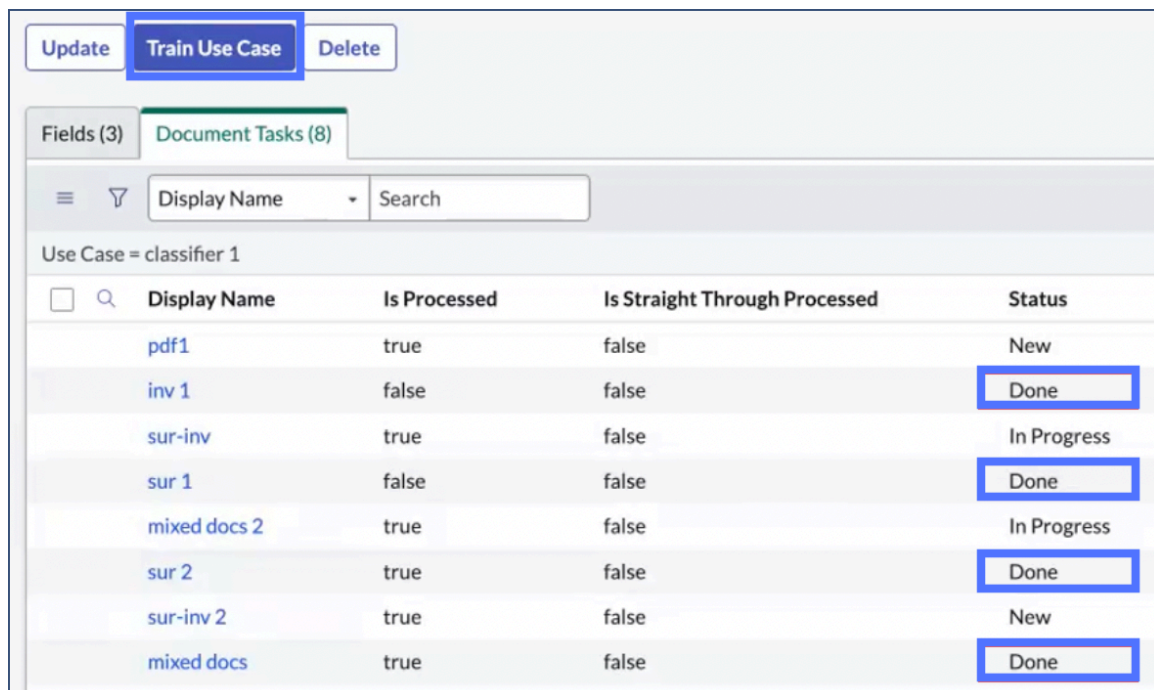
Document classification use cases don't begin with pre-trained AI models, so it's important to train the models with user input from completed document tasks.

**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.

## Manage document classification use cases

Manage your Document Intelligence use cases to efficiently support your document classification requirements.

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.

#### 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**.

#### 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.

#### 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.

Field	Description
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.
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 <a href="#">Domain separation and Document Intelligence</a> .

**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.

## 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.

### 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**.

Home Use Cases Settings

Invoice Processing with Tables - TOI

Target Table: Invoice Extraction mode: Autofill

**Overall Performance**  
Document Intelligence is learning on the go. The more extractions you complete, the better Document Intelligence becomes.

**Accuracy of Extraction**  
Line chart showing accuracy percentage from Mar 9 to May 18. Accuracy starts at approximately 45% and gradually declines to near 0%.

**Agent effort**  
Line chart showing average keystrokes from Mar 9 to May 18. The effort starts at approximately 30, drops to near 0 by mid-April, and then rises to about 15 by late May.

Fields Document tasks **Integrations**

Last refreshed 1m ago. New Integration

Name	Target Table	Flow	Conditions
1	sn_shop_invoice	DocIntel Task Processing Flow - Invoice Processing with Tables - TOI - 1	
2	sn_shop_invoice	DocIntel Extract Values Flow - Invoice Processing with Tables - TOI - 2	

Showing 1-2 of 2 | 20 rows per page

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. <b>Note:</b> 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 <b>Process Task</b> or <b>Extract Values</b> . Document classification use cases also have a <b>Map To Document Data Extraction Use Case</b> option.  The <b>Process Task</b> type creates an integration point to automatically create and process DocIntel document tasks based on specific triggers happening in the target table.  The <b>Extract Values</b> type creates an integration point to automatically propagate the extracted values to the target

Field for data extraction	Field for document classification	Description
		<p>table when extraction has been completed in DocIntel (that is, when the document task status changes to Done).</p> <p>The <b>Map To Document Data Extraction Use Case</b> type creates an integration point that allows a processed document classification value to automatically create a new data extraction task.</p>
Conditions		<p>Filters used to select certain fields to use as specific triggers for the integration.</p> <p>Process Task only.</p>
Create Flow		<p>Select this option to create a flow for this integration in Workflow Studio.</p> <p>This option should be selected, unless you're planning to write your own custom script to set up the integration.</p>

**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.

### New Process Task integration

#### Create Integration

Integrate DocIntel with other applications. [Learn more](#)

Name your Integration \*

Use case \*

Where do you want to take the documents from and store the extracted data? \*

What type of integrations you want to set ? \* ⓘ

Create Flow

### New Extract Values Integration

✕

#### Create Integration

Integrate DocIntel with other applications. [Learn more](#)

Name your Integration \*

Use case \*

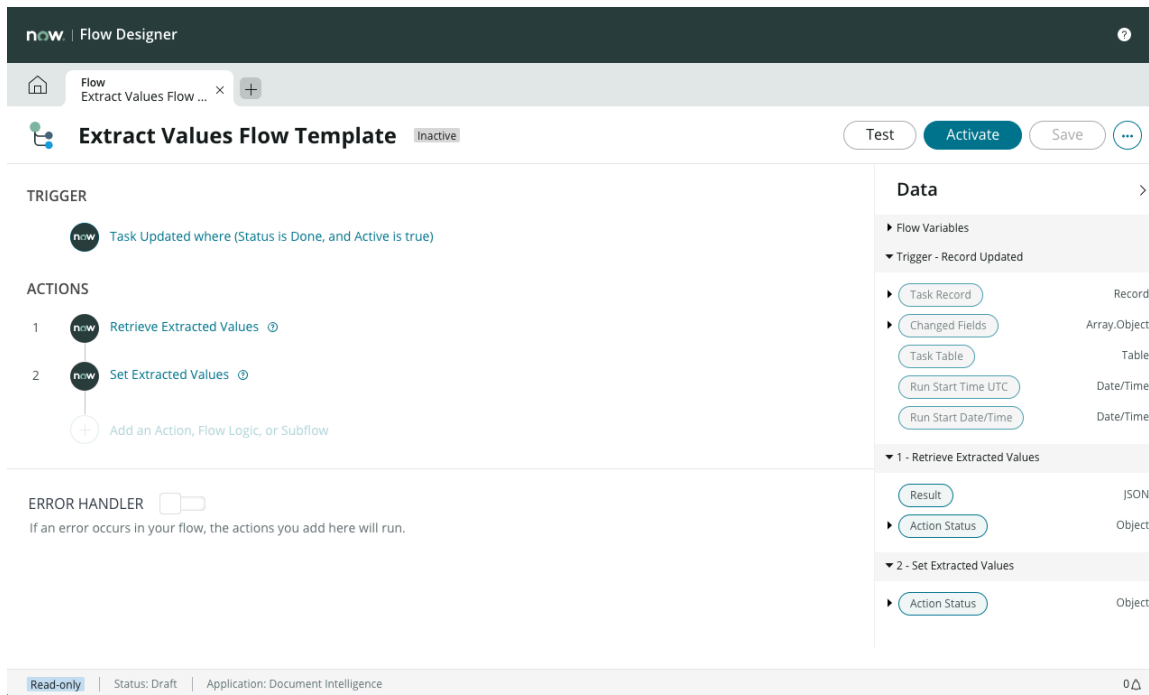
Where do you want to take the documents from and store the extracted data? \*

What type of integrations you want to set ? \* ⓘ

Create Flow

### What to do next

If you selected to create a flow, finish the activation in Workflow Studio.



For more information, see [Building flows in Flow Designer](#).

## 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.

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](#) .

## Integrate with Financial Services Operations

Document Intelligence provides document extraction capabilities to Financial Services Operations (FSO).

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).

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.

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.

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.

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.

### 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](#).

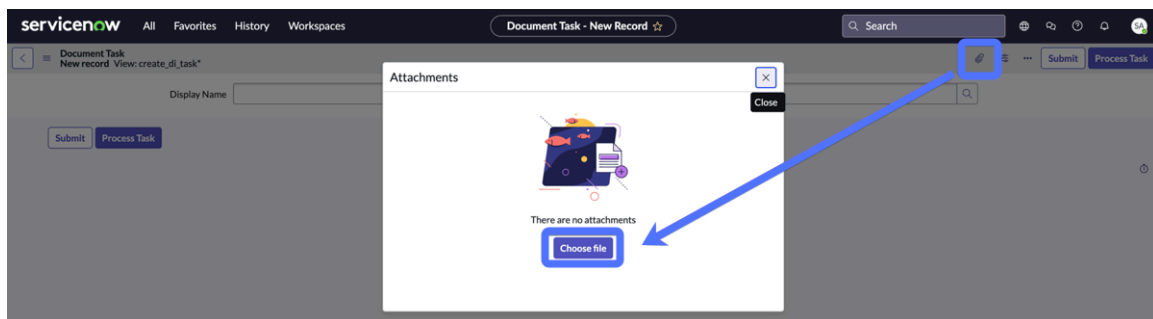
### Procedure

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.



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.

### 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. The Is Processed field for the PO-1 task is true, which 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**.

#### Document tasks list

Display Name	Active	Status	Processing start date	Error Message	Is Processed	Is Trained	Locked
process task - pre-built demo 1	true	Failed	(empty)	No documents attached to task.	false	false	false
test	true	Setup	(empty)		false	false	false
PurchaseOrderHAC	true	Setup	(empty)		false	false	false
PO-410087	true	Failed	(empty)	No documents attached to task.	false	false	false
PO-1	true	In Progress	2022-05-02 19:43:46		true	true	true
PO-2	true	New	2022-05-02 19:48:37		true	true	true
PO-3	true	New	2022-05-02 20:02:35		true	true	true

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

Document Task  
Degas survey 3923-B View: edit\_sys\_di\_task

Manage Attachments (1): Degas Survey 3923-B.pdf [rename][download]

Display Name: Degas survey 3923-B

\* Use Case: Degas survey

Source Record: (empty)

Is Processed:

Is Straight Through Processed:

\* Status: In Progress

Warning Message:

Error Message:

Buttons: Update, Show In DocIntel, Delete

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.

