Washington DC Enable AI

Last updated: 12/16/2025

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

servicenow

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

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

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

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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

servicenow

Table of Contents

Enable AI experiences **4**

Now Assist 8

Exploring Now Assist 10

Configuring Now Assist settings and features 41

Analyzing Now Assist performance 62

Now LLM Service updates 65

Now Assist reference 67

Generative AI Controller 72

Natural Language Understanding 119

Exploring Natural Language Understanding 120

Model management 133

Multilingual model management 227

Virtual Agent and NLU Workbench integration 242

NLU Workbench - Advanced Features 244

Natural Language Query 288

Exploring Natural Language Query 288

Using Natural Language Query 290

Configuring NLQ 292

Natural Language Query References 298

Predictive Intelligence 300

Explore Predictive Intelligence 301

Install Predictive Intelligence 304

Configure Predictive Intelligence 306

Creating and training solutions 333

Using Predictive Intelligence 368

Predictive Intelligence references 402

Document Intelligence 407

Exploring Document Intelligence 408

Configuring Document Intelligence 411

Integrating Document Intelligence with other applications 431

Using Document Intelligence 437

Monitoring Document Intelligence performance 456

Document Intelligence references 459

Task Intelligence 476

Exploring Task Intelligence 477

Configure Task Intelligence 480

Manage machine learning models with Task Intelligence 481

Task Intelligence references 486

Additional resources for AI products and solutions 488

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Enable AI experiences

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

Get started

Choose one of these tiles to get started.

Now Assist



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

Natural Language Understanding



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

Natural Language Query



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

Predictive Intelligence



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

Document Intelligence



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

Task Intelligence



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

Additional Resources

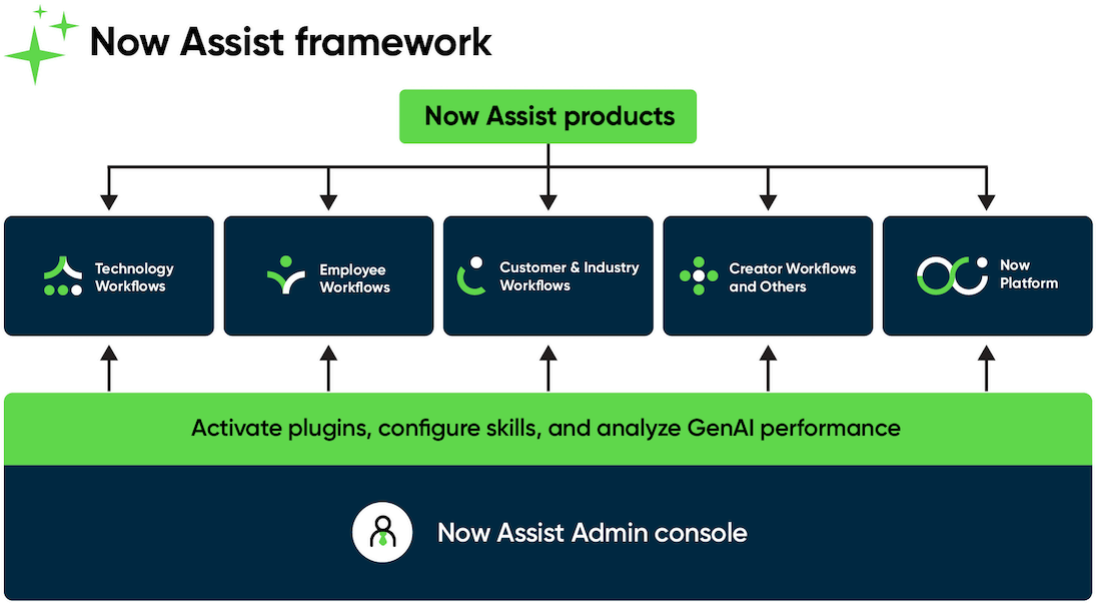


Upgrade your AI experience to the next level

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

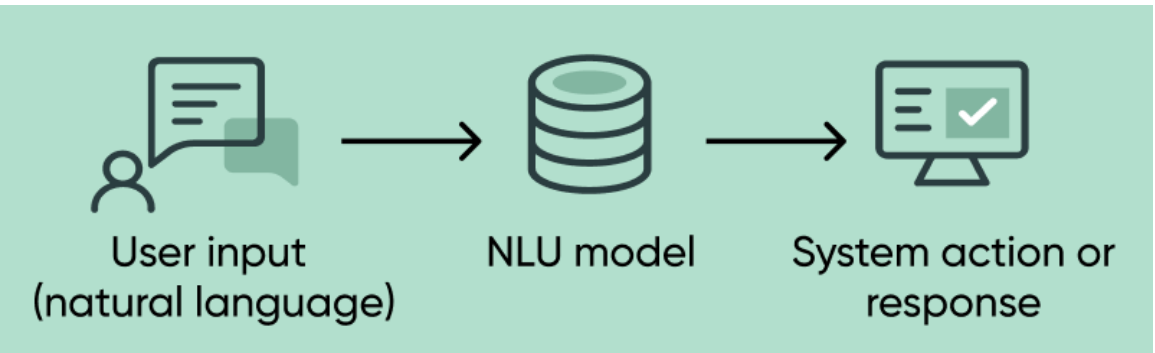
Now Assist

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.



Natural Language Understanding

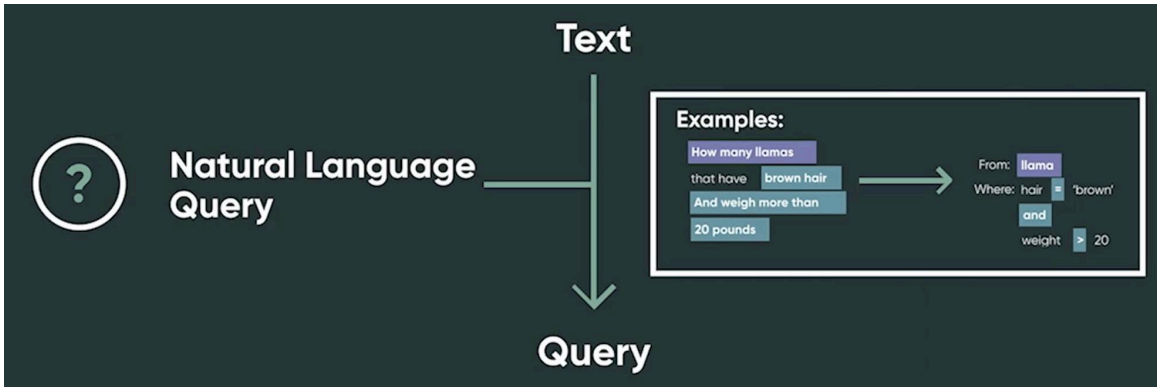
Help users communicate with your system in naturally-expressed language, using Natural Language Understanding. NLU enables your system to perform intelligent actions in response to human language input in 17 supported languages. Start from the provided pre-built models and expand them further, or build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



Natural Language Query

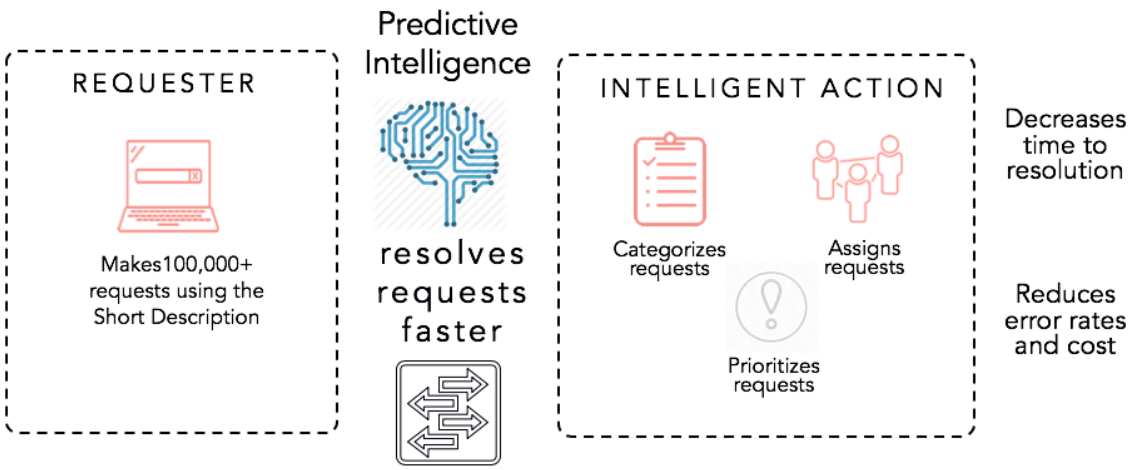
Transform natural-language questions into formal database queries with Natural Language Query (NLQ). Get data from your instance by using plain language requests in the supported languages American English, French, French Canadian, German, Japanese, and Spanish. NLQ is consumed by several other applications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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Predictive Intelligence

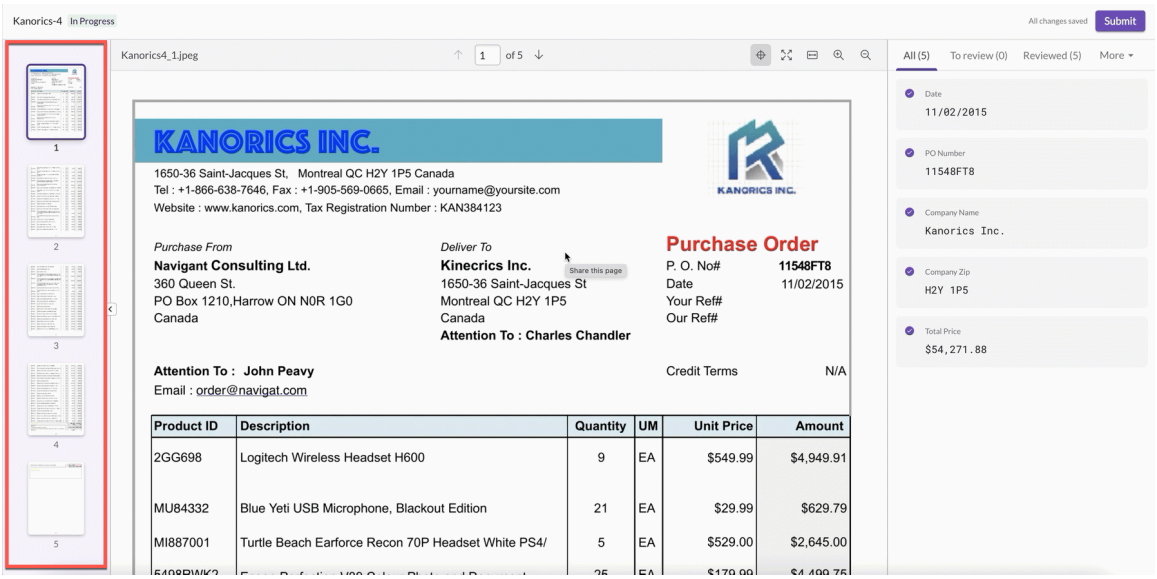
Train your system to resolve requests swiftly with the Predictive Intelligence Workbench. There are four separate machine-learning frameworks that support all NLQ languages along with Dutch, Italian, and Brazilian Portuguese. These frameworks help your AI to handle a range of business challenges. Reduce 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.



Document Intelligence

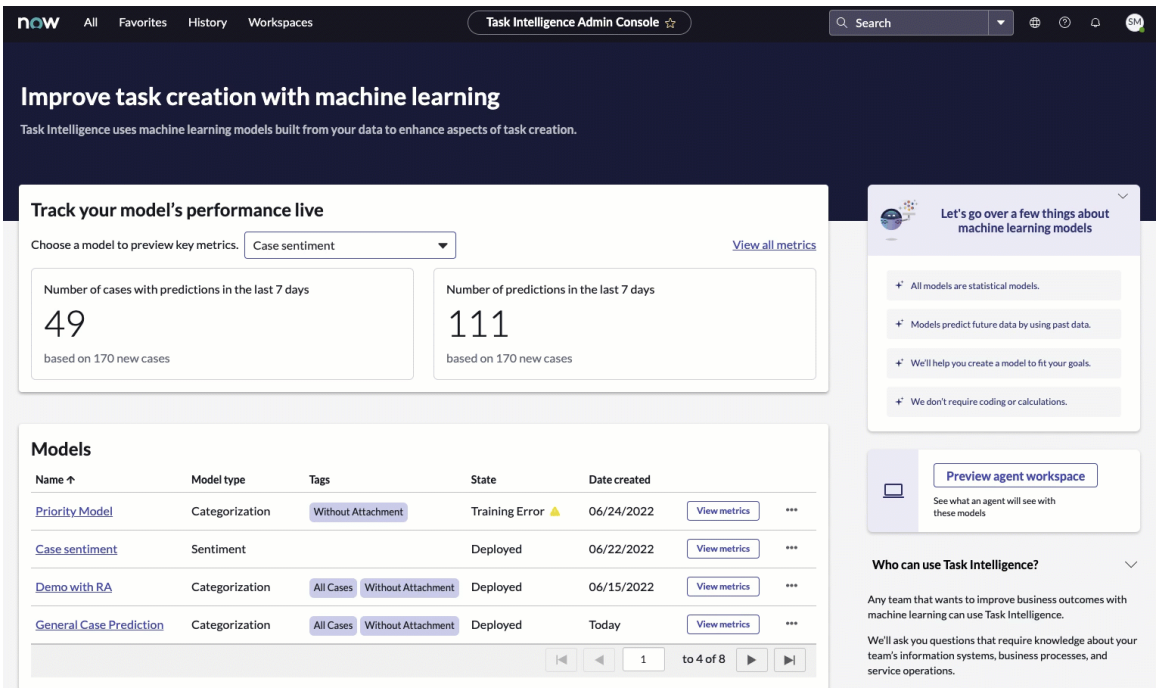
Get clear, concise data from your records quickly, and integrate it where you need it most. Document Intelligence can process single or multiple-page documents in various popular formats. More than just an optical character recognition (OCR) solution, this tool uses a continual-learning AI to identify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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Task Intelligence

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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

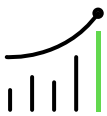
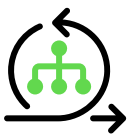
- Task Intelligence
- Document Intelligence

Now Assist

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

Put Generative AI to Work with Now Assist on the ServiceNow AI Platform.

Get started

Explore  Learn more about Now Assist on the ServiceNow AI Platform.	**Configure**  Explore, configure, use, and analyze generative artificial intelligence (AI) features and skills.	**Analyze**  Analyze and monitor skill performance.
Now LLM Service model updates  Learn about the Now LLM Service models.	**Reference**  Learn about user roles, data usage policy, and domain separation in Now Assist.	**Generative AI Controller**  Customize your generative AI experience on the ServiceNow AI Platform with the Generative AI Controller.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist products

Now Assist for Customer Service Management (CSM)



Summarize case information, generate case resolution notes, and summarize the chat for an interaction.

Now Assist for HR Service Delivery (HRSD)



Summarize case information, generate resolution notes, and summarize the chat for an interaction.

Now Assist for IT Service Management (ITSM)



Summarize incident information, generate incident resolution notes, and summarize the chat for an interaction.

Now Assist for Field Service Management (FSM)



Summarize work order tasks and enable users to generate work order task summaries.

Now Assist for IT Operations Management (ITOM)



Simplify alerts in Event Management.

Now Assist for Creator



Use generative AI skills that can make developing on the ServiceNow AI Platform more efficient.

Now Assist for Strategic Portfolio Management (SPM)



Analyze and extract actionable insights from the feedback of your customers in a summary

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist platform experiences

Now Assist in AI Search



The power of search with the Now LLM Service generative AI models.

Now Assist in Virtual Agent



Use generative AI skills and large language model (LLM) topic authoring in your conversational experiences.

Troubleshoot and get help

- ServiceNow Community on AI and Intelligence
- Issues and mitigation for Now Assist application and plugin upgrades
- 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

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. The following generative AI features are currently offered:

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Case or incident summarization
- Chat summarization
- Code generation
- KB generation
- Playbook generation
- Resolution notes generation
- Work order task closure summarization
- Now Assist in AI Search
- Flow generation
- Now Assist in Virtual Agent

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.

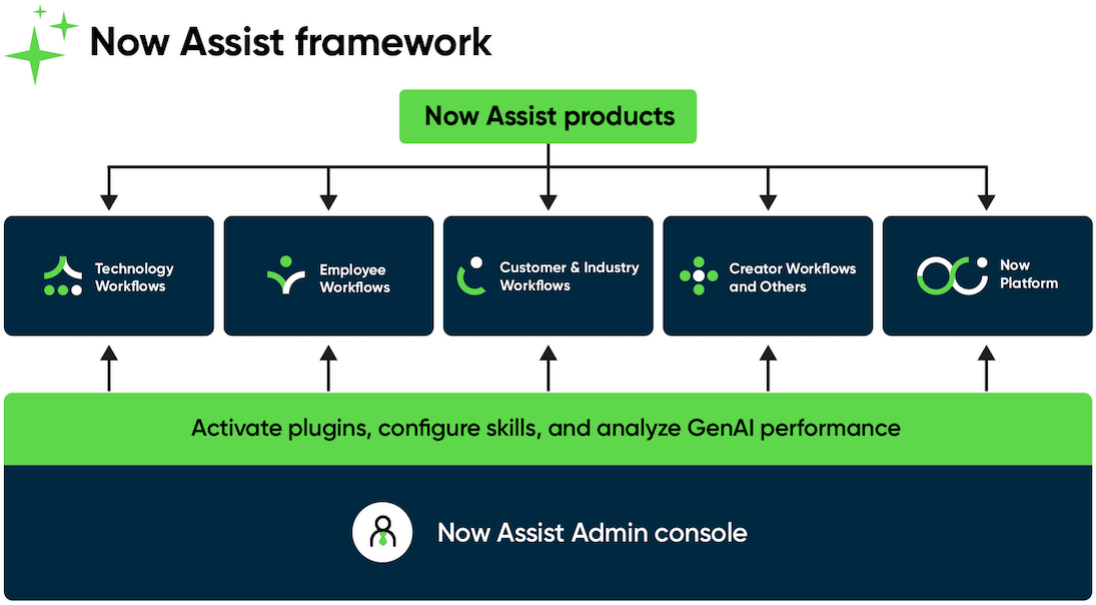
- Now Assist products:
  - Now Assist for Creator
  - Now Assist for Customer Service Management (CSM)
  - Now Assist for Field Service Management (FSM)
  - Now Assist for HR Service Delivery (HRSD)
  -
  - Now Assist for IT Service Management (ITSM)
  - Now Assist for Strategic Portfolio Management (SPM)
- ServiceNow AI Platform tools and skills for Now Assist:
  - Configuring Now Assist settings and features: Enable your administrators to install plugins, manage skills, and analyze usage and performance with the Now Assist Admin console.
  - Now Assist panel: Enable your users to take advantage of Now Assist skills through a panel interface on the instance.
  - Use Now Assist in AI Search to generate answers for AI Search.
  - Generative AI Controller: Enable your developers to integrate generative AI features in custom flows and conversations by using your own third-party large language model (LLM) licenses.

Some Now Assist products include the foundational platform tools for Now Assist. For details, see Now Assist skills in the Platform workflow.

The following diagram shows what's available in the Now Assist framework.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist framework



Now Assist benefits

Benefit	Feature	Users
Leverage the power of search with the Now LLM generative AI model to answer questions in user searches with actionable AI-generated summaries of relevant knowledge articles.	Now Assist in AI Search	Everyone
Install and configure Now Assist applications and the skills they provide.	Now Assist Admin console	Administrators
Choose which skills to turn on, and which users can access them.	Now Assist Admin console	Administrators
Monitor the usage and performance of generative AI skills.	Analyzing Now Assist performance	Administrators
Access generative AI skills in context through a user-friendly interface.	Now Assist panel	Everyone
Customize your workflows and use your own third-party LLM license.	Generative AI Controller	Administrators or developers
Use Now Assist in other platform features.	Now Assist skills in the Platform workflow	Administrators or developers
Monitor Now Assist consumption on your instance.		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.

Introduction to the Now Assist Admin console.

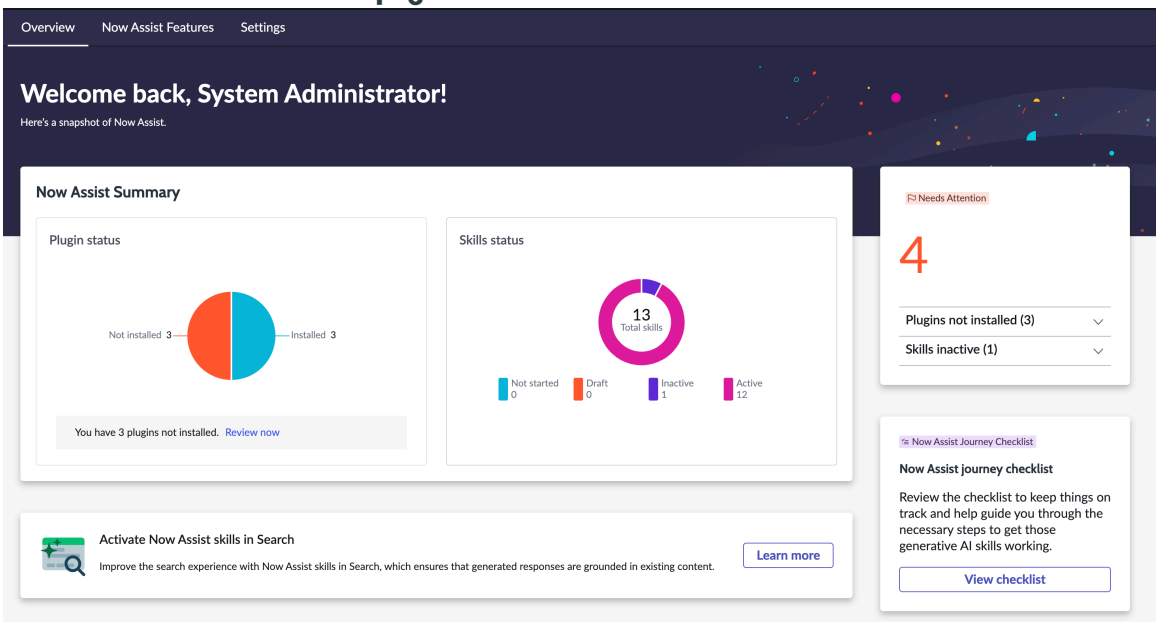
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist Admin 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 Overview page



Now Assist Admin workflow

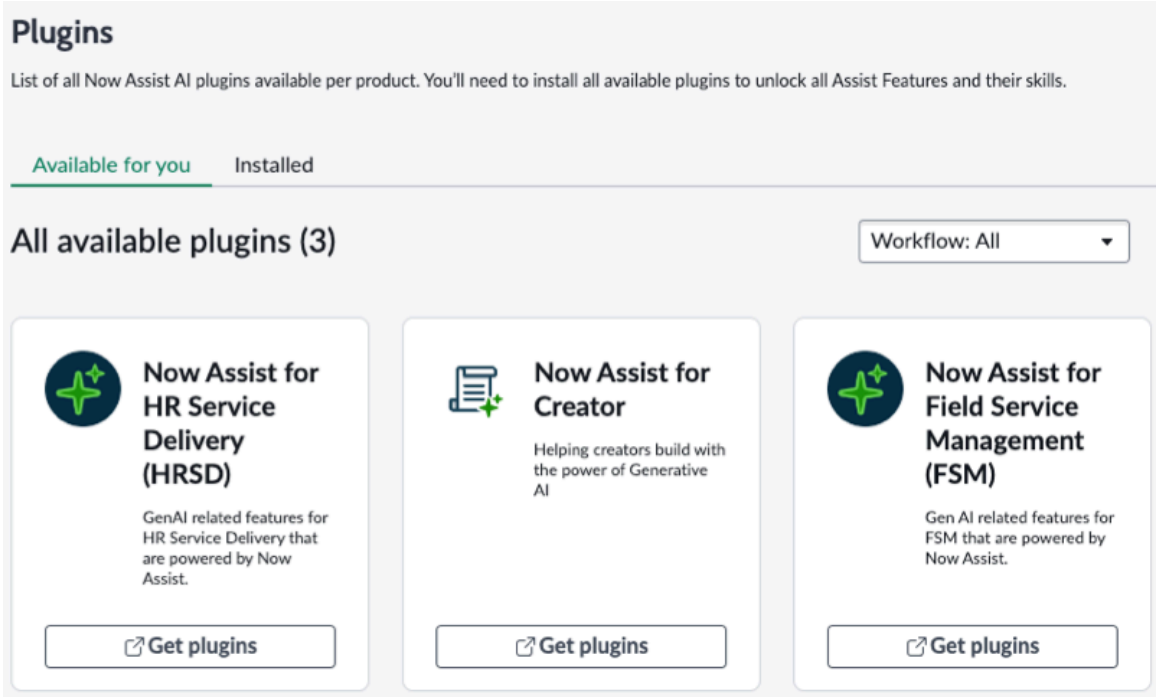
Take five steps to begin using the Now Assist Admin console.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Available 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.

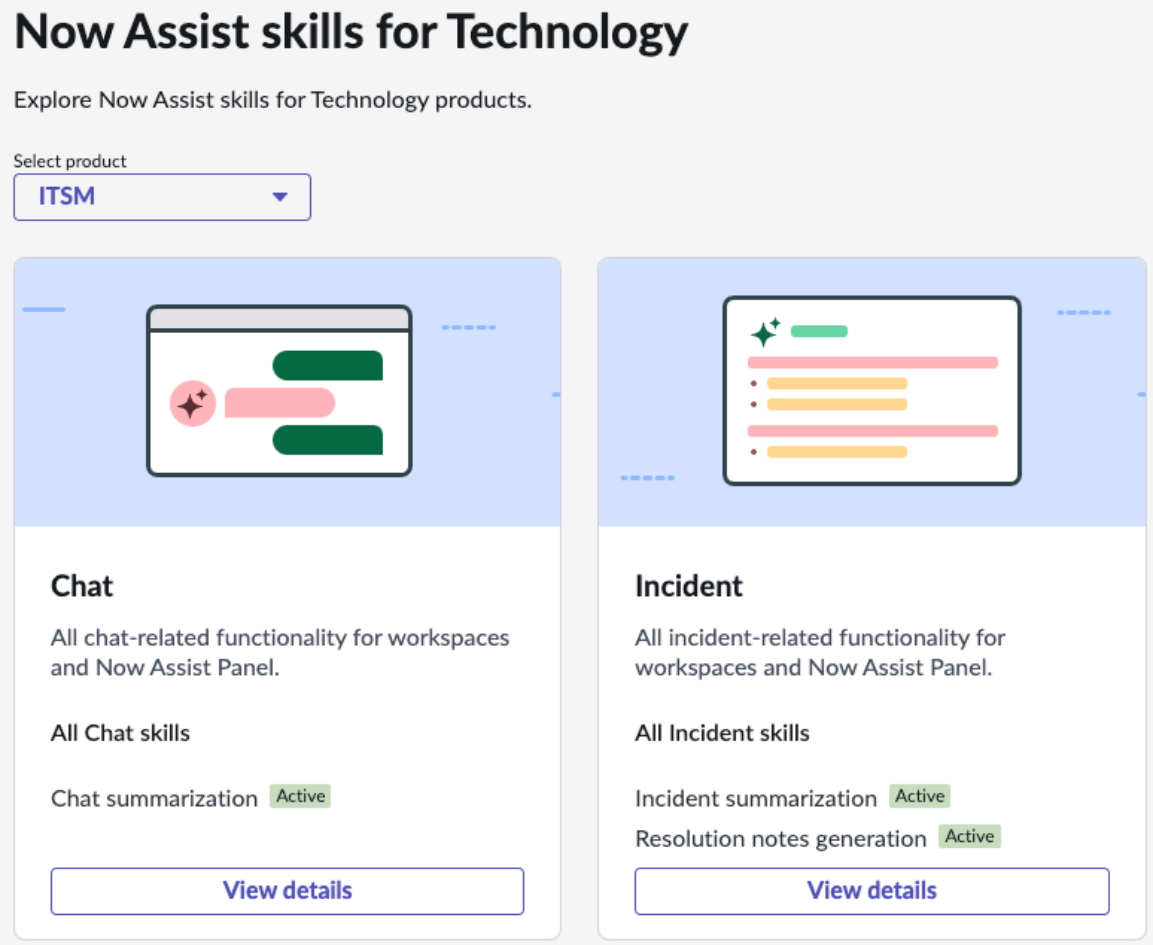
This step is optional because the skills can also display in-product in the Core UI and in Workspaces.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Available Now Assist features and skills in the Technology workflow

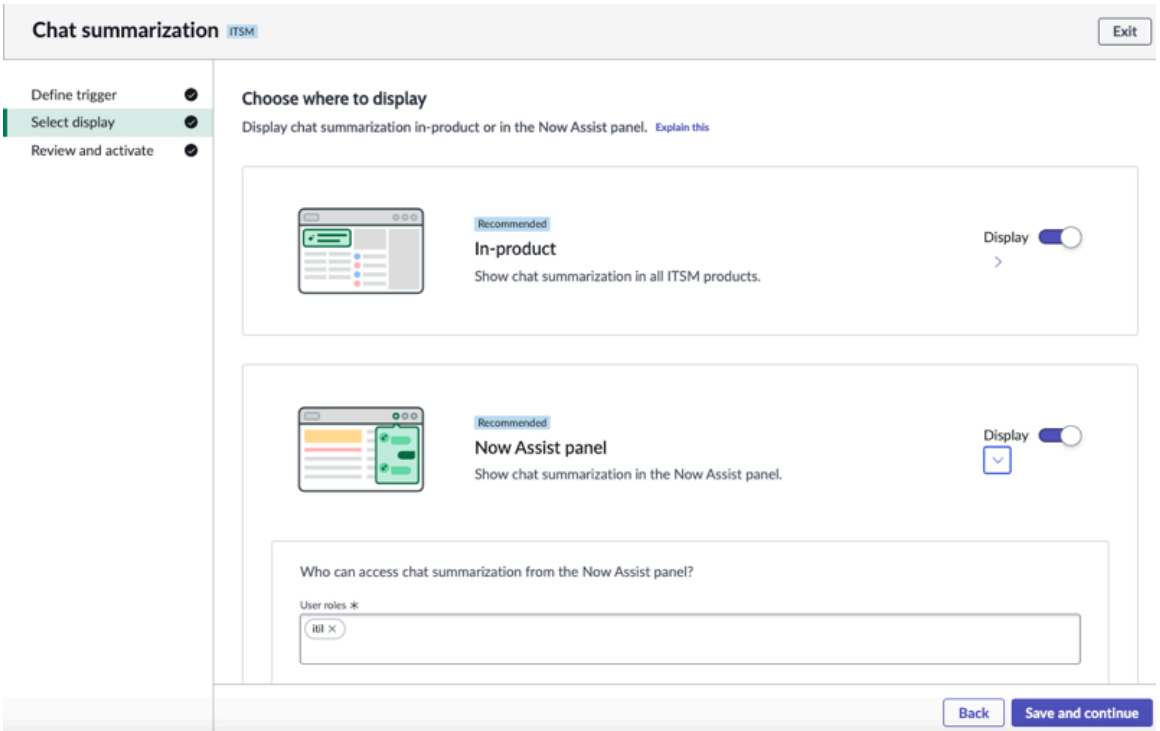


After deciding which skills best fit your business needs, you can activate them from the console. Some skills require configuration so that you can customize the skill to your needs, such as determining the skill inputs and triggers. You can select the skills that you want to configure in the Now Assist Admin Features page.

The following example shows a step in the guided setup of the skill activation process, which is choosing where to display the skill. Skills can be displayed in-product, including the Core UI and Workspaces, or in the Now Assist panel. You can choose either or both locations. The following diagram shows the process for chat summarization activation in Now Assist Admin.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

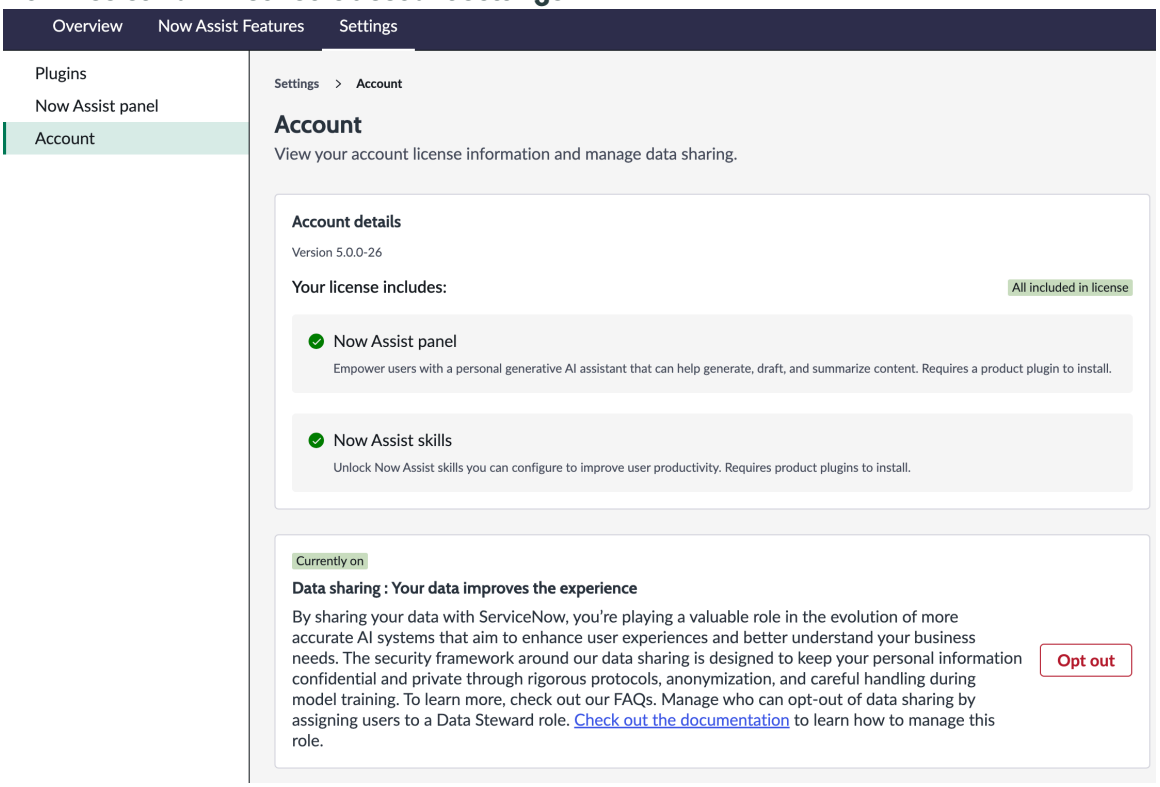
Chat summarization activation in Now Assist Admin



4. Review your Now Assist account settings.

The Now Assist Admin console Settings page enables you to set up language support, if you have Dynamic Translation enabled on your instance, and review your account details. Get up-to-date information about what plugins are available to you and the status of data sharing on your instance.

Now Assist Admin console account settings

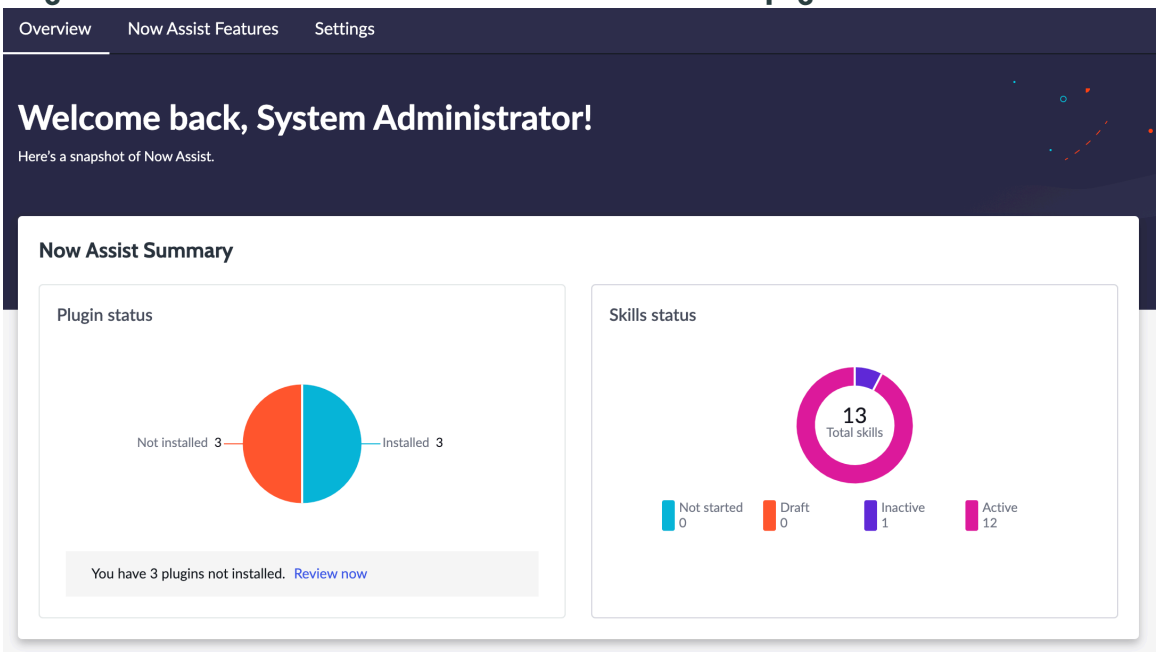


5. Monitor and analyze skill performance.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use the 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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Learn more about the Now Assist journey

- Get an overview of the Now Assist framework.
- Discover the Now Assist panel.
- Explore the Now Assist skills available.
- Activate and configure a Now Assist skill using the Now Assist Admin console.
- Analyze and monitor Now Assist skill usage and performance.

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.

Note:
Conversational aspects of the Now Assist panel, such as skill detection, are powered by Now LLM Service.

Now Assist Overview

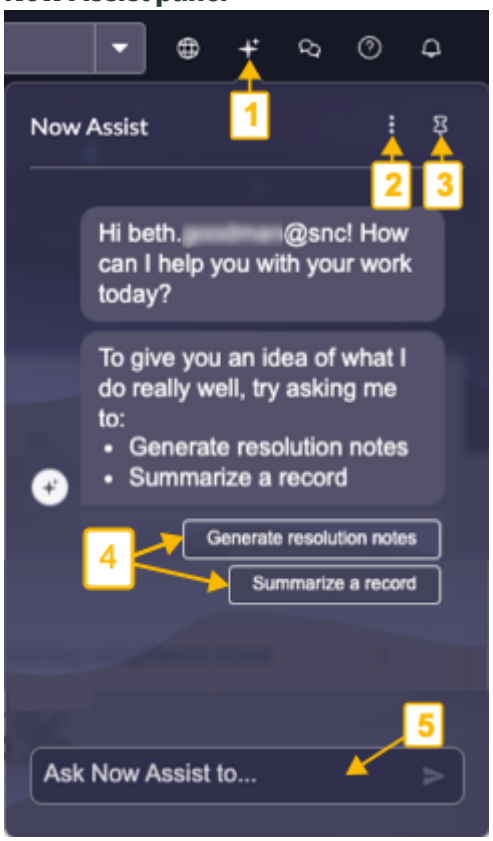
Agents can use the Now Assist panel to interact with and get assistance from generative AI. On the Now Assist panel, you can increase your productivity and efficiency by using the generative AI experience to summarize a chat, case, incident or generate resolution notes.

Note:
Next Experience must be enabled to use the Now Assist panel. For more information, see Considerations for activating Next Experience.

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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist panel



The Now Assist panel includes:

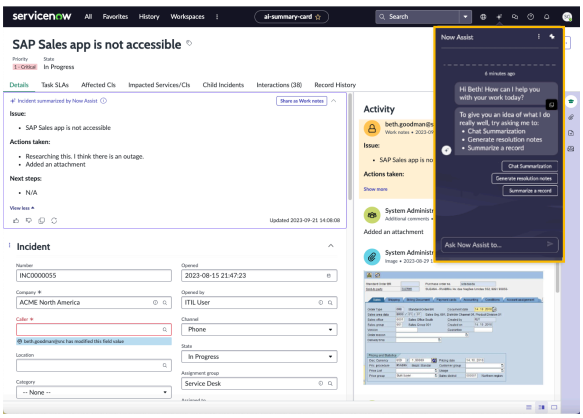
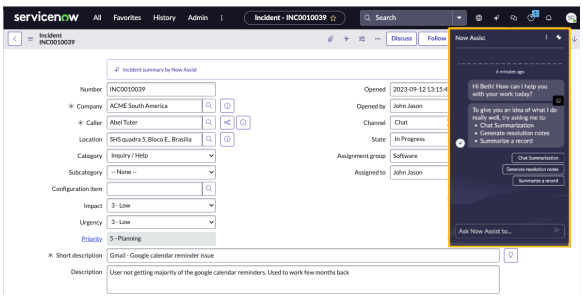
Item number	Description
1 - ✦	Now Assist icon Displays the Now Assist panel. If a number in a square appears, it indicates the number of messages you have missed when the Now Assist panel was closed.
2 - ⋮	Reset icon Select the reset icon if you think that Now Assist doesn't understand what you want 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.
3 -	Pushpin icon Permanently pin the Now Assist panel to the screen.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Item number	Description
4 - Option buttons	Option buttons
5 - Ask Now Assist to..... field	Field where the user can enter actions.

Now Assist is available in both Next Experience and Core UI. The following screenshots show the Now Assist panel on Next Experience and Core UI screens.

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.

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:

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- For 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 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	Available skills	Available products
Technology	• Alert simplification • Chat summarization • Feedback summarization • Incident summarization • KB generation • Resolution notes generation	• • Now Assist for IT Service Management (ITSM) • Now Assist for Strategic Portfolio Management (SPM)
Customer	• Chat summarization • Incident summarization • KB generation • Resolution notes generation • Work order task closure summarization	• Now Assist for Customer Service Management (CSM) • Now Assist for Field Service Management (FSM)
Employee	• Chat summarization • Case summarization	Now Assist for HR Service Delivery (HRSD)

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Available skills by workflow (continued)

Workflow	Available skills	Available products
	• KB generation • Resolution notes generation	
Creator	• App generation • Catalog generation • Code generation • Flow generation • Flow recommendations • Playbook generation	Now Assist for Creator
Platform	• Now Assist Q&A Genius Results • Now Assist Multi-Turn Catalog Ordering • Now Assist Topics	The following plugins are included with some Now Assist products: • Now Assist in AI Search • Now Assist in Virtual Agent

Alert simplification

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 .

Availability

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

Now Assist products and workflows

Workflow	Product
Technology	

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

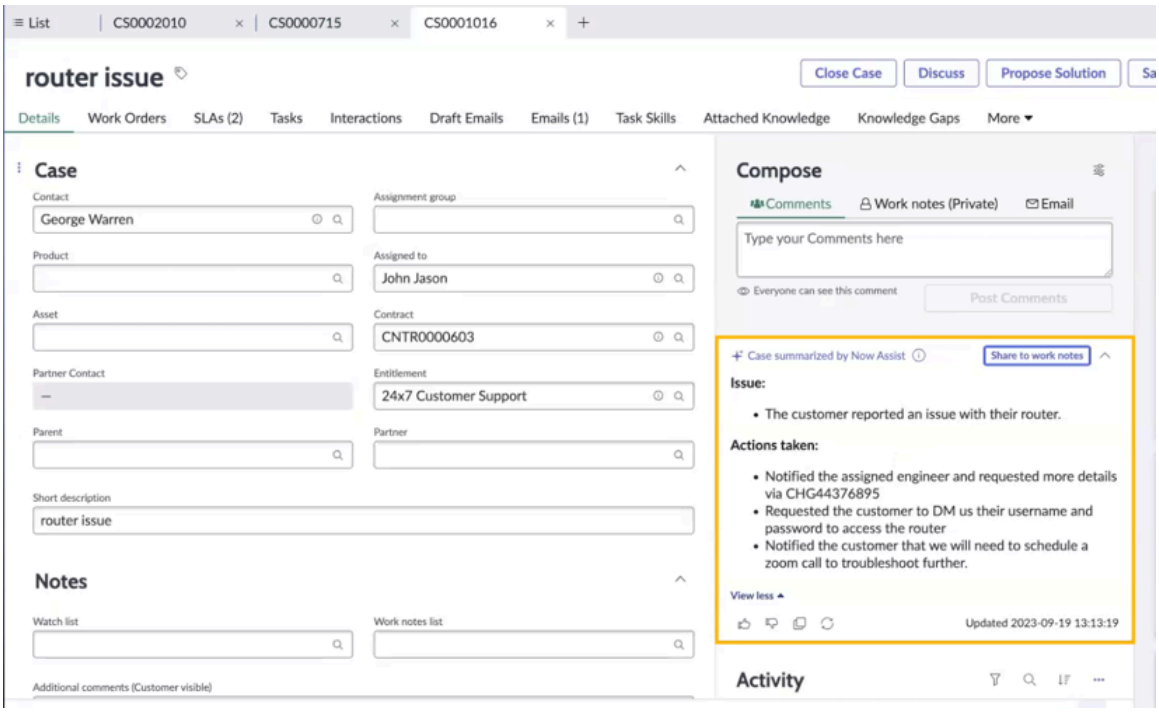
Record summarization overview

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.

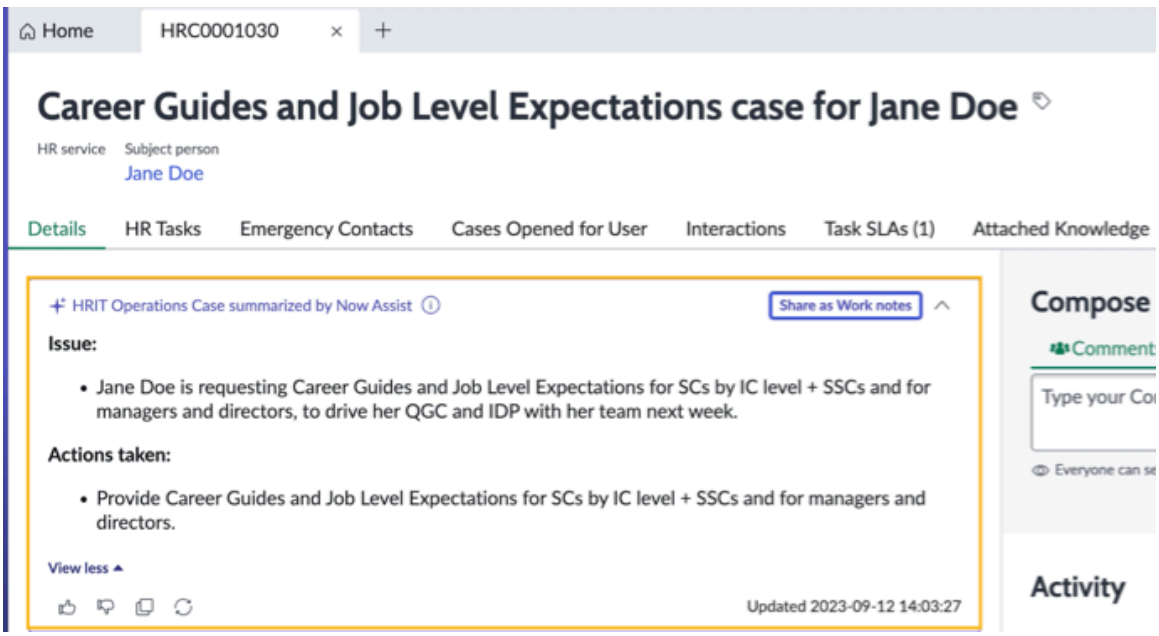
Note:
The ServiceNow® Large Language Model (Now LLM Service) is the provider for this Now Assist skill.

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



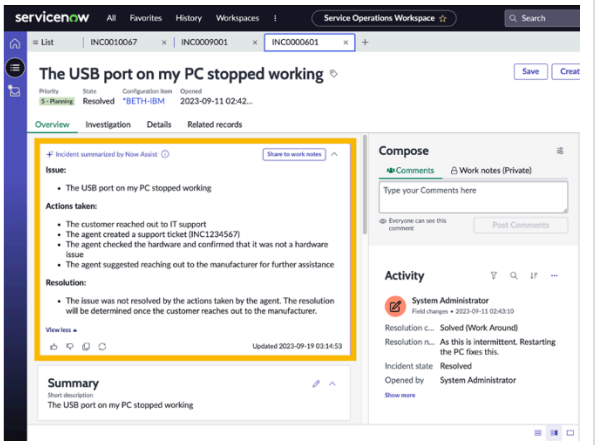
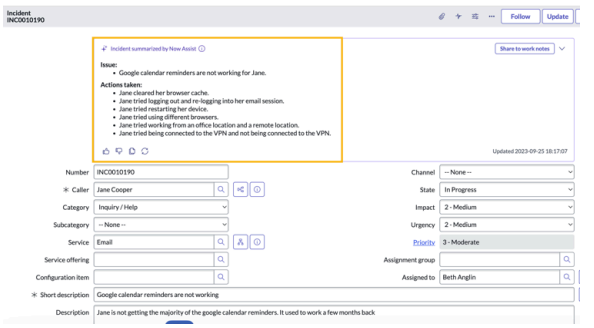
Case summarization in Now Assist for HRSD



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In the Now Assist 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
	

Generating a case or incident summarization

You can generate a case summarization for Now Assist for CSM, Now Assist for HRSD, or Now Assist for ITSM:

- To generate a case summarization for Now Assist for CSM, see Summarize a case by using Now Assist for Customer Service Management (CSM).
- To generate a case summarization for Now Assist for HRSD, see Summarize a case by using Now Assist for HR Service Delivery (HRSD).
- To generate an incident summarization for Now Assist for ITSM, see Summarize an incident 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	Now Assist for IT Service Management (ITSM)
Customer	Now Assist for Customer Service Management (CSM)
Employee	Now Assist for HR Service Delivery (HRSD)

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.

Summarize a chat with Now Assist for Customer Service Management

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

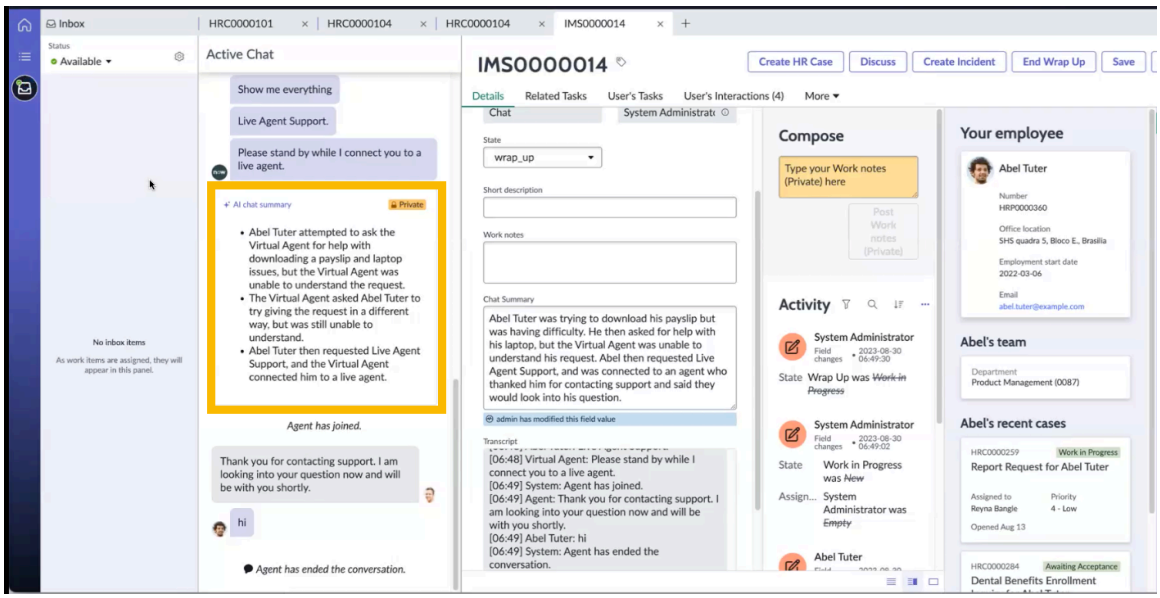
Overview

The chat summarization is 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 for Now Assist for Customer Service Management (CSM), Now Assist for HR Service Delivery (HRSD), or Now Assist for IT Service Management (ITSM).

- To generate a chat summarization for Now Assist for CSM, see Summarize a chat conversation by using Now Assist for Customer Service Management (CSM).
- To generate a chat summarization for Now Assist for HRSD, see Summarize a chat conversation by using Now Assist for HR Service Delivery (HRSD).
- To generate a chat summarization for Now Assist for ITSM, see 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist products and workflows

Workflow	Product
Technology	Now Assist for IT Service Management (ITSM)
Customer	Now Assist for Customer Service Management (CSM)
Employee	Now Assist for HR Service Delivery (HRSD)

Feedback summarization

Summarize the details of the feedback from your customers by using the 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 of your product.

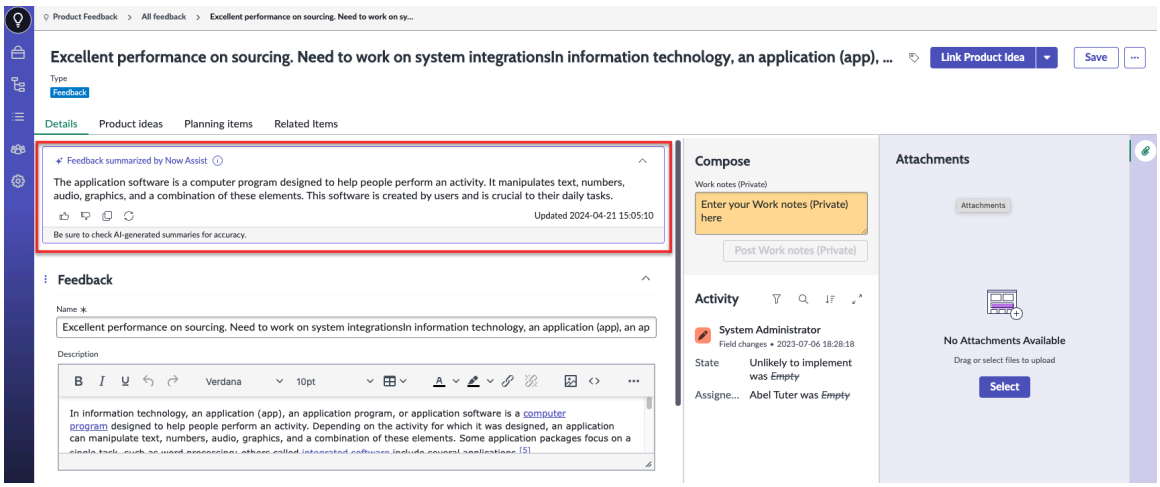
Overview

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 a feedback record so that they can quickly understand the feedback context. With feedback summarization, your product managers can eliminate the need to analyze feedback manually.

Note:
Now LLM Service is the default provider for 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 feedback summarization in Now Assist for Strategic Portfolio Management (SPM).

Feedback summarization in Now Assist for SPM



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).

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Availability

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

Now Assist product and workflow

Workflow	Product
Technology	Now Assist for Strategic Portfolio Management (SPM)

KB generation

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

Overview

Now Assist can create drafts of knowledge articles on how to resolve a case or an incident 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, classic environment, or in the Now Assist panel.

The fields used as inputs are:

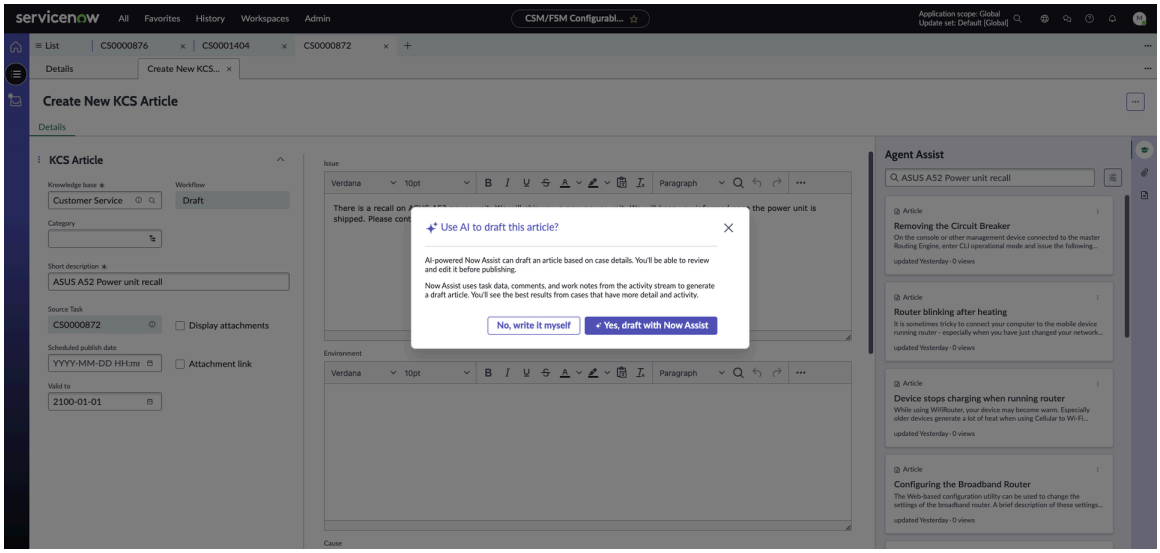
- Short description
- Description
- Resolution notes
- Close notes
- Work notes
- Comments

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

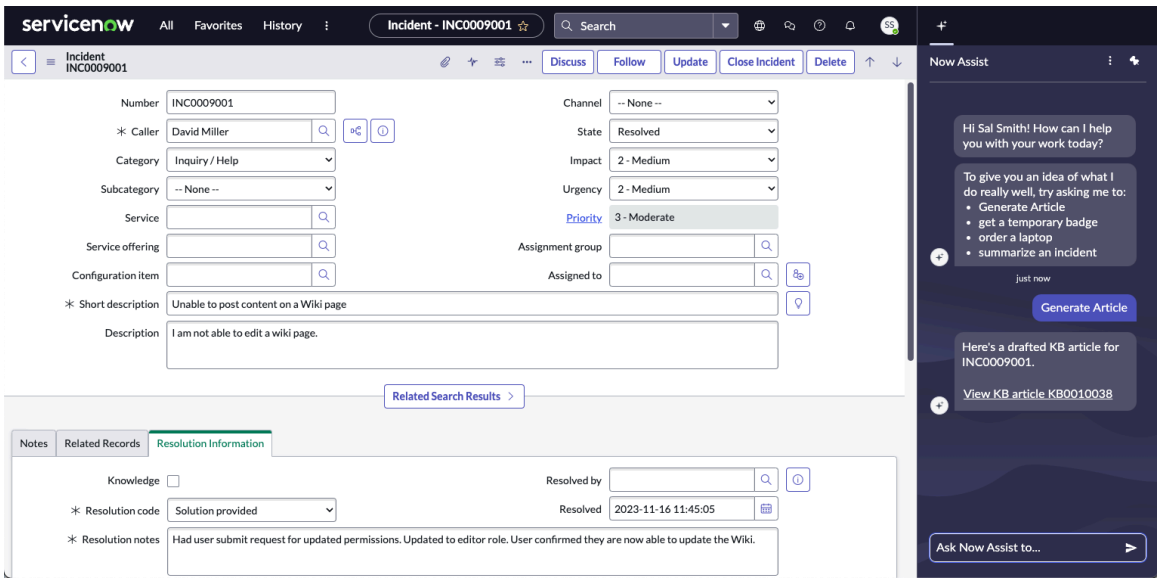
Once the Now Assist application is installed, the agent has the option to use Now Assist to generate a knowledge article from a case or an incident.

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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



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



Generating knowledge articles

You can generate knowledge articles in Now Assist for Customer Service Management (CSM), Now Assist for IT Service Management (ITSM), and Now Assist for HR Service Delivery (HRSD) in several ways:

- To generate knowledge articles based on cases in CSM Configurable Workspace, see Generate a knowledge article from the CSM Configurable Workspace and classic environment with Now Assist.
- To generate knowledge articles based on incidents in Service Operations Workspace for ITSM, see Generate a knowledge article from the Service Operations Workspace for ITSM and classic environment with Now Assist.
- To generate knowledge articles based on cases in Agent Workspace for HR Case Management, see Generate a knowledge article from Now Assist for HR Service Delivery (HRSD).

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- To generate knowledge articles based on cases or incidents in the classic environment, see Generate a Knowledge article from the classic environment with Now Assist.
- To generate knowledge articles based on cases or incidents from the Now Assist panel, see 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	Now Assist for IT Service Management (ITSM)
Customer	Now Assist for Customer Service Management (CSM)
Employee	Now Assist for HR Service Delivery (HRSD)

Resolution notes generation

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

Overview

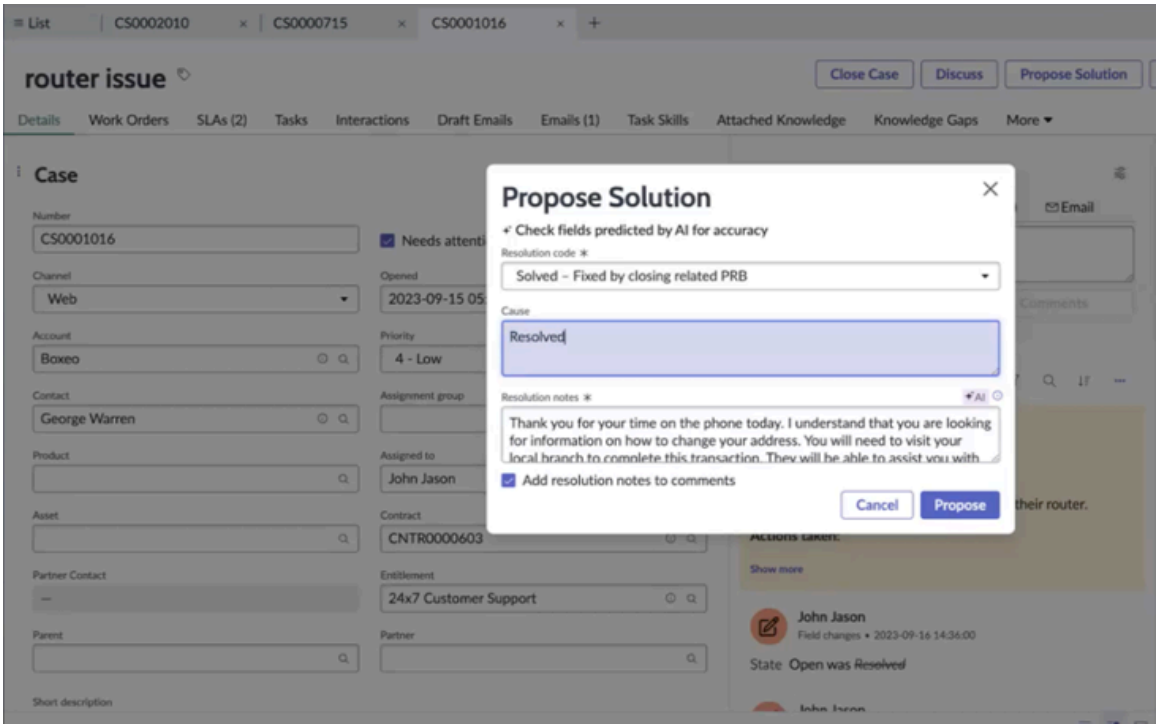
Generating the resolution notes creates a condensed version of the incident's resolution notes that are generated by Now Assist. The generated resolution notes can provide information to other agents who might encounter similar interactions.

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

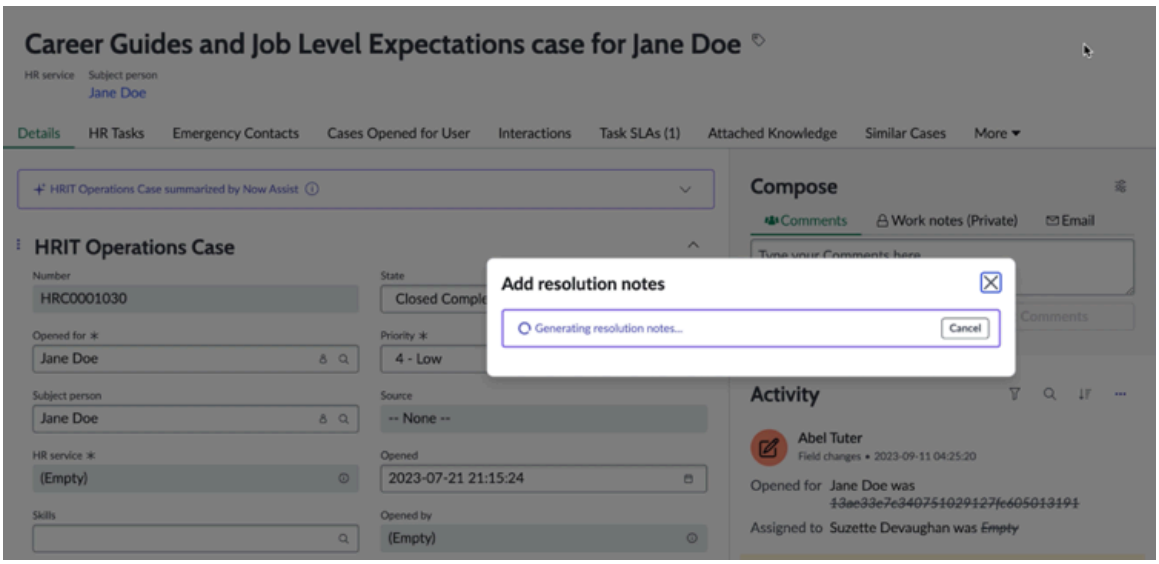
The following examples show Now Assist generated resolution notes in Now Assist for Customer Service Management (CSM), Now Assist for HR Service Delivery (HRSD), and Now Assist for IT Service Management (ITSM).

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Generated resolution notes in Now Assist for CSM

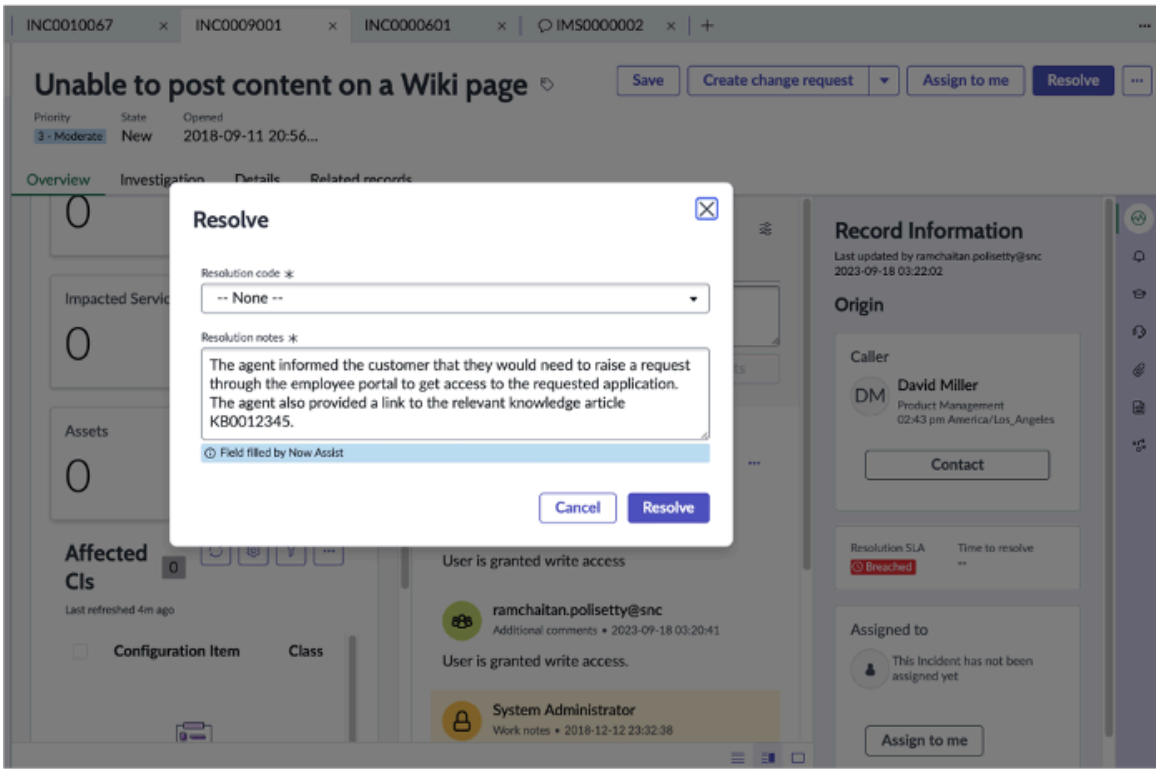


Generated resolution notes in Now Assist for HRSD



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Generated resolution notes in Now Assist for ITSM



Generating resolution notes

You can generate resolution notes for Now Assist for CSM, Now Assist for HRSD, or Now Assist for ITSM:

- To generate resolution notes for Now Assist for CSM, see Generate the resolution notes for a case by using Now Assist for Customer Service Management (CSM).
- To generate resolution notes for Now Assist for HRSD, see Generate the resolution notes for a case by using the Now Assist for HR Service Delivery (HRSD).
- To generate resolution notes for Now Assist for ITSM, see Generate the resolution notes for an incident 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	Now Assist for IT Service Management (ITSM)
Customer	Now Assist for Customer Service Management (CSM)
Employee	Now Assist for HR Service Delivery (HRSD)

Work order task closure summarization

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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

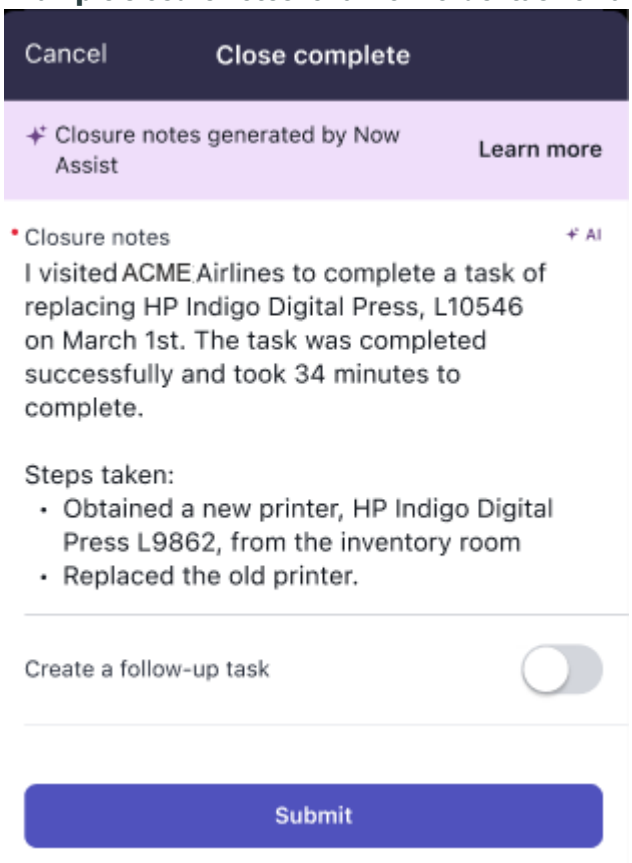
Overview

Now Assist can generate a work order task summary only if the task has at least 200 words in the fields that are used for the input data. The 200-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



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	Now Assist for Field Service Management (FSM)

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see Now Assist for Creator.

App generation

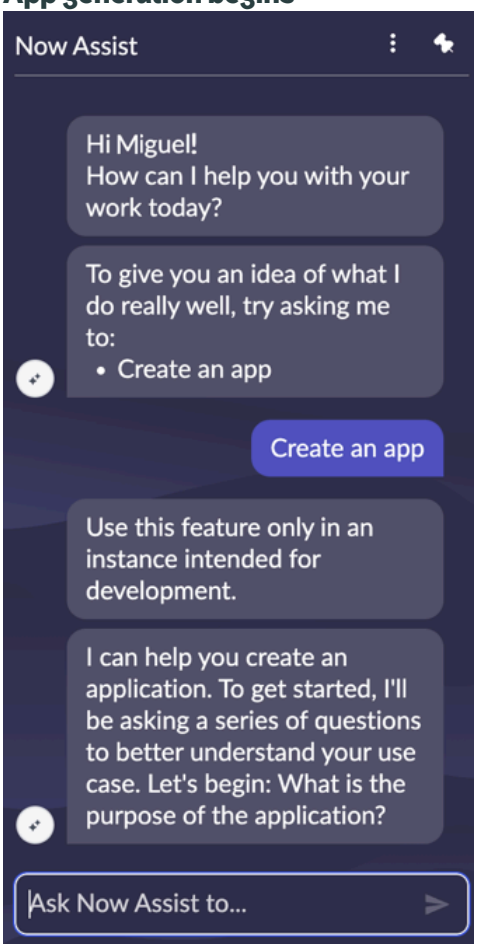
App generation enables developers to kickstart application building through conversation.

Overview

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 Generate apps with Now Assist for Creator.

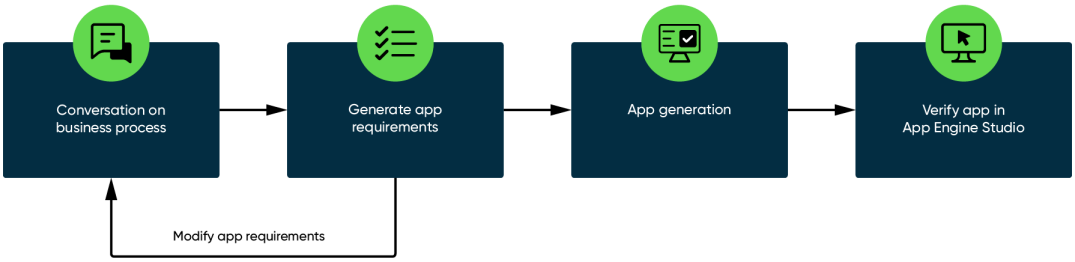
App generation begins



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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

App generation workflow



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.

The developers can open and verify everything that is generated in App Engine Studio, which is available with Now Assist for Creator. The developer can modify the app to make it suit your organization's needs. For example, the app'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.

Availability

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

Now Assist products and workflows

Workflow	Product
Creator	Now Assist for Creator

Catalog item generation

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

Overview

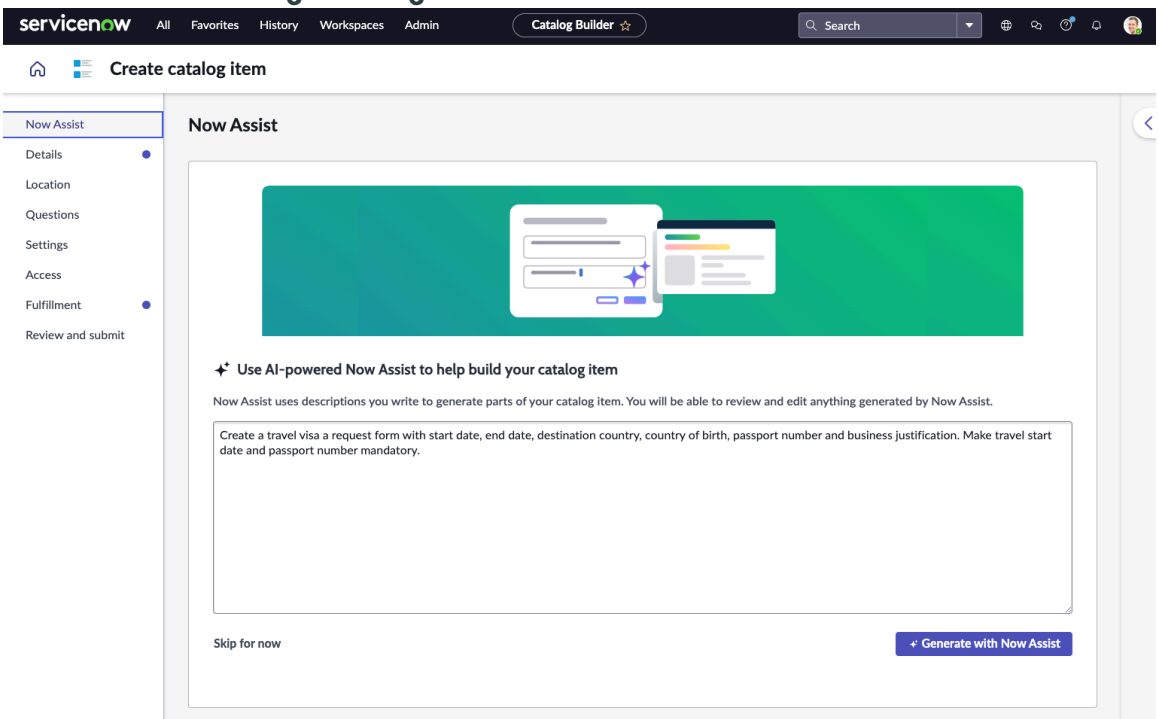
The Catalog item generation skill helps reduce the time needed to create catalog items. While creating a catalog item, you can describe the catalog item you want, and Now Assist generates the catalog item for you.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see 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	Now Assist for Creator

Related topics

Code generation

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.

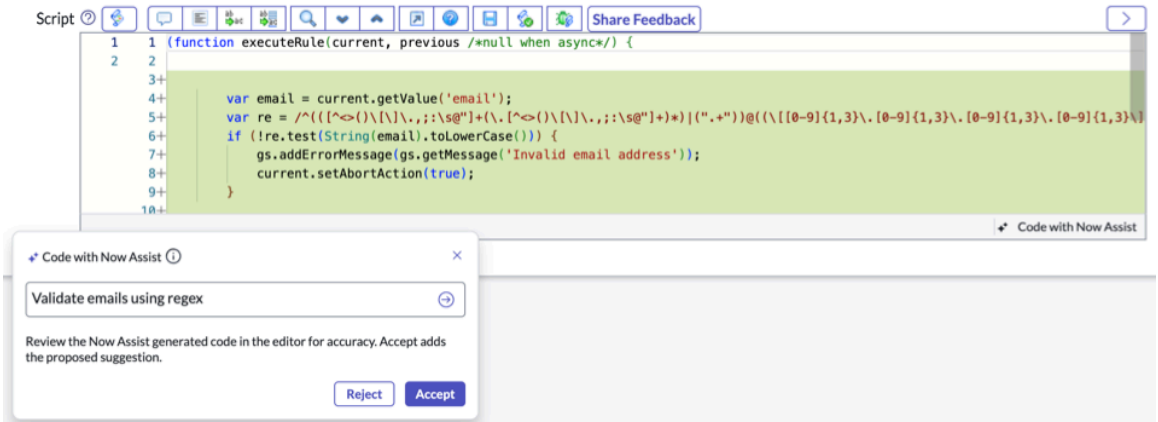
Overview

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 Flow Designer. 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To 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	Now Assist for Creator

Flow generation

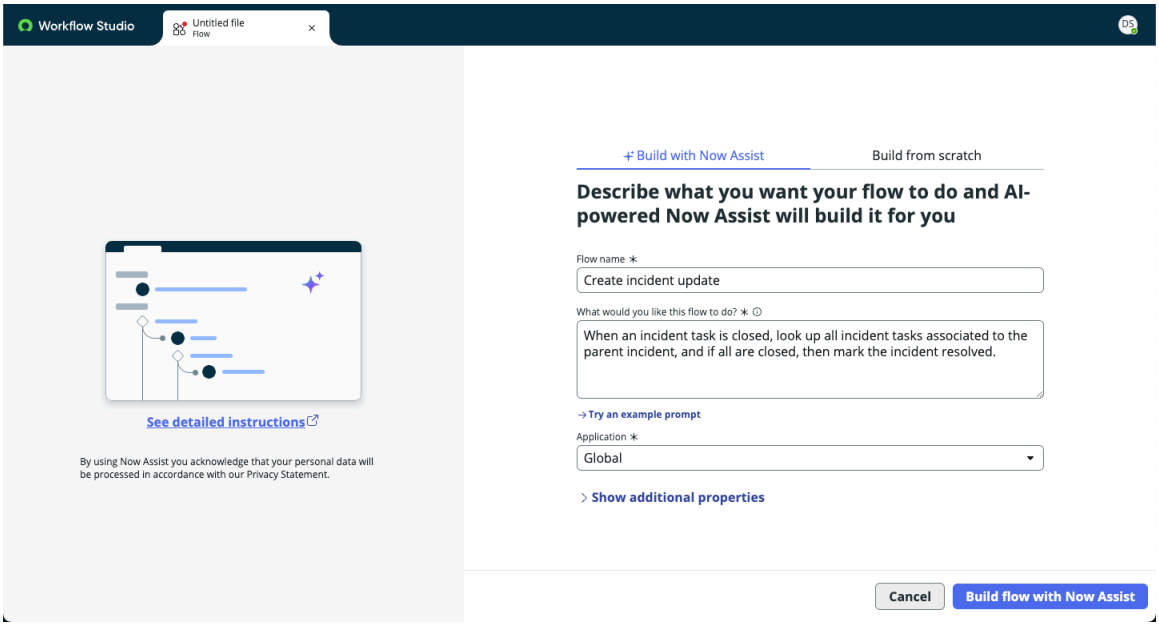
Use the flow generation skill in Flow Designer to create new flows.

Overview

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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use 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	Now Assist for Creator

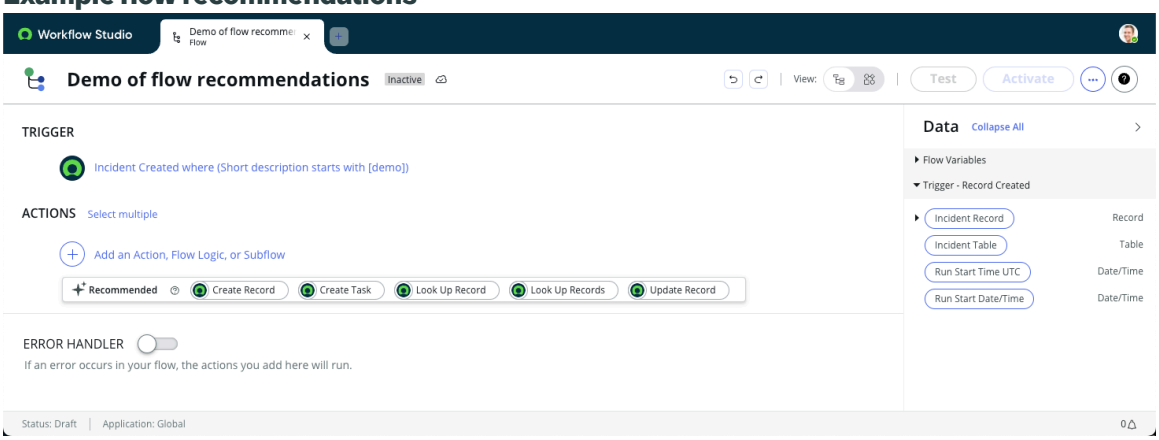
Flow recommendations

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

Overview

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



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see Flow recommendations.

Availability

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

Now Assist products and workflows

Workflow	Product
Creator	Now Assist for Creator

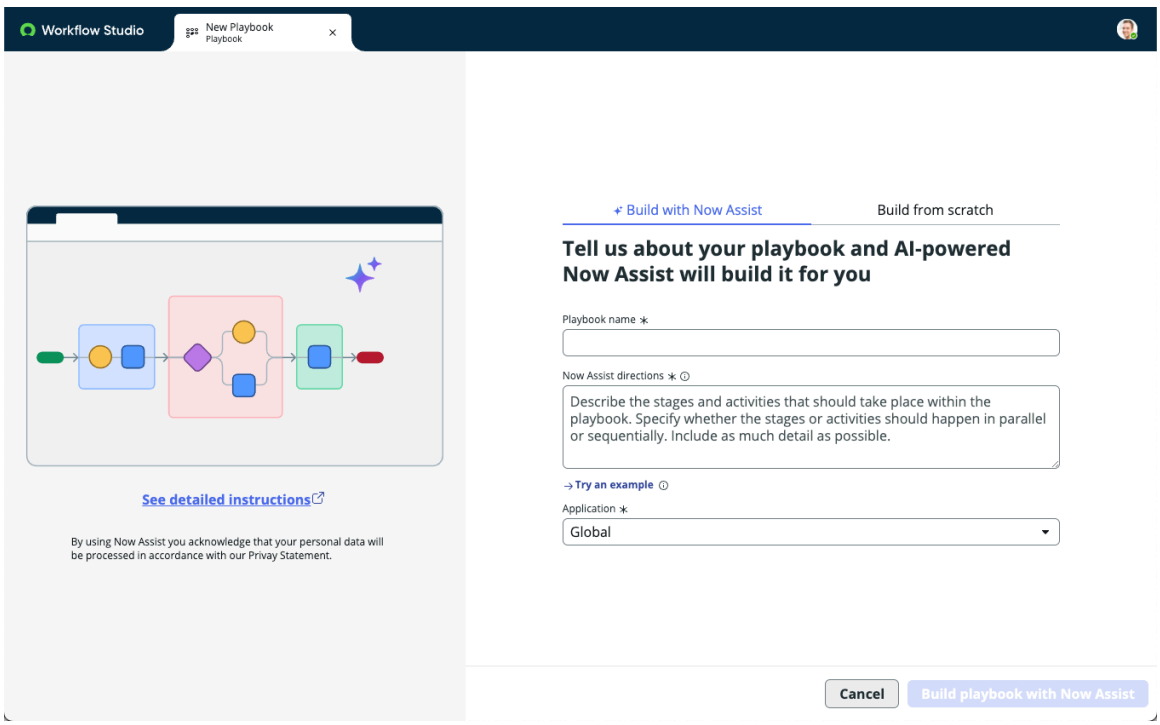
Playbook generation

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.

Overview

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Now Assist products and workflows

Workflow	Product
Creator	Now Assist for Creator

Now Assist skills in the Platform workflow

Some Now Assist generative AI products include skills for the conversational user experience.

For more information, see Now Assist in Virtual Agent and Now Assist in AI Search.

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.

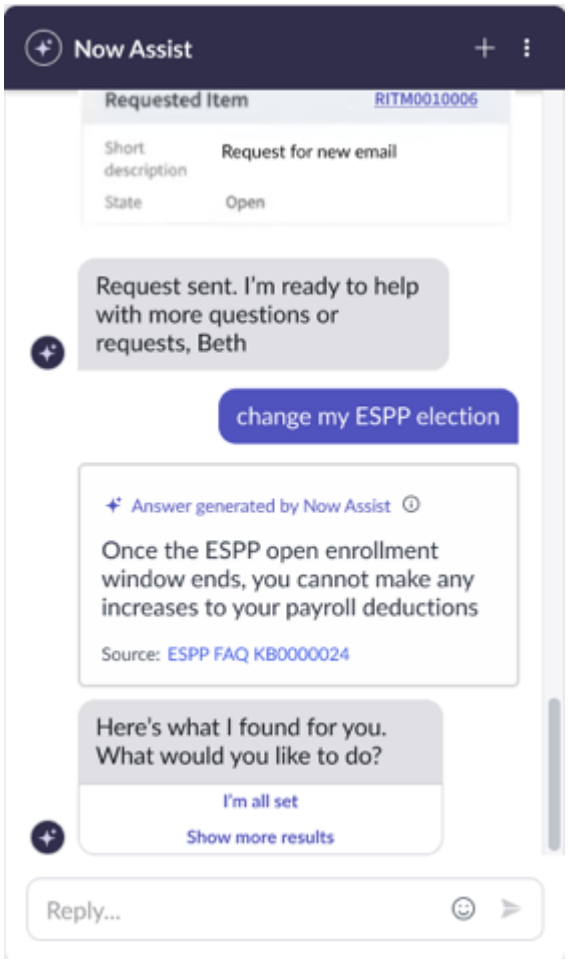
The following skills require Now Assist in AI Search and Now Assist in Virtual Agent. Both are available in some Now Assist products.

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.

Now Assist Q&A Genius Results

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



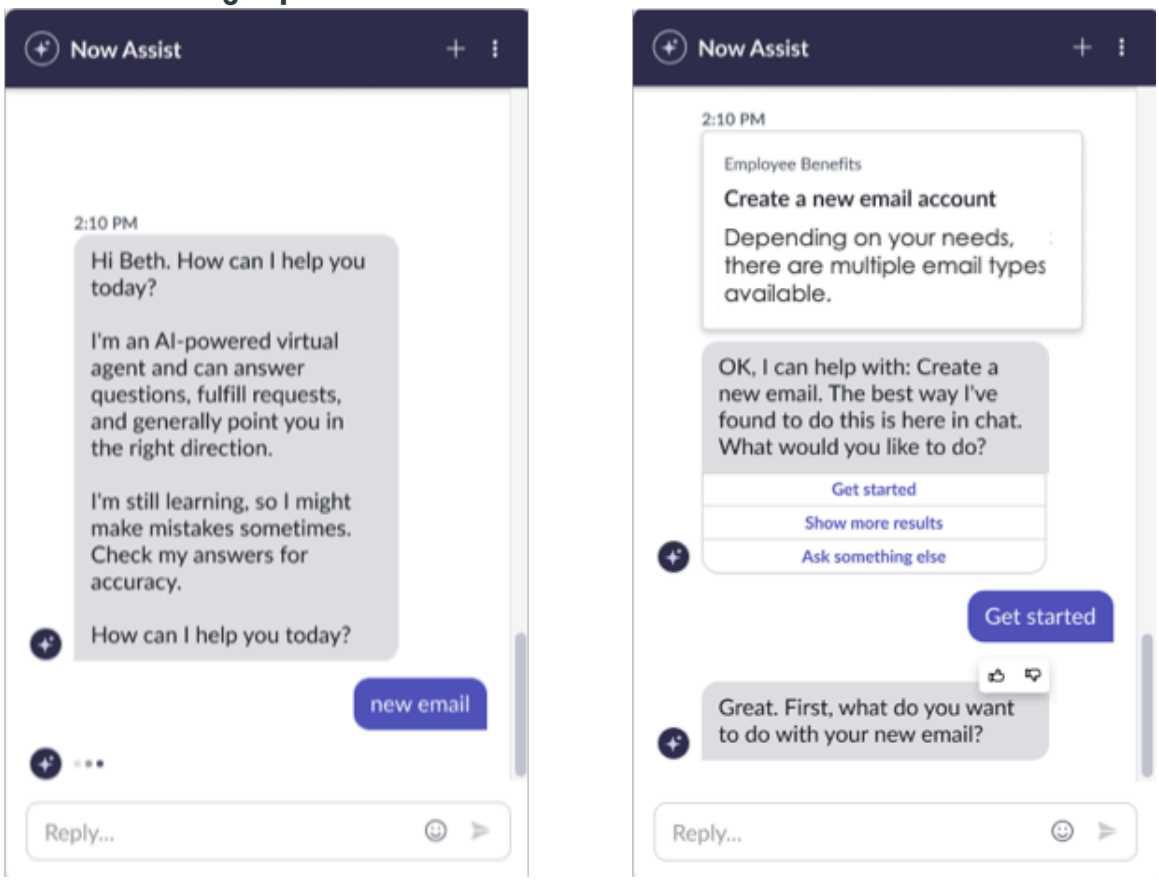
For more information, see Now Assist in AI Search.

Now Assist Multi-Turn Catalog Ordering

The Now Assist Multi-Turn Catalog Ordering skill gives users 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

New email catalog request from the chat window



Now Assist Topics

The Now Assist Topics skill lets Virtual Agent admins create topics that use large language model (LLM) discovery. 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

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.

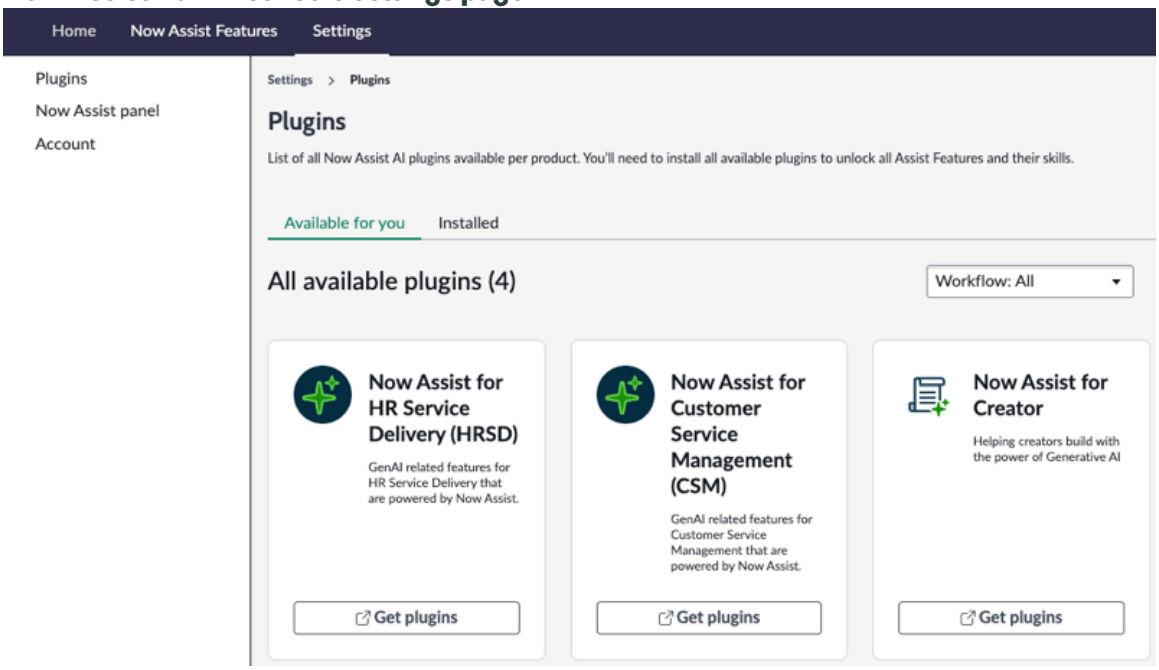
Configuration overview

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The following example shows 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



Install Now Assist plugins

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

Before you begin

Role required: admin

About this task

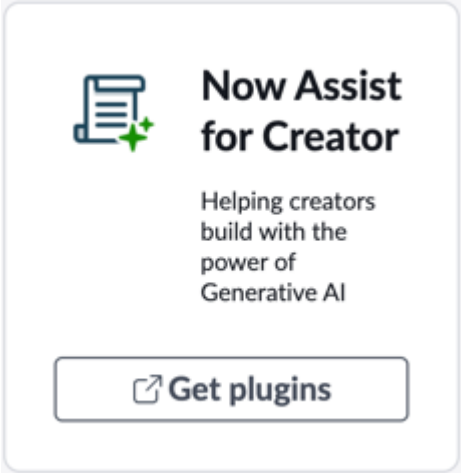
To get started with Now Assist, you must install at least one Now Assist application on your instance. The Now Assist Admin console can guide your implementation, starting with installation. Check out the Now Assist Journey Checklist for more information.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



3. Select **Get plugins** on the card for the plugin that you want to install.
4. In the confirmation window, select **Install Plugin** to open the ServiceNow Store page for the plugin in a new browser tab.
5. Install the plugin from the ServiceNow Store page.
   Some applications may require you to request the app from the ServiceNow® Store first. After you've requested the application from the ServiceNow® Store page, navigate to **All** > **System Applications** > **All Available Applications** > **All** to finish the installation.
6. Return to the Now Assist Admin console.
7. In the dialog box, select **Refresh**.
8. Either close the dialog box to review all of your plugins or select **View all <Plugin> Assists and Skills** to review the features of your new plugin.

Result
Your plugin is successfully installed.

If you encounter issues installing or updating applications, see this knowledge article for steps that may address your issue. Otherwise, you can make a Support case.

What to do next
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 summarization, in a side panel on the user interface.

Before you begin
You must install at least one Now Assist application before you can turn on the Now Assist panel.

Role required: admin

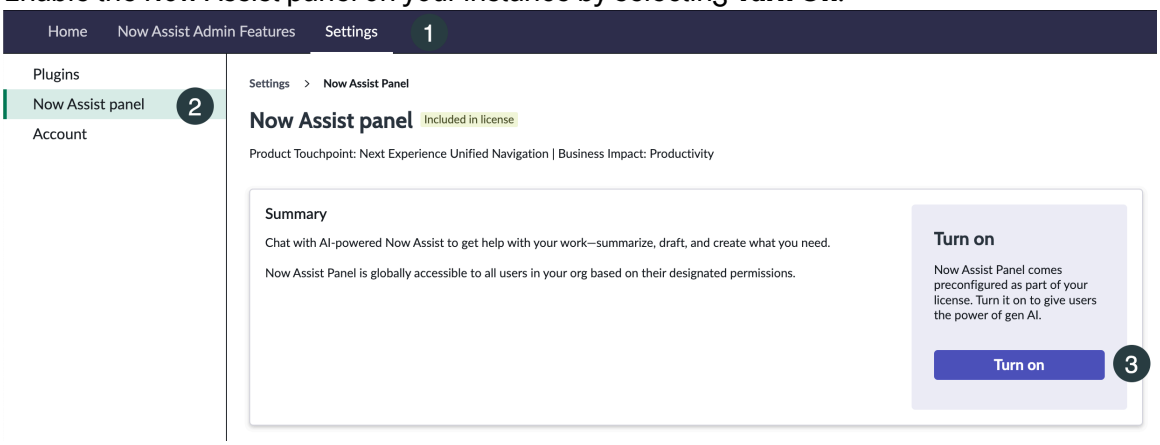
About this task
You must have the now_assist_panel_user role to have access to the Now Assist panel once you turn it on.

To learn more about the Now Assist panel, and how it can assist your agents, see Now Assist panel.

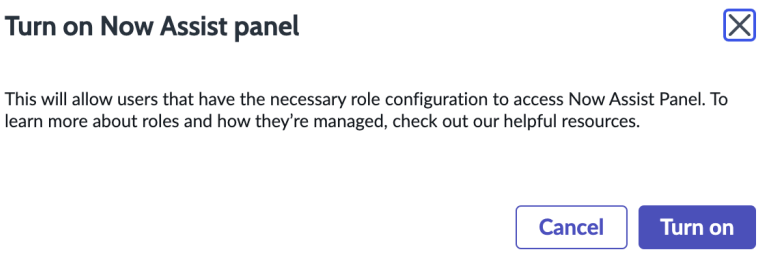
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Settings**.
   If you're already in Now Assist Admin, select the **Settings** 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**.



Result
The Now Assist panel is active on your instance for users with the now_assist_panel_user role.

What to do next
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: 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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

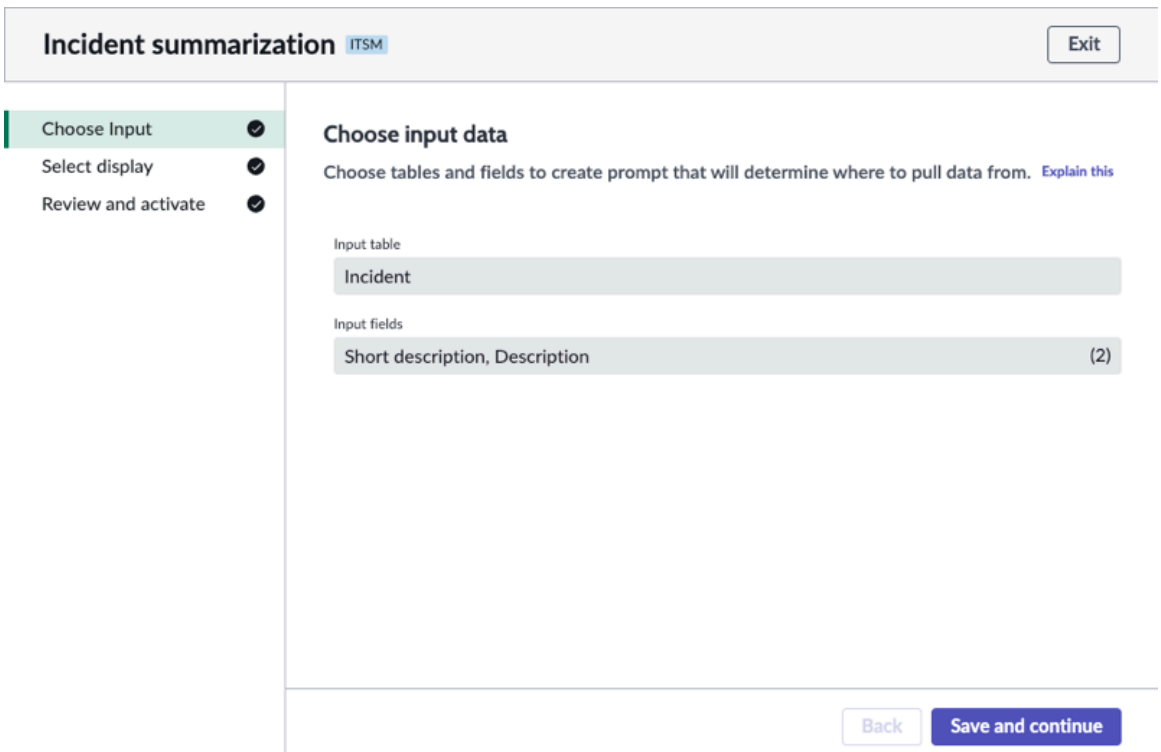
with record or chat summarization or resolution note generation. Turn on the skills that are most relevant to your use cases and business needs.

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.

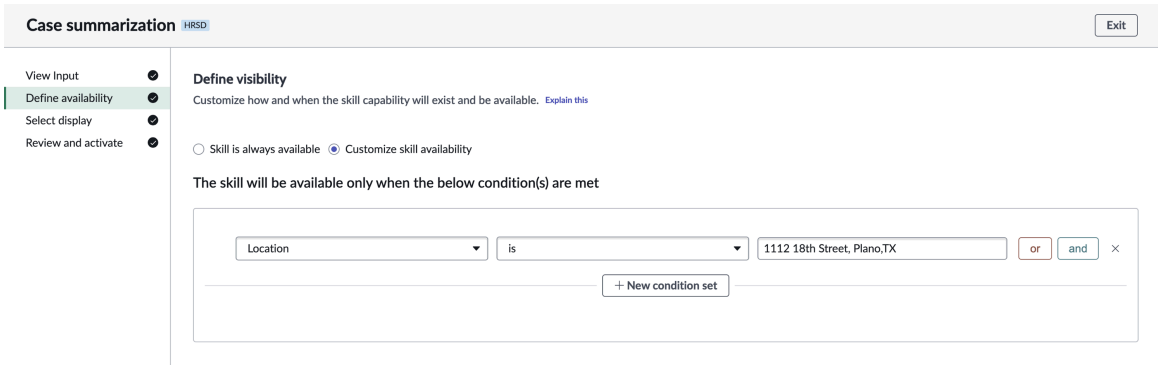
   Each skill configuration has steps that are shown in the guided setup. The exact steps vary from skill to skill. A symbol next to each step indicates whether the step is completed, partially completed, or not completed. After configuring a step, select **Save and continue** to go to the next step. Return to a previous step by selecting **Back**.



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.
   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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



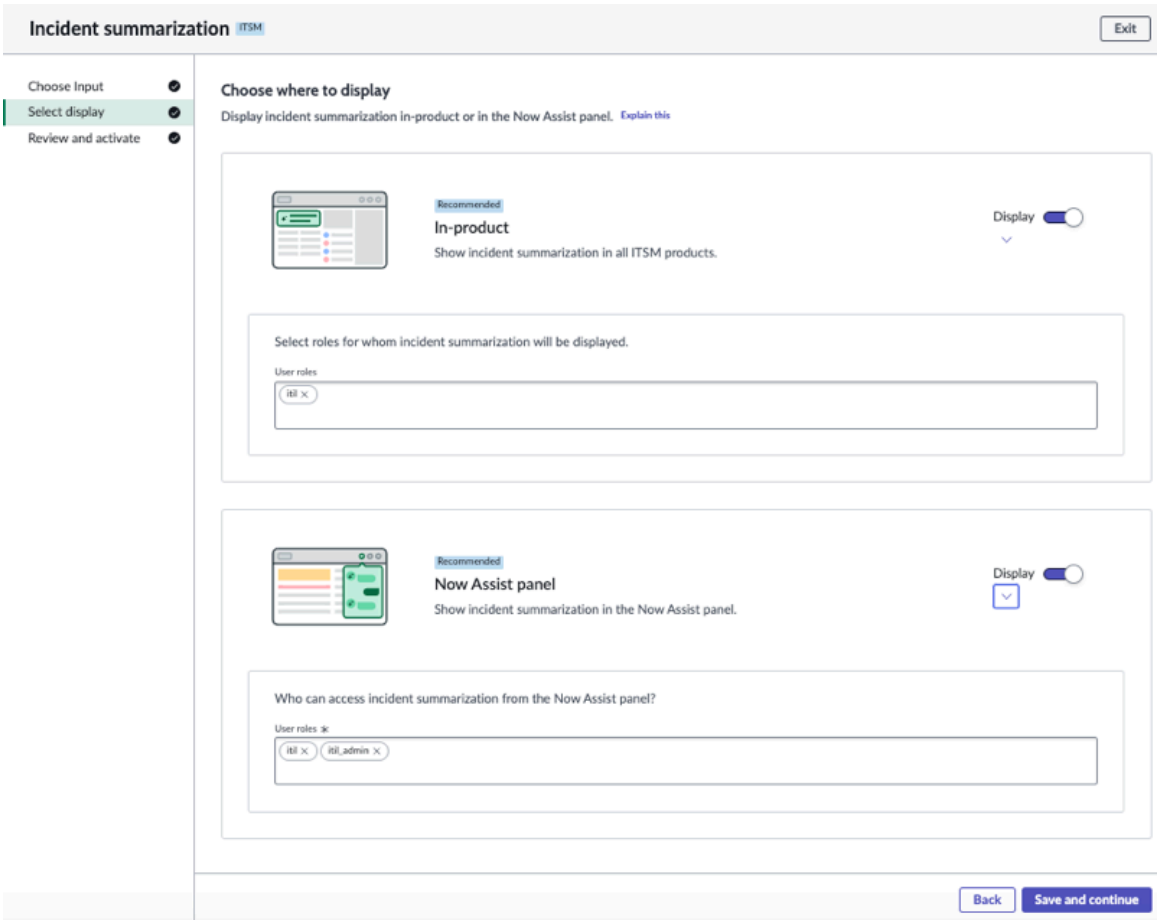
8. In the next step of the skill configuration, select where you'd like to display the skill. You can select both in-product, Now Assist panel, or both.
   - **In-product**: When selected, Now Assist skills are displayed on forms and workspaces.

     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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



9. 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: admin

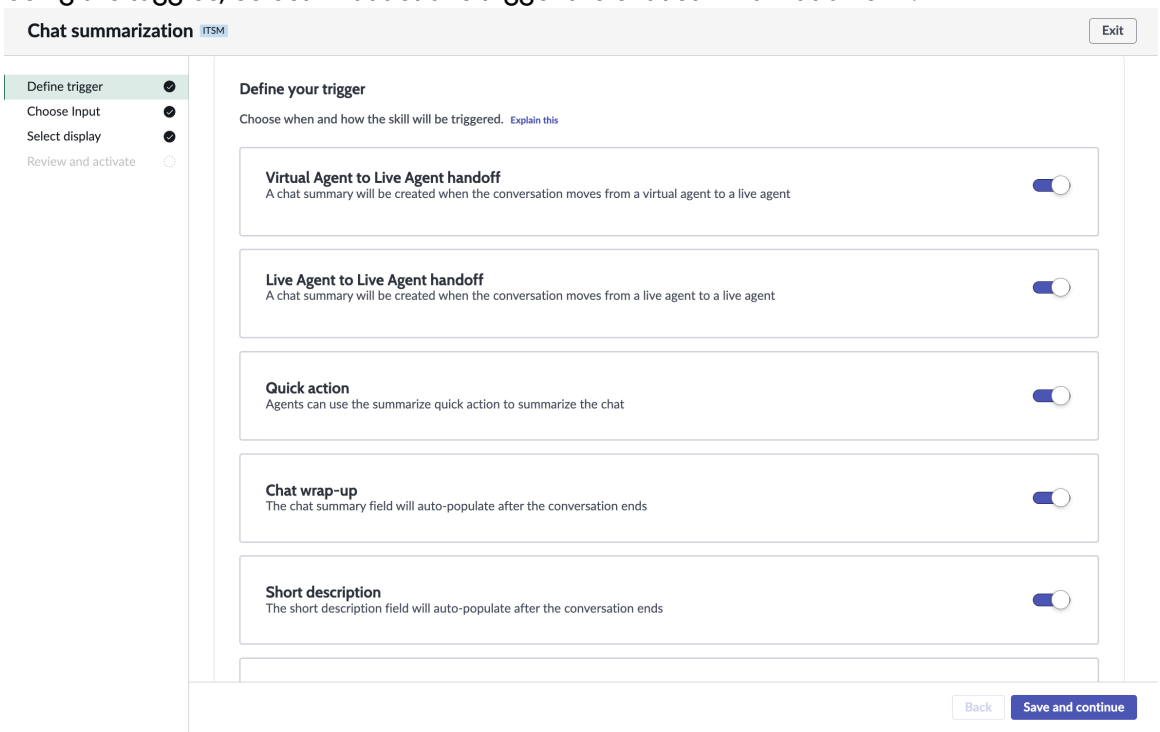
Procedure

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**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Each workflow contains feature sets.

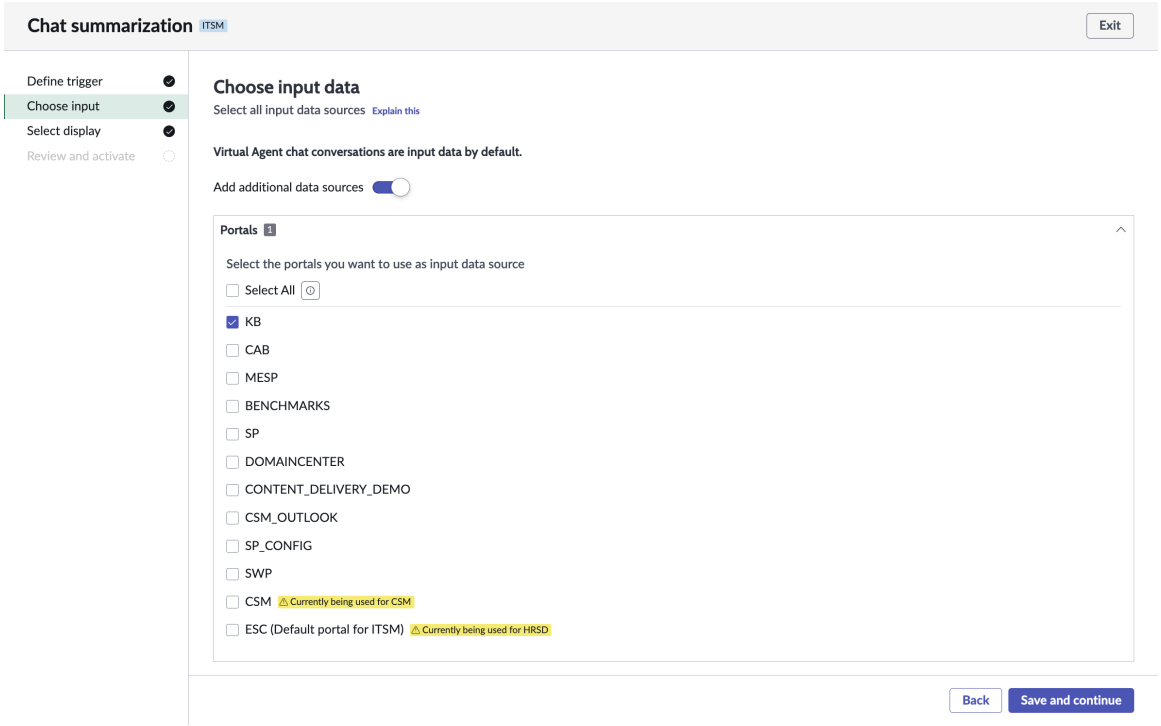
3. On the Chat feature card, select **View details**.
4. In the All available Chat skills section of the chat summarization card, select **Activate Skill**.
5. Go to **Define Trigger**, the first step in the guided setup.
   By default, many of the options in the setup are configured for the most common 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can use additional data sources to add context to the chat conversation. By default, the summary takes the information from the conversation as the source of the input data.

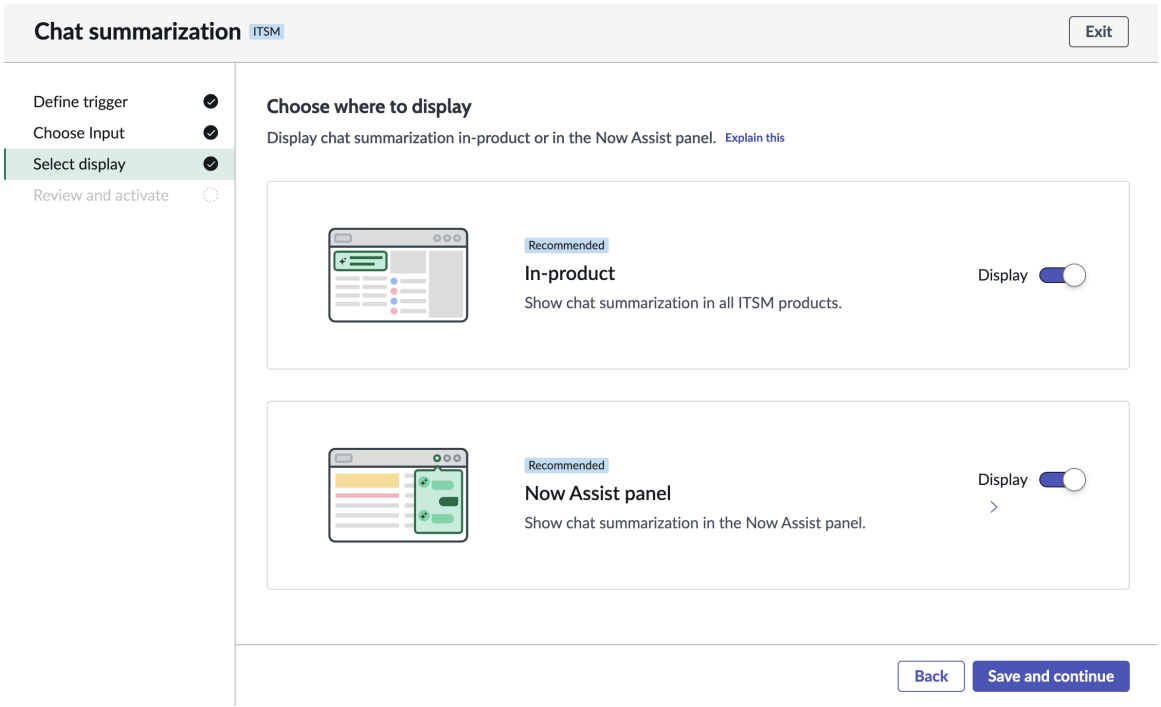


10. Select **Save and continue**.
11. Go to **Select display**, the last step, and select where you would like to display the skill. You can select both in-product, Now Assist panel, or both.
    - **In-product**: When selected, Now Assist skills are displayed on forms and workspaces.
    - **Now Assist 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.

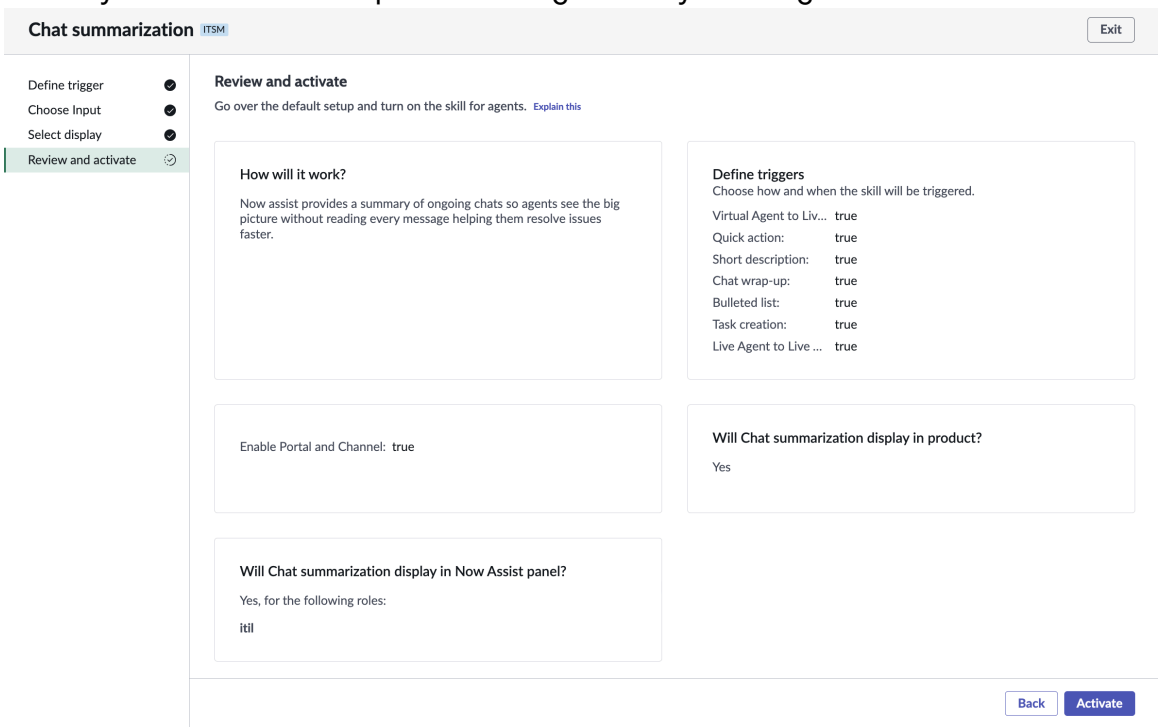
      For the skills that appear in the Now Assist panel, select the down arrow to identify the roles that can use the skill.

    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 have at least one role specified, but you can add as many roles as you like.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



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

Enable chat summarization for chat messages from Microsoft Teams

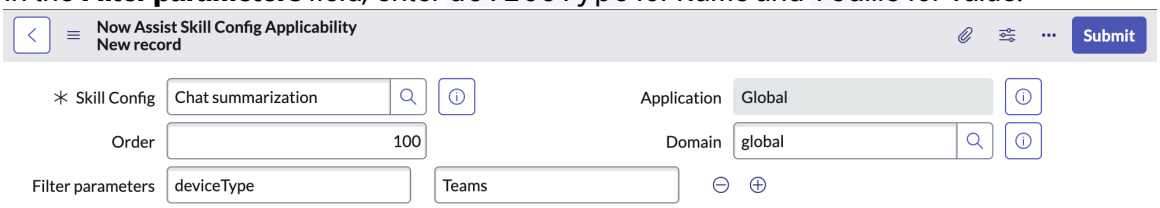
Enable chat summarization for messages sent in Microsoft Teams.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: nsa_admin

Procedure

1. Go to the Now Assist Skill Config Applicability (sn_nowassist_skill_config_applicability) table by entering `sn_nowassist_skill_config_applicability.list` in the navigator.
2. Select **New** to create a record.
3. In the `Skill config` field, enter `Chat summarization`.
   If you have multiple Now Assist applications installed, there might be more than one record called Chat summarization. If you want to check the application of the skill config, select the record, and then use the record preview icon ( ) to look at the **Application** field value, such as Now Assist for IT Service Management (ITSM).
4. In the **Filter parameters** field, enter `deviceType` for Name and `Teams` for Value.



5. Select **Submit** to create the record.
6. **Optional:** If you want to configure chat summarization for all of your Now Assist applications, repeat the process for each one.

Result
Chat summarization is enabled for Microsoft Teams messages for the Now Assist application that you configured.

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.

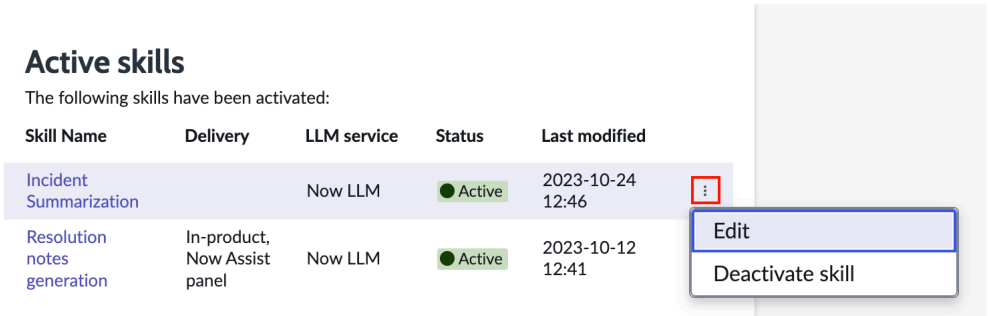
Before you begin
Role required: admin

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**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

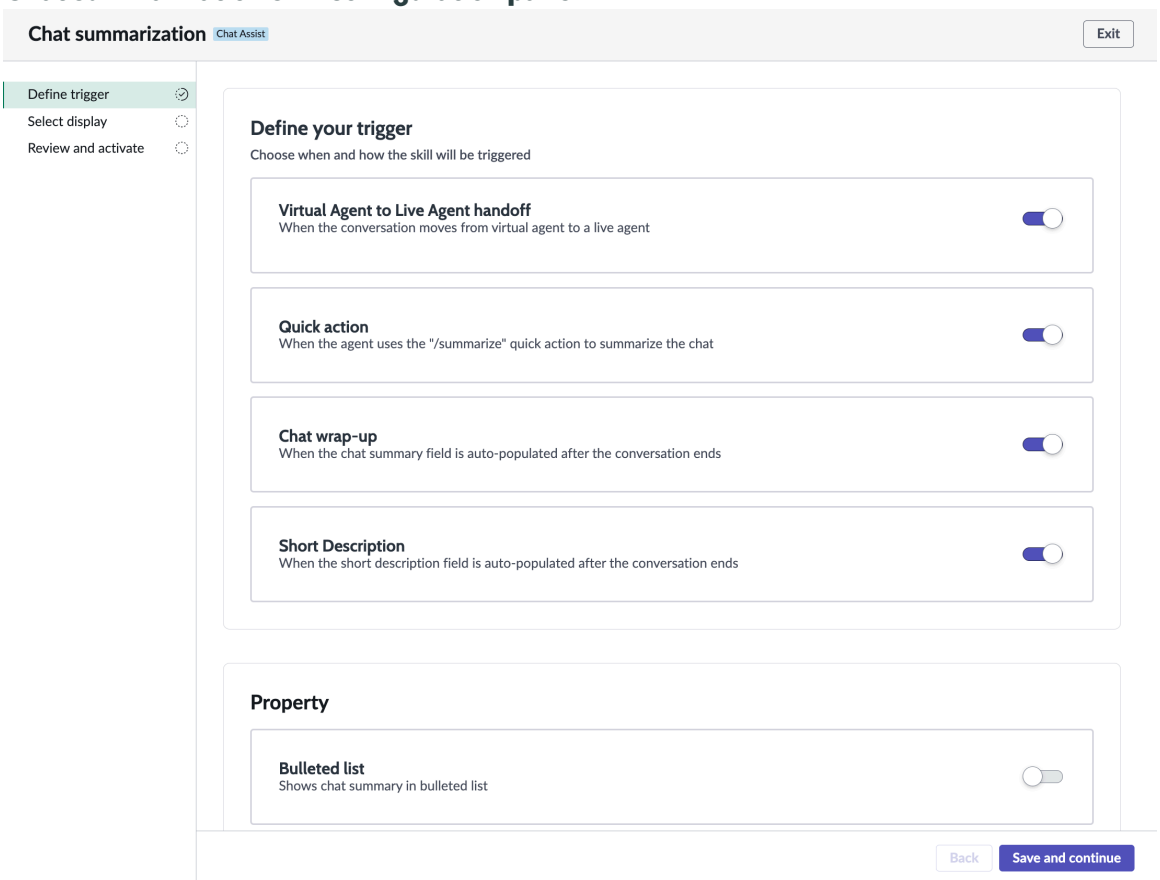


The first step in the guided setup for the skill is displayed.

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



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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Make a copy of a Now Assist skill

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: 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.

Procedure

1. Navigate to **All** > **Now Assist Admin** > **Features**.
   If you're already in the Now Assist Admin console, select the **Now Assist Features** tab.
2. In the navigation pane, select the workflow of the skill that you want to copy, such as Technology or Customer.
3. On the feature card that contains the default skill, select **View details**.
4. In the All available skills or Active skills section, select the more options icon next to the skill that you want to make a copy of and select **Make a copy**.

   Note:
   Only one version of a skill can be active at a time. If you create and activate a copy of the skill, any previously activated version of the skill is deactivated.

5. In the modal, select **Make a copy**.

Result

A copy of the skill is generated and you're taken to the guided setup.

What to do next

Continue the steps in the guided setup to activate the skill. For more information, see Activate a Now Assist skill.

If you're making a copy of the case or incident summarization skill and would like to learn more about your options, see the documentation for configuring record summarization.

Configure case or incident summarization in the Now Assist Admin console

Configure case or incident summarization by using the guided setup in the Now Assist Admin console. You can choose the input tables and fields as well as customize the prompt output for copies of the record summarization skills.

https://player.vimeo.com/video/996395898?h=e609c55303&badge=0&autopause=0&player_id=0&app_id=58479

Before you begin

You can only customize the input data and prompt output for a copy of a record summarization skill. To learn more about making a skill copy, see Make a copy of a Now Assist skill. After you create a skill copy, you can learn the steps to complete the skill setup here.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: nsa_admin

About this task

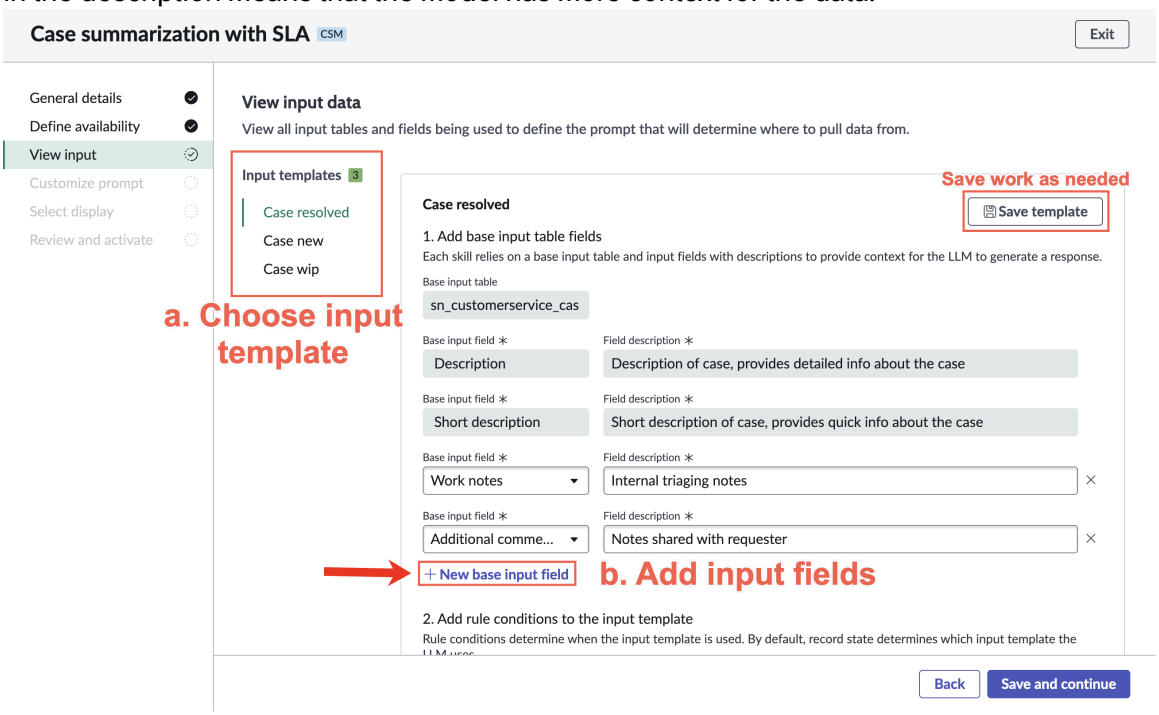
By default, many settings for Now Assist record summarization are optimized for general use cases. Review your goals for incorporating generative AI on your instance to determine whether you want to make changes and what those changes are. After you have made a plan, you can create a copy of a skill 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.

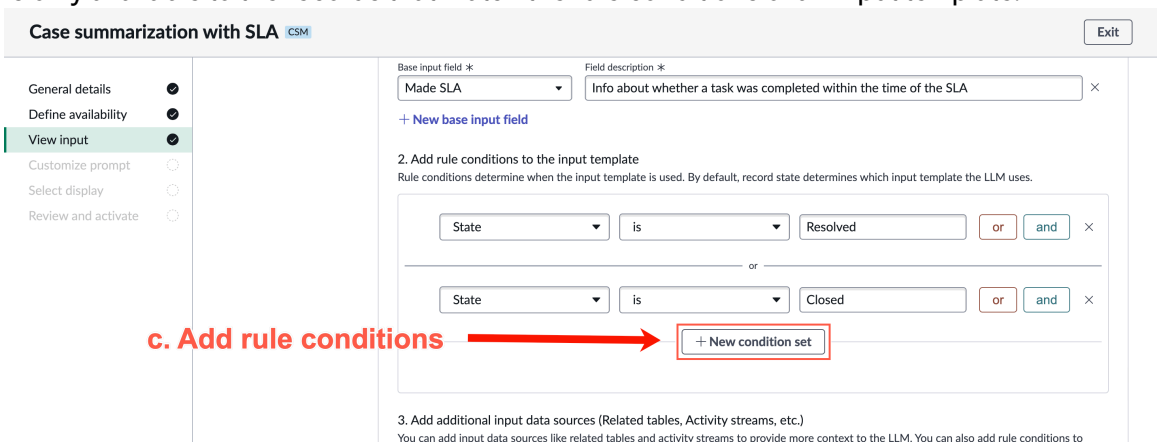
   b. Add the base input table fields by selecting **New base input field**, choosing a field, and entering a field description.
   Each base input table field requires a description. The description informs the large language model (LLM) what the field is for and how the information should be interpreted. The more information that you put in the description means that the model has more context for the data.



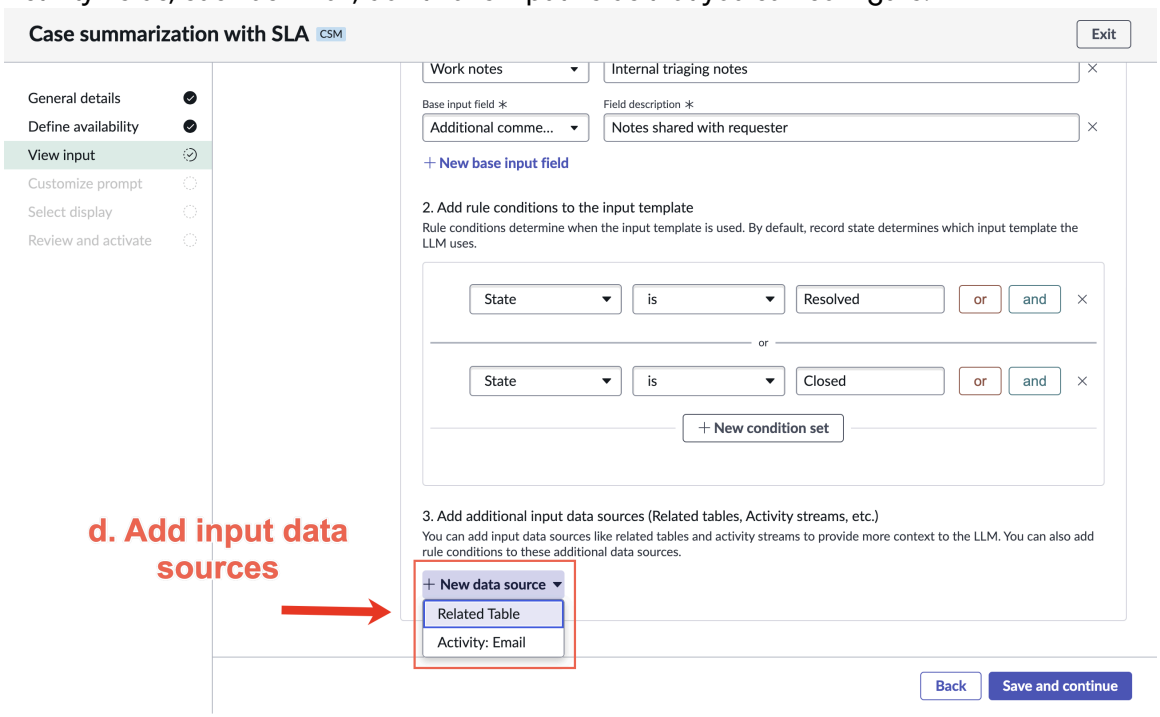
   c. Add or modify the rule conditions for the base table.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The 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.



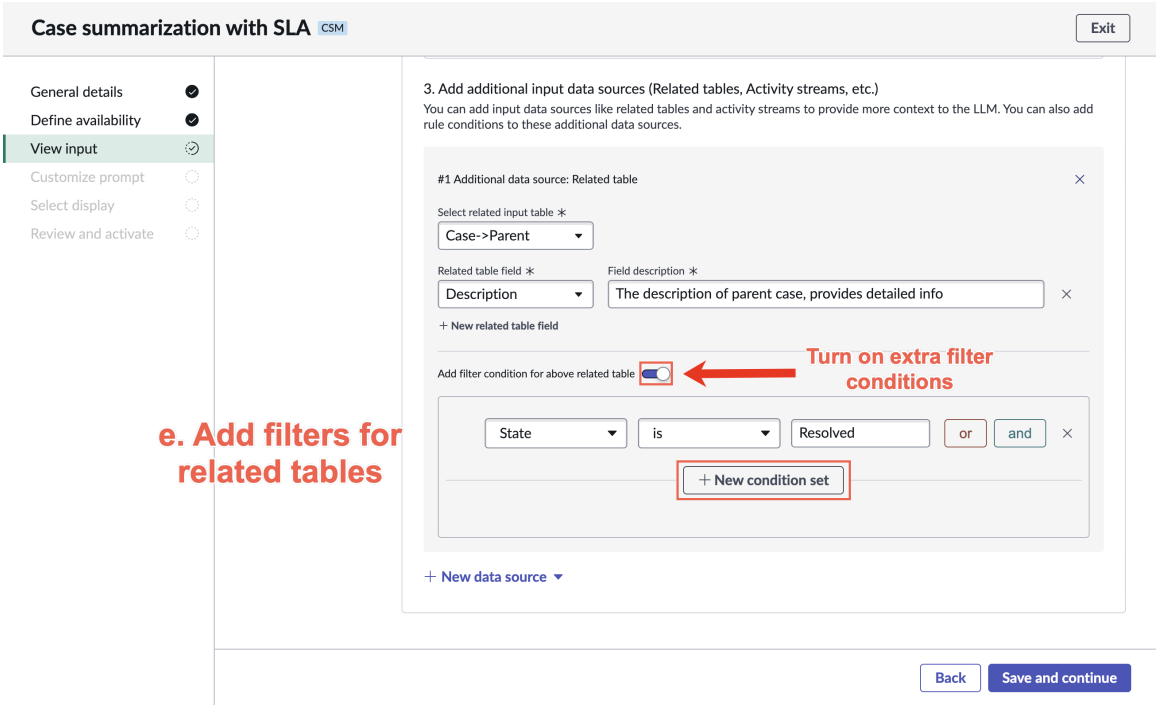
d. Optional: Add additional input data sources by selecting **New data source** and choosing either **Related Table** or **Activity: Email**.
Each related table is configured with input fields and descriptions. More specific descriptions for related tables help provide more context to the LLM. Activity fields, such as Email, don't have input fields that you can configure.



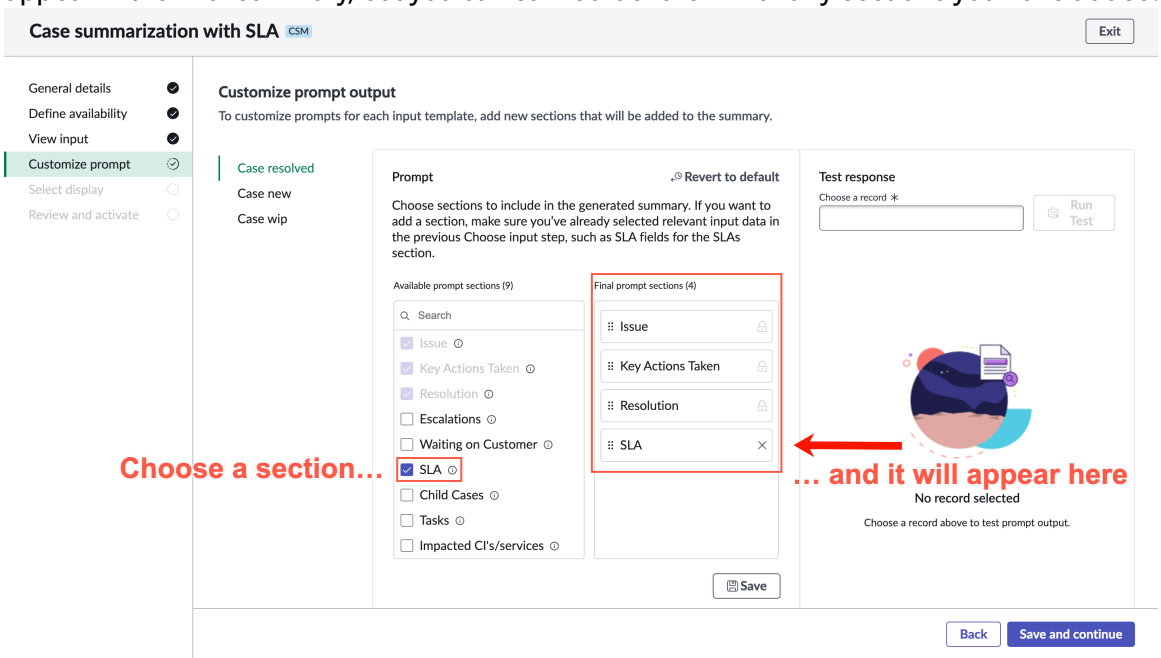
e. Optional: Add a filter condition to the related table.
You can add more rule conditions to the related table. These rule conditions determine whether the data from the additional data source is incorporated into

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

the summary. You can generate summaries on cases that don't match additional data source rule conditions as long as the base table rule conditions are met.



5. Select **Save and continue**.
6. Choose prompt output sections to appear in summaries by moving a prompt section in the Available prompt sections list to the Final prompt sections list.
   You can reorder sections by dragging the boxes in the Final prompt sections list. Some input templates have sections that are marked with the lock icon ( ). These sections must appear in the final summary, but you can still reorder them with any sections you have added.



7. In the Test response panel, select a record from the **Choose a record** field.
8. Generate a summary for the chosen record by selecting **Run Test**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

> **Important:**
> Each time that you test your prompt output, the operation counts as an assist that is tracked by your Now Assist subscription. To track your Now Assist usage, see Monitoring Now Assist usage.

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 .

Enable Dynamic 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.

If you do not 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 will install or update the required dependencies to the latest version.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: admin

About this task

Enabling Dynamic Translation enables you to support users who speak different languages in Now Assist applications. After you have enabled Dynamic Translation, translations will be available for in-product experiences, Virtual Agent, and the Now Assist panel.

All languages that you have enabled through Dynamic Translation 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.

The default translation service provider for Now Assist applications is Microsoft Azure OEM.

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 card, toggle the switch to **On**.

Result

Dynamic Translation with the default language service provider is enabled for Now Assist applications on your instance.

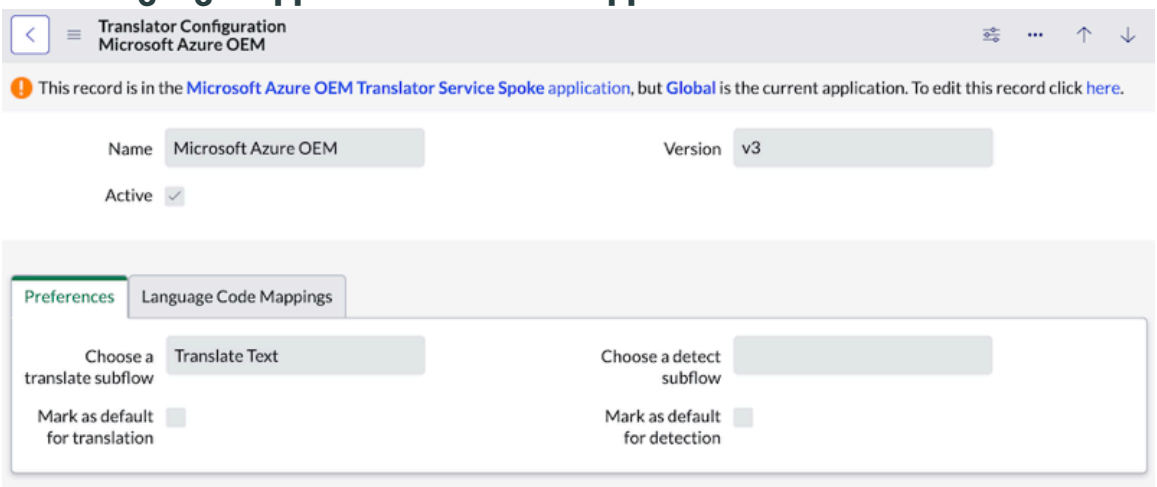
What to do next

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information about enabling dynamic translation in this context, see Enable Dynamic 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".

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: admin

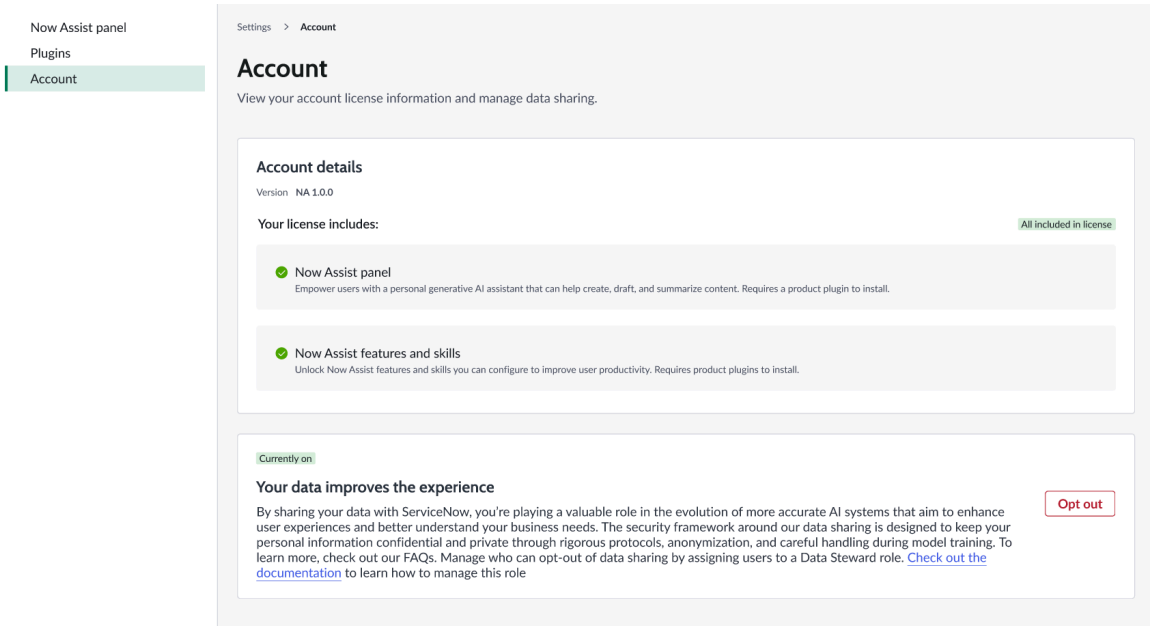
About this task

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 > **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**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



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

About this task
A data steward is responsible for determining whether prompts, inputs, and responses with generative AI in Now Assist applications on your instance are shared with ServiceNow.

The ServiceNow customer data sharing program allows you to provide data to improve ServiceNow AI products. This data helps improve prediction accuracy, enhance user experience, tailor products to your business needs, and reduce hallucinations for your activated Now Assist skills.

The following steps explain how to assign a role to a specific user. Another option is to make a group, such as "AI Data Stewards," assign the role to the group, and then put the user in that group. See Assign a role to a group for more information.

Procedure
1. Navigate to **All** > **User Administration** > **Users** and then open a user record.
2. In the Roles related list, select **Edit**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. In the 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 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.

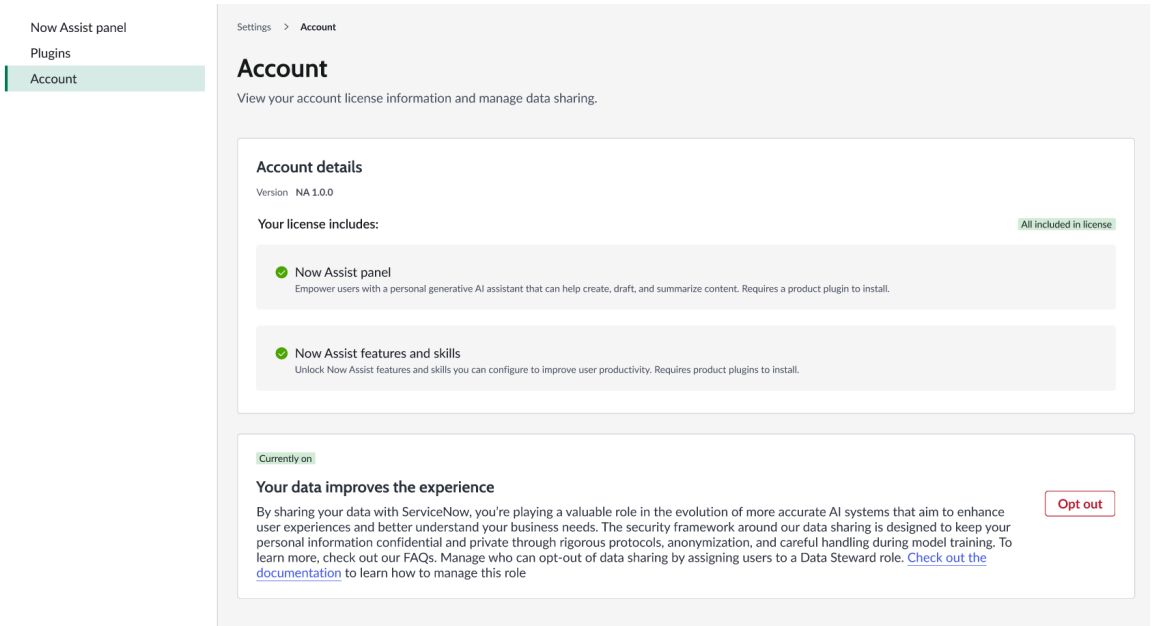
You can choose to opt out a ServiceNow instance from data sharing from the Now Assist Admin console if you don't want to participate. Repeat the opt-out process for all instances that use the Now Assist functionality.

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**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. In the 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.

Analyzing Now Assist performance

Use the Now Assist analytics and monitoring tools in the Overview page to review the summaries, performance information, and issues that need your attention.

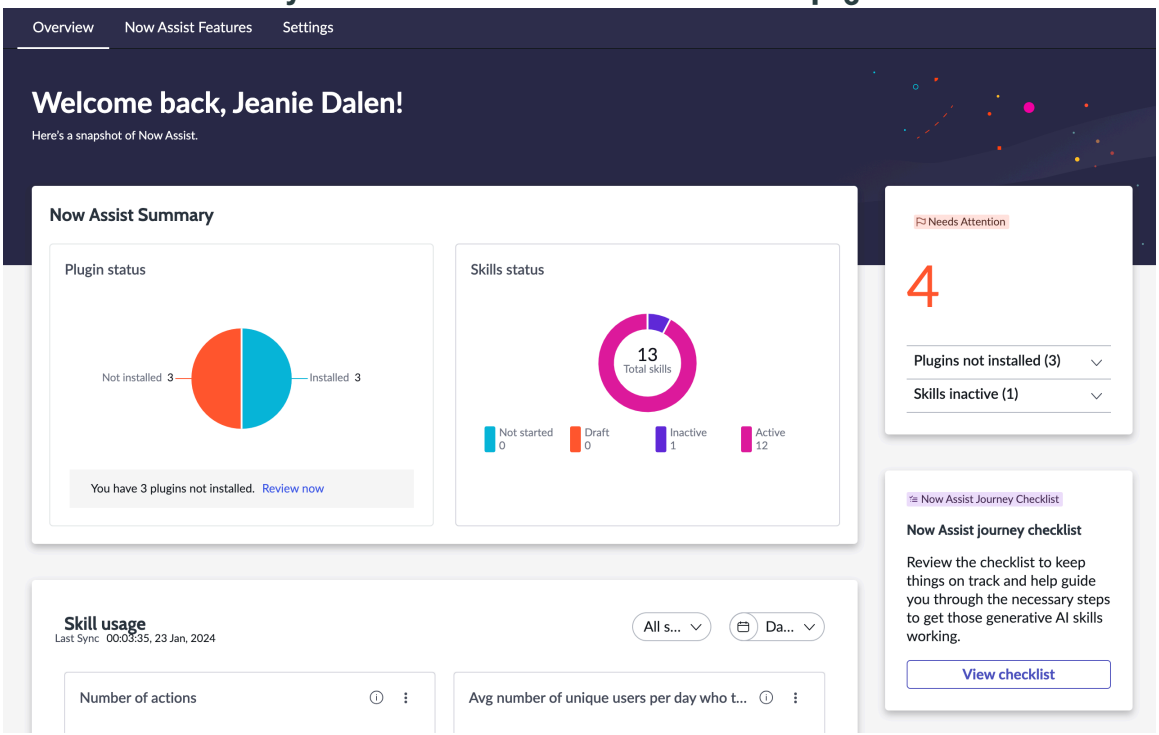
Now Assist Admin Overview page

After you unlock and begin using the skills, you can measure their performance and usage over time with graphs that display important metrics. The console Overview page contains the information about skill performance, including the task completion over time, and a summary of which skills are active or available.

The following example shows the Now Assist Admin Overview page.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Performance and analytics on the Now Assist Admin Overview page



Now Assist Summary section

The Now Assist Summary section displays information about the status of your skills and plugins in a graphical format.

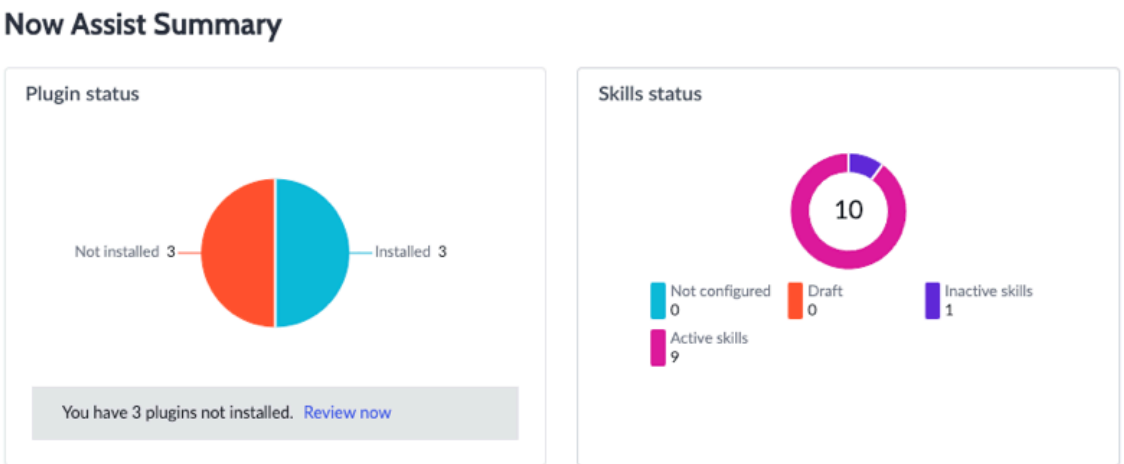
Plugin status

Displays a pie chart that shows the number of plugins that are installed and the number of uninstalled plugins that are available. Select the **Review now** link to review and install additional plugins.

Skills status

Displays a donut chart that shows your skills according to the state of their configuration: active, inactive, draft, or not configured.

The following diagram shows the plugin status that displays as a pie chart, and the Skills status that displays as a donut chart.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Skills performance section

The Skills usage section displays the metrics on your installed active skills. Select the info icon ⓘ to see the information about what each card's metrics represents, or select the option icon ⋮ to refresh the card.

Use the configuration controls to configure the charts.

All skills

View all skills, or select one or more skills individually to change the view.

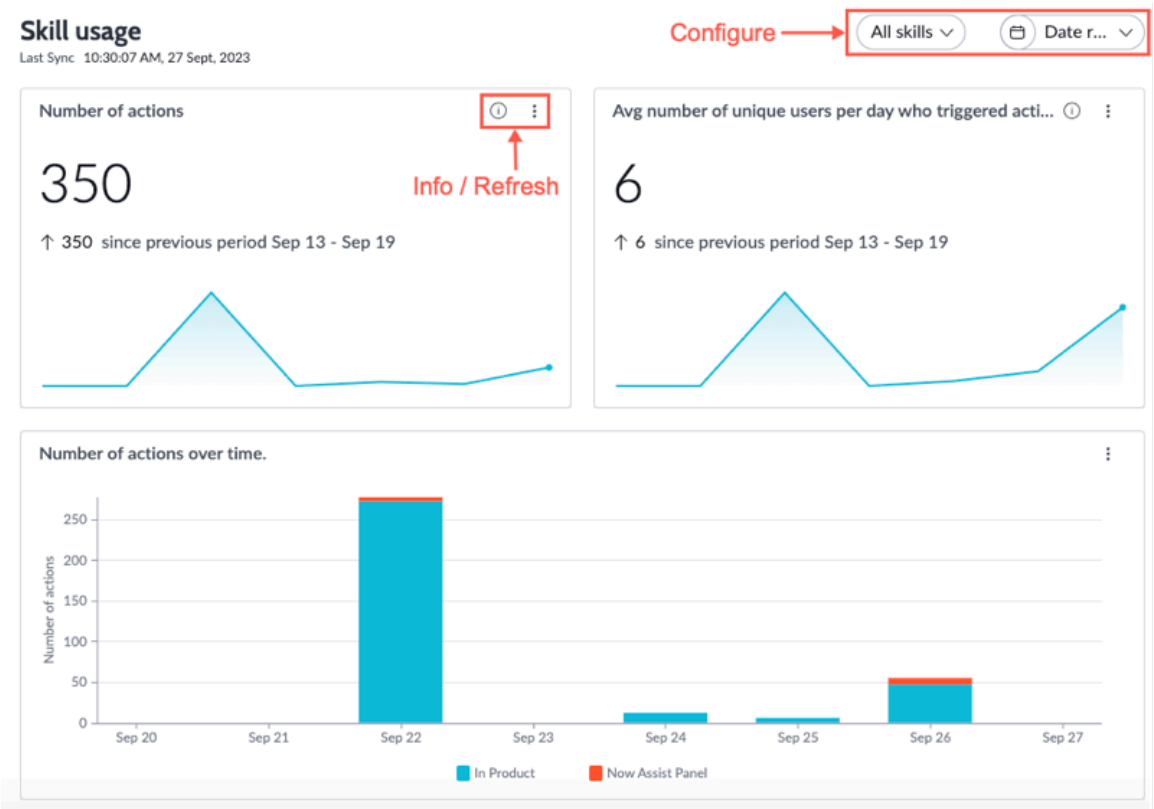
Date range

Change the date range for the analytics.

The following diagram shows the skills usage graphs and controls, which include the number of actions, average unique users per day, and the number of actions over time.

Data is collected once a day and uses Performance Analytics and Reporting (PAR) for data collection and visualization.

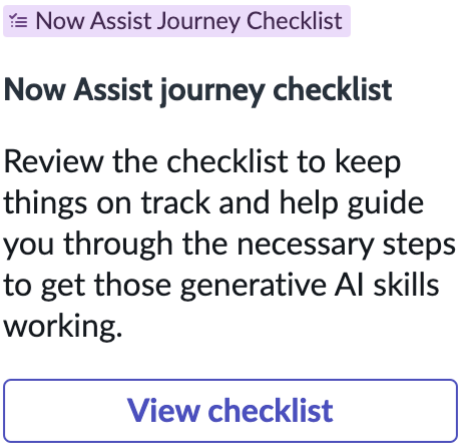
Skills usage graphs and controls



Now Assist journey checklist

The Now Assist journey checklist provides you with a guide of the Now Assist workflow. If you are confused as to what your next step should be when implementing Now Assist on your instance, access the checklist at any time by selecting **View checklist**.

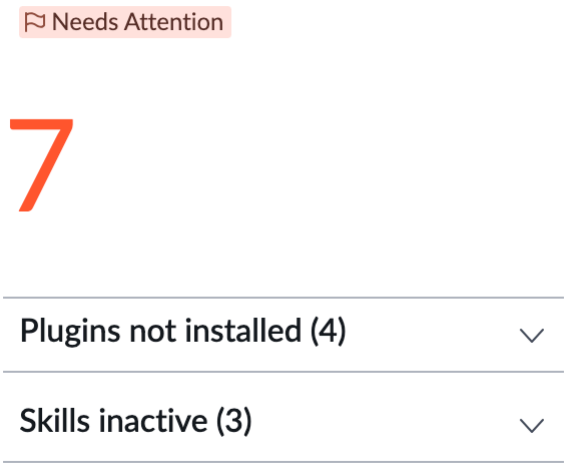
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Additional resources on the Overview page

The Helpful resources and FAQs sections provide links to documentation and answer common questions.

The Needs Attention section displays a count of items that need your attention. Below the count are categorized links to those items.



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 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Model 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.

Now LLM Service for Agent

This document describes changes to the Now LLM Service for Agent by release.

May 2024

- Released KB Generation, which enables agents to use genAI to draft a KB article from a case or incident.
- Improved the completeness of task summarization and resolution note generation by incorporating emails related to the task record.
- Enabled chat summarization in agent chat for live agent-to-live agent transfers.

March 2024

- Improved framework by adding language detection to prepare for multi-lingual summarization.
- Added standard responses for Actions Taken and Resolution sections if no activities are detected.

February 2024

- Fixed hallucination seen in Task Summary, related to case state, resolution section, and activity stream.
- Improved relevance of content in task summary by removing content from system-generated worknotes.
- Fixed issues in chat summarization related to HTML formatting in the prompt.
- Fixed incorrect email ordering for summarization with email content.

Now LLM Service for Q&A

This document describes changes to the Now LLM Service for Q&A by release.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

March 2024

- >2X improvement in average latency and >3X improvement in the 90th percentile latency.
- ~25% relative improvement in the model's quality metrics (accuracy, completeness, conciseness, and formatting).

February 2024

No updates.

Now LLM Service for Flow Designer

This document describes changes to the Now LLM Service for Flow Designer by release.

February 2024

Enhanced out-of-the-box components index to include integration hub components for a better customer experience in building flows that need app integration.

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 [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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

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..

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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.

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.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Overview

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**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. 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 aren't 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 done by definition and skill config so you can monitor executions of large language model (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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Generative AI 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

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 Flow Designer 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.

Explore  Explore Generative AI concepts and terminology.	Configure  Configure Generative AI Controller provider capabilities.
Use  Use Generative AI Controller provider capabilities.	Reference  Reference for Generative AI Controller.

Troubleshoot and get help

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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

AI limitations

This application uses artificial intelligence (AI) and machine learning, which are rapidly evolving fields of study that generate predictions based on patterns in data. As a result, this application may not always produce accurate, complete, or appropriate information. 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 Flow Designer flows, Virtual Agent topics, and scripting like background scripts and business rules.

Generative AI Controller benefits

Benefit	Feature
Generate text to summarize complex information.	Flow Designer Actions to Generate Content and Summarize
Analyze user sentiment to identify and alleviate end-user concerns.	Sentiment Analysis Flow Designer Action
Query a large language model (LLM) directly.	Generic Prompt Flow Designer Action
Write scripts for AI model capabilities that are designed to increase the accuracy and scalability of your custom content.	Script with generative AI
Integrate with Now LLM Service or third-party AI service providers to customize your AI experience	OpenAI, Azure OpenAI, Google AI, Aleph Alpha, IBM watsonx

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get started 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

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

How 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 Flow Designer, 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Microsoft Azure OpenAI
- OpenAI
- Google AI Studio
- Google Vertex AI
- Aleph Alpha
- IBM watsonx

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, 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.

Add semantic filtering and error handling for sensitive information

Configure semantic filtering of sensitive topics and redirect users to another Virtual Agent topic if sensitive material is brought up in interactions.

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

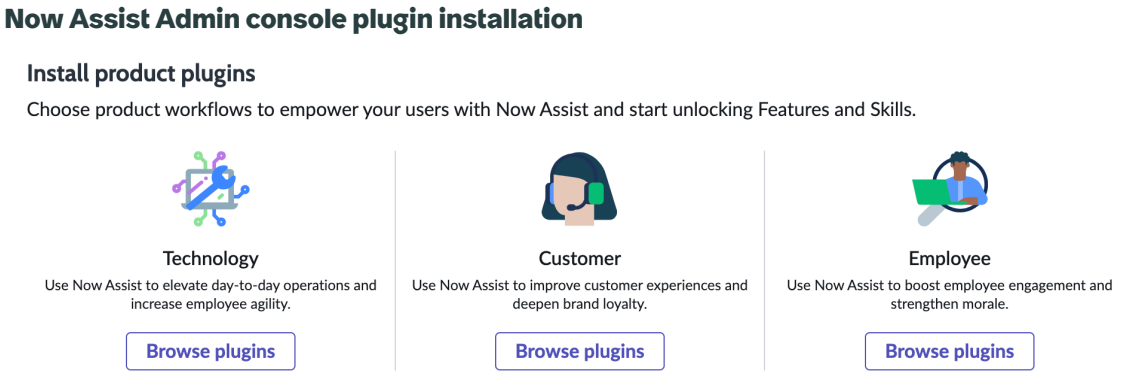
You can install the Generative AI Controller application (sn.generative.ai) with Now Assist applications if you have the admin role.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



The 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
- 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

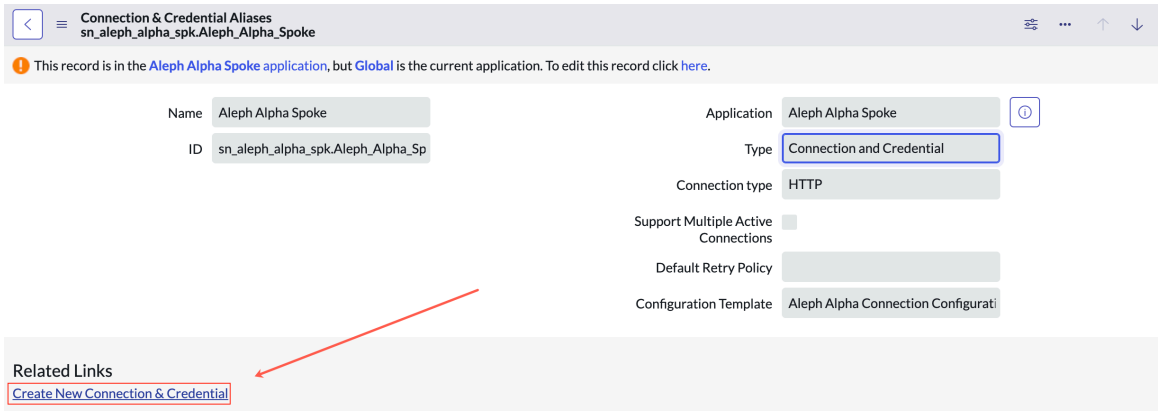
About this task

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.

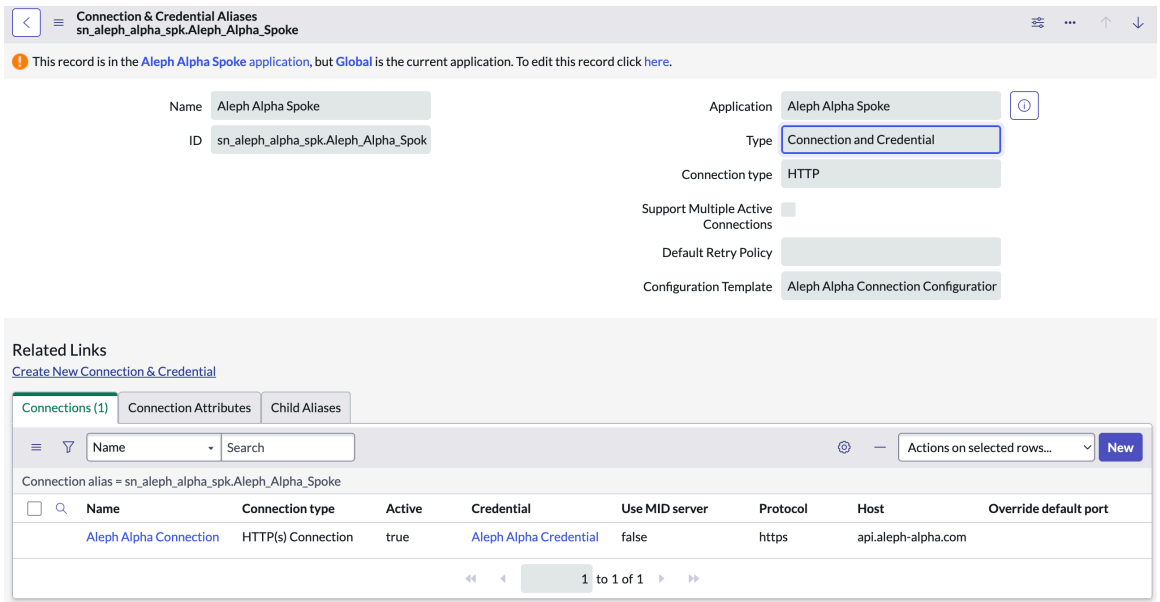
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



3. Select the **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 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

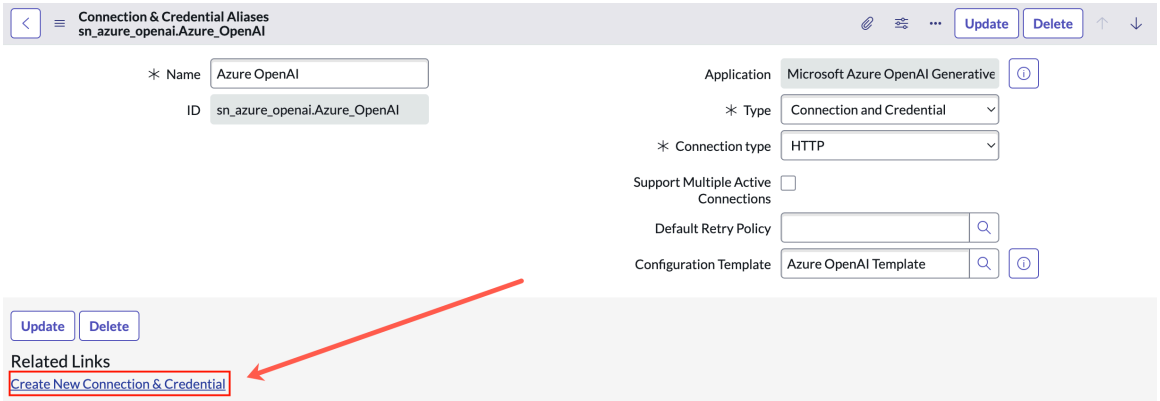
About this task

In order to use models with Azure OpenAI as your LLM provider for Generative AI Controller capabilities, you must have an active connection configured.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Connections & Credentials** > **Connections & 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.

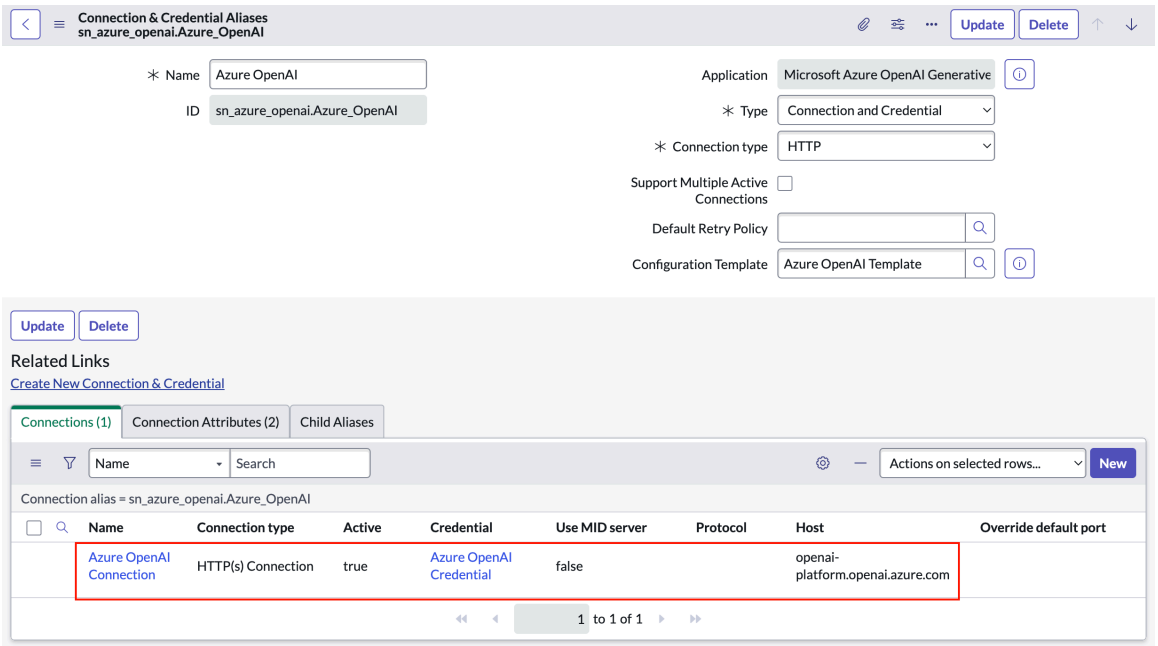
5. In the API key field, enter the API key for the provider.

   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.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next
If you want to use generative AI capabilities through your MID Server, open the new Connection record, select the **Use MID server** check box, and save the record.

Configure API credentials for Google AI Studio

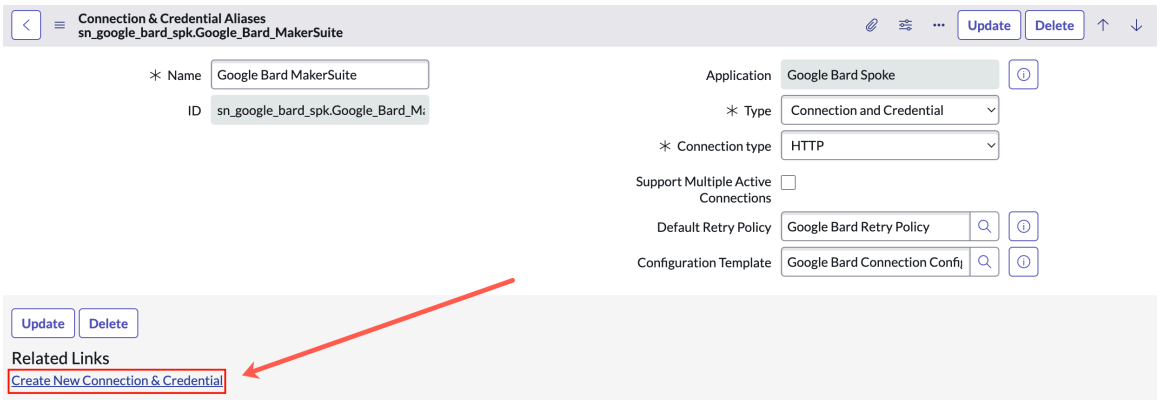
Configure your API credentials to use Google AI Studio in custom workflows and Virtual Agent Designer topics.

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.

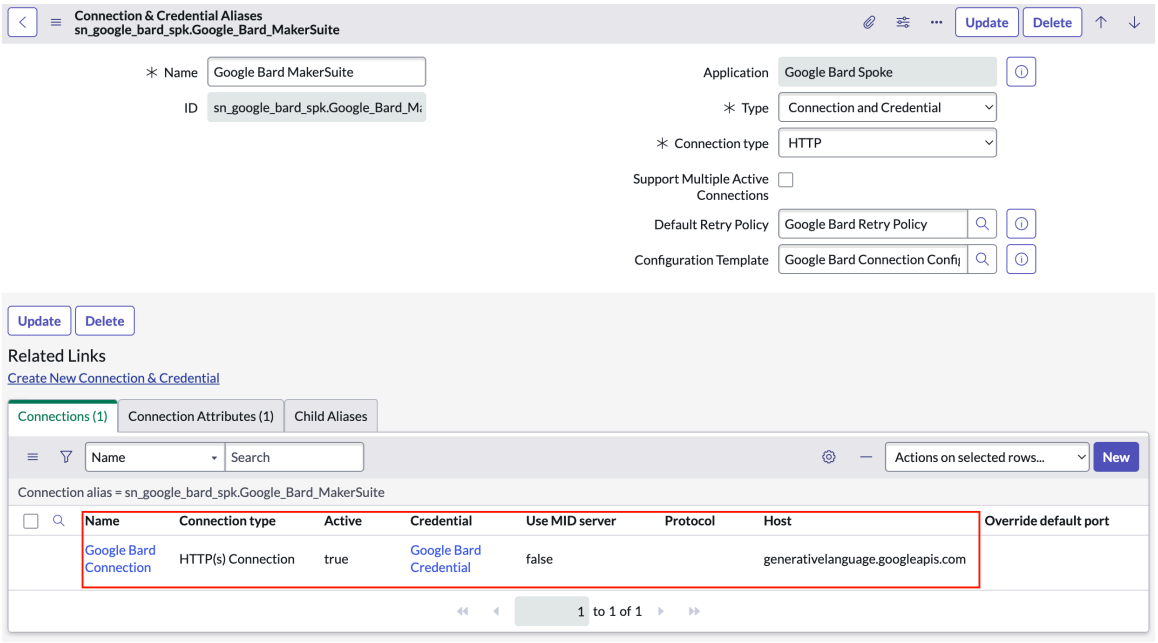
Procedure
1. Navigate to **All** > **Connections & Credentials** > **Connections & Credentials Aliases**.
2. Open the record for Google Bard MakerSuite.
3. Select the **Create New Connection & Credential** related link.



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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



What to do next

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.

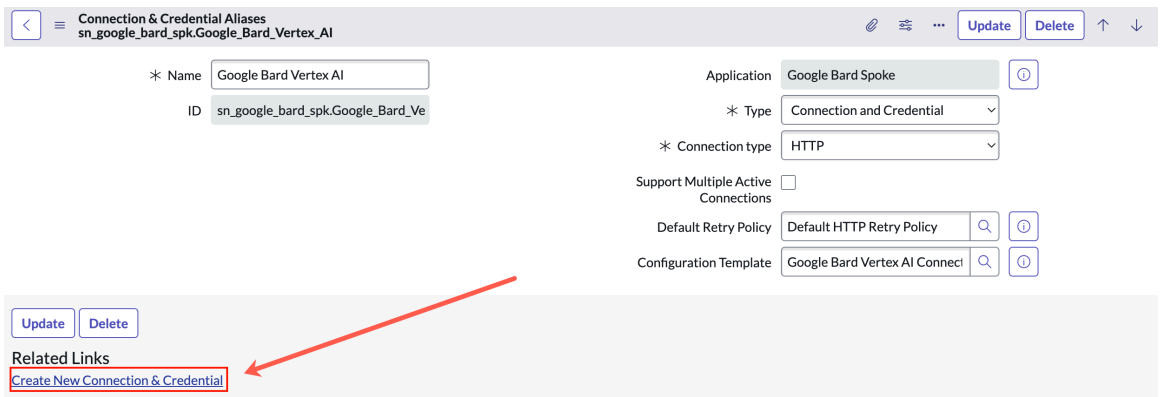
Role required: admin

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

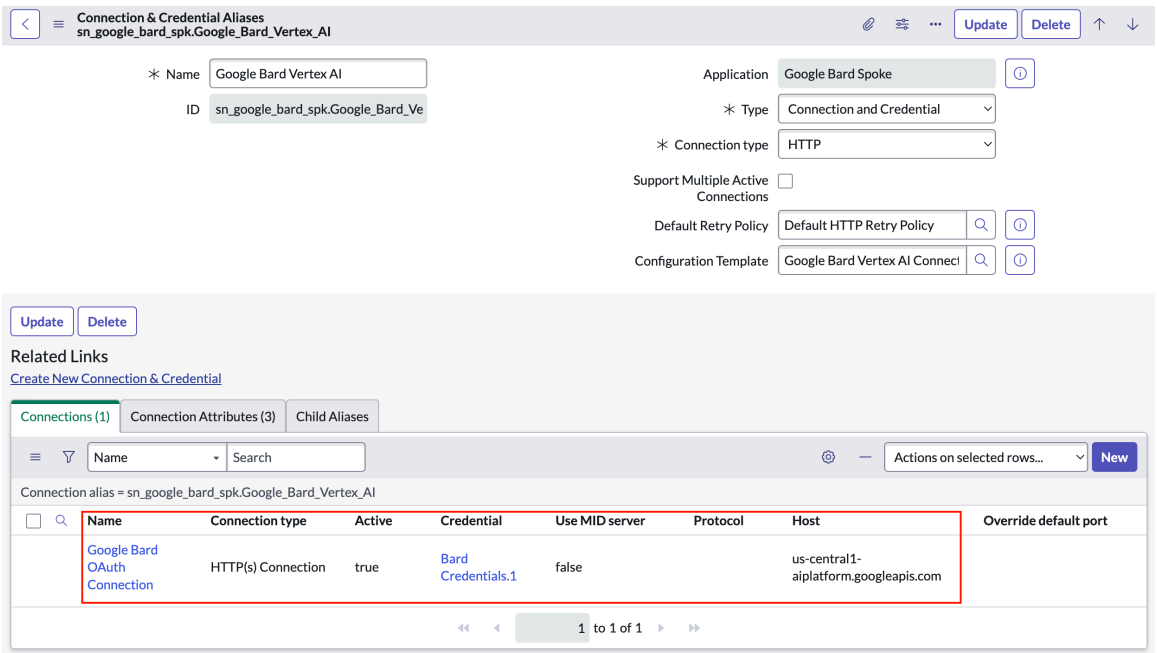
Google OAuth Connection

Field	Value
Project ID	The Project ID found in the Google Cloud console
Credential Name	The name of your credential, such as `Google OAuth Credential`
OAuth Name	The name of your OAuth authentication, such as `Google Registry`
OAuth Client ID	To get the OAuth Client ID, create a new OAuth Client ID with the Google Cloud console with the following attributes: **a.** Application type: `Web application` **b.** Authorized redirect URI: URL in the OAuth Redirect URL field, usually `<instance>.service-now.com/oauth_redirect.do` For more information, see the Google documentation for creating OAuth client IDs. Once you have created the OAuth client, a pop-up window will have the Client ID and Client secret for you to copy into your clipboard.
OAuth Client Secret	Client secret from your OAuth Client ID found in the Google Cloud console

5. In the pop-up window, log in to a Google Account with access to the project.
6. When prompted for Google Cloud access for gsuite spokes, select **Allow**.

Result

You can now use Completions – Vertex AI and Chat Completions – Vertex AI in Flow Designer, Virtual Agent Designer, and scripts to create custom experiences with generative AI.



What to do next

Use your LLM provider to create flows with Flow Designer, topics with Virtual Agent Designer, or scripts to provide the benefits of generative AI to your users.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Configure API credentials for 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.