## 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:

- 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:

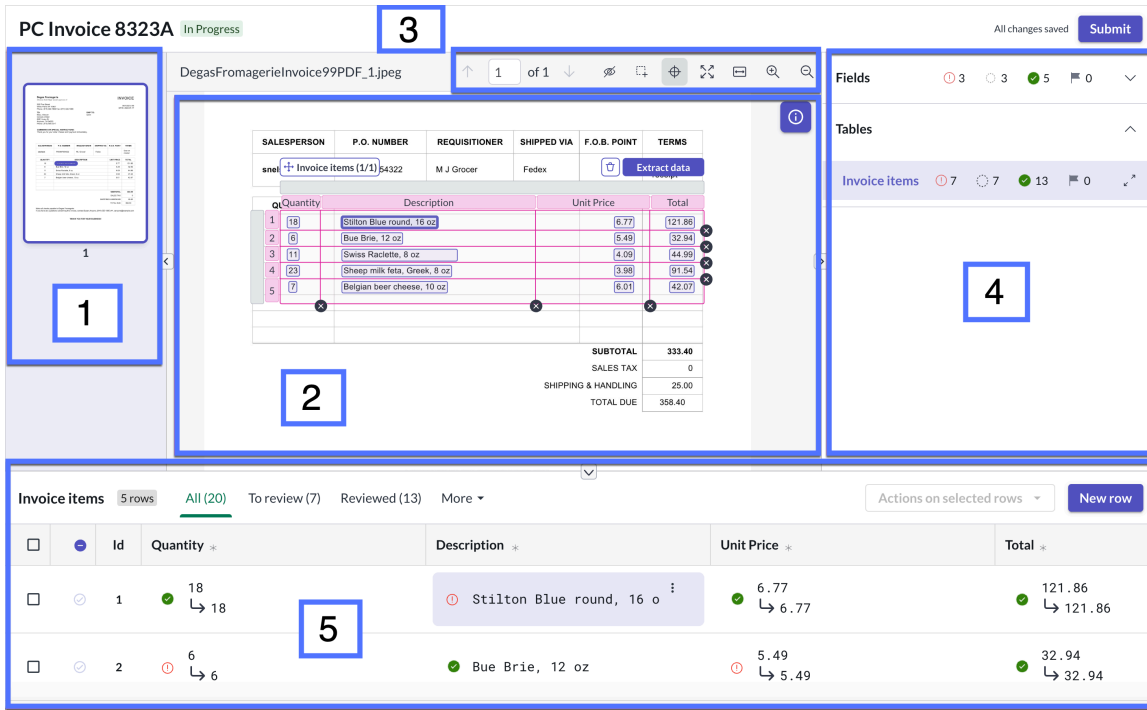
- [Create a document task](#) for a document data extraction use case, upload a document, and process the task.
- After DocIntel has processed the task, you can begin using the workspace. See [Complete a document task](#).

## 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

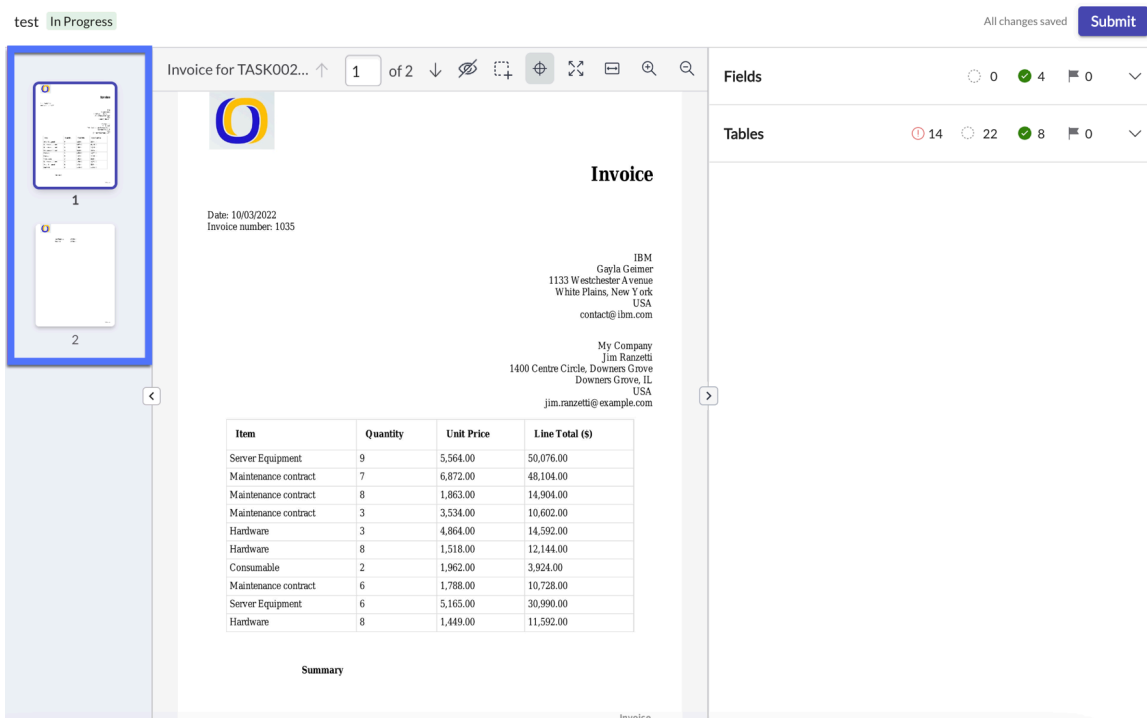
**i Note:** In this view, the document fields panel is expanded so that the fields are visible.



## 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.



## 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.

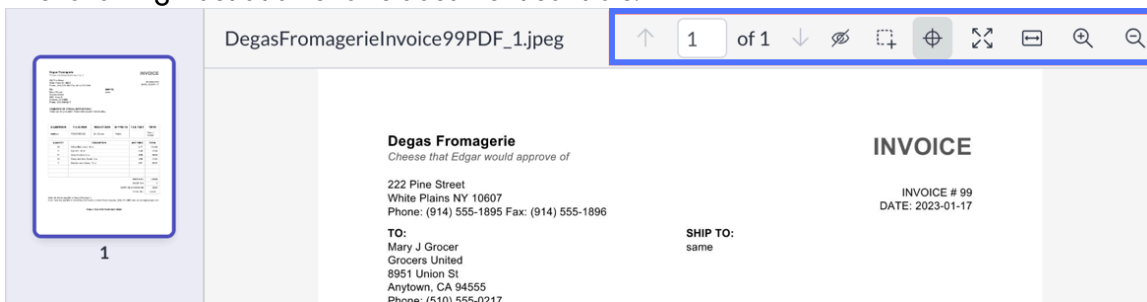


You can also extract information directly from the document image using the draw tool. For more details, see [Extract fields using the draw tool](#).

## 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](#).

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.



## 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.

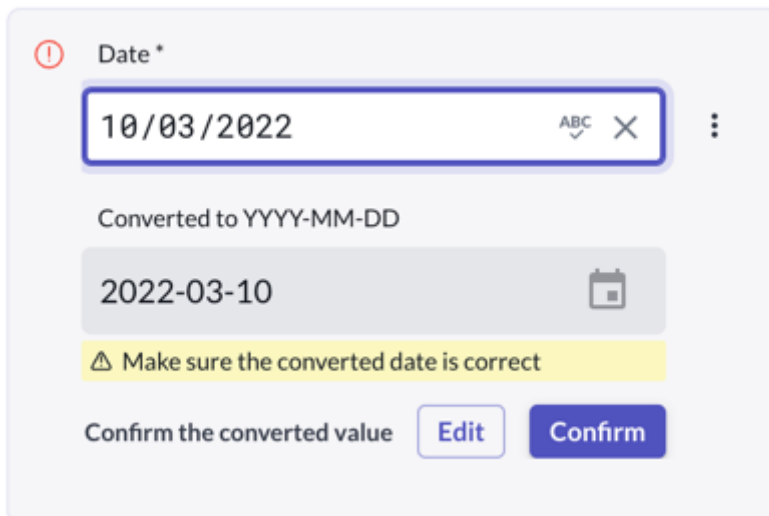
### 4. Optional: If needed, filter or flag the recommendations

- 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.

(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.



**Date \***

10/03/2022

Converted to YYYY-MM-DD

2022-03-10

**⚠ Make sure the converted date is correct**

Confirm 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.

#### **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.

- 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.

### Extract fields using the draw tool

Use the draw tool to extract information directly from text and tables on the document image.

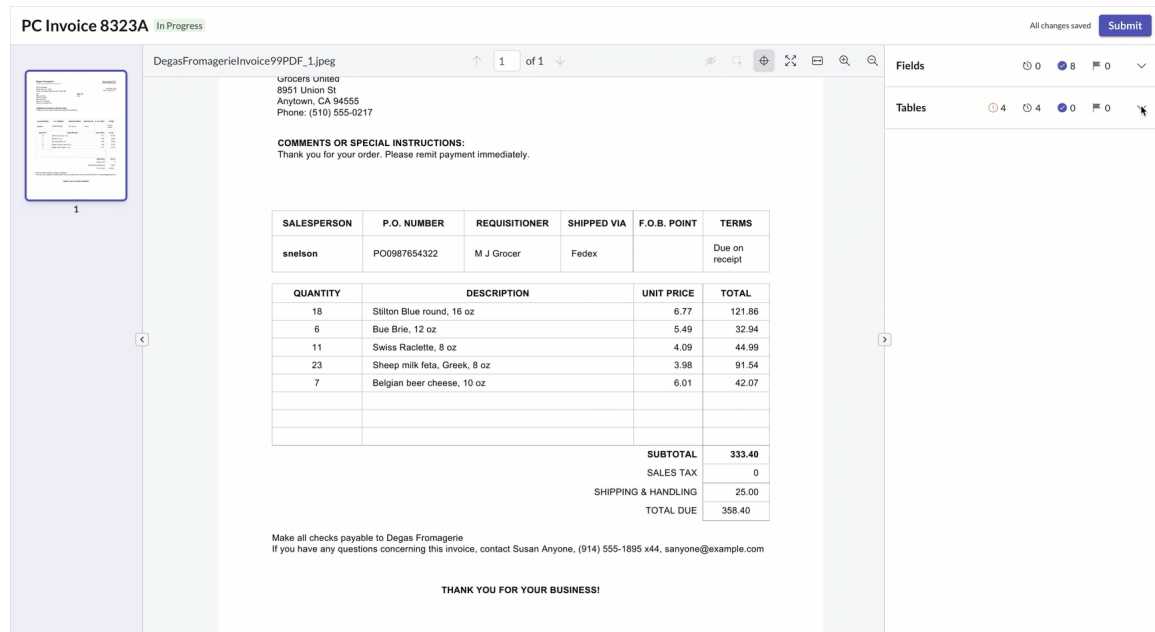
### Overview of the draw tool

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.