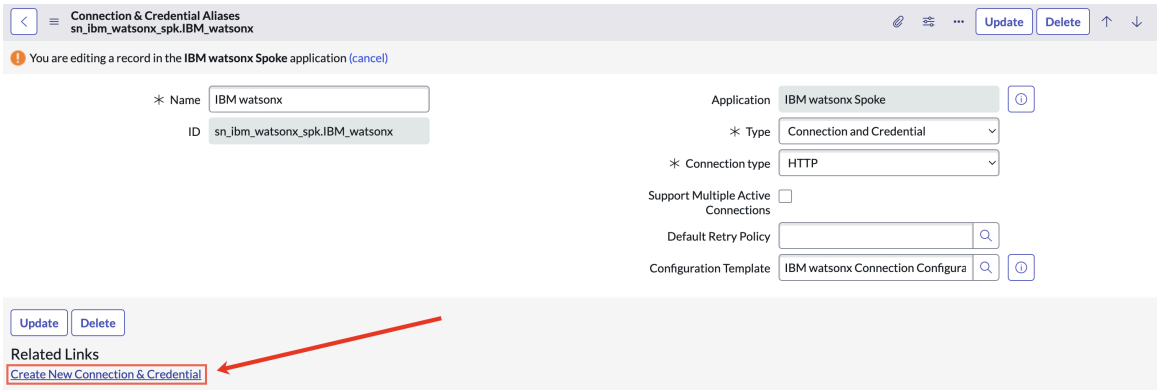
Role required: admin

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.



4. In the API key field, enter your API key.

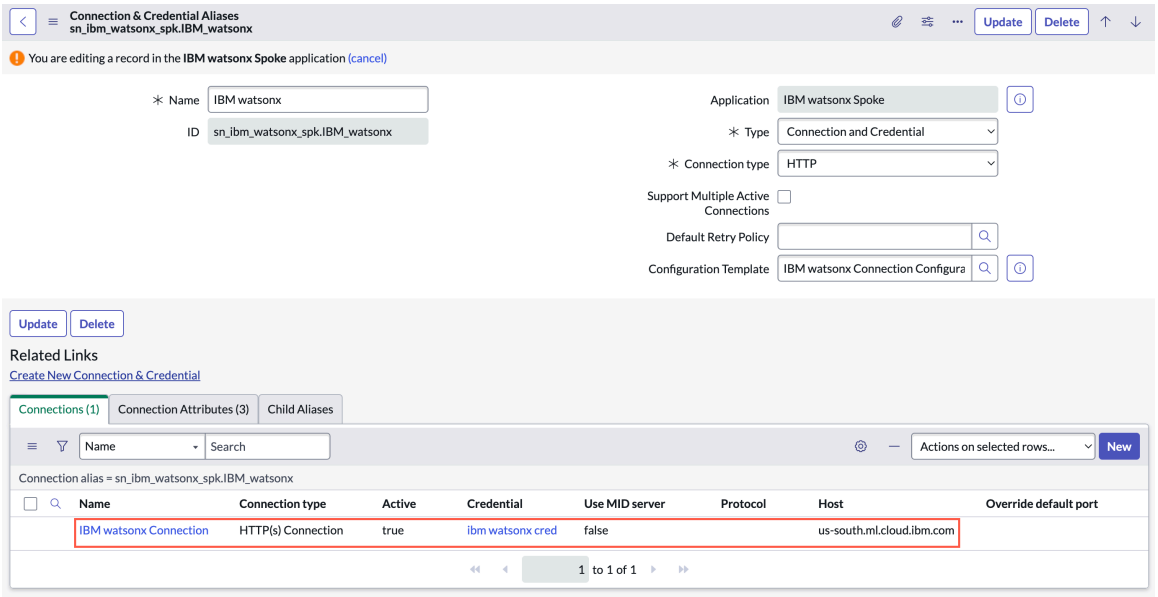
   For more information about generating an API key, see IBM's documentation for generating API keys for authentication.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



What to do next

If you want to use generative AI capabilities through your MID Server, open the new Connection record, select the **Use MID server** check box, and save the record.

Configure API credentials for 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.

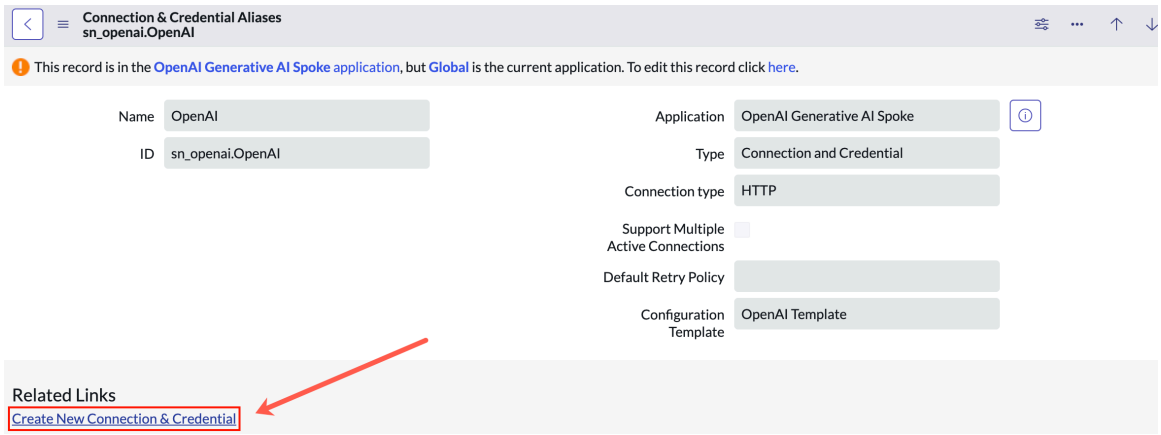
Role required: admin

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.

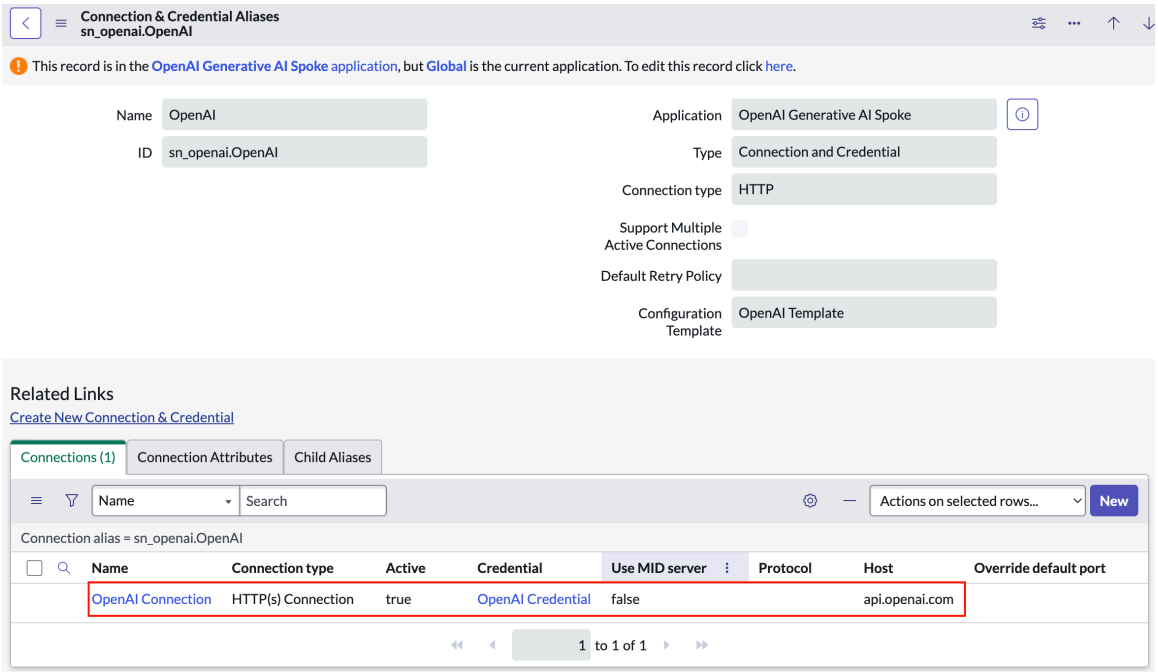
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For 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

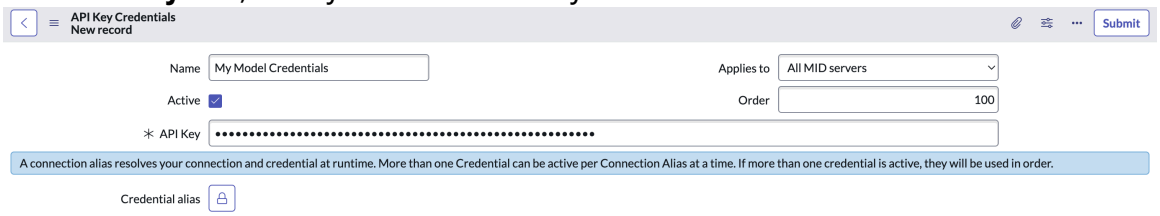
You can connect an external LLM to the ServiceNow AI Platform by creating a connection alias to the model, create a model record that points to the alias, create a prompt record for the model, and then create a transformer record to transform the request or response from the LLM.

Procedure

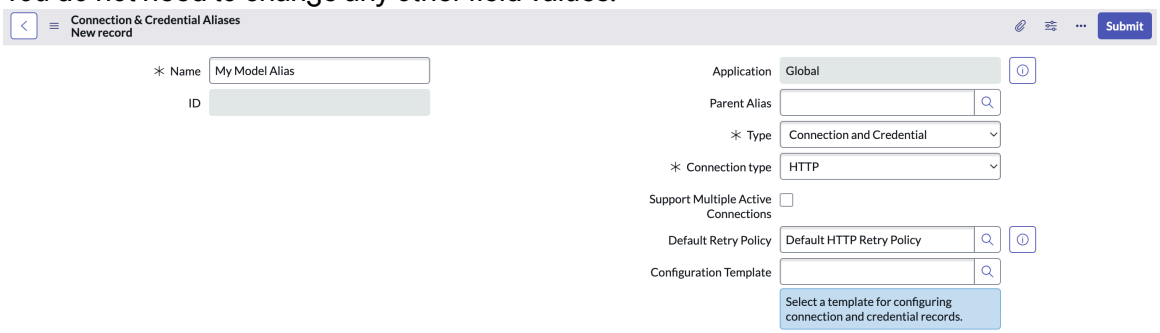
1. Navigate to **All** > **Connections & Credentials** > **Credentials**.
2. Select **API Key Credentials**
3. In the **Name** field, give your API credential a name.
   Including the name of the model in the API Key Credential makes it easier to identify later.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

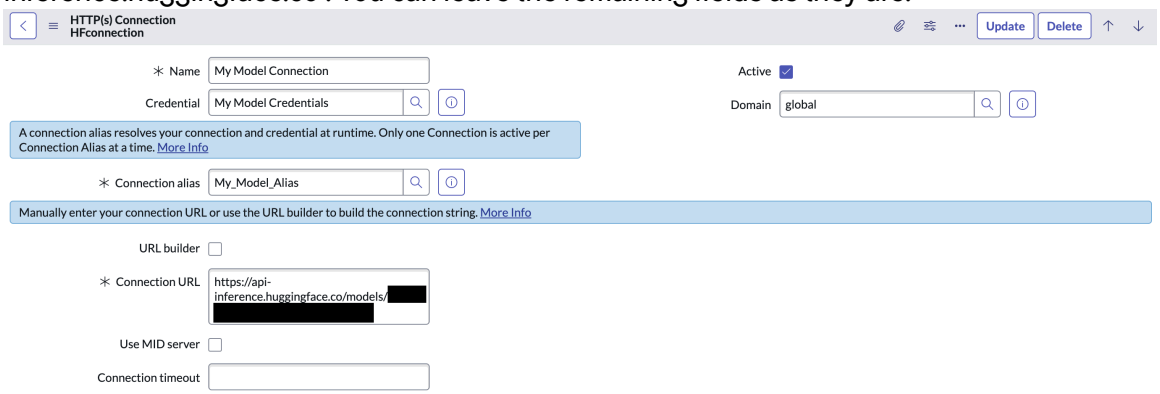
4. In the **API Key** field, enter your model's API key.



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.

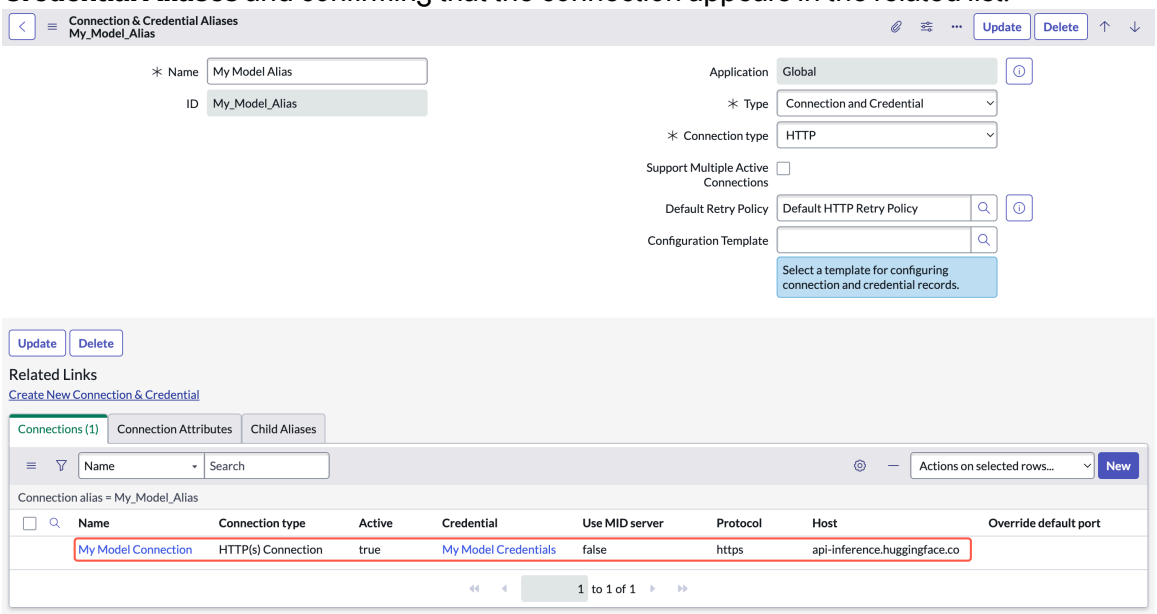


16. Select **Submit**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

You have the connection and credential alias to use for connecting a generic LLM to use generative AI capabilities on the ServiceNow AI Platform. You can confirm this by navigating to **All** > **Connections & Credentials** > **Connections & Credential Aliases** and confirming that the connection appears in the related list.



What to do next

For more information on configuring a generic LLM, see configure a generic LLM connector

Configure a generic large language model (LLM) connector

Connect an external LLM to the ServiceNow AI Platform by using a generic LLM connector. With a connector, you can write your own prompts to send to the LLM and create your own generative AI capabilities.

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.

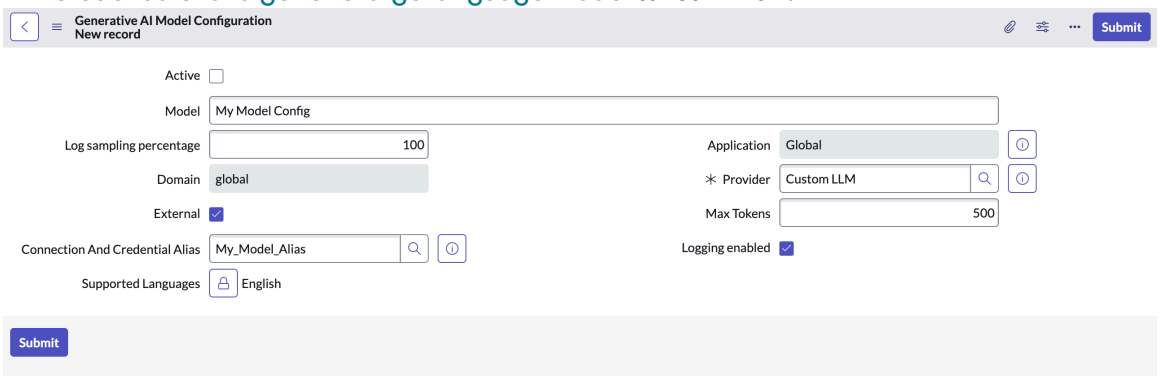
Note:
Generative AI Controller only supports text generation.

Procedure

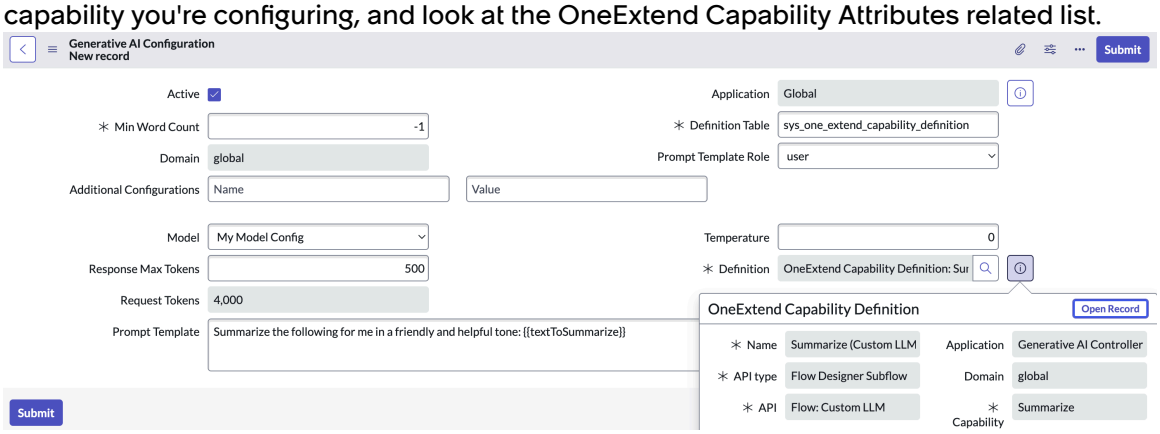
1. In the navigation filter, go to the Generative AI Model Configuration [sys_generative_ai_model_config] table by entering `sys_generative_ai_model_config.list`.
2. Select **New**.
3. In the **Model** field, enter the name of the model.
   If you are using Azure OpenAI, then the model name is the deployment name of your resource.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

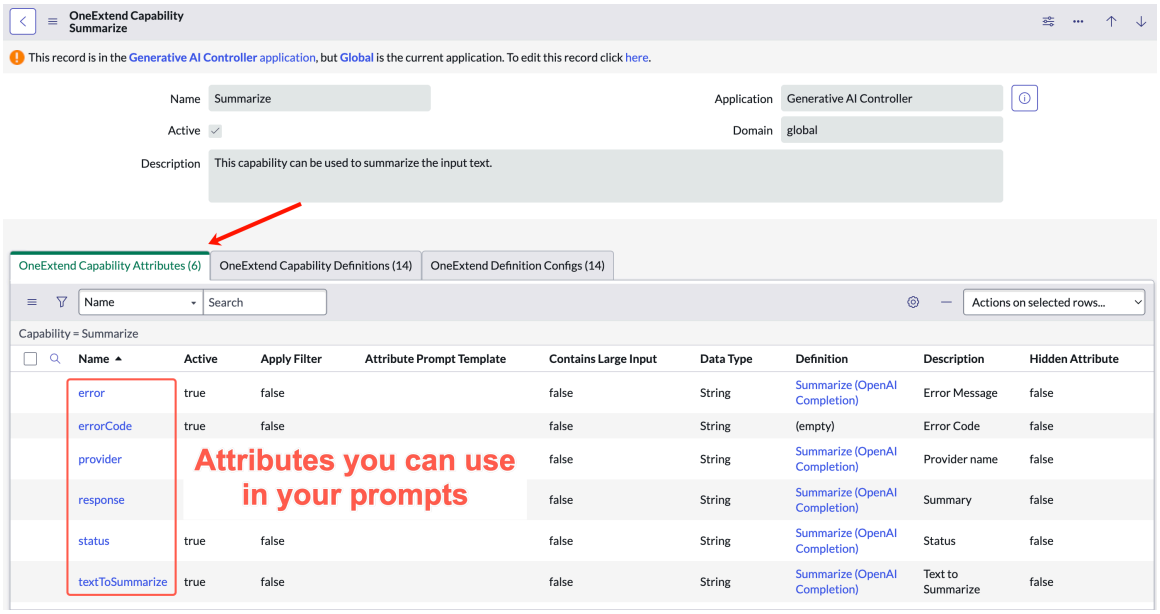
4. In the **Provider** field, select **Custom LLM**.
5. In the **Connection and Credential Alias** field, select the alias that you created when you set up your credentials.
   If you haven't already configured your API credentials, see configure API credentials for a generic large language model to 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.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



You 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
```

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
        }],
        max_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;
```

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
})(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.

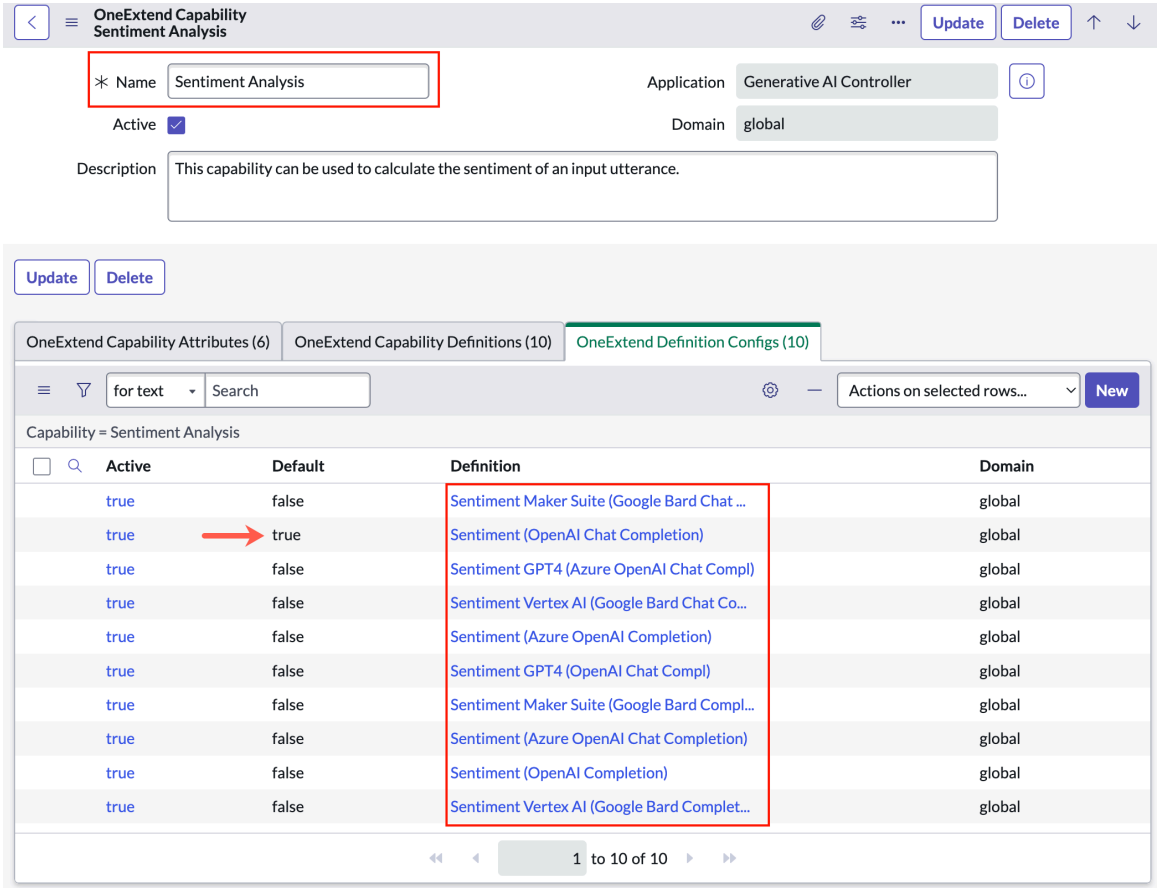
Generative AI Controller capabilities

Capability	Definition
Generate Content	Generate texts about a given topic with Flow Designer and Virtual Agent Designer.
Generic Prompt	Create your own use case and prompt.
Sentiment Analysis	Analyze user sentiment with Flow Designer and Virtual Agent Designer.
Summarize	Create summaries of topics with Flow Designer and Virtual Agent Designer.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. In the 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).



Model Options

You have the option to choose between several different models for each capability.

Capability Definition	Model
○ OpenAI Completion ○ Azure OpenAI Completion	GPT-3
○ OpenAI Chat Completion ○ Azure OpenAI Chat Completion	GPT-3.5
○ GPT4 (OpenAI Chat Compl) ○ GPT4 (Azure OpenAI Chat Compl)	GPT-4
○ AI Studio (Google Cloud Completion) ○ AI Studio (Google Cloud Chat Completion)	Google Gemini

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Capability Definition	Model
○ Vertex AI (Google Cloud Completion) ○ Vertex AI (Google Cloud Chat Completion)	
Aleph Alpha Completion	Luminous
IBM watsonx	Granite

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.

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

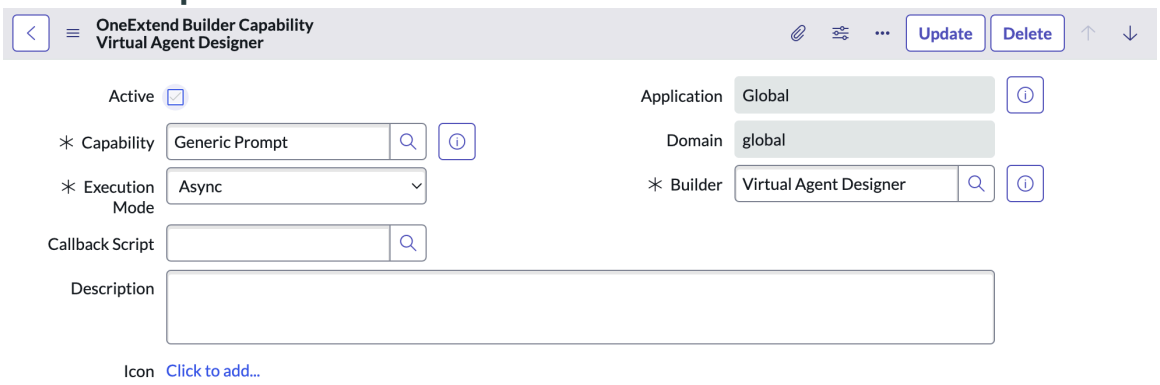
Procedure
1. Go to the OneExtend Builder Configs table by entering `sys_one_extend_builder_config.list` in the navigation filter.
2. Open the Virtual Agent Designer record.
3. In the OneExtend Builder Capabilities related list, select **New**.
4. On the form, fill in the fields.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Generic Prompt builder definition form



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 Flow Designer 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.

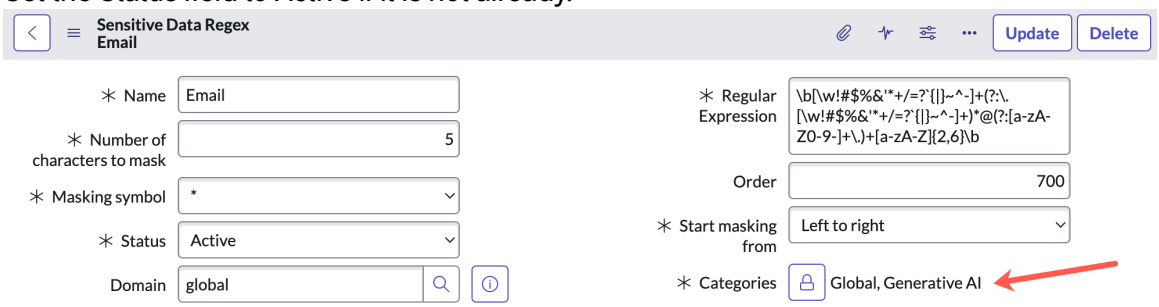
There are some considerations for using the Sensitive Data Handler with generative AI.

- 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **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.

Add semantic filtering and error handling for sensitive information

Create semantic filters for generative AI requests and responses to help filter out sensitive topics during interactions with generative AI in Virtual Agent.

Before you begin

You must have a Virtual Agent topic to switch to when a filter detects unwanted material. Examples include a text output topic signaling that an error has occurred or a topic that transfers the user to a Live Agent. The topic must be published and active.

Role required: admin

About this task

To help prevent generative AI conversations about sensitive topics, you can use semantic filtering to help identify phrases in Virtual Agent conversations. Identify material that you don't want to be included in a request or response of an LLM and add filters to remove that material. AI Semantic Search can be used to help determine if sensitive content is being discussed and trigger a switch to a designated Virtual Agent topic.

Procedure

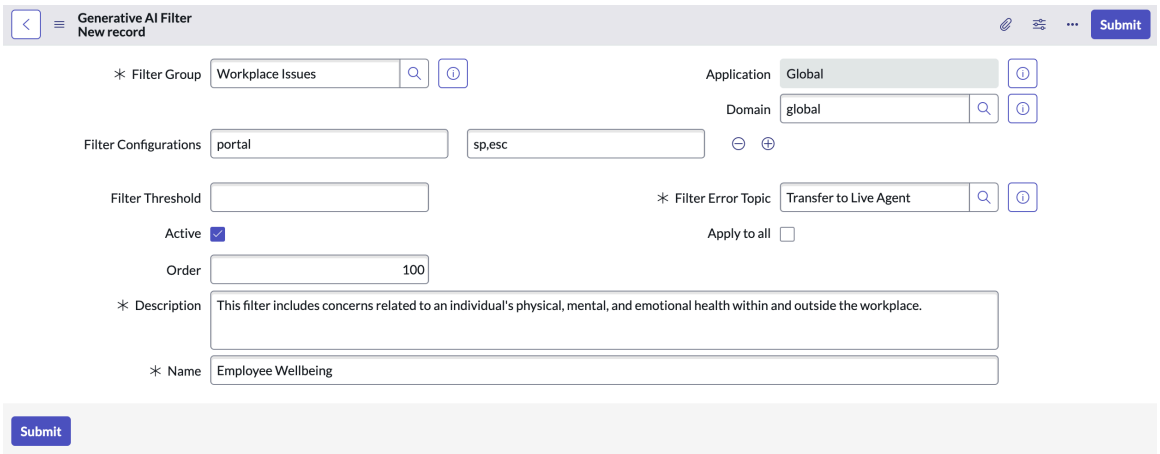
1. In the navigator, go to the Generative AI Filter Group [sys_gen_ai_filter_group] table by entering `sys_gen_ai_filter_group.list`.
   If you have Now Assist for HR Service Delivery (HRSD), there may be groups already created for you. You can add to these groups or create your own. If you don't have Now Assist for HRSD, you can still create your own.
2. Select **New**.
3. In the **Name** field, enter a name for the group.
4. In the **Description** field, enter a description for the group.
   A filter group description could start with "This filter includes workplace issues/concerns such as employee well-being, fairness, and ethics." You could then add examples of these workplace issues/concerns in the description for this category of sensitive material.
5. Select **Submit** to create the group.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. In the navigator, go to the Generative AI Filter [sys_gen_ai_filter] table by entering `sys_gen_ai_filter.list`.
7. Select **New**.
8. In the **Filter group** category, choose a Generative AI Filter Group.
9. **Optional:** In the **Filter configuration** field, use the "portal" value to determine when the filter should be run, such as setting “portal” to “sp” to only apply the filter on the sp portal.
   Use a comma-separated list to select multiple values.
10. In the **Filter Threshold** field, set the cutoff for determining when the LLM should generate the filter error.
    The lower the value, the more sensitive the filter. Experiment with different filter thresholds to help set your preferred boundaries.
11. 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.
12. In the **Description** field, enter a description of the filter.
    Like the Generative AI Filter Group record, the more descriptive you are, the more likely that the LLM will understand the purpose of the filter.
13. Set the **Filter Error Topic** to the topic that should run when the filter detects unwanted content, such as a Transfer to Live Agent topic or Error Handling Topic.
14. Select **Apply to all** to apply the filter to all capabilities (Summarize, Generate Content, etc.).

> **Important:**
> If you did not select **Apply to all**, you need to create a Generative AI Filter Config record to map the filter to the capability. See the additional steps below to create the mapping.

15. Select **Submit** to create the filter.



16. To map a filter to a capability, create a Generative AI Filter Config record.

    a. In the navigator, go to the Generative AI Filter Config [sys_gen_ai_filter_config] table by entering `sys_gen_ai_filter_config.list`.

    b. Select **New**.

    c. In the **Config Table** field, select `OneExtend Capability [sys_one_extend_capability]`.

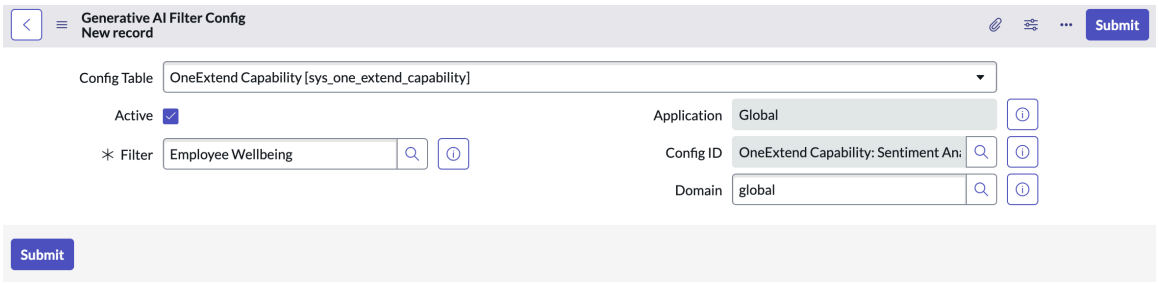
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

d. In the **Config ID** field, select the list lookup icon ( ) to select a Config.

e. In the **Document** field, select the capability you would like the filter to apply to.
If you want the filter to apply to some but not all capabilities, you need to create a Generative AI Filter Config record for each capability.

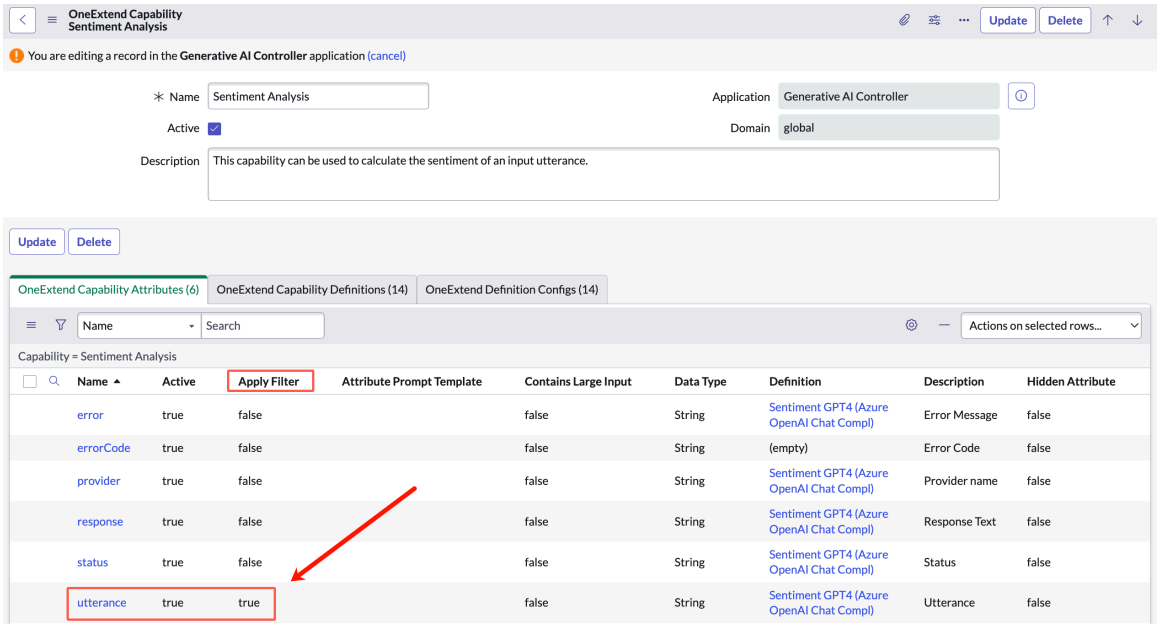
f. Select **OK** to make your selection and close the modal.

g. Select **Submit**.



17. In the navigator, go to the Generative AI Filter Sample [sys_gen_ai_filter_sample] table by entering `sys_gen_ai_filter_sample`.
18. Select **New** to create a sample.
19. In the **Sample Text** field, enter an example of what you want the filter to detect.
Each sample can be short, and having more sample records helps AI Semantic Search determine what the filter should catch. If you find that certain terms or statements aren't being caught by your filter, you can try adding more examples to see if the filter improves.
20. To enable the filter on an attribute, go to the OneExtend Capability [sys_one_extend_capability] table by entering `sys_one_extend_capability.list`.
After you have created the Generative AI Filter Group, Generative AI Filter, and Generative AI Filter Samples, you need to enable the filter on an attribute for the capability you are configuring.
21. Open the record for the capability.
22. Set the Apply Filter field to true to apply the filter.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



23. Repeat for each capability you want the filter to be on.

Result
Semantic filtering is now available for the capability on the attribute you configured.

What to do next
Configure sensitive data handling to replace personally identifiable information with placeholder text in requests sent to the LLM.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

4. In the 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

If you have Dynamic Translation installed and active for Virtual Agent, Dynamic Translation is enabled by default for LLM Virtual Agent conversations. You can disable these dynamic translations using a system property.

Before you begin
Role required: admin

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task
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 not translated into English before being sent to Now LLM and output messages are not translated back into the requester's preferred language.

Using Generative AI Controller

Use Generative AI Controller to integrate generative AI with Flow Designer, Virtual Agent Designer, and scripting.

Use Generative AI Controller with Flow Designer

Use the Generative AI Controller capabilities in flows with Flow Designer.

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 Flow Designer

Use the Analyze Sentiment capability to determine the user attitude and the use that results in a flow with Flow Designer.

Use a script with Generative AI Controller

Write scripts to harness the power of generative AI directly.

Use Generative AI Controller with Flow Designer

Summarize the information automatically and add it to a form with the Generative AI Controller Summarize action in Flow Designer. You can build a flow to summarize any field.

Before you begin
Role required: admin

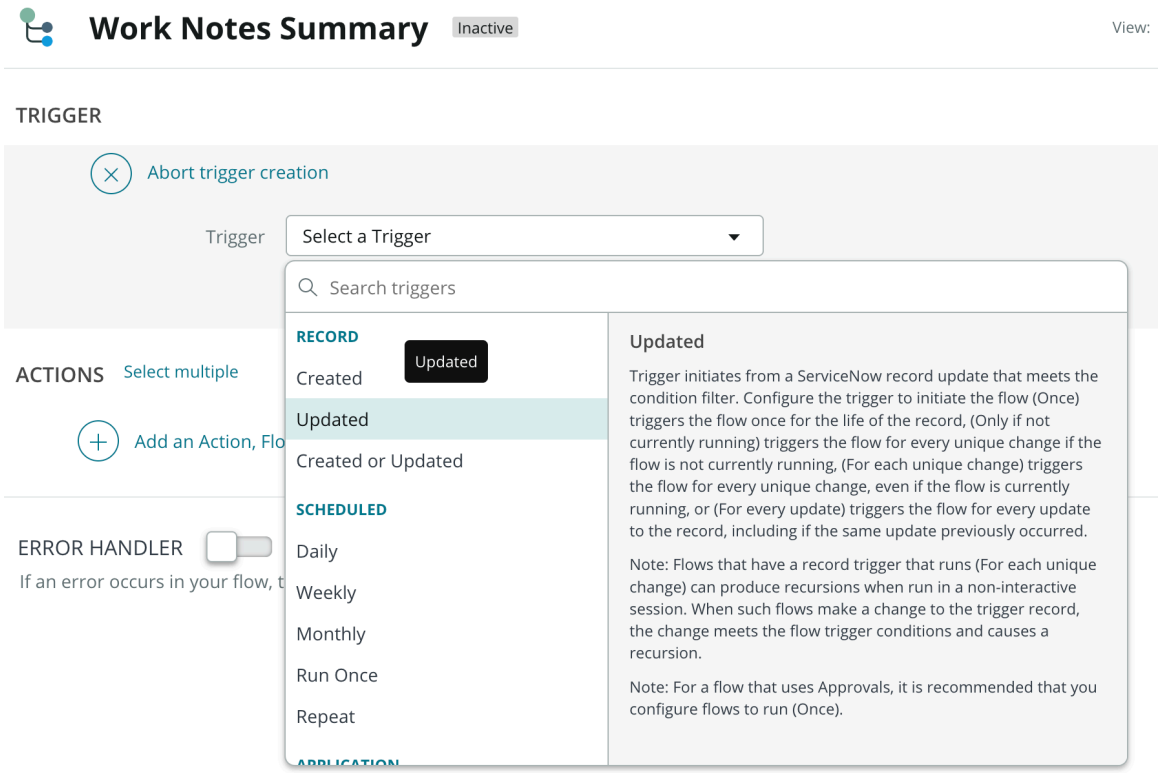
About this task
In this example, you can build a flow to summarize the comments in the work notes of an incident and use that summary in the resolution notes.

For more information on tracking Generative AI Controller usage, see Monitoring Now Assist usage.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

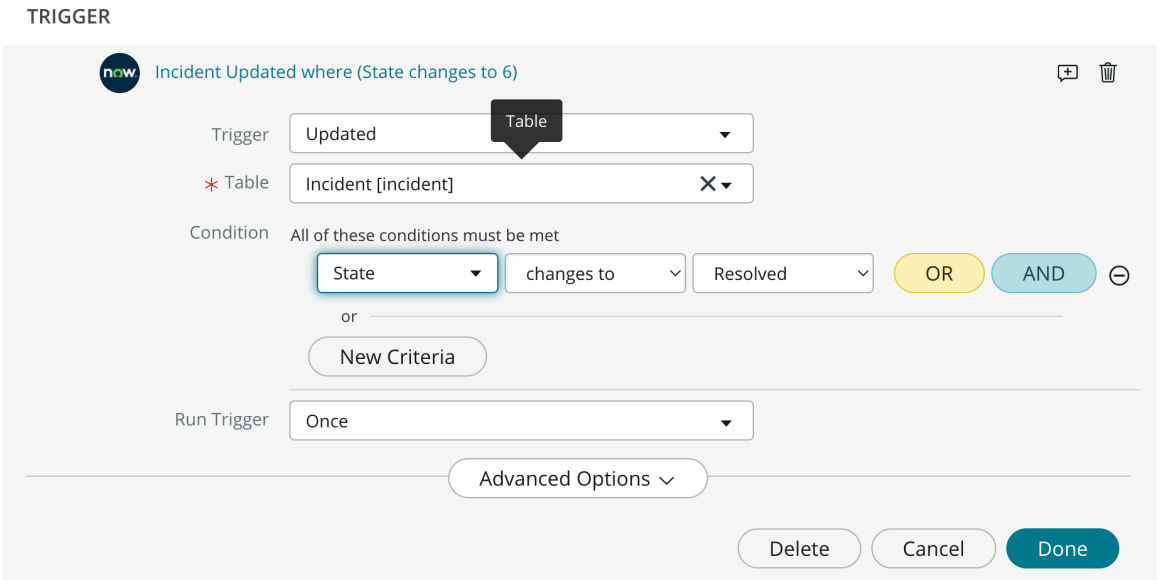
1. Navigate to **All** > **Process Automation** > **Flow Designer**.
2. Select **New** > **Flow**.
3. Enter the name of the flow, such as `Summarize Comments and Work Notes`.
4. Select **Submit**.
5. Select **Add a trigger**.
6. Select when you would like the flow to run.
   To follow this example, select **Record** > **Updated**.



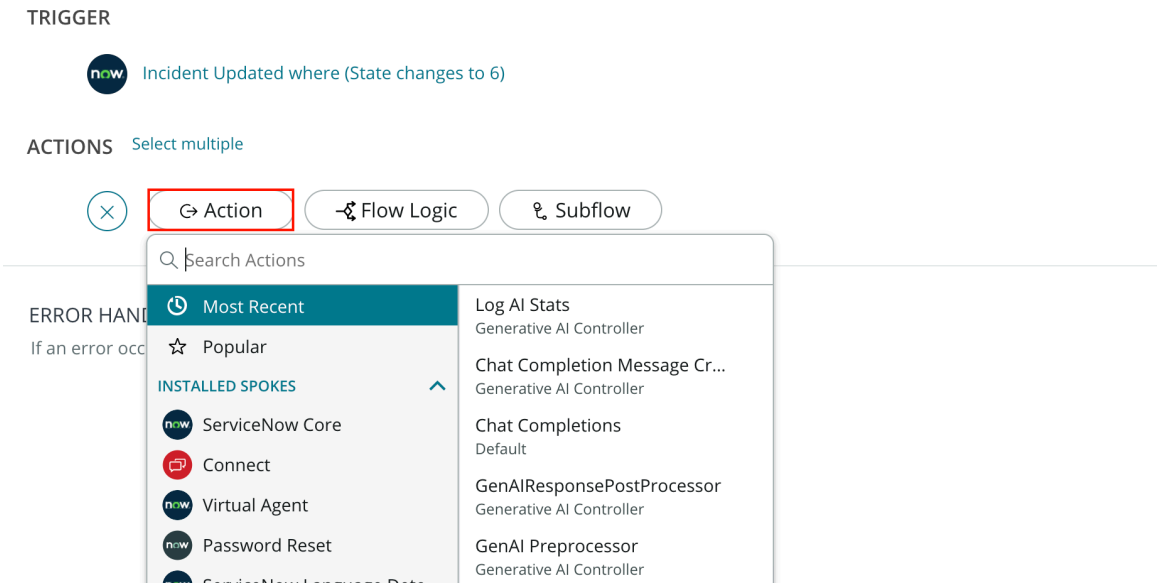
7. Select the table that the flow should run on.
   For this example, select the Incident table.
8. Select the **Add filters** button to add a condition for the flow.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The Condition is `State changes to Resolved`.

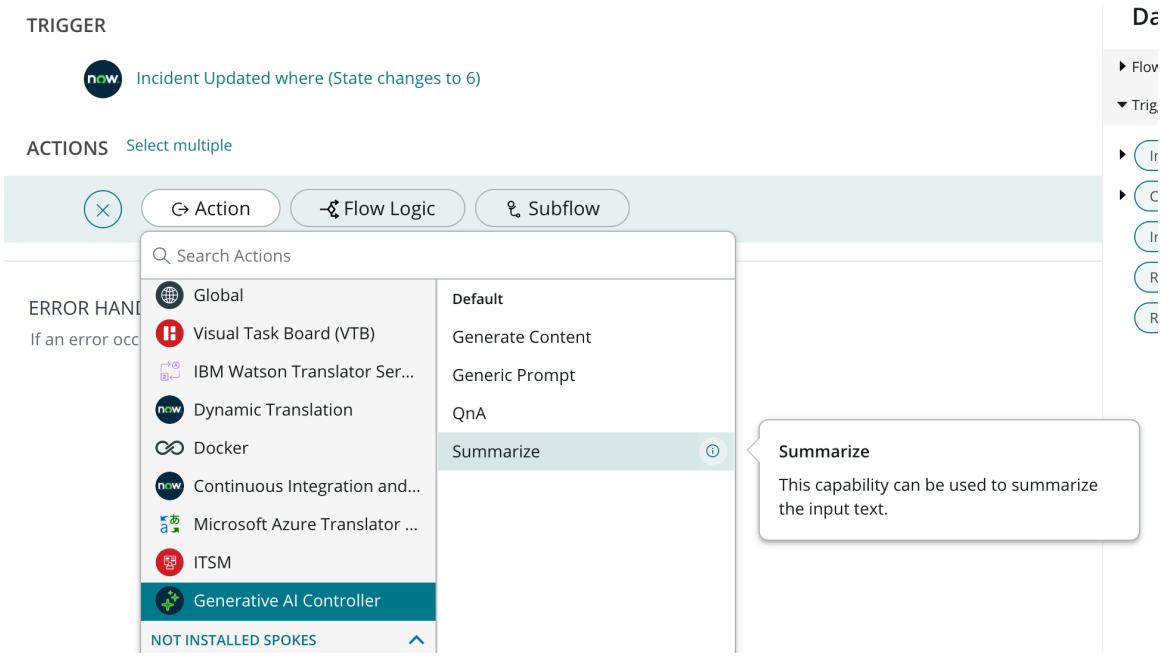


9. Select **Done**.
10. Under Actions, select **Add an Action**, **Flow Logic**, or **Subflow** and then select **Action** to open the Actions panel.



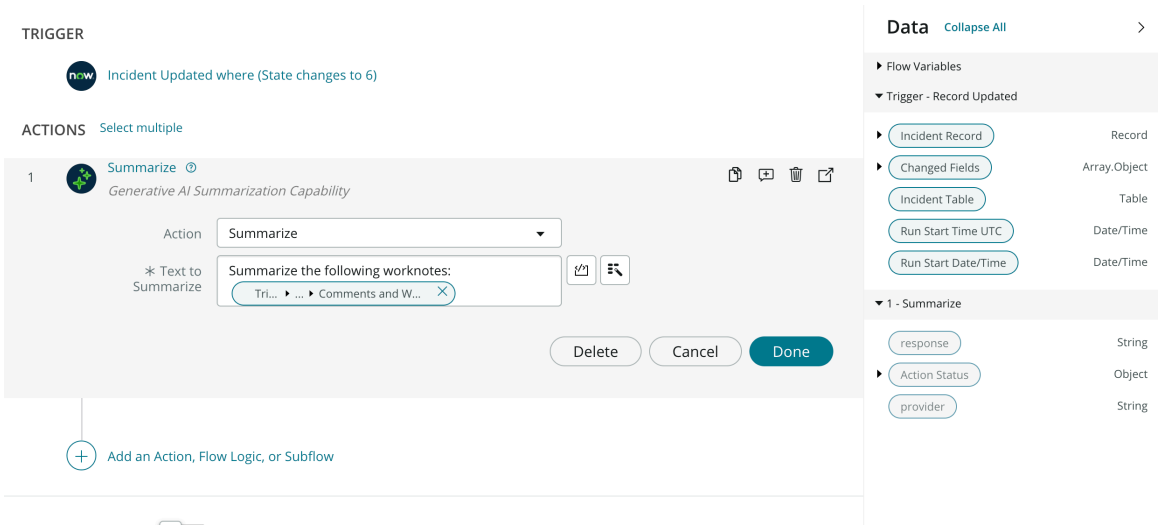
11. Select the **Generative AI Controller** > **Summarize** action and select it.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



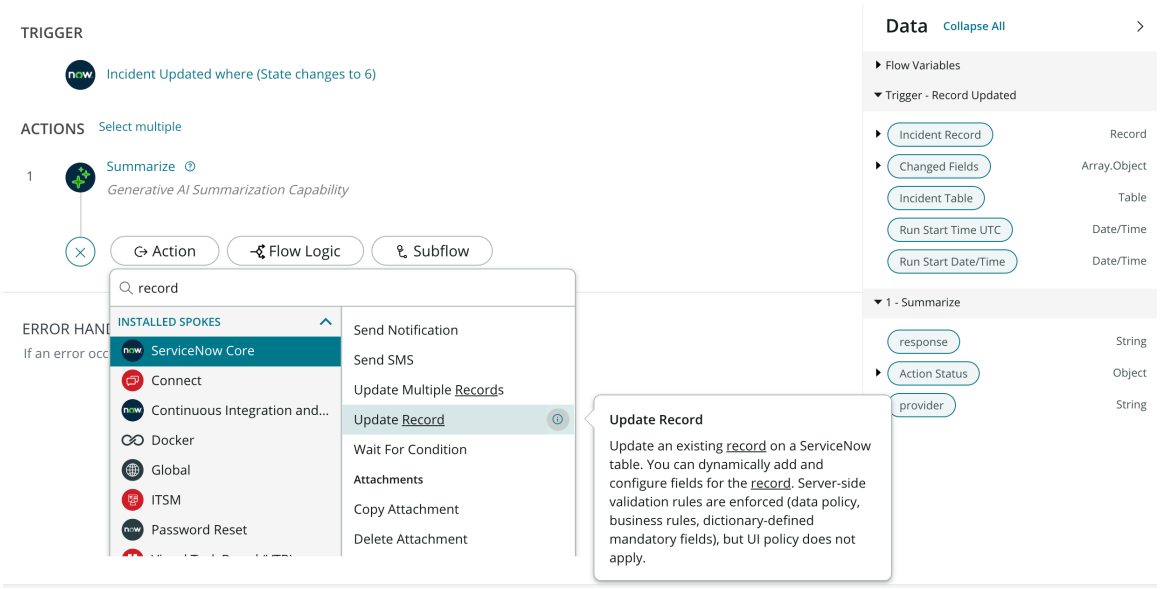
12. Drag a data pill from the Data panel or select the data pill picker icon ( ) to search. In this case, drag the **Trigger** > **Incident Record** > **Comments and Work notes** pill.

> **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.



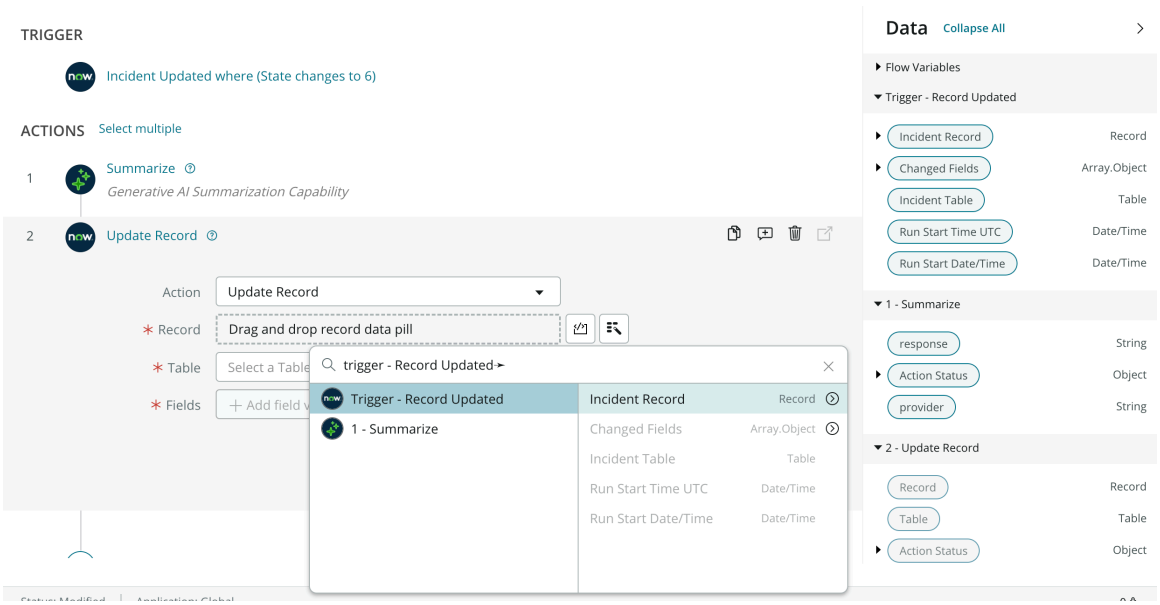
13. Select **Done**.
14. Add an action by selecting **Add an Action, Flow Logic, or Subflow**.
15. In the action selector, search for **Update Record** and select the action.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



16. Drag the **Trigger** > **Incident Record** data pill into the **Record** field or select the icon to search.

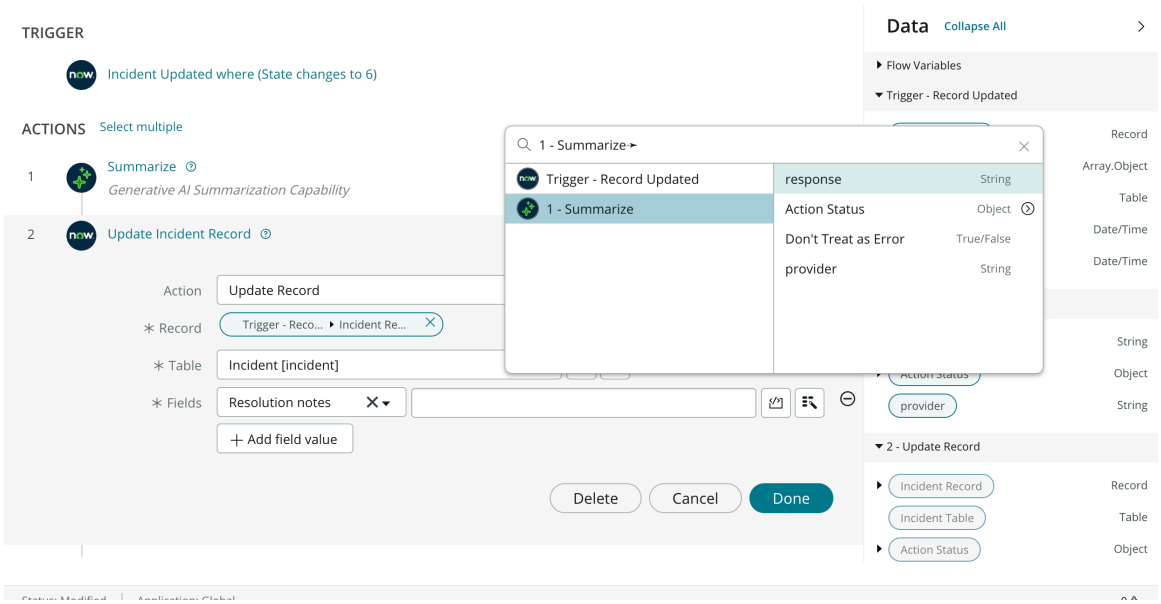
The Table field is filled in automatically based on the Trigger record.



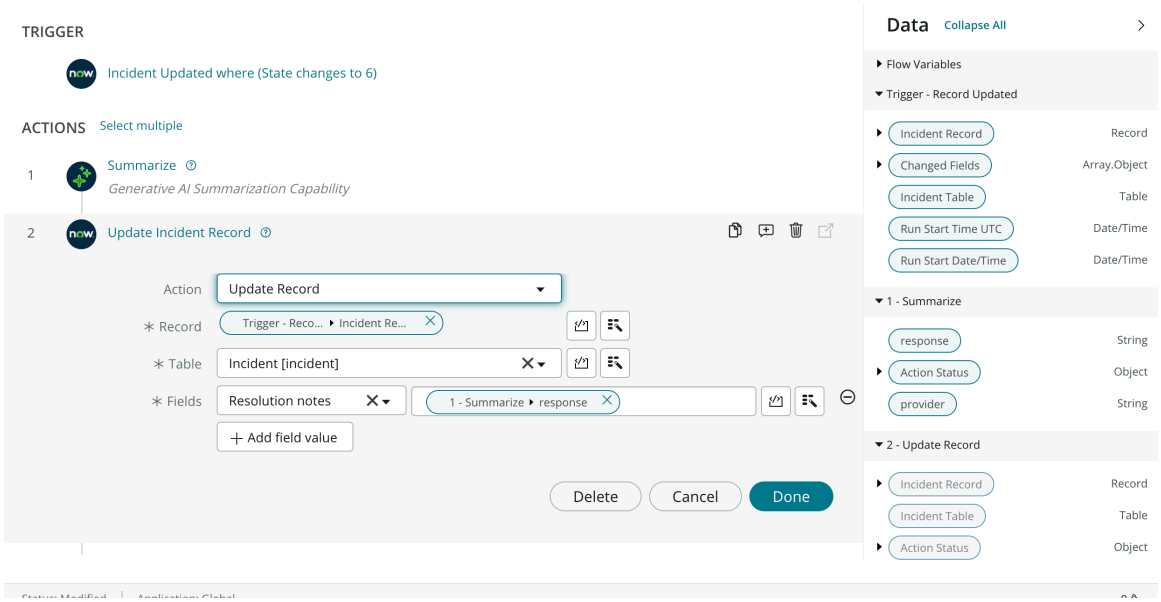
17. Select **Add field value** and search for the fields that you want to update.
In the following example, select **Resolution notes**.

18. Add the **Summarize** > **Response** data pill into the data field by dragging it from the Data panel or with the data pill picker.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



19. 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information on tracking Generative AI Controller usage, see Monitoring Now Assist usage.

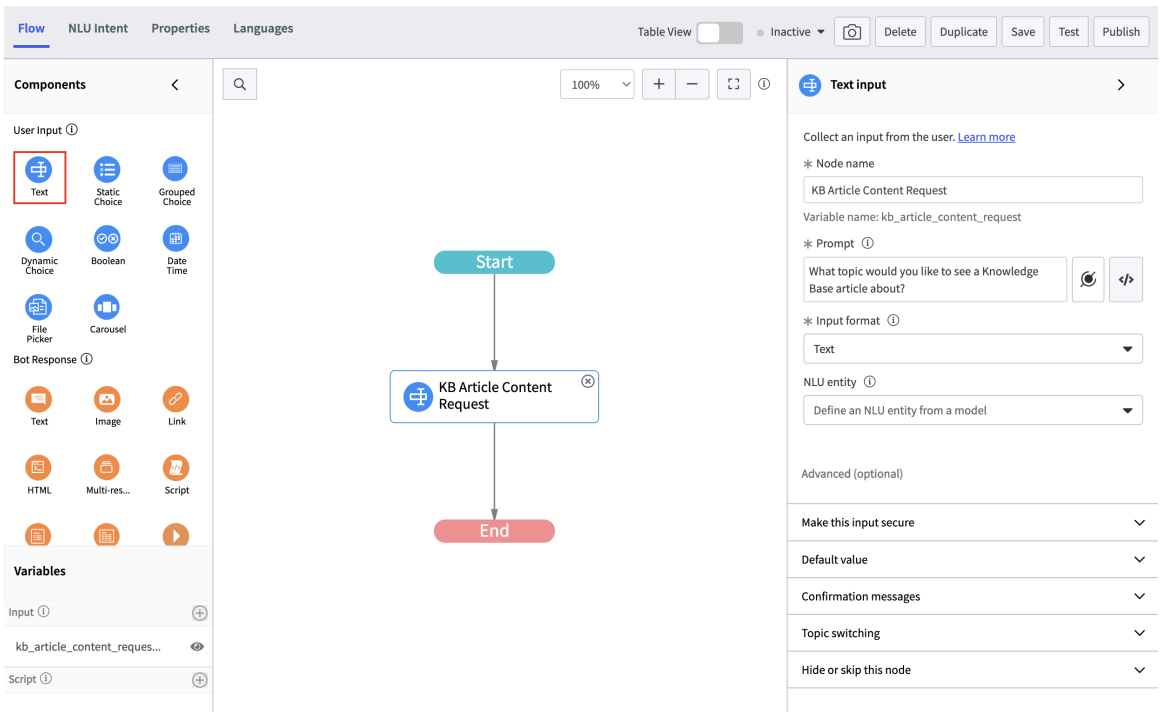
Procedure

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`.



7. In the prompt field, enter a message to display to the user.

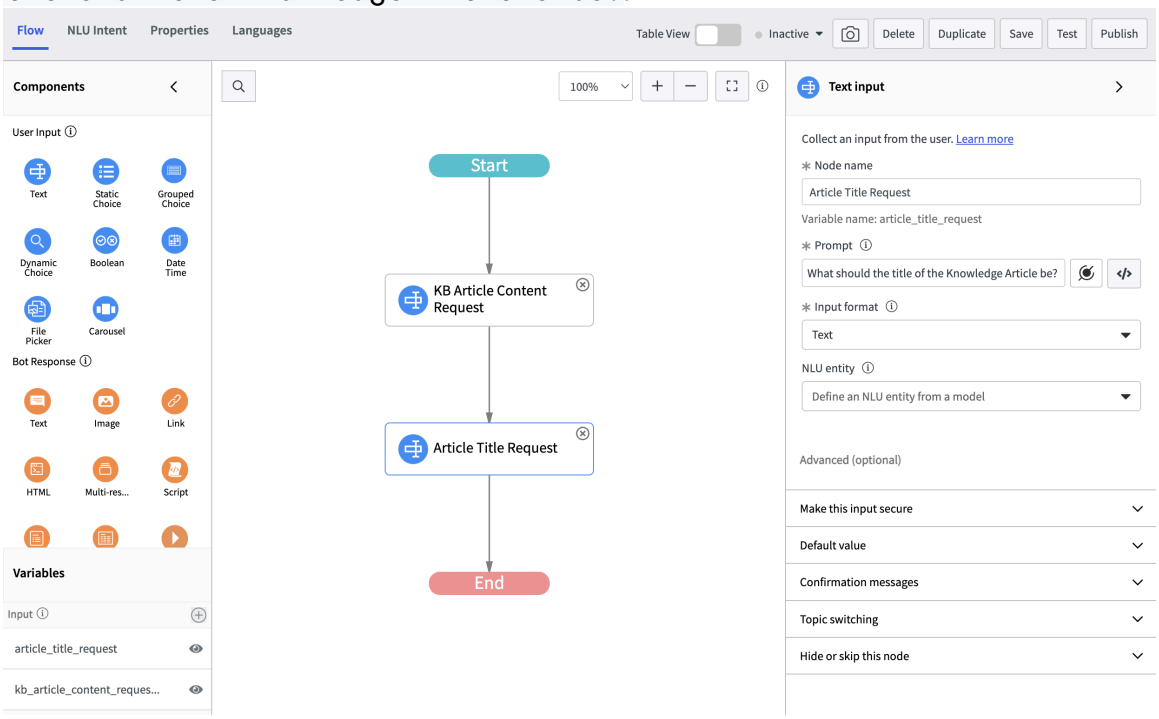
   Because the user supplies the topic for the knowledge base (KB) article content, the prompt in this example is `What topic would you like to see a Knowledge Base article about?`.
8. Add User Input nodes to fill in all required fields.

> **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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Add another **User Input** > **Text** node and name it `Article Title Request`.

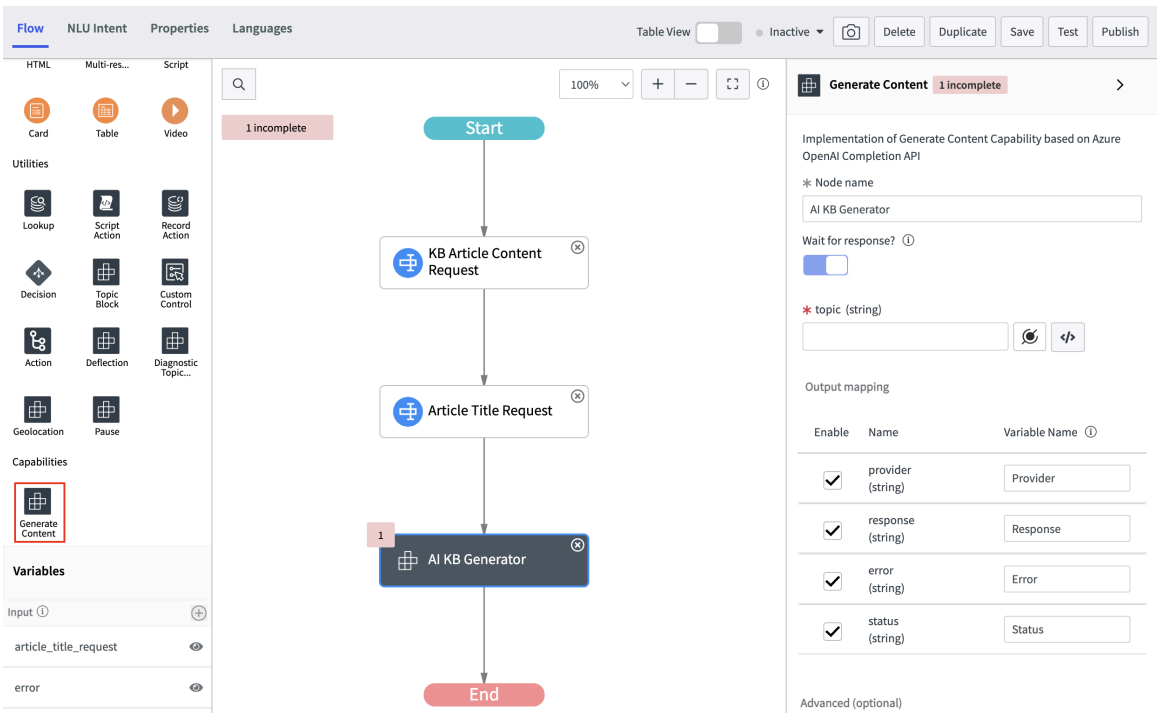
b. Add a prompt for the required `Article title` field such as `What should the title of the Knowledge Article be?`.



9. On the Components panel, scroll down until you see the Capabilities section and drag the Generate Content action next in the Topic flow.

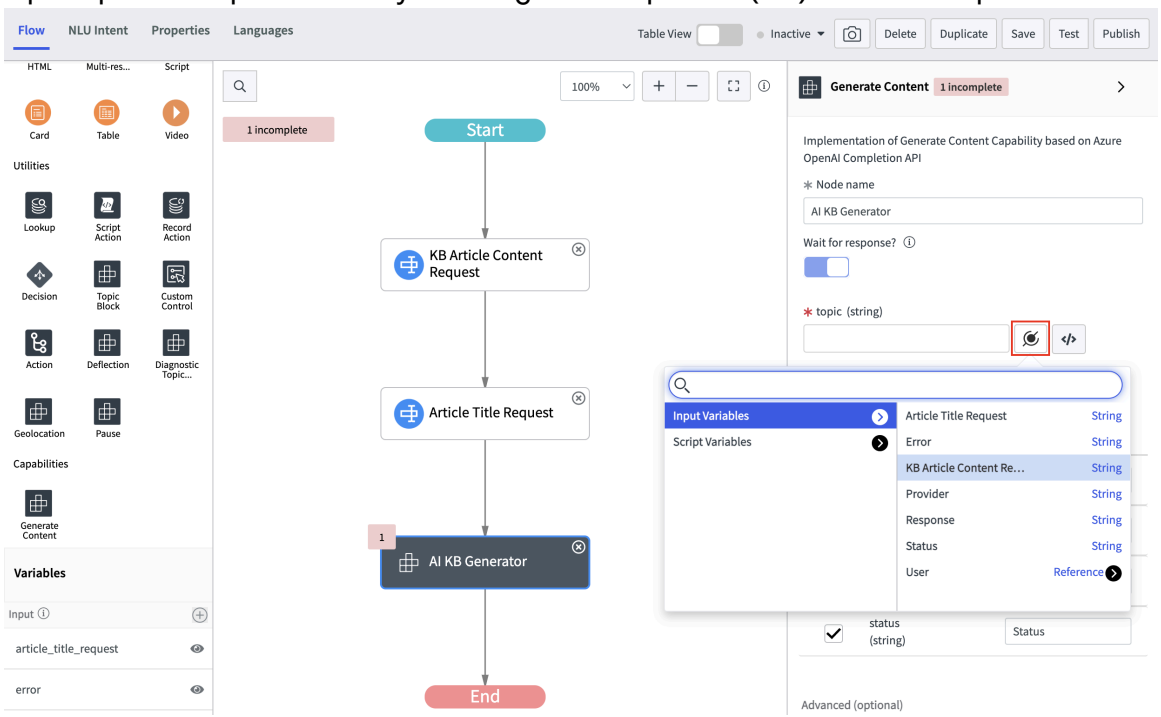
10. In the Info panel, name the node.

The node name that is used in the example is `AI KB Generator`.

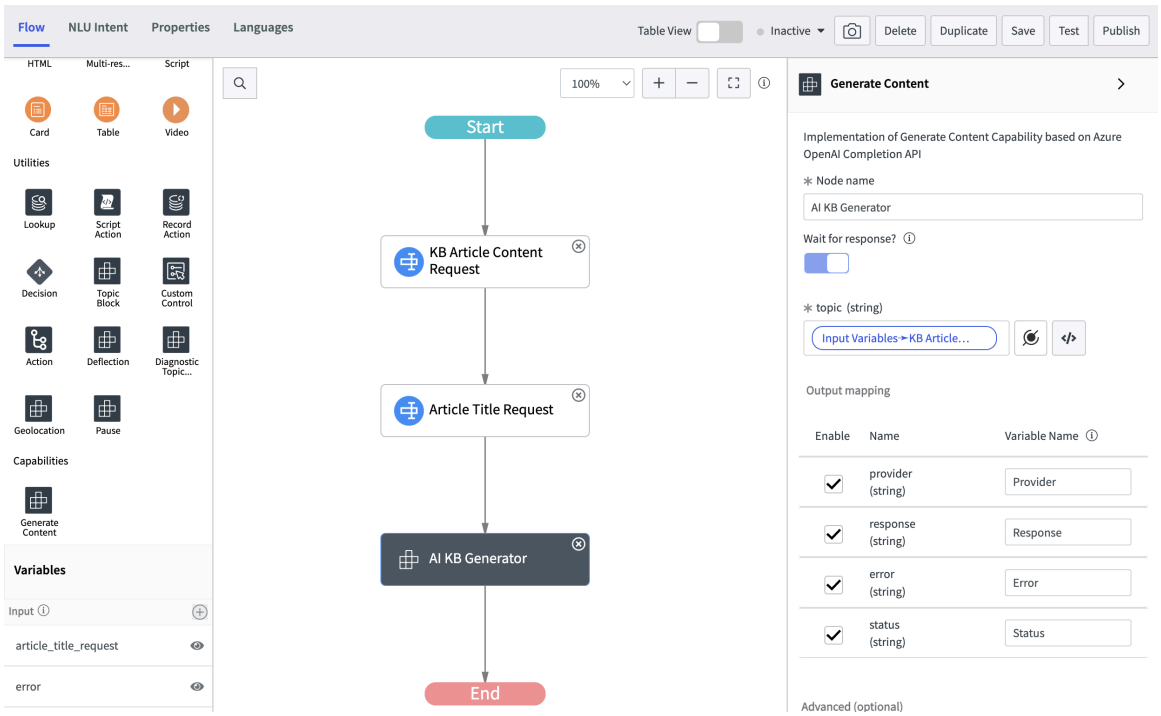


© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

11. Open up the data pill selector by selecting the data pill icon ( ) next to the topic field.

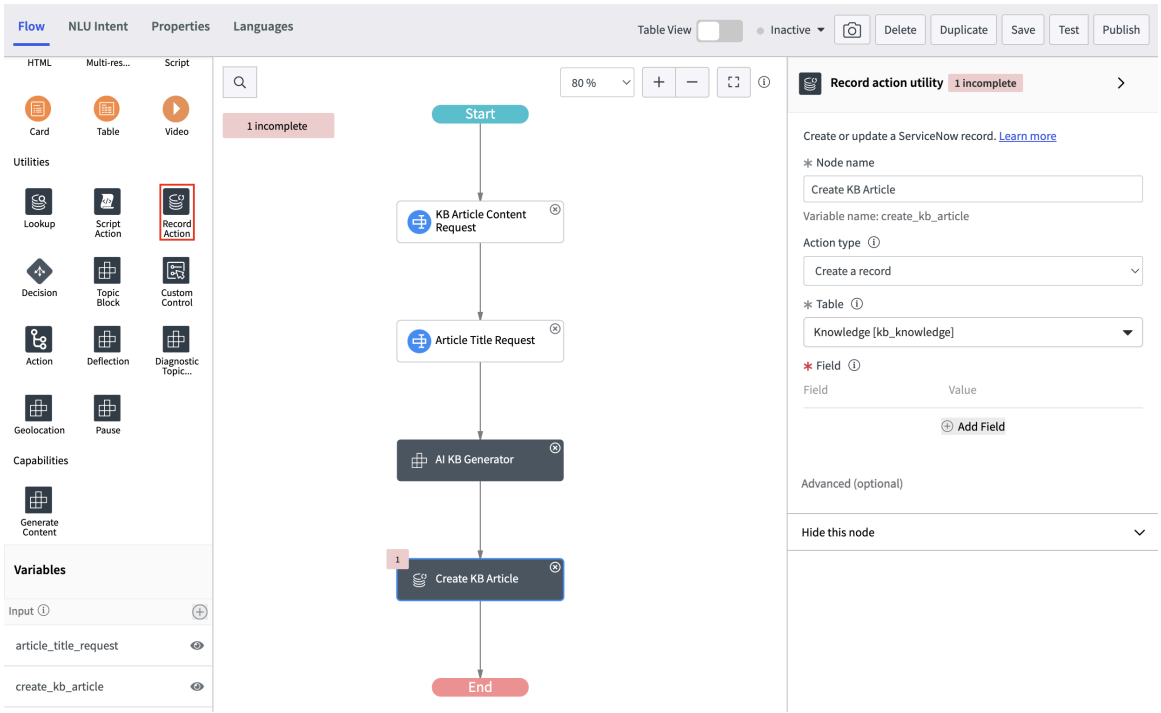


12. Use the text that was created in the interaction by selecting the **Input Variables** > **Response** variable.



13. Drag the **Utilities** > **Record Action** component from the Components panel after the Generate Content node.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



14. Enter a name for the node.

    For this example, the node name is `Create KB Article`.

15. Fill in the other required fields.

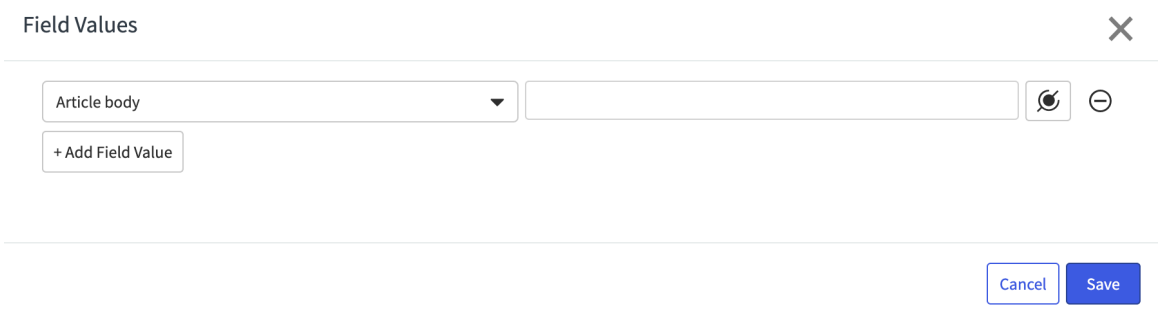
    a. For the Action type, select whether you want to create a record or update an existing one.

       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**.

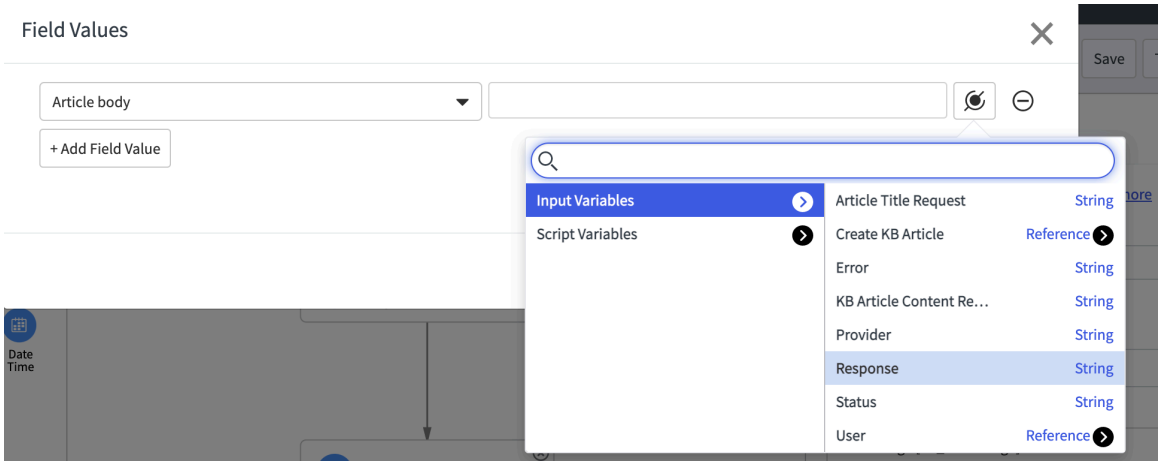


18. Select the field that you would like to update from the list.

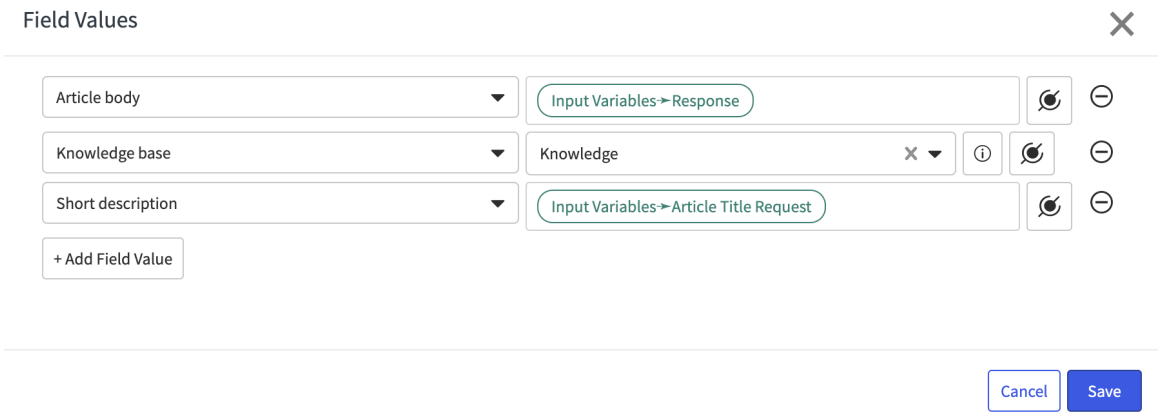
    You can add as many fields as you like. In this example, the fields are `Article body`, `Knowledge base`, and `Short description`.

19. Select the data pill icon to open the data pill selection menu.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



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 Flow Designer

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

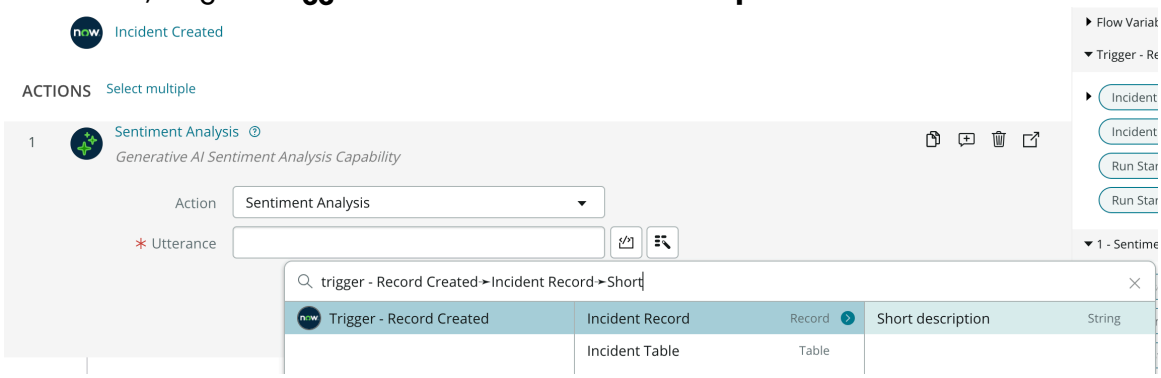
The Sentiment Analysis capability helps determine whether the user-generated text has a positive or negative sentiment. If a user is expressing negative statements such as "Level 1 support couldn't help me," you could upgrade the assignment group automatically for the case to help the user have a better experience.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

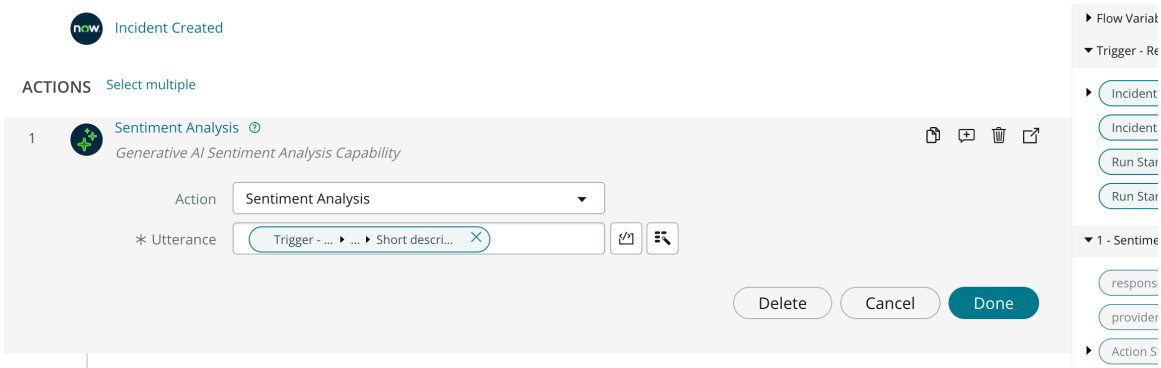
For more information on tracking Generative AI Controller usage, see Monitoring Now Assist usage.

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**.



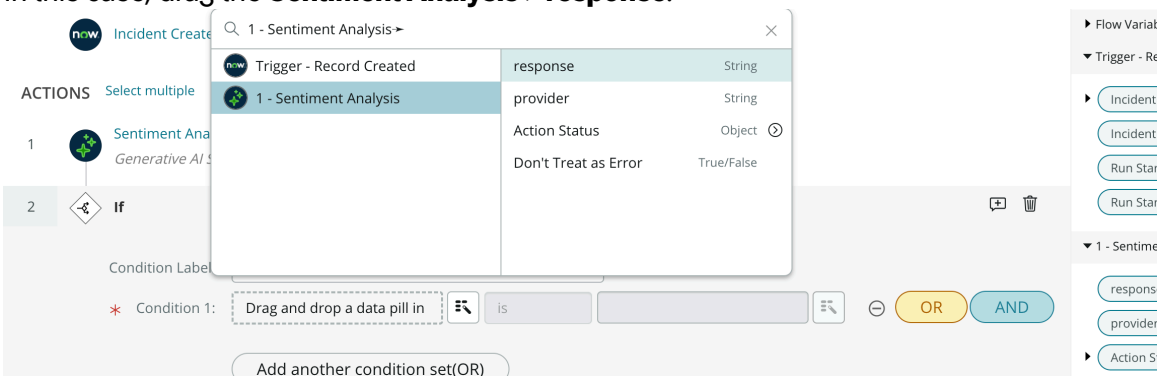
12. Select **Done**.



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.

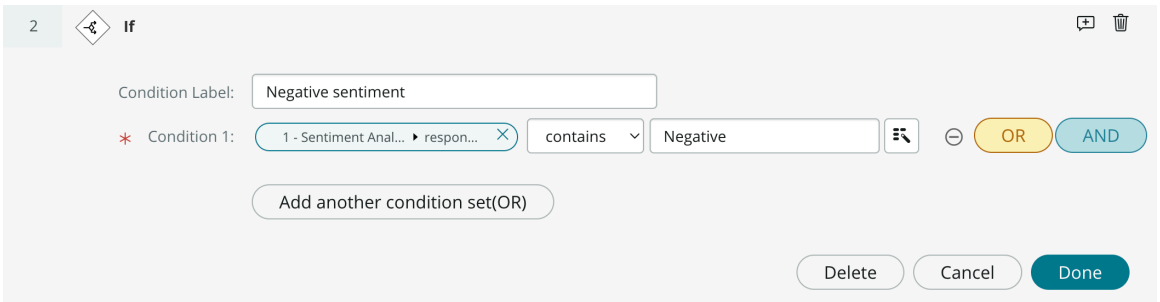
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In this 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`.

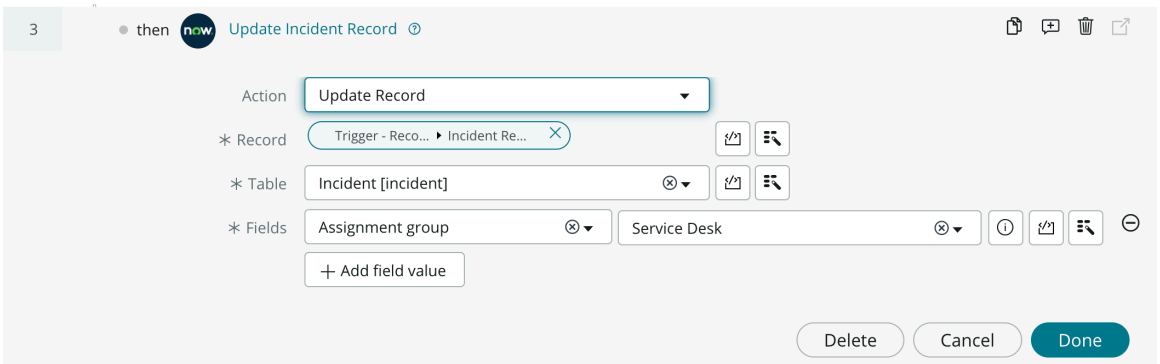


18. Search for **Update Record** and select the action.
19. Drag the **Trigger** > **Case Record** data pill into the **Record** field or select the icon to search.

    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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use 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	prompt: Prompt for the generative AI to respond to.
Generate Content	topic: Subject for content generation.
Sentiment Analysis	utterance: Text to analyze for user sentiment.
Summarize	textToSummarize: 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

6. **Optional:** 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 Flow Designer subflows.
OneExtend Capability Definition Attribute	sys_one_extend_definition_attribute	Input and output variables for Flow Designer 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 Flow Designer and Virtual Agent Designer.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Name	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 Flow Designer or Virtual Agent Designer topic, as well as any scripts such as business rules or UI actions.

Properties

Generative AI Controller properties

Property	Description
com.sn.generative.ai.provider	Default provider when capability definition has no default. Type: choice list No default value
com.sn.generative.ai.ais.message	Message that is displayed when AI Search fails to find an answer to a query. Type: string Default value: No answer found.
com.sn.generative.ai.log_prompt	Prompt that determines whether generative AI API calls are logged. Type: true | false Default value: true
com.sn.generative.ai.moderation.message	Message that is displayed if the OpenAI or Azure OpenAI moderation tools identify the content that goes against their terms of service. Type: string Default value: The response cannot be displayed as it’s deemed inappropriate by OpenAI.
com.glide.one.extend.token.buffer	Buffer that checks the request for the number of tokens before a OneExtend capability is executed. The maximum allowed request tokens are calculated based on the maximum tokens that are permitted by the AI provider's API minus the response token and buffer value that is specified in this system property. Type: integer

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Generative AI Controller properties (continued)

Property	Description
	Default value: 250

External links

Provider	Data policy	Usage policy
OpenAI	API data usage policies	Usage policies
Microsoft Azure OpenAI	Data, privacy, and security for Azure OpenAI Service	Code of conduct for Azure OpenAI Service
Google Cloud	Google Cloud Platform Terms of Service	Google Cloud Platform Terms of Service
Aleph Alpha	Data privacy	Terms and Conditions

Generate Content action

Use the Generate Content action with Flow Designer 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*

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Details about the success of the action.

Generic Prompt action

Use the Generic Prompt action for Flow Designer 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 Flow Designer 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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 Flow Designer 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Provider

Data type: *String*

Name of the AI model that generated the response.

Action Status

Data type: *String*

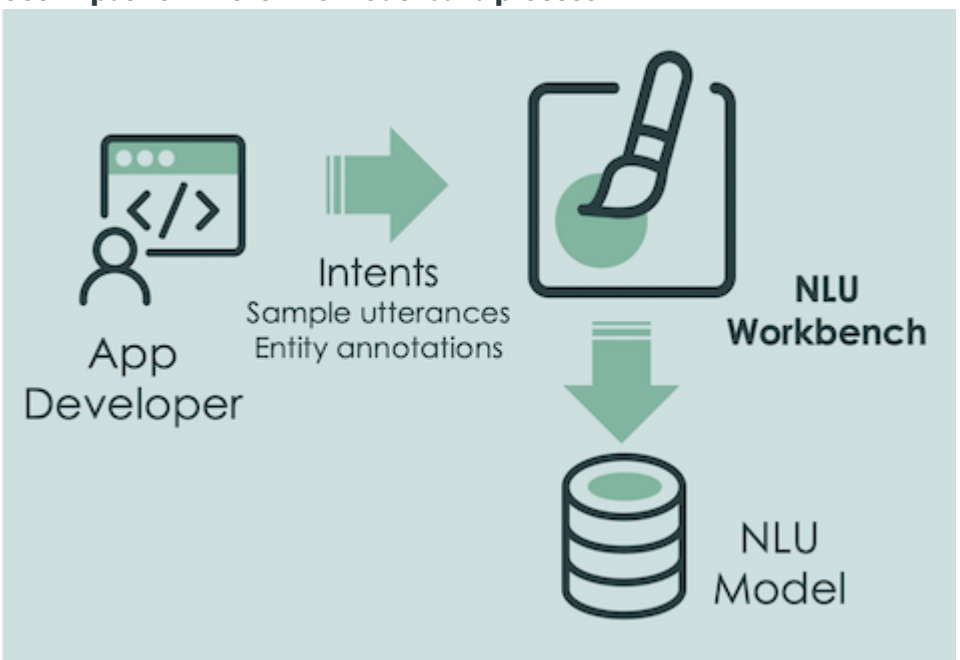
Details about the success of the action.

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

User input flow in the NLU model build process



This image shows you the user input flow in the NLU model build process.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Get started

Explore	Use	Integrate
 Learn about NLU concepts and features.	 Create, test, translate, and publish your NLU models.	 Learn how Virtual Agent administrators can update NLU models from within Virtual Agent Designer .
		Reference  Learn about using models in different languages for use in other applications.

Troubleshoot and get help

- Virtual Agent & NLU community page
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Natural Language Understanding

ServiceNow® Natural Language Understanding (NLU) provides an NLU Workbench and an NLU inference service that you can use to enable the system to learn and respond to human-expressed intent. By entering natural language examples into the system, you help it understand word meanings and contexts so it can infer user or system actions.

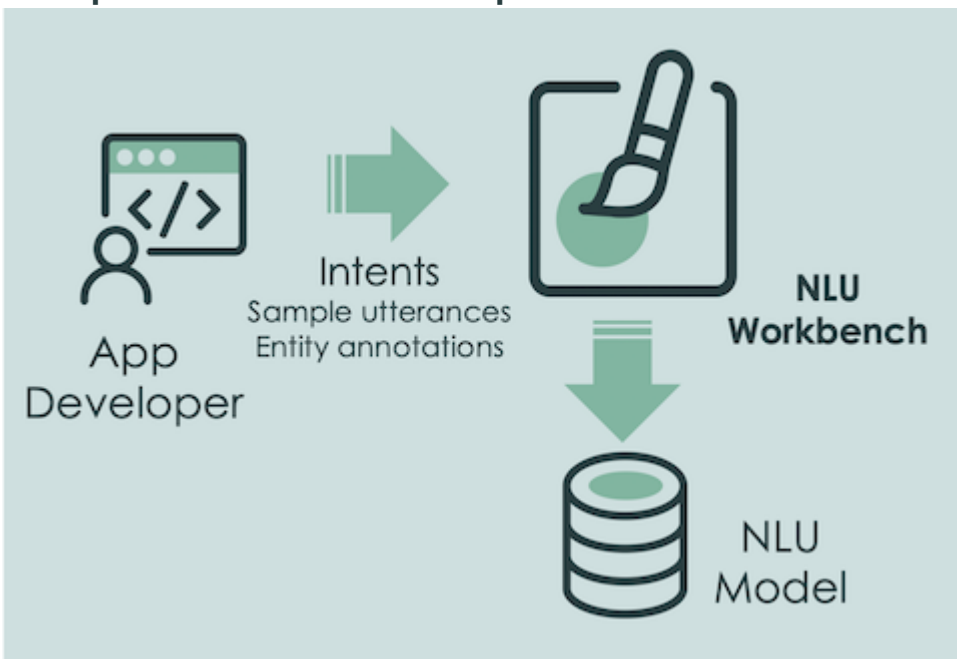
Natural Language Understanding overview

Natural Language Understanding (NLU) is a component of artificial intelligence natural-language processing that deals with machine-reading comprehension. Predictive Intelligence uses NLU so the system can understand word meanings and word contexts to infer user or system actions.

Natural Language Understanding workflow

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

User input flow in the NLU model build process



This image shows you the user input flow in the NLU model build process.

Natural Language Understanding benefits

Benefit	Feature	Users
Use the NLU Workbench to create models for Virtual Agent and AI Search.	NLU Workbench	NLU Administrator
Build the model's content by adding intents, entities, vocabulary, and test set utterances. Your NLU model content determines how the model responds to user inputs.	Model Management	NLU Administrator
Add a language to an existing NLU model by translating it. Use one of several translation options to add a secondary model in a supported language.	Model Management	NLU Administrator
Virtual Agent administrators can access and update their NLU models from within the Virtual Agent Designer user interface.	Virtual Agent Integration	Virtual Agent Administrator

Activate the NLU Workbench

Activate the following plugins to activate the NLU Workbench.

Before you begin

Role required: admin

About this task

Activate the following plugins if they aren't already active in your instance.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugin	Description
NLU Workbench - Core com.glide.nlu	Installs the required tables for persisting NLU models that are created using the NLU Workbench.
NLU Workbench com.snc.nlu_studio	Enables the creation of Natural Language Understanding (NLU) models. These models can understand the intent (action) and entities (details about the action) for a given user utterance. Any ServiceNow application can invoke an NLU model. Requires the NLU Workbench - Core plugin and the NLU Common Model plugin.
Predictive Intelligence com.glide.platform_ml	Enables the NLU Service APIs used for model creation and inference. Enables the creation of machine learning solutions that are trained on data in your instance. Provides frameworks for classification, similarity, and clustering. A trained solution can be invoked by any application by using a prediction API.
NLU Common Model com.glide.nlu.model	Packages all language NLU common models. Also includes commonly used pattern entities that can be imported and used in any NLU model in the NLU Workbench. Commonly used patterns like email, phone, and ServiceNow specific pattern entities such as INT, RITM are made available. Requires the NLU Workbench - Core plugin.
NLU Active Learning-Properties com.glide.nlu.active_learning_properties	Enables the nlu_admin to configure the system properties for the Expert Feedback Loop application. If you don't use the Expert Feedback Loop application, this plugin is not used and can safely be ignored. For more information about the system properties for the Expert Feedback Loop, see the NLU Expert Feedback Loop documentation.

The following two plugins are for apps associated with a for-fee subscription and are available on the ServiceNow Store. Installing these apps adds additional features to the NLU Workbench. Contact your account manager if you are interested in these apps.

Plugin	Description
Intent Discovery sn_nlu_discovery	Delivers the Intent Discovery feature, which identifies user intents by analyzing incident/case data. Use this application to help identify which intents to model and build for Virtual Agent conversations to attain maximum deflection. For more information, see Install Intent Discovery.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Plugin	Description
NLU Workbench - Advanced Features sn_nlu_workbench	Delivers the Model Performance, Multi-model Batch Testing, Cross-model Conflict Review, and Expert Feedback Loop features. For more information, see Install NLU Workbench - Advanced Features.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the following plugins using the filter criteria and search bar: NLU Workbench - Core (com.glide.nlu), NLU Workbench (com.snc.nlu_studio), Predictive Intelligence (com.glide.platform_ml), and NLU Common Model (com.glide.nlu.model).

   You can search for the plugin by its name or ID. If you cannot find a plugin, you might have to request it from ServiceNow personnel.
3. Select **Install** to start the installation process.

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

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

NLU Workbench roles

NLU Workbench is installed with these roles.

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

NLU user [nlu_user]

Users who have read access to models in NLU Workbench.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Important:
Avoid granting an admin role when more specialized roles are available.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This 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

Important:
Avoid granting an admin role when more specialized roles are available.

This role is installed with NLU Workbench - Core.

NLU administrator [nlu_admin]

Users who can use NLU Workbench to manage NLU models.

Contains Roles

List of roles contained within the role.

- sn_ace.ace_user
- sn_nlu_workbench.nlu_feedback_admin
- nlu_editor
- nlu_user

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Important:
Avoid granting an admin role when more specialized roles are available.

This role is installed with NLU Workbench - Core.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

NLU 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

Important:
Avoid granting an admin role when more specialized roles are available.

This role is installed with NLU Workbench - Advanced Features. For more information, see NLU Workbench - Advanced Features.

This role is used in Expert Feedback Loop modules. For more information, see NLU Expert Feedback Loop or Issue Auto Resolution Tuning in NLU.

NLU Workbench properties

Refer to these system properties for the Natural Language Understanding (NLU) application.

NLU Workbench properties and their usage

To access your system properties, use the admin or nlu_admin role and the following path in the application navigator: **All** > **NLU Workbench** > **Settings**.

Model Settings

Label and Name	Default value	Plugin	Recommen usage
Maximum number of utterances per intent *glide.nlu.utterances_per_intent.value_limit*	200	NLU Workbench	Use fewer than 200 utterances per intent t keep your model well balanced in terms of intent size.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Model Settings (continued)

Label and Name	Default value	Plugin	Recommen usage
			Note: Value must b greate than 5 and le than o equal 300.
Maximum number of records in a Table vocabulary source *glide.platform_ml.api.max_nlu_lookupsource_records*	100,000	NLU Workbench	Keep the value unde 100,000.
Maximum number of values in a List vocabulary source *glide.nlu.static_lookup.value_limit*	1,000	NLU Workbench	Keep the value unde 1,000.
Enable pre-built vocabulary for software names *glide.mlpredictor.option.nlu.@LookupSources:software*	enabled	NLU Workbench	Enable pre-built vocabulary so the system ca recognize software names.
Enable pre-built vocabulary for hardware names *glide.mlpredictor.option.nlu.@LookupSources:hardware*	enabled	NLU Workbench	Enable pre-built vocabulary so the system ca recognize hardware names.

Advanced Settings

Label and Name	Default value	Plugin	Recommended usage
Maximum number of records for Intent Discovery classification *sn_nlu_discovery.intent_discovery_max_classification_limit*	300,000	Intent Discovery	Keep the number of records less than 500,000.
Minimum number of records for Intent Discovery classification *sn_nlu_discovery.intent_discovery_min_classification_limit*	10,000	Intent Discovery	Use at least 10,000 records to get high quality results.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced Settings (continued)

Label and Name	Default value	Plugin	Recommended usage
Minimum number of records for NLU performance analysis *sn_nlu_workbench.glide.nlu.performance.min_clustering_records*	5,000	NLU Workbench - Advanced Features	Use at least 5,000 records to get high quality results.
NLU Conflict Detection - Moderate Threshold *sn_nlu_workbench.glide.nlu.conflict.moderate_threshold*	.85	NLU Workbench - Advanced Features	Must be a decimal between 0 and 1. Keep this threshold less than the Critical Threshold.
NLU Conflict Detection - Critical Threshold *sn_nlu_workbench.glide.nlu.conflict.critical_threshold*	.95	NLU Workbench - Advanced Features	Must be a decimal between 0 and 1. Keep this threshold greater than the Moderate Threshold.
The maximum number of rows in a batch test import file *sn_nlu_workbench.glide.nlu.batch_test.max_import_rows*	10,000	NLU Workbench - Advanced Features	Make sure your batch test import file has no more than 10,000 rows.
The maximum number of utterances to display for feedback in the expert feedback loop *glide.mlpredictor.option.nlu.activeLearning.label_candidate_table.max_respon*	300	NLU Workbench - Advanced Features	Pull no more than 300 utterances from your users' Virtual Agent chat logs to display for feedback in the Expert Feedback Loop application.The minimum umber of utterances a user should review before tuning the model
The minimum number of utterances a user should review before tuning the model *sn_nlu_workbench.glide.nlu.optimize.min_labeled_data*	100	NLU Workbench - Advanced Features	Provide and save feedback for at least 100 utterances from your users' Virtual Agent chat logs so you can execute the Tune Model feature in the Expert Feedback Loop application.
The maximum number of records to fetch from Virtual Agent chat logs *glide.mlpredictor.option.nlu.activeLearning.va_chat_logs.max_row_limit - 3000*	3,000	NLU Workbench - Advanced Features	If there is high NLU usage, increasing the default value to a maximum of 50,000 records will increase the data available for the active learning job to filter up on and display in the Expert Feedback Loop application to give feedback on.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Advanced Settings (continued)

Label and Name	Default value	Plugin	Recommended usage
Size limit on Label Candidate Table (used for pruning the table) *glide.mlpredictor.option.nlu.activeLearning.label_candidate_table.max_data_s - 10000*	10,000	NLU Workbench - Advanced Features	The recommended usage for this property is the same as the property above.
Size limit on Labeled Data Table (used for pruning the table) *glide.mlpredictor.option.nlu.activeLearning.label_table.max_data_size - 10000*	10,000	NLU Workbench - Advanced Features	The recommended usage for this property is the same as the property above.
Enable this property to unblock your instance during NLU model training. The training will be scheduled for an off-peak time, and we will notify you when it's done. *glide.mlpredictor.scheduled.nlu.model.training*	False	NLU Workbench - Advanced Features	False

To get more feedback data from Virtual Agent (VA) chat logs, refer to the **Procuring additional VA feedback data on demand** section in the Expert Feedback Loop documentation.

NLU language support

The NLU Workbench application provides support for creating NLU models in different languages for use in other applications, such as Virtual Agent.

The platform supports NLU for 17 languages. Ten of these languages available to your models have both intent and entity support; the remaining 7 languages have intent-only support. (Japanese entity support includes character annotation.)

Available in Tokyo, Utah, and Vancouver	Intent and entity support (in Utah and Vancouver)	Intent-only support (in Utah and Vancouver)
• Brazilian Portuguese • Chinese (simplified) • Danish • Dutch • English • Finnish • French • French Canadian • German • Italian • Japanese • Korean	• Brazilian Portuguese • Dutch • English • French • French Canadian • German • Italian • Japanese • Portuguese • Spanish	• Chinese (simplified) • Danish • Finnish • Korean • Norwegian • Polish • Swedish

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Available in Tokyo, Utah, and Vancouver	Intent and entity support (in Utah and Vancouver)	Intent-only support (in Utah and Vancouver)
• Norwegian • Polish • Portuguese • Spanish • Swedish		

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This 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

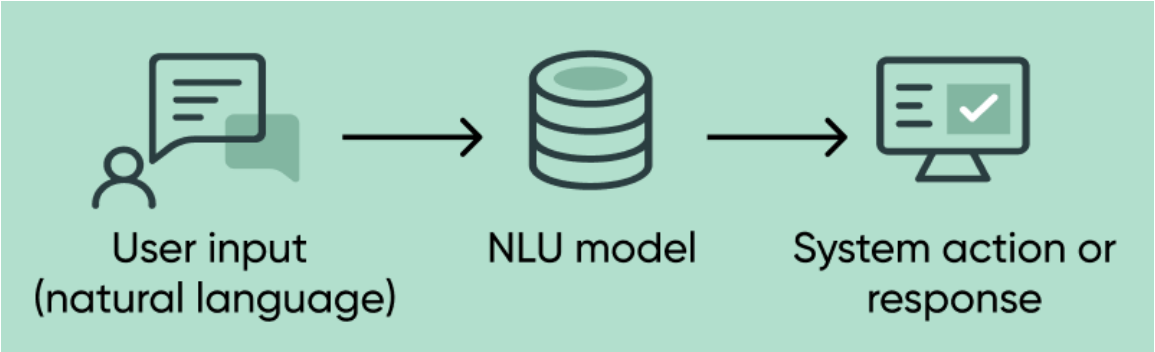
- Created the new DATE-TIME system entity (English only) for use in Virtual Agent.
- Added vocabulary source and entity support (simple, mapped, and open ended) for the Dutch and Italian languages, and system entity support for Portuguese and Brazilian Portuguese.
- Upgraded the ServiceNow Language Model 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Overview of how user input is transformed to a system action



A model is a collection of utterances, intents, entities, and vocabulary that the system uses to respond to natural language inputs from your users. The model takes the natural language input from your users and matches it to actions to be performed by the system.

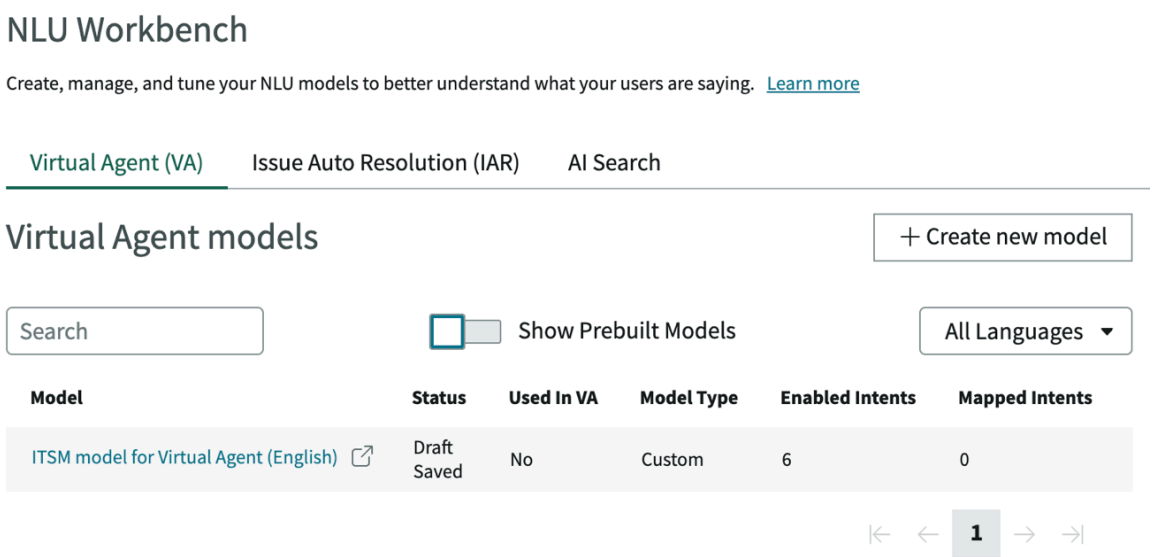
The utterances in your model data are examples of what your users might ask for. These examples are used to train your model to recognize which system actions to perform in response. These system actions are called intents.

Usage, roles, and navigation

Use the nlu_admin role to access the NLU Workbench.

Navigate to **NLU Workbench** > **Models**.

There are three tabs on the NLU Workbench homepage: Virtual Agent, Issue Auto Resolution, and AI Search. The Virtual Agent tab opens by default. Select the correct tab for your application.



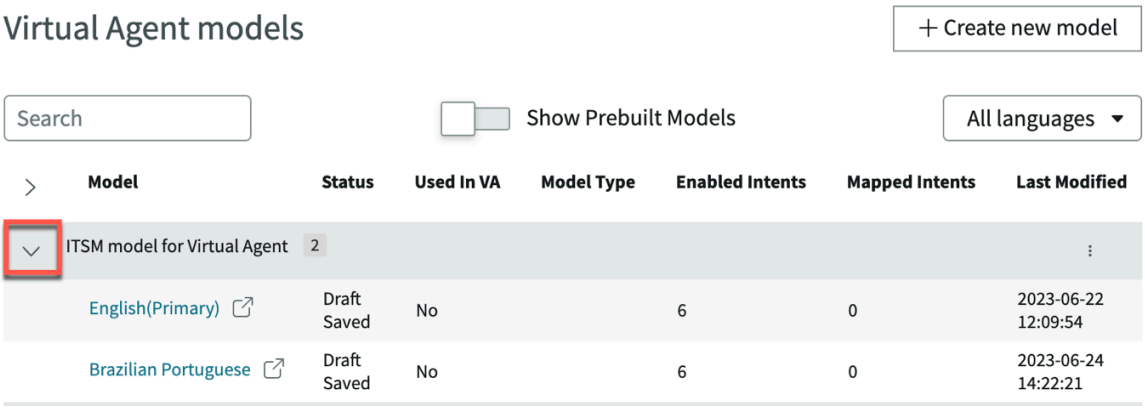
You can use the **Search** field to look for models by name.

The default setting of the **All Languages** filter displays all models regardless of their language. You can filter models by language using the list in **All Languages**. To learn more about the available languages, see NLU language support.

If you have a multilingual grouped model, its row in the list of models has an arrow on the left side. Select the arrow to expand the group so that all language versions of the model are

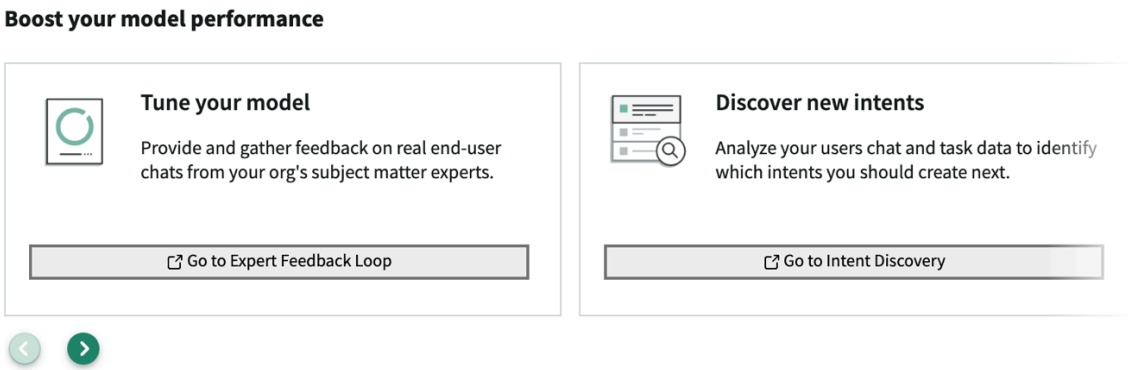
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

displayed. For more information on model languages and grouping, see Multilingual model management.



The **Used In VA** column indicates whether a model has already been linked to another application such as Virtual Agent. For AI Search, this column indicates whether the model has already been linked to a Genius Result. Similarly, the **Mapped Intents** column displays a count of intents that have been mapped or linked to Virtual Agent topics.

Scroll down the NLU Workbench homepage to view the **Boost your model performance** section. This section displays cards you can use to access functions that are available to your model type.



Model applications

The output of your NLU models can be consumed by the ServiceNow applications Virtual Agent, Issue Auto Resolution, and AI Search. The NLU Workbench homepage organizes each application's models by tab. In the list of models, different columns are displayed depending on the application.

NLU Models for Virtual Agent

For more information on Virtual Agent, see Virtual Agent.

- The tab for Virtual Agent displays a list of models. Select a model name to open an overview page for that model.
- The **Show Prebuilt Models** toggle is off by default. Switch this toggle on to display read-only models that can be copied and used as the basis for your custom models.
- The **Boost your model performance** section displays a total of 5 cards for functions available to Virtual Agent.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

NLU Models for IAR

Fewer functions and columns are displayed for IAR than for the other applications in NLU Workbench. This is because the IAR model is prebuilt and is configured in the IAR Admin console.

- When you select the name of your IAR model in NLU Workbench, you are taken to IAR Tuning rather than to a model overview. For more information, see Issue Auto Resolution Tuning in NLU.
- The **Show Prebuilt Models** toggle is not displayed.
- The **Create new model** button is not displayed because the prebuilt model is used directly.
- The section **Boost your model performance** is not displayed.

For more information on setting up IAR, see Using Issue Auto Resolution.

NLU Models for AI Search

- There is one prebuilt model for AI Search. You can copy this model as the basis for a custom model.
- The **Show Prebuilt Models** toggle is on by default.
- The **Create new model** button is available.
- The **Boost your model performance** section displays two cards for functions available to AI Search.

For an introduction to managing AI Search in the NLU Workbench, watch the following video. A video on how to manage AI Search in the NLU Workbench.

For more information, see the AI Search landing page.

What you can do next

To learn more about model content and the model life-cycle, see Model management.

Model management

Manage your NLU model's life cycle in the NLU Workbench. Model management phases guide you through the iterative process of building, testing, and publishing your model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Bringing your NLU model from creation to deployment requires multiple steps, separated into phases. You can return to earlier phases when you want to adjust and maintain your model.

The phases available for your model depend on the model's application. The system will display a phase, button, or function only when it applies to your model's application.

To use the model management phases, ensure you have all the necessary NLU plugins. For more information, see the following resources:

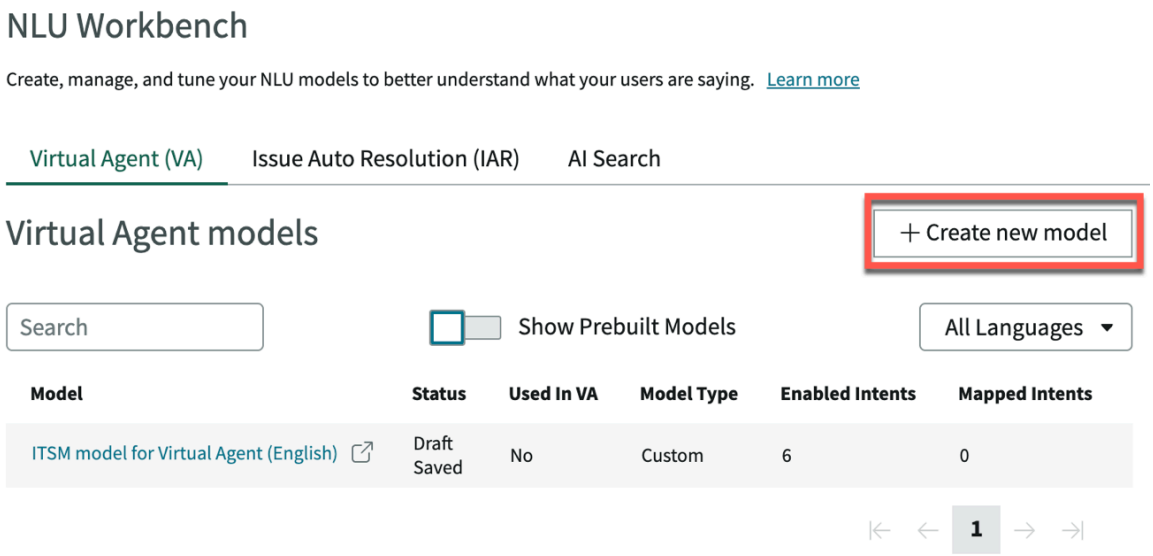
- Activate the NLU Workbench
- Install NLU Workbench - Advanced Features
- Install Intent Discovery

Note:
NLU Workbench - Advanced Features and Intent Discovery are available from the ServiceNow Store. NLU model testing and performance monitoring require NLU Workbench - Advanced Features.

Create a model

To create a model for Virtual Agent or AI Search, navigate to **NLU Workbench** > **Models**. The Virtual Agent tab opens by default. Select the appropriate tab for the model you want to create.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



You can choose different ways to create a model:

- **Use prebuilt model**: Copy one of the included read-only models, and add content specific to your business.
- **Import data from CSV**: Upload a CSV file that contains training utterances and matched intents.
- **Start from blank**: Go through the process of setting up a new model from scratch.

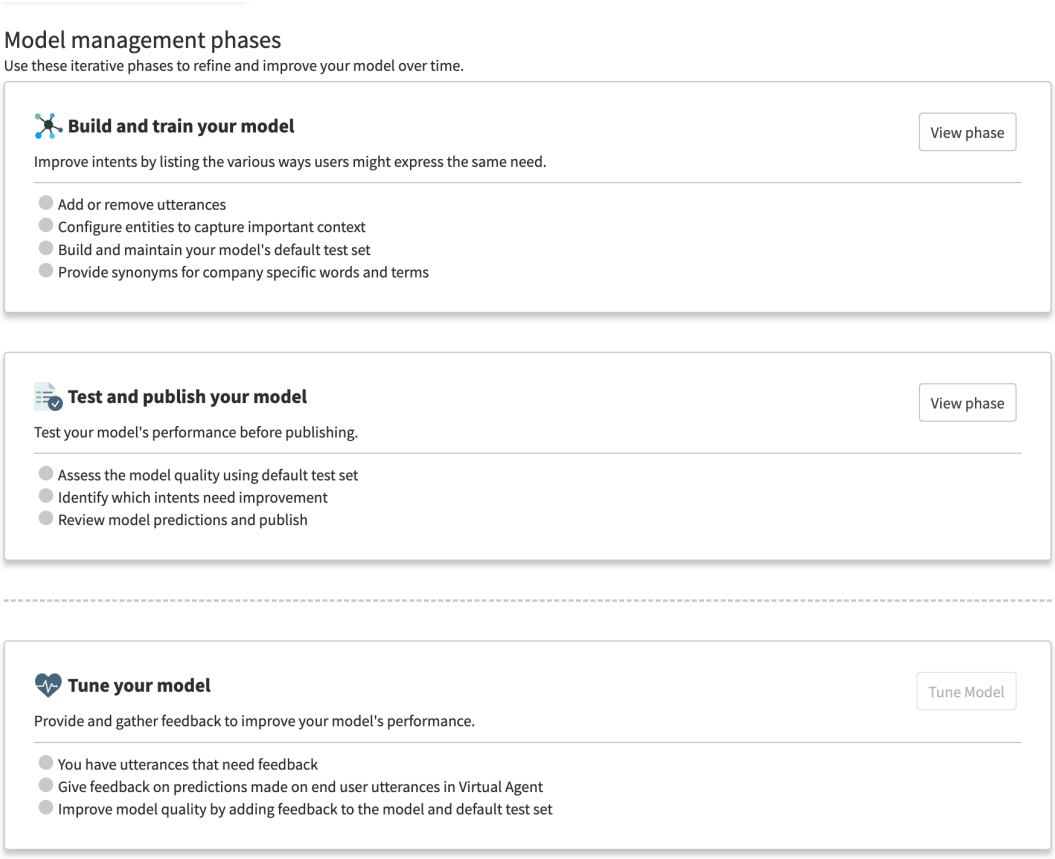
To get started, see Creating models.

Model management phases

After creating a model, access its management phases by navigating to **NLU Workbench** > **Models**. Select the tab for your model's application, then the name of the model to open the **Model details** page on the model overview.

There are three phases on a Virtual Agent model's overview page: Build and train your model, Test and publish your model, and Tune your model. These phases guide you as you build and improve your model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Build and train your model

Build the model by adding and managing content:

- Intents: Add more intents to broaden the range of user requests that your model can understand.
- Entities: Add more entities so that your model can extract more contextual details from your users' requests.
- Vocabulary: Add vocabulary to enable the model to better understand words and phrases that are specific to your business, such as industry terms and acronyms.
- Test set: Add test utterances and their expected intents to your model's default test set.

To learn more, see Build and train your model.

Train your model using utterances that the model is likely to encounter from your users. To learn more, see Train and try your NLU model.

Test and publish your model

Test your model to gauge the performance and identify areas for improvement.

> **Note:**
> Model testing requires the NLU Workbench - Advanced Features store application. For more information, see NLU Workbench - Advanced Features.

For more information on testing and thresholds, see Test and publish your model.

When you're satisfied with the results of testing, publish your model to make it available for use by other applications. For more information, see Publish your NLU model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Tune your model

If NLU Workbench - Advanced Features is installed, and your model is created for Virtual Agent, the **Tune your model** phase is enabled. With this phase, you can use **Expert Feedback Loop** to incorporate actual user utterances into your model.

For more information, see Tune your model.

If your model is created for Issue Auto Resolution, you will be taken to IAR Tuning by selecting the name of your model in the IAR tab of the NLU Workbench homepage. For more information, see Issue Auto Resolution Tuning in NLU.

Model settings

Use the **Settings** page of the model overview to change the name and description of the model. You can also modify the confidence threshold of the model. The confidence threshold determines how confident the model must be to predict an intent.

For more information, see NLU model settings.

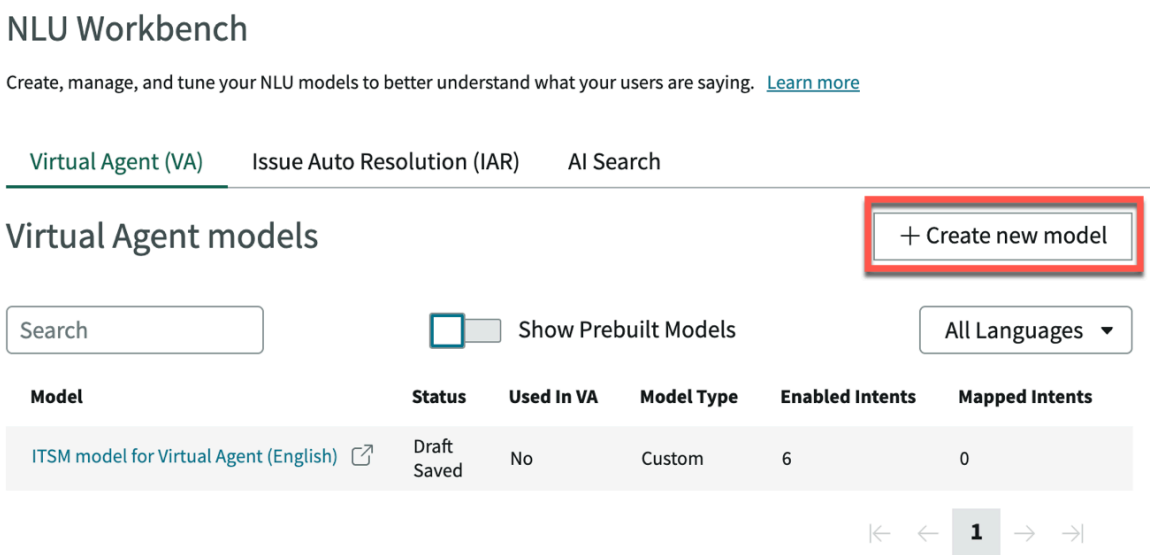
Creating models

Creating models is the first step to taking advantage of Natural Language Understanding (NLU) in your instances. Create models for Virtual Agent and AI Search in the NLU Workbench.

You create NLU models for consumption by the ServiceNow applications Virtual Agent or AI Search.

Note:
For Issue Auto Resolution, a prebuilt model is provided for you to configure.

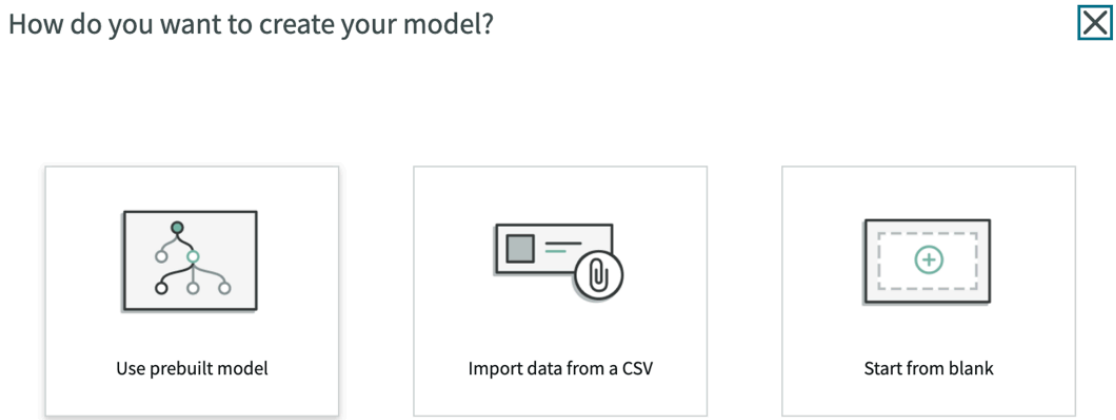
To start creating models, set your scope to the application scope you want for your new model. Then navigate to **NLU Workbench** > **Models**. The Virtual Agent tab opens by default. Select the appropriate tab for the model you want to create.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

New model creation

When you select the **Create new model** button, a modal opens to display your model creation options. Start by selecting one of the



icons:

- Create an NLU model using a pre-built model: Copy a prebuilt model and its contents as a starting point for your new model.
- Create an NLU model from a CSV file: Upload a CSV file containing a list of intents and corresponding utterances.
- Create an NLU model from blank: Build a model from scratch and add intents and utterances as you go.

After creating, add content to your model. The intents, utterances, entities, and vocabulary you add helps improve the model's ability to interpret natural language. See Build and train your model.

Duplicating, exporting, and updating

After creating a model, you have options to use that model across other models and instances. With the NLU Workbench, you can perform the following actions with your models:

- Duplicate an NLU model: Copy a model to create a model with the same content.
- Export an NLU model: Export a model as a CSV file containing the associated utterances and intents. Share the model or use it to create one.
- Add an NLU model to an update set: Add a model and its artifacts to an update set to transfer the model across instances.

Create an NLU model from blank

Create a Natural Language Understanding (NLU) model from scratch. Start with an empty model to have full control over the model's content.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- You can create NLU models for Virtual Agent and AI Search.
- Role required: nlu_admin or admin

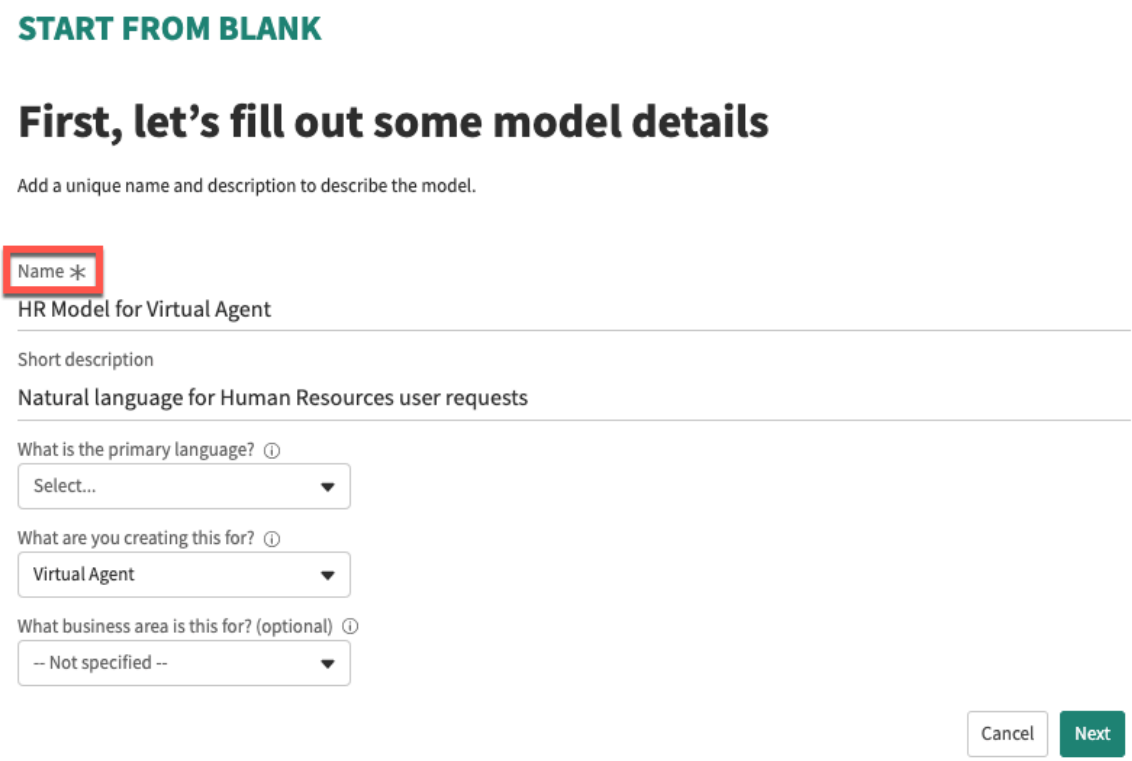
About this task

In this example procedure, you're building an NLU model to help Virtual Agent understand human-expressed intent regarding Human Resources.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Set your scope to the application scope you want for your model.
2. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
3. Select the tab for the type of model you want to create, such as AI Search.
4. Select the **Create new model** button.
5. In the **Add some details** window, select the **Start from blank** button.
6. In the define details page, enter a unique **Name** and **Short description**.

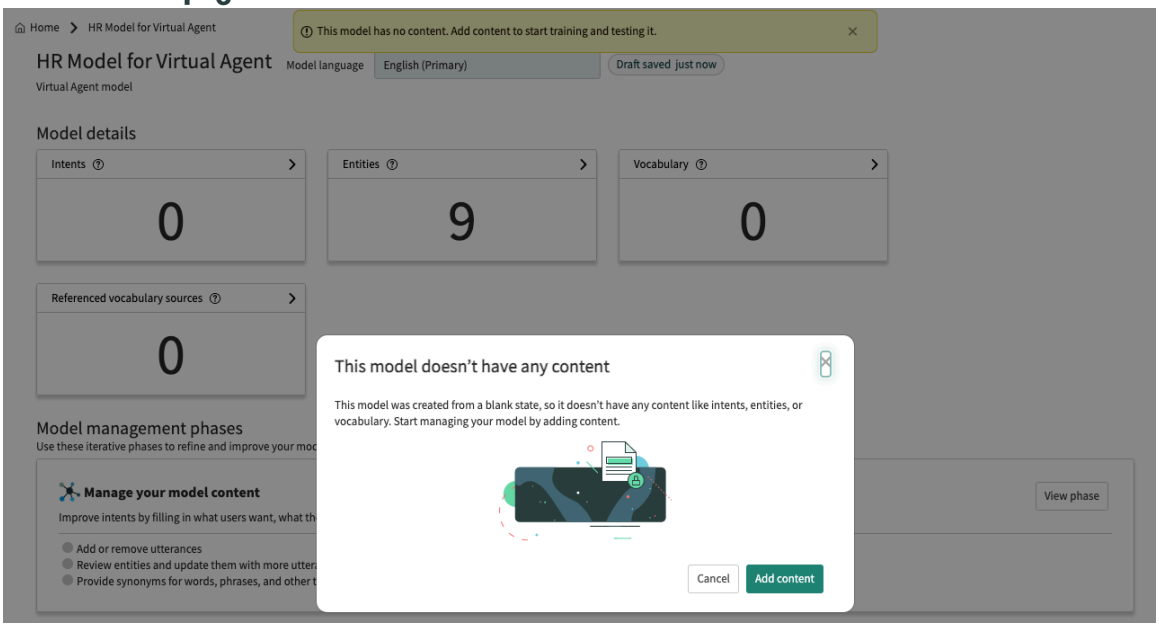


   In this example scenario, you enter `HR Model for Virtual Agent` for the name and `Natural language for Human Resources user requests` for the description.
7. Select the language and purpose from the drop-down lists.
   In this example scenario, select `English` and `Virtual Agent`.
8. **Optional:** Select a business area for your model.
9. Select **Next**.
   Your model starts building. When complete, select **View model** to open the **Model details** page for your new model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

Model details page



Your new model contains no content. Select **Add content** to begin adding intents, entities, and vocabulary. For more information, see Build and train your model.

Add test utterances and intents to build the model's default test set. For more information, see Test set creation and management.

Create an NLU model from a CSV file

Upload a CSV or XLSX (Excel Workbook) file containing utterances and their intents to create a Natural Language Understanding (NLU) model. Use this method to quickly create models from your data or other exported models.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- You can create NLU models for Virtual Agent and AI Search.
- Role required: admin or nlu_admin

About this task

In this example procedure, you're building an NLU model to help Virtual Agent understand user requests regarding calendars.

In a CSV file, you've listed the intents and utterances in two columns.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Sample CSV setup

Intent	Utterance
schedule	when is my next meeting
schedule	what time is my next meeting
schedule	schedule a meeting for 11am
schedule	cancel my 4:30 meeting
holiday	when is the next holiday
holiday	do we get Black Friday off
timeoff	request friday off
timeoff	how many PTO days do i have

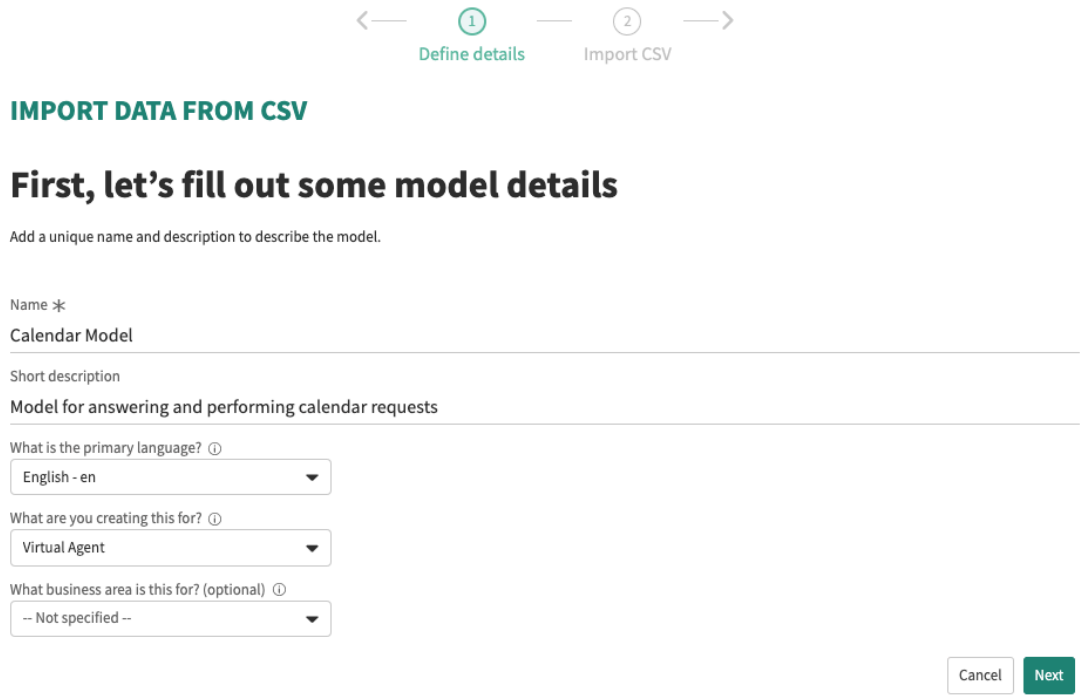
Note the following for creating NLU models by CSV import:

- A model needs at least 1 intent with a minimum of 5 training utterances in each intent. For optimum performance, aim to have 15 training utterances per intent.
- Utterances should not contain a comma.
- Importing with a CSV file does not preserve entities. Make sure to annotate utterances as needed after importing.

Procedure

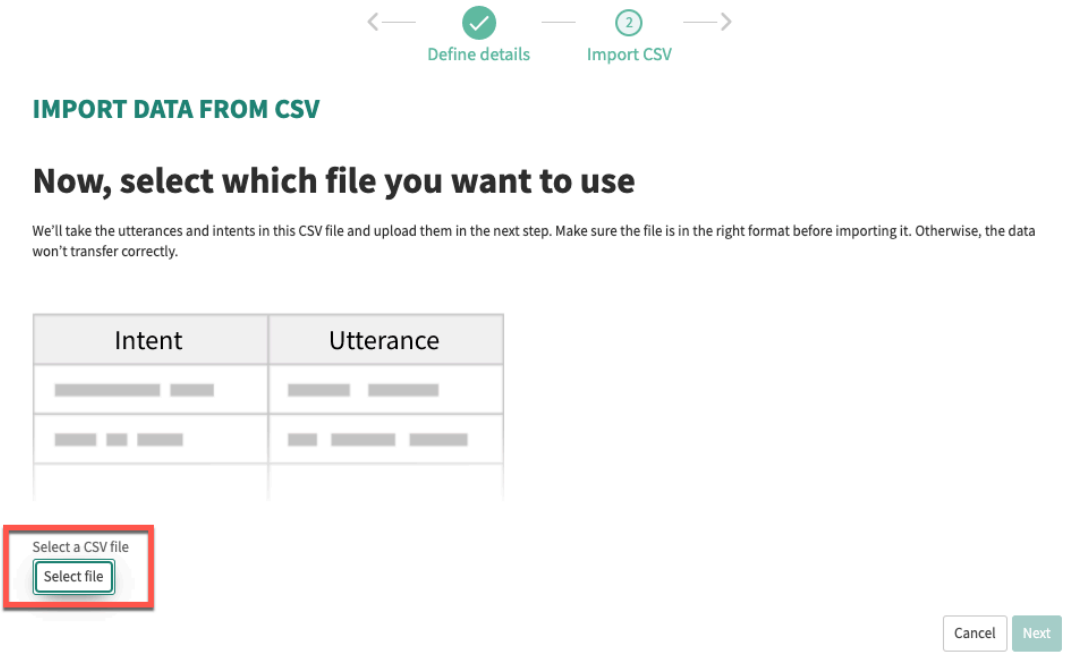
1. Set your scope to the application scope you want for your new model.
2. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
3. Select the tab for the type of model you want to create, such as AI Search.
4. Select the **Create new model** button.
5. In the **How do you want to create your model?** window, select **Import data from a CSV**.
6. In the **Add some details** window, add the **Name** and **Short description** for the model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



In this example scenario, you enter `Calendar Model` for the name and `Model for answering and performing calendar requests` for the short description.

7. Select the language and purpose from the drop-down lists.
   In this example scenario, you select `English` and `Virtual Agent`.
8. Click **Next**.
9. On the **Import CSV screen**, click **Select file**.



10. Choose the CSV or XLSX (Excel Workbook) file from the pop-up.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

11. Select **Next**.
    Your model starts building. After completion, select **View model** to open the model details page.

What to do next

Add intents and training utterances to continue building the model. Add entities and vocabulary to help the model understand inputs from your users. For more information, see Build and train your model.

Add test utterances and intents to build the model's default test set. For more information, see Test set creation and management.

Create an NLU model using a pre-built model

Use the prebuilt Natural Language Understanding (NLU) models in the NLU Workbench to create a new one.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- You can create NLU models for Virtual Agent and AI Search.
- Role required: nlu_admin or admin.

About this task

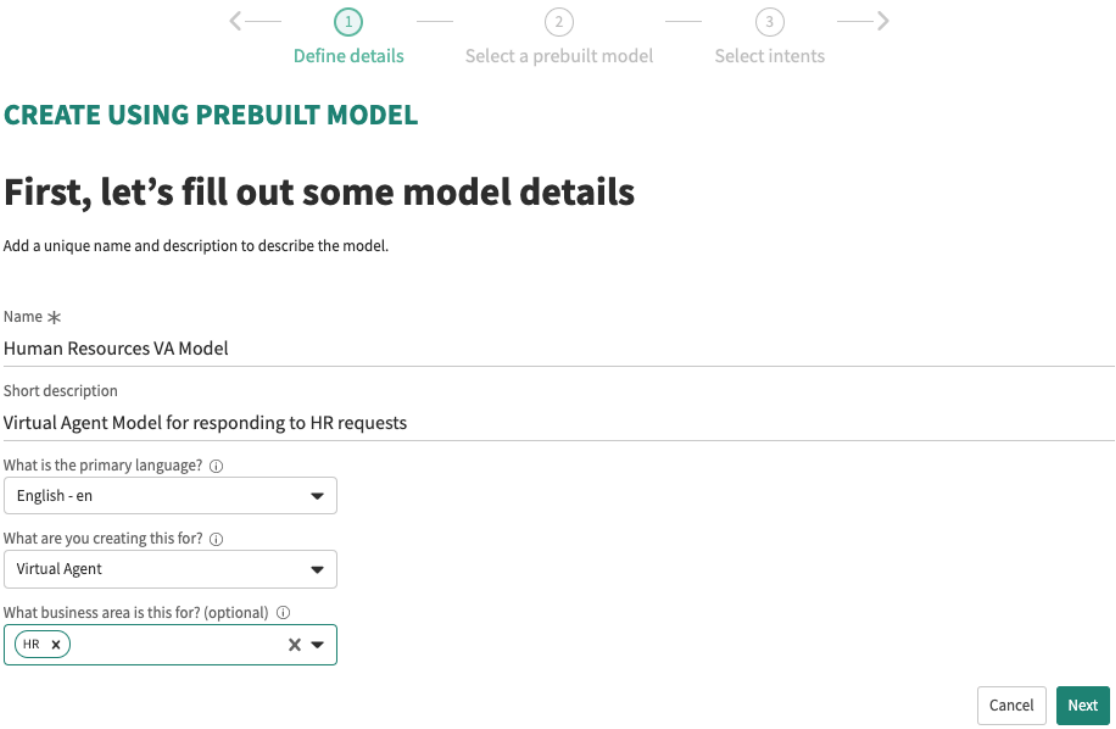
In this example procedure, you're building an NLU model based on the prebuilt HR NLU for VA model. You base your model on the prebuilt model, then edit or add to the model content to meet your business requirements.

Your new model contains intents, entities, and vocabulary from the prebuilt model. The new model also contains an empty default test set, ready to be populated with test utterances.

Procedure

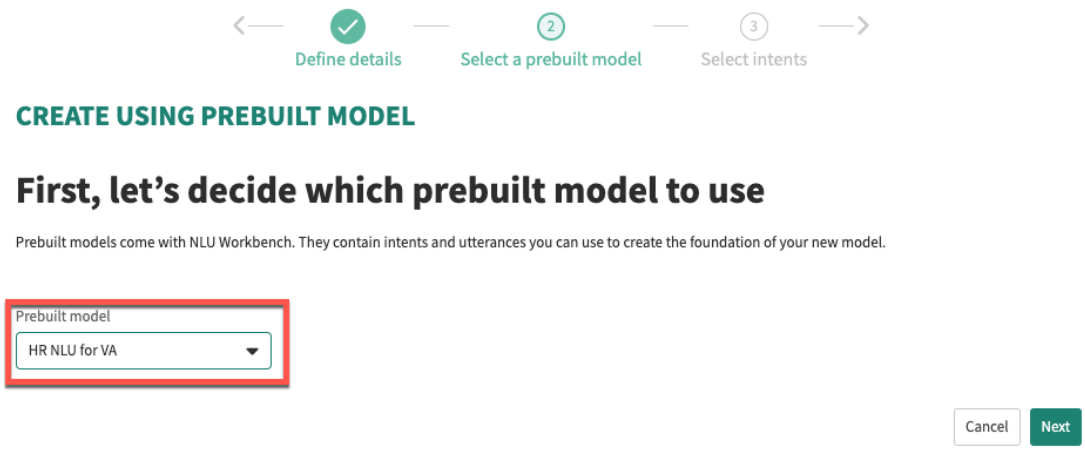
1. Set your scope to the application scope you want for your new model.
2. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
3. Select the tab for the type of model you want to create, such as AI Search.
4. Select the **Create new model** button.
5. In the **How do you want to create your model?** window, select **Use prebuilt model**.
6. On the **Add some details** window, fill in the **Name** and **Short description** for the model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



In this example scenario, you enter `Human Resources VA Model` for the name and `Virtual Agent Model for responding to HR Requests` for the short description.

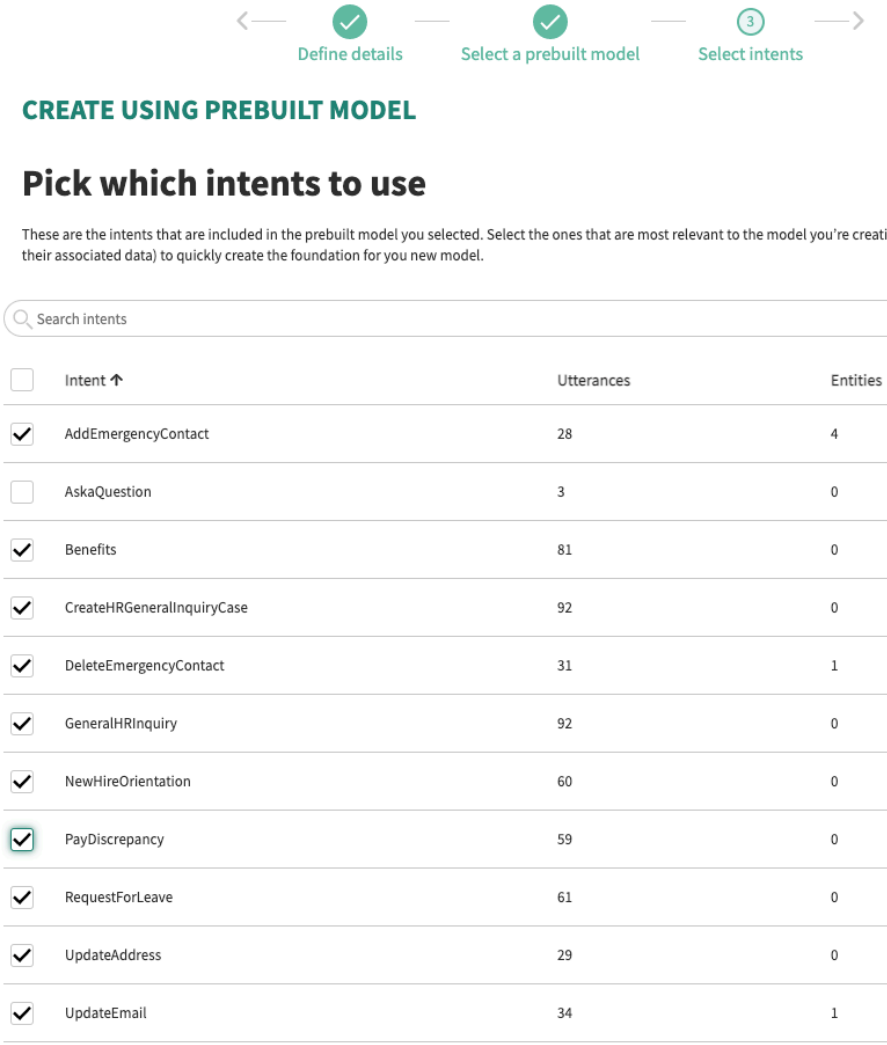
7. Select the language and purpose from the drop-down lists.
   In this example scenario, you select `English` and `Virtual Agent`.
8. **Optional:** Select the business area for the model.
   In this example scenario, you select HR.
9. Click **Next**.
10. Select the prebuilt model from the drop-down list.



In this example scenario, you select `HR NLU for VA`.

11. Select **Next**.
12. On the **Select intents** screen, select the prebuilt intents to add to your model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



In this example scenario, you select 10 intents from the list.

Note:
You can select all the prebuilt intents by checking the **Intent** box at the top of the list.

13. Select **Next**.
    Your model starts building. After completion, select **View model** to open the **Model details** page for your new model.

What to do next

Add more intents and training utterances to continue building the model. Update entities and vocabulary to help the model understand inputs from your users. For more information, see Build and train your model.

Add test utterances and intents to build the model's default test set. For more information, see Test set creation and management.

Duplicate an NLU model

Duplicate an existing Natural Language Understanding (NLU) model to create a new one. Duplicating a model copies the settings and contents of the original model, including its default test set.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Identify an existing NLU model to copy from. The source model should have the language and purpose you want for the new model. If a suitable model does not exist, create and train one.
- Role required: nlu_admin or admin.

About this task

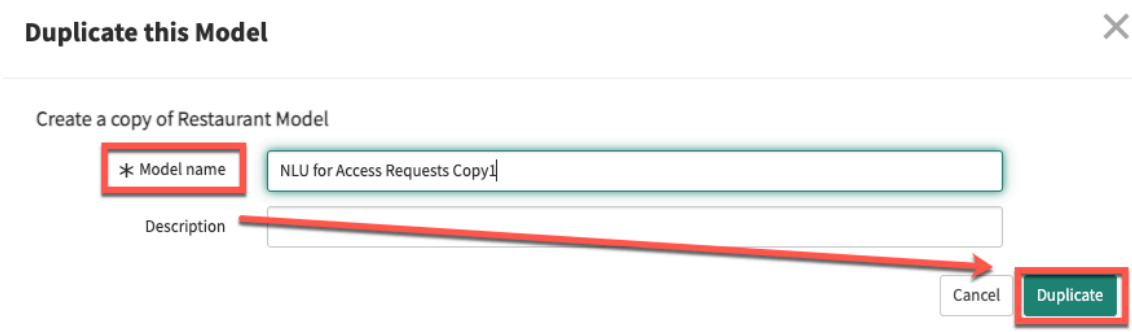
In this example procedure, you've already created a model that's titled NLU for Access Requests. You want to customize it for another department. You duplicate the model to create a new, separate model. To customize the new model, you then add content related to the other department.

Procedure

1. Set your scope to the scope of the model you want to create.
   After creation, these model settings cannot be changed: language, purpose, scope.
2. Navigate to **All** > **NLU Workbench** > **Models**.
3. On the far right column of the model list, select the **More options** menu for the model you want to duplicate.



4. Select **Duplicate this model**.
5. In the **Duplicate this Model** window, enter a name and description for the new model.



   In this example scenario, you enter `NLU for Access Requests Copy1` for the description.

6. Click **Duplicate**.
   The system duplicates the model. The **Model details** page loads for your new model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

You can customize the new model by adding new content to it. Train, test, and publish the new model. For more information, see Build and train your model.

To duplicate an entire model group, or one of the models in a group, see Multilingual model management.

To transfer a model to a different instance, see Add an NLU model to an update set.

Export an NLU model

Export a Natural Language Understanding (NLU) model to create a CSV file of the intents and utterances. You can then use the CSV file to edit, share, and import.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Role required: nlu_editor, nlu_admin, or admin

About this task

Exporting an NLU model creates a CSV file. The file contains a table of the utterances from the model and the matched intents. The data comes from the **Utterances** tab for each intent in the model.

Note:
The file doesn't contain the sources of the utterances. Also the file doesn't transfer entities or vocabulary associated with the model. To export all model data, see Add an NLU model to an update set.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, and find your model in the list.
3. Scroll to the far right column of your model's row, and select the more options menu of the model you want to export.



4. In the more options menu, click **Export model as CSV**.
   The model CSV file downloads to your browser.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

You can 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.

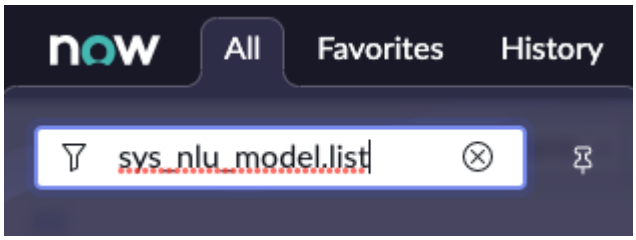
When models (either Global or scoped) are moved using update sets, their training and publishing state remains the same after transfer. So a model that is trained and published before it is added to an update set does not require retraining or republishing on the target instance.

For information about parent-child update sets, see Update set batching. The following video provides a demonstration of this process. A video walkthrough of migrating NLU models by update sets in the Vancouver release.

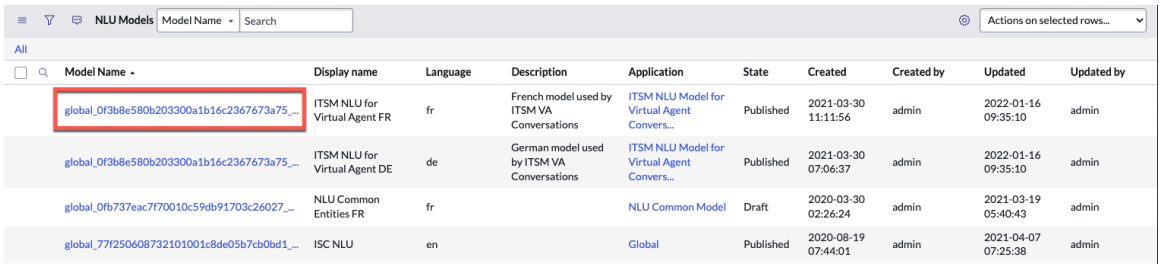
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Select **All** and enter `sys_nlu_model.list` into the navigator.

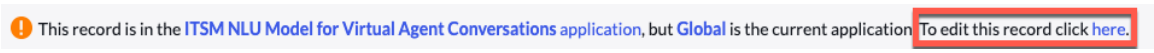


2. Select the **Model Name** of the model.

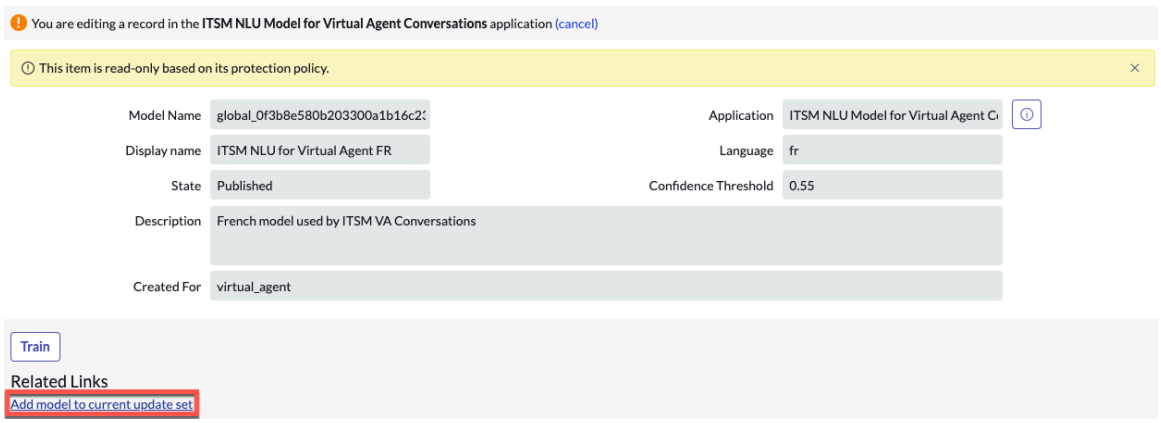


The **Display name** is the name you gave the model when creating it. The **Model name** is given to the model by the system.

3. **Optional:** If Global is your current application, follow the prompt at the top to edit the record.

(Optional) 

4. In the **Related Links** section, click **Add model to current update set.**

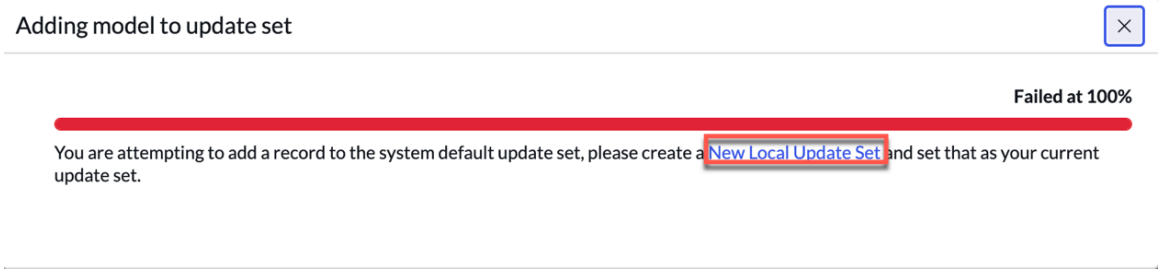


If the model's records are all in one scope, the system adds the model to the current update set. Once completed, the record can be found in the sys_update_set table.

However, the system displays an error if the designated scope is the Default scope. Continue with the following steps.

5. If the system displays the error **You are attempting to add a record to the system default update set**, continue with the remaining steps in this procedure.

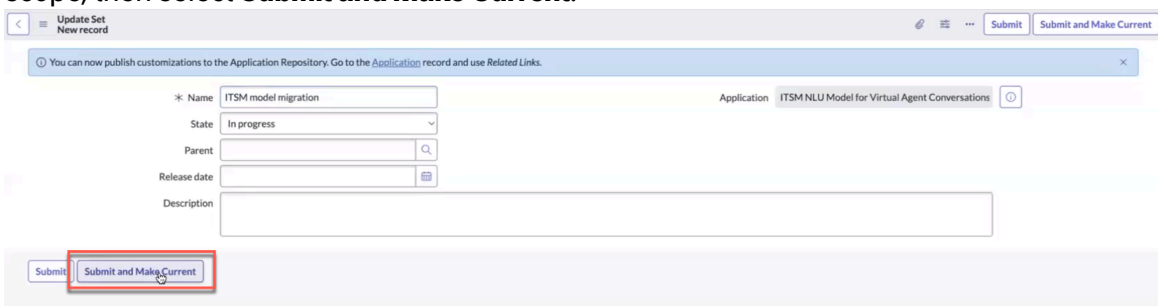
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



In the error banner, select the **New Local Update Set** link to create a new update set which is not Default. In the case of a parent-child update set, the parent is first created in the model's scope. (Any child update sets are created in a later step).

New Local Update Set opens a new record in the sys_update_set table.

6. For the new update set record, provide a name, review other values including Application scope, then select **Submit and Make Current**.



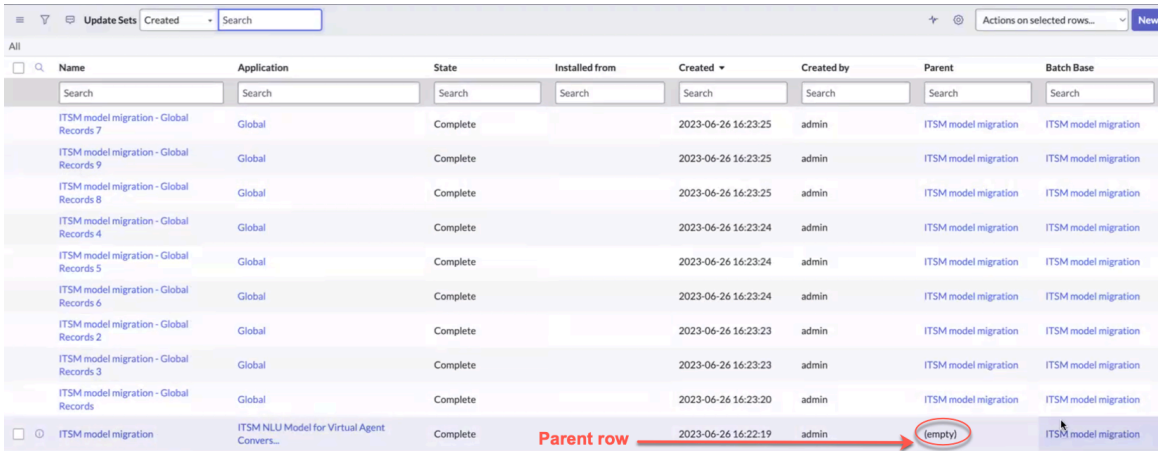
A new update set is created, and the screen re-opens to the model's record in the sys_nlu_model table.

7. In the model's record in sys_nlu_model, select **Add model to current update set** under **Related Links**.
When the model's records are in multiple scopes, this step creates child update sets that have Global scope.

> **Note:**
> If the model's records are all in one scope, the system adds the model to the current update set. Once completed, the record can be found in the sys_update_set table.

8. If the model is in a non-Global scope: In the Update Sets table sys_update_set, locate the parent update set and its child update sets (refresh the page if necessary).
   - The parent update set has the value **(empty)** in the Parent column.
   - Child update sets have the name of the parent in the Parent column, and **Global** in the Application column.
   - Both parent and child update sets contain the same value in the Batch Base column.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



9. Open the parent's record in sys_update_set and set the **State** field to `Complete`.
   A Confirmation dialog box pops up. Select `Yes` to confirm. This sets the parent and all child update sets to Complete.
10. To migrate the update set file, select **Export Update Set Batch to XML** in the parent's record in sys_update_set.
    The parent and child update sets are packaged together and exported.

What to do next

In the target instance, navigate to **Retrieved Update Sets** and select **Import Update Sets from XML**. When unpackaged, the parent and all children will be listed in the Retrieved Update Sets table. Open the parent and select **Preview Update Set Batch**.

If a referenced record in the source instance is not present on the target instance, you may encounter errors while applying the update set. For example, the origin field on sys_nlu_intent could refer to an intent from a different model that doesn't exist on the target instance. You can click **Accept remote update** on the failed records to commit the update set anyway.

Delete an NLU model

Delete a Natural Language Understanding (NLU) model permanently.

Before you begin

Role required: admin or nlu_admin

About this task

Deleting an NLU model removes the model and its contents, including its default test set, from the system. You cannot delete a model if it meets any of the following criteria:

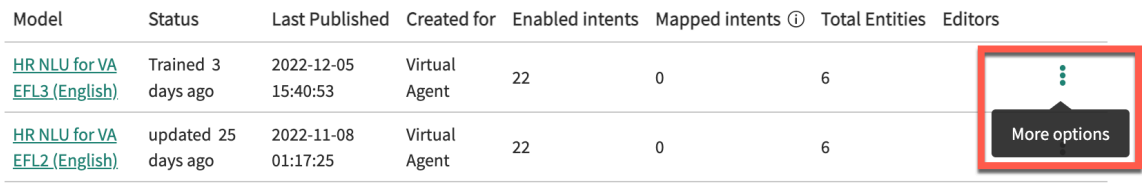
- The model is pre-built
- The model contains at least one intent that is mapped to a Virtual Agent topic

> **Warning:**
> Deleting a model cannot be undone.

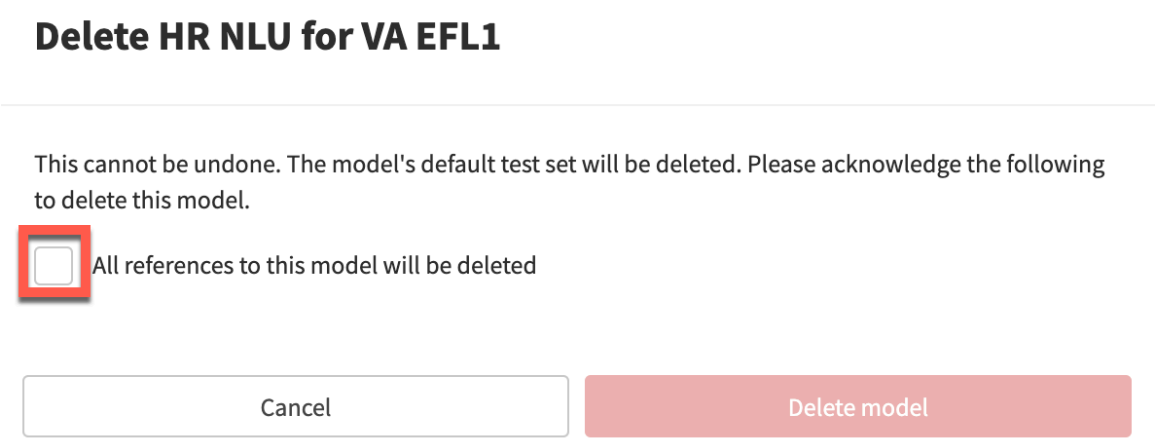
Procedure

1. Set your scope to the scope of the model you want to delete.
2. Navigate to **All** > **NLU Workbench** > **Models**.
3. On the far right column of the model list, select the more the **More options** menu for the model you want to delete.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. Select the option to **Delete this model**.
5. Select the check box to acknowledge that all references to this model will be deleted.



6. Select **Delete model**.
   The system deletes your model and reloads the page.

Build and train your model

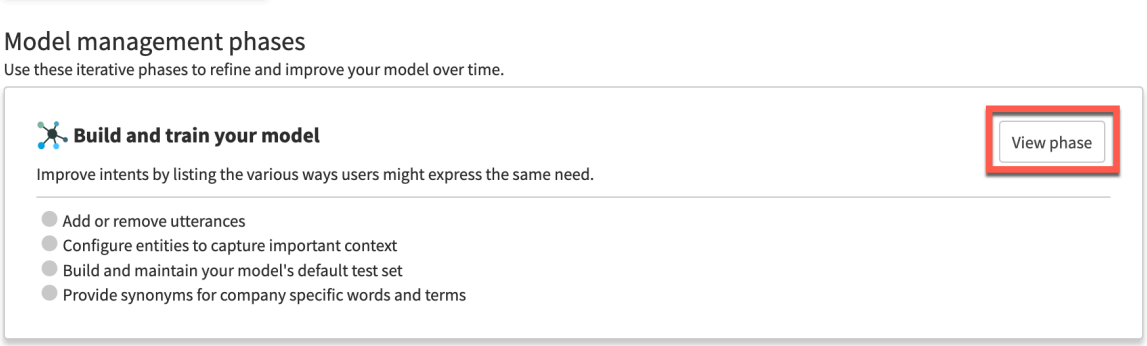
After creating a model, build the model's content by adding intents, entities, vocabulary, and test set utterances. Your NLU model content determines how the model responds to user inputs.

Models are made up of the following content:

- Intents: An action the user wants to do or wants the application to do.
- Entities: Object or context for an action.
- Vocabulary: Add vocabulary to help your model understand the range of words in your users' utterances.
- Test set: To assess model performance, add test utterances and the intents that you expect to be predicted for those utterances.

To access the model content, navigate to **NLU Workbench** > **Models**. The Virtual Agent tab opens by default. Select the tab for your model's application and then select the name of the model to open the **Model details** page. In the **Build and train your model** card, select **View phase**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Intents

When your model receives user input, it uses an intent to perform a system action. For example, a user types in `I have a critical issue with a slow laptop`. The model matches the utterance input to the intent #TroubleshootSlowComputer. If the intent is linked to a Virtual Agent topic, it triggers further action.

Intents contain training utterances, or examples of user inputs that would trigger the system action. Provide realistic utterances that the model might encounter from your users. The quality of training utterances affects the accuracy of your model.

For more information, see NLU intents.

Entities

Your intents use entities to provide additional context for the model when receiving inputs. In the computer example, the laptop is the entity, or object of, the action.

NLU entities fall into two categories: system-defined and user-defined. System entities such as DATE, TIME, and LOCATION are available by default in your instance. You can create your own user-defined entities to provide additional associations and meaning for your business requirements.

All entities are reusable across other NLU models. However, you must add them to a training utterance for each model to use them.

For more information, see NLU entities.

Vocabulary

Your users' input may contain a wide variety of words and phrases. Also, your model may not understand some terms used in specialized domains or business areas.

To improve your model's ability to understand a wide range of user input, you can define synonyms by creating vocabulary items.

For example, your model includes an entity for the term computer. When a user types in `I need a new computer`, the model knows how to respond. However, if a user enters `laptop` or `workstation`, the model might fail to predict the intent. You can add vocabulary to the model to train it to understand synonyms and variations.

You can also use tables and lists as vocabulary sources. Your models can look up the vocabulary sources when predicting intents.

For more information, see NLU vocabulary.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Test set

Your model contains a default test set that you can use to evaluate the model's performance. Initially the test set is empty, ready to be populated with your content. Add test utterances and their expected intents to build the test set.

For more information, see Test set creation and management.

Test panel

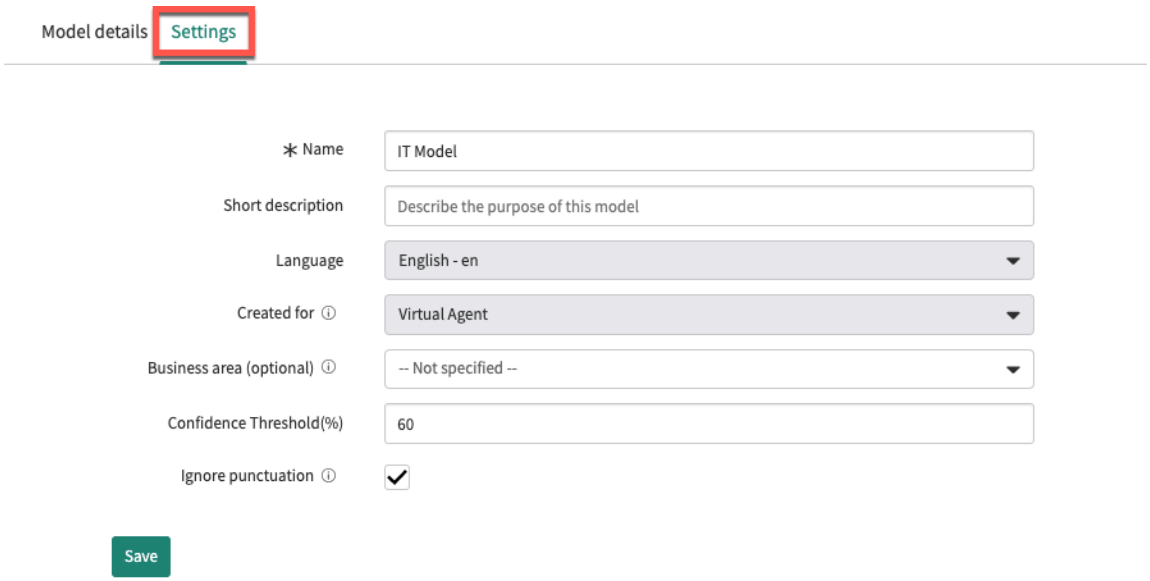
Access the test panel by clicking **Train model** or **Try model** in the **Build and train your model** phase. Training incorporates new content into your model. With **Try model**, you can manually enter individual utterances to test what intents the model predicts for them.

For more information, see Train and try your NLU model.

You can also use the test panel to provide feedback on your model's predictions. Your feedback helps improve intent prediction. See Test panel feedback.

Settings

Use the **Settings** tab to edit the name, short description, and confidence threshold of the model. You can't change the language or purpose of the model.



For more information on the confidence threshold, see Test and publish your model. For more information on Settings, see NLU model settings.

NLU intents

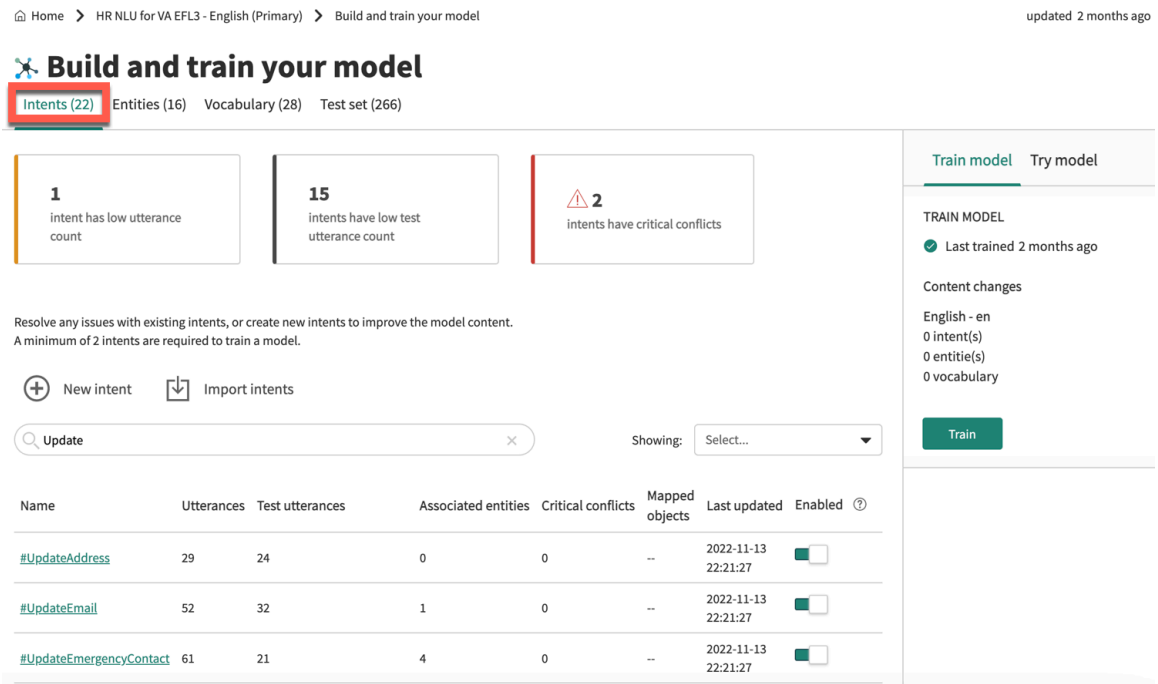
Intents drive your models' responses by matching a system action to user inputs. Models with good intents help Virtual Agent and Search respond to your users accurately.

Think of intents as the core that drives the comprehension of natural language. Intents help the models translate utterances from your users into actions that the system can perform. Intents perform best if your utterances, entities, and vocabulary work together to support your model. Using vocabulary and realistic utterance examples can help the system to be more accurate when predicting intents.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When creating your Natural Language Understanding (NLU) model, you add intents to match user inputs with system actions. The more intents that a model has, the more actions it can take when it receives utterances from your users. Model intents can perform various actions, from creating hardware service requests to adding users to groups.

To access your model's intents, navigate to a model's overview page. In the **Build and train your model** phase, select **View phase**. The **Intents** tab shows by default.



A model can contain up to 750 intents. However, models with over 300 intents or 4500 utterances, whichever comes first, take longer to train, test, and publish.

Note:
As you create intents, keep in mind that they can sometimes impact each other. For example, you build and test an intent that works on one or two utterances. But when you test it in a larger intent environment, it might behave differently. To reduce the likelihood of such an event, you may want to create at least five intents in a model before you start proper testing.

To start adding content to your model, see Create an NLU intent.

The **Enabled** column shows whether or not the intent is active in predictions. An NLU admin can deactivate an individual intent but keep it in the model.

However, if an intent is mapped to a published Virtual Agent topic:

- You cannot deactivate the intent.
- You cannot delete the intent.

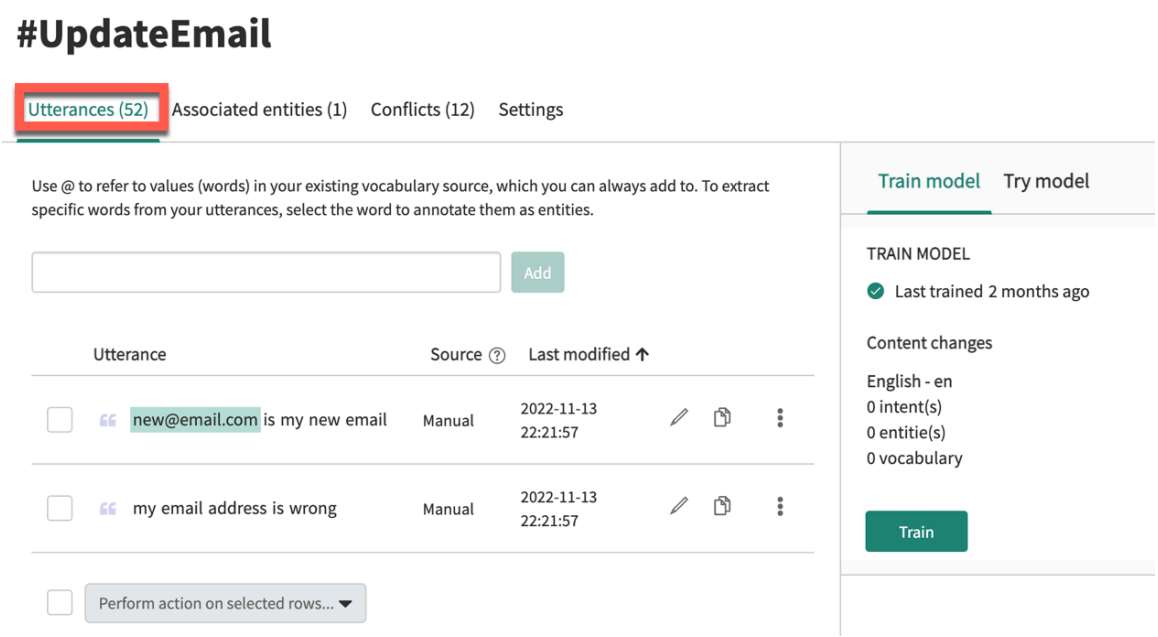
After changing the **Enabled** status of an intent, retrain the model.

Utterances

Intents contain training utterances, which are examples of inputs that the model may see from your users. Each intent in a model has its own utterances. When trained, the model learns to recognize similar utterances from your users and then respond with the matching intent.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Select an intent to access the intent details page. The **Utterances** tab lists all the utterances currently in the intent.



After adding utterances, you can edit, copy, move, or delete the utterances using the icons in the right column. You can move or delete multiple utterances by selecting the box on the left first and using the **Perform action on selected rows** button.