### Procedure

1. In the document fields panel, expand **Fields**.
2. Select a field.

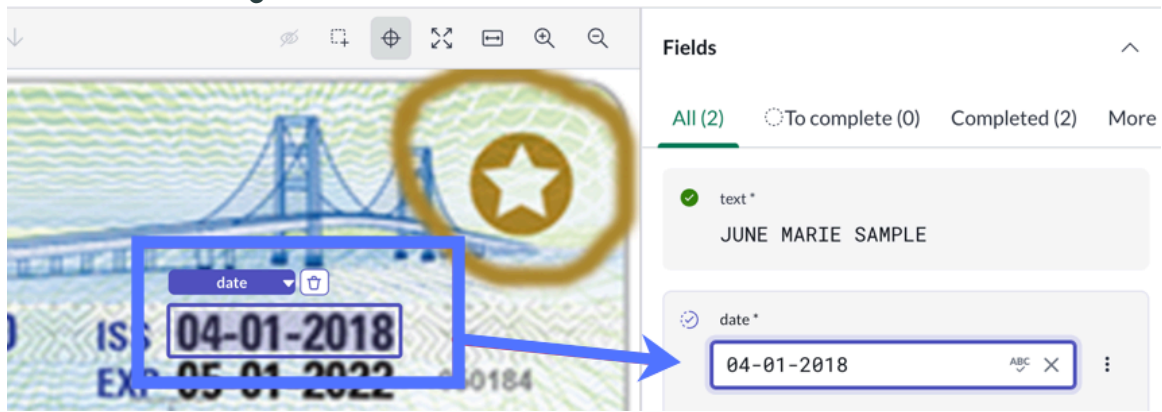
The Draw Tool icon (  ) is enabled in the document controls toolbar.

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
<b>Resize the box</b>	(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.
<b>Remove the box</b>	(Optional) Select the Remove the box icon (🗑️).  The box is deleted and the extracted value is removed from the field.
<b>Extract to a different field</b>	(Optional) Select the field name next to the box and choose another option from the list.

**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.

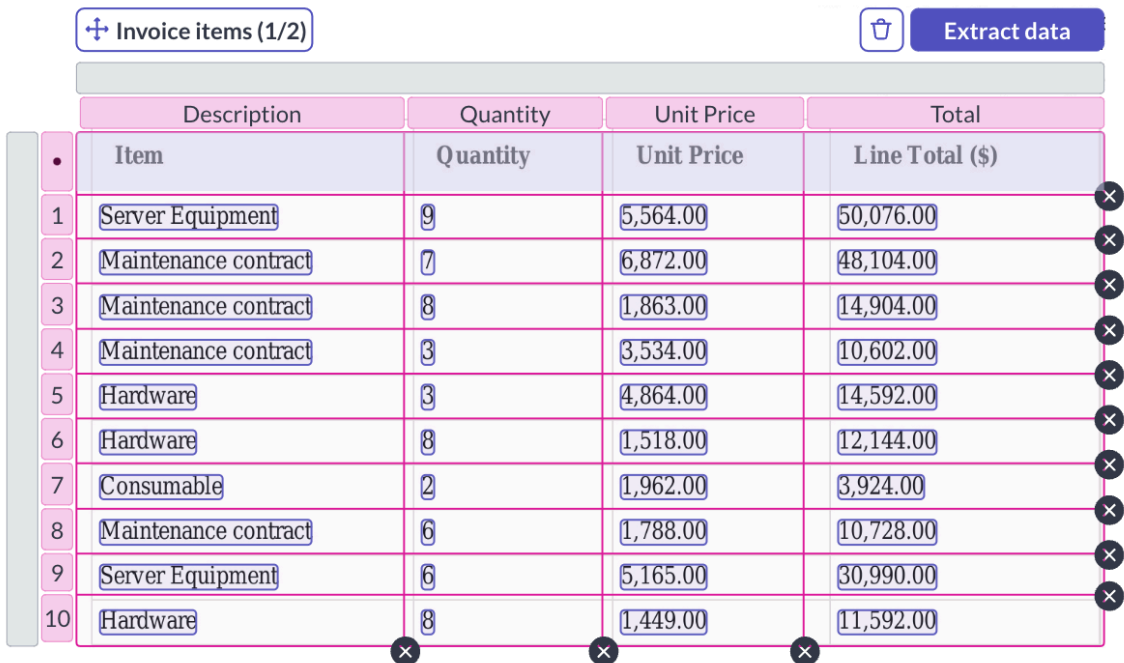
 **Tip:** Collapse the Table panel using the Collapse table panel icon (  ), if needed.

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.




	Description	Quantity	Unit Price	Total
•	Item	Quantity	Unit Price	Line Total (\$)
1	Server Equipment	9	5,564.00	50,076.00
2	Maintenance contract	7	6,872.00	48,104.00
3	Maintenance contract	8	1,863.00	14,904.00
4	Maintenance contract	3	3,534.00	10,602.00
5	Hardware	3	4,864.00	14,592.00
6	Hardware	8	1,518.00	12,144.00
7	Consumable	2	1,962.00	3,924.00
8	Maintenance contract	6	1,788.00	10,728.00
9	Server Equipment	6	5,165.00	30,990.00
10	Hardware	8	1,449.00	11,592.00

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.**

- o Select **Extract data**.

This button displays if no data has been extracted yet.

- o Select **Update fields**.

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.

**On-premise Software**  
 Except for perpetual licenses, these licenses to Products and Services and any applicable Support terminate on the identified License Term End Date. The License Term commences from the date of delivery of the Products and Services. The dates below represent best estimates of Start Date of the License Term. "Support" services for the Products and Services are described here:

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

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
<b>Move the grid</b>	Drag the Move icon (+) to the desired position.
<b>Change a column header</b>	Select the column header and select the column where the data will be extracted to in the table fields.
<b>Make the first row a header row</b>	Select the row header and check the <b>This is a header row</b> check box.

Option	Description
<b>Move a column or row border</b>	Drag the border to the desired position.
<b>Add a new column or row</b>	Go to the outside border of the grid and select the Add a line icon (⊕). Adding a line splits the affected cells so they can be extracted to separate fields.
<b>Remove a column or row border</b>	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.
<b>Remove the grid</b>	Select the Remove the grid icon (🗑️). This action only removes the grid, not the extracted data.
<b>Ignore a table column or row</b>	<ol style="list-style-type: none"> <li>Select the column or row heading.</li> <li>Select the <b>Do not extract this column/row</b> check box.</li> <li>Select <b>Save</b>.</li> </ol>

### Review a field on the document image

Use the draw tool to review auto-filled fields directly from the document image panel.


### 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

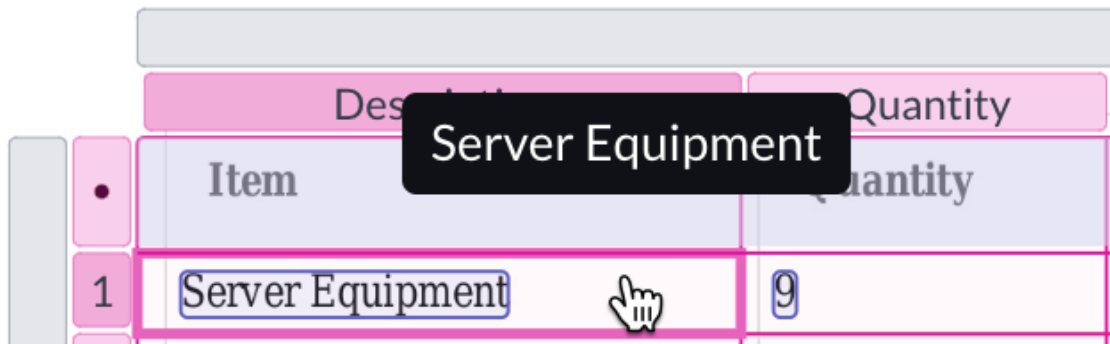
### Procedure

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

**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.

**PC Invoice 8323A** In Progress

DegasFromagerieInvoice99PDF\_1.jpeg ↑ 1 of 1 ↓

SALESPERSON	P.O. NUMBER	REQUISITIONER	SHIPPED VIA	F.O.B. POINT	TERMS
snel	+ Invoice items (1/1) 4322	M J Grocer	Fedex		<a href="#">Extract data</a>

Quantity	Description	Unit Price	Total
18	Stilton Blue round, 16 oz	6.77	121.86
6			32.94
11			44.99
23			91.54
7			42.07

Row 1 - Description

Stilton Blue round, 16 oz

- ✓ Stil ... 6 oz 50% p1
- Stilto ... round 50% p1
- Stilton 50% p1
- Stilton Blue 50% p1

Make all checks payable to ☐  
If you have any questions con

TOTAL DUE 358.40

x44, sanyone@example.com

---

**Invoice items** 5 rows All (20) To review (7) Reviewed (13) More ▾

<input type="checkbox"/>	<input type="checkbox"/>	Id	Quantity *	Description *
<input type="checkbox"/>	<input checked="" type="checkbox"/>	1	18 ↳ 18	Stilton Blue round, 16 o
<input type="checkbox"/>	<input checked="" type="checkbox"/>	2	6 ↳ 6	Bue Brie, 12 oz

**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.

**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:

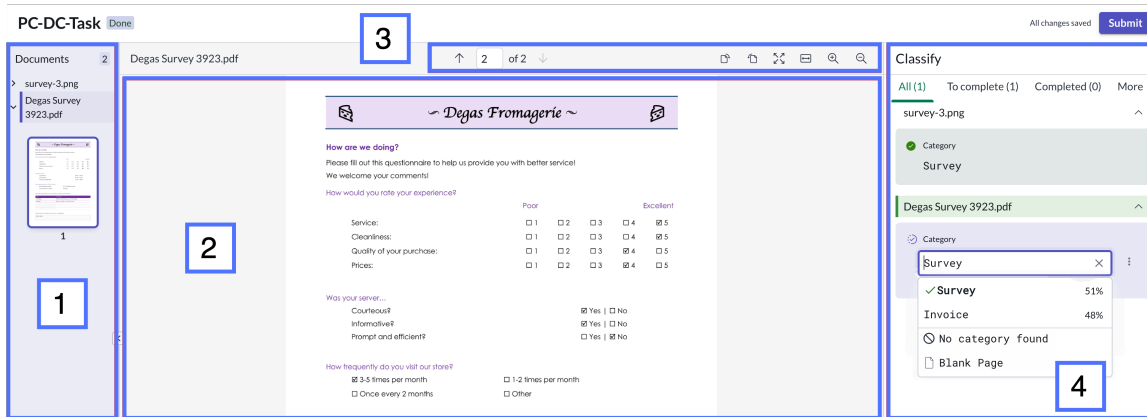
- [Create a document task](#) for a document classification use case, upload a document, and process the task.
- After DocIntel has processed the task, you can begin using the workspace. See [Complete a document task](#).