Here are some things to consider when adding utterances to intents:

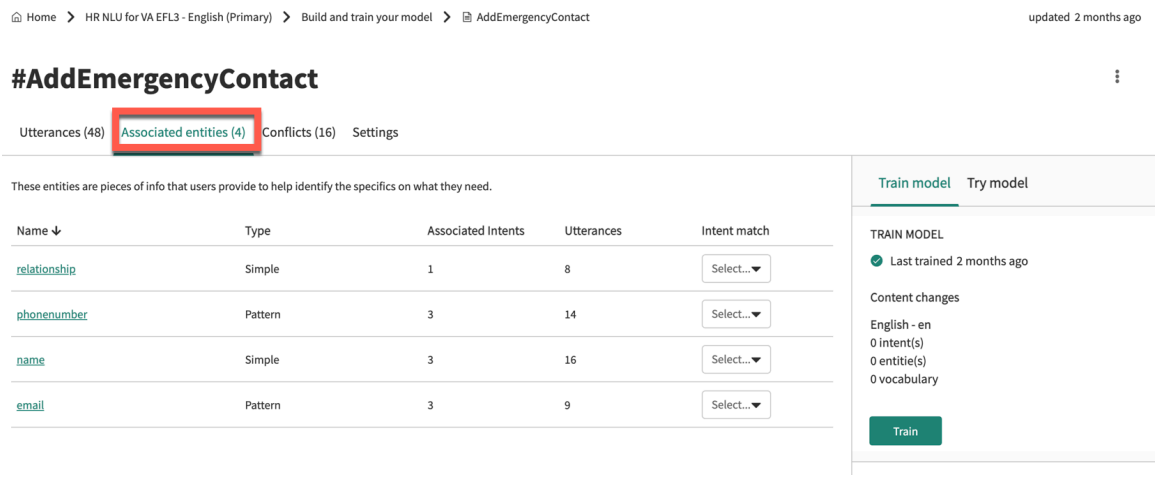
- A model must have at least 1 intent, with a minimum of 5 utterances in each intent.
- An intent needs at least 5 utterances to begin training.
- The system currently supports utterances up to 25 words or 200 characters. Utterances that exceed that limit fail to return an intent prediction.
- The system currently supports up to 20,000 utterances in a single model.
- Models with more than 4,500 utterances take longer to train, test, and publish.

Provide vocabulary for any words or phrases that are relevant to your organization or domain when adding utterances to your model. The vocabulary helps with intent prediction for words and phrases that your users are likely to use. You can use the @ symbol when adding an utterance to call on a vocabulary source. For more context and examples, see NLU vocabulary.

Associated entities

Your model uses entities to provide additional context and meaning when predicting user input. You add entities to the training utterances of your intent to provide the system with more information to perform the intended action.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

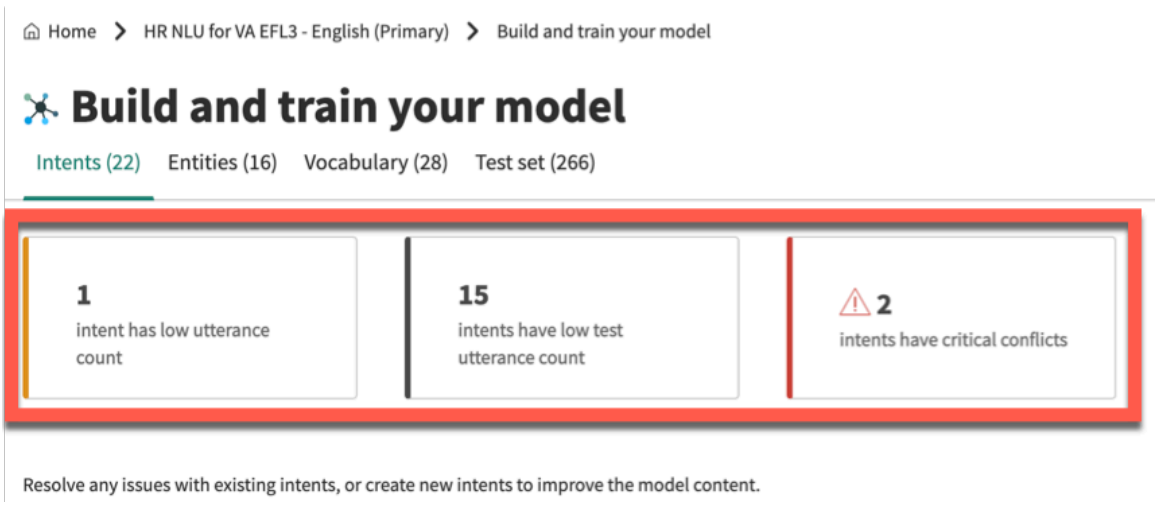


For more information, see NLU entities.

Intent issues

Building large models increases the chance that intents overlap, conflict, or fail to contain enough training utterances. For example, the utterance examples in one intent may end up identical to the examples in another intent. If your intents conflict, the model may not know which intent to predict when receiving user input.

If your model has issues or conflicts, the **Intents** page displays cards showing the number of intents affected.



Select a card to see a filtered list of intents with that issue. Resolving issues ensures that your intents meet the requirements and work as intended.

For more information, see Resolve intent issues.

Create an NLU intent

Create an intent for your Natural Language Understanding (NLU) model. Intents provide your model with a system action to perform when it receives user input.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Workbench - Advanced Features plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Ensure that you are in the same application scope as your model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- You can create intents for Virtual Agent and AI Search models in NLU Workbench.
- Role required: admin or nlu_admin

About this task

This procedure shows you how to create an intent. To reuse intents from other models, see Reusing intents from prebuilt NLU models.

Here's an example of how intents can interact with the vocabulary in their training utterances.

- Intent: #AddMembersToDistributionList
- Utterance A: "Please add Carlos Santana to the uxinfodev list"
- Utterance B: "I'm mistakenly removed from the arlo-drury-directreports group"
- Result: The system doesn't recognize uxinfodev or arlo-drury-directreports and can't use these words to predict the intent.
- Solution: Add uxinfodev and arlo-drury-directreports as vocabulary items and add synonyms to them. The synonyms you provide help add more context to the utterance and the intent in which they reside. Your intent prediction confidence can be even higher if you also mark them as entities.

Note:
Training utterances and utterances from your users have a limit of 25 words or 200 characters. Utterances that exceed that limit fail to return an intent prediction.

Do not include unrealistic terms such as "OrderLaptop" or "sfsdfasdfas" in training utterances. Utterances should be correct and natural examples in the model's language.

In the following example procedure, you're creating an intent and adding utterances that users might say when requesting information on payment. You've already created an NLU model that you've titled HR Model for Virtual Agent and you're creating an intent in that model.

Procedure

1. Set your application scope to your model's scope.
2. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
3. Select the tab corresponding to your model's application, then select the name of the model you want to add an intent to.
   In this model example, you select your HR Model for Virtual Agent model.
4. In the **Build and train your model** card, select **View phase**.
5. Select **Intents**.
6. Select **New Intent**.
7. In the **Create an intent** window, add a name and description for the intent.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



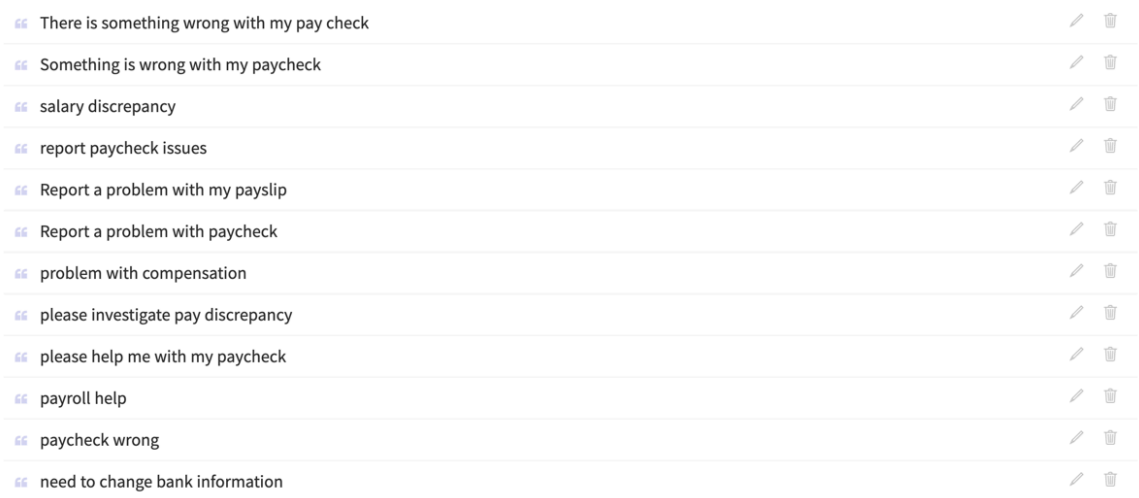
In this example scenario, you enter `PayDiscrepancy` for the name. You can also add a description.

Note:
When you create an intent, the system adds a hashtag to the intent name.

8. Select **Add intent**.
   The #PayDiscrepancy intent screen appears, including sections for its utterances, associated entities, and settings. The intent draft status is also shown in the upper right corner of the Intent screen.
9. In the **Utterances** tab, enter training utterance examples that are relevant to the intent.

   Note:
   The utterance examples that you provide must be unique and contain fewer than 25 words or 200 characters. Aim to add at least 15 utterances with as much variety between them as possible. You must add at least 5 utterances to begin training the model.

   In this scenario, you add the following utterances into the field and select **Add**



As you build and retrain your model iteratively, you can check how your updated intent affects the model's predictions. See Train and try your NLU model.

What to do next

Train your model to save your updates. For issues with intents, see Resolve intent issues.

To improve your utterances, add entities to provide context. See NLU entities.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The available Intent Discovery feature can help identify intents that would be possible to add, based on your historical data.

Reusing intents from prebuilt NLU models

Reuse Natural Language Understanding (NLU) intents by importing them from a prebuilt NLU model to other models. Reusing intents saves time when building your models.

Prebuilt Virtual Agent NLU models provide language understanding needed for chatbot conversation flows in ITSM, CSM, and HR topics. Each NLU intent in these models maps to a single Virtual Agent conversation topic created in Virtual Agent Designer accordingly.

The prebuilt models are read-only, but contain validated intents that you can reuse in your NLU models. For example, the prebuilt HR model for Virtual Agent contains several intents regarding HR requests. Adding the intents to one of your models makes that model ready to respond to user inputs regarding HR.

Prebuilt Virtual Agent model content

Each of the prebuilt Virtual Agent models uses language that's pertinent to their respective business unit (BU). The NLU that processes this language, built from a word corpus of 3 million words, is context-aware of general linguistic patterns and both ServiceNow and user-defined vocabularies.

General Patterns	ServiceNow Vocabulary	User Defined Vocabulary

• Date
• Time
• Duration
• Location
• Phone number
• Phone extension
• Email
• Email domain
• Website address
• Computer drive
• Executable file
• IP address
• Bytes
• Proper Names
	• Task number
	• Incident number
	• Change number
	• Case number
	• KB article number
	• Problem number
		• Standard phrases
		• List – ServiceNow Table: Field Examples
		  • Change request: Type
		  • Task: Impact
		  • Task: Urgency
		• List – user defined
		• Regular expressions

The prebuilt models map to common conversation topics used by Virtual Agent chatbots across each of the three BUs. Those topics correspond with intents that perform BU-specific actions.



For more NLU and Virtual Agent context, refer to the following product documentation:

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- For ITSM context, see ITSM Virtual Agent.
- For HR context, see Virtual Agent for HR Service Delivery.
- For CSM context, see Customer Service Virtual Agent conversations.

Virtual Agent and NLU plugins and roles

To access the prebuilt Virtual Agent models in the NLU Workbench, you must install and activate these plugins using the admin role.

Plugins	Descriptions
• CSM Virtual Agent conversations [com.sn_csm.virtualagent] • Human Resources Scoped App: Virtual Agent Conversations [com.sn_hr_virtual_agent] • ITSM Virtual Agent conversations [com.snc.itsm.virtualagent]	Enables Pre-built Virtual Agent topics, topic blocks, and ServiceNow NLU models for the Customer Service Management, HR Service Delivery, and IT Service Management applications. • Requires the Glide Virtual Agent (com.glide.cs.chatbot) plugin to be activated. • Automatically activates the NLU Workbench (com.snc.nlu_studio) plugin.

Creating NLU models that reuse predefined Virtual Agent intents

The prebuilt Virtual Agent models are set to read-only and can't be edited. However, the intents in these models can be imported into your new model, alongside any new intents you've created in your model. You can also copy the prebuilt model and use the intents as a foundation for a new model.

To reuse the intents from a prebuilt Virtual Agent model, choose the prebuilt model when importing intents. See Import an NLU intent.

NLU models and their intents and entities are associated to an application scope. The scope can't be changed later, so verify your target application scope before you begin.

Import an NLU intent

As you create intents for your Natural Language Understanding (NLU) model, you can also import and reuse intents from other models in the same application scope. Reusing intents saves time when building new models.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- You can import intents for Virtual Agent and AI Search models.
- Verify that the application scope of your source model and target model are the same.
- Role required: nlu_editor, nlu_admin, or admin

About this task

When you import an intent, the following associated items are also imported:

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- training utterances, with their entity annotations
- entities

The following items are not included when importing an intent:

- test utterances in the default test set
- regular and pattern vocabulary

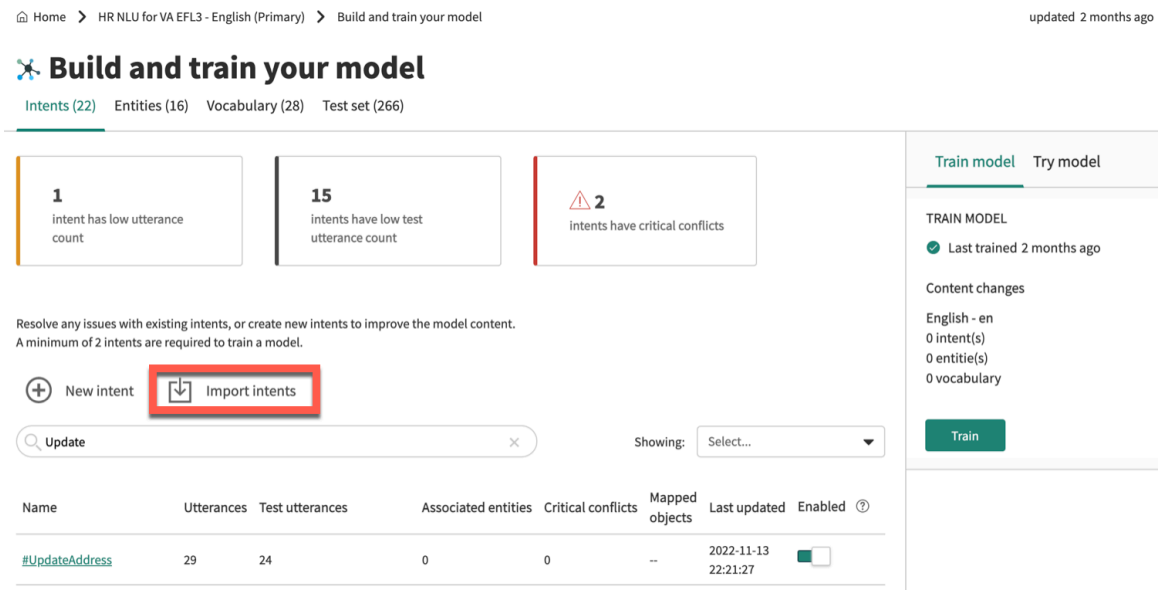
Note:
You cannot import intents in the following situations:

- When the model already has an intent with the same name as the intent you are trying to import.
- When the model already has an entity with the same name but different attributes as what you are trying to import.

In this example scenario, you are building the content of your NLU model and you want to enhance the model by reusing intents from another model.

Procedure

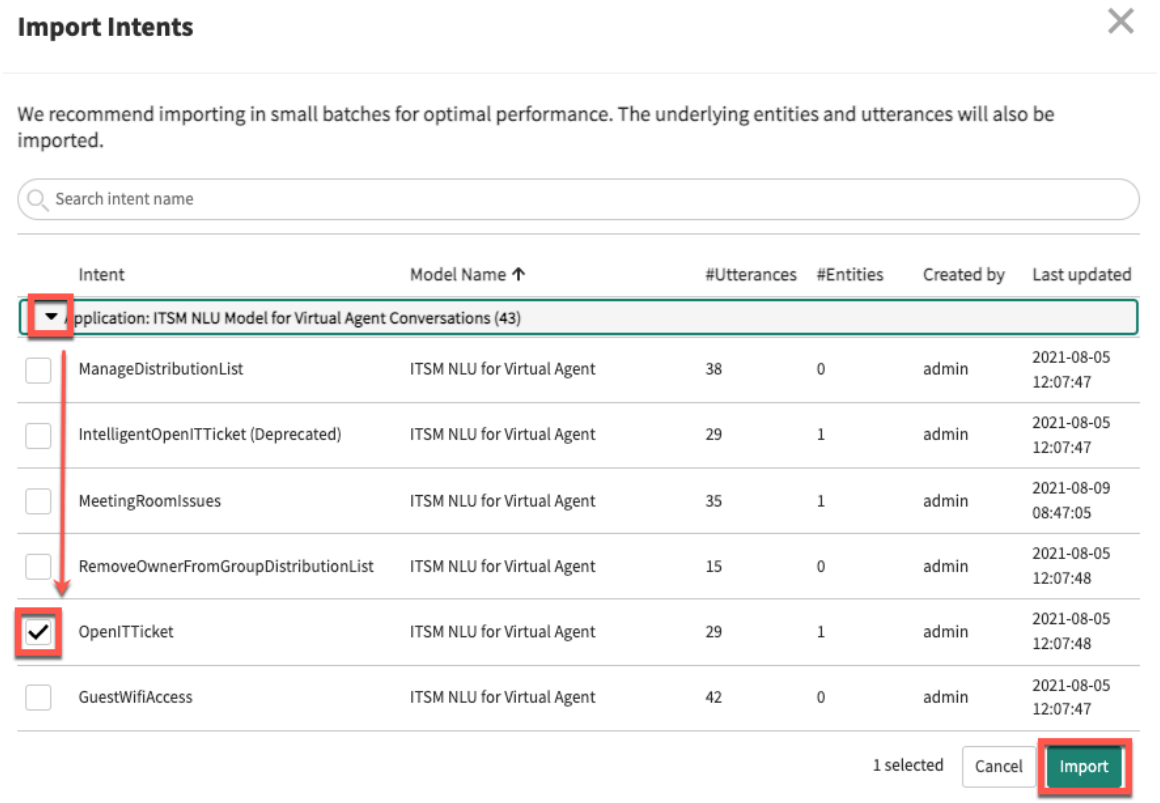
1. Ensure that you are in the target application scope, and navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, then select the name of the model you want to add an intent to.
3. In the model overview page, find the **Build and train your model** phase and select **View phase**.
4. In the **Intents** tab, select **Import intents**.



   The **Import intents** screen appears, showing a list of NLU models from which you can select intents to import.

5. Locate the intents you want to import, and select their check boxes.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



In this scenario, you select the **OpenITTicket** intent in the ITSM NLU Model for Virtual Agent Conversations model.

6. Select **Import**.

   Result: The **#OpenITTicket** intent appears in the intents list of your target model. By importing this particular intent, you have added 29 additional training utterances to your model.

What to do next

Review the annotations of the newly imported training utterances. Add vocabulary annotation if needed.

Train your model to save your updates. For issues with intents, see Resolve intent issues.

If there are any secondary language models associated to your target model, consider importing and translating the new content. See Import primary model content to a secondary model.

For information about prebuilt Virtual Agent models that you can import from, see Reusing intents from prebuilt NLU models.

Resolve intent issues

Use the issue cards to identify intents that have conflicts, need reviewing, or need more utterances. Resolving intent issues ensures the intents in your Natural Language Understanding (NLU) models work properly.

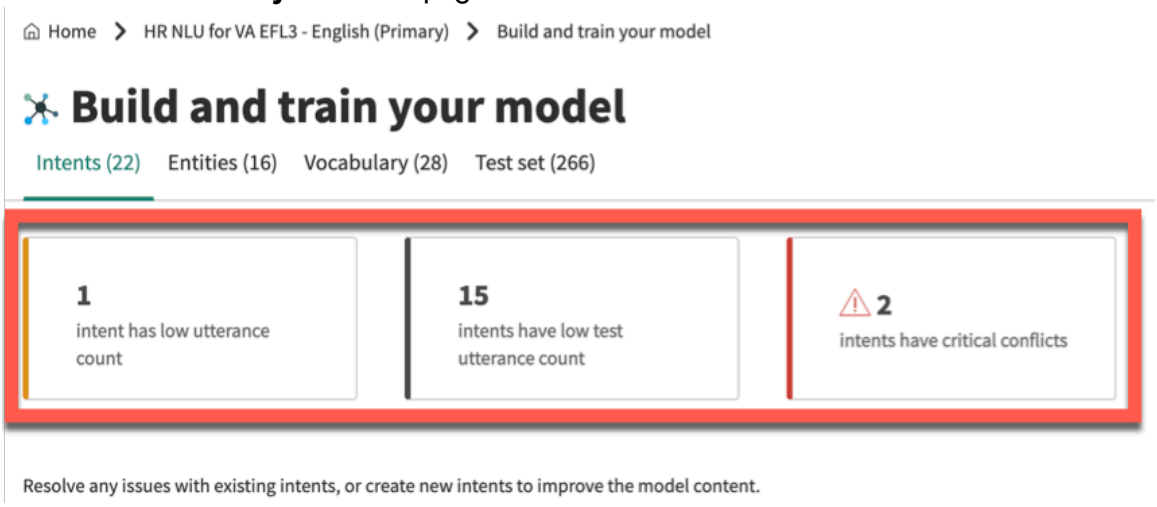
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Workbench - Advanced Features plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- You can resolve conflicts for Virtual Agent and AI Search models.
- Role required: nlu_admin or admin. When assigned to a model, the nlu_editor can resolve intent issues for that model.

About this task

When a trained model contains intents with issues, the issue cards are displayed on the intents tab of **Build and train your model** page. Select each card to filter the list of intents for that issue.



The cards display different issues:

- **intents have low utterance count**: Intent doesn't contain the required minimum of 5 training utterances. This card also displays when the intent contains less than the recommended count of 15 utterances.
- **intents have low test utterance count**: The default test set does not have enough test utterances for the intent. This card displays when the count of test utterances for an intent is below the recommended 15.
- **intents have critical conflicts**: Intent contains training utterances that overlap with utterances in another intent.

  Note:
  Utterances which are marked as `Not relevant`may also conflict with utterances assigned to intents. For more information, see Irrelevance detection in NLU.

- **intents need review**: Intent was translated but must be reviewed by an nlu_admin or nlu_editor.

The issue cards are hidden under the following circumstances:

- The model is a pre-built model
- The model is in a draft state and not trained
- The model contains no conflicts
- A successful conflict report doesn't exist for the latest trained version
- The NLU Workbench - Advanced Features plugin isn't installed

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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 Enable or disable a secondary model intent.

In this example procedure, you are resolving critical conflicts in one of your intents. You built a Virtual Agent model for handling HR requests, but two of your intents, GeneralHRInquiry and CreateHRGeneralInquiryCase included similar utterances.

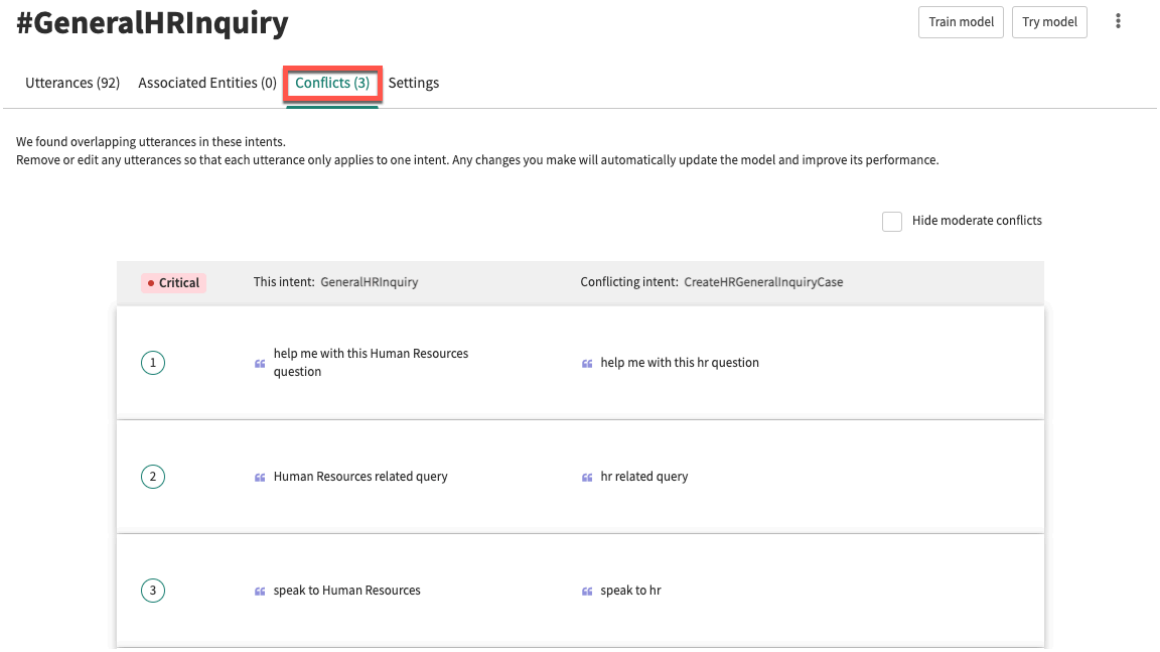
Procedure

1. Ensure that you are in the same application scope as your model, and navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab corresponding to your model's application, then select the name of your model.
3. In the **Build and train your model** card, click **View phase**.
4. On the intents tab, click the **intents have critical conflicts** card.
   The intents list filters and shows the intents with critical conflicts.
5. Click the intent name you want to resolve.
   In this example, you click the GeneralHRInquiry intent.
6. Click the **Conflicts** tab.

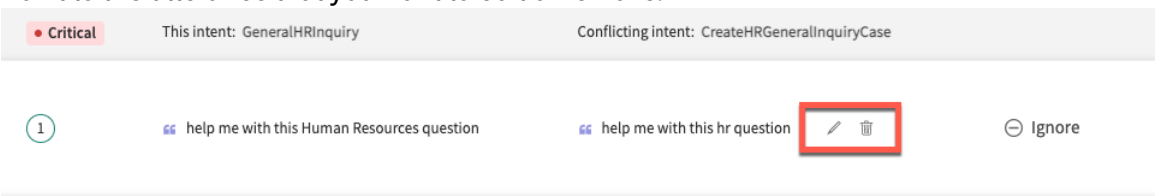
   The conflicts tab appears and lists the conflicting utterances. The utterances for the current intent show on the left. The conflicting utterances for the other intent appear on the right.

   Note:
   You can hide moderate conflicts to focus on the critical conflicts. You can also ignore conflicts. However, be sure to fix all conflicts for best performance.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



7. Point to the utterance that you want to edit or remove.



In this example procedure, click the trash icon to remove the utterance from the CreateHRGeneralInquiryCase intent.

The utterance is removed from the intent.

> **Note:**
> Any changes you make automatically save.

What to do next

Utterances designated as `Irrelevant` or `Not relevant`may also conflict with training utterances. Irrelevant utterances are displayed as though they are in their own intent named **NO_INTENT**. However, these utterances cannot be edited or deleted in Conflicts. Copy the utterance to the Irrelevance Detection module to edit or delete. For more information, see Irrelevance detection in NLU.

Continue to resolve all conflicts. When finished, train your model to see a list of updated conflicts, if any. Resolve all intent issues prior to testing and publishing your model.

NLU entities

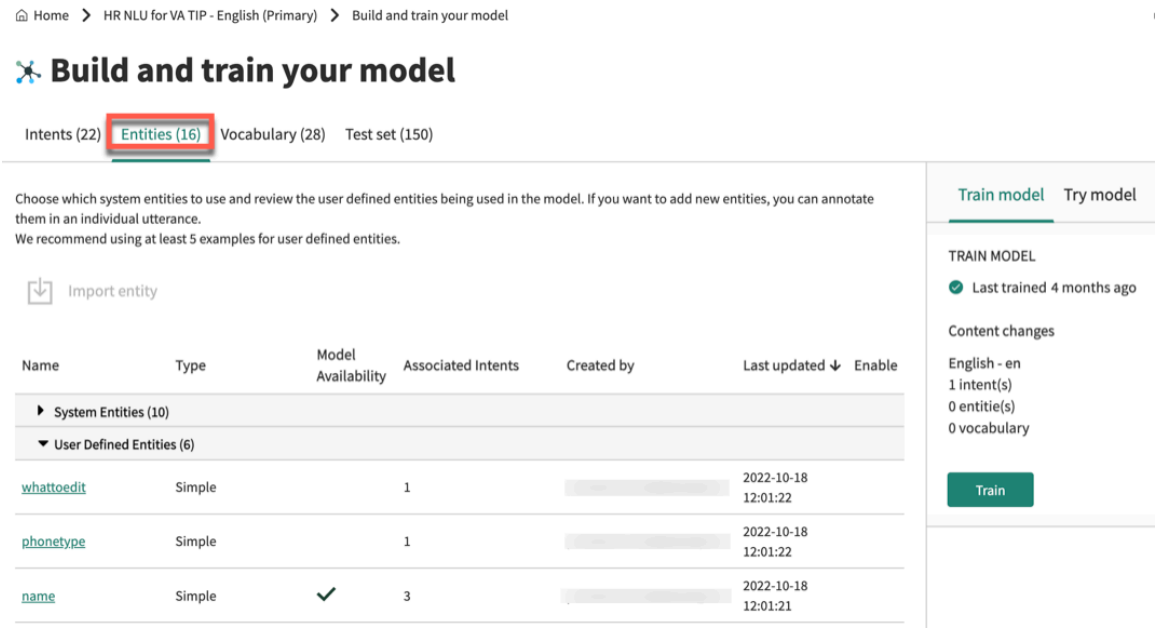
Entities provide your model with additional context when receiving user input. Add entities to your utterances and intents to improve the predictions of your Natural Language Understanding (NLU) model.

Think of entities as the object of the action that the user wants to perform. The model interprets the utterance by matching it with an intent, but also uses entities to gather more details on the request.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, if a user types in `I have a critical issue with a slow laptop`, the system predicts a match to the intent #TroubleshootSlowComputer. From this utterance, the model also identifies the following entities:

- HARDWARE (entity) - laptop (value)
- urgency (entity) - High (value)



NLU entities fall into two categories: system and user-defined. System entities such as DATE, TIME, and LOCATION are available by default in your instance. You can create your own user-defined entities to provide context relevant to your business.

There are five types of user-defined entities. Select one of the following links to learn more.

- Create a simple entity
- Create a mapped entity
- Create a pattern entity
- Create a system-derived entity
- Create an open-ended entity

When you create entities, annotate them on utterances to provide examples to help your model learn. By annotating entities, you provide your model with linguistic associations and meaningful context for the system vocabulary. Annotations strengthen the relevance of entities and help your model perform the correct action in response to your users' inputs.

You add entities to your utterances when you are creating the intents. The entities then become associated to that intent, giving you the **Associated Intents** number.

Regular expressions

Regular expressions (regex) help your model establish patterns that improve that model's ability to locate, match, and manage text. Use regular expressions with pattern entities to help your model understand formats such as email addresses, phone numbers, and incident numbers.

To learn more, see Using regular expressions in entities.

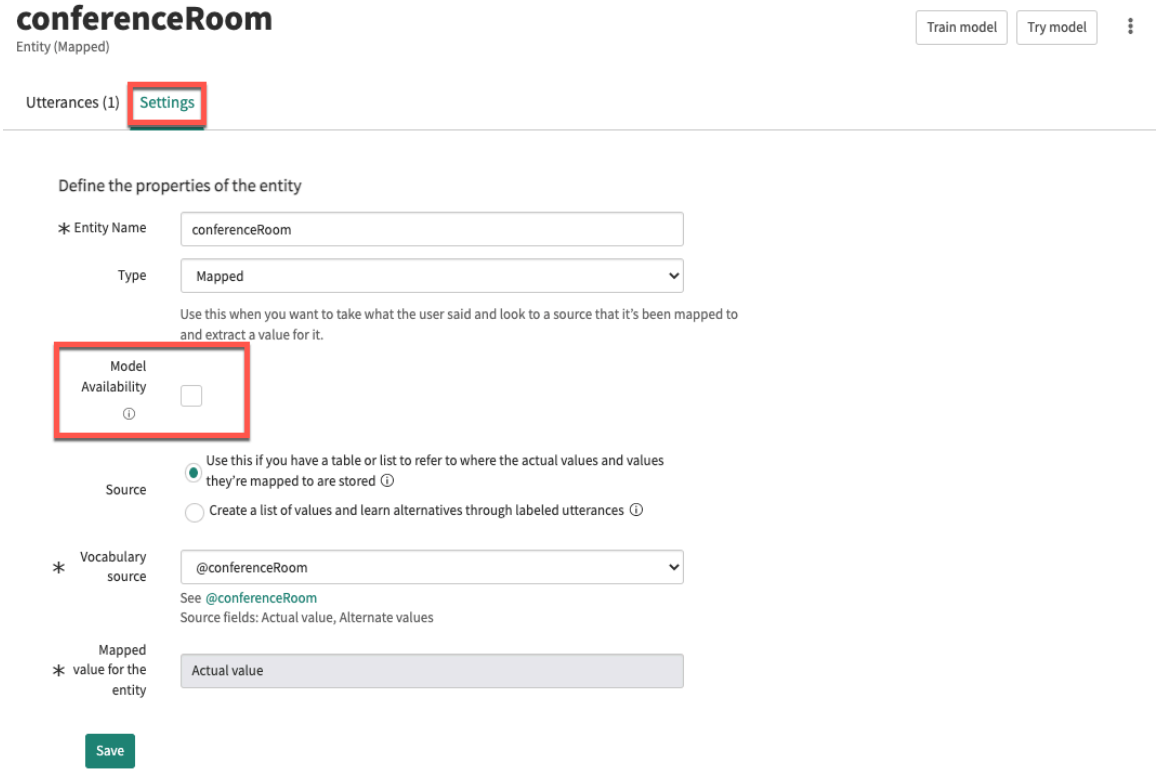
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Model availability

When you create an entity, you can choose to make the entity available for reuse by other intents in the model. If you didn't select the **Model availability** box when creating the entity, you can edit the entity afterwards.

On the Model details page, select **Entities**. Select the name of the entity to bring up the entity details page. Then, select the **Settings** tab.

Entity settings page



Create a simple entity

Create one or more simple entities from words in your utterance examples. An entity is an object of, or context for, an action.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model for Virtual Agent or AI Search.
- Create or use an existing intent.
- Role required: nlu_editor, nlu_admin, or admin. The nlu_editor must be assigned to the model.

About this task

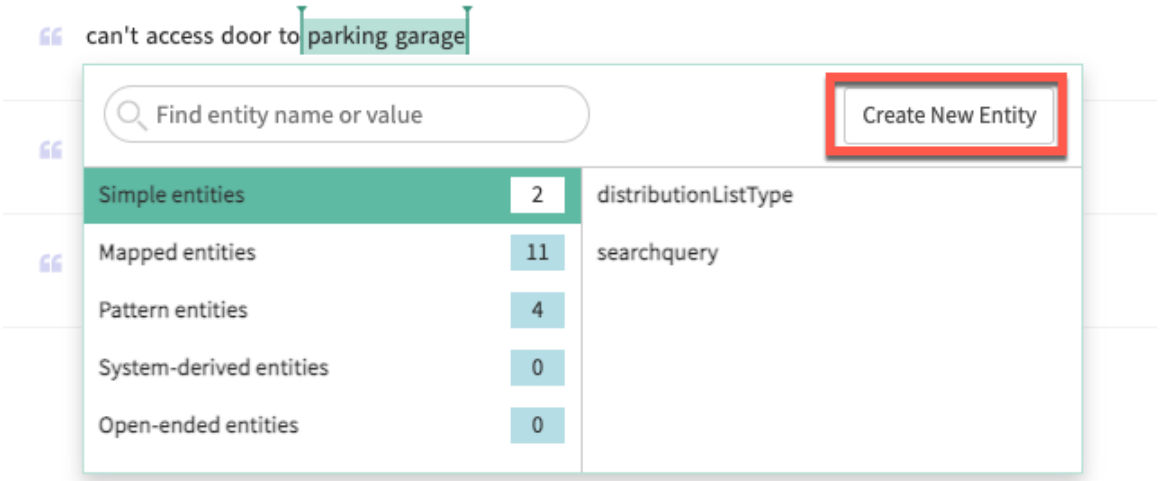
Simple entities are words or phrases whose value can be extracted by your model. Simple entities are identified based on the context in which the entity is used in an utterance.

For the following example procedure, you've already created an intent that's titled SubmitAccessRequest and you're creating a simple entity for the type of access the user is requesting.

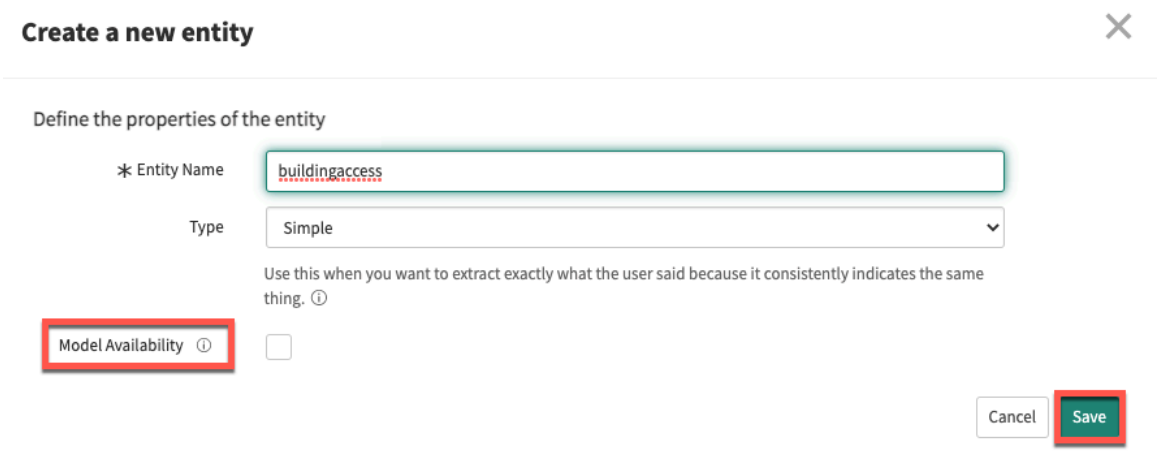
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab corresponding to your model's application, then the name of the model you want to add an entity to.
   For this example, you choose the NLU for Access Requests model.
3. In the **Build and train your model** card, click **View phase**.
4. In the intents tab, click the name of the intent.
   For this example, you click SubmitAccessRequest.
5. Click a word in one of the utterances to bring up the entity window.
   For this example, click the word parking garage.



6. Click **Create New Entity**.
7. In the Create a new entity screen, enter a name and select the **Simple** entity type.
   For this example, you enter `buildingaccess` for the **Entity Name**.



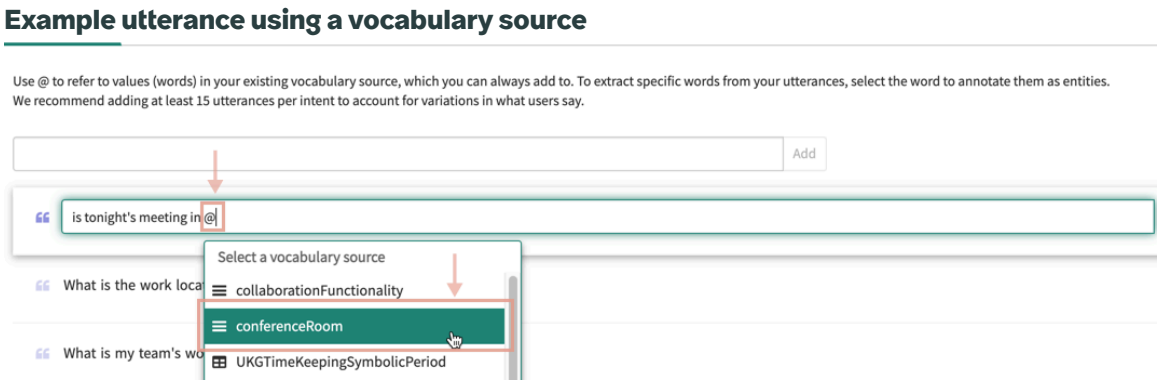
Note:
Select **Model availability** to make the entity available to every intent in the model. If you check the box, the entity shows in the **Associated Entities** tab.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. Click **Save**.
   The 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.



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.

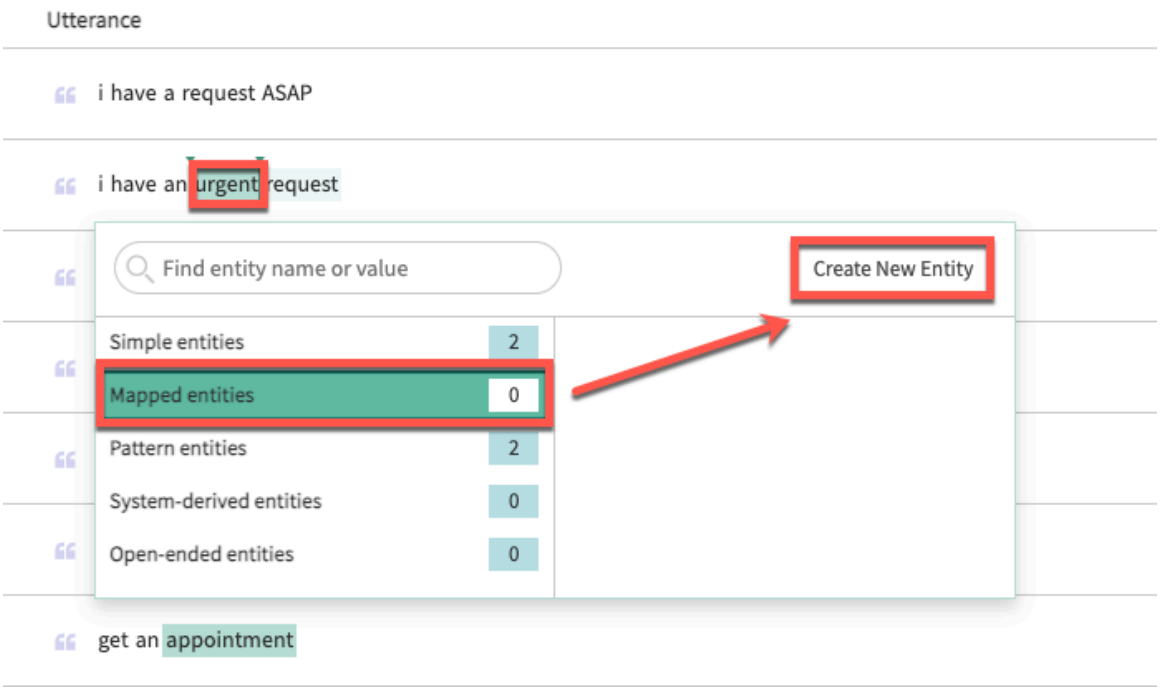
In this example procedure, you create a mapped entity for urgency.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, then the name of your model.
3. On the model details page, select the **Intents** tab.
4. In the Intents section of the model, select the name of an intent.
   For this example procedure, you select **#SubmitRequest**.
5. In the Utterances tab, select a word in an utterance

   In this scenario, you select the word urgent in the utterance I have an urgent request.
6. Select **Mapped entities**.
7. Select **Create New Entity.**



8. On the form, configure the fields.

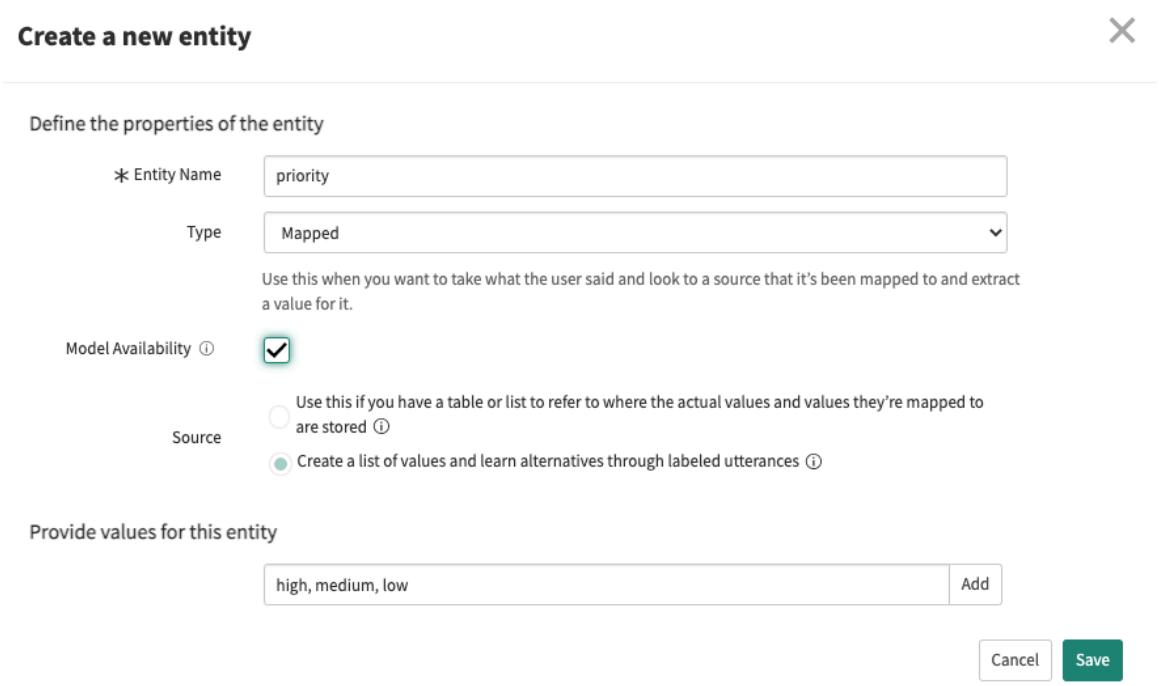
Field	Description
Entity Name	Name for the entity.
Type	Type of entity.
Model availability	Select this option if you want this entity to be included in all intents in your model.
Source	Source of the entity values.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Provide values for this entity	Values used to provide context for the model.

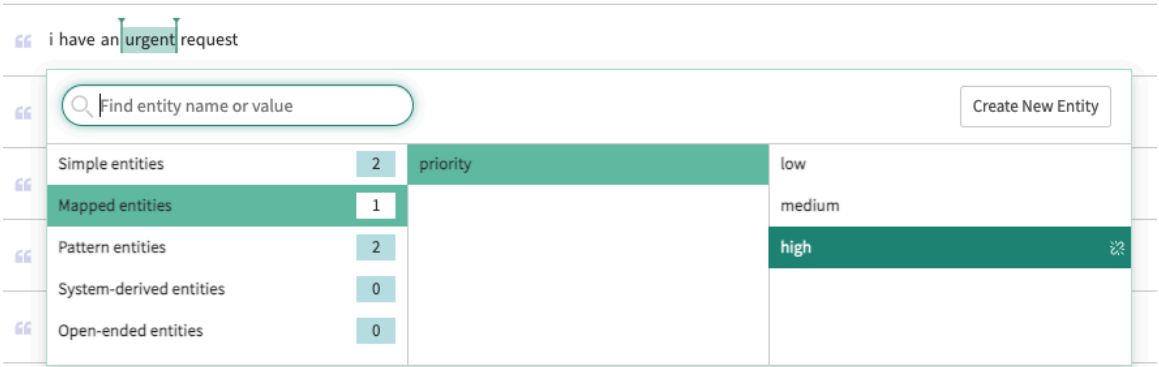
For this example procedure, use the following configurations:

- **Entity Name**: `priority`
- **Type**: Mapped
- **Model availability**: Select the check box
- **Source: Use this if you have a table or list to refer to where the actual values and values they're mapped to are stored**
- **Mapped value for the entity**: `high, medium, low`.



9. Click **Save**.

Result: Your mapped entity saves. The entity appears on the **Associated entities** tab. Now the model can leverage machine learning and use the values provided to identify possible values.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

You can create a mapped entity using a vocabulary source to use the values in the source as the mapped entity.

Related topics

Create a table vocabulary source

Create a list vocabulary source

Create a pattern entity

Create a pattern entity from a word or phrase with repeatable patterns, such as email addresses and phone numbers. These patterns help the system to recognize similar utterances based on the patterns.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model for Virtual Agent or AI Search.
- Create or use an existing intent.
- Role required: nlu_editor, nlu_admin, or admin. The nlu_editor must be assigned to the model.

About this task

Pattern entities help your model identify, contextualize, and govern similar types of content. Inputs such as phone numbers and email address follow the same pattern. By creating a pattern entity, you can train the model to recognize all variations of those inputs.

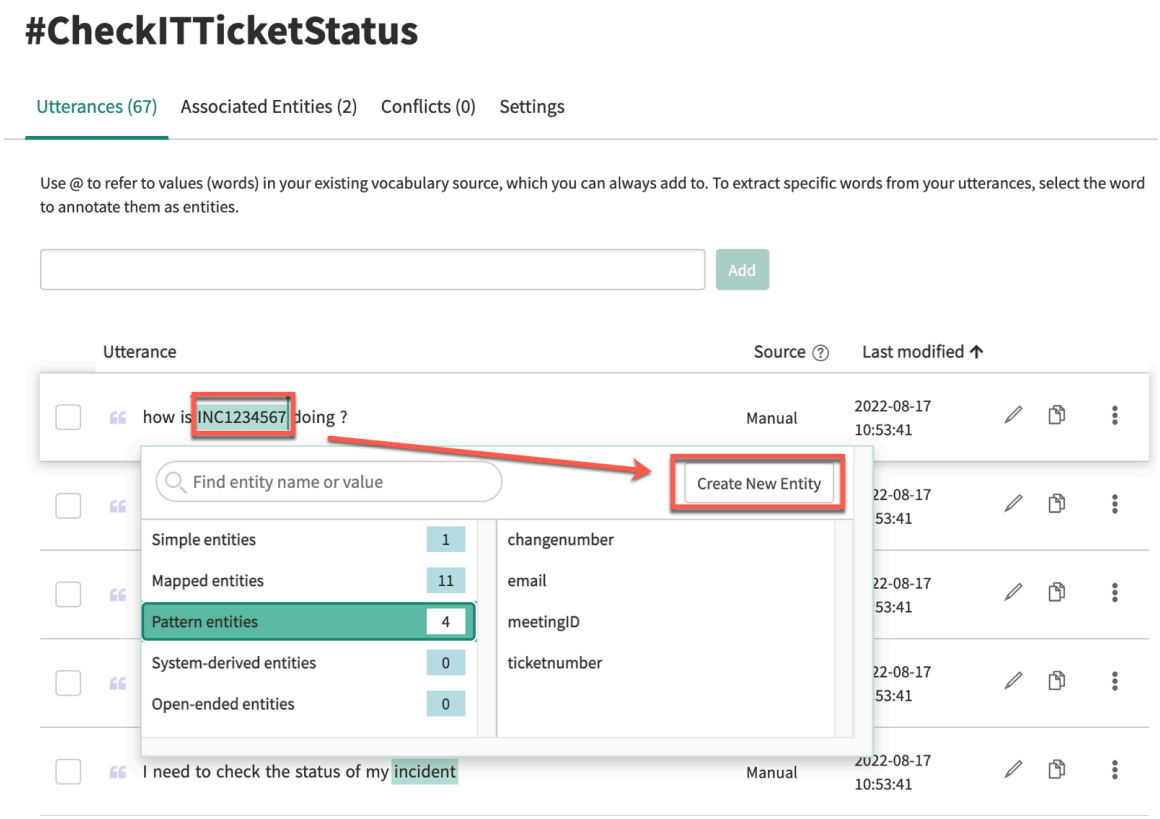
Note:
Pattern entities use regular expressions (regex). The regex field value is a Java regular expression. For more information on how to use regex, see Using regular expressions in entities.

In this example scenario, you've created an intent that's titled #CheckITTicketStatus. In this example procedure, you're creating a pattern entity for incident record numbers.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, then the name of the model.
3. On the model details page, select **Intents**.
4. Select the name of an intent.
   In this example scenario, you select **#CheckITTicketStatus**.
5. Select one of the words of an utterance.
   In this scenario, you click **INC1234567**.

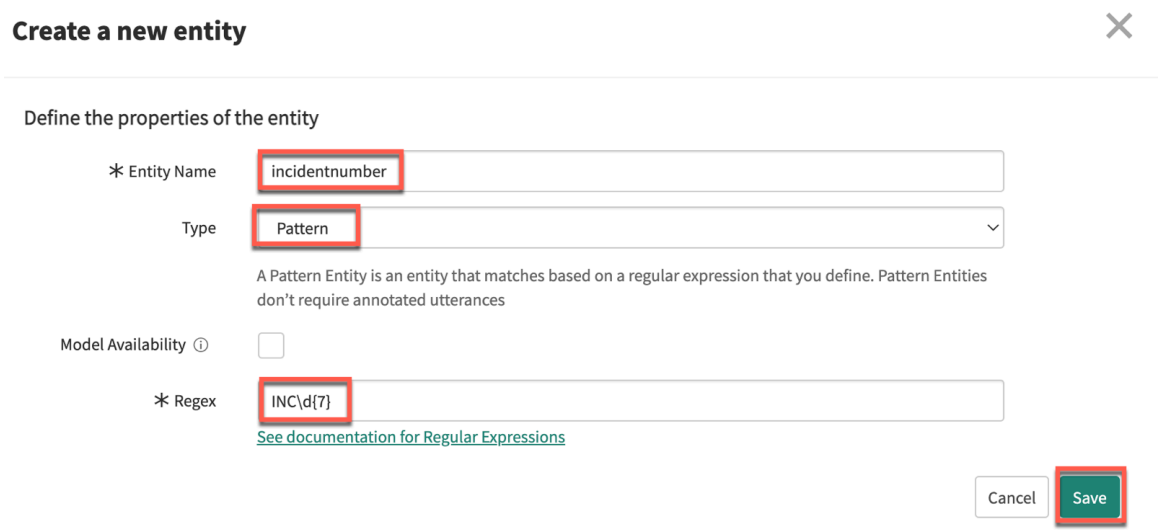
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



6. Select **Create New Entity**.
7. In the Create a new entity window, configure the fields.

   For this example procedure, use these configurations:

   - **Entity Name**: Enter `incidentnumber`
   - **Type**: Select **Pattern**
   - **Model Availability**: Select this box if you want this entity to be included in all intents in your model
   - **Regex**: Enter `INC\d{7}`



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. Click **Save**.
   The pattern entity saves and appears in the **Associated Entities** tab. The model can use the pattern entity to interpret variations of the annotated utterance. For example, the model can match the #CheckITTicketStatus intent if a user enters any incident record number.

Create a system-derived entity

Create a custom entity that's derived from a default system entity such as date, time, duration, or location.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model for Virtual Agent or AI Search.
- Create or use an existing intent.
- Role required: nlu_editor, nlu_admin, or admin. The nlu_editor must be assigned to the model.

About this task

System entities are pre-built in your ServiceNow instance by default, and appear on the Entities section of the Model screen. These entities, such as DATE, TIME, and DATE_TIME, are enabled by default. You can disable and re-enable them if needed by clicking **Enable**.

System-derived entities extend your system entities, providing more context. For example, your model already understands date formats thanks to the system entity DATE. However, you can create system-derived entities, such as startdate and enddate, to extract more information from user utterances regarding dates.

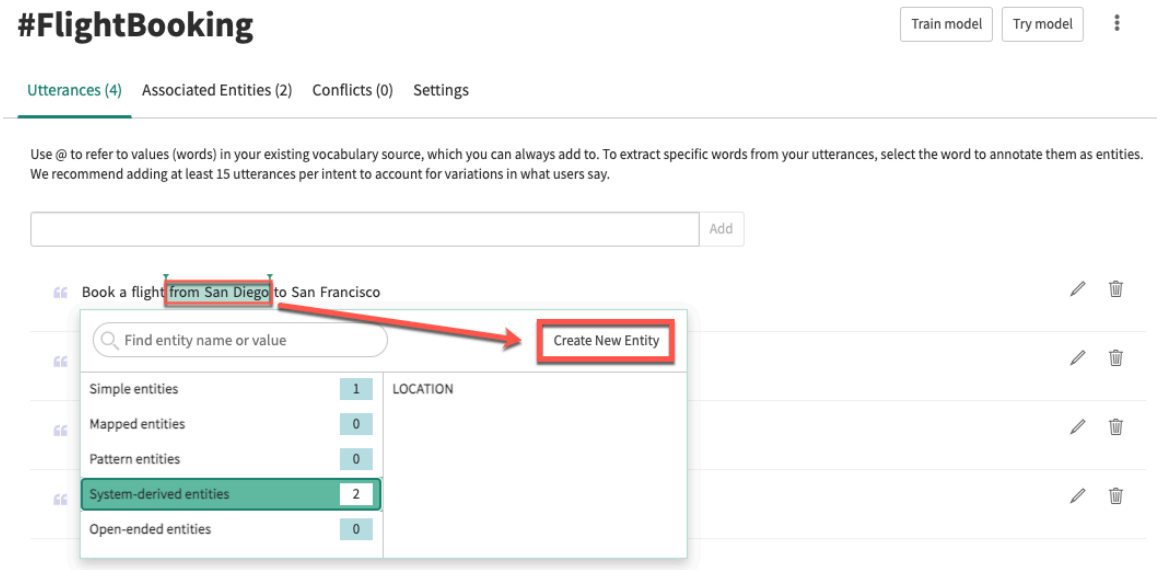
For the following example scenario, you're creating a model for booking flights, cars, hotels, and events. You want your intent #FlightBooking to interpret utterances with your users flight requests. Your model contains the system entity LOCATION, but flight plans usually involve two locations.

In this example procedure, you create two system-derived entities to gather the departure and arrival locations of the flight.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, then the name of your model.
3. In the model details page, select **Intents**.
4. Select the name of the intent.
   For this example, you select the **#FlightBooking** intent.
5. On the **Utterances** tab, select a word or phrase from one of the utterances.
   In this example scenario, you have an example utterances `book a flight from San Diego to San Francisco`. For this step, click **from San Diego**.

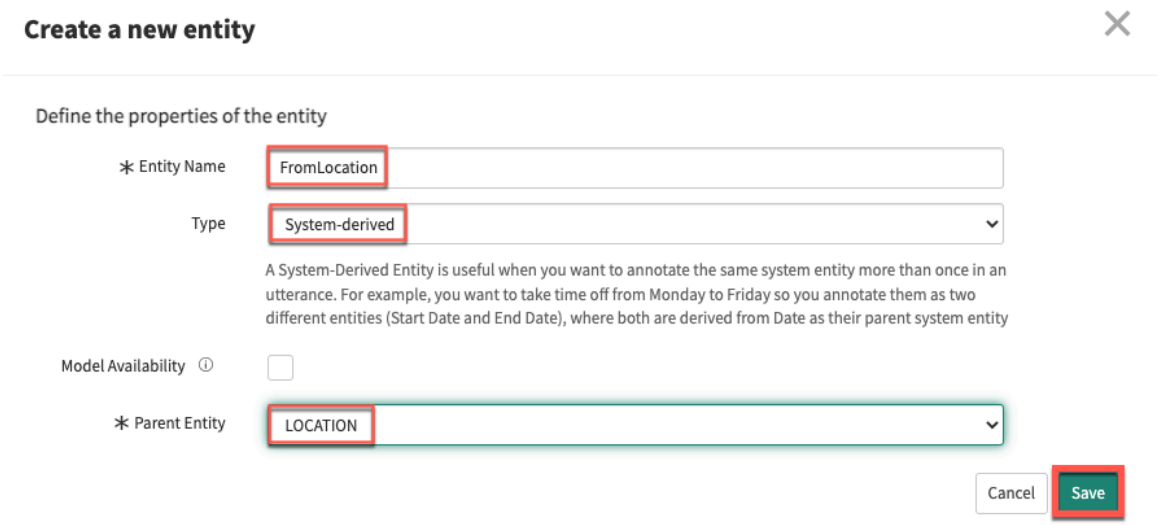
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



6. Select **Create New Entity**.
7. In the Create a new entity screen, configure the fields.

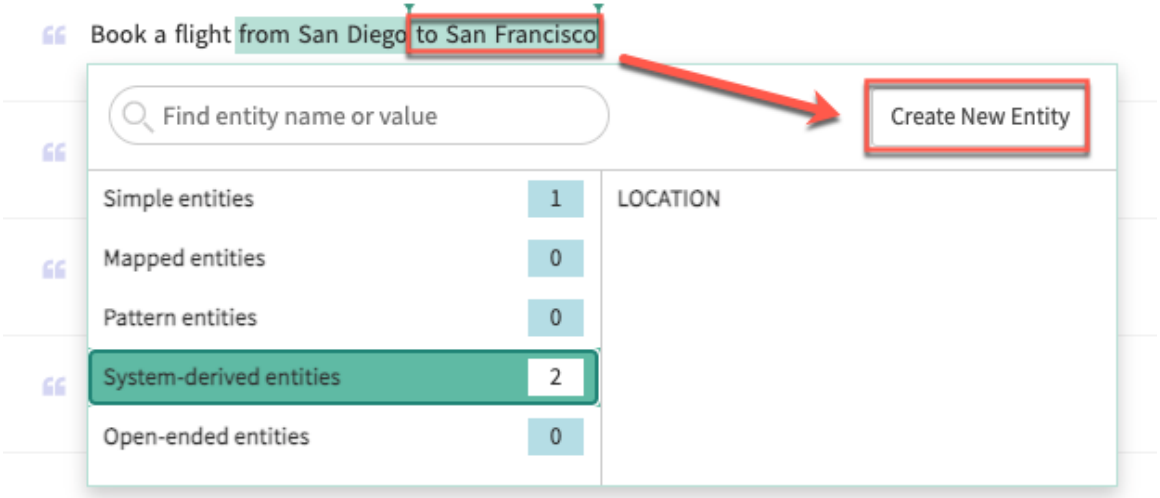
   For this example, use these configurations.

   - **Entity Name**: Enter `FromLocation`
   - **Type**: Select **System-derived**
   - **Model Availability**: Select this check box if you want this entity to be included in all intents in your model
   - Parent Entity: Select **LOCATION**



8. Select **Save**.
9. On the same utterance, select another word or phrase.
   In this example scenario, you click **to San Francisco**.

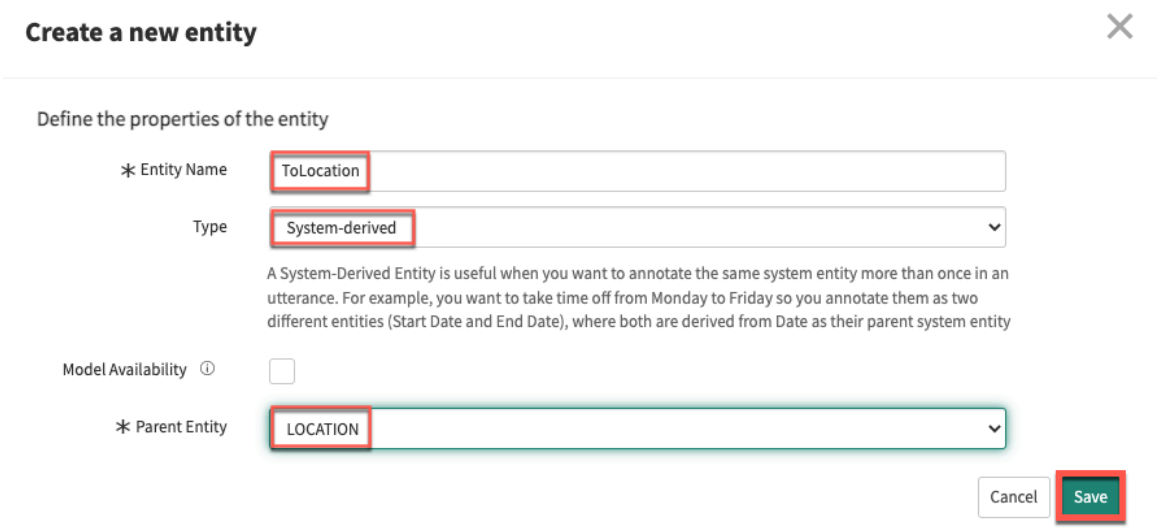
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



10. Select **Create New Entity**.

For this example, use these configurations.

- **Entity Name**: Enter `ToLocation`
- **Type**: Select **System-derived**
- **Model Availability**: Select this check box if you want this entity to be included in all intents in your model
- **Parent Entity**: Select **LOCATION**



11. Select **Save**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result: Your two system-derived entities appear in the entity window.



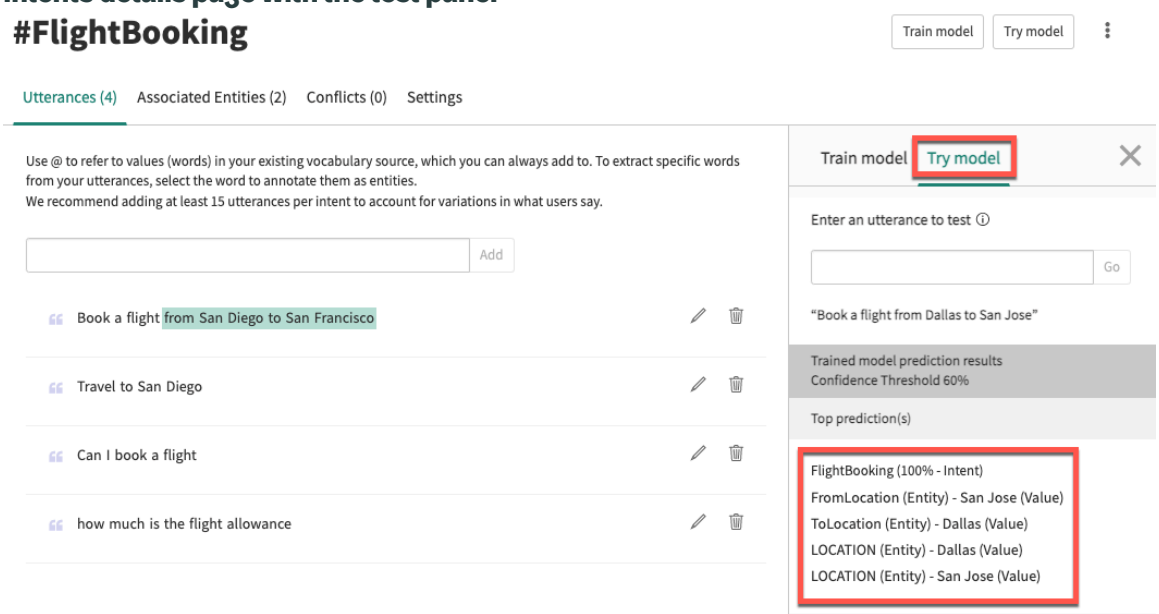
What to do next

Train your model to save the entities. You can try your model to see if it recognizes and interprets variations of the entities that you created.

For this example, you can test your model with a different departure and arrival location.

1. Select **Try model**.
2. Enter `book a flight from Dallas to San Jose`.
3. Select **Go**.

Intents details page with the test panel



The model predicts the intent, showing the entities it used to determine the values. Notice how it uses both the system entity and system-derived entity when predicting.

Create an open-ended entity

Use an open-ended entity when you want to improve intent prediction accuracy. Open-ended entities help your model focus on the context of the utterances.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model for Virtual Agent or AI Search.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Create or use an existing intent.
- Role required: nlu_editor, nlu_admin, or admin. The nlu_editor must be assigned to the model.

About this task

Open-ended entities tell the model to focus on the context of the entity rather than the entity itself. When you mark a word or phrase as open-ended, the system skips the entity and predicts the intent from the context that precedes or follows the entity in the utterance.

For example, in the utterance `I want to order an iPhone`, you annotate the words "an iPhone" as an open-ended entity. The model focuses on the context, predicting the user wants to order something. Since there are numerous things the user could want to order, naming all of them would be an unbearable task for the model author.

Using an open-ended entity instead of a simple entity helps the model focus on the rest of the utterance and not the entity. In the iPhone example, the entity itself is less relevant; so you want the system to ignore it.

In other scenarios you should use a simple entity, as there could be multiple intents where you shouldn't have the system ignore the entity.

Note:
You can't annotate a vocabulary source (referenced by @vocab_source in an utterance) as an open ended entity. You can only annotate a vocabulary source as a simple entity or a mapped entity. For example, if the utterance is “I want to order a laptop”, then the word “laptop” can be annotated as an open ended entity. However, if the utterance is “I want to order @laptop” where @laptop refers to a table vocabulary source or a list vocabulary source, it can't be annotated as an open ended entity.

For this example scenario, you've created an NLU model with an intent for your users to order company merchandise.

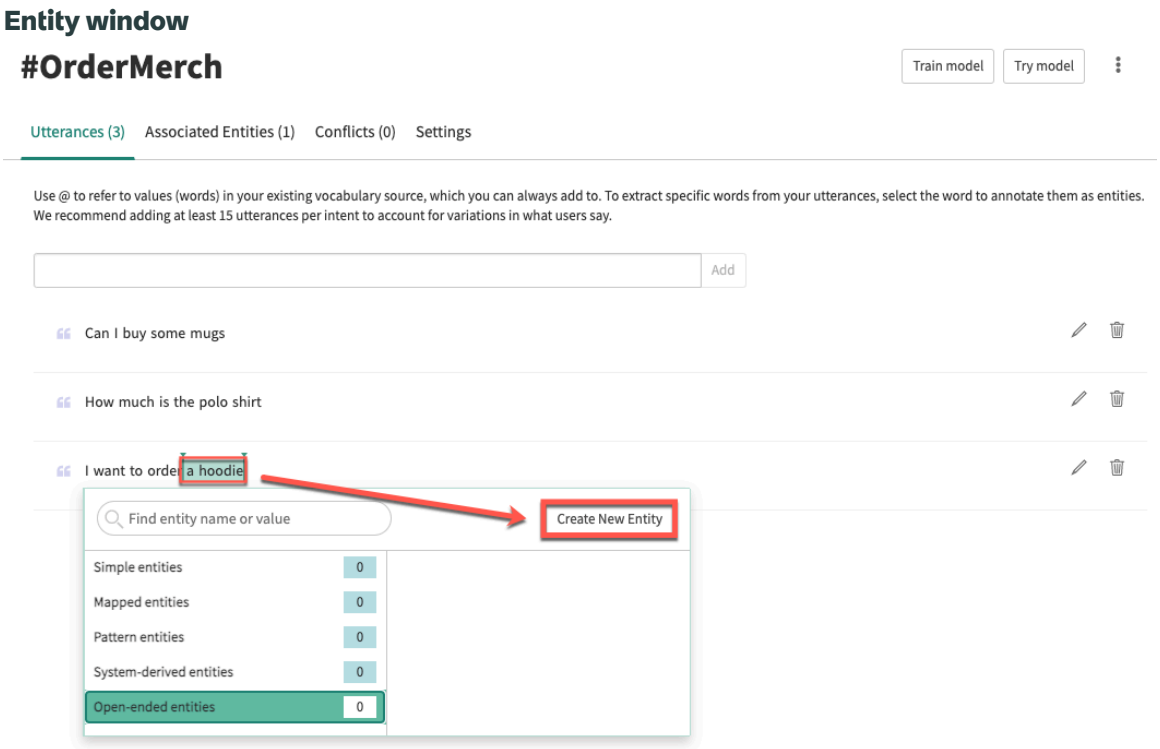
In the following example procedure, you create an entity from one of your utterances so the system can recognize it as open-ended and reusable in other NLU models in your instance.

Note:
You can use only one open-ended entity per intent.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, then select the name of your model.
3. On the Model details page, click **Intents**.
4. Select the name of the intent you want to add the entity to.
   For this example, you select the #OrderMerch intent.
5. In the Utterances tab, select a word or phrase of one of the utterances to bring up the entities window.
   For this example, you select a hoodie.

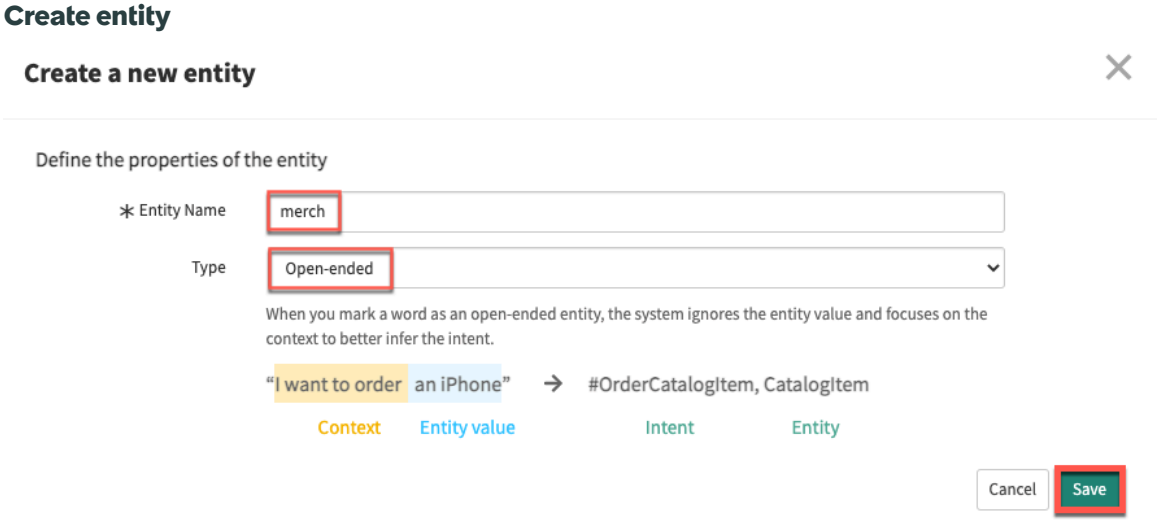
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



6. Select **Create New Entity**
7. On the Create a new entity screen, configure the fields.

   For this example, use the following configurations:

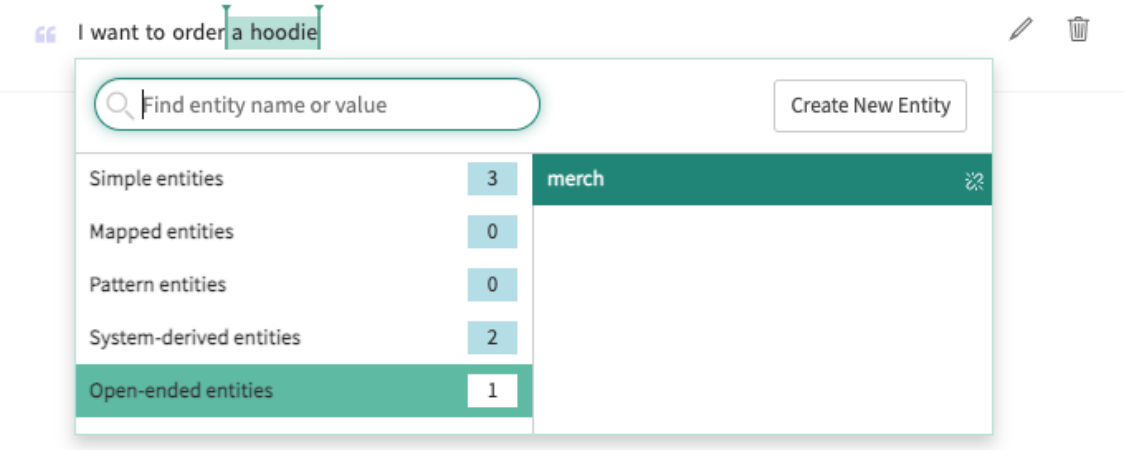
   - **Entity Name**: `merch`
   - **Type**: Select **Open-Ended**



8. Select **Save**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The merch open-ended entity is annotated in the Utterances section of your model's Intent screen. When you point to its name, you can see that it persists as a new entity in the annotation details. This entity is reusable in all other NLU models in your instance.

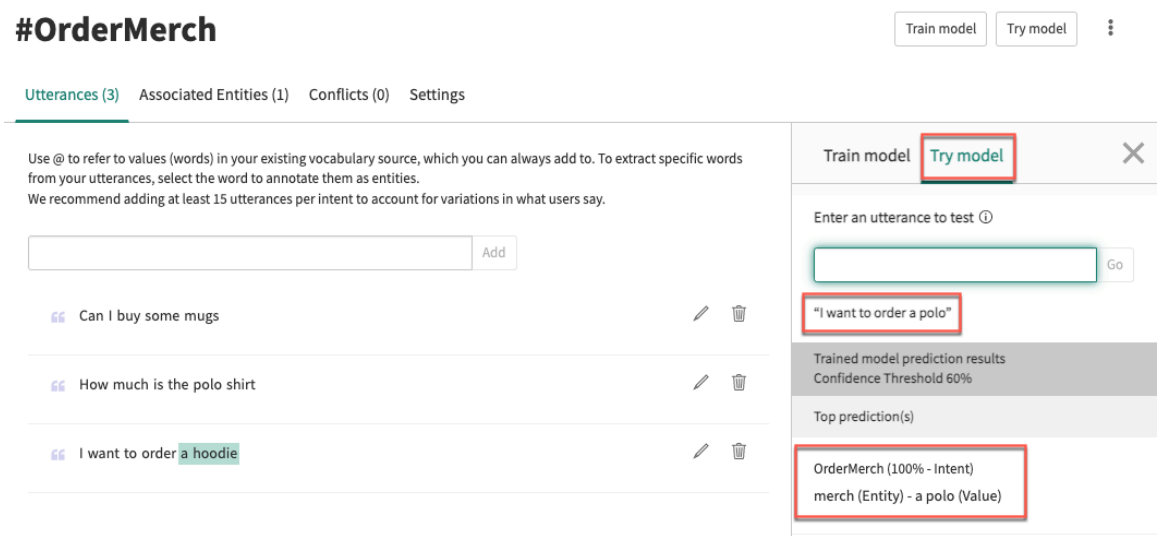


What to do next

Train your model to save the entities. You can try your model to see if it interprets the utterance based on the context of the entity, rather than the entity itself.

For this example, you can test your model with a different merchandise item.

1. Select **Try Model**.
2. Enter `I want to order a polo`.
3. Select **Go**.



The model predicts the intent and shows that it used the merch entity for the a polo value.

Import entities

Reuse entities that you have created across your other Natural Language Understanding (NLU) models. Importing entities saves time and helps improve the intents in your model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model.
- Create or use existing intents and entities.
- Role required: nlu_editor, nlu_admin, or admin. The nlu_editor must be assigned to the model.

About this task

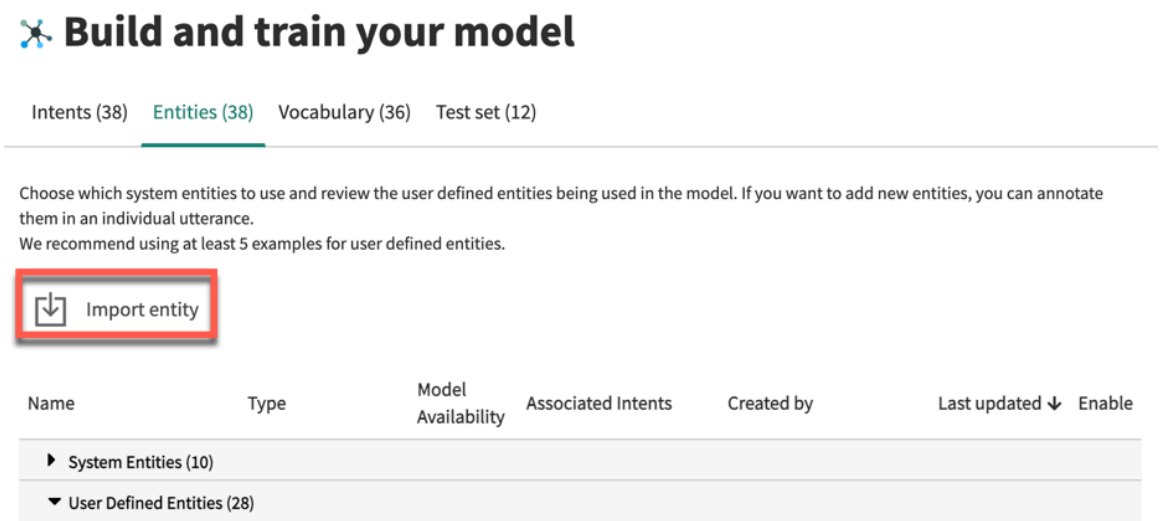
New entities are usually created inside an individual intent. To help save time and improve your intents, you can also reuse entities from your other models by importing them. After importing entities into a model, annotate the model's utterances using the new entities.

Note:
All models include several system entities by default.

In this example procedure, you are importing several entities to a new model.

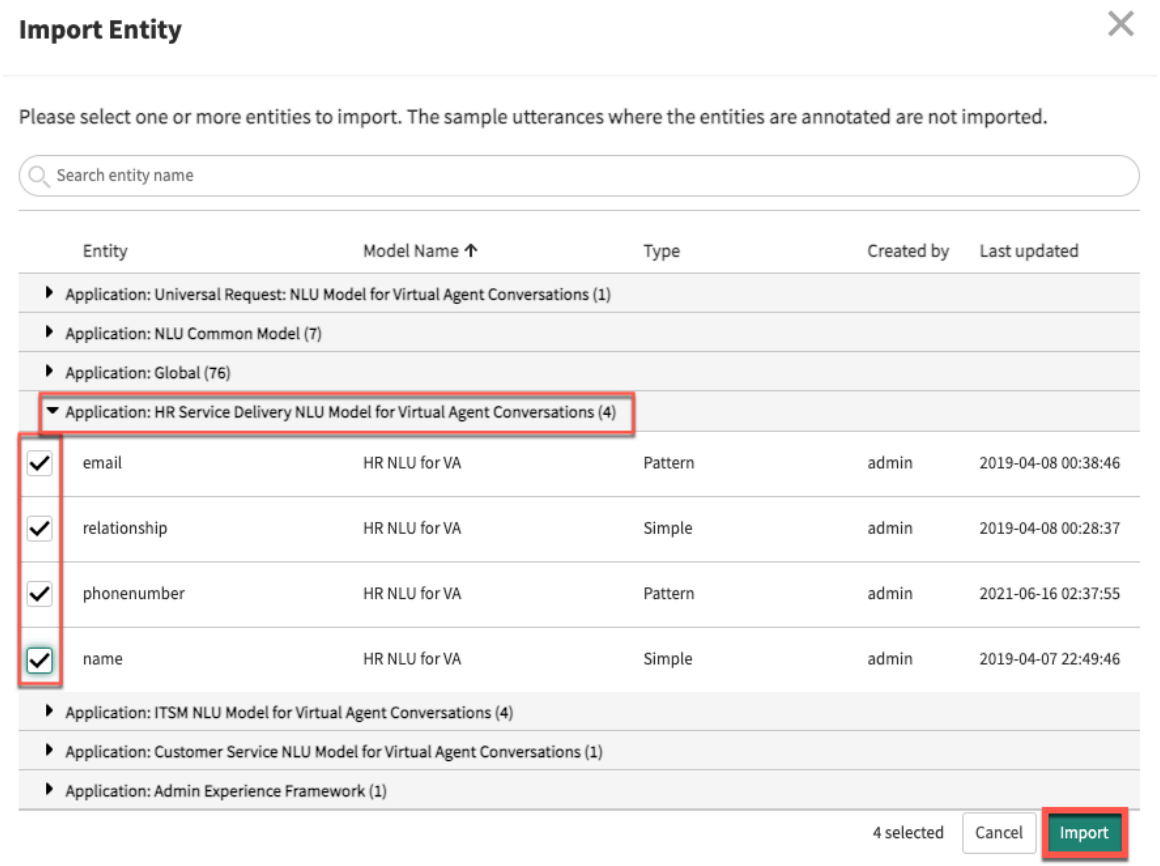
Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab corresponding to your model's application, then select the name of the model you want to add entities to.
3. In the **Build and train your model** phase, select the **Entities** tab.
4. Select **Import entity**.



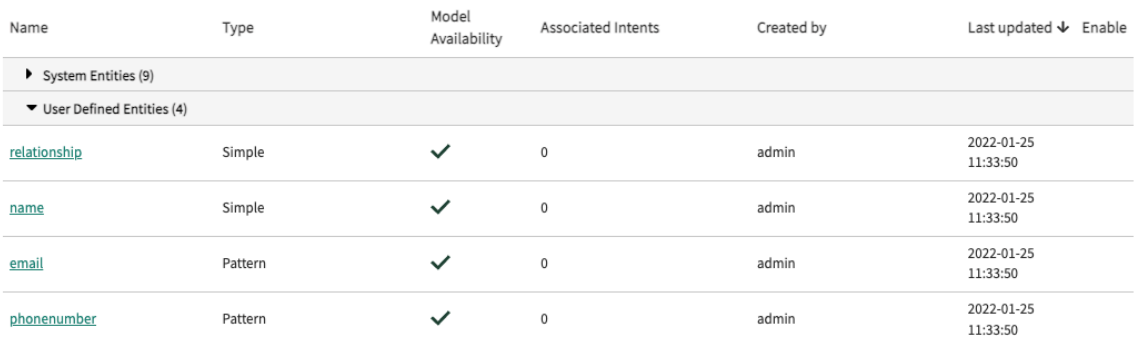
5. In the Import Entity window, select the entities you want to import.
   In this example, you click the drop-down arrow for Application: HR Service Delivery NLU Model for Virtual Agent Conversations. Select the four entities.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Note:
Importing entities does not import any utterances that are annotated with those entities. You must annotate your current model's utterances using the new imported entities, then retrain the model.

6. Click **Import**.
   The selected entities import to the model. The entities appear under **User Defined Entities** in the Entities tab. By default, **Model Availability** will be enabled so that all intents in the model can use the entity.



What to do next

Use the imported entities to annotate the utterances of the model. Train your model to save your changes.

Related topics

Import an NLU intent

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Using regular expressions in entities

Learn how to use regular expressions in your NLU entities to establish patterns that help the system locate, match, and manage text.

Pattern entities use regular expressions (regex) to match any pattern of text, such as the format of an email address, a phone number, or an incident or case ID.

Note:
ServiceNow uses and supports Java regex exclusively and not other vendor options, such as Perl regex.

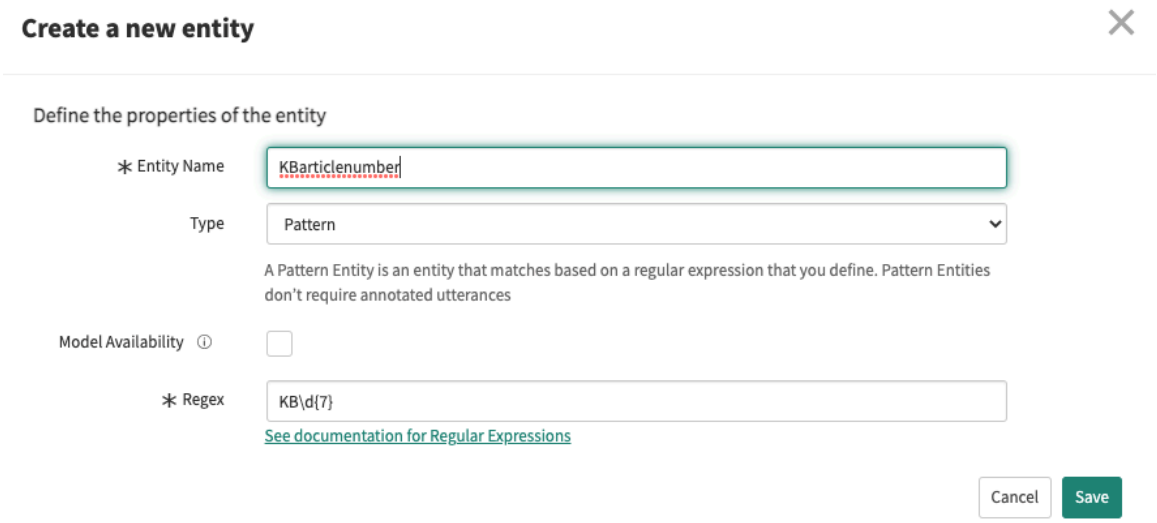
Regex examples

For a full example of regex code used in a pattern entity, see Create a pattern entity.

The following four examples provide the format and regex code for different pattern entities.

Knowledge base article

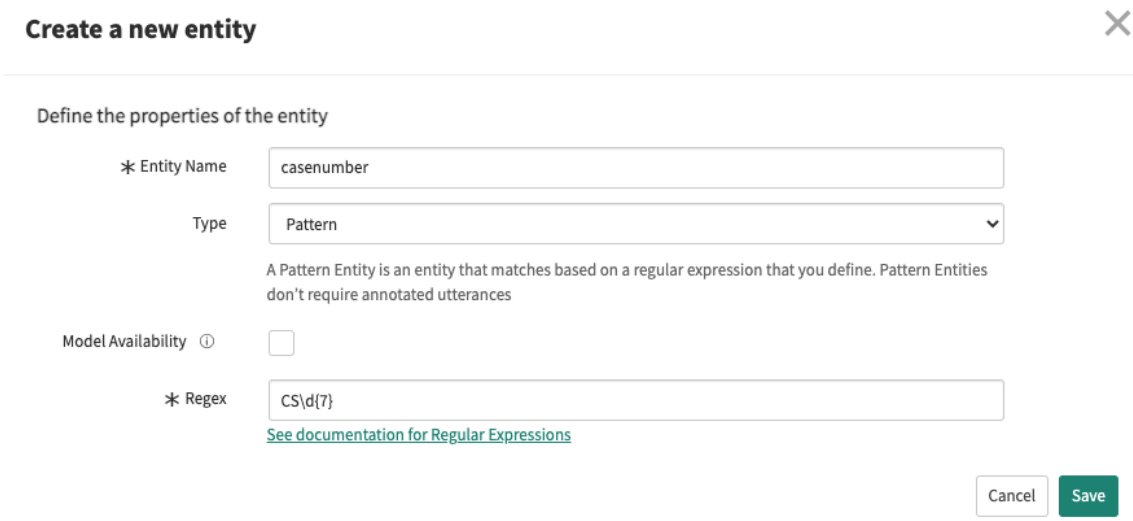
ServiceNow Knowledge Base (KB) articles use the format KB1234567. The regex code is KB\d{7}, where KB = knowledge base record and d{7} = 7 digits.



Case number

Case tasks use the format CS1234567. The regex code is CS\d{7} where CS = case record and d{7} = 7 digits.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Email address

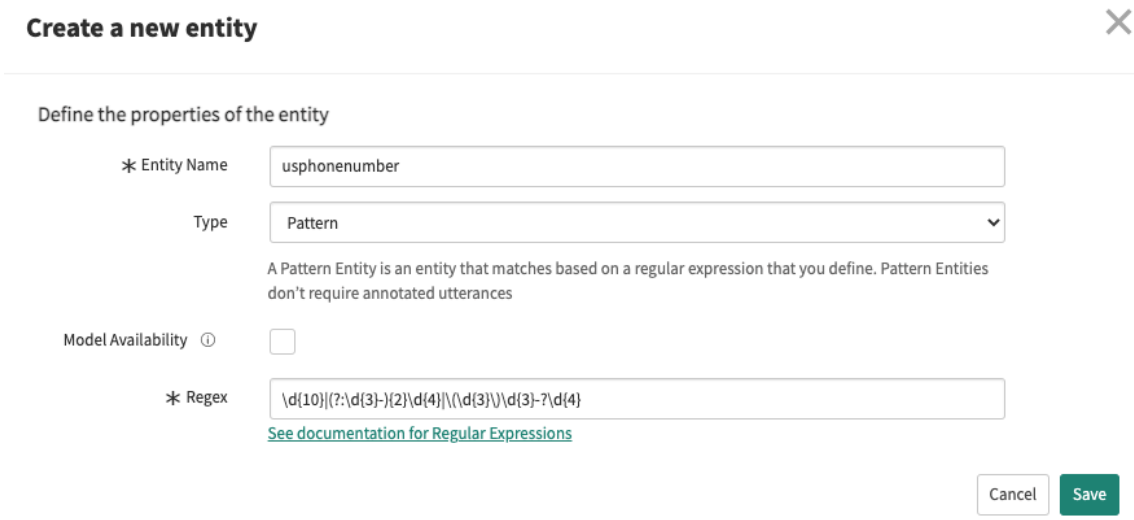
Email addresses use the format name@servicenow.com. The regex code is \b[a-zA-Z0-9&*/_.+-]+@[a-zA-Z0-9-]+\.[a-zA-Z0-9-.]+\b, which supports these formats: name@servicenow.com and name@servicenow.co.in.



Phone number

United States phone numbers use the format 555-123-4567. The regex code is \d{10}|(?:\d{3}-){2}\d{4}|\(\d{3}\)\d{3}-?\d{4}, which supports these formats: 5108882062,510-888-2062, and (510)888-2062.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Regex resources

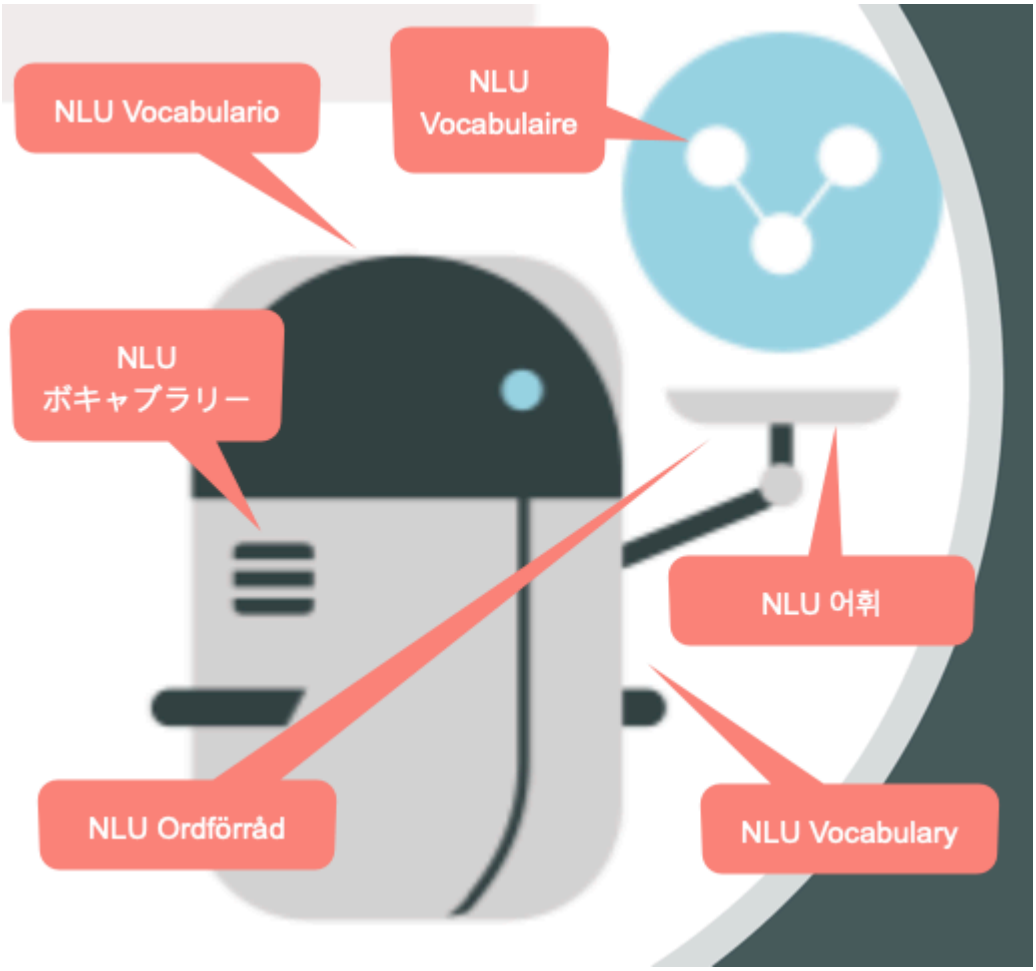
For further information on regular expressions, see the following external links.

- Java Regular Expressions
- Java Regular Expression Tester
- Pattern (Java Platform SE7)
- Java regex match abbreviations

NLU vocabulary

Use NLU vocabulary items to help the system recognize the various ways your users express their requests. Use vocabulary sources to help the system recognize objects in tables or lists, such as names of conference rooms or catalog items.

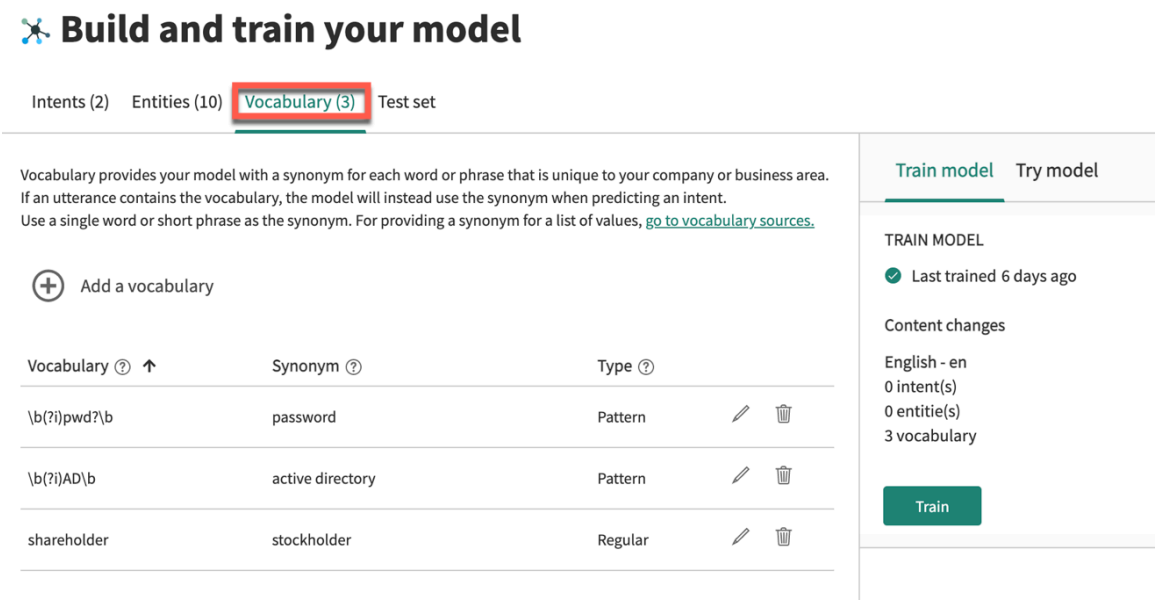
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Vocabulary usage and context

Vocabulary helps your model with the various words and phrases that it may encounter from your users. Since humans are not computers, they may come up with different ways to say the same thing. For example, if one user inputs an acronym rather than the full phrase, the system might be able to predict the correct intent by using the context in the utterance. However, by defining the acronym with a vocabulary item, you raise the model's confidence level and capability to predict intents correctly.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Vocabulary items cover the varying words or phrases that might appear in utterances. The vocabulary items are mapped to the synonyms that you provide, for intent prediction. This helps the system recognize the various ways your users express their requests, while at the same time ensuring the consistency, confidence, and accuracy of predictions.

Vocabulary item types

NLU provides the following vocabulary types that you can use to create and configure a vocabulary item.

Type	Definition
Regular	A word or phrase that is not commonly known, such as a business or industry specific term or acronym. Regular vocabulary is case insensitive, so all case variations will be captured by the vocabulary you create. See Create a regular vocabulary item.
Pattern	A regular expression (regex) that can capture specific formats such as email addresses and phone numbers. See Create a pattern vocabulary item.

Vocabulary sources

You can also use vocabulary sources to cover various objects that your users might request. Vocabulary items and vocabulary sources differ in their usage:

- Use a vocabulary item for an individual word, phrase, or pattern that can easily be mapped to a single synonym for the model to use.
- Use a vocabulary source to reference a ServiceNow table or list so that the values can all be replaced by the single synonym you define. Vocabulary sources can be reused across all your NLU models.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use vocabulary sources when your user utterances refer to objects in a set. The sources can be referenced from a list of values from a ServiceNow table or a static list that you define, such as a list of catalog items or conference rooms. Once you create and save these sources, you can use the @ symbol to specify them in training utterances. You can also use these sources as entity values.

For example, say you already have a list of all the conference rooms for your office. You can create a vocabulary source to reference that list rather than typing each conference room name into your intent's training utterances.

To create vocabulary sources, refer to Create a table vocabulary source and Create a list vocabulary source.

Vocabulary usage in relation to an intent

Here's an example of how an intent can interact with the vocabulary in its utterance examples.

Note:
In NLU vocabulary, the synonym replaces the vocabulary that appears in the utterance.

Intent: OrderSoftware

Utterance	Issue and Solution
"I need to access sfcrm"	• Issue: the system doesn't recognize the acronym sfcrm, so can't accurately predict the intent. • Solution: Add sfcrm as a regular vocabulary item, and provide a synonym such as CRM software.
"I need to install Word"	• Issue: The term word is very common and does not necessarily indicate a software product. The term install may provide helpful context, but the confidence score may be too low to predict the OrderSoftware intent. • Solution: Create a pattern vocabulary item with a regex for capitalized Word, so that the system can recognize it as a software product. This added constraint makes it more likely that the OrderSoftware intent will be predicted. **Note:** To extract the specific software name to use in a Virtual Agent conversation topic, annotate it as an entity in the utterance.

Use a single word or short phrase as a synonym for best results. You can map multiple vocabulary items to one synonym. Do not map one vocabulary item to multiple synonyms.

For more context and examples, see Create an NLU intent.

A Regex example for a pattern vocabulary item

Let's say that you want to use a vocabulary item to identify the acronym IT, and map it to the synonym information technology.

Regular vocabulary items are case-insensitive by default. This means that a regular vocabulary item would match both the acronym IT and the common word it. So you decide to use a pattern vocabulary item.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The regex pattern that you can use is `\bIT\b`. The `\b` is a word boundary marker. This ensures that the pattern would not match ITSM or JIT, for example. The default case sensitivity of pattern vocabulary items means that `\bIT\b` would not match the common word it.

Regex details: You can turn off case sensitivity in a pattern by using `(?i)`. You can end that mode by using a minus sign as in `(?-i)`. For example, `(?i)te(?-i)st` should match both test and TEst, but not teST or TEST.

For more information, see Using regular expressions in entities.

Related topics

Create an NLU intent

Create a regular vocabulary item

Add a word or phrase that your users might use, and match that vocabulary item to a synonym. Your model uses the synonym during intent prediction.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model for Virtual Agent or AI Search.
- Role required: nlu_editor, nlu_admin, or admin. The editor must be assigned to the model.

About this task

Regular vocabulary items provide the model with a synonym for words or phrases your users might use in an utterance. The model uses the synonym to replace the vocabulary when predicting the intent. Use a single word for the synonym when possible.

Regular vocabulary items are case-insensitive by default. If you need to create a case-sensitive vocabulary item, use a pattern vocabulary item. For more information, see Create a pattern vocabulary item.

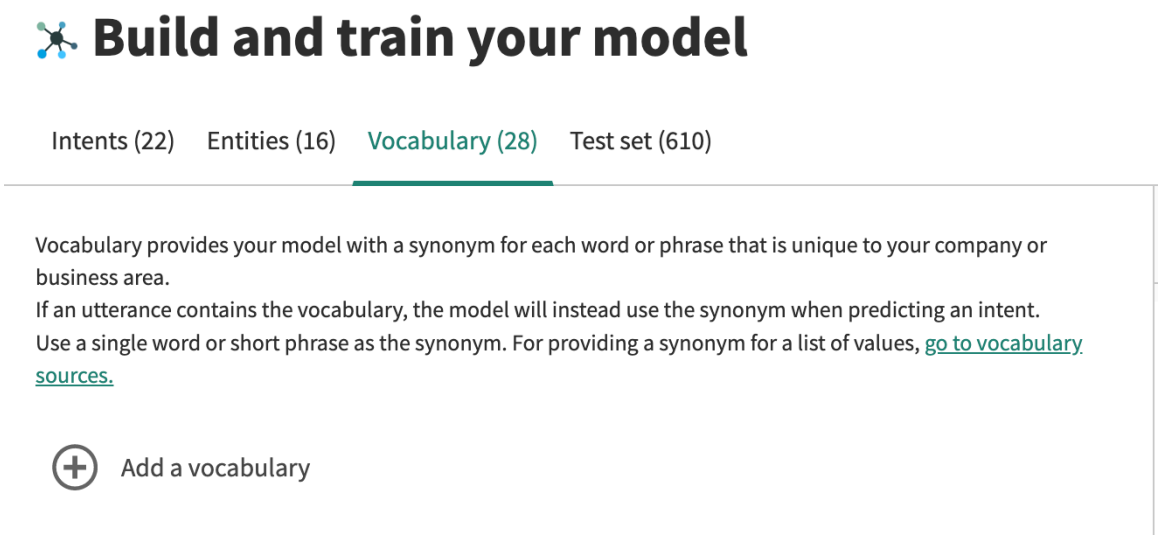
Note:
Choose a synonym that is a commonly-occurring word in the same language as your model.

In this example scenario, you are adding a vocabulary item for the word credentials. Say that your users may use credentials to refer to their password. By creating a vocabulary item, you can make sure that the system correctly predicts the intent for an utterance such as reset my credentials.

Procedure

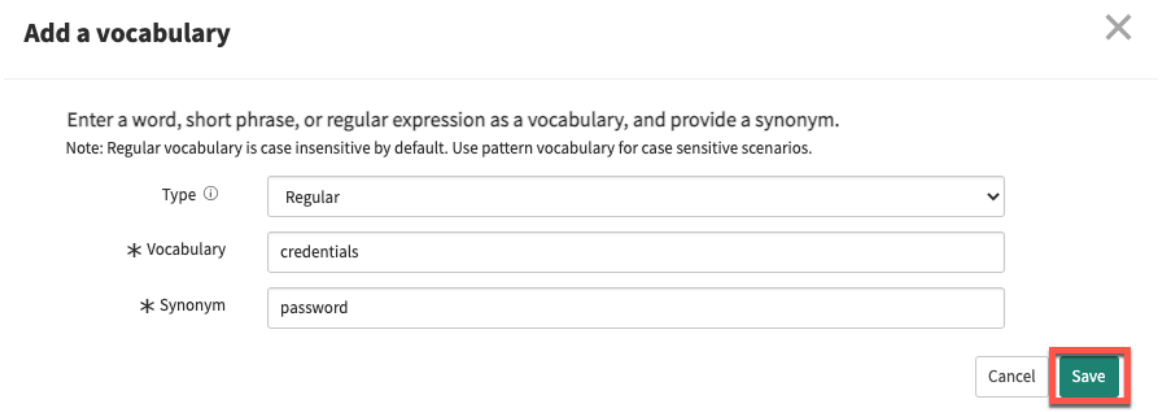
1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab corresponding to your model's application, then select the name of your model.
3. On the Model details tab of the model overview, select the **Vocabulary** card.
4. In the Vocabulary tab, select **Add a vocabulary**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



5. In the **Add a vocabulary** window, select **Regular** as the Type.
6. Add a vocabulary word or phrase that your users might use, and then add the synonym that the model should use for intent prediction.

   In this example procedure, add credentials as the vocabulary and password as the synonym.



7. Select **Save**.

What to do next

To deploy your new vocabulary item, train and publish your model again.

Add more vocabulary items to improve model coverage and accuracy.

Create a pattern vocabulary item

Use regular expression (regex) encoding to establish a pattern format for vocabulary items such as email addresses, phone numbers, and record naming conventions. You can create your own patterns for the vocabulary data in your instance.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Create or use an existing NLU model for Virtual Agent or AI Search.
- Role required: nlu_editor, nlu_admin, or admin. The editor must be assigned to the model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

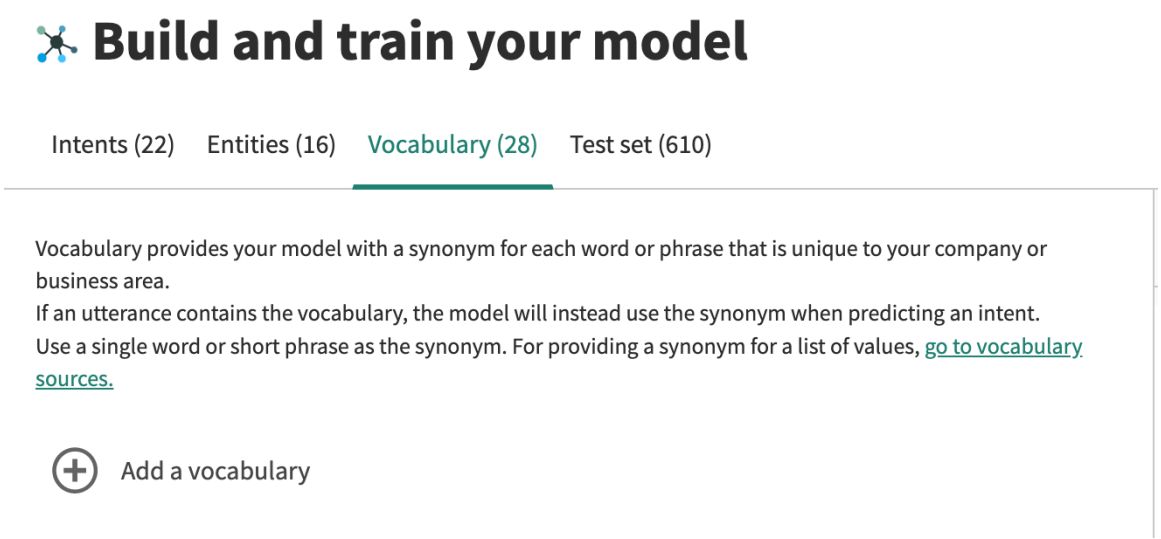
Pattern vocabulary items use regular expressions (regex) to find variations of words and phrases in your users' utterances. The pattern then maps these variations to the synonym you provide. The synonym is used during intent prediction.

In this example scenario, you've created an NLU Virtual Agent model for requests related to authentication. To refer to multi-factor authentication, your users often use "MFA", "mfa", or "mfa authentication". You can provide a pattern to cover these variations. They are replaced with the matching synonym multi-factor authentication when predicting an intent.

See Using regular expressions in entities for further details on regular expressions.

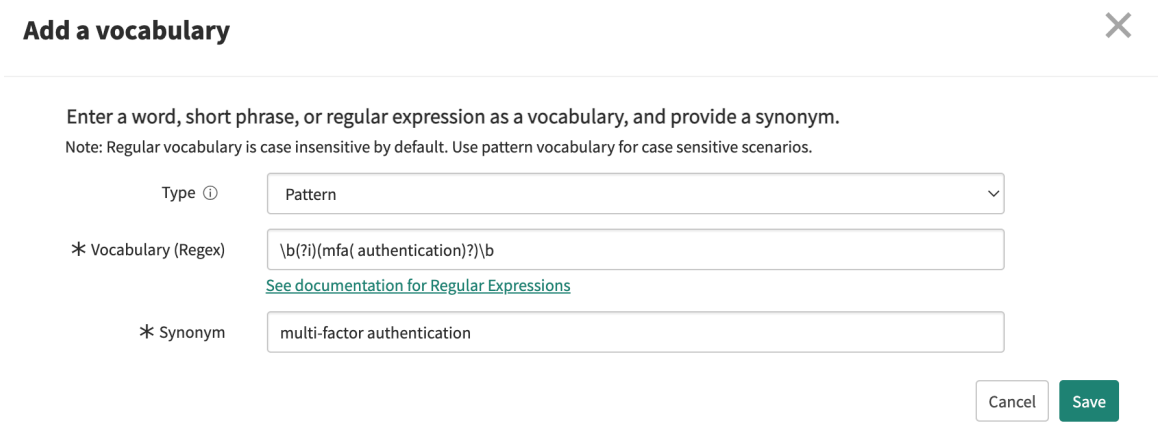
Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab corresponding to your model's application, then select the name of your model.
3. On the Model details tab, select the **Vocabulary** card.
4. Click **Add a vocabulary**.



5. In the **Type** field of the **Add a vocabulary** window, select **Pattern**.
6. Add a regular expression that covers words or phrases in your users' utterances, and a synonym the model uses to predict the intent.
   In this example, the regular expression covers several variations of "mfa authentication".

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



7. Click **Save**.

   Result: Your pattern vocabulary item appears in the Vocabulary section of your Model screen. Each time a user's utterance contains one of the spelling variations, the system recognizes the meaning as multi-factor authentication, in this example.

Related topics

NLU vocabulary

Using regular expressions in entities

Create a list vocabulary source

Create a list of words or phrases to act as a vocabulary source. The values in the list source are replaced by the synonym if they are detected in a user utterance.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Role required: admin, nlu_admin, or nlu_editor

About this task

Creating a list vocabulary source enables your models to interpret all the alternate and actual values of the list if they appear in a user utterance. The model interprets those values as the synonym you provide when predicting an intent.

> **Note:**
> During intent prediction, the synonym you provide replaces the value. During entity prediction, the actual value is used as the entity if the actual value or one of the alternate values are detected in the utterance.

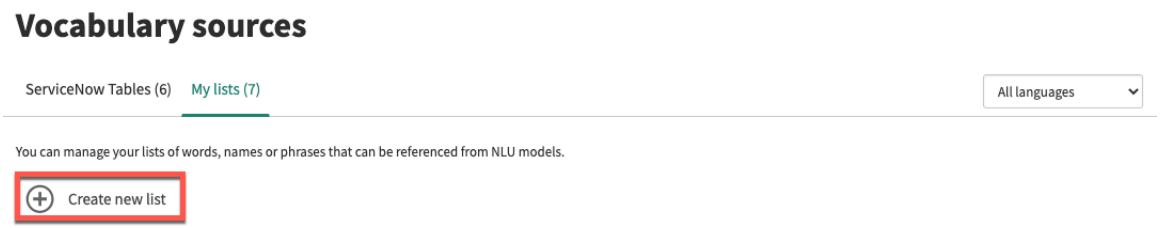
In this example procedure, you're creating a list vocabulary source for your company's meeting rooms. You add the names of the rooms and alternatives your users might call them. Once you create the list, you add it to an utterance in an intent for the model. The model interprets the intent (for example, #BookMeetingRoom) and uses the name the user entered as an entity (for example, Everest).

Procedure

1. Navigate to **All** > **NLU Workbench** > **Vocabulary Sources**.
2. Click the **My lists** tab.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Click **Create new list**.

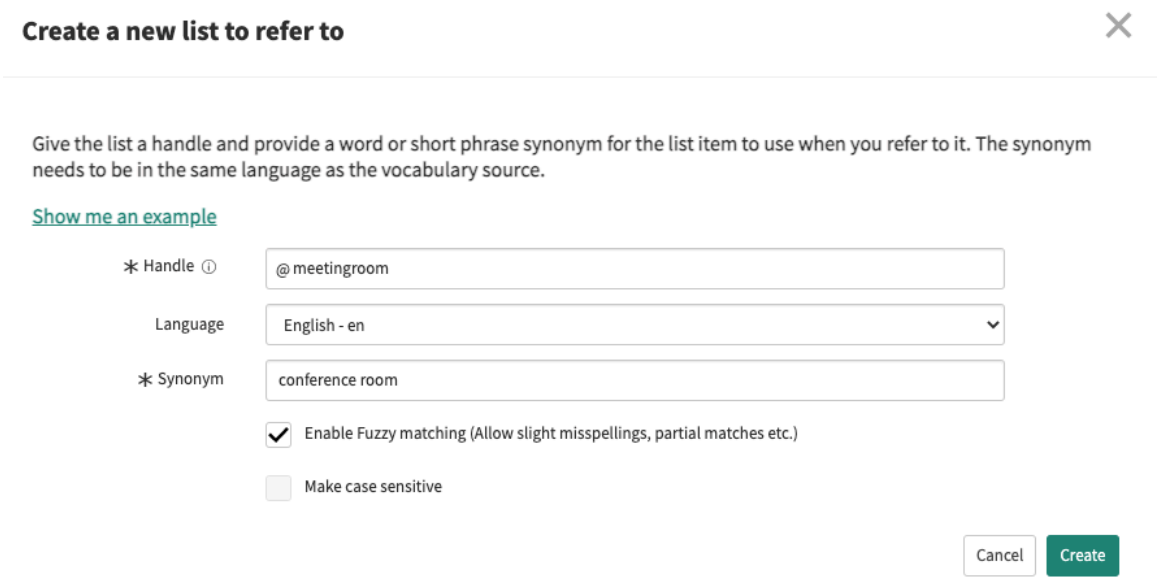


4. In the **Create a new list to refer to** window, configure the fields.

Field	Description
Handle	Name for the vocabulary source. Used to refer to in an utterance.
Language	Language of the vocabulary source. Synonym must be in the same language.
Synonym	Word or phrase the model uses during intent prediction. Choose a commonly-occurring word in the same language as your model.
Enable Fuzzy Matching	Check this box if you want the items on the list to match when a user utterance has slight misspellings.
Make case sensitive	Check this box to make the values in the list case sensitive. Utterances with wrong cases won't match.

For this example, use the following configurations:

- **Handle**: `@meetingroom`
- **Language**: `English - en`
- **Synonym**: `meeting room`
- **Enable Fuzzy Matching**: Select the box.
- **Make case sensitive**: Leave the box clear.

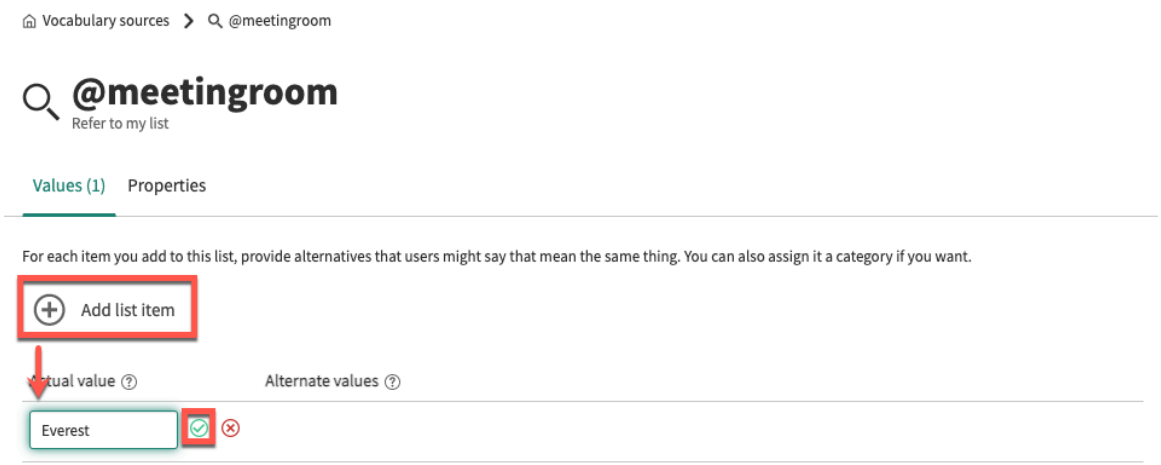


5. Click **Create**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Your 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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

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

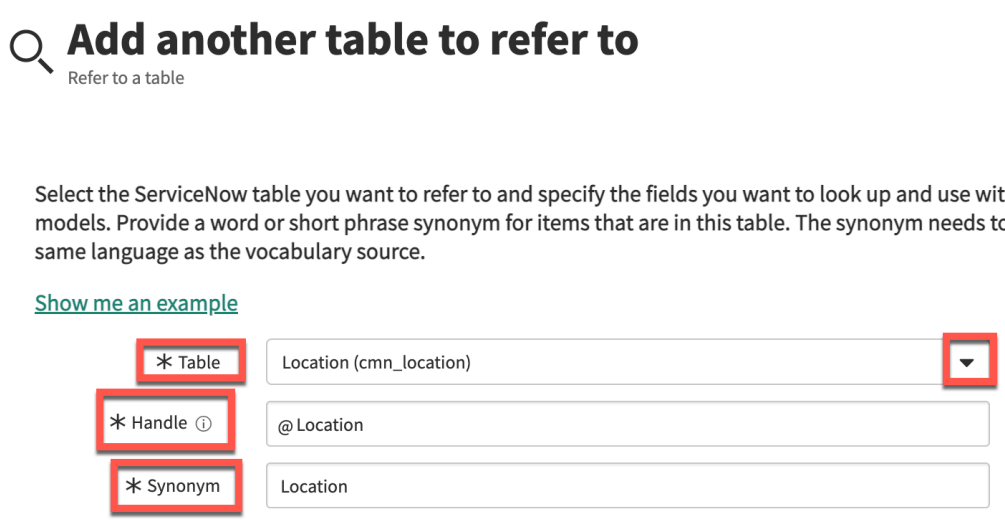
1. Navigate to **All** > **NLU Workbench** > **Vocabulary Sources**.
2. Click **ServiceNow Tables**.
3. Click **Add another table**.
4. On the Add another table to refer to page, configure the fields.

   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.



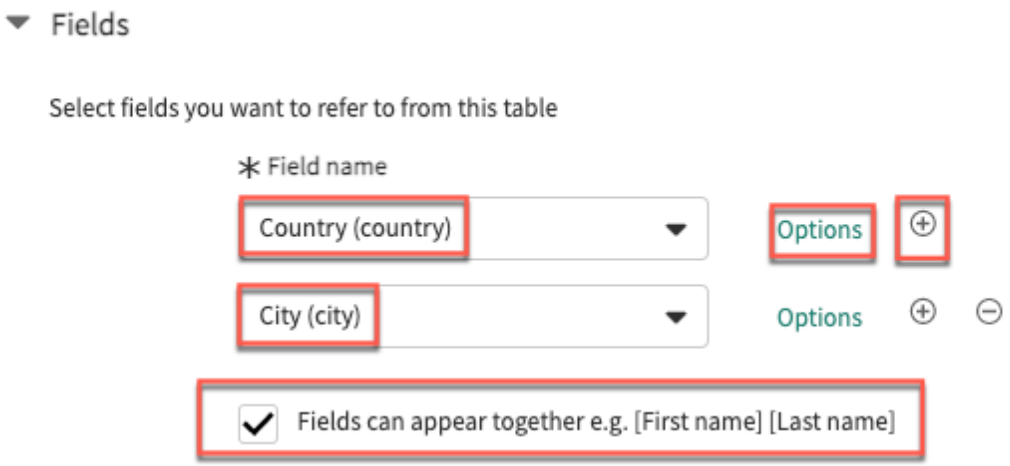
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

b. Select 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-separate values in this field, for example, if your city column has multiple names, such as NYC, New York, and New York City.
- **Fields can appear together**: Leave the check box empty. This field is typically used in an NLU Search model, where users can enter words next to each other in an utterance to find a record. For example, you can enter `Pierre Development` while searching for an employee named Pierre who works in the Development team, where **Name** and **Team** are two fields in the table.



c. Configure the advanced options.

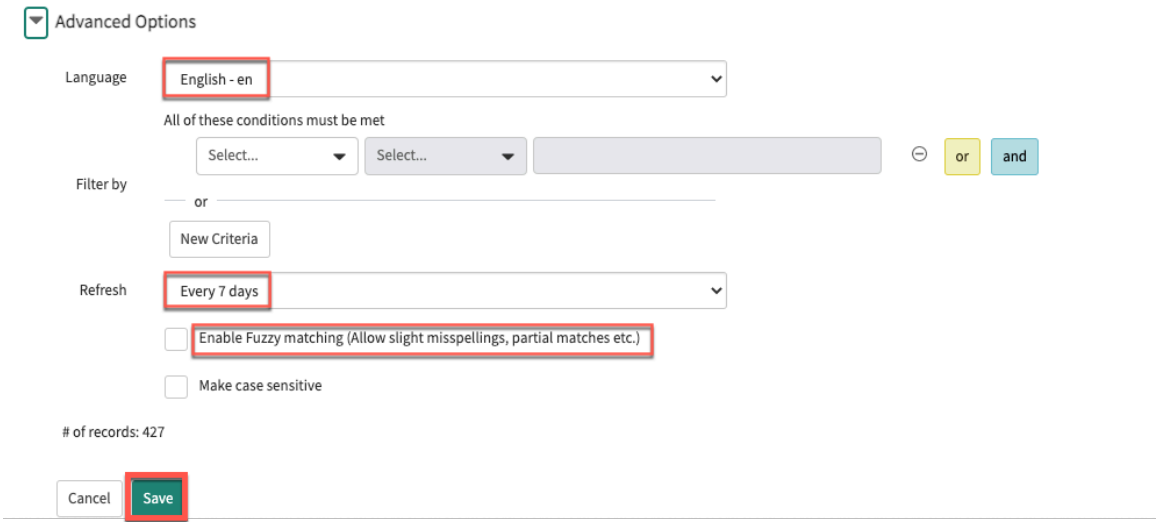
In this example, use the following configurations:

- **Language**: Select **English - en**.
- **Filter by**: You can use the condition builder to filter the values in the source table for the vocabulary source. Leave empty for this scenario.
- **Refresh**: Select **Every 7 days**. By selecting this, you set the system to get new values from the table every seven days. There are different refresh options you can choose based on how often the source table is updated.
- **Enable Fuzzy matching**: Checking this box make it so a record can be matched even if an utterance contains a slightly misspelled word or part of a word. For example, when searching for a city such as Kansas City, you might enter Kansas City or just Kansas, and the system will still be able to match with the correct location record. Fuzzy matching may

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

return false matches in some cases. Use fuzzy matching sparingly and test out your model before using it.

- **Make case sensitive**: If you don't enable fuzzy matching, you can check this box to make the values in the source table case sensitive. If an utterance doesn't use the correct cases, the model doesn't predict the intent.



5. Click **Save**.

   Result: Your **@Location** table vocabulary source appears on the Vocabulary sources screen, where it begins to sync with its source table.

6. If the sync doesn't start immediately, select the **Sync Lookup** icon on the far right to sync the source table data with the NLU Service.
   Result: When the sync completes, you can use the vocabulary source in your models. You don't need to retrain the model after updating a table vocabulary source.

What to do next

Add a vocabulary source to a model by annotating an utterance. Use the @ symbol when adding an example utterance to an intent.

Sync a table vocabulary source

Synchronize your table vocabulary sources to obtain the latest changes to the ServiceNow source table. Synchronizing your vocabulary sources ensures your NLU models have the latest values when predicting intents.

Before you begin

- Make sure that the NLU Workbench plugin, NLU Workbench - Core plugin, NLU Common Model plugin, and Predictive Intelligence plugin are all installed and activated on your instance.
- Role required: admin or nlu_admin

About this task

When you reference a vocabulary source in an utterance, it pulls the values at the time the model is trained. However, if the table values change over time, the model still references the values from the last training session.

You can select a schedule to automatically refresh the vocabulary values used by NLU. This schedule can be edited later. For more information, see Create a table vocabulary source.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

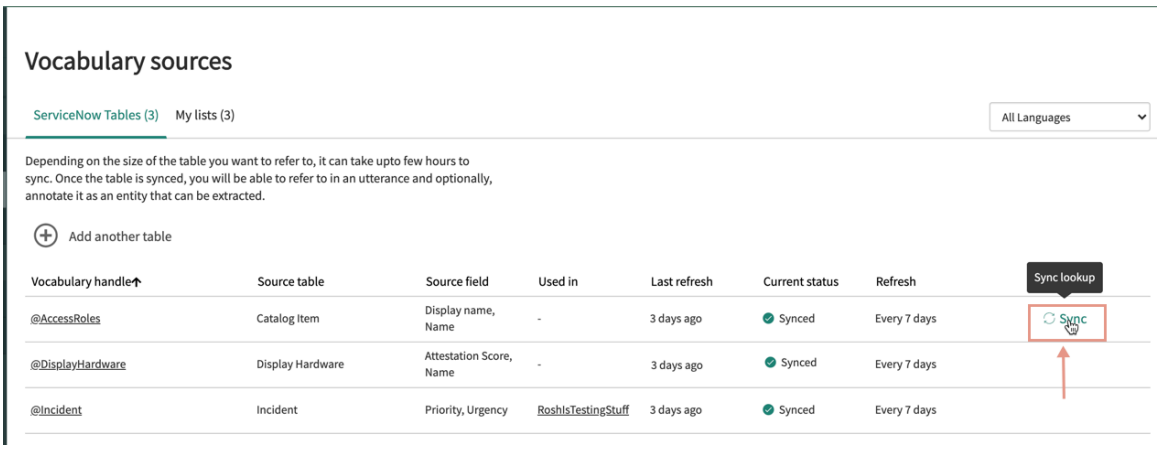
You can also manually sync a table vocabulary source, such as before model training.

Note:
If an utterance references a table vocabulary source that has never been synchronized, the model fails to train. Check the vocabulary source's current status, and sync manually if **Never synced**.

In the following example scenario, you're manually syncing the @AccessRoles vocabulary source to the Catalog Item table.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Vocabulary Sources**.
2. In the ServiceNow Tables tab, point just to the right of the **Refresh** column to invoke the **Sync lookup** icon.

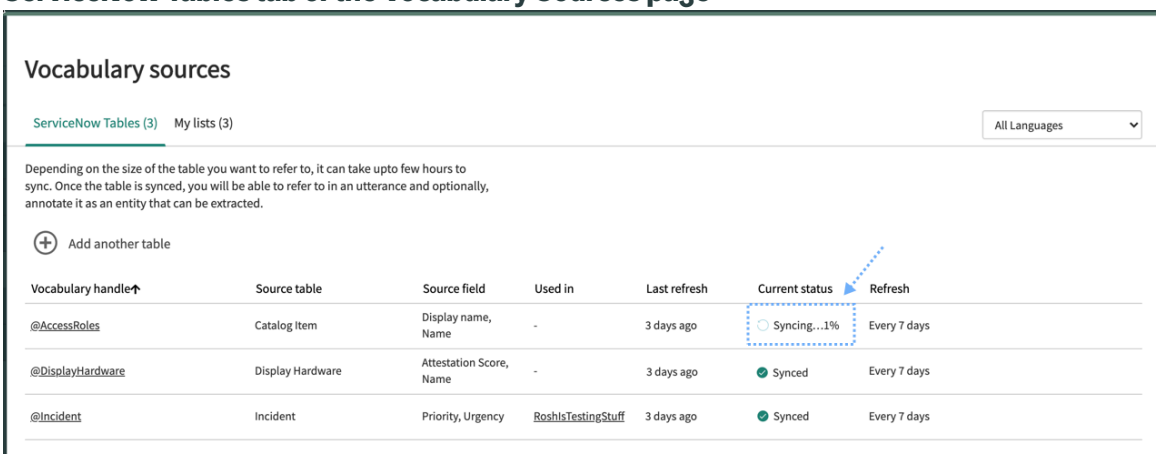


3. Select **Sync**.
   The vocabulary source starts synchronization.

What to do next

The syncing operation can take some time, depending on the source table size.

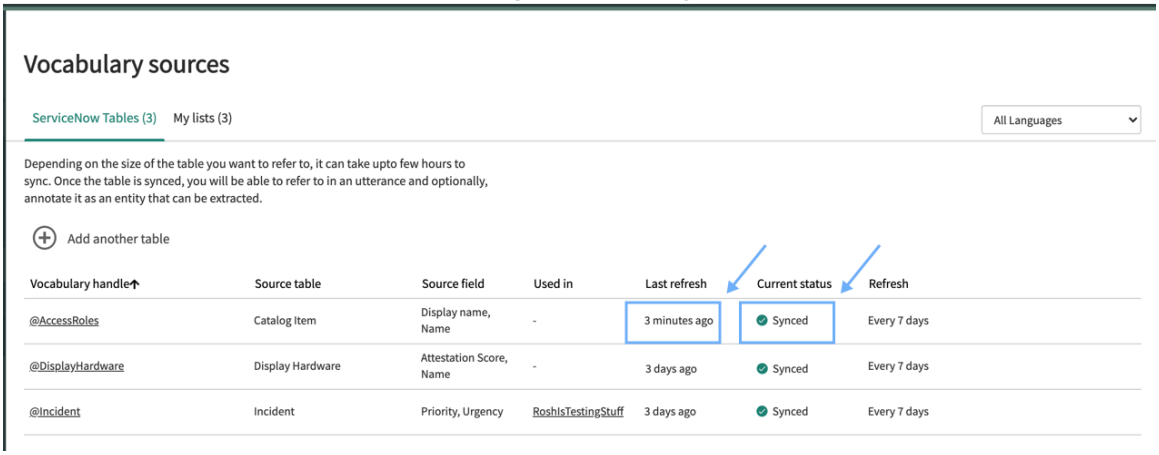
ServiceNow Tables tab of the Vocabulary Sources page



The values in the Last refresh and Current status columns reflect the current status of the vocabulary source.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ServiceNow Tables tab of the Vocabulary Sources page



Proceed to train your model. For more information on model training, see Train and try your NLU model.

For information about errors when syncing a table vocabulary source, see article KB1588239 in the Now Support Knowledge Base.

Related topics

Create a table vocabulary source

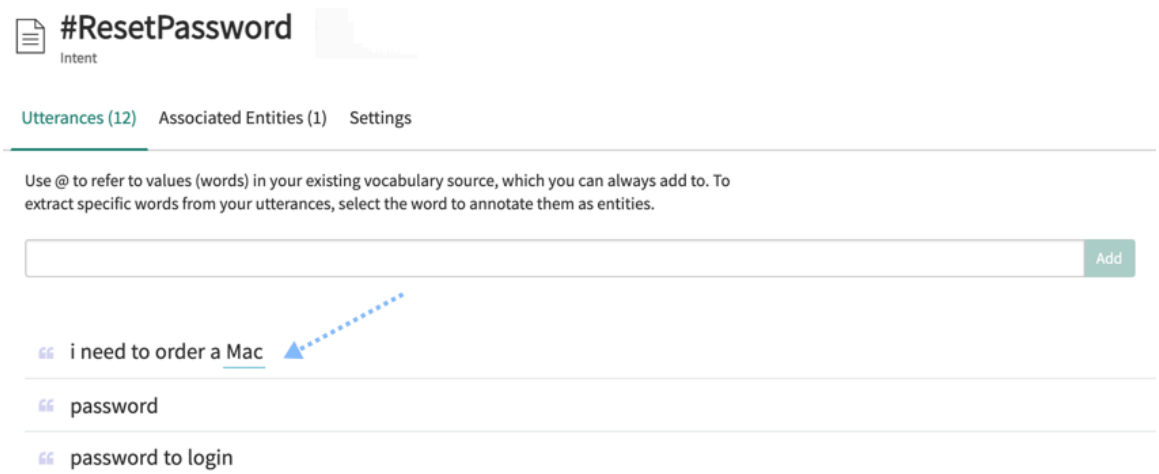
Pre-built vocabulary

Use ServiceNow® pre-built vocabulary for software and hardware terms so the system recognizes their multiple variations in utterances.

Your Natural Language Understanding models contain pre-built vocabulary settings you can use when adding an example utterance. The vocabulary provides definitions for software and hardware terms whether they're expressed in slang or professional usage. The vocabulary can also recognize product misspellings.

For example, for one of your example utterances, you enter `I need to order a Mac`. When the system recognizes a pre-built vocabulary item, the term has a blue line under it.

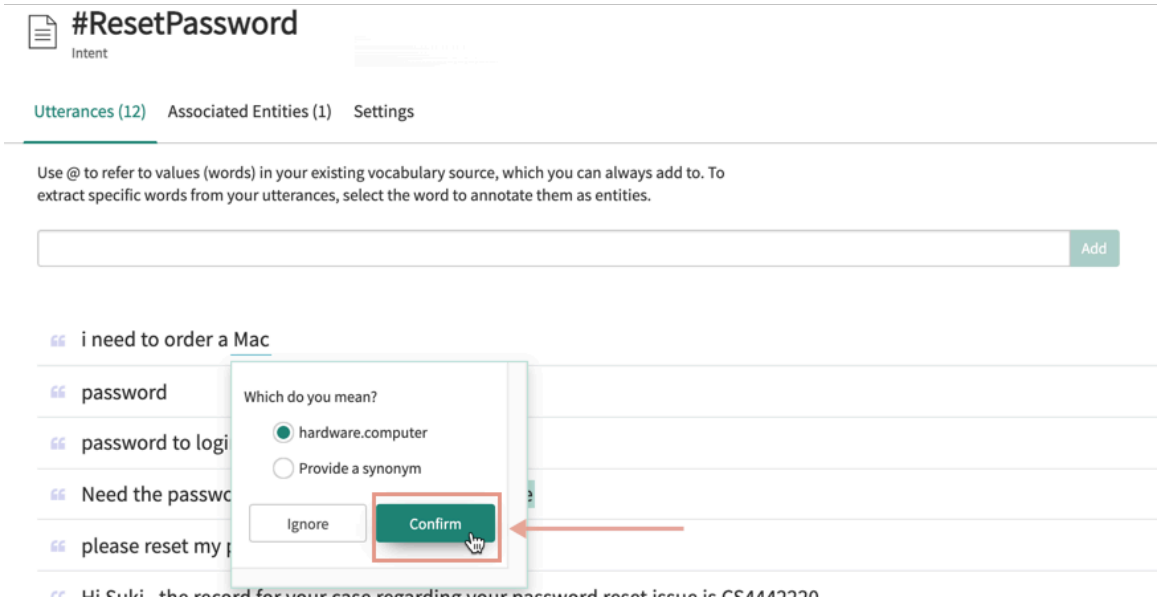
Utterance tab of the Intent details page



When you click the word, a window appears with two options to choose for the word:

- A pre-built suggested definition for the word
- An option to add a synonym

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



If you select the first option and click **Confirm**, the system uses the pre-built suggested definition and the blue line disappears.

If you select the second option, enter a synonym for the word, and click **Confirm**, the word and the synonym are added to the model vocabulary.

If you select either of the two options and click **Ignore**, the blue line disappears and the word remains as it was previously.

Test set creation and management

Use the default test set of your NLU model to test the model's performance and accuracy. Manage your test set over time by building or updating its content in the NLU Workbench.

Embedded video: Test set management in NLU Workbench

Note:
To test your model, install the ServiceNow® Store application NLU Workbench - Advanced Features. For more information, see Install NLU Workbench - Advanced Features.

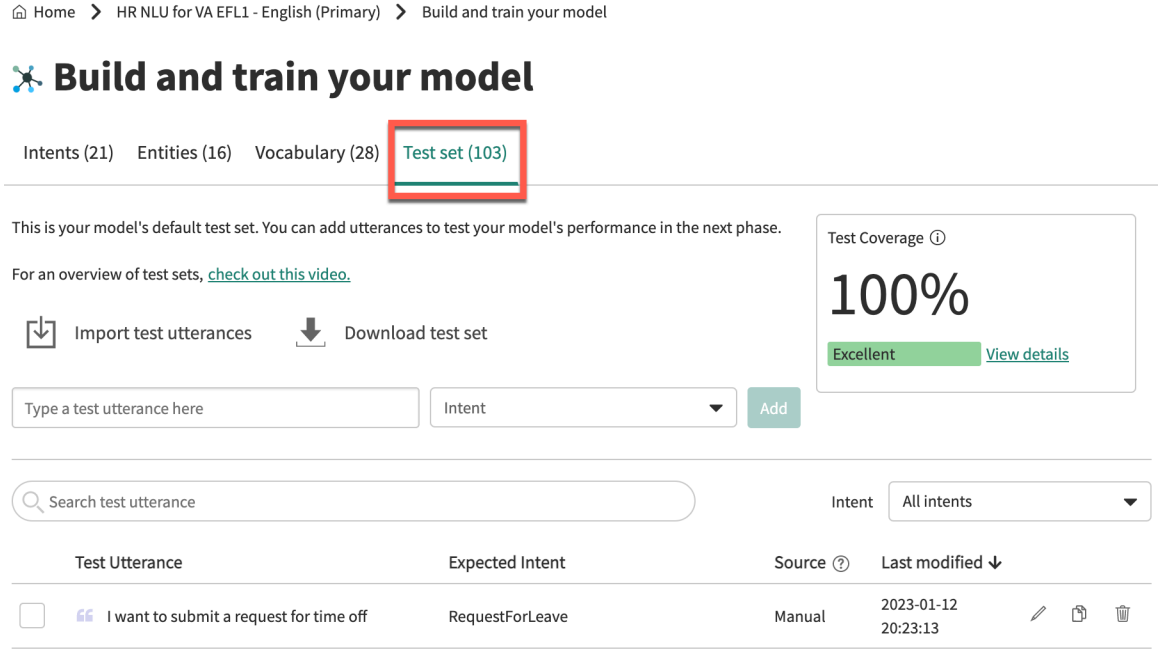
When you create an NLU model for Virtual Agent or AI Search, a default test set is created and associated to the model. You can use the default test set to evaluate the model's performance. Initially the test set is empty, ready to be populated with your content.

Access your default test set

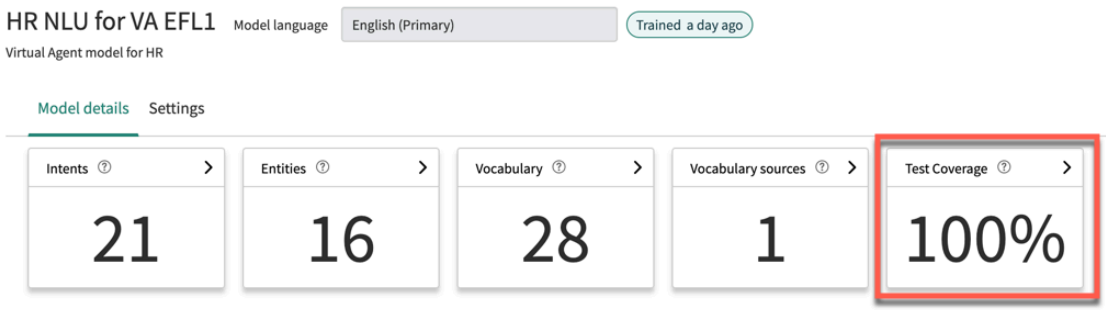
Access your default test set with one of the following methods.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Navigate to **All** > **NLU Workbench** > **Models**. Select the tab for your model's application, then the name of your model from the list. On the model's overview page, find the **Build and Train your model** card and select its **View phase** button. Then select the **Test set** tab.



- Navigate to **All** > **NLU Workbench** > **Models**. Select the tab for your model's application, then the name of your model from the list. On the model's overview page, select the **Test Coverage** tile.



- Navigate to **All** > **Multi-model Batch Testing** > **Test sets** tab. Find the name of your model. Default test sets are labeled as **Default**.



Add content to your default test set

Add utterances and their expected intents to build and manage your test set over time. You can add content to the default test set with the following methods:

- Add test utterances and their expected intents manually. From the model's overview page navigate to **Build and train your model** > **Test set** tab. Type your input into the Type a test utterance here field, select an appropriate intent, then select the **Add** button.

  These test utterances are assigned a source of Manual.

- Import test utterances and their expected intents from a CSV file or from other models. To import content to a default test set, from the model's overview page navigate to **Build and train your model** > **Test set** tab. Select **Import test utterances**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Imported test utterances are assigned a source of Manual.

- The Expert Feedback Loop feature lets you add actual user utterances from Virtual Agent chat logs to the test set.

  These test utterances are assigned a source of Expert Feedback. For more information, see NLU Expert Feedback Loop.

Test Coverage

The Test Coverage score is the percentage of a model's enabled intents that have test utterances in the default test set. Before testing your model, ensure that there is at least 60% coverage. The higher the Test Coverage score, the more accurate the performance testing results.

Your test coverage needs to be at least 60%, with at least 5 test utterances per intent, in order for the system to provide an optimal confidence threshold during batch testing. For more information about the confidence threshold, see NLU model settings.

Aim to have about 10 percent of a model's test utterances marked as "not relevant", meaning that there is no intent associated. This helps assess how the model handles irrelevant utterances which should not have any intent predicted. For more information about irrelevant utterances, see Irrelevance detection in NLU.

Use the test set

To use the default test set from the **Test and publish your model** phase, see Test and publish your model.

To use the test set in Multi-model Batch Testing, see Multi-model Batch Testing.

Characteristics of default test sets

When an instance is upgraded, default test sets are created for any existing models that don't already have them.

When you copy a model using **Duplicate this model**, the original's default test set is copied into the new model. For more information, see Duplicate an NLU model.

The utterances in the test set shouldn't be the same as the utterances in the training set.

Default test sets can't be deleted separately from their models.

Test set utterances should be in the same language as their model.

Test sets are available for Virtual Agent or AI Search models.

Downloading or moving default test sets

Default test sets can be downloaded or moved as follows.

- Default test sets can be separately downloaded in CSV format. To download the test set, from the model's overview page navigate to **Build and train your model** > **Test set** tab. Select **Download test set**.

  Note:
  Test sets that are downloaded from **Download test set** contain test utterances and their expected intents, but not the sources.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Default test sets can be moved with update sets. When you add an NLU model to an update set, its default test set is added, including test utterances, expected intents, and sources. For more information, see Add an NLU model to an update set.
- When using the **Export model as CSV** function in the **All existing models** table, the default test set is not included. For more information, see Export an NLU model.

Train and try your NLU model

Train and try your model iteratively so that its intents and entities are validated, compiled, and saved to your model.

Before you begin

- Make sure that the NLU Workbench - Core plugin, NLU Workbench plugin, and Predictive Intelligence plugin are all installed and activated.
- Create an NLU model. For more information, see Creating models.
- Create one or more NLU intents and their associated entities for your model. For more information, see NLU intents.
- If any utterance references a table vocabulary source, ensure that the source has been synchronized so that its values are available to your model. For more information see Sync a table vocabulary source.
- Role required: nlu_editor, nlu_admin, or admin. The NLU editor must be assigned to the model.

About this task

Training your model saves any changes you made to the content, and checks for conflicts or errors. Training also makes a model available for publishing.

After training, you can try your model by manually entering individual utterances to see what intents are predicted.

Note:
To run a test of your model against a list of test utterances, see Test and publish your model.

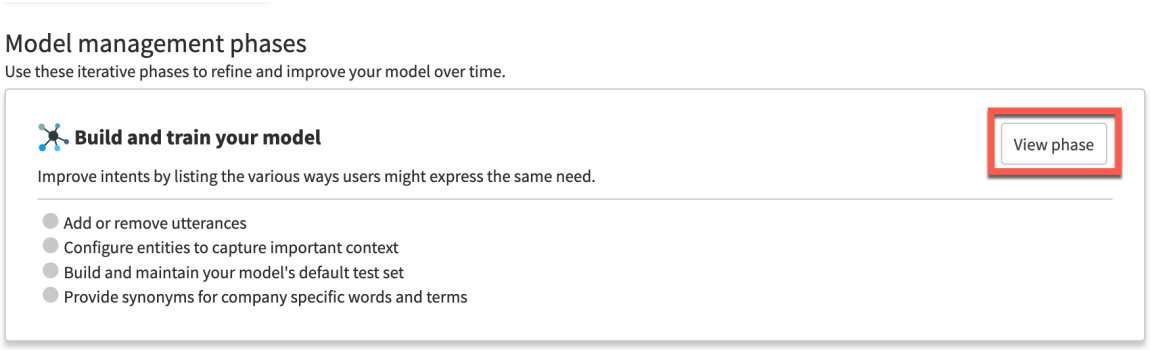
The mid-conversation responses of Dialog Acts can't be tried or tested in NLU Workbench.

In this example scenario, you've already built sufficient model content by adding intents, utterances, entities, and their associated annotations. Following the example procedure, you first train your NLU model. Then you try your model by manually entering utterances so you can check the prediction results and confidence scores.

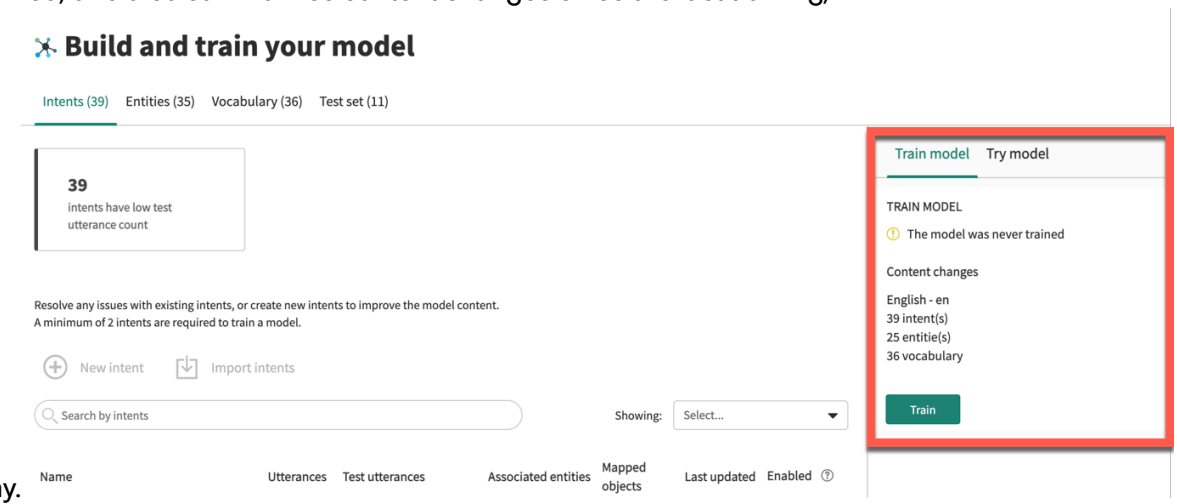
Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab for your model's application, then select the name of your model.
3. In the **Model details** tab of the model overview, make sure there is enough content in Intents, Entities, and Vocabulary.
4. On the **Build and train your model** card on the model overview, select **View phase**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



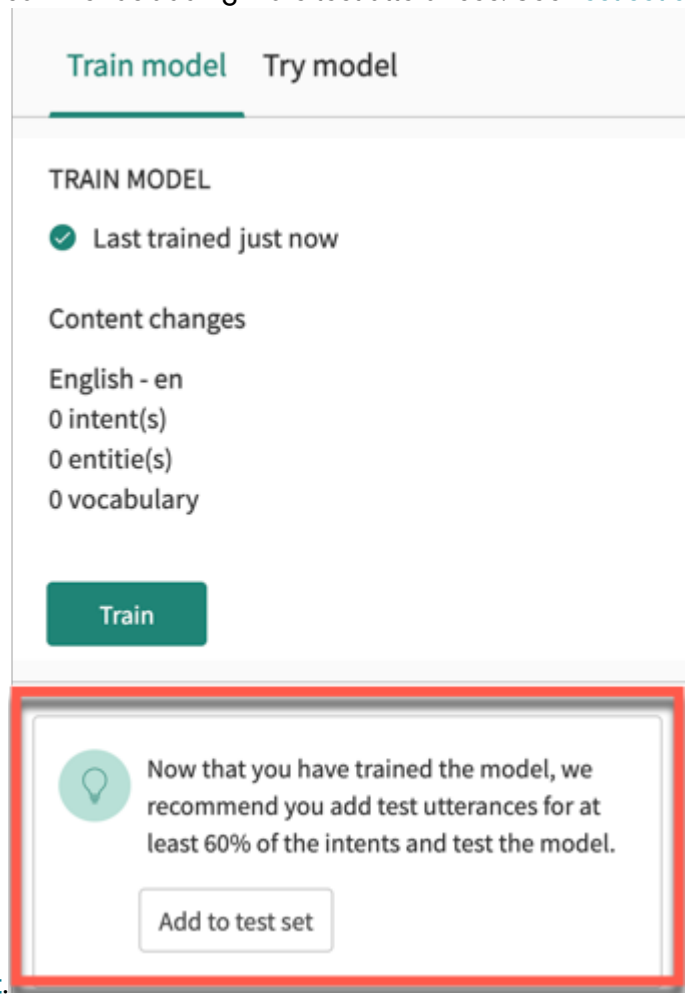
5. When the **Build and train your model** phase opens, ensure that the **Train model** tab is selected.
   Result: The **Train model** tab displays the last time the model was trained, and also summarizes content changes since the last training, if any.



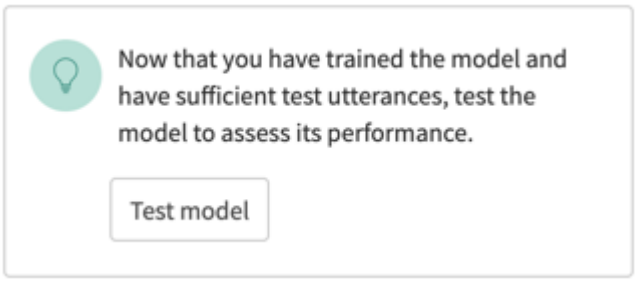
6. Select the **Train** button.
   Result: The system displays a progress bar during training. When finished, the system displays one of two recommendations:

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- When less than 60% of the model's intents are covered in the default test set, the system recommends adding more test utterances. See Test set creation and management.



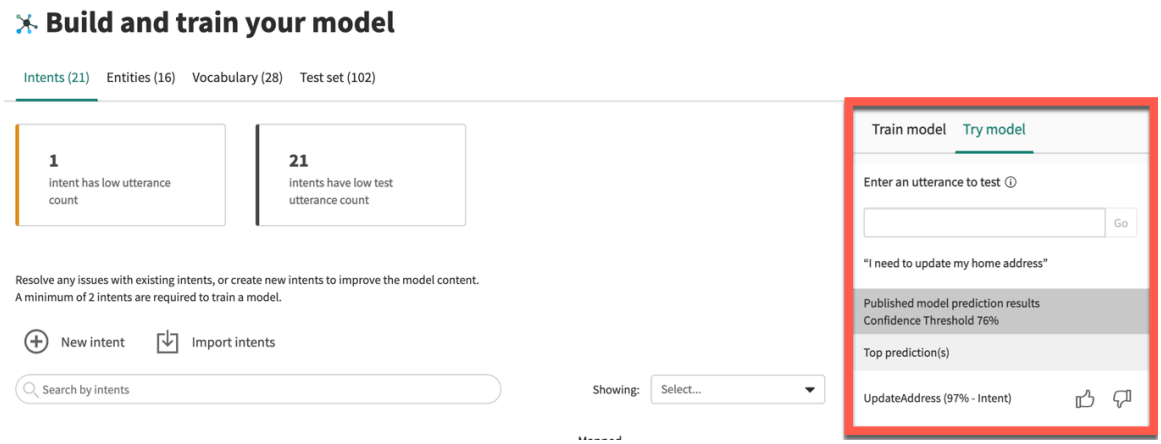
- When over 60% of the model's intents are covered in the default test set, the system recommends proceeding to testing. See Test and publish your model.



7. To manually try individual utterances, select the **Try model** tab.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. In the text field under **Enter an utterance to test**, type an utterance and select **Go**.



Result

In this example, you entered `I need to update my home address` as the utterance to try.

1. The system displays the model's confidence threshold, which is 76% in this example.
2. Under **Top prediction(s)**, the system displays all intents that were predicted with a confidence score greater than the threshold.
3. In the example, the intent UpdateAddress is predicted with a confidence score of 97%, which is greater than the threshold of 76%.

The **Try model** results also display thumbs-up and thumbs-down icons for you to provide feedback. For more information, see Test panel feedback.

What to do next

- Continue trying various utterances to check that your updates to model content are effective. See Compare draft and published versions of your NLU model.
- To test your model against a list of test utterances, use its default test set in the Test and publish your model phase, or navigate to Multi-model Batch Testing.
- To adjust the model's confidence threshold, use the **Settings** tab on the model's overview page. For more information, see NLU model settings.
- If you're satisfied with the results of your testing, Publish your NLU model.

Test panel feedback

When testing your NLU model on the Try model section of the test panel, use this feature to provide feedback on the model's intent predictions.

Summary context

When a model is trained and tested for an utterance and the model returns an intent prediction, you can provide a thumbs up or thumbs down rating on the predicted intent it returns. Marking a different intent prediction as correct adds the utterance to the corrected intent. All other feedback is captured for continual learning. The system then incorporates your feedback to optimize the model predictions. This feature requires the nlu_admin role to access and test the model. NLU editors can also access the test panel if an NLU admin assigns them to it.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Providing prediction feedback

The ratings you provide help the system to match an intent to an utterance. These ratings are essential for the system to continuously learn, evolve, and improve the accuracy of the intent predictions based on user input. They also enable you to notify the system if the intent prediction is correct or not.

The following scenarios below show examples of how to interact with your model test panel and provide prediction feedback to the system. In all scenarios, you use these four steps:

1. In the **Build and train your model** phase of your model, select **Try model** to open the test panel.
2. In the test panel's **Enter an utterance to test** field, enter a brief utterance that's similar to a training utterance in one of the intents.
3. Click **GO**.

   Result: The system returns its predictions for your test utterance in the **Top Predictions(s)** section of the test panel.

4. Click the **Thumbs Up** icon or the **Thumbs Down** icon.

   If you want the system to know it has predicted the correct intent for your utterance, select the **Thumbs Up** icon.

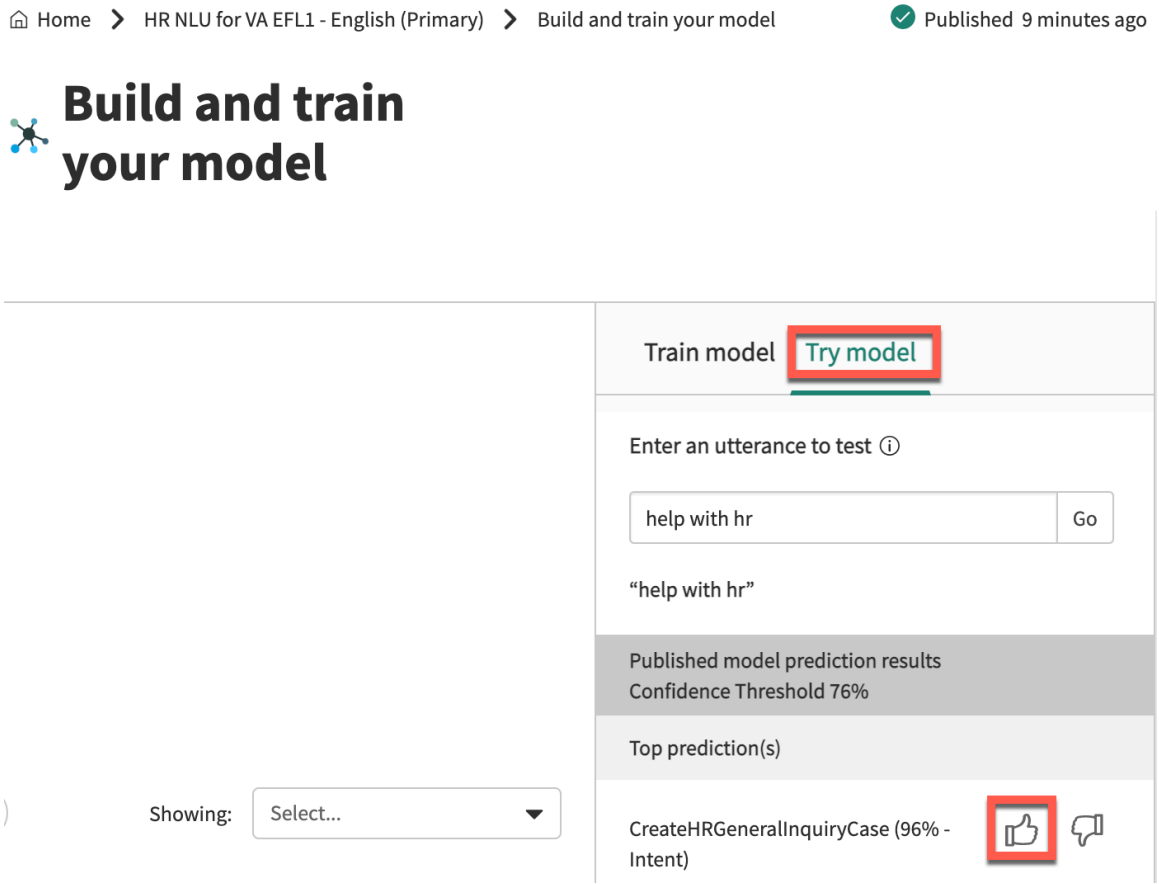
   In all other cases, select the **Thumbs Down** icon, which opens the **Provide feedback to improve this prediction** section. Here you can choose an intent other than the top predicted intent.

Scenario 1: On the Try model section of the test panel, you enter `help with hr` as the utterance. When the top prediction results appear, you're confident that the predicted intent is the correct match to your utterance. So in this case, you click the **Thumbs Up** icon.

Results:

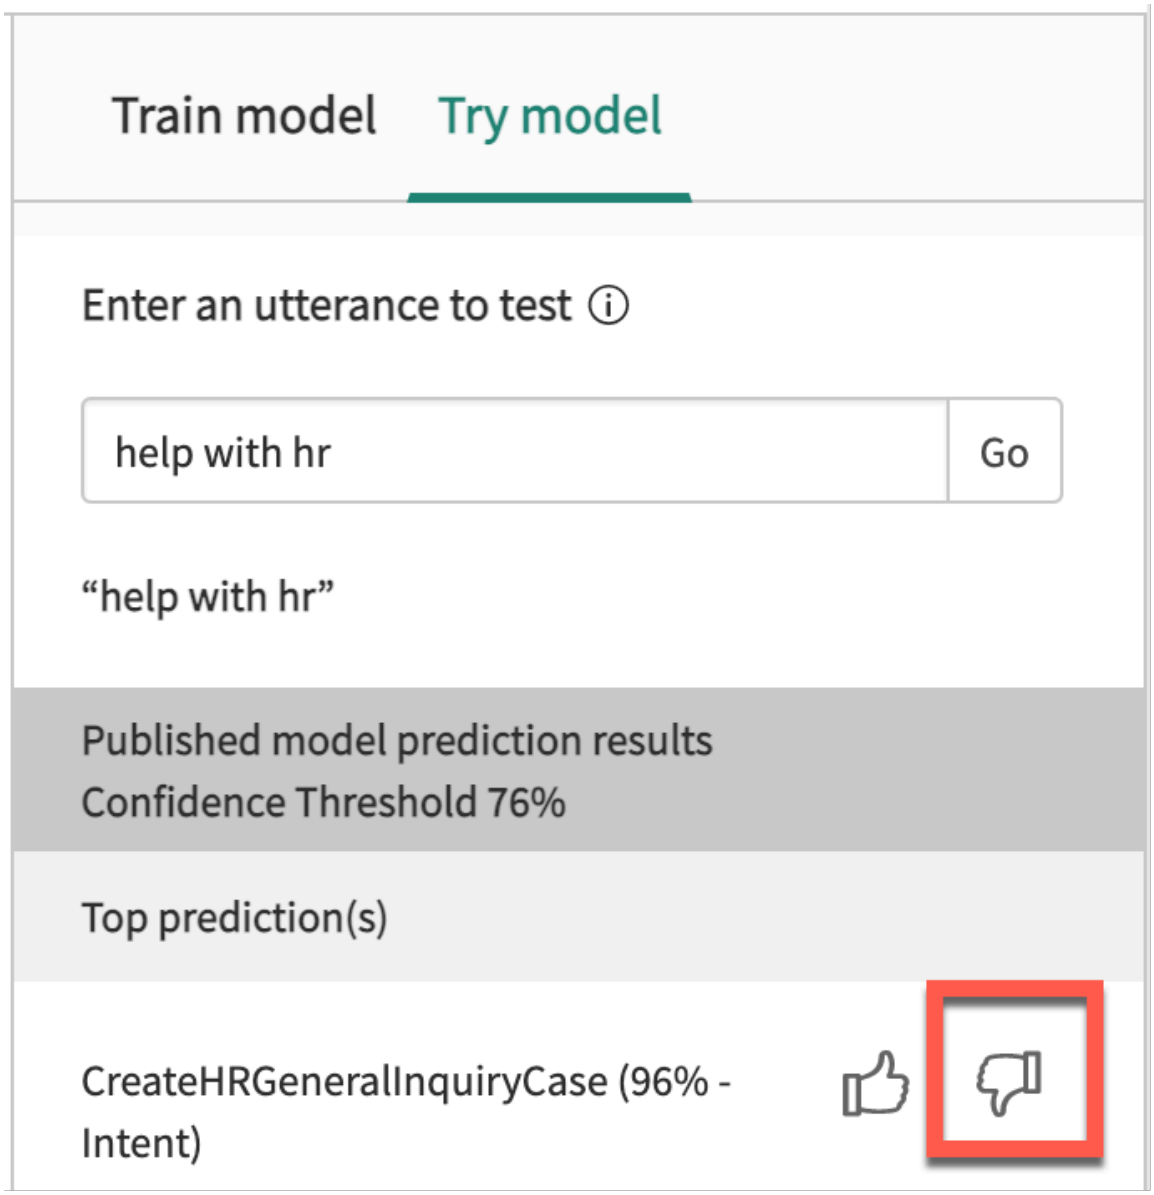
- The system predicted the correct intent, which in this case is **#CreateHRGeneralInquiryCase**.
- Your feedback notifies the system that it has matched the correct intent to your test utterance.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Scenario 2: In a separate model on a separate instance, a different user enters the same `help with hr` utterance. The system responds with the top prediction results for the intent, but the user isn't sure if it's the correct intent or not. So this user clicks the **Thumbs Down** icon, as shown in the image below.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

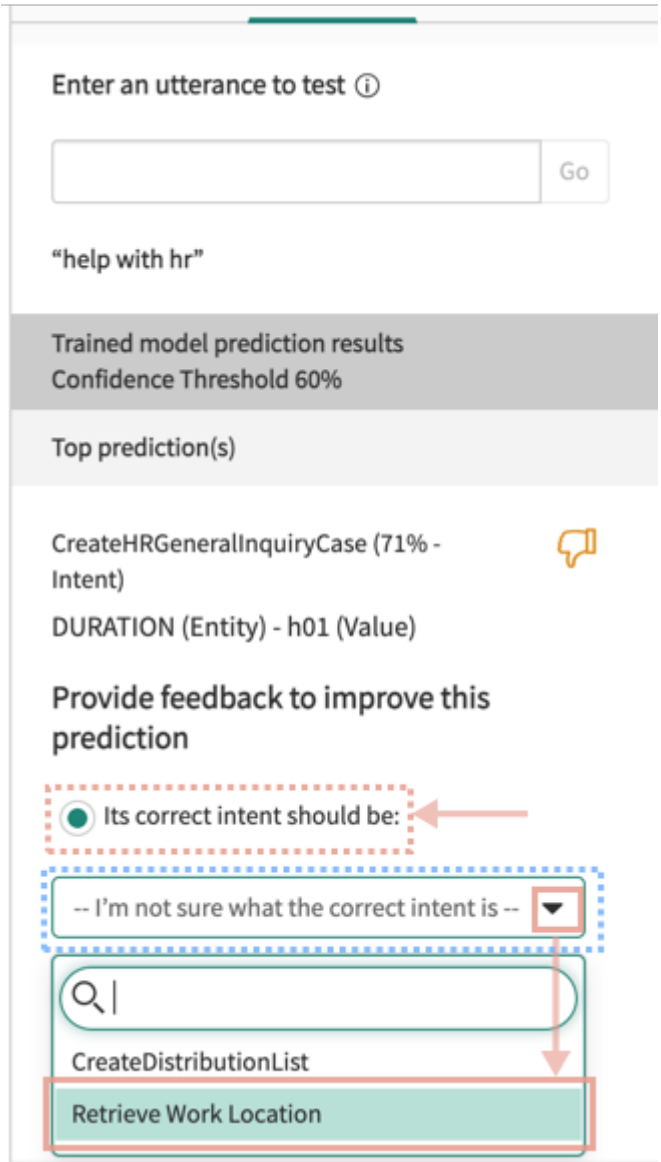


Result: The panel expands to show the **Provide feedback to improve this prediction** section where users can submit feedback that may help to improve the intent prediction.

There are two options here:

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

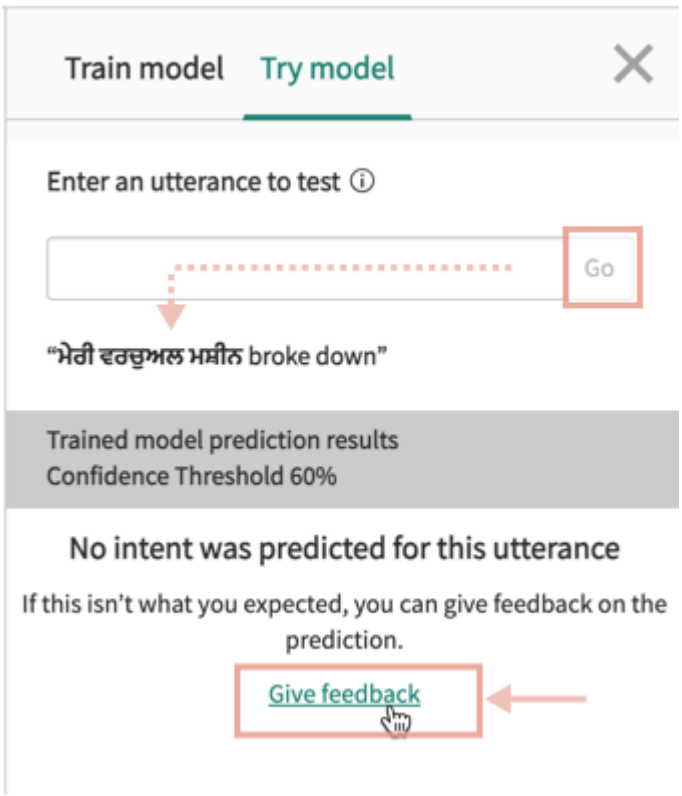
- If users click the **Its correct intent should be:** button, a list appears where they can choose a more appropriate intent for the test utterance. In this example scenario, a user selects the **Retrieve Work Location** intent, as shown in the image below.



- If you click the **I'm not sure what the correct intent is** prompt, instead of returning a top prediction, the system shows the next best intent predictions available.

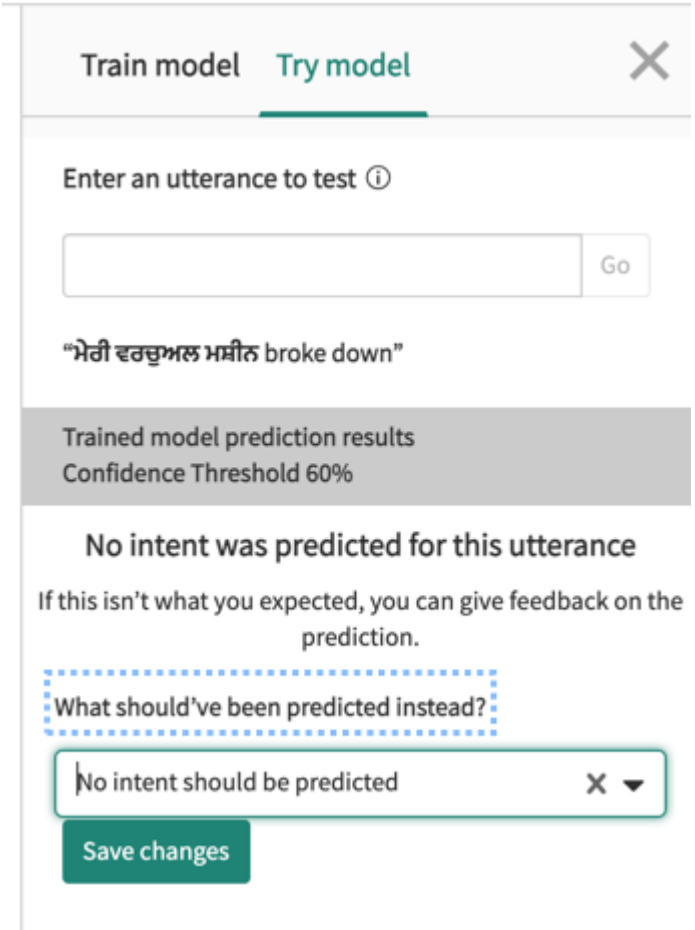
Scenario 3: In a separate model on a separate instance, another user submits an utterance that uses gibberish, or uses a language that's different from the language the model uses. For example, a user mistakenly submits an utterance comprised of both non-English and English languages, as shown in the image below.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



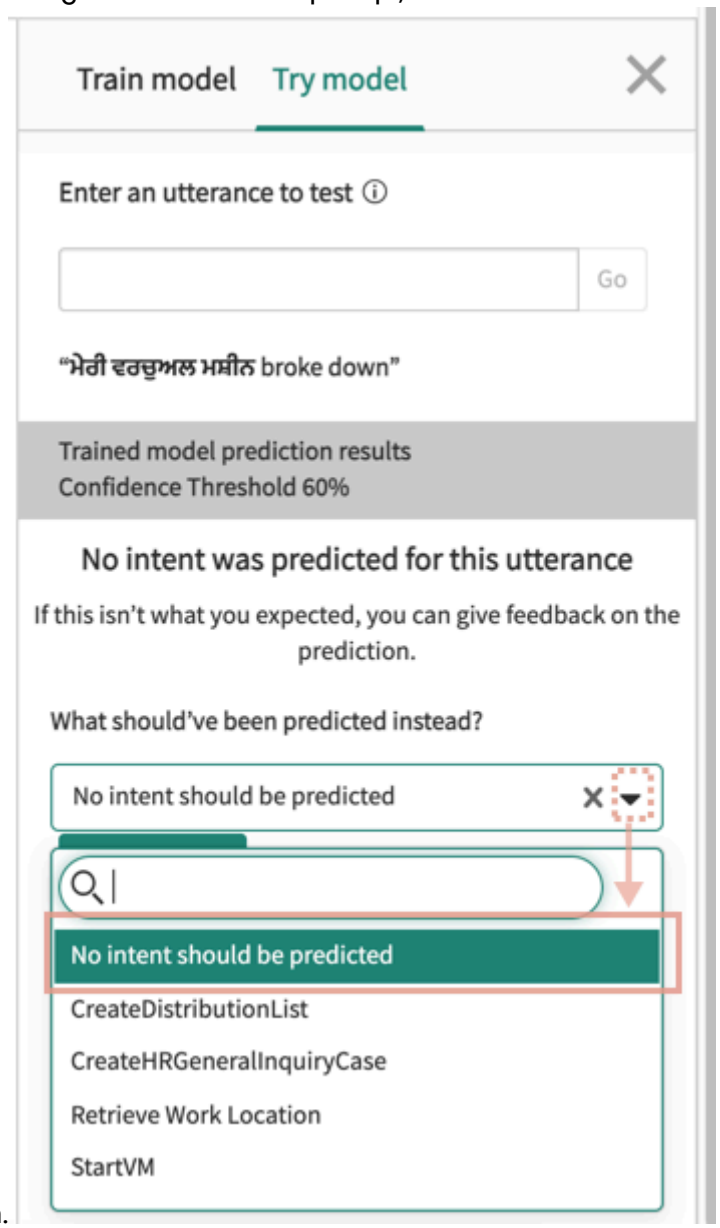
Result: The system doesn't return a prediction because the utterance uses two different languages together. Since no intent was predicted, the user clicks the **Give feedback** option which expands the Try model section to show other intent alternatives.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

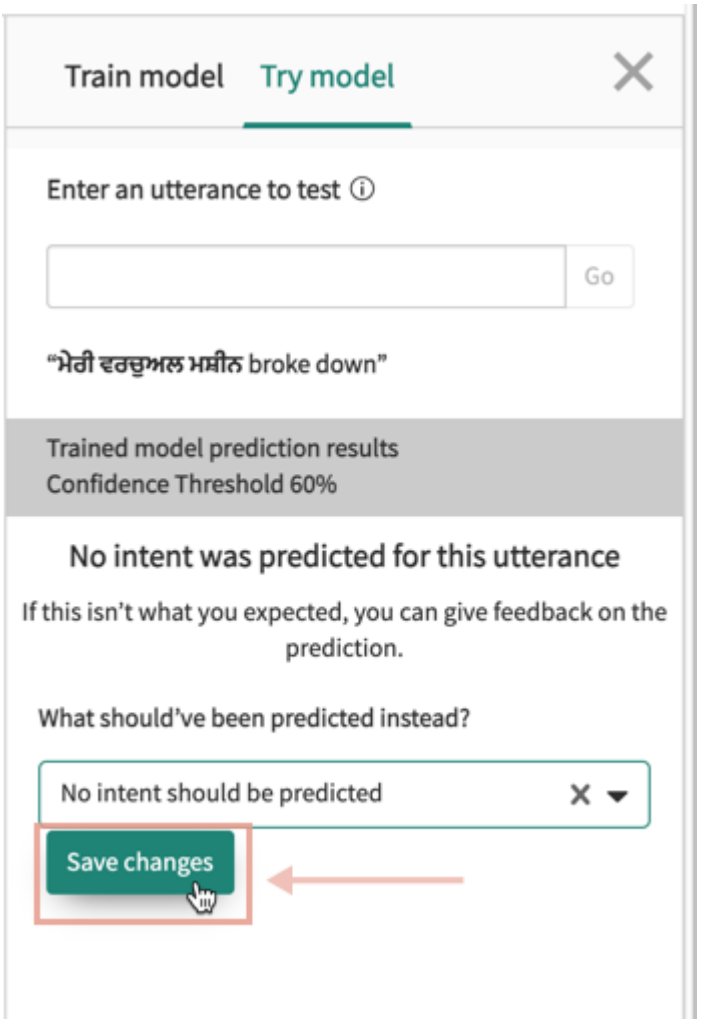
So instead of choosing an intent from the prompt, this user selects the **No intent should be predicted** option.



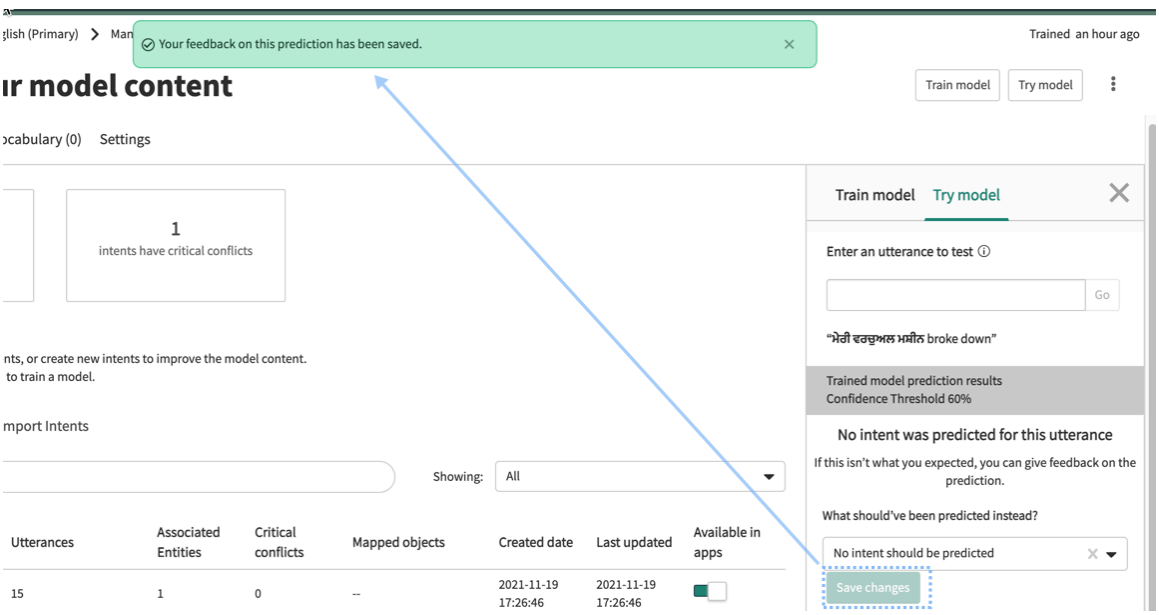
Note:
When you choose and save **No intent should be predicted**, the utterance is removed from all intents which it is a part of.

Scenario 4: Along with choosing from a list of your model's intents for a prediction, you can also directly notify the system that the utterance is irrelevant to the model. To do this, you click the **Exclude this model's predictions for this utterance** button, then click **Save changes**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Result: A banner appears at the top of the screen confirming the user feedback for the prediction is saved, as shown in the image below.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Accessing your feedback records

Your feedback data is stored in the **ml_labeled_data** table, which is also used by other ServiceNow products. This table can also house multiple sources, such as Virtual Agent chat logs that can be used for future predictions.

Test and publish your model

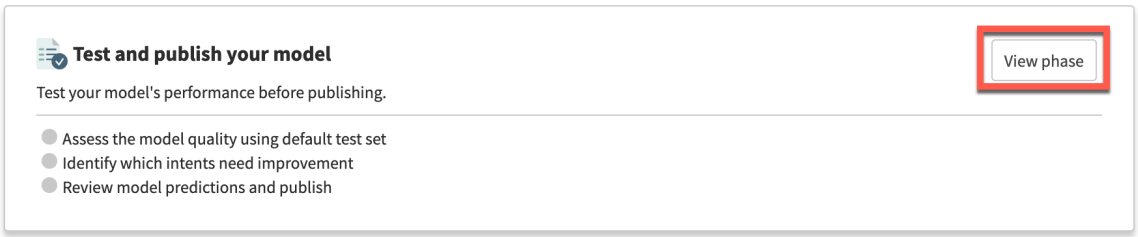
Assess the performance of your NLU model to identify areas for improvement. Then publish your model to make it available to other applications such as Virtual Agent.