**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.

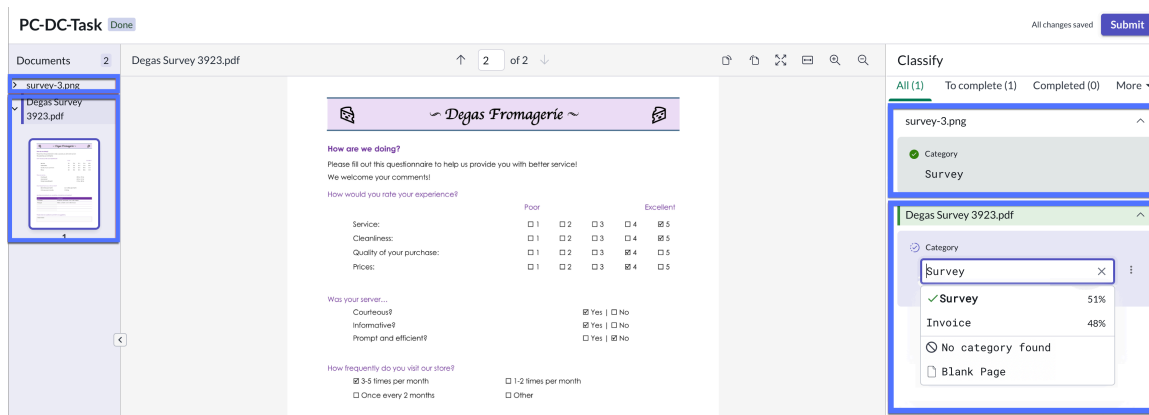
## Document controls

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.



## 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.

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.

### View reports on the Document Intelligence Admin home page

Monitor document extraction performance in the Admin experience.

#### Before you begin

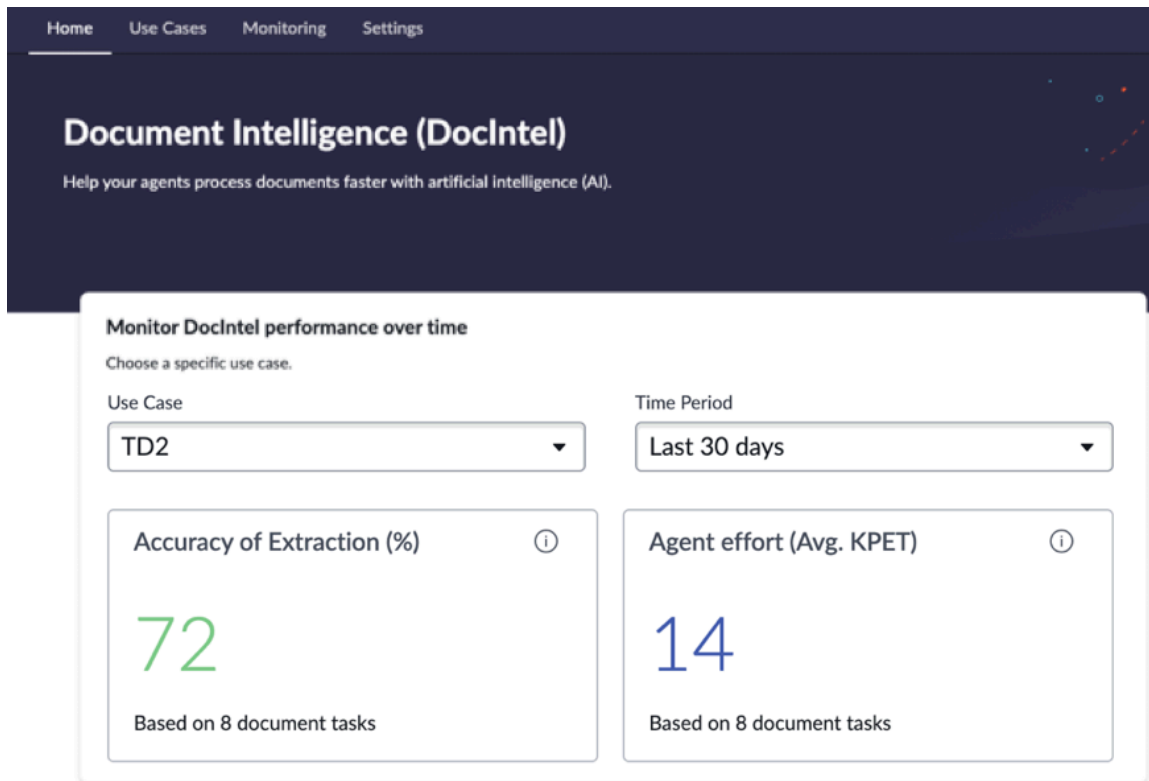
- Ensure that the Document Intelligence application (sn\_docintel) and Document Intelligence Admin (com.snc.docintel\_admin) ServiceNow<sup>®</sup> Store application is installed and active. For more information, see [Install Document Intelligence](#).
- Have an active use case with multiple completed 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.



### Procedure

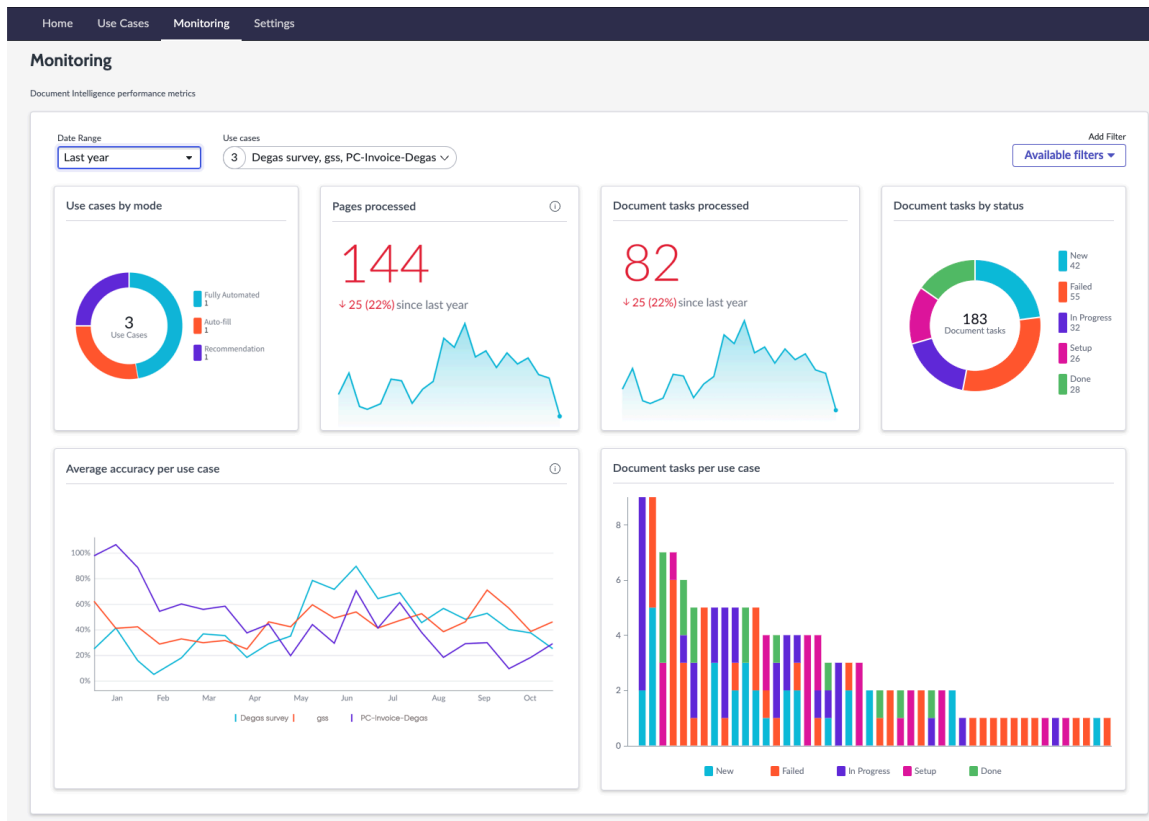
1. Navigate to **All > Document Intelligence > Document Data Extraction Administration > Home**.
2. Expand the **Use Case** list and select your use case.
3. Expand the **Time Period** list and select a date range.
4. Review the results displayed in the **Accuracy of Extraction** and the **Agent effort** widgets.

### Document Intelligence monitoring dashboard

Monitor the overall performance of Document Intelligence over time in the Document Intelligence monitoring dashboard.

### Overview of the Document Intelligence monitoring dashboard

The Document Intelligence monitoring dashboard provides a high-level overview of your Document Intelligence usage and value.



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.

## 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

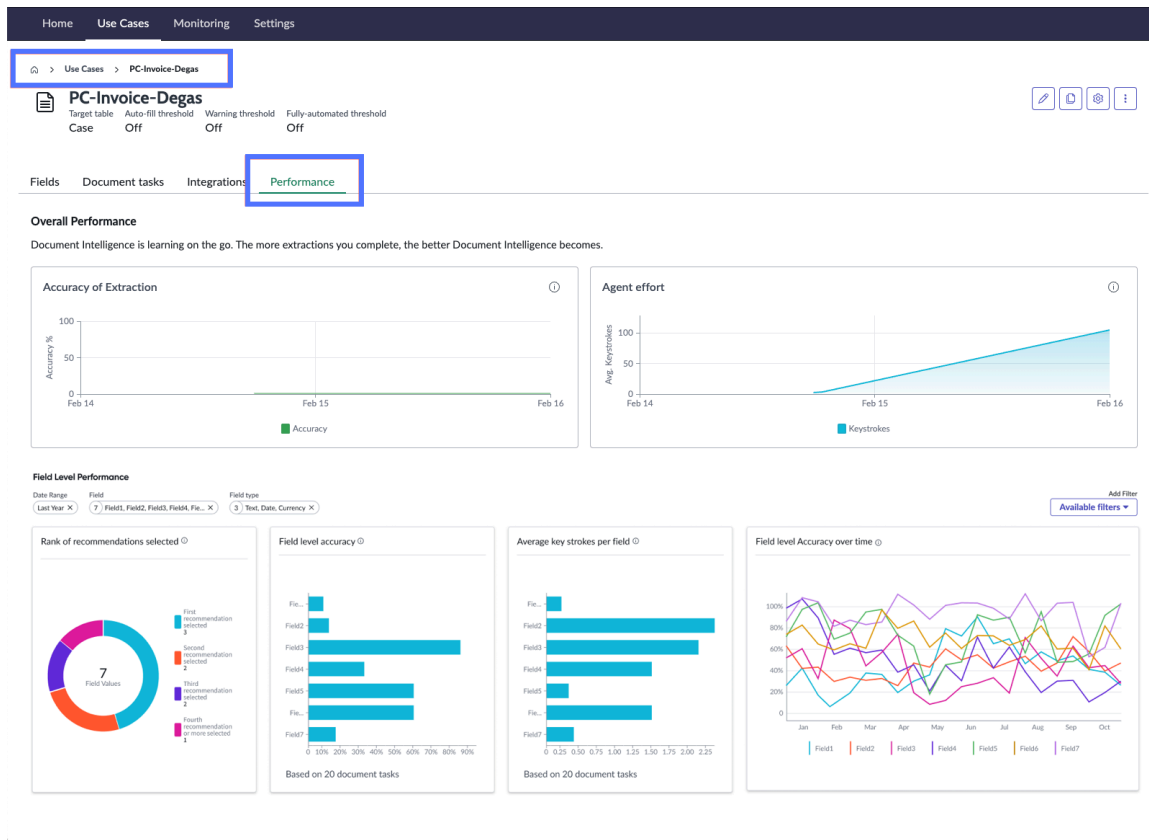
Name	Description
Use cases by mode	The number of document extraction use cases in this instance, grouped by extraction mode.
Pages processed	The number of pages processed for data extraction using DocIntel in this instance over the selected date range.
Document tasks processed	The number of DocIntel 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.