Summary usage

Test your Virtual Agent or AI Search model against its default test set to see how the model responds. Test results provide information you can use to improve your model.

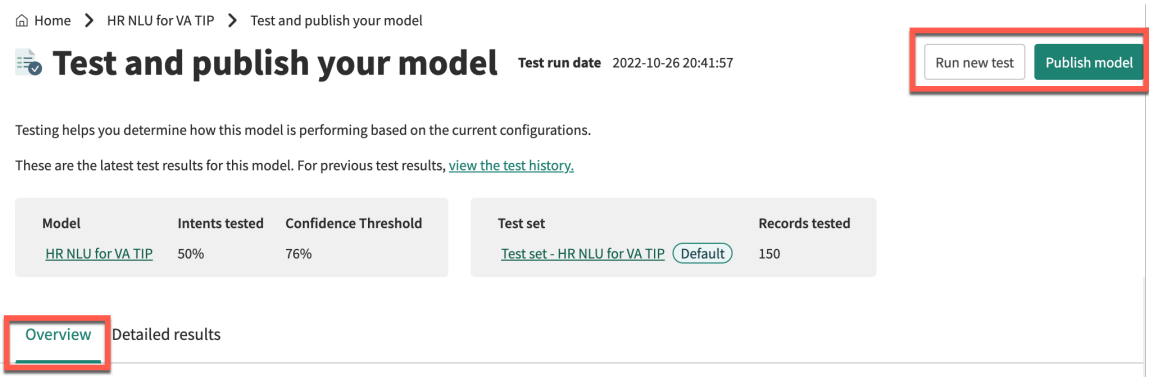
Note:
Testing your model requires the Multi-model Batch Testing feature, available with the NLU Workbench - Advanced Features application from ServiceNow® Store. For more information, see Install NLU Workbench - Advanced Features.

To test your model, navigate to **NLU Workbench** > **Models**. Select the tab for your model's application, then select the name of the model. In the **Test and publish your model** card, select **View phase.**



Overview

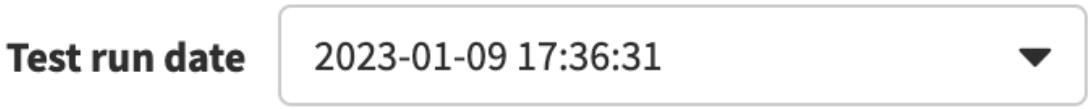
The **Test and publish your model** phase opens in the **Overview** page by default. Buttons for **Run new test** and **Publish model** are located here.



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.

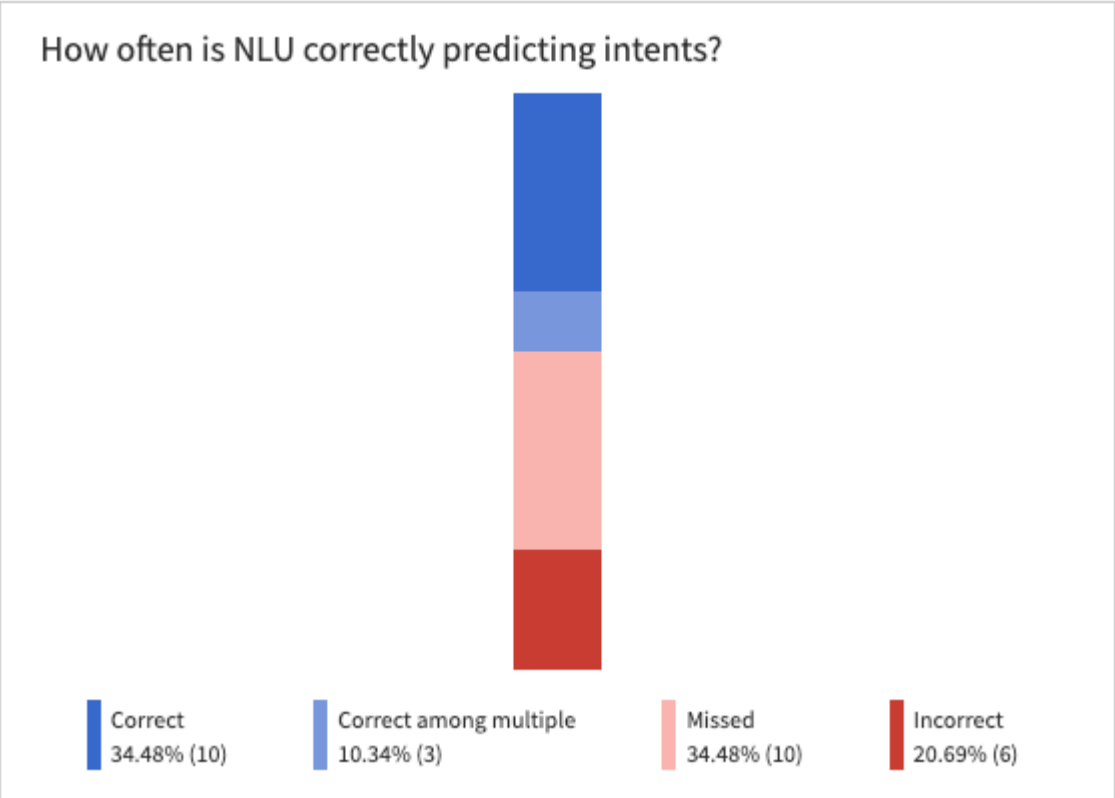
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



To drill down into the test results table, select the **Detailed results** tab. Each test utterance is listed in **Detailed results**, with its prediction.

Understanding test results

The test results show how your model responded to the utterances in the test set.



The bar chart shows the prediction percentages for correct, correct among multiple, missed, and incorrect:

Percentage	Description
Correct	The percentage of utterances for which your model correctly predicted the intent. When the model predicts no intent for utterances marked as `Not relevant`, that result is counted as Correct.
Correct among multiple	For utterances that had more than one intent predicted. The percentage of utterances for which the model correctly predicted the intent or intents, but also predicted intents that did not belong to the utterance.
Missed	The percentage of utterances for which your model did not predict an intent, even though there was an expected intent.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Percentage	Description
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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

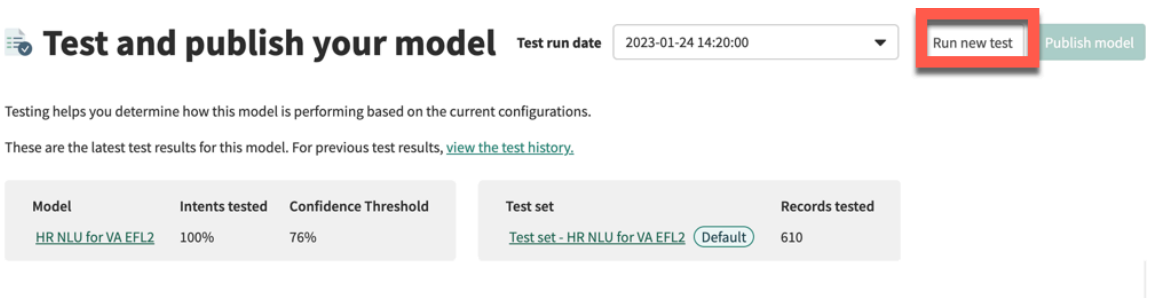
- If an expected intent in your test set does not correspond to any intent in the model, the utterances with those intents are not used for testing. They are not included in the test results.
- The mid-conversation responses of Dialog Acts can't be tried or tested in NLU Workbench.
- When the model returns 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.

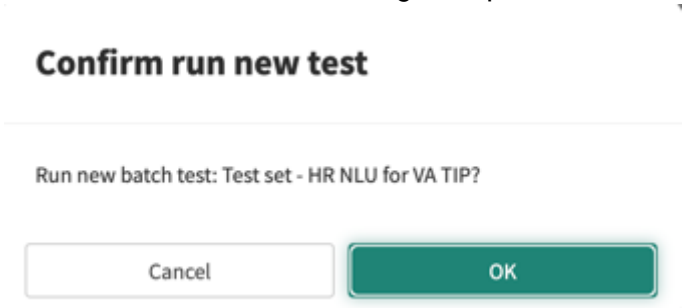
Procedure

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can 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

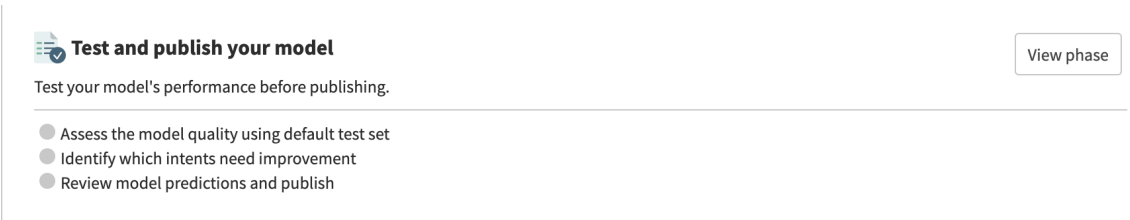
For this procedure, you've already trained, tested, and refined your NLU model and are satisfied with the results.

If your model is already published, you can publish it again. However, you must train the model again before republishing.

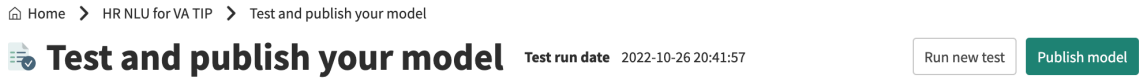
Procedure

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**.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:
If 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

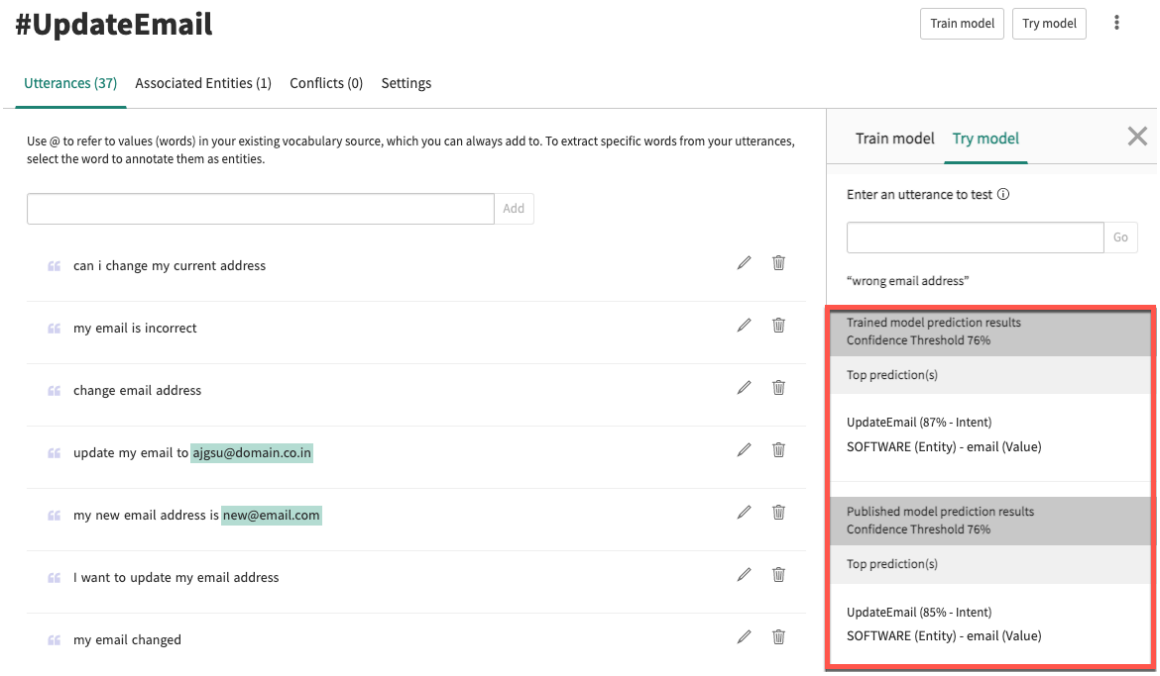
1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default.
2. Select the tab corresponding to your model's application, then select the name of your published model.
3. From the model's overview page, locate the **Build and train your model** card and click its **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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. 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.



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`.

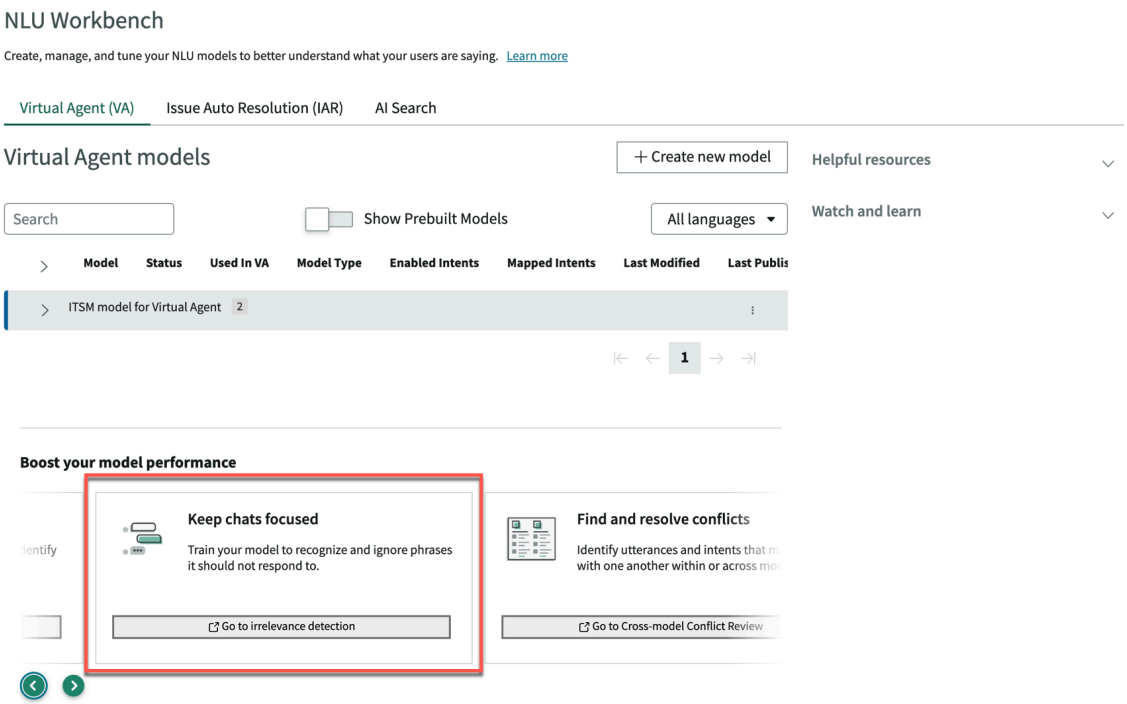
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Roles, 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

These 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**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note 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.

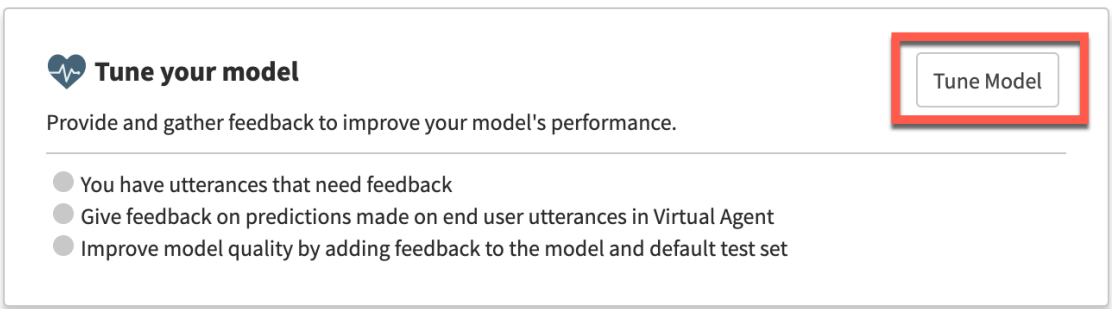
More information

- 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® Store application NLU Workbench - Advanced Features must be installed to view **Tune your model**. The Virtual Agent plugin is also required.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

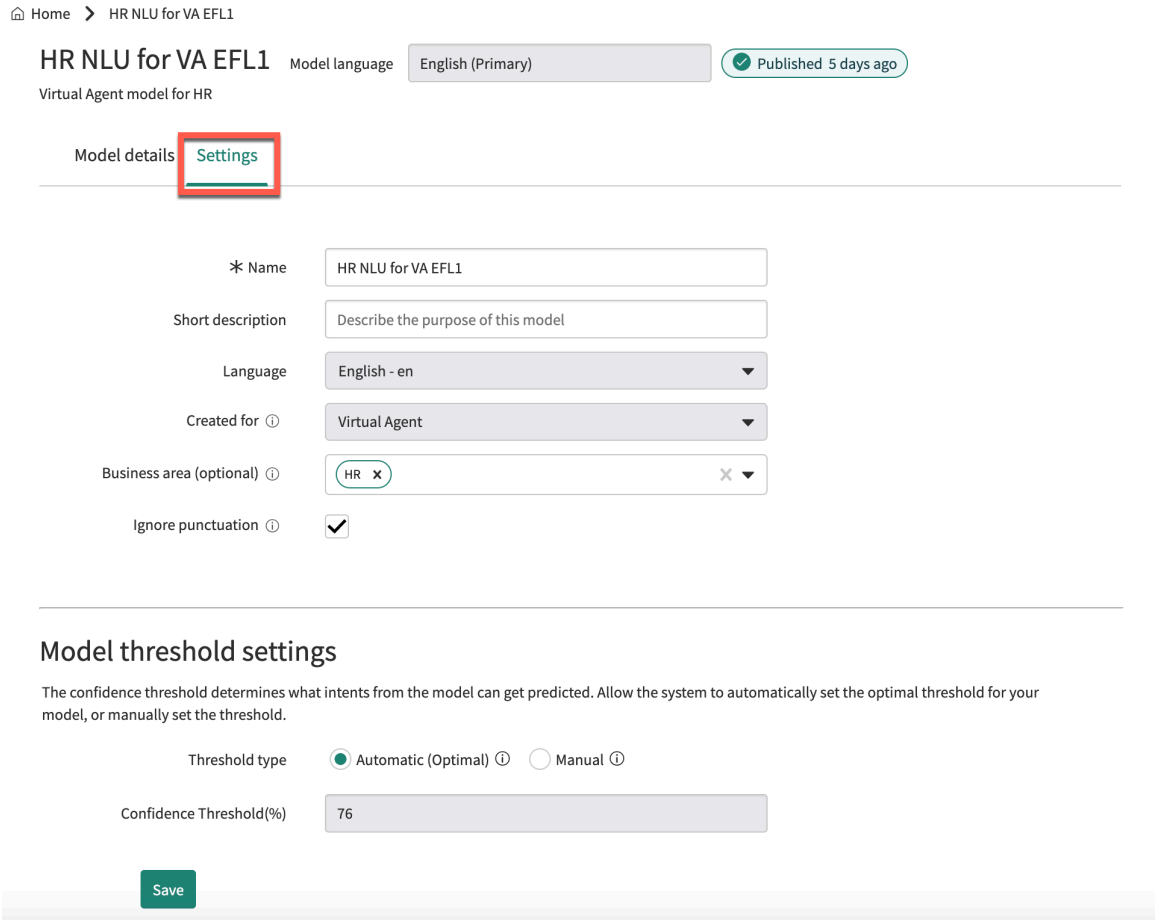
When the Expert Feedback Loop has feedback waiting for your review, **You have utterances that need feedback** is highlighted on the card.

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.



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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, if the 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.

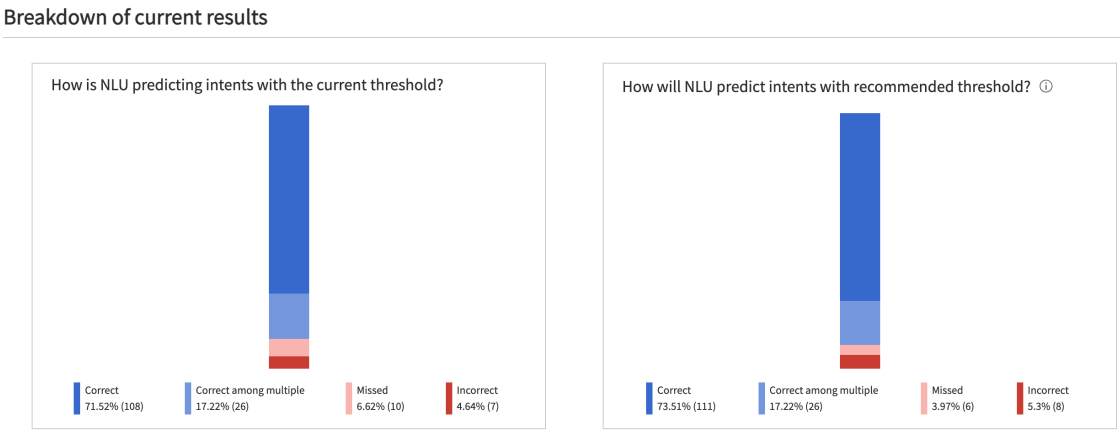
There are two types of model threshold settings:

- **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.



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

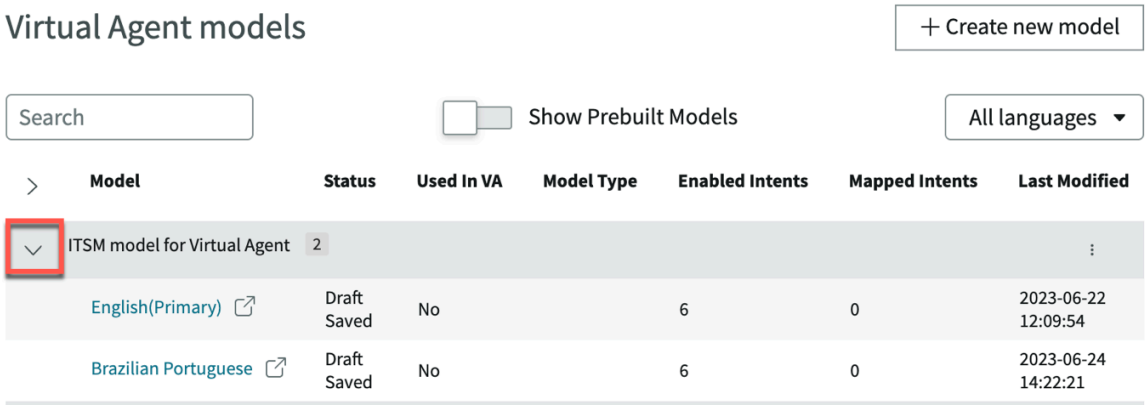
Use multilingual Natural Language Understanding (NLU) models for the system to understand user input in several languages. The NLU Workbench helps you manage and maintain a consistent structure for content across languages to provide a unified experience.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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.



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.

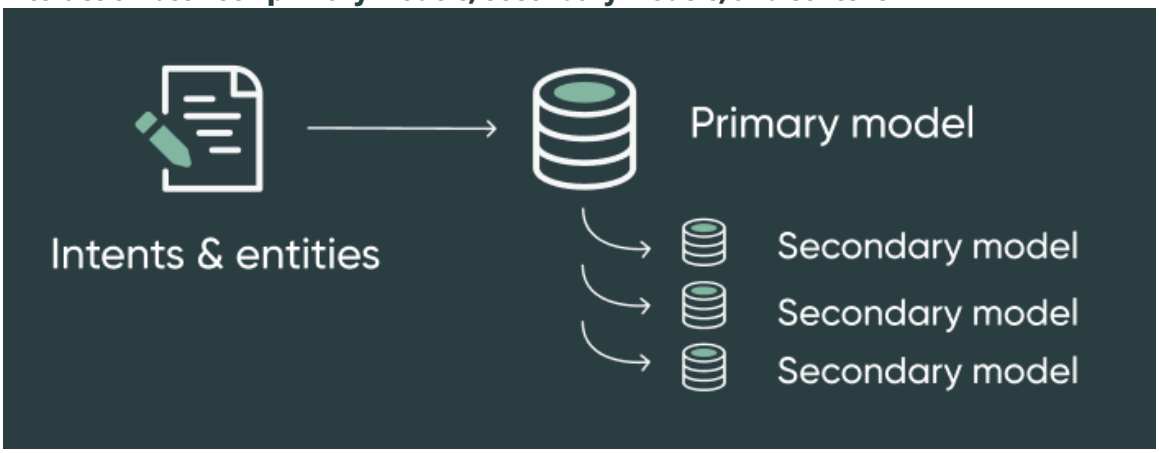
Primary and secondary model interactions

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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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



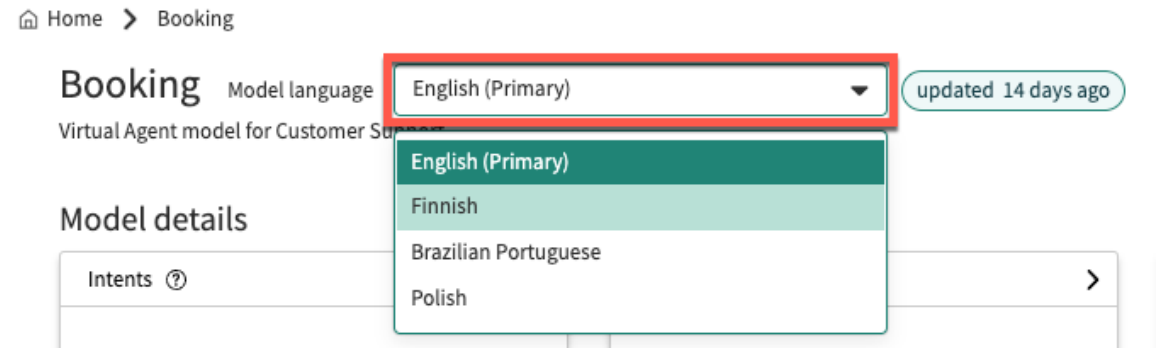
If every intent in a secondary model is disabled, then the **Train** and **Try** buttons are also disabled in the model. However, even if only one intent is enabled in the model, then you can train and test the model. Ensure that some intents are enabled in your secondary models to translate them.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When 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



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.

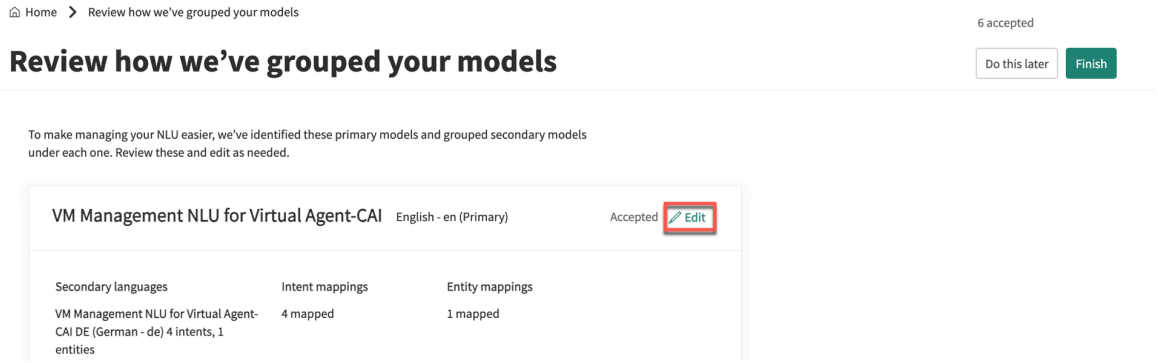
If you have added a new language to a model, a language group is created automatically, so you may already have existing language groups. For more information on adding and translating a new language model, see Translate a multilingual model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

To view current language groups and inspect possible new language groups, navigate to:

`<instance_name>.service-now.com/$nlu-studio.do#/model-migration`.

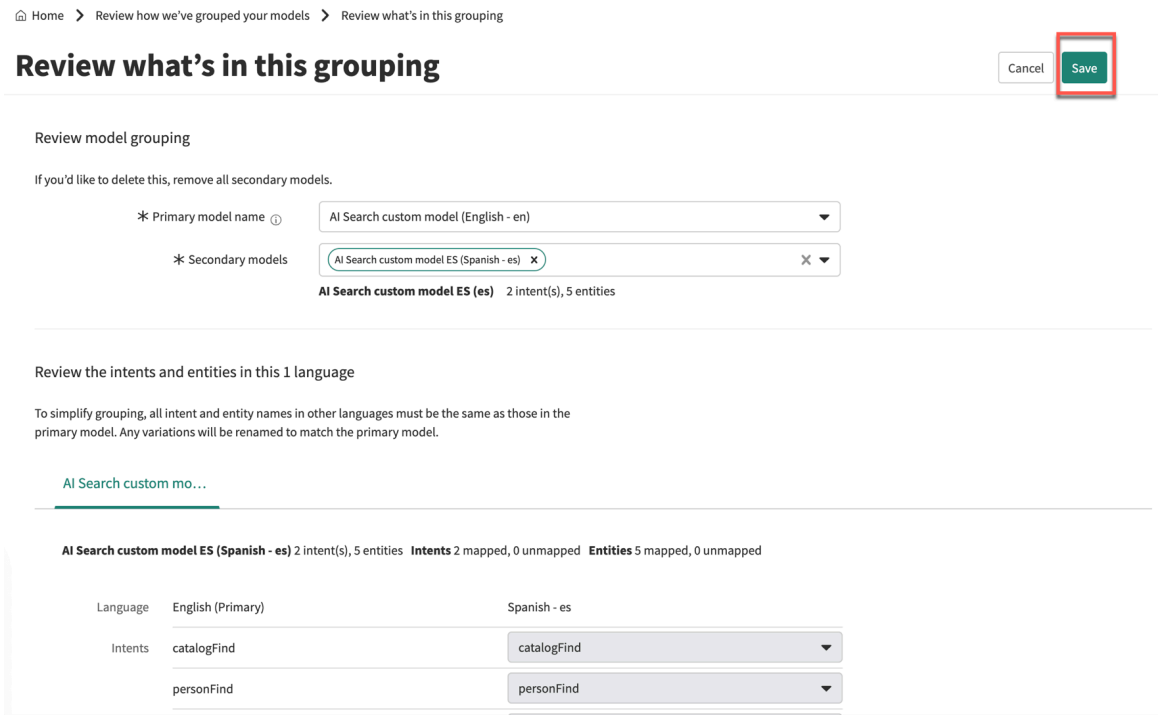
The page titled **Review how we've grouped your models** opens.



Reviewing your grouped models

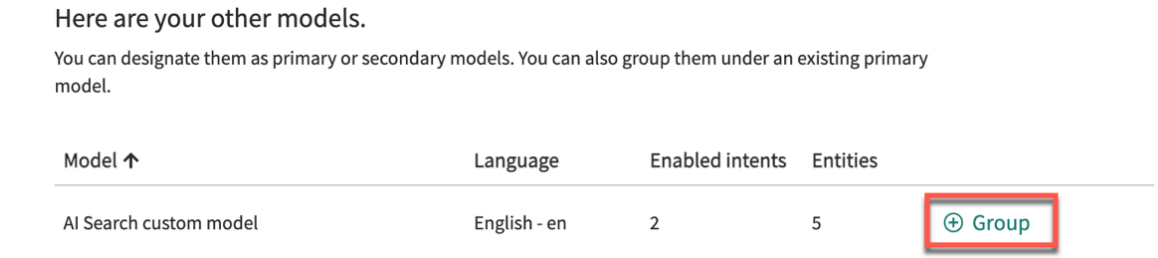
There are two things you can do on the **Review how we've grouped your model** screen when using the nlu_admin role.

- 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.

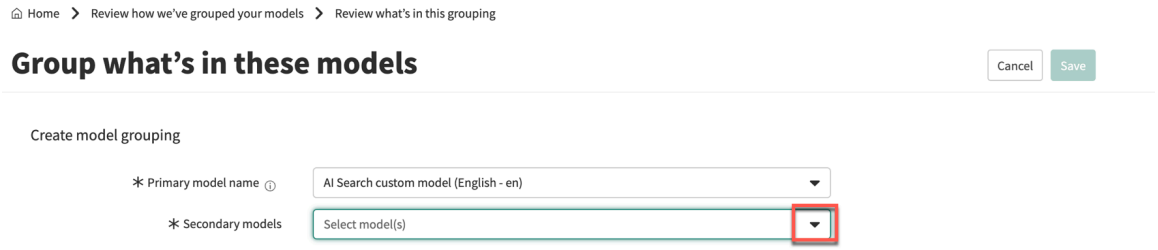


© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- You can 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.



The **Group** function opens the page titled **Group what's in these models**. You can expand the list of possible Secondary models, and then select the **Save** button. Any new groups are displayed in the upper section of **Review how we've grouped your models**, after refreshing the page.



Note:
The time it takes the system to group your models depends on the number and size of your models.

Translate a multilingual model

Add a language to an existing NLU model by translating it. Use one of several translation options to add a secondary model in a supported language.

Before you begin

- Activate the Localization Framework (com.glide.localization_framework.installer) plugin. See Localization Framework.
- In your target application scope, create a primary language model, or use an existing primary model.
- Multilingualism is available for Virtual Agent and AI Search models.
- Role required: nlu_admin or admin.

Note:
The nlu_editor role cannot initiate a model translation, because the nlu_editor does not have permission to create a new model. See Assign an NLU editor to a model.

About this task

Translating a model creates a secondary model of the language you choose, in the same application scope as the primary model.

The secondary model contains the same content as the primary. However, the intents of the secondary model are disabled and must be reviewed prior to publishing the secondary model.

To translate a primary model and create a secondary model, you must choose a translation mode. The modes appear in the user interface as individual cards that you use in Step 4 of this

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

procedure. 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
Use AI	Translate using machine-intelligence providers such as Google. Selecting the card creates a translated secondary model. Machine translations may be incorrect or too literal, so review the translated content for accuracy. For more information, see Dynamic Translation. To activate the Use AI mode, Localization Framework must be configured and activated.
Use a third party	Request translation by a professional translation team. When you select this mode, an LF task is created per the LF language group settings for the secondary language. When the translations are completed, the translated model is created and ready for review. For more information, see Localization Framework and Translation modes.
No translation	Default. Manually translate your primary language into a secondary language. Selecting this mode creates a draft secondary model with the intent names the same as that of the primary model. These intents are disabled by default, and you must enter your utterances manually. Names of intents and names of entities imported from the primary model aren't editable. Names of entities that you create directly in the secondary model are editable. This mode is also helpful if you want to draft and save the primary model first, and then translate your utterances later.

Multilingual models are supported in 17 languages.

Note:
Some of these languages support only intents, not entities. If the language is intent-only, your primary model entities will not be translated or copied. For more information, see NLU language support.

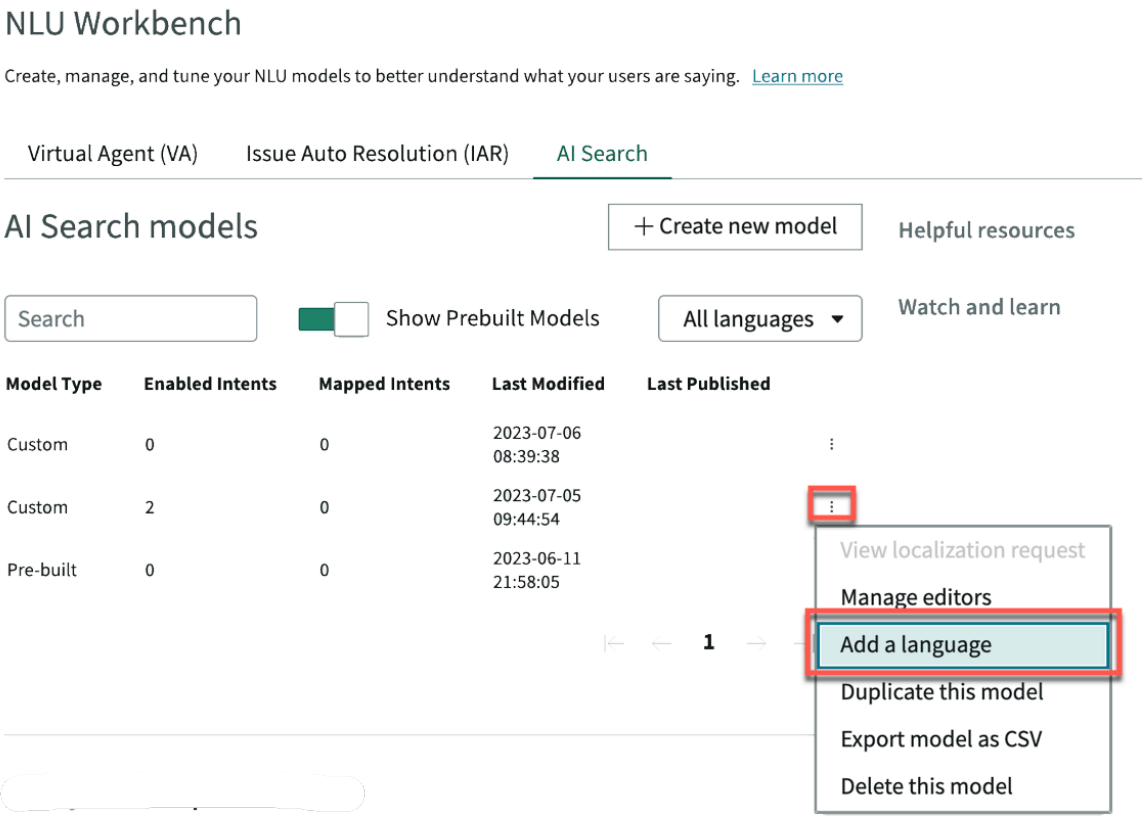
In this example scenario, your primary model uses the English language and you're adding a version in the French language, without translating utterances.

Procedure

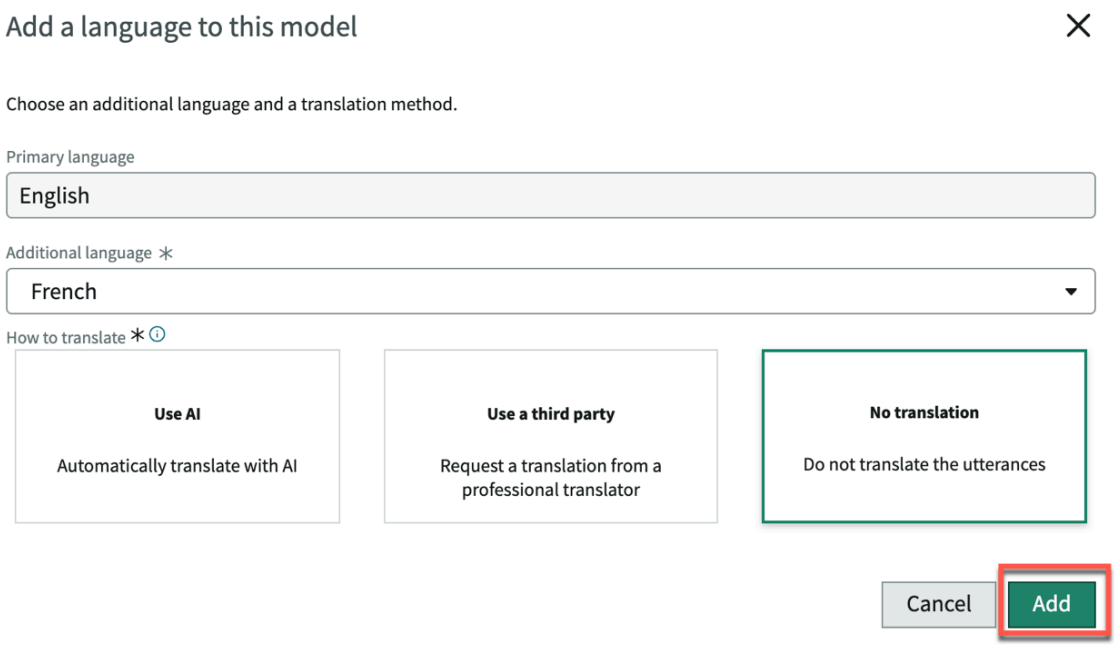
1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default. If appropriate, select the AI Search tab.
2. Identify a primary model to which you want to assign a secondary language.
3. In the row of the primary model, select the more options menu, then select **Add a language**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:
You may need to scroll to the right to find the more options menu.



4. In the **Add a language to this model** window, choose the desired language for the secondary model and the translation method.



5. Select the **Add** button.
The secondary model is created and the **Build and train your model** phase of the secondary model loads.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

If you chose to enable the model without translation, you must manually enter training utterances for each intent. If you choose to use software or a third-party translator, you must review the translations before you can continue working on the model.

Entering training utterances and reviewing secondary models can be delegated to users who have the nlu_editor role. For more information, see Assign an NLU editor to a model and Resolve intent issues.

Enable or disable a secondary model intent

Enable and disable intents in your Natural Language Understanding (NLU) models to make them active or inactive. Disable intents while editors or admins edit, review, or update its content and translations.

Before you begin

Role required: nlu_admin or nlu_editor

About this task

NLU admins and NLU editors can enable or disable intents in secondary models. If an intent is disabled, the model doesn't use it to make predictions. However, it also can't be accessed by any ServiceNow application using the model, such as Virtual Agent or Search.

An intent's state is represented in the **Enabled** column.



Enabling of intents works in the following ways:

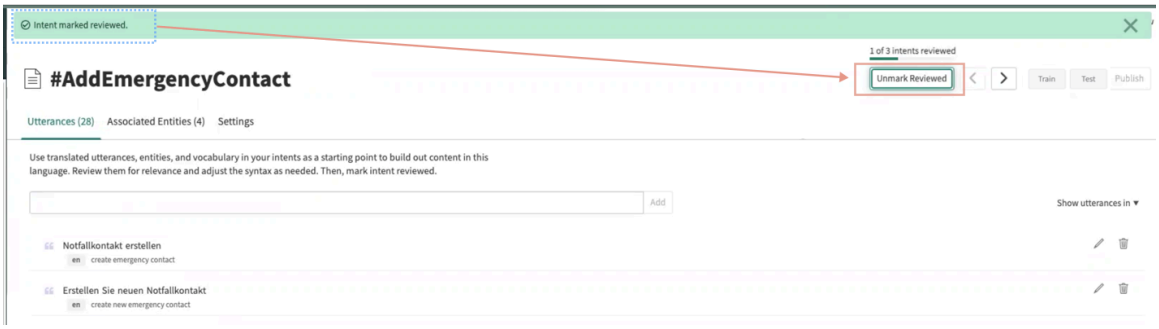
- Intents created in primary models are active by default.
- Intents from primary models are automatically created, but not enabled, in their secondary models (pending import).
- Disabling an intent in the primary model disables the corresponding intent in all secondary models.
- If an intent is disabled in the primary model, the intent can't be re-enabled in a secondary model.
- Enabling a disabled intent in the primary model does not enable the corresponding intents in the secondary models.
- Disabling intents in secondary models has no effect on the primary model or other secondary models.
- Secondary model intents that are marked with Needs review are disabled by default.
- If you disable a primary model intent that is mapped to a Virtual Agent (VA) topic, a confirmation message appears.
- If you disable a secondary model intent that is mapped to a VA topic, a confirmation message appears.
- NLU admins can enable and disable intents in any model.
- NLU editors can enable and disable intents only for models they are assigned to.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Disabling intents gives editors time to review the intent translations and update them if needed. When you're satisfied with the content, enable the intent to make it accessible to other NLU models and ServiceNow applications.

In this review example, you have a list of disabled intents in the Needs review state. The goal of this task is to review the translated content for a secondary model. When you complete your review, or if the content is fine as it is, you click Mark as reviewed. This moves the intent into the Reviewed state. You can also undo the Mark as reviewed state for an intent by clicking the **Unmark Reviewed** button, but only if you remain on the Intent screen. If you leave the screen prematurely, the **Unmark Reviewed** button disappears, and you won't be able to retrieve it.

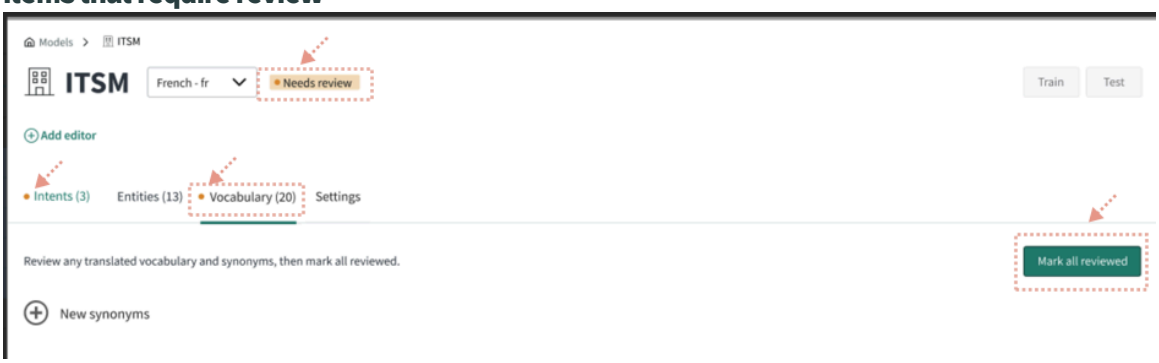
Switch between the Reviewed and Unmark Reviewed states for an intent



The editorial review also includes the Vocabulary section on the Model screen, where vocabulary that has been translated from the primary model to the secondary model is marked with an orange dot and the Needs review state. You can edit or delete the vocabulary items, even if they are in review. After you review the vocabulary and update the translations if needed, you can mark all of them at once by clicking the Mark all reviewed button. This action makes the orange dot disappear from the Vocabulary section of the Model screen.

Note:
The **Train** and **Try** buttons on the Model screen are disabled until the vocabulary is reviewed, and the intents in the Needs review state are either reviewed or disabled.

Items that require review



In this example scenario, you have applied the Use Software translation mode to translate a secondary model into the Japanese language. Because the translation hasn't been reviewed yet, the system marks the intents and the model with the Needs review state. When a model is in the Needs review state, it can't be trained, tested, or published. When it moves out of the Needs review state, you can then train and test the model.

Note:
An NLU editor can't create or publish a model. Only the NLU admin has permission to do this.

In this example procedure, you're reviewing the intents in a model one at a time.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

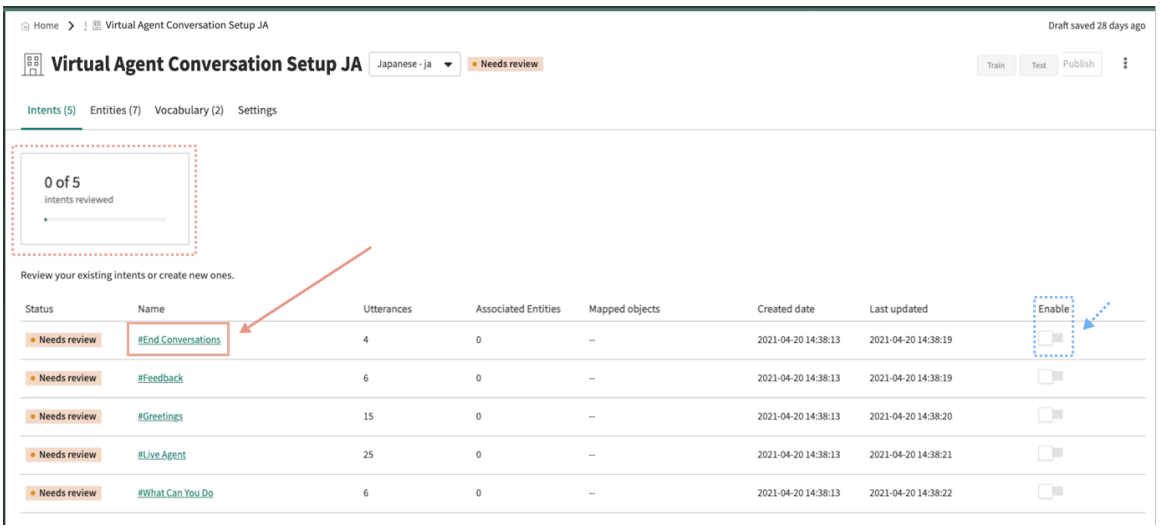
Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
2. Locate one of your secondary models that has already been translated.

   In this example scenario, you locate the Virtual Agent Conversation Setup JA model, as shown in the image below.

   In the Intents section of the model, there's a list of five intents that are marked with the Needs review state. Above the list of intents, there's also a box that shows a running count of the intents that need review.

The needs review state



3. Click an intent Name so you can access the intent content.
   In this scenario, you click the **#End Conversations** intent name.
4. Review the secondary language translations of the model utterances, entities, and vocabulary, then update the content as needed.
   When you finish your review of the intent, the Mark as reviewed button appears.
5. Click **Mark as reviewed**.
   Result: The system updates the **#End Conversations** intent from the Mark as reviewed state to the Reviewed state. In addition, the reviewed intents box changes its running count from 0 of 5 to 1 of 5.
6. If you're satisfied with your intent review, click the **Enable** button so the intent is available and can be accessed by any application, such as Virtual Agent (VA) or Search.

What to do next

Repeat Steps 1 through 6 for the four remaining intents in the list. As each intent review completes and is marked with the Reviewed state, the running count in the intents reviewed box increases toward 5 of 5.

Assign an NLU editor to a model

Assign an editor to review your Natural Language Understanding (NLU) model translations and edit model content. Delegate the maintenance, testing, and optimization of model content to an editor.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Create a model, or use an existing one. For translation review, create both primary and secondary (translated) models. For more information on multilingual model groups, see Multilingual model management.
- Assign the nlu_editor role to users. See Assign a role to a user.
- Role required: nlu_admin. The nlu_editor can't assign another editor to models.

About this task

The nlu_admin can delegate model content work to an nlu_editor. Assign language specialists as editors to review the consistency between primary and secondary (translated) models, or further localize the translation by editing intents and utterances. Delegate the maintenance of an existing model's content, without affecting the model's published status.

The nlu_editor can do the following, when assigned to a model:

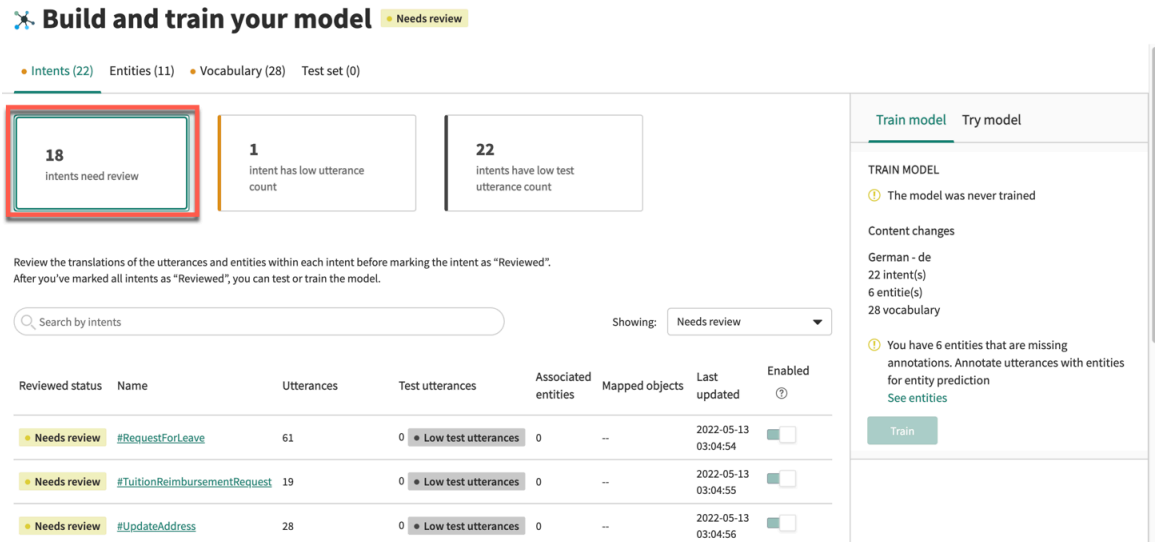
- Train and test the model.
- Export an NLU model as CSV.
- When a model or intent is in the **Needs review** state, the editor can mark it as **Reviewed**.
- Modify the model's confidence threshold.
- Edit the model name and description in **Settings**.
- Review, edit, and approve translated utterances.
- Add, edit, and delete training utterances.
- Annotate training utterances (add references to NLU vocabulary and NLU entities).
- Update entity properties (except the entity name in a secondary model).
- Add, edit, or delete the model's vocabulary.
- View the **Vocabulary sources** module (read only).
- Add, activate, and de-activate NLU intents in primary and secondary models.
- Resolve conflicting intents; also access the Cross-model Conflict Review module.
- View the Irrelevance detection in NLU module.
- Add, edit, or delete a test utterance in a test set; also access the Multi-model Batch Testing module.

Note:
The editor must be assigned to a model and in the model's application scope to take these actions, unless otherwise indicated.

The nlu_editor can't do the following:

- Create, add, publish, or delete a model.
- Duplicate a model.
- Translate a model (initiate a translated secondary model).
- Import primary model content to a secondary (translated) model.
- Sync or edit vocabulary sources.
- View model performance.
- Provide feedback in NLU Expert Feedback Loop.
- Manage other editors, such as assigning an editor to a model or removing an editor from a model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



First, an admin assigns the nlu_editor role to a user. Next, the NLU admin assigns the editor to a model.

When you assign an editor to a model, you must do so at the individual model level. You can't assign an editor at the model group level.

The following models can't be assigned to an editor:

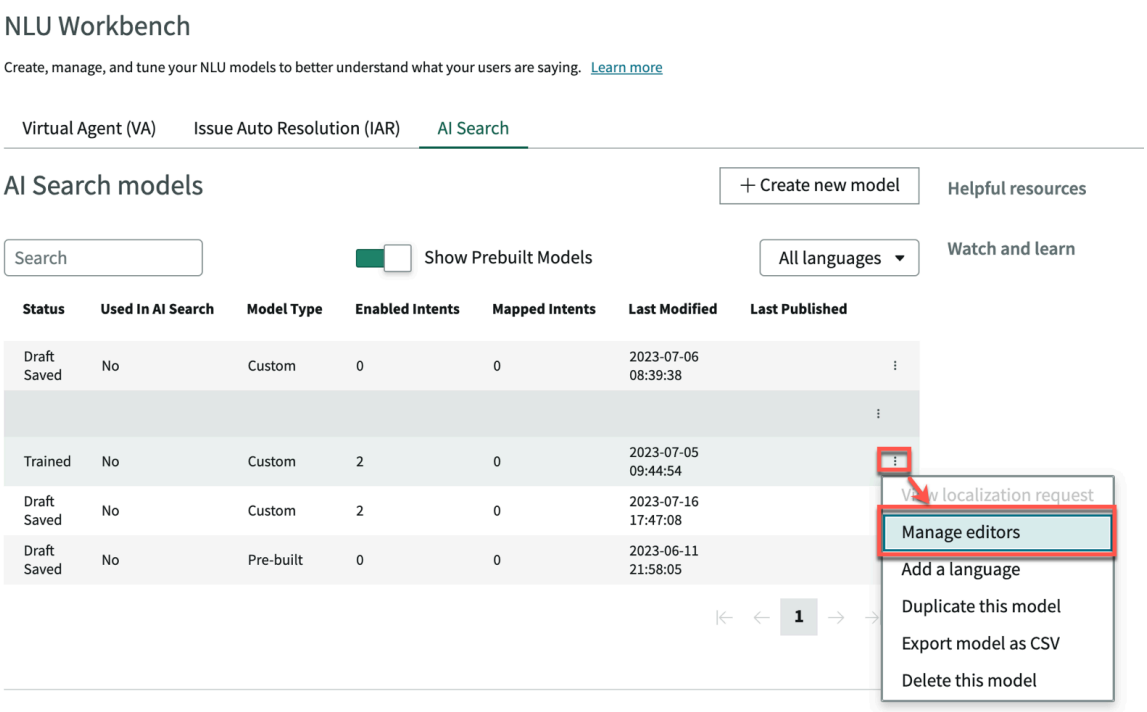
- Prebuilt models
- Models that have a translation in progress

In this example procedure, you add an NLU editor to a secondary model that needs reviewing.

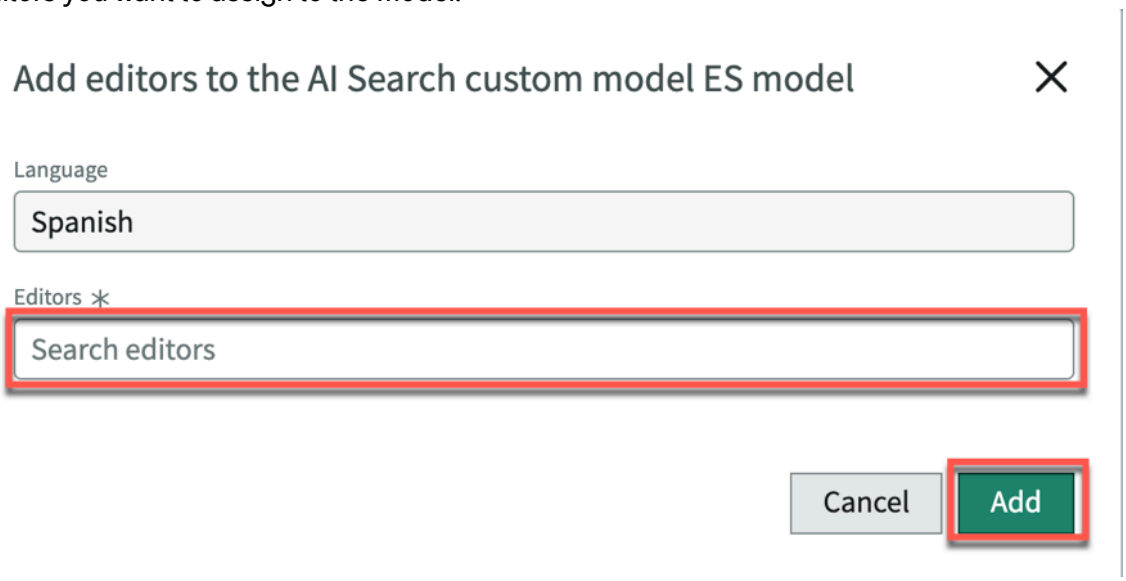
Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default. Select the appropriate tab for your model.
2. Locate the name of your NLU model and scroll to the right end of its row.
   If your model is in a multilingual group, expand the group to find the individual model. Editors cannot be assigned to groups.
3. Select the model's **More options** menu, then select **Manage editors** in the list.

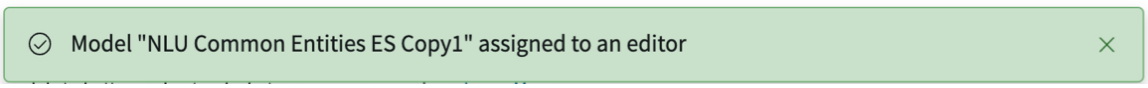
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. In the **Add editors** window, use the drop-down list and search bar to add the names of the editors you want to assign to the model.



5. Select **Add**.
   The editor is assigned to the model. A banner acknowledgement displays at the top of the screen. To add more editors, or remove editors, repeat this procedure.



What to do next

Assign more editors to the model, if necessary. You can assign up to four editors to a model at one time. To add more editors, or remove editors, repeat the above procedure.
As an NLU admin, ensure that your editors have what they need to properly review and edit models.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Import primary model content to a secondary model

When content in your primary NLU model is updated, you can import the updates directly to secondary models.

Before you begin

- Activate the Localization Framework (com.glide.localization_framework.installer) plugin. See Localization Framework.
- For Virtual Agent or AI Search models, create or use an existing primary model that has at least one secondary model.
- Role required: nlu_admin or admin.

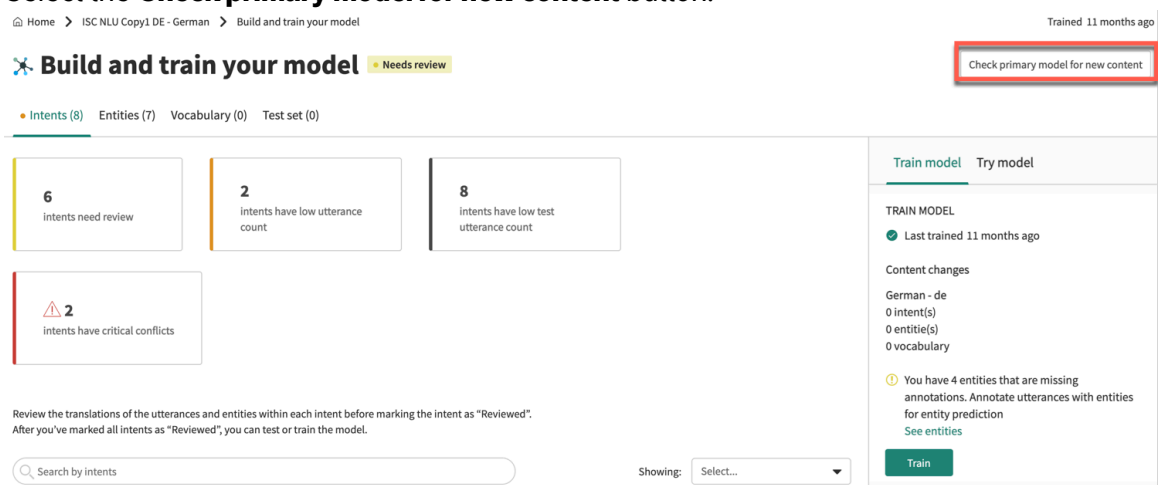
About this task

You can import content from a primary model to a secondary model after the initial translation. If you modify the primary model, those changes can be imported to your secondary models without the need to translate the entire primary model once again.

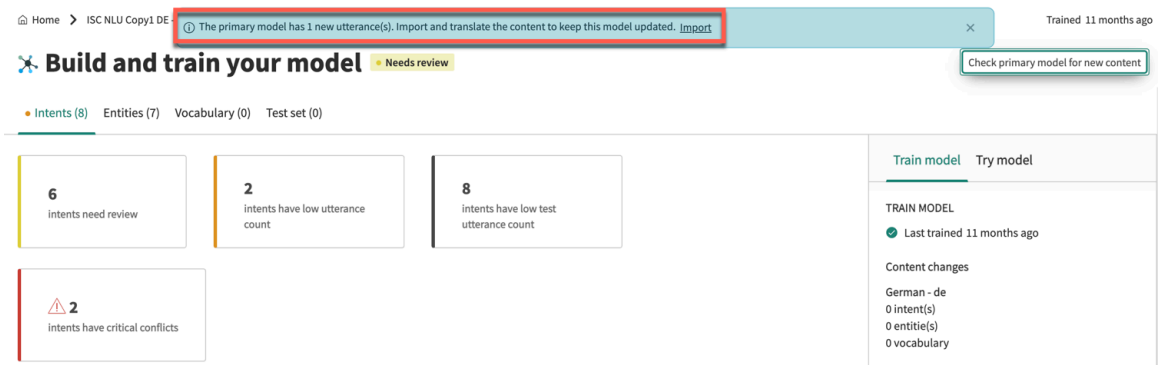
You can import primary model content to only one secondary model at a time.

Procedure

1. Navigate to **All** > **NLU Workbench** > **Models**.
   The Virtual Agent tab opens by default. Select the appropriate tab for your model.
2. Access a secondary model that has been translated under the primary model, and on **Build and train your model**, select **View phase**.
   By default you are taken to the **Intents** tab.
3. Select the **Check primary model for new content** button.



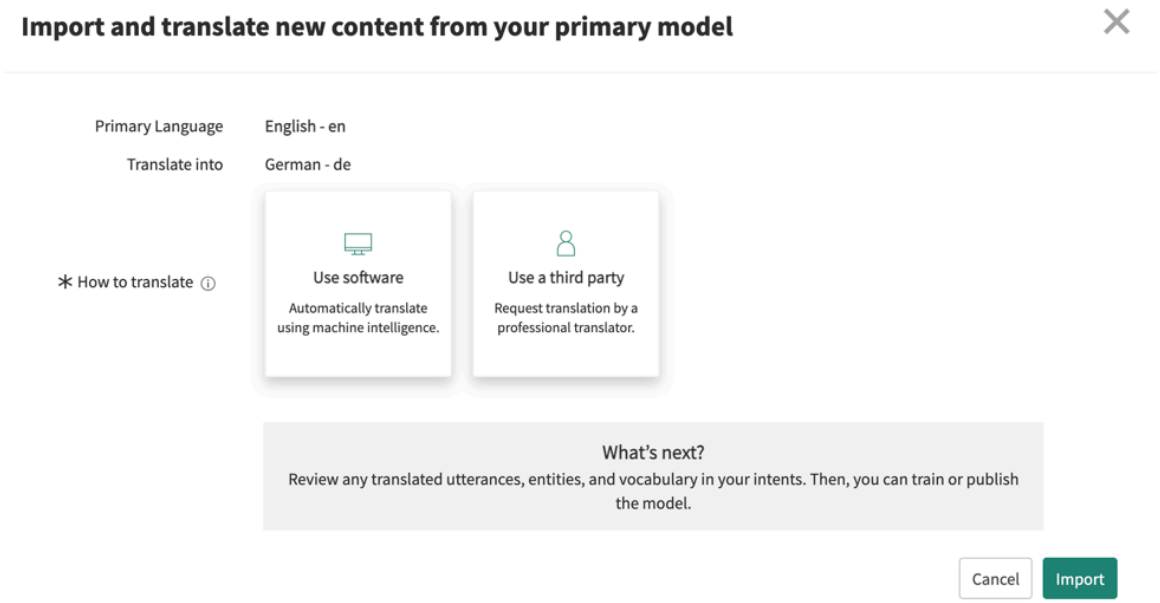
The banner Import and translate the content to keep this model updated appears at the top of the screen, along with a summary of the updates in the primary model.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Note:
By default, newly imported intents are disabled in secondary models. Activate any imported intents manually.

4. In the banner, select **Import**.
5. In the **Import and translate new content from your primary model** window, choose a translation method and select **Import**.



The new content begins translating. When finished, the translated content must be reviewed.

What to do next

Make sure that each of your secondary models receive the new content from the primary model. Review all the new, translated content before publishing the secondary models.

Virtual Agent and NLU Workbench integration

Virtual Agent administrators can access and update their NLU models from within the Virtual Agent Designer user interface.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Note:
If you have Now Assist in Virtual Agent, you can continue to use your existing NLU topics and migrate them into new LLM topics using the topic migration feature within Virtual Agent Designer. For more information on topic migration, see Migrate NLU topics to LLM topics.

Integration setup tasks, roles, and details

As Virtual Agent administrators create and configure their conversation topics, they must first create their NLU model and its associated intents in the NLU Workbench. This action requires they use the NLU Workbench and the admin or nlu_admin role.

In addition, they must also complete the following tasks in Virtual Agent **General Settings**.

- Enable NLU.
- Select the NLU service provider.
- If using language-specific NLU models, enable the languages for those models.

Virtual Agent administrators must also apply their NLU model to a conversation topic by completing the following tasks in Virtual Agent Designer:

- In **Topic Properties**, select the NLU model, the NLU intent, and the topic switching behavior.
- For input controls used in the topic flow, set the NLU properties for entity extraction.

Optionally, admins can activate Dialog Acts to enable Virtual Agent to respond flexibly when users make a modification in mid-conversation. Currently available response types are Modify, Affirm, and Negate, based on the last 5 exchanges in the conversation. Dialog Acts can be configured for English only, in Topic Properties. For more information see Dialog Acts for Virtual Agent.

After the NLU model is complete and associated with a Virtual Agent conversation topic, administrators with the virtual_agent_admin or admin role can do the following from within the Virtual Agent Designer user interface.

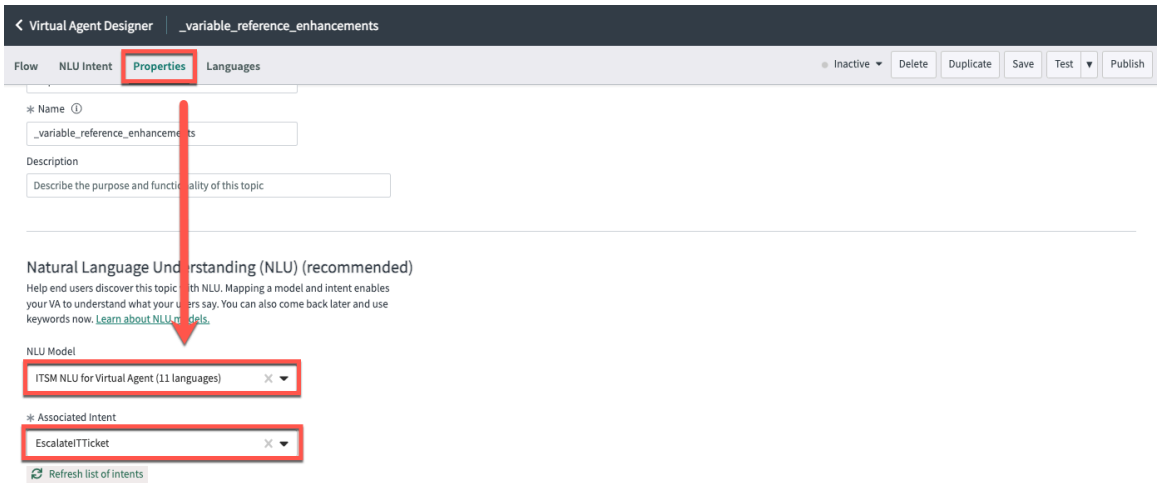
- Update NLU intent utterances.
- Train, test, and publish the NLU model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see Natural Language Understanding (NLU) topic discovery in Virtual Agent.

Publishing topics from Virtual Agent

When your model is published in the NLU Workbench, it's ready to use in Virtual Agent Designer. When editing a topic, click the **Properties** tab to select a model and intent to map to that topic.



When you click **Publish**, the model and intent are mapped to that topic and published seamlessly.

Publishing a topic with a mapped intent fails for the following reasons:

- Model isn't trained, or training is in progress.
- The last trained model is already published with a VA topic.
- The intent is not enabled in the model.

NLU Workbench - Advanced Features

NLU Workbench - Advanced Features expands the functionality of NLU Workbench to help you manage and improve your models.

Summary usage

The NLU Workbench - Advanced Features application includes the following features:

- Multi-model Batch Testing: Test large groups of utterances against your NLU models to see how the model predicts intents.
- Cross-model Conflict Review: Identify conflicting intents within or across models to improve model performance.
- NLU Model Performance: Use NLU Model Performance to see how well your models predicted intents in Virtual Agent based on end-user confirmation.
- NLU Expert Feedback Loop for VA: Provide feedback on Virtual Agent chat log utterances to help the system continuously learn and to better predict user input.
- Irrelevance detection in NLU: train your Virtual Agent models to ignore utterances that are not relevant.
- Issue Auto Resolution Tuning in NLU (Expert Feedback Loop for IAR): provide feedback to your ITSM Issue Auto Resolution model to improve its predictions.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

After the application is installed and activated, Model Performance, Expert Feedback Loop for VA, and Expert Feedback Loop for IAR appear under **All** > **NLU Workbench**. Multi-model Batch Testing and Cross-model Conflict Review appear under **All** > **NLU Workbench** > **NLU Advanced Features**.

Installation

NLU Workbench - Advanced Features is available from the ServiceNow Store. For instructions on how to purchase and download, see Install NLU Workbench - Advanced Features.

Related features

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 Intent Discovery.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System 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 **Ignore**.
NLU Batch Test Result [nlu_batch_test_result]	Contains the parsed results of a batch test including each utterance and predicted intent.
NLU Conflict Execution [nlu_conflict_execution]	Stores the information related to each execution of a conflict analysis run for a given definition.
NLU Batch Test Run Definition	Contains base information related to the batch test execution.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
[nlu_batch_test_run_definition]	
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.
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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Summary 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.

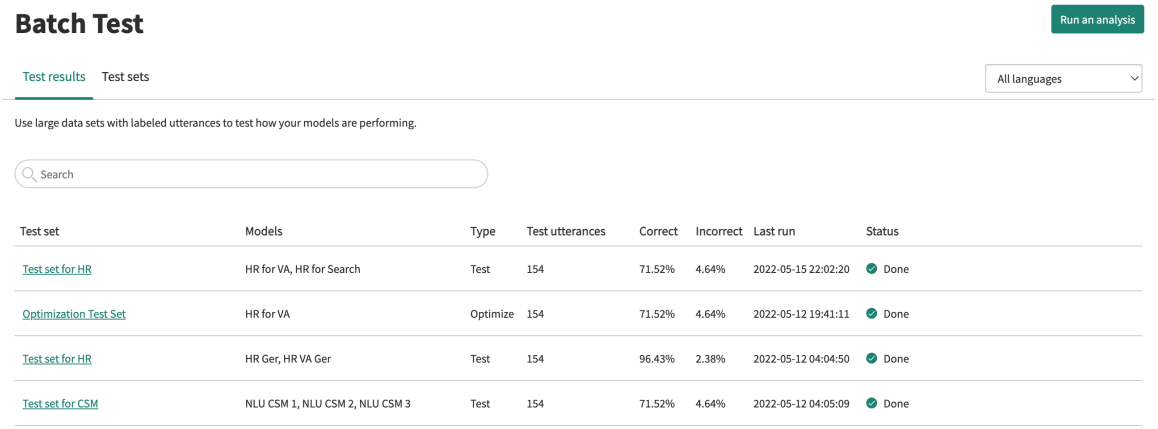
After you have a test set, you can test trained NLU models. To begin testing, see Run a multi-model batch test.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



To see 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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.

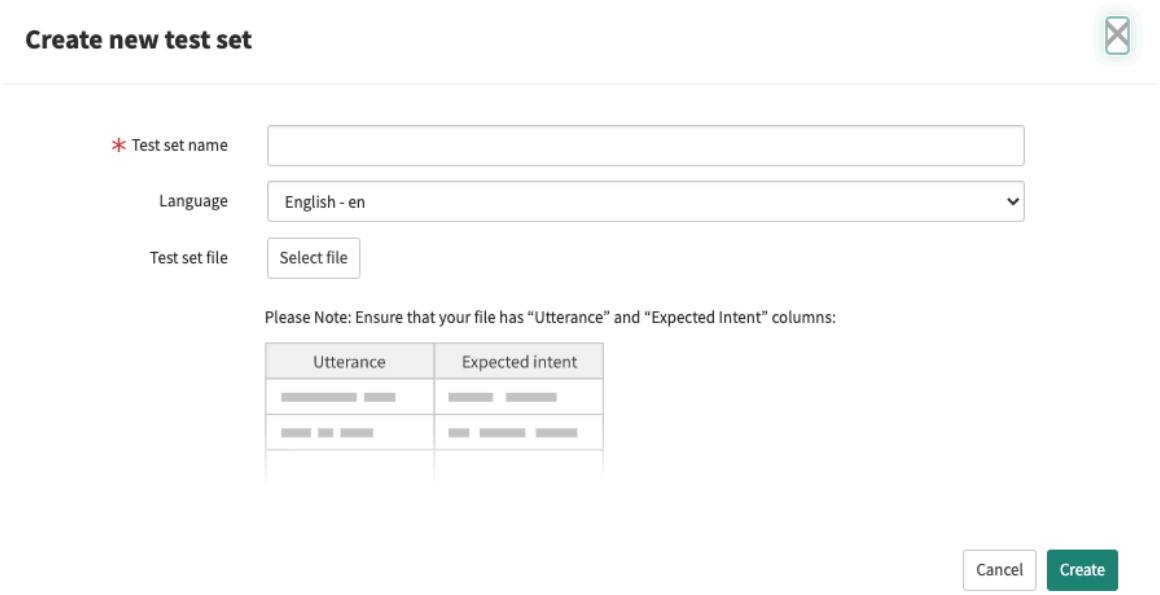
Note:
To indicate that a test utterance in your file has no expected intent, the value for "Expected intent" should be empty.

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**.



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.

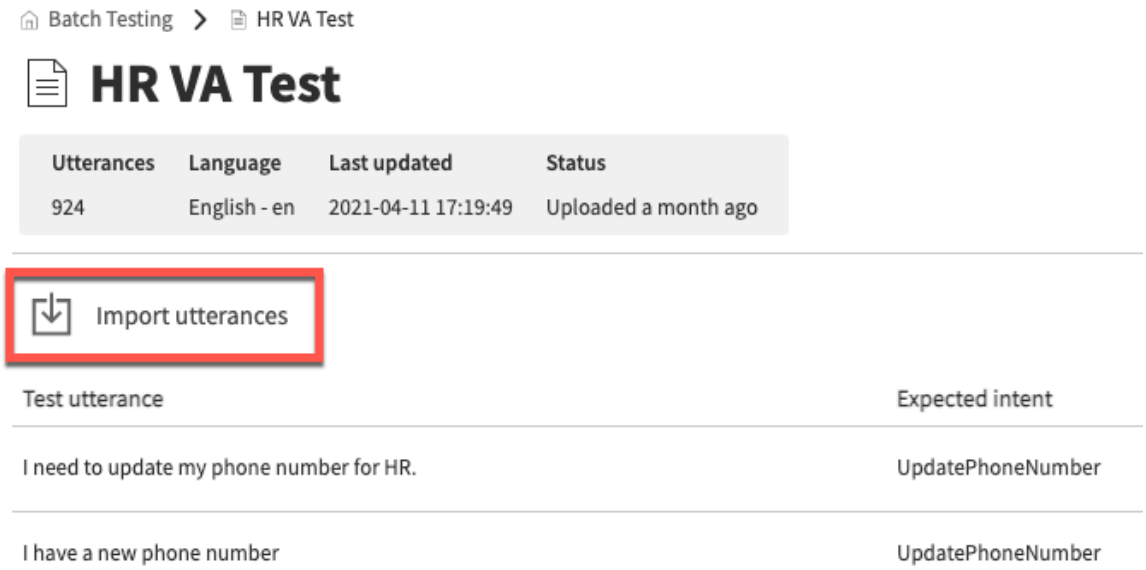
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

What to do next

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.

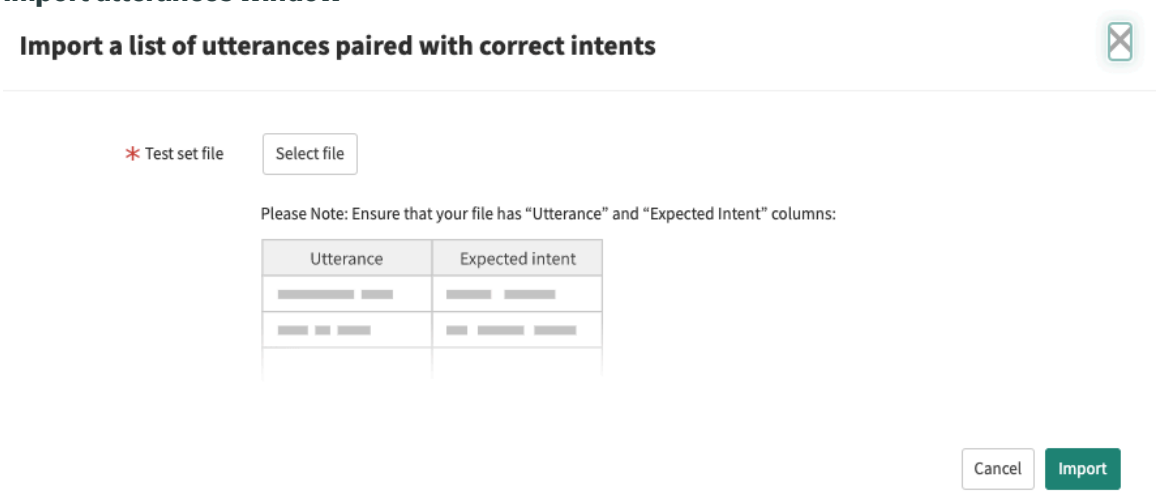
After you have created a test set, you can add more utterances to it. When viewing a test set, click **Import utterances**.

Test set page with import utterances button



Select a CSV or XLSX (Excel Workbook) file with your additional utterances to import.

Import utterances window



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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- Make sure that the NLU Workbench - Core plugin, NLU Workbench plugin, 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.

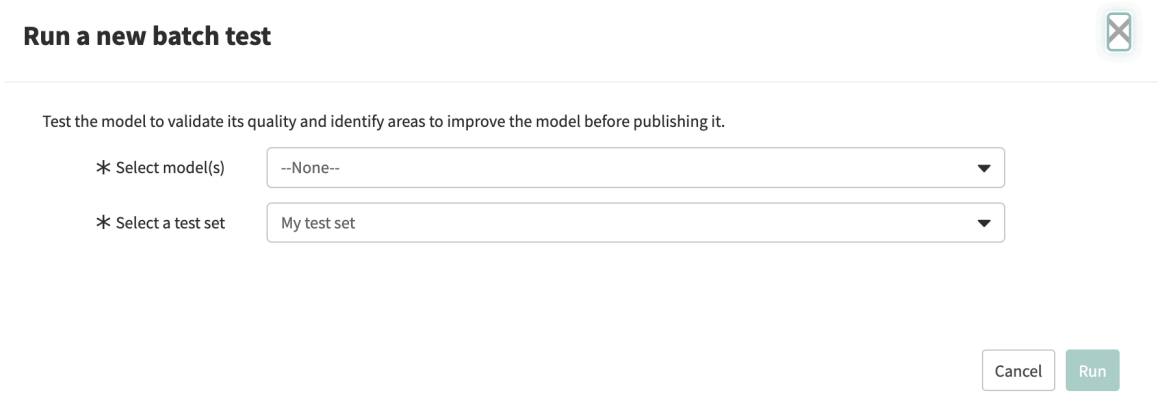
When testing multiple models, your test set must cover at least 25% of the total intents of all the models. Use test sets that contain utterances the models are likely to encounter in Virtual Agent or AI Search.

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.

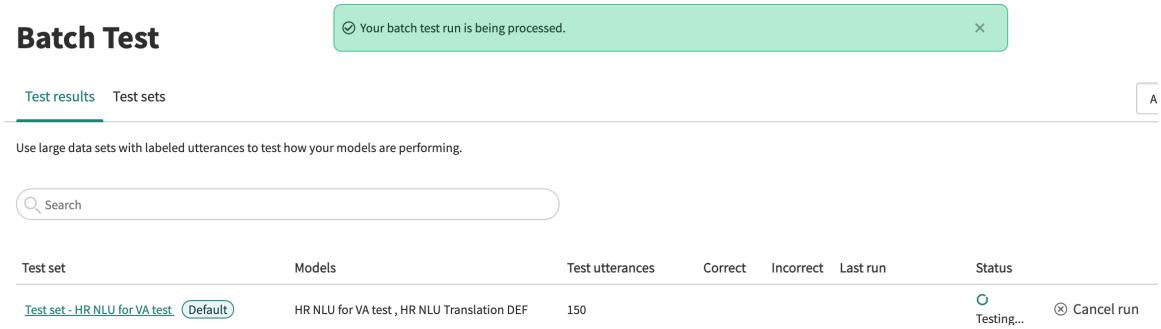
Procedure

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.



4. Select a test set from the list.
5. Click **Run**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Your 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,

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

you'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.

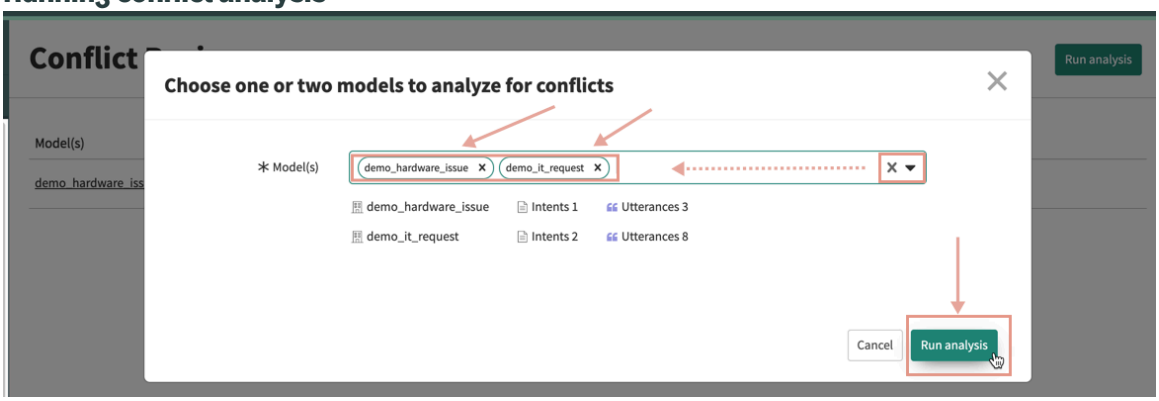
1. Navigate to **All** > **NLU Workbench** > **NLU Advanced Features** > **Cross-model Conflict Review**.

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

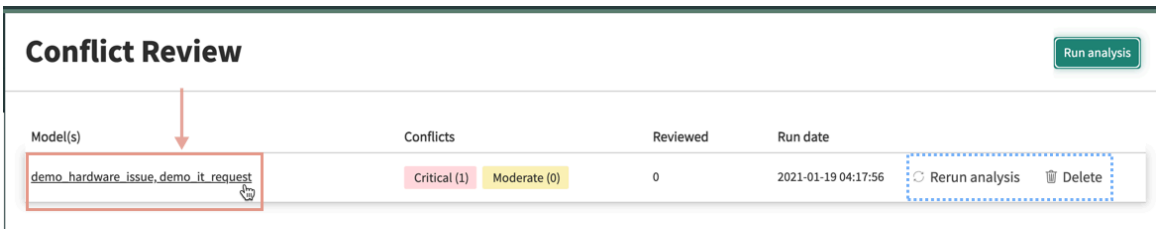


The Conflict Review screen refreshes to show the analysis, including the two models you selected for analysis, the counts of Critical and Moderate conflicts under review, the number of reviews that have been completed, and the run date for the analysis. If you point to the far right column on the screen 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

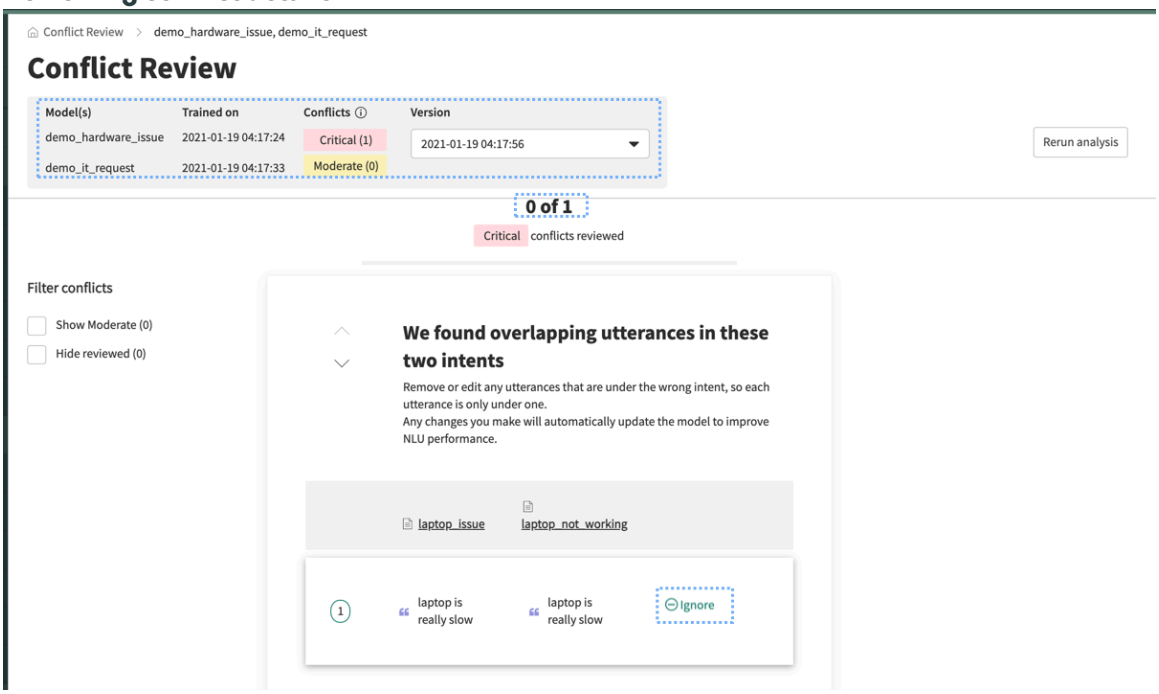
Conflict review



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



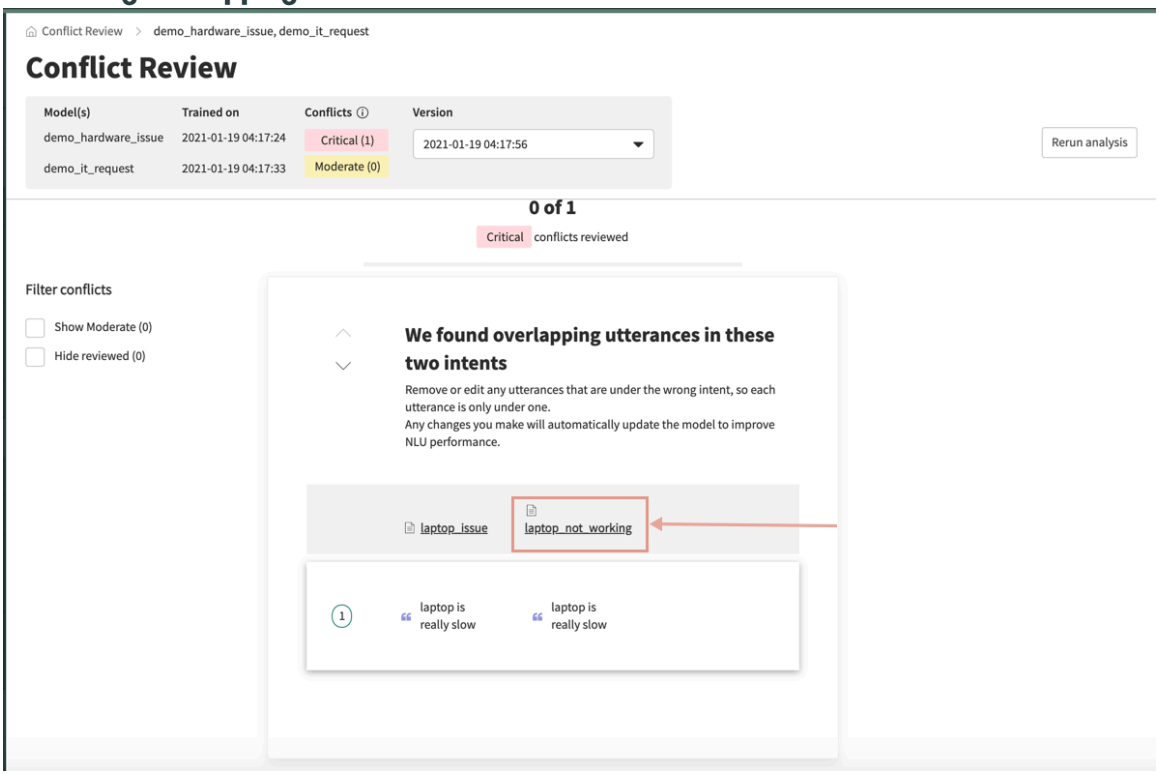
To determine how you will resolve this conflict, consider the 2 intent names and the identical utterances they share. Consider which intent is more likely to use the laptop is really slow utterance. If you compare the 2 intent names closely, you might realize that a laptop that's really slow isn't the 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**.