#### Overview of the use case performance dashboard

The use case performance dashboard provides quality metrics for the use case.



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.
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.

### Components installed with Document Intelligence

Several types of components are installed with activation of the Document Intelligence plugin, including tables and user roles.

### 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.	<ul style="list-style-type: none"> <li>platform_ml_di.admin</li> <li>action_designer</li> <li>flow_designer</li> <li>sn_ace.ace_user</li> <li>canvas_user</li> <li>usage_admin</li> </ul>
DocIntel Viewer [sn_docintel.viewer]	Has view-only access on Document Intelligence document tasks that they are authorized to view.	<ul style="list-style-type: none"> <li>snc_read_only</li> <li>platform_ml_di.viewer</li> </ul>
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.	<ul style="list-style-type: none"> <li>platform_ml_di.manager</li> <li>action_designer</li> <li>flow_designer</li> <li>sn_ace.ace_user</li> <li>canvas_user</li> <li>usage_admin</li> </ul>

### 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.
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	[Internal table] Contains a log of the metrics daily aggregation jobs that happened in the instance.

Table	Description
[sys_di_metrics_job_log]	
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.

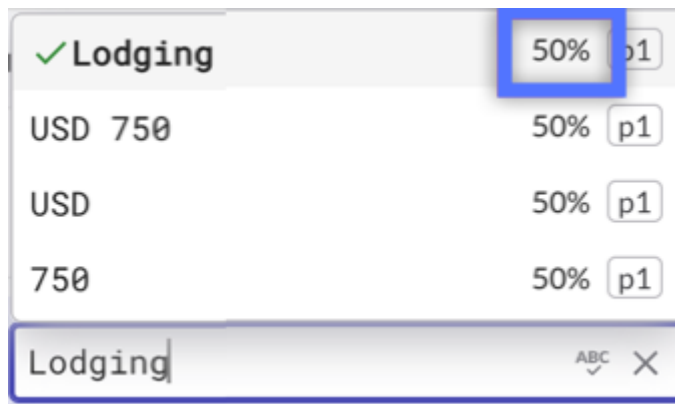
### Confidence scores

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.

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.

DocIntel uses the following extraction modes.

Extraction mode	Description
Recommendation	<p>DocIntel provides recommendations for the fields in the Document Intelligence workspace. Choose the recommendation or manually enter the value. All fields must be reviewed.</p> <p>Recommendations are ordered based on how confident the AI is in the prediction. As DocIntel continues processing your documents, recommendations improve over time.</p>
Auto-fill	<p>DocIntel auto-fills the fields in the Document Intelligence workspace. All fields must be reviewed.</p> <p>Auto-fill works only if the AI has enough confidence to make the prediction. You can change the confidence threshold by updating the <b>Auto-fill threshold</b> field in the use case.</p>
Fully automated	<p>DocIntel automatically extracts the data for all fields and processes the document task if</p>

Extraction mode	Description
(Straight-through processing)	<p>the confidence scores for all required fields are above the defined confidence threshold. Fields don't need to be reviewed.</p> <p>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.</p>

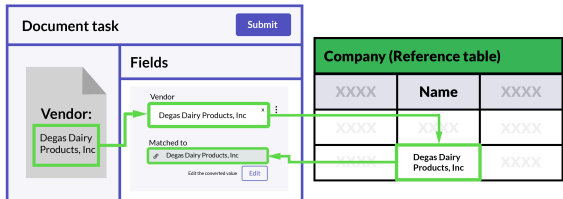
## Data normalization

Certain types of data extracted from documents are converted into a standard format so that they appear the same across all fields.

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.

## Field types

The following field types are converted to support data normalization:

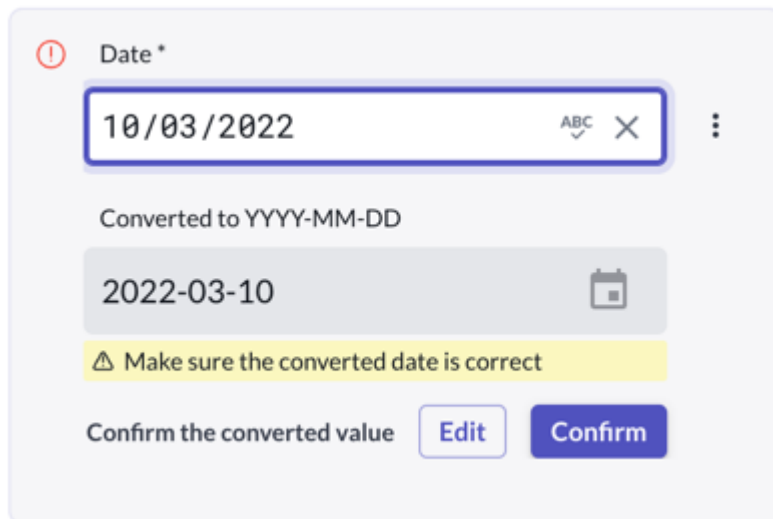

Field type	Description
Date	Standard date format. For example, YYYY-MM-DD.
Reference field	<p>A field that uses a field in another table as a standard. DocIntel matches the extracted data to the standard.</p> <p>For example, a use case has a reference field called <b>Vendor</b> that points to the Name 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 <b>Vendor</b> 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.</p>  <p>The diagram illustrates the matching process. On the left, a 'Document task' form shows a 'Vendor' field containing the text 'Degas Dairy Products, Inc'. A green box highlights this field. An arrow points from this field to a 'Company (Reference table)' on the right. The reference table has a 'Name' column. A green box highlights the entry 'Degas Dairy Products, Inc' in the 'Name' column. Another arrow points from this entry back to the 'Vendor' field in the document task, indicating a successful match.</p>
Integer	Whole number. For example, 12.

Field type	Description
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.



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.

#### 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: <ul style="list-style-type: none"> <li>• A required field must be filled in or marked as missing in the document.</li> <li>• An extracted value has a low confidence score and should be reviewed.</li> <li>• An extracted value is ambiguous and should be verified.</li> </ul>

### Document Intelligence forms


Use forms to view and update Document Intelligence information.

### Check box list form

The Check box list form enables you to define a check box list for extraction.

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.  <b>Note:</b> 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 <a href="#">Integrate with a custom application or workflow</a> .
This field is required for extraction	Option to make a field required.  Required fields can't be left unreviewed.  Required fields affect how document tasks are processed in the Fully automated extraction mode. For more information, see <a href="#">Configure data extraction modes</a> .
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

**Check box list form (continued)**

Field	Description
	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.

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	<p>The type of the field. For example, a text or date field.</p> <p>Some field types convert the extracted value into a standard format. For more information, see <a href="#">Data normalization</a>.</p>
Converted date format	<p>The date format that the extracted value is converted to for data normalization.</p> <p>This field is available when the <b>Type</b> field is set to <b>Date</b>.</p>
Target table	The table that stores the document processing results for this use case.
Target field	<p>Field on the target table that you want to align this field with.</p> <p><b>Note:</b> The use case must have a target table selected.</p> <p>This field is used for integration with other applications. For more information, see <a href="#">Integrate with a custom application or workflow</a>.</p>
When the date is ambiguous in a document, DocIntel will interpret it in the following order	<p>The default interpretation of the date format.</p> <p>For example, if you select <b>Month first</b> 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</p>

Single field form (continued)

Field	Description
	<p>select Day first, it will interpret it as February 1.</p> <p>This field is available when the <b>Type</b> field is set to Date.</p>
<p>When the number is ambiguous in a document, DocIntel will interpret it as</p>	<p>The default interpretation of the number format.</p> <p>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.</p> <p>This field is available when the <b>Type</b> field is set to Integer, Decimal, or Float.</p>
<p>Reference table</p>	<p>The table that stores the reference column. It's automatically populated based on the selected target field.</p> <p>This field displays when the <b>Type</b> field is set to Reference field.</p>
<p>Reference column</p>	<p>The column in the reference table that contains the referenced data.</p> <p>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 <a href="#">Data normalization</a>.</p> <p>This field is available when the <b>Type</b> field is set to Reference field.</p>
<p>Distinguisher(s)</p>	<p>Additional columns in the reference table that help the user to distinguish between similar records.</p> <p>This field is available when the <b>Type</b> field is set to Reference field.</p>
<p>This single field is required for extraction</p>	<p>Option to make the field required.</p> <p>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.</p>

### Single field form (continued)

Field	Description
	<p>If it's a reference field type, the required field must have a valid, exact match. By default, DocIntel uses the first matched record.</p> <p>Required fields affect how document tasks are processed in the Fully automated extraction mode. For more information, see <a href="#">Configure data extraction modes</a>.</p>
<p>Select one of these options for cases when:</p> <ul style="list-style-type: none"> <li>the date/number is ambiguous in the document</li> <li>there are multiple reference matches in the document</li> </ul>	<p>Option for agent review in situations when DocIntel encounters an ambiguous value in a required field.</p> <p>In such cases, the selected default interpretation will apply to the extracted value.</p> <p>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.</p>
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.

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	<p>Field on the target table that you want to align this field group with.</p> <p><b>Note:</b> You must first select a target table.</p>
This field group is required for extraction	Option to make the field as required.


Single field group form (continued)

Field	Description
	<p>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.</p> <p>If it's a reference field type, the required field must have a valid, exact match. By default, DocIntel uses the first matched record.</p> <p>Required fields affect how document tasks are processed in the Fully automated extraction mode. For more information, see <a href="#">Configure data extraction modes</a>.</p>
Fields	
Field name	<p>The name for the field as it appears in the Document Intelligence workspace.</p>
Type	<p>The type of the field. For example, a text or date field.</p> <p>Some field types convert the extracted value into a standard format. For more information, see <a href="#">Data normalization</a>.</p>
Select Target Field	<p>Field on the target table that you want to align this field with.</p> <p>This field is used for integration with other applications. For more information, see <a href="#">Integrate with a custom application or workflow</a>.</p>
When the date is ambiguous in a document, DocIntel will interpret it in the following order	<p>The default interpretation of the date format.</p> <p>For example, if you select <b>Month first</b> 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 <b>Day first</b>, it will interpret it as February 1.</p> <p>This field is available when the <b>Type</b> field is set to <b>Date</b>.</p>
When the number is ambiguous in a document, DocIntel will interpret it as	<p>The default interpretation of the number format.</p>

Single field group form (continued)

Field	Description
	<p>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.</p> <p>This field is available when the <b>Type</b> field is set to Integer, Decimal, or Float.</p>
Reference table	<p>The table that stores the reference column. It's automatically populated based on the selected target field.</p> <p>This field displays when the <b>Type</b> field is set to Reference field.</p>
Reference column	<p>The column in the reference table that contains the referenced data.</p> <p>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 <a href="#">Data normalization</a>.</p> <p>This field is available when the <b>Type</b> field is set to Reference field.</p>
Distinguisher(s)	<p>Additional columns in the reference table that help the user to distinguish between similar records.</p> <p>This field is available when the <b>Type</b> field is set to Reference field.</p>
This field is required for extraction	<p>Option to make a field required.</p> <p>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.</p> <p>If it's a reference field type, the required field must have a valid, exact match. By default, DocIntel uses the first matched record.</p> <p>Required fields affect how document tasks are processed in the Fully automated extraction mode. For more information, see <a href="#">Configure data extraction modes</a>.</p>

### Single field group form (continued)

Field	Description
<p>Select one of these options for cases when:</p> <ul style="list-style-type: none"> <li>the date/number is ambiguous in the document</li> <li>there are multiple reference matches in the document</li> </ul>	<p>Option for agent review in situations when DocIntel encounters an ambiguous value in a required field.</p> <p>In such cases, the selected default interpretation will apply to the extracted value.</p> <p>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.</p>
New single field	<p>Option to add a field to the group.</p> <p>Use the reorder icon (  ) to reorder a field in the group.</p>
Create multiple field groups	<p>Option to keep the pop-up window displayed on the screen. Enable this option if you're adding more than one single field group to the use case.</p>

#### Table form

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.
Target table	The table that stores the document processing results for these table fields.
Parent mapping to field	<p>Field on the target table that you want to align this table with.</p> <p> <b>Note:</b> You must first select a target table.</p>
This table is required for extraction	Option to make the table fields required.


Table form (continued)

Field	Description
	<p>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.</p> <p>Required fields affect how document tasks are processed in the Fully automated extraction mode. For more information, see <a href="#">Configure data extraction modes</a>.</p>
Columns	
Column title	Name of the column header in the table.
Type	<p>The type of the field in the table column. For example, a text or date field.</p> <p>Some field types convert the extracted value into a standard format. For more information, see <a href="#">Data normalization</a>.</p>
Select target field	<p>Field on the target table that you want to align this field with.</p> <p>This field is used for integration with other applications. For more information, see <a href="#">Integrate with a custom application or workflow</a>.</p>
When the date is ambiguous in a document, DocIntel will interpret it in the following order	<p>The default interpretation of the date format.</p> <p>For example, if you select <b>Month first</b> 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 <b>Day first</b>, it will interpret it as February 1.</p> <p>This field is available when the <b>Type</b> field is set to <b>Date</b>.</p>
When the number is ambiguous in a document, DocIntel will interpret it as	<p>The default interpretation of the number format.</p> <p>For example, if you select <b>1,00</b> 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.</p>

**Table form (continued)**

Field	Description
	<p>This field is available when the <b>Type</b> field is set to Integer, Decimal, or Float.</p>
Reference table	<p>The table that stores the reference column. It's automatically populated based on the selected target field.</p> <p>This field displays when the <b>Type</b> field is set to Reference field.</p>
Reference column	<p>The column in the reference table that contains the referenced data.</p> <p>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 <a href="#">Data normalization</a>.</p> <p>This field is available when the <b>Type</b> field is set to Reference field.</p>
Distinguisher(s)	<p>Additional columns in the reference table that help the user to distinguish between similar records.</p> <p>This field is available when the <b>Type</b> field is set to Reference field.</p>
This field is required for extraction	<p>Indicates whether a column required. It's automatically selected or cleared based on whether the table is required.</p> <p>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.</p> <p>If it's a reference field type, the required field must have a valid, exact match. By default, DocIntel uses the first matched record.</p> <p>Required fields affect how document tasks are processed in the Fully automated extraction mode. For more information, see <a href="#">Configure data extraction modes</a>.</p>

**Table form (continued)**

Field	Description
<p>Select one of these options for cases when:</p> <ul style="list-style-type: none"> <li>the date/number is ambiguous in the document</li> <li>there are multiple reference matches in the document</li> </ul>	<p>Option for agent review in situations when DocIntel encounters an ambiguous value in a required field.</p> <p>In such cases, the selected default interpretation will apply to the extracted value.</p> <p>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.</p>
New column	<p>Option to add a column to the table.</p> <p>Use the reorder icon (  ) to reorder a column in the table.</p>
Create multiple tables	<p>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.</p>

**Document Intelligence properties**

Document Intelligence (DocIntel) system properties control the behavior of the Document Intelligence application.

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.	<p>Allowed: Format of [ integer ] px</p> <p>Default: 416px</p>
sn_docintel.default_image_fit	The default image fit for the document panel	<p>Allowed: fit_to_page, fit_to_width</p>

Property name	Description	Values
	viewer in the Document Intelligence workspace.	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 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  Minimum number of days between trainings for a given document classifier use case.	Default: 30
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

Property name	Description	Values
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
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.

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

## Groups

List of groups this role is assigned to by default.

None.

## Special considerations

**i 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

**i 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

**i 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.

## Groups

List of groups this role is assigned to by default.

None.

## Special considerations

**i 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

**i 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.

### Updated terminology

As of Document Intelligence v2.3, DocIntel features include updated terminology.

#### Terminology updates

New term	Old term
Document task	Task, extraction task
Field	Key, attribute
Field group	Key group

**Terminology updates (continued)**

New term	Old term
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.

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**

In Document Intelligence, a bit of text found on a document. The recommendation includes information about its location in the document, meaning the page of a document and the particular 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.

**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 <b>Process Task</b> . 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.
Done	The task is completed after having been automatically processed in Fully Automated mode or after a user has reviewed the fields and selected <b>Submit</b> .

**Statuses for document tasks (continued)**

Status	Description
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.

**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.

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.

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. There are also some limitations with Now Assist in Document Intelligence that differ from the limitations in Document Intelligence.

## Document Intelligence limits

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 <a href="#">Languages supported by Document Intelligence</a> .
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.
Character set	Document Intelligence detects the following characters: <ul style="list-style-type: none"> <li>• a-z, A-Z, 0-9</li> <li>• á à â ã ä å ß ç é è ê ë ì í î ï ñ ó ò ô õ ö ú û ü œ</li> <li>• Á À Â Ã Ä Å ß Ç È É Ê Ë Ì Í Î Ï Ñ Ó Ò Ô Õ Ö Ú Û Ü Û Œ</li> <li>• , ; . ' \ " ! ? ¿ ¡ + - * ( ) [ ] { } &amp; % @ # /   ~ ^ &lt; &gt; ` ± = _ \$ £ ¢ €</li> </ul>
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.

### Now Assist in Document Intelligence limits

The following table is a list of the important limitations in Now Assist in Document Intelligence that differ from the limitations in Document Intelligence.



### Limits in Now Assist in Document Intelligence

Limit	Description
File formats	<p>For document extraction, the supported file format is PDF. The PDF may include pages that are created from digital or scanned paper documents.</p> <p>For document analysis, the supported file format is docx (Microsoft Word document).</p>
File size limits	<p>The file size limit is 5 MB.</p> <p>The file size limit is also determined by the <code>com.glide.attachment.max_size</code> system property.</p> <p>For more information, see <a href="#">Configure attachment system properties</a>.</p>
Page count limit per document task	<p>The page count limit is 15 pages per file.</p>

## 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.

### Get started

<p><b>Explore</b></p>  <p>Learn about the generative AI skills that are available in Now Assist in Document Intelligence</p>	<p><b>Configure</b></p>  <p>Activate Now Assist in Document Intelligence and configure the generative AI skills</p>
---	--

**Use**

Use the generative AI capabilities that are offered by Now Assist in Document Intelligence

**Reference**

Get details about Now Assist in Document Intelligence

***i* 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](#).

**Disclaimer**

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. Further, 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](#).

## Exploring Now Assist in Document Intelligence


With Now Assist in Document Intelligence, your agents can use generative AI to help review documents for key information and extract document data for use in your workflows.

### Now Assist in Document Intelligence overview

Now Assist in Document Intelligence makes the following generative AI capabilities available to an 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 Now Assist in Virtual Agent](#) .

### Now Assist in Document Intelligence skills

For the Xanadu release, Now Assist in Document Intelligence includes the contract metadata extraction skill and the contract analysis skill. These skills are available with the Now Assist in Contract Management workflows. For more information, see [Now Assist in Contract Management](#) .

## 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.

For the Xanadu release, there is no stand-alone configuration of Now Assist in Document Intelligence. The contract metadata extraction skill and the contract analysis skill contain the Now Assist in Document Intelligence feature. These skills are available with the Now Assist in Contract Management workflows. For more information, see [Now Assist in Contract Management](#) .

## 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.

For the Xanadu release, Now Assist in Document Intelligence includes the contract metadata extraction skill and the contract analysis skill. These skills are available with the Now Assist in Contract Management workflows. For more information, see [Now Assist in Contract Management](#).

### Extracting document information with the Now Assist in Document Intelligence workspace

Use the Document Intelligence workspace to review the information that was extracted from a document by Now Assist in Document Intelligence. You can also search for additional recommendations, flag fields, and identify missing information in the document.

#### Document extraction overview

Now Assist uses generative AI capabilities to provide recommendations for the field values that were extracted from the document. The Document Intelligence workspace provides document management features that enable you to quickly review and process text extraction.

With the Document Intelligence workspace, you can:

- Efficiently review the Now Assist recommendations, and confirm your document's extracted text.
- Flag fields, identify missing fields, and review pending fields.

For the Xanadu release, Now Assist in Document Intelligence includes the contract metadata extraction skill and the contract analysis skill. These skills are available with the Now Assist in Contract Management workflows. For more information, see [Now Assist in Contract Management](#).

#### Document Intelligence workspace example

The following example shows the Document Intelligence workspace in the Now Assist in Contract Management workflow. Now Assist in Document Intelligence provides the field values that were extracted from the contract document.

## Document Intelligence workspace using Now Assist in Document Intelligence for extraction

The screenshot displays a document workspace for a Non-Disclosure Agreement (NDA). The document is titled "demoNDA (1)\_1.jpeg" and is currently in an "In Progress" state. The main content area shows the text of the NDA, including sections for "[Disclosing Party]", "[Receiving Party]", and "NON-DISCLOSURE AGREEMENT". The document is divided into three pages, with the first page visible. A "Fields" panel on the right side of the workspace shows extracted fields: "Document Name \*", "Parties \*", "Agreement Date \*", "Effective Date \*", and "Expiration Date \*". The "Document Name" field is highlighted, and a tooltip is visible with options "Flag for follow-up" and "Missing in the document".

For more information on the Document Intelligence workspace, see [Extract fields using the Document Intelligence workspace](#).

## Reviewing a legal document with Now Assist in Document Intelligence

Use Now Assist in Document Intelligence so that generative AI can provide you with a document analysis for contract reviews.

Now Assist uses generative AI capabilities to identify the specified clauses in the document text and provide recommendations for the standardized text.

For the Xanadu release, Now Assist in Document Intelligence includes the contract metadata extraction skill and the contract analysis skill. These skills are available with the Now Assist in Contract Management workflows. For more information, see [Now Assist in Contract Management](#).

## Now Assist in Document Intelligence reference

The following topics provide additional information about Now Assist in 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. Extraction modes only apply to use cases for the document extraction skill.

Now Assist in Document Intelligence uses the following data extraction modes.

Extraction mode	Description
Default	<p>Now Assist auto-fills each field in the Document Intelligence workspace if it recognizes a value for it in the document.</p> <p>Now Assist may not always produce accurate, complete, or appropriate information. All fields should be reviewed for accuracy.</p> <p>A reviewer can accept the predicted value or manually enter a value.</p>
Full automation	<p>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.</p> <p>By turning on Full automation mode, the agent review used to check the accuracy of the predicted values is bypassed.</p> <p>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.</p>

## Document and visual insights AI agent

The document and visual insights AI agent gathers context from user input and image attachments, generates the requested information based on the image 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 using existing use cases.

#### Attachment summarization

Summarize the key topics of document and image attachments.

#### Document 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 agent actions

When used, the agent may attempt the following actions:

- Determine the type of task to perform based on information provided by the user.
  - Data extraction
  - Attachment summarization
  - Document question answering (Q&A)
- Retrieve the relevant record details.
- Initiate the process based on the type of task.
- Provide the task results as directed.
- Notify the user of task completion and next steps.







## 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

<p style="text-align: center;">Explore</p>  <p style="text-align: center;">Explore Task Intelligence and machine learning models.</p>	<p style="text-align: center;">Install</p>  <p style="text-align: center;">Install Task Intelligence and related applications.</p>	<p style="text-align: center;">Configure</p>  <p style="text-align: center;">Configure machine learning models.</p>
<p style="text-align: center;">Create models</p>  <p style="text-align: center;">Create and edit Task Intelligence models.</p>	<p style="text-align: center;">Analyze models</p>  <p style="text-align: center;">Monitor and analyze model performance.</p>	<p style="text-align: center;">References</p>  <p style="text-align: center;">References for components installed with Task Intelligence and language support.</p>

## 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

The screenshot displays the 'Task Intelligence' section in the ServiceNow Admin Console. At the top, there's a navigation bar with 'servicenow', 'All', 'Favorites', 'History', and a search bar. Below this, a header reads 'Improve task creation with machine learning' with a sub-header 'Task Intelligence uses machine learning models built from your data to enhance aspects of task creation.'

A central card titled 'Let's go over a few things about machine learning models' contains four bullet points:
 

- All models are statistical models.
- Models predict future data by using past data.
- We'll help you create a model to fit your goals.
- We don't require coding or calculations.

Below this is a 'Models' table with the following data:
 

Name ↑	Model type	Tags	State	Date created	
<a href="#">Case sentiment</a>	Sentiment	Sentiment	Training completed	10/17/2023	<a href="#">View metrics</a>

The bottom section, 'Set up your machine learning models', features three cards:
 

- Autofilled Prediction:** Predict field choices to reduce handle time. Includes a 'Set up model' button.
- Case sentiment:** Predict case sentiment to improve CSAT. Includes an 'Edit model' button.
- Case language:** Predict case language to improve assignments. Includes a 'Set up model' button.

The Admin Console supports the following model types:

Model	Application	Description
Incident field prediction	<a href="#">Task Intelligence for ITSM</a>	<a href="#">Create an incident field prediction model</a>
Similarity model management	<a href="#">Task Intelligence for ITSM</a>	<a href="#">Create a similar records prediction model in Task Intelligence for ITSM</a>
Case field prediction	<a href="#">Task Intelligence for CSM</a>	<a href="#">Predicts output fields on case forms, including categorization. Record categorization</a>

Model	Application	Description
Sentiment prediction	<a href="#">Task Intelligence for CSM</a>	<p>Predicts current and trending user sentiment.</p> <p><b>Note:</b> Model is pre-trained but can be tested and edited. Only one sentiment model can be configured.</p>
Language detection	<a href="#">Task Intelligence for CSM</a>	<p>Detects language which language is being used.</p> <p><b>Note:</b> Model is pre-trained but can be tested and edited. Only one language detection model can be configured.</p>

### 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	<a href="#">Record categorization</a>	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	<a href="#">Sentiment Analysis</a>	Agents
Identify the language used to create customer service cases	<a href="#">Language detection</a>	Agents

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.

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 comprises the sample records that describe application features for the 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.

**Create a model to detect case language** 

Edit and test the pre-trained model to detect the language used to create customer service cases.

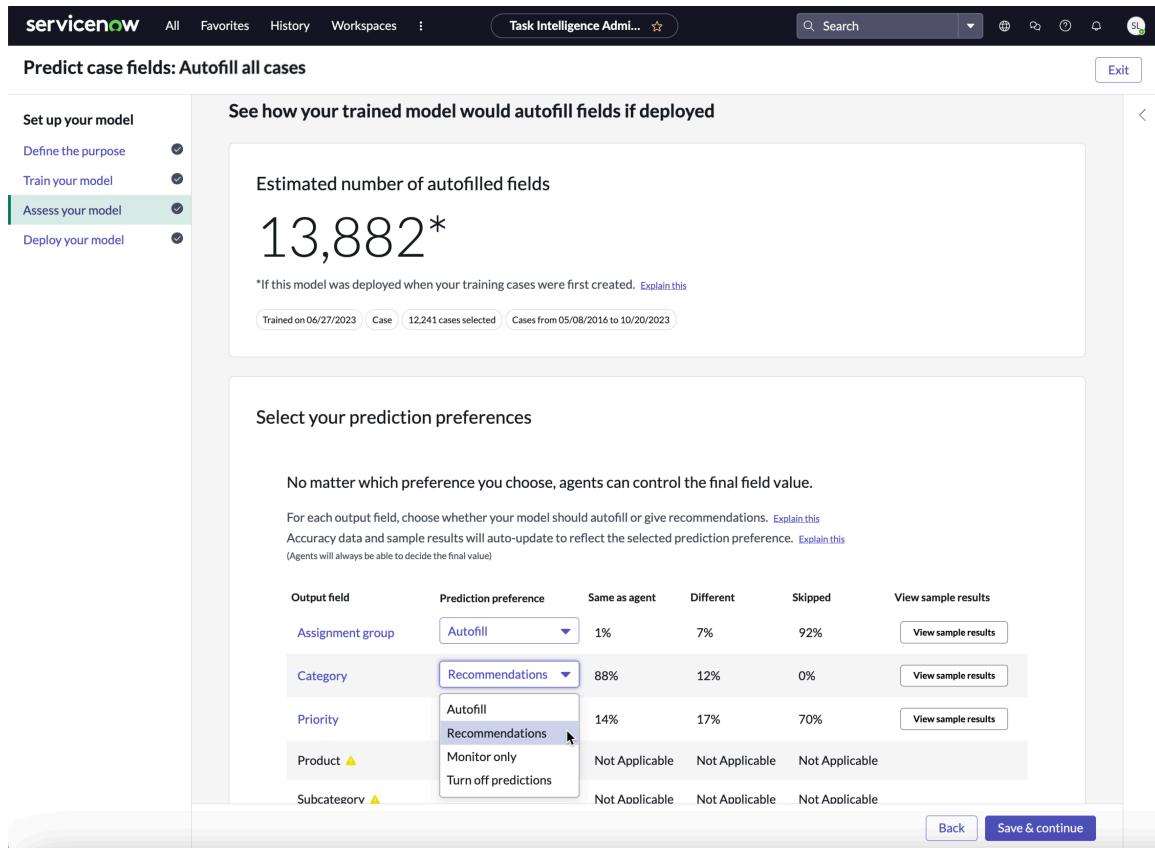
**Assess a Task Intelligence model**

Assessing a machine learning model's performance helps determine how to use and train the model to achieve your desired outcomes.

**Overview of assessing a Task Intelligence model**

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.

## 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

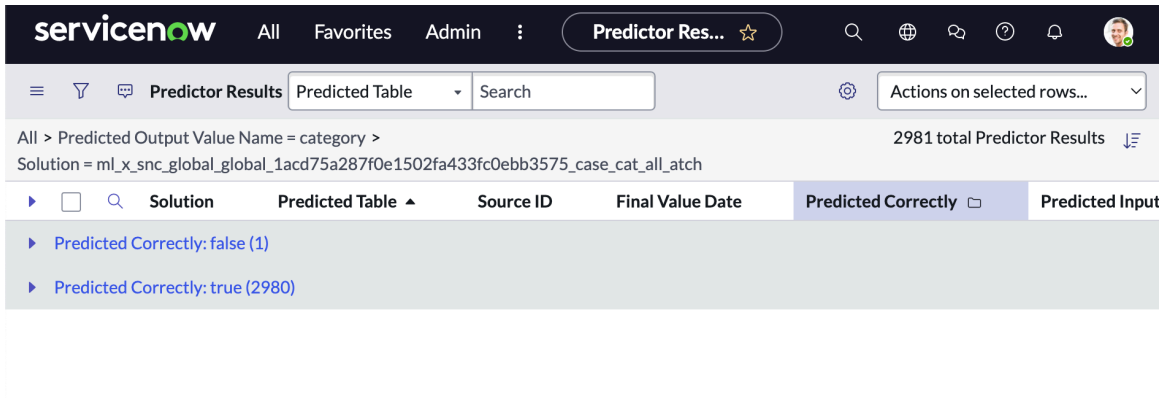
After you have created and tested your model, assess its performance to evaluate how well it's predicting what you want it to predict. For multi-output models that predict multiple fields, you may want to assess the performance of each field individually.

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

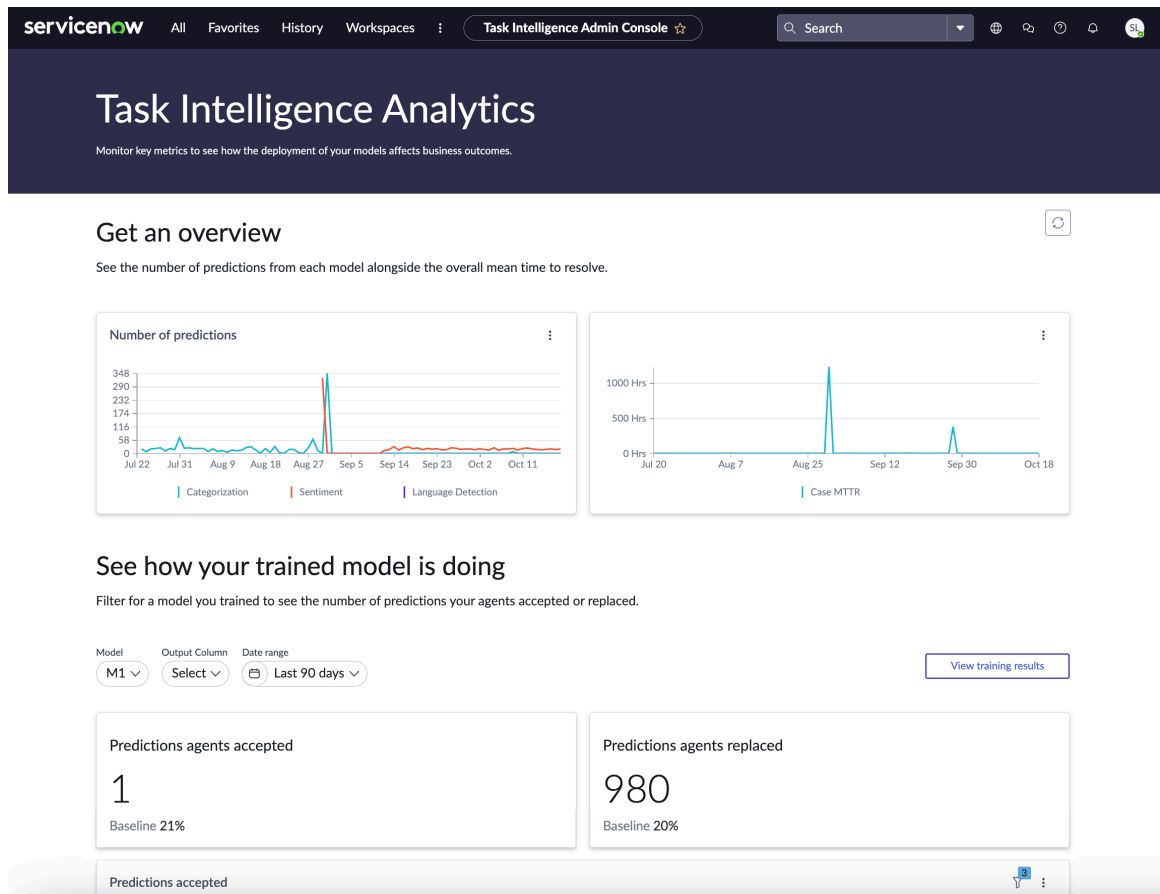
1. Navigate to the **Predictor Results** (`ml_predictor_results_task`) by entering `ml_predictor_results_task.list` in the Application navigator.
2. Filter the **Predicted Output Value Name** for the name of the field you want to assess, such as "product" or "category."
3. Group the list by **Predicted 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.

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

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.

**Models**

Name ↑	Model type	Tags	State	Date created	
<a href="#">Case sentiment</a>	Sentiment		Deployed	06/22/2022	<a href="#">View metrics</a> ...
<a href="#">Priority Case Model</a>	Categorization	<a href="#">All Cases</a> <a href="#">Without Attachment</a>	Deployed	06/15/2022	<a href="#">View metrics</a> ...
<a href="#">General Case Prediction</a>	Categorization	<a href="#">All Cases</a> <a href="#">Without Attachment</a>	Deployed	Today	<a href="#">View metrics</a> 
<a href="#">Specific Case Model</a>	Categorization	<a href="#">Without Attachment</a>	Training Error 	06/24/2022	<a href="#">View metrics</a> <a href="#">Edit model</a> <a href="#">Export model</a> <a href="#">Delete model</a>

1 to 4 c

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.

## 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

**i 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.

None.

#### Special considerations

**i 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.

**i 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

Language name	Language code	Incident fields (ITSM)	Case fields (CSM)	Sentiment prediction (CSM)	Language detection (CSM)
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
Russian	ru				X
Spanish	es	X	X		X
Swedish	sv	X			X
Thai	th				X
Turkish	tr				X

## 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	<a href="#">Now Assist AI Agents Release Notes</a> 

**Resource links (continued)**

AI features or products	Resource type	Resources
AI Agents	Getting started	<a href="#">Getting started with AI Agents</a> . <a href="#">Identifying agentic workflows</a> .
AI Agents	Best practices	<a href="#">Build custom skills</a> <a href="#">Assess the AI agents' and agentic workflows' performance</a>
AI Agents	Glossary	<a href="#">Now Assist AI Agents Glossary.</a>
AI Agents	Training	<a href="#">AI Agents Essentials (2025 Q1 Store Release)</a> <a href="#">AI Agent Readiness</a> <a href="#">AI Essentials</a> <a href="#">Initiating AI Agents</a> <a href="#">Introduction to Agentic AI</a> <a href="#">Now Assist AI Agents learning path</a>
AI Control Tower	Getting started	<a href="#">Overview of Advanced AI and Data Products</a>
Document Intelligence	Release notes	<a href="#">Document Intelligence release notes</a>
Document Intelligence	Getting started	<a href="#">Getting started with Document Intelligence</a> <a href="#">Document Intelligence quick start guide (ServiceNow Community)</a>
Document Intelligence	Tips and examples	<a href="#">Document Intelligence examples</a>
Document Intelligence	Troubleshooting	<a href="#">Troubleshooting and best practices for Document Intelligence</a>
Document Intelligence	FAQs	<a href="#">Document Intelligence FAQs</a>
Document Intelligence	Training	<a href="#">Document Intelligence Overview</a>

**Resource links (continued)**

AI features or products	Resource type	Resources
		<a href="#">Document Intelligence Fundamentals</a> 
Natural Language Understanding and Natural Language Query	Getting started	<a href="#">Getting started with NLU</a>  <a href="#">NLU quick start guide (ServiceNow Community)</a> 
Natural Language Understanding and Natural Language Query	Best practices	<a href="#">Best practices for NLU</a> 
Natural Language Understanding and Natural Language Query	Tips and examples	<a href="#">NLU examples</a> 
Natural Language Understanding and Natural Language Query	Troubleshooting	<a href="#">Troubleshooting for NLU</a> 
Natural Language Understanding and Natural Language Query	FAQs	<a href="#">NLU FAQs</a> 
Natural Language Understanding and Natural Language Query	Training	<a href="#">Natural Language Understanding (NLU) Fundamentals</a> 
Now Assist	Release notes	<a href="#">Now Assist release notes</a> 










Resource links (continued)

AI features or products	Resource type	Resources
		<p><a href="#">Now Assist for Creator release notes</a></p> <p><a href="#">Now Assist for Customer Service Management (CSM) release notes</a></p> <p><a href="#">Now Assist for Field Service Management (FSM) release notes</a></p> <p><a href="#">Now Assist for Financial Services Operations (FSO) release notes</a></p> <p><a href="#">Now Assist for IT Service Management (ITSM) release notes</a></p> <p><a href="#">Now Assist for Security Operations release notes</a></p> <p><a href="#">Now Assist for Telecommunications, Media and Technology (TMT) release notes</a></p> <p><a href="#">Now Assist in Virtual Agent release notes</a></p> <p><a href="#">Generative AI Controller release notes</a></p>
Now Assist	Getting started	<p><a href="#">Getting Started with Now Assist (YouTube)</a></p> <p><a href="#">Getting started with Now Assist for Code (YouTube)</a></p> <p><a href="#">Getting Started with Now Assist for Creator (YouTube)</a></p> <p><a href="#">Getting Started with Now Assist for CSM (YouTube)</a></p> <p><a href="#">Getting Started with Now Assist for ITSM (YouTube)</a></p> <p><a href="#">Getting Started with Now Assist for AI Search (YouTube)</a></p> <p><a href="#">Now Assist quick start guide (ServiceNow Community)</a></p> <p><a href="#">Now Assist for Code quick start guide (ServiceNow Community)</a></p> <p><a href="#">Now Assist panel fundamentals (YouTube)</a></p>
Now Assist	Best practices	<p><a href="#">Now Assist best practices (ServiceNow Community)</a></p> <p><a href="#">Knowledge article best practices (ServiceNow Community)</a></p>
Now Assist	Tips and examples	<p><a href="#">Now Assist conversational catalog request in practice (ServiceNow Community)</a></p> <p><a href="#">Now Assist examples</a></p>
Now Assist	FAQs	<p><a href="#">Generative AI Controller FAQ</a></p> <p><a href="#">Now Assist FAQ</a></p>

Resource links (continued)

AI features or products	Resource type	Resources
		<a href="#">Now Assist in AI Search FAQ</a>
Now Assist	Training	<a href="#">Introduction to Generative AI</a> <a href="#">Now Assist Essentials (November 2024 Store Release)</a> <a href="#">Now Assist in Virtual Agent Implementation (November 2023 Store Release)</a> <a href="#">Now Assist in Virtual Agent Implementation (August 2024 Store Release)</a> <a href="#">Now Assist in Search Implementation (Nov 2023 Store Release)</a> <a href="#">Now Assist in Search Implementation (Feb 2024 Store Release)</a> <a href="#">Now Assist for Creator Implementer</a> <a href="#">Now Assist for Field Service Management (FSM) Implementer</a> <a href="#">Now Assist for Customer Service Management (CSM) Implementer</a> <a href="#">Now Assist for HR Service Delivery (HRSD) Implementer</a> <a href="#">Now Assist for IT Service Management (ITSM) Implementer</a>
Predictive Intelligence	Release notes	<a href="#">Predictive Intelligence release notes</a>
Predictive Intelligence	Getting started	<a href="#">Getting started with Predictive Intelligence</a> <a href="#">Predictive Intelligence quick start guide (ServiceNow Community)</a>
Predictive Intelligence	Best practices	<a href="#">Predictive Intelligence best practices</a>
Predictive Intelligence	Tips and examples	<a href="#">Predictive Intelligence examples</a>
Predictive Intelligence	FAQs	<a href="#">Predictive Intelligence FAQs</a>

**Resource links (continued)**

AI features or products	Resource type	Resources
Predictive Intelligence	Training	<a href="#">Predictive Intelligence Fundamentals</a> 
Task Intelligence	Getting started	<a href="#">Getting started with Task Intelligence</a>  <a href="#">Task Intelligence quick start guide (ServiceNow Community)</a> 
Task Intelligence	Best practices	<a href="#">Troubleshooting and best practices for Task Intelligence</a> 
Task Intelligence	FAQs	<a href="#">Task Intelligence FAQs</a> 
Task Intelligence	Training	<a href="#">Task Intelligence for ITSM - Quick Start Guide</a>  <a href="#">Task Intelligence for CSM - Quick Start Guide</a>  <a href="#">AI Academy Get Started with Task Intelligence -use AI to categorize cases</a>  <a href="#">AI Academy Task Intelligence for ITSM</a>  <a href="#">AI Academy Task Intelligence - Customer Service Management</a> 