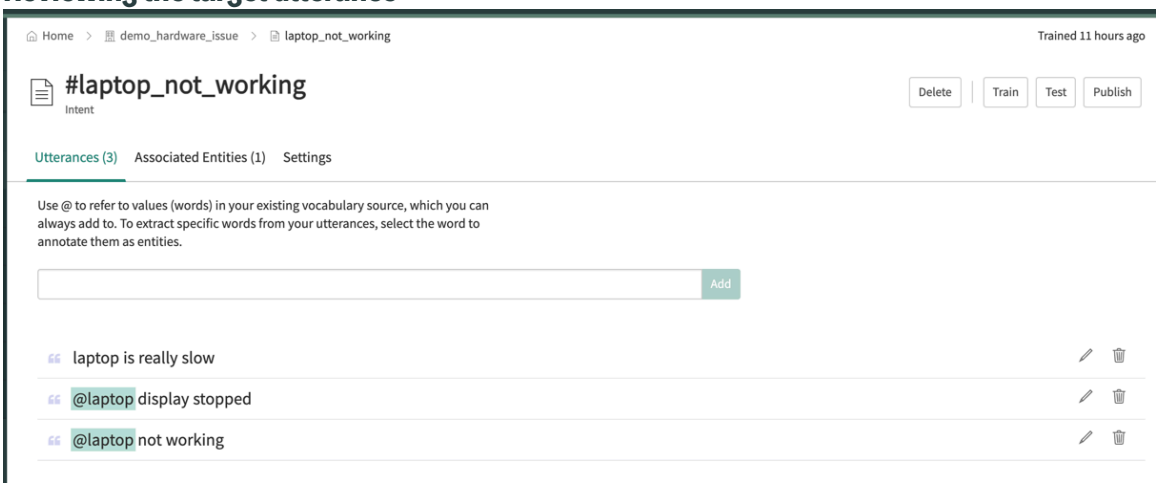
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Reviewing overlapping utterances



The #laptop_not_working Intent screen appears, showing its current 3 utterances. You make the decision to delete the laptop is really slow utterance from the #laptop_not_working intent.

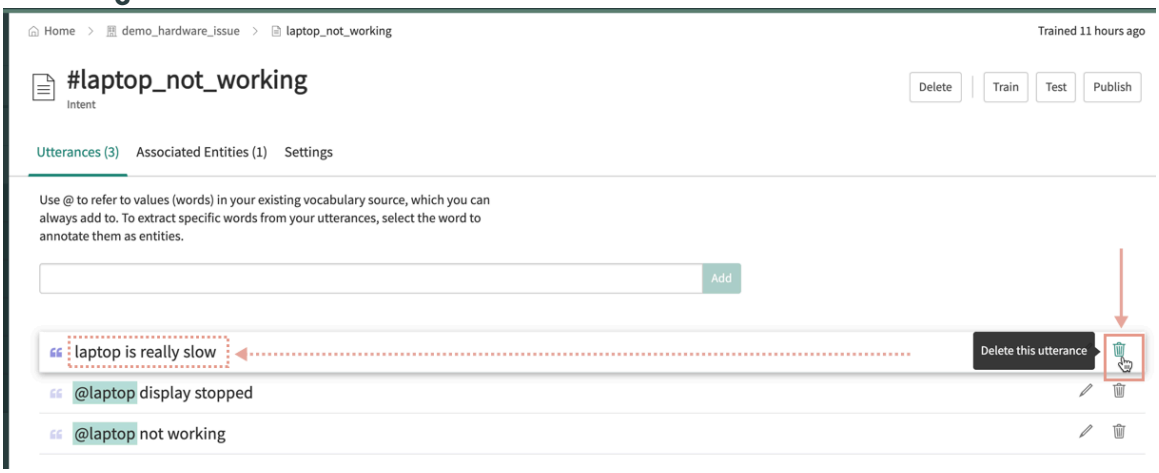
Reviewing the target utterance



7. Click the **Delete this utterance** trash can icon.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

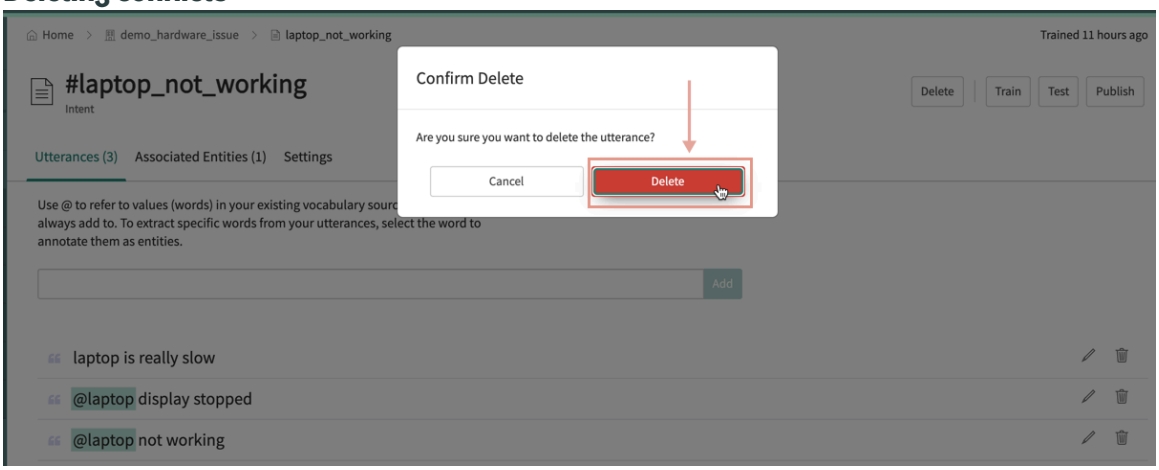
Resolving conflicts



The Confirm Delete screen appears.

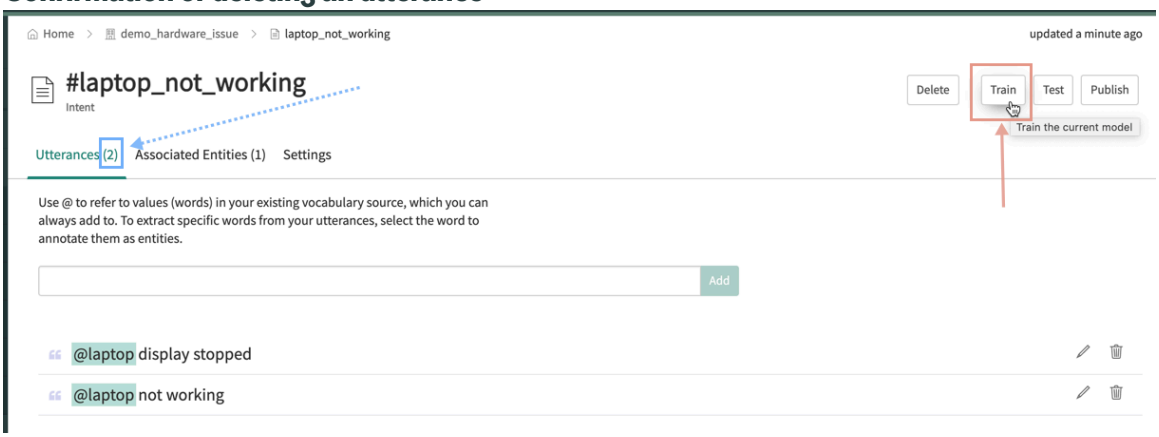
8. Select **Delete**.

Deleting conflicts



The Confirm Delete screen disappears, and the Utterances count drops from 3 to 2 because you've deleted the laptop is really slow utterance from the intent.

Confirmation of deleting an utterance



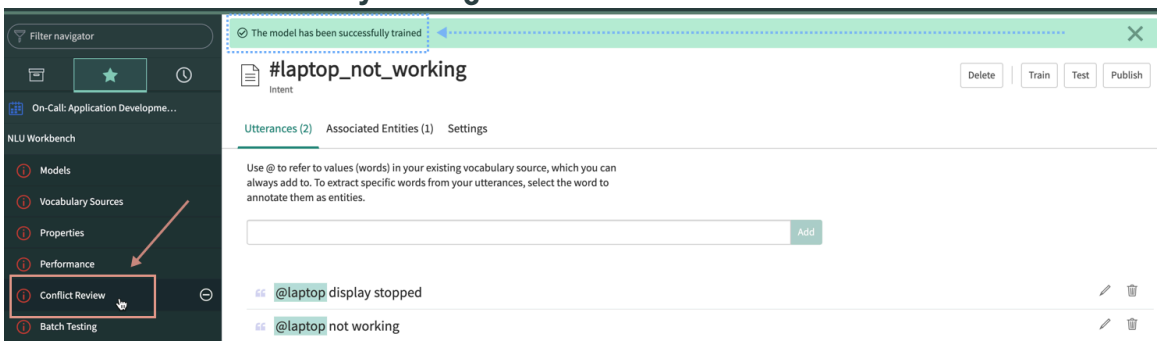
9. Select **Train**.

A banner appears on the Intent screen, confirming the model is successfully trained.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

10. Select **Conflict Review** in the navigator.

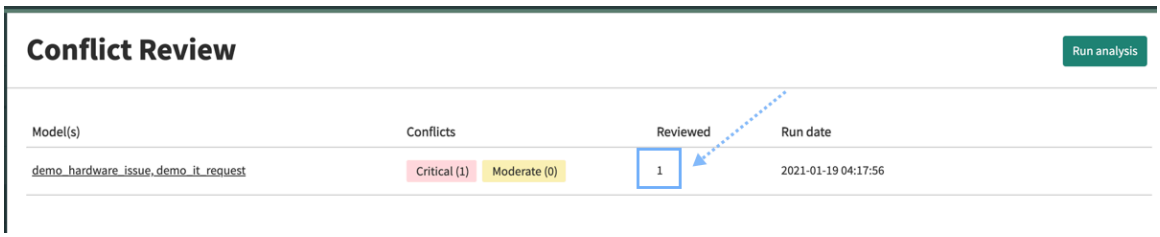
Confirmation of successfully training a model



Result:

The Conflict Review list screen appears, showing your conflict review analysis is complete, and that it's been reviewed.

Confirmation that a conflict has been reviewed



NLU Model Performance

Use NLU Model Performance to see how well your models predicted intents in Virtual Agent (VA) based on end-user confirmation.

Summary usage

The NLU Model Performance application provides an analysis and report of how well your NLU models predict VA users' intents from their utterances. As NLU models for VA are updated, published, and deployed, you can use NLU Model Performance to review the efficacy of the intent predictions they make. If these predictions are skipped, it's because they're unable to predict an intent with a high enough confidence score for the model. To improve VA topic prediction, the system groups unsupported VA utterances into clusters for an analysis and then generates a report that identifies the outcome of the analysis.

To access this application, use the nlu_admin or admin role and navigate to **All** > **NLU Workbench** > **Performance**.

Note:
To use NLU Model Performance, you must first have at least 5,000 VA utterances in your instance. With the nlu_admin role, you can modify this limit by resetting the *sn_nlu_workbench.glide.nlu.performance.min_clustering_records* property.

Installation

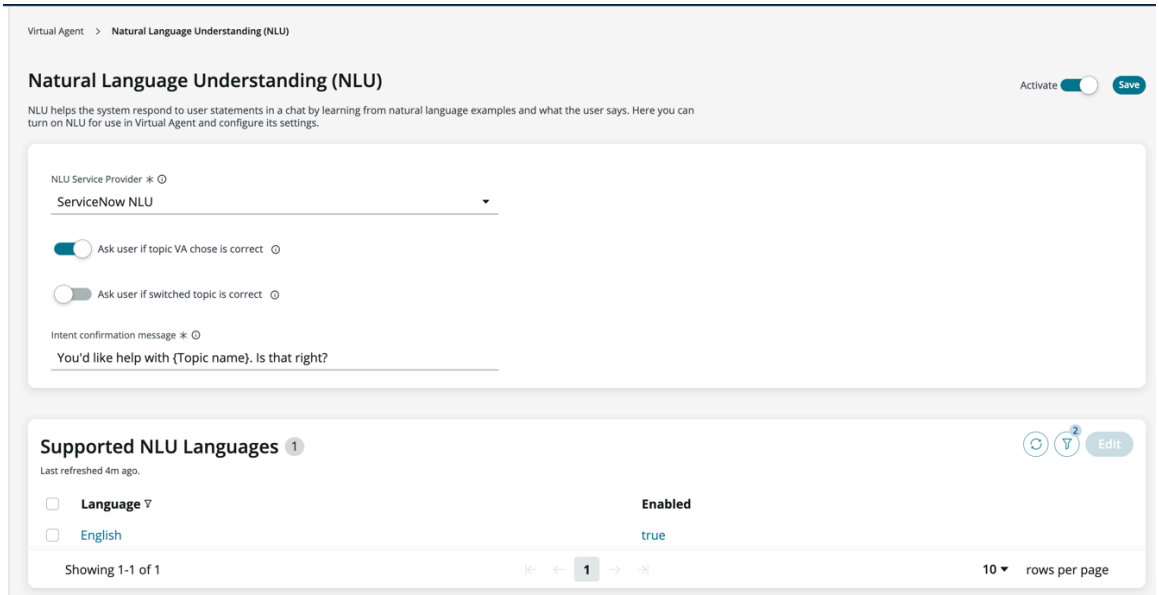
NLU Model Performance is an application available on the ServiceNow® Store. To use this application, ensure that the NLU Workbench - Advanced Features (sn_nlu_workbench) plugin is active on your instance. For more information, see Install NLU Workbench - Advanced Features and Activate the NLU Workbench. Because NLU Model Performance relies on data coming in

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

from VA, ensure that the Glide Virtual Agent (com.glide.cs.chatbot) plugin is also active. See Activate Virtual Agent.

Configuring Conversational Interfaces settings for ServiceNow NLU

To help with tracking NLU performance, you must first configure the Conversational Interfaces (CI) settings in Virtual Agent. To access these settings, navigate to the top of the **NLU Model Performance for Virtual Agent** landing page and click **Conversational Interfaces Settings**. This action takes you to the CI settings page in Virtual Agent, as shown in the image below.



To configure the settings, you need to use the virtual_agent_admin or admin role. For configuration guidance, see Implement NLU in Virtual Agent and Enable NLU languages in Virtual Agent settings.

Reviewing the user interface (UI)

The NLU Model Performance UI is comprised of the listed parts below, which are highlighted in the following image that shows the NLU Model Performance default landing page.

- The **Performance** tab and its colored chart which shows data describing scenarios that occur when users interact with VA. This chart is segmented in shades of three colors, each representing a user scenario. The rows in the chart that have text and percentages are referred to as the legend. The three scenarios in the first row of the legend have two outcome scenarios beneath them in the second row. The bars underneath the legend correspond to the colors of the scenarios above them. These bars range in size based on the percentage value for each scenario.
  - For example, the dark blue bar shows you that 10% of this batch of VA users confirmed that the system presented them the correct topic to use for their VA chatbot conversation, and that 58 % of the users (represented by the dark red bar) confirmed that the topic presented to them was incorrect. Residing between these two scenarios is a scenario showing that 32% of the users (represented by the dark brown bar) didn't make any confirmation at all.
  - It's a good idea to toggle the **Show all bars** filter back and forth at any time to see or hide the colored bars and their associated scenarios in the chart. Note also that when you first access NLU Model Performance, the default view of the chart shows only the first row of scenarios and the first row of bars.
  - There are also other views you can see depending on which scenario you're viewing. For example, when you first access NLU Model Performance and you click a scenario in the first

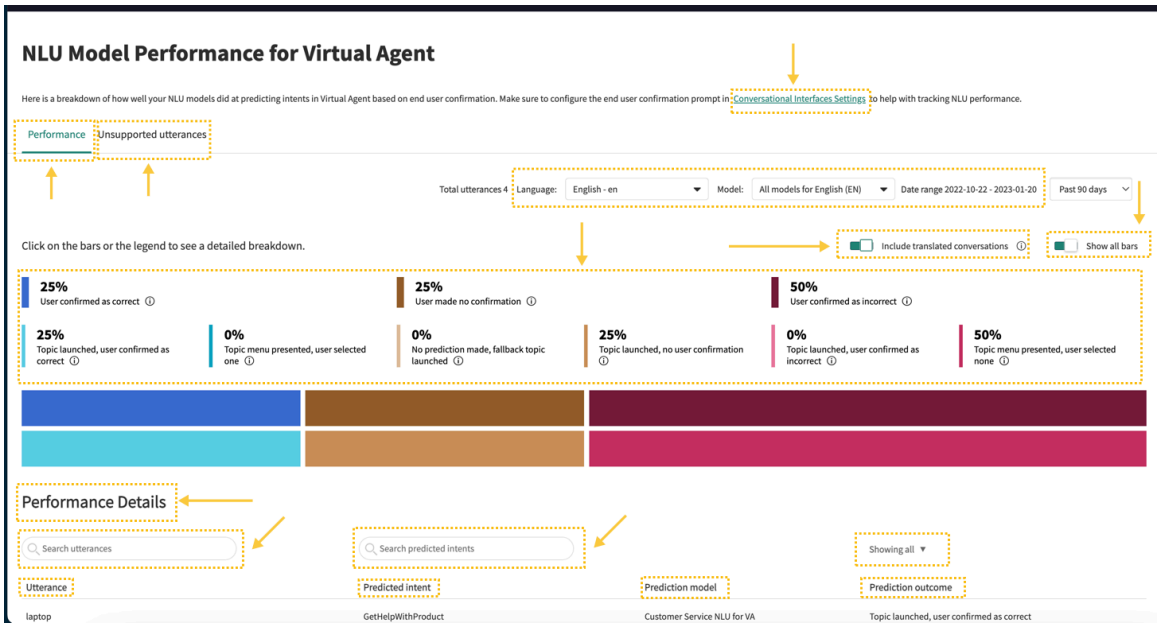
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

row of the legend, the system hides the bars of the other two scenarios. This helps to keep the UI less cluttered because it partially isolates the scenario you're focusing on from the rest of the scenarios. These actions don't change any data, they simply show you different views of the legend and bars within the chart.

- The **Unsupported utterances** tab, which shows utterances that are grouped into clusters for use in the performance analysis. This section of the UI is where you can navigate to and return from the Expert Feedback Loop application, and also where you can run the analysis.

  Unsupported utterances are different from irrelevant utterances. For more information, see Irrelevance detection in NLU.

- The **Performance Details** section located underneath the bar chart. This section has these four columns: **Utterance**, **Predicted intent**, **Prediction model**, and **Prediction outcome**. The details shown in this section interact with the legend data and bar data above them.



Reviewing the VA user scenarios

In the legend section of the UI, the text and percentages are accompanied by Information icons. Point to an icon to invoke the definition for its user scenario. Refer to the table below to see the definitions.

Scenario Definitions

Scenario	Definition
User confirmed as correct	The correct topic was presented to the end user and the end user has confirmed it is correct.
Topic launched, user confirmed as correct	The topic was launched and the user confirmed it is the topic they needed.
Topic menu presented, user selected one	Multiple topics were shown as a menu to the user and the user chose one of the topics to address their need.
User made no confirmation	The topic was launched but the user did not confirm whether or not it met their need.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

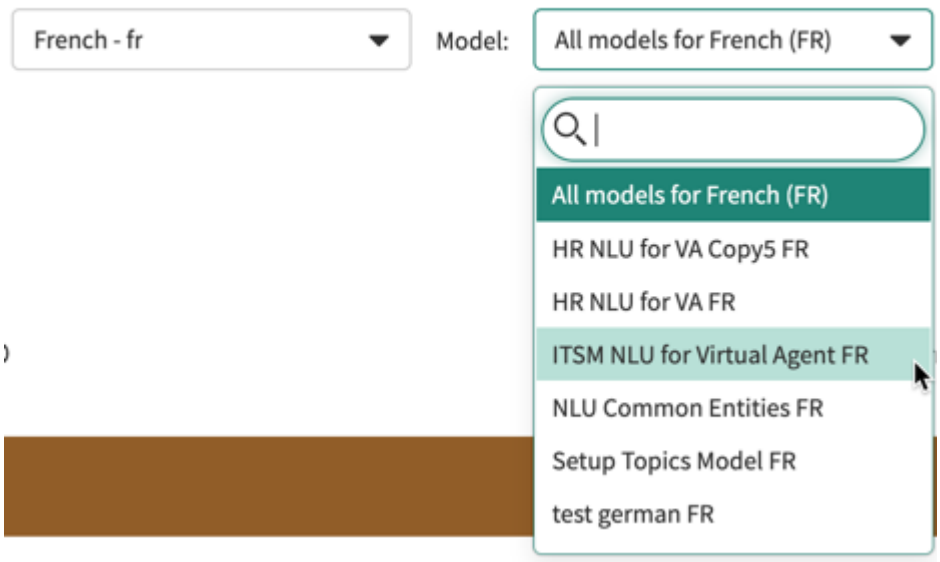
Scenario Definitions (continued)

Scenario	Definition
No prediction made, fallback topic launched	No predictions were made and the fallback topic was launched.
Topic launched, no user confirmation	The topic was launched but the user did not confirm whether it was correct.
User confirmed as incorrect	The topic(s) presented did not address the end user's needs, and the user has confirmed it is not correct.
Topic launched, user confirmed as incorrect	The topic was launched and the user confirmed it was not correct.
Topic menu presented, user selected none	Multiple topics were shown as a menu to the user and the user decided none of the topics were relevant and did not choose any topic.

Using the Language and Model filters

On the Performance tab, you can see the Language and Model filters. Next to them, you can also see the most recent Date range values you set above the legend and bars of the chart. Click the **Language** filter to see all languages that are available in NLU. Click the **Model** filter to see all prediction models that are available in your instance. The default setting for the **Language** filter is **All languages**, and the default setting for the **Model** filter is **All models**.

The Language and Model filters interact with each other. For example, if you choose `French - fr` in the **Language** filter, the **Model** filter automatically shows all prediction models in your instance that use the French language, as shown in the following image.



There are other interactions. For example, from the default view of the NLU Model Performance landing page:

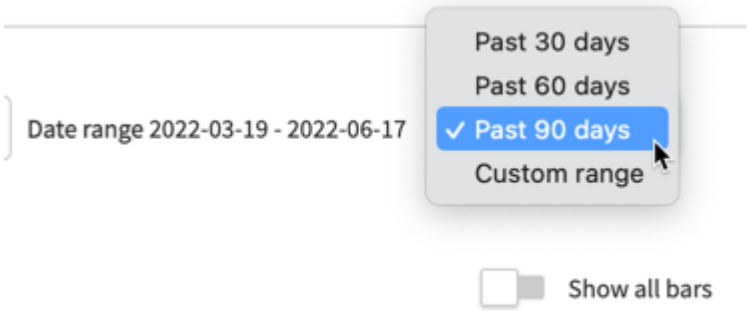
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If you select a specific model from the **Model** filter, the **Language** filter value is updated to display the language of the selected model.
- If you select a specific language from the **Language** filter, the **Model** filter only displays models of that selected language.

Underneath the Language and Model filters, you can see the **Include translated conversations** switch. To include the performance of VA conversations that were translated to your primary language using dynamic language translation, toggle the switch to the right so it changes color from grey to green. When you do this, the bars in the charts legend may also change position and colors. You can manage dynamic language translation in the **Conversational Interfaces Settings**.

Setting the Date range

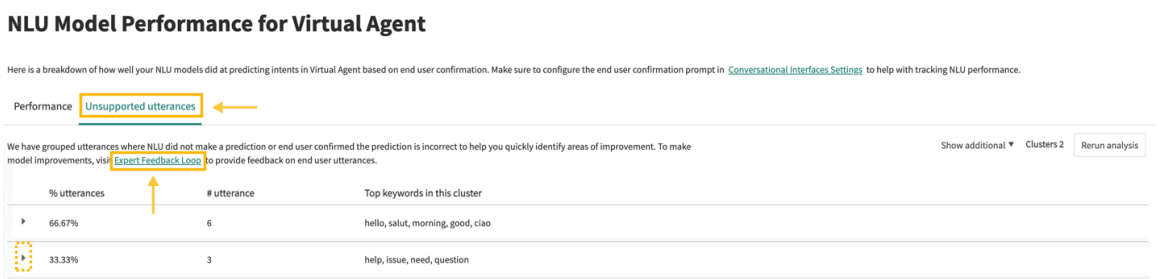
On the Performance tab, use the **Date range** to define how far back you want the system to pull VA data from. Choose **Past 30 days**, **Past 60 days**, **Past 90 days**, or **Custom range**. The further back you go in the date range, the more data you will have in your analysis.



Running an analysis

To run a performance analysis, click the **Unsupported utterances** tab. This section of the UI shows rows of expandable clusters containing VA utterances where NLU didn't make a topic prediction, or where the VA end-user confirmed that the predicted topic was incorrect. The next thing you want to do is to click **Expert Feedback Loop**. This action takes you to the NLU Expert Feedback Loop application where you review and provide feedback on the utterances that were pulled in from VA.

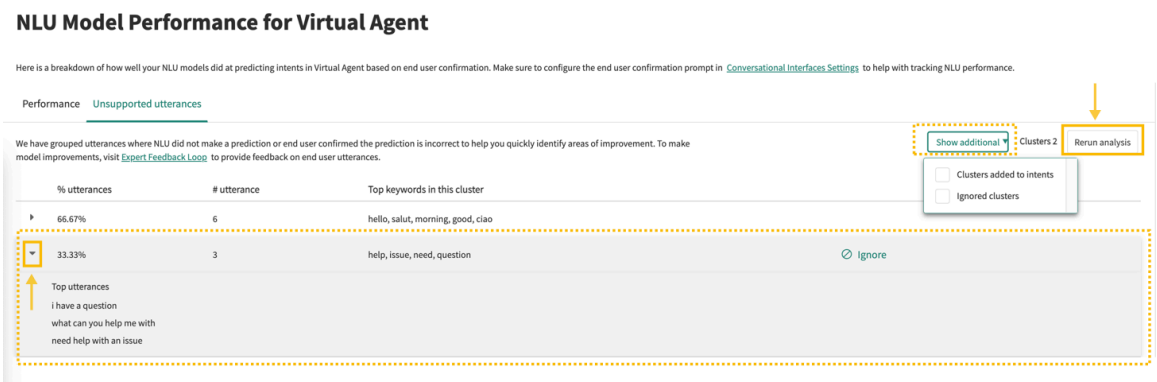
Reviewing unsupported utterances



When you leave the Expert Feedback Loop application and return to NLU Model Performance, click the **caret** icon in any cluster to open it. Within the cluster you can then see the top most representative VA utterances, as shown in the following image.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Reviewing the top representative utterances in a cluster



Click **Run analysis** or **Rerun analysis**, whichever is available. Each time you run an analysis, the system pulls the most recent VA utterances into the analysis.

Related topics

- Virtual Agent
- Natural Language Understanding in Virtual Agent

NLU Expert Feedback Loop

Provide feedback on Virtual Agent chat log utterances to help the system continuously learn and to better predict user input.

Embedded video: NLU Expert Feedback Loop, presented by the ServiceNow Virtual Agent Academy

Summary usage

The Expert Feedback Loop takes data from your instance and provides it to you for feedback. This data comes from your users' Virtual Agent (VA) chat logs and includes utterances from those logs. Using the nlu_admin role, navigate to **All** > **NLU Workbench** > **Expert Feedback Loop** and mark each utterance in a model by confirming whether the utterance is correct (match) or incorrect (mismatch) to an intent. This feedback helps the model to continuously improve the model performance.

Installation

Expert Feedback Loop is part of the NLU Workbench - Advanced Features app available on the ServiceNow® Store.

To use Expert Feedback Loop, ensure that the NLU Workbench - Advanced Features (sn_nlu_workbench) plugin is active on your instance. For more information, see Install NLU Workbench - Advanced Features and Activate the NLU Workbench.

Because Expert Feedback Loop relies on the data from VA chat logs, make sure that the Glide Virtual Agent (com.glide.cs.chatbot) plugin is also active. See Activate Virtual Agent.

Importing expert feedback loop data between instances

Before you start an import, make sure you have access to the data in your instance and have enough data to proceed with your feedback.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

If you're working in a sub-production instance, you must import the feedback data from the **open_nlu_predict_intent_feedback** table in your production instance to your sub-production instance. For guidance on how to import, see Importing from another ServiceNow instance.

Here are a few data scenarios regarding the system behavior for the NLU models that house the feedback loop utterances.

- If you move a model to a different instance, the feedback data persists.
- If you upgrade the instance, the feedback data persists.
- If you clone an instance, the data does not persist, so you need to follow the procedure referenced above to import the data onto the cloned instance.

Feedback context and access

The Expert Feedback Loop provides a mechanism for you to improve NLU models deployed to VA through feedback that you provide on a select subset of utterances. For each utterance, you are asked to confirm the predicted intent or to provide the correct intent that the utterance belongs to. After this feedback is complete, the data is used to tune the model performance, resulting in an improved model that can now be redeployed to gather more end user data. This is an iterative cycle that allows the system to continuously collect data, learn from the feedback you provide, and use that feedback to further improve your NLU models.

Once every 30 days, the system pulls up to 300 utterance samples from VA chat logs to the Expert Feedback Loop. The utterances are selected for feedback based on how well they represent all the utterances in the logs. Every utterance sampled from VA chat logs has a predicted intent picked by the system.

You can also set the number of utterances you want to be pulled from VA chat logs by adjusting the setting for the *glide.mlpredictor.option.nlu.activeLearning.label_candidate_table.max_respons* system property.

Further information on the process of pulling utterances from VA chat logs can be found in the sections Procuring additional VA feedback data on demand and Using the Active Learning job, below.

As you review the utterances, decide whether each utterance belongs to its predicted intent or should be moved to a different intent. If you aren't sure about the correct intent, the utterance can be marked to revisit for further review. After you have marked at least 100 utterances with your feedback, the system uses all the marked utterances to tune and improve the model.

Reviewing your VA chat log utterances

Now that you're in the loop, you can see in the following image a list of NLU utterances that support the intents in the VA models. Each page of the list shows approximately 20 utterances per page. Your goal is to review each utterance and mark its relationship status to a given VA intent. You mark each utterance with either the **NLU_Match**, **Mismatch**, or **Unsure** value.

Marking an utterance as **NLU_Match** means that you agree with the NLU prediction for that utterance. Marking an utterance as **Mismatch** means that the utterance belongs to a different intent, and if you select **Mismatch**, a dropdown of intents appears so you can select the correct intent. If you aren't sure which intent is correct, you can mark it as **Unsure**.

Utterances that are marked and saved with the **NLU_Match** value or the **Mismatch** value are moved to the **Completed reviews** section. Utterances that are marked and saved with the **Unsure** value are moved to the **Needs further review** section.

Here's a scenario for the basic steps you use to complete your utterance reviews:

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

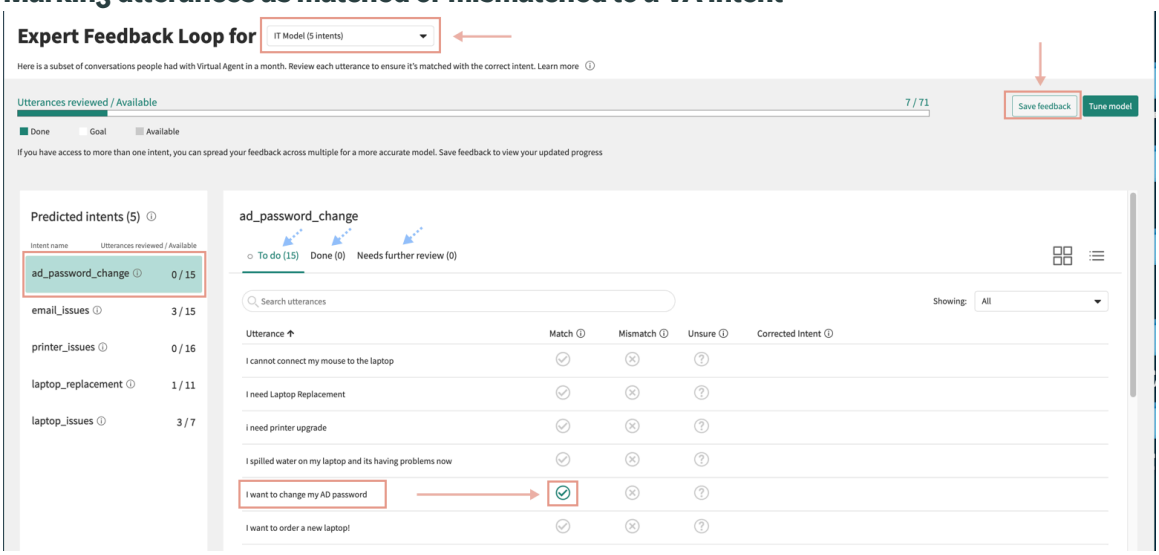
1. At the top of the Export Feedback Loop screen, select a model you want to review in the **Expert Feedback Loop for** prompt. Within this feature, in some ways it doesn't matter which model you choose because these models serve mainly to house the utterances you are marking in your reviews. So in this example scenario, you decide to select the **IT Model** model, which has five intents and many utterances within each of those intents. The number of utterances you need to review are shown next to the name of the predicted intent, as shown in the following image. The intents in these models are listed in the **Predicted intents** column on your screen. Although you can review and mark utterances in any of the predicted intents, it's a better idea to complete the reviews consecutively in the order the intents first appear. For example, begin with the **ad_password_change** intent and then onward to other predicted intents.
2. Click the **ad_password_change** intent so it loads its utterances into the **To do** section for your review.
3. Review the 15 utterances in the intent. In this scenario, you've started by correctly marking the I want to change my AD password utterance as a match to the **ad_password_change** intent.
4. Click **Save feedback**.

   Result: The system moves the utterance you marked over to the **Done** section of the screen. At the same time, the count of **To do** utterances drops from 15 to 14, while the count on the **Done** section rises from 0 to 1. If you had also marked another utterance with the **Unsure** value and saved that change, the count of utterances in the **Needs further review** section would also rise from 0 to 1.

   Note:
   If you don't know which intent best matches the utterance, mark it with the **Unsure** value. This action moves the utterance over to the **Needs further review** section, giving you time to mark other utterances in the **To do** section that may be easier to match or mismatch to an intent. You can always return to the **Needs further review** section to mark utterances that are more difficult to pin down to a specific intent.

5. Repeat steps 1 through 4 as you move through the remaining intents in the **Predicted intents** column.

Marking utterances as matched or mismatched to a VA intent



The utterances in the **ad_password_change** intent are easy to mark as a match to its predicted intent because there's at least one utterance within the intent that has the word password in it.

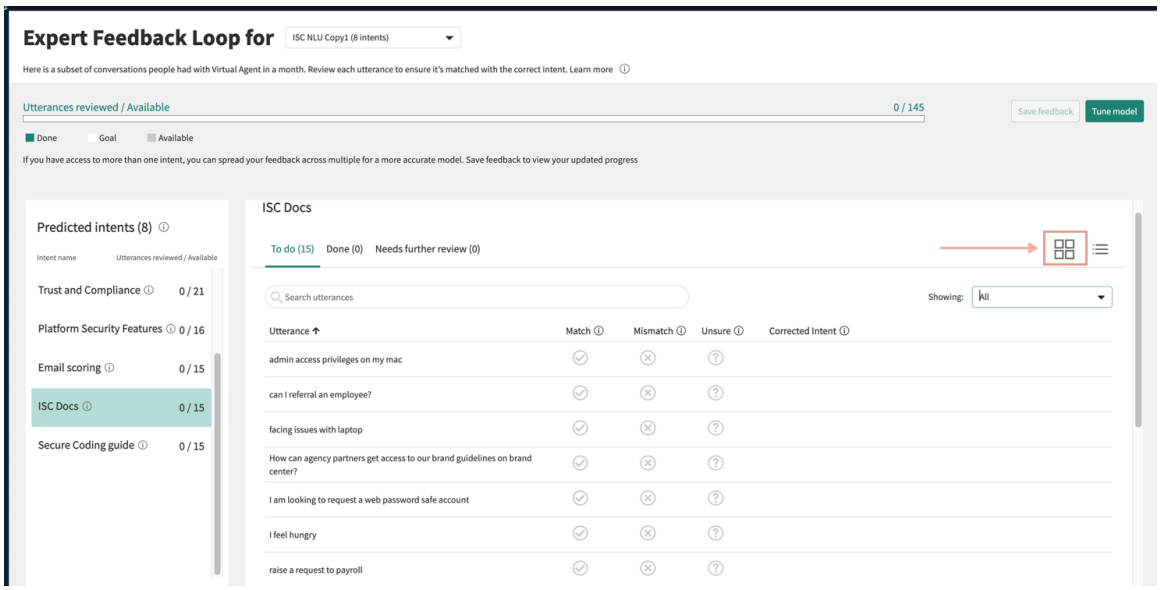
When you have finished reviewing the utterances in the predicted intent and you click **Save feedback**, the screen refreshes to highlight the next predicted intent in the model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

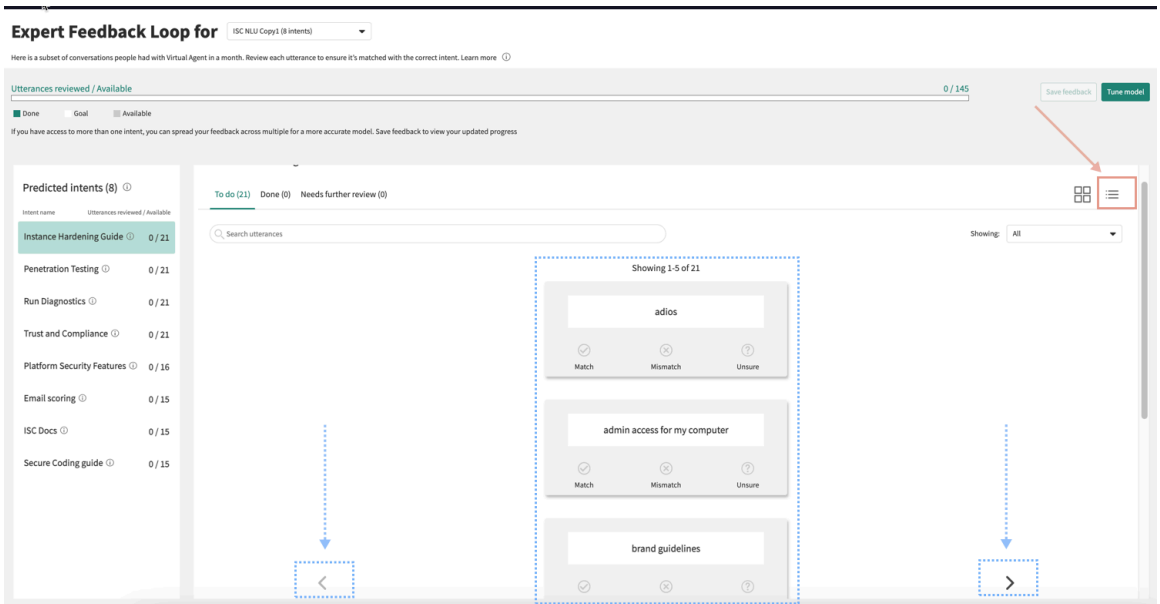
If you want to see your utterances grouped by their value, click the **Showing** filter and select one of the options in its prompt. The values are **All**, **NLU_Match**, and **Unsure**.

Using the card view option to review your utterances

Your Expert Feedback Loop utterances appear on the **To do** section in a list view by default, as shown in many of the images you've seen so far above. However, you can also choose a different view that shows each utterance framed within a card. If you choose this option, each utterance appears on the page in groups of five. While still within the list view, click the **card** icon.



Result: The card view replaces the list view. Note the left and right arrows on either side of the page of utterance cards. When you click the right arrow, it turns the page to the next set of five cards. When you click the left arrow, it turns the page to the previous set of five cards. If you want to switch back to the list view instead of the card view, click the **list** icon as shown in the upper right hand corner of the image below.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Changing an utterance value

If you mark an utterance but then change your mind on the value you chose, you can resolve the issue by clicking on **MisMatch** and selecting another intent. Click **Save feedback** to save the change.

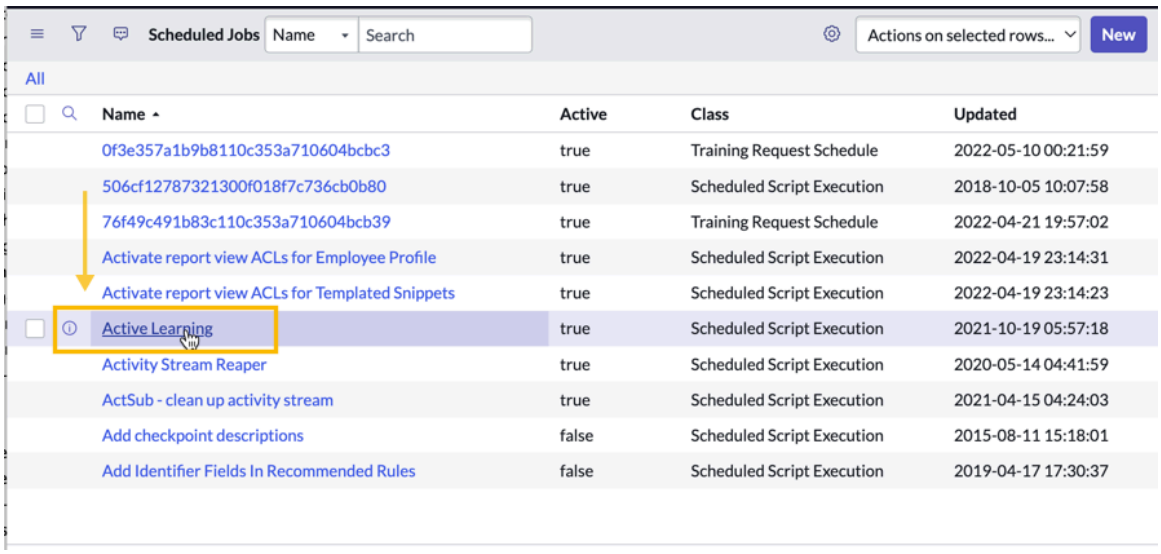
Unsaved feedback

If you finish and log out of a session of marking your utterances, but forget to save your changes, select **Unsaved changes** from the **Showing** prompt. This action displays all the utterances you have given feedback for but haven't saved yet.

Procuring additional VA feedback data on demand

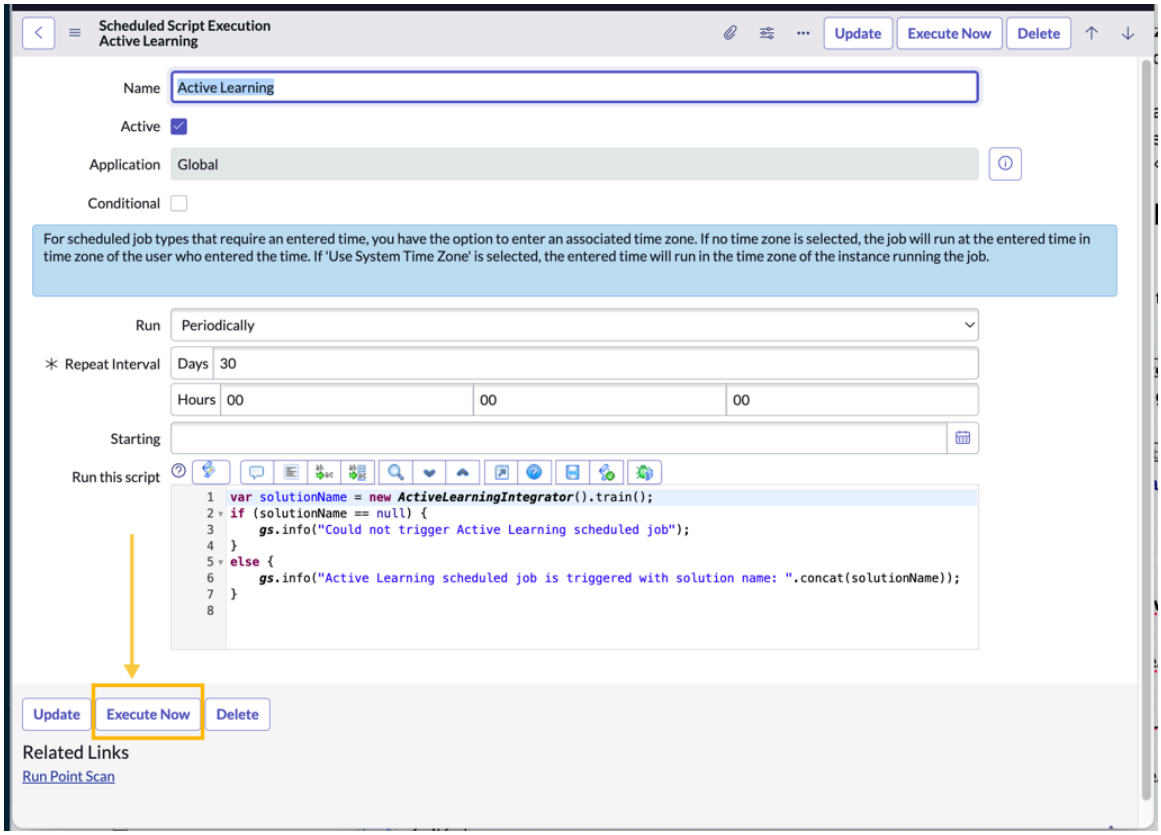
To get more data from the VA chat logs to use in the Expert Feedback Loop, take the following steps.

1. Use the nlu_admin role and navigate to **All** > **System Definition** > **Scheduled Jobs** > **Active Learning**.
2. Click **Active Learning**.



3. Click **Execute Now**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. Increase or set the values in the following four NLU system properties.

- *glide.mlpredictor.option.nlu.activeLearning.va_chat_logs.max_row_limit - 3000*
- *glide.mlpredictor.option.nlu.activeLearning.label_candidate_table.max_data_ - 10000*
- *glide.mlpredictor.option.nlu.activeLearning.label_table.max_data_size - 10000*
- *glide.mlpredictor.option.nlu.activeLearning.label_candidate_table.max_respo*

To see how these properties work within the context of other NLU properties, see NLU Workbench properties.

Reviewing uncategorized utterances

As part of the feed from VA chat logs to the Expert Feedback Loop, the system collects and displays in a list any utterances in your instance that aren't part of a VA intent. You access these utterances by selecting **Uncategorized** in the **Expert Feedback Loop for** prompt at the top of your screen.

When the screen refreshes, these utterances are shown in the **Utterance** column of the **Uncategorized utterances** section of your screen.

Important:
It is extremely important to provide feedback on this set of utterances because the system is suggesting there is no associated intent for these utterances. By confirming this lack of association or by associating these utterances with an existing intent, you are helping the model to learn correctly.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

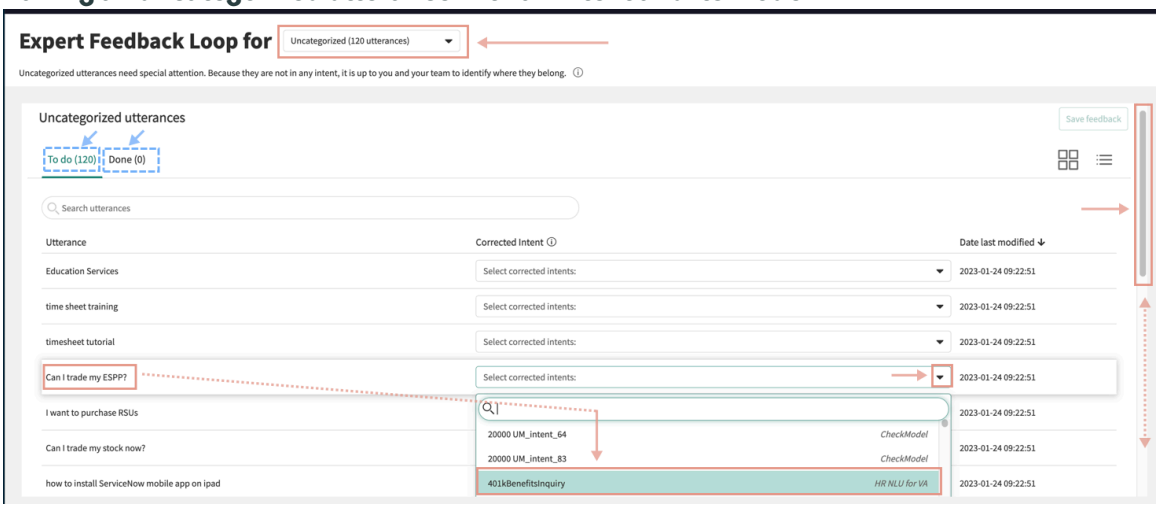
Here are the steps you use to connect an uncategorized utterance to an intent:

1. In the **Utterance** column, select an utterance from the list.
2. In the **Corrected Intent** column, search and select an intent and its model from the prompt that you think is the best match for the utterance.

For example, a good pairing for this connection is the Can I trade my ESPP? utterance from the **Utterance** column and the 401kBenefitsInquiry intent from the **Corrected Intent** column, as shown in the image below.

Note:
The **Corrected Intent** column enables you to search for and use all intents across all models in your instance.

Pairing an uncategorized utterance with an intent and its model



Regarding the image above, let's cover a few items that help describe how the uncategorized utterances page works, as follows. The **To do** section collects utterances that must be connected to an intent. On this particular day we see there are 120 such utterances. Note that the **Done** section on the screen has no utterances because you've not taken any action yet.

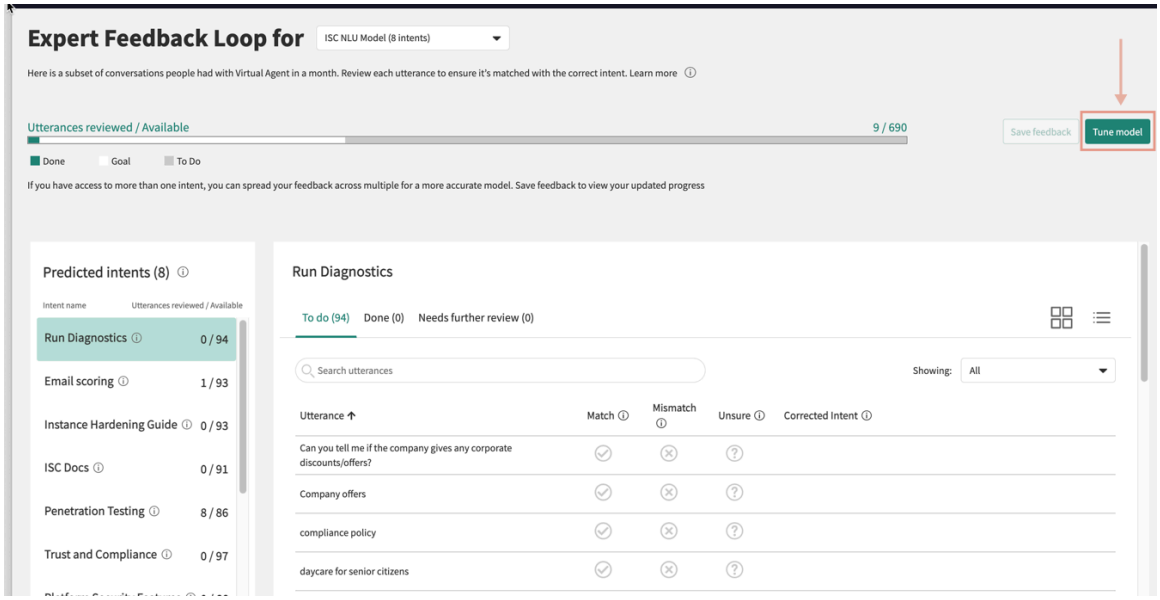
Note also that there's a vertical bar to the far right of the screen where you can scroll up and down to navigate the list of utterances. At the bottom of this scrolling screen, there are arrows that point left and right so you can to turn to the next or previous page of the full utterance list for the day.

Note also that when you do take action to make a connection, the **Save feedback** button, normally active by default, deactivates because the system recognizes your action and automatically saves your change.

Expert Feedback Loop data in the Tune Model phase

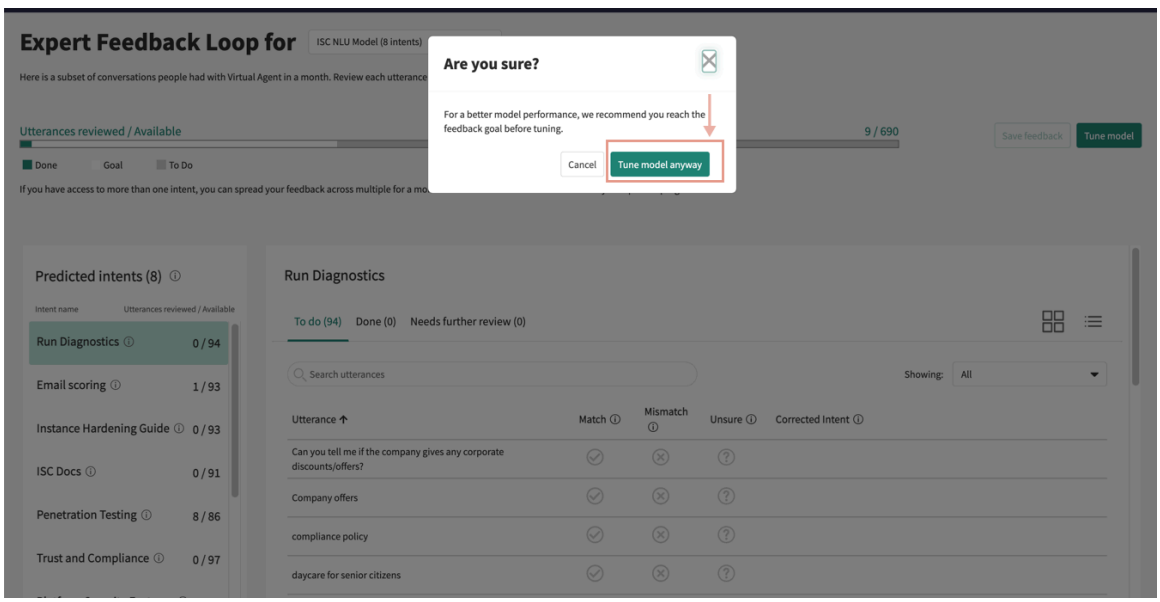
The **Tune model** button in the Expert Feedback Loop is always enabled and supports the experience of moving utterances from the Expert Feedback Loop **Done** tab to the model and its test set. Use the nlu_admin role to click this button any time you need to tune or retune your model.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



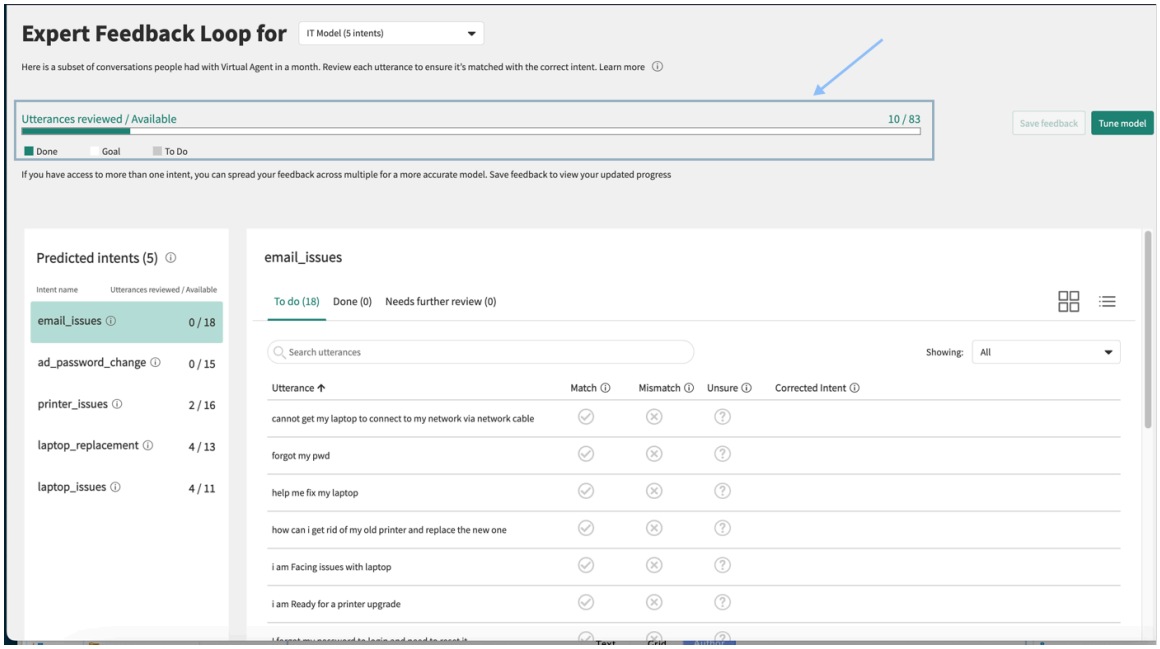
After reviewing utterances in the Expert Feedback Loop, you can push a portion of the feedback data to the default test set of your model. These utterances are then directly added to the **Test utterances** tab of your model. This helps you to continuously maintain and update your test set with real end user utterances. The system tracks the source of the test utterances for visibility into whether they came from the Expert Feedback Loop or from another source.

If you click the **Tune model** button before you have marked and saved at least 100 utterances, the screen refreshes to remind you that your model performance increases in quality when you have passed the 100 utterance feedback goal, as referenced in the text in the image below.



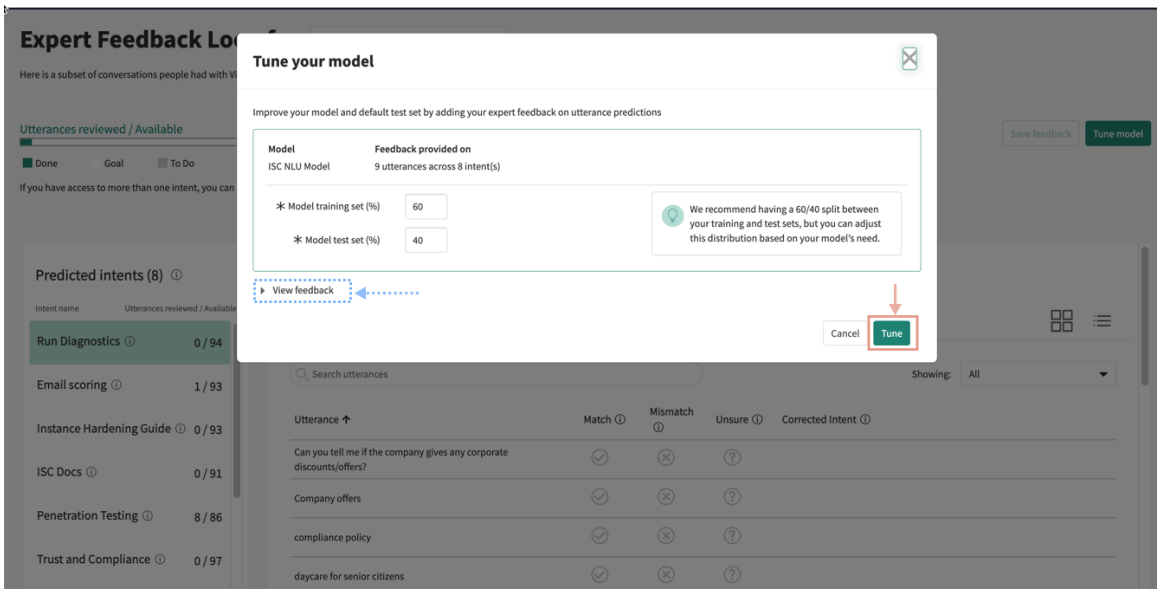
As you keep marking and saving your feedback data, you can see the ratio of reviewed utterances (in green), and those that are still available for review (in white) on the progress bar image below.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



On the **Tune your model** image below, you can see the default 60/40 percentage split for your model training set and your model test set. You can adjust these default values if needed by selecting your own chosen numbers in each of the two **(%)** boxes. Once you're satisfied with your percentage split, click the **Tune** button to tune your model based on the percentage values you chose.

Note:
If you click the **View feedback** caret, it takes you to an Expert Feedback Loop screen where you can continue reviewing, marking, and saving your feedback utterances.



Using the Active Learning job

Scheduled jobs, also known as batch jobs or batch scheduling, are automated pieces of work that are performed at a specific time or on a recurring schedule. Many jobs run on schedules but we're focusing now on the Active Learning job as an example.

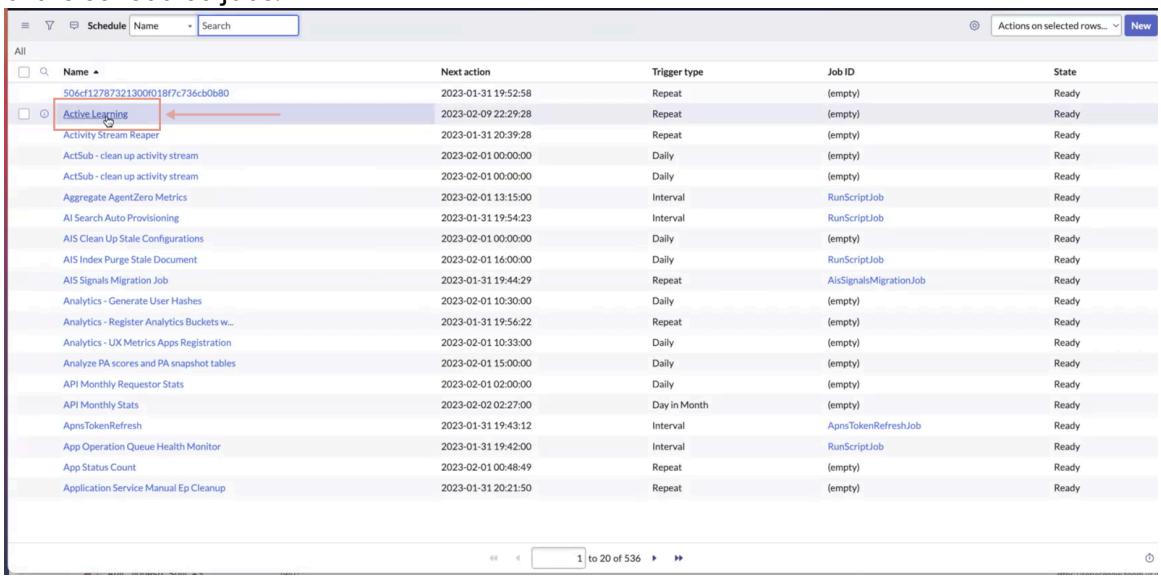
Here are three things you can do when using the Active Learning job.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

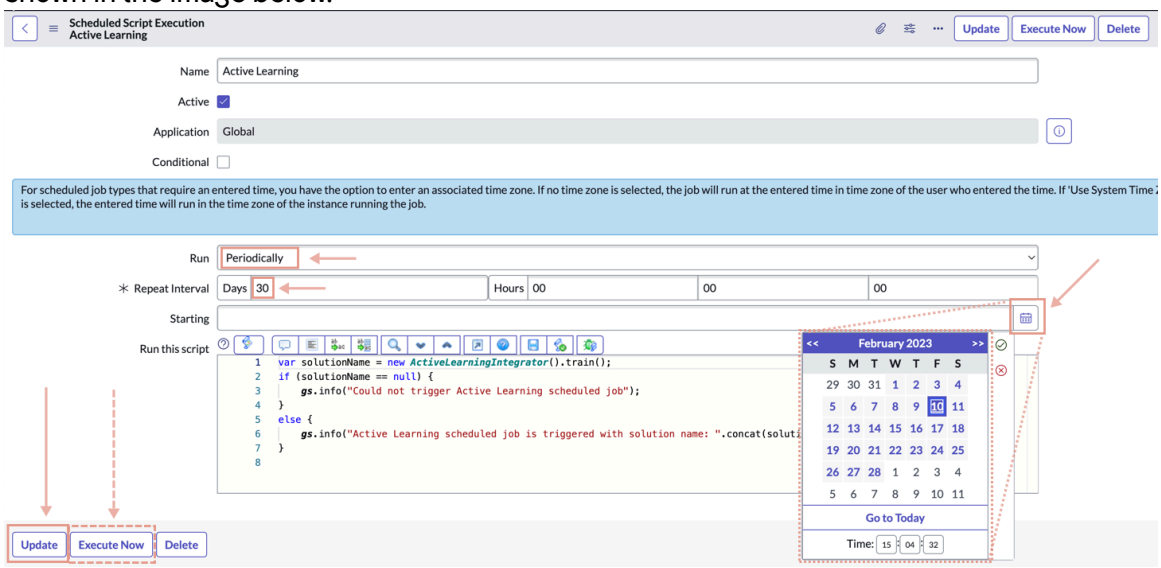
- Change the repeat frequency interval with which the Active Learning job will be set to run.
- Check when the next scheduled run is set for the Active Learning job.
- Execute the Active Learning job whenever you want (on demand).

Here are some steps to use to get started.

1. Using the nlu_admin role, navigate to the **All** field and type **sysauto_script.list**, then press the **return** key on your keyboard. The screen refreshes to show the Schedule page, which lists all of the scheduled jobs.



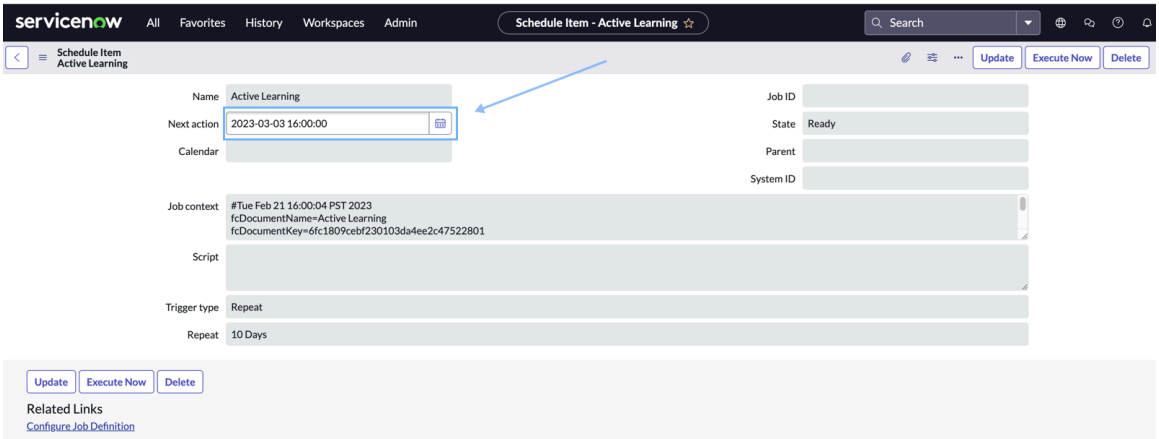
2. On the Schedule page, click **Active Learning**. A record for the Active Learning job appears, as shown in the image below.



3. On the Active Learning record, configure the following fields.
   - In the **Run** field, select one of the eight Repeat Interval options for the job. The options are: Daily, Weekly, Monthly, Periodically, One, On Demand, Business Calendar:Entry Start, and Business Calendar:Entry End.
   - In the **Repeat Interval** field, enter the number of days you want between now and the next repeat interval for the Active Learning job.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- In the **Starting** field, click the Calendar icon and choose the day and time for the next Active Learning job interval.
- Click the **Update** button to save your configurations or click the **Execute Now** button to start the Active Learning job run.
- To verify when the next Active Learning job runs, navigate to the **All** field and type **sys_trigger.list**, then press the **return** key on your keyboard. The Schedule page appears. Click **Active Learning**. The Schedule Item/Active Learning record appears and populates the **Next action** field with the date and time for the next Active Learning job run.



Detailed information about NLU's Active Learning job can be found in this KB article on the Support portal: KB1633901.

Issue Auto Resolution Tuning in NLU

Use the NLU Workbench homepage to support Issue Auto Resolution (IAR) tuning in NLU.

A video introducing Issue Auto Resolution tuning in the NLU Workbench

Summary usage and roles

Use the nlu_admin or admin role to access IAR Tuning in the NLU Workbench. IAR Tuning in the NLU Workbench requires at least the nlu_feedback_admin role. Note that the nlu_admin role contains the nlu_feedback_admin role. Also, the virtual_agent_admin role contains the nlu_admin role.

If you click on the IAR model name, you will be taken straight to tuning, in the product. You will not be taken to a model overview page, so this behavior differs from Virtual Agent or AI Search models in the NLU Workbench.

The IAR Tuning workflow

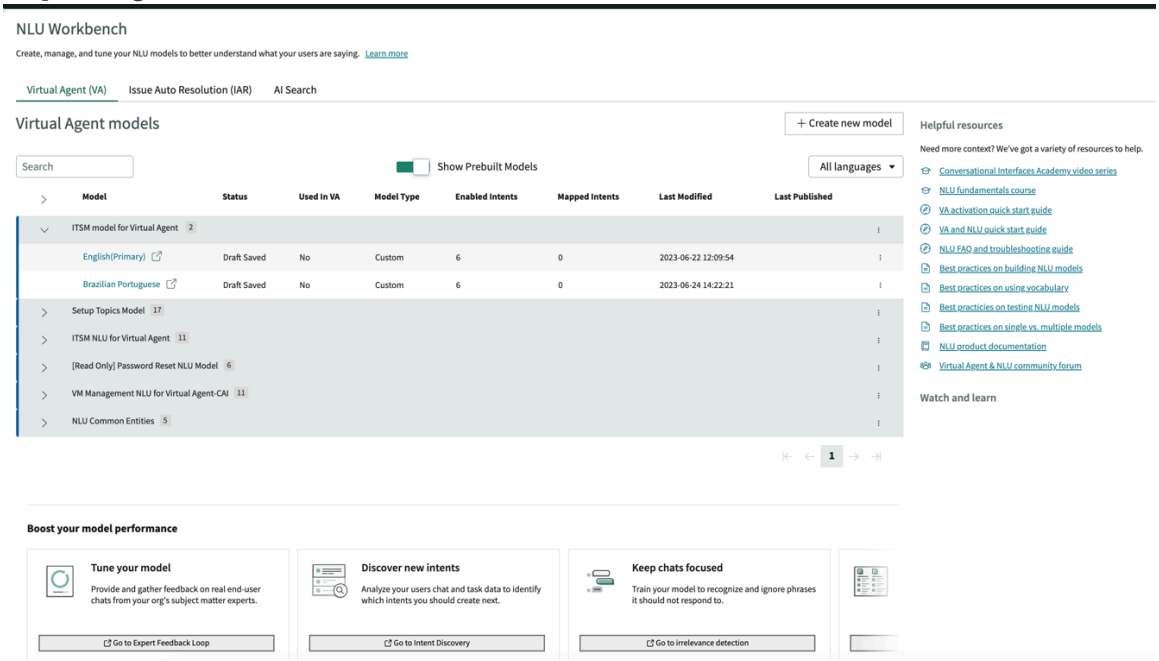
IAR admins begin their model tuning journey in the IAR Admin Console, and then land in the NLU Workbench to tune their ITSM model. If they haven't trained the ITSM model in the console yet, the workflow sends them to the Expert Feedback Loop documentation under the **Boost your model performance** section of the NLU Workbench.

How IAR models differ from NLU models

Unlike the Virtual Agent and AI Search tabs, the IAR tab doesn't use a Create new model button. The IAR-ITSM model that IAR admins use is a prebuilt model. IAR models can't be moved using update sets.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Exploring the NLU Workbench



At the top of the NLU Workbench page are three tabs that group Virtual Agent, Issue Auto Resolution, and AI Search models separately. Below those tabs are a list of models colored grey. In the Model column of the list, if you click the caret to the left of the model name, the model changes color from grey to white and opens to show the model's languages; status; usage; model type; number of enabled intents and mapped intents and the date when the model was last modified or last published.

Issue Auto Resolution tuning options

When you are tuning your Issue Auto Resolution model in NLU Workbench, you can adjust the output for several goals: precision, automation, or a balance of the two. Compare how your choice of tuning options affect match rate and coverage, before committing.

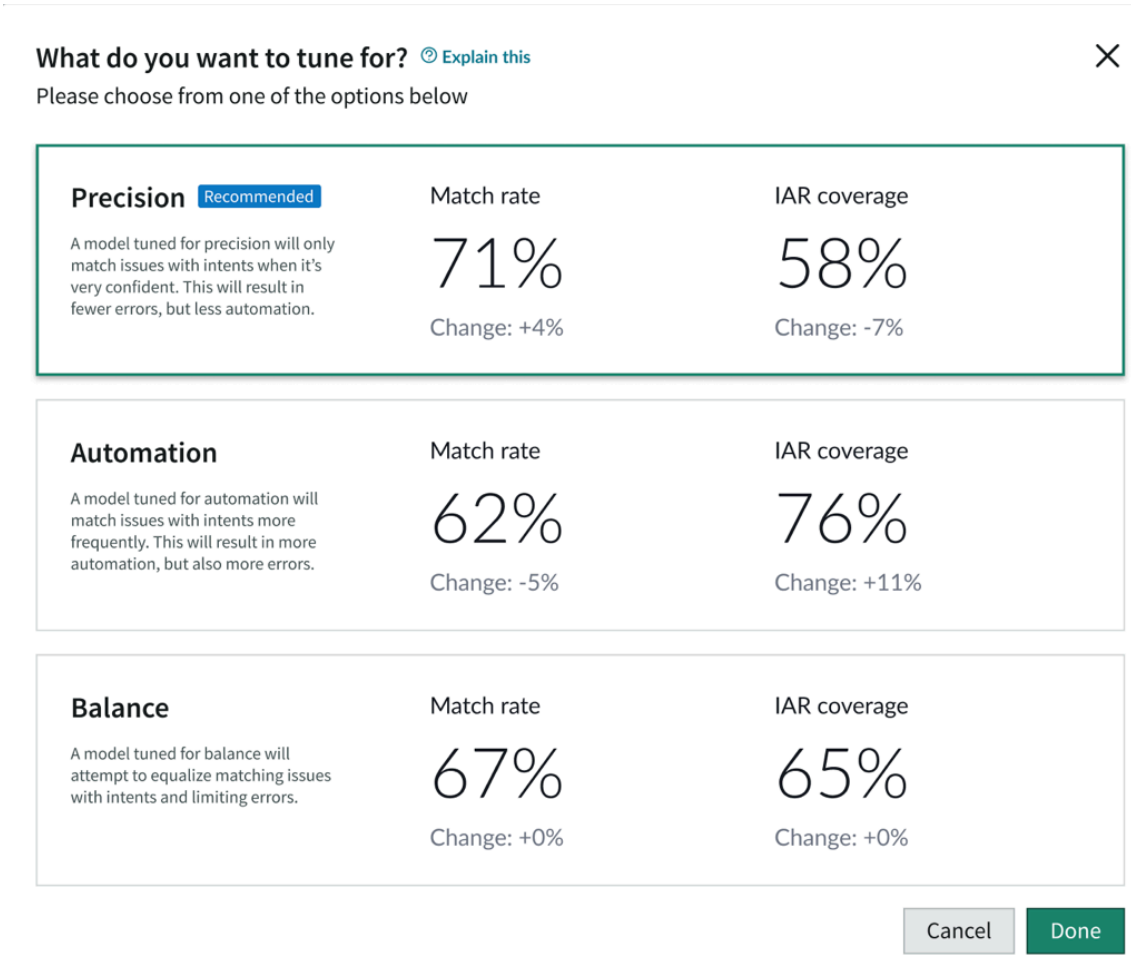
Summary usage

By default, Issue Auto Resolution tuning in NLU Workbench optimizes for precision. Depending on your business requirements, you can also tune the model for other objectives. In the **Analyze** step of Issue Auto Resolution tuning, the tuning goals list enables you to adjust for **Precision**, **Automation**, or **Balance**. As you select one of these options, the projected Match rate and IAR coverage percentages change accordingly, so you can compare possible outcomes.

To access the **Analyze** step of IAR tuning, use the nlu_admin role and navigate as follows.

1. Navigate to **All** > **NLU Workbench** > **Models**.
2. Select the Issue Auto Resolution tab, then select the model name. The tuning experience opens to step 1 (Feedback) initially.
3. Provide feedback, then select the **Analyze** button. Step 2 (Analyze) opens.
4. In the section **Here are your tuning options and projected results**, using the list **You can tune for precision, automation, or balance**, select options to see projected scenarios. You can also select the link **Learn about tuning goals** to open the following window.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Precision

When tuned for precision, the IAR model makes predictions only when its confidence is relatively high. This results in lower error rates, but also in fewer incidents resolved.

Precision is the recommended tuning option for the IAR ITSM model, so this option is selected by default.

Automation

When tuned for automation, the IAR model makes predictions at a lower confidence threshold. This results in more predictions, so more incidents are resolved. However, higher error rates are possible.

Balance

When tuned for balance, the IAR model attempts to strike a balance between precision and automation.

Match rate

The match rate is defined as the number of Incidents where the intent was predicted correctly, divided by the number of predictions for that intent. This ratio is averaged across all intents except for **NO_INTENT**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

IAR Coverage

Coverage is defined as the percentage of Incidents that would be resolved because the model was able to make predictions above its confidence threshold. The predictions may contain some errors.

Using tuning options

Select several different tuning options to compare projected results. Depending on the option you select, the system presents scenarios for projected Match rates and IAR coverage rates. Also the system displays how much these rates change according to your selection.

Review further information in the **Here's a detailed breakdown** section of Analyze. Here you can drill down into results that are specific for each intent in the model.

Note that the intents are grouped into mapped and unmapped intents, depending on whether they have been mapped to Virtual Agent topics. After providing feedback in IAR Tuning, you may wish to activate some intent-to-topic mappings. To do so, expand **See unmapped intents**, then select the **Map more intents** button. This opens the IAR Admin Console.

When you have decided the optimum tuning option for your requirements, select the **Save choice** button in the **Learn about tuning goals** window. Then, select the **Tune and publish model** button to advance to the next step.

Intent Discovery

Use the Intent Discovery application to help identify opportunities for incident deflection. For example, you can use it to identify which Virtual Agent conversations to activate next.

Summary usage

For applications that consume NLU, such as Virtual Agent and AI Search, Intent Discovery helps you to better understand which prebuilt intents you can benefit from, and which custom intents would be useful to create.

Intent Discovery provides an analysis that you run on historic incident data or other task data. You can also group the run's remaining records into different clusters so you can manually add utterances to NLU intents. In addition, you can use specific clusters to create new intents in a model.

In this example scenario, you're using Intent Discovery to identify the top intents in your instance, and how much coverage they can provide across your historic incident records.

Installation

Intent Discovery is available from the ServiceNow Store. For more information, see Install Intent Discovery.

After Intent Discovery is installed and activated, it appears under **All** > **NLU Workbench** > **NLU Advanced Features**.

Note:
Although organized under NLU Advanced Features in the navigation pane, Intent Discovery is a separate application that is not included when installing NLU Workbench - Advanced Features.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Intent Discovery report details

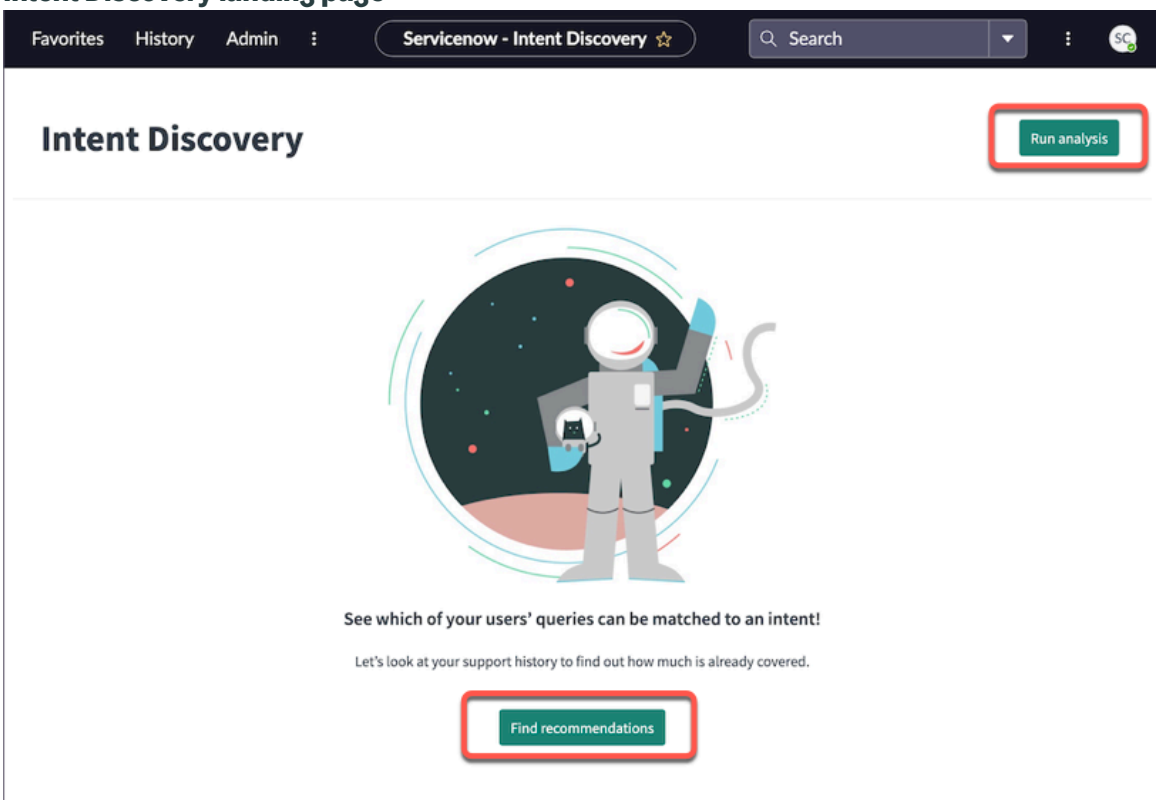
- When **Taxonomy** is selected, the generated report contains intent recommendations against the selected taxonomy. A taxonomy is a prebuilt library of intents in a specific domain. While you don't have access to the underlying intents, when you run Intent Discovery against a specific taxonomy, data that maps to any intent in the taxonomy will be identified.
- **Unmatched records** are the utterances which couldn't match to any intent in the taxonomy.
- **Recommended intents** are the intents which are found from utterances that data was run on.
- The percentage of **Unmatched records (clustered)** are the records that aren't classified (records that don't belong to any of the recommended intents).
- The percentage of unmatched records and the number of recommended intents don't need to match. It's a coincidence if they match.

Creating an Intent Discovery report

1. Using the admin or nlu_admin role, navigate to **All** > **NLU Workbench** > **NLU Advanced Features** > **Intent Discovery**.

2. Select either **Run analysis** or **Find recommendations**.

Intent Discovery landing page



Running an analysis on the report

1. For this example report, you configure the following fields on the Intent Discovery > Create new screen.

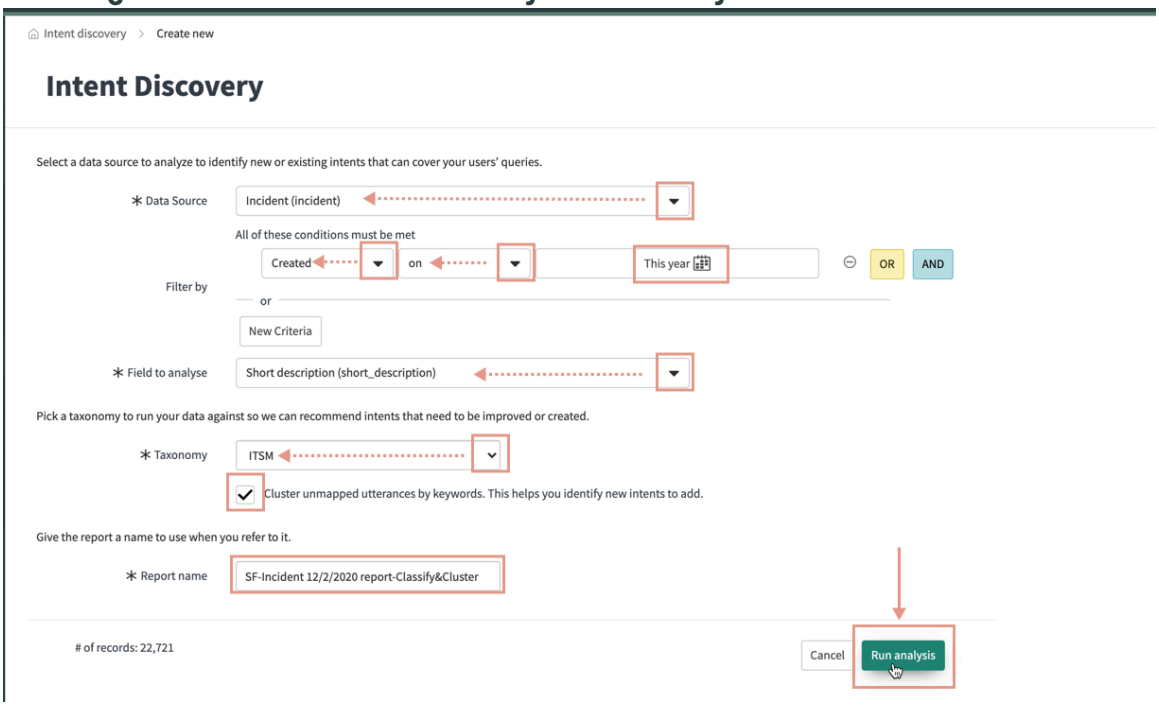
- Data Source: Select the **Incident (incident)** table.
- Filter by: **[Created] [on] [This quarter]**

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Field to analyze: **Short description (short_description)**. You choose Short description because it's a highly used string field that references words that can help the system identify an intent.
- Taxonomy: Select **ITSM**. This field tells the system to run classification processing on your ITSM incident records. It has 3 options: Classification, ITSM, or blank, which defaults as Classification.
- Cluster unmapped utterances by keywords... : Select the **check box**. When you check this box, the system groups your incident records that weren't classified into clusters.
- Report name: The field automatically defaults to **Incident <month/day/year>**. You can edit the name if you prefer. In this example scenario, you enter `Incident 12/16/2020 - SF Test`.

2. Select **Run analysis**.

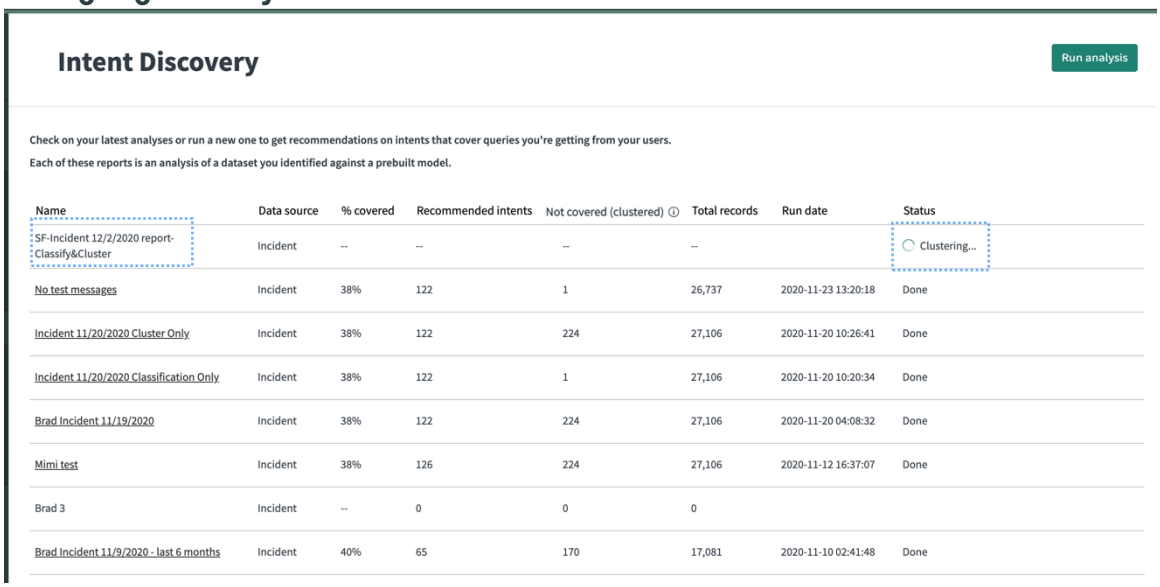
Selecting data sources in Intent Discovery for a run analysis



Result: Your report appears on the Intent Discovery screen, showing its status as the analysis begins. The subsequent status values appear in the following order during the analysis: Preparing to run, Work in progress, Clustering, and Done. This can take from 5 minutes to 30 minutes to complete. The fewer the records you have in a cluster, the less time it takes. Turning clustering off can also speed up the process.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

An ongoing run analysis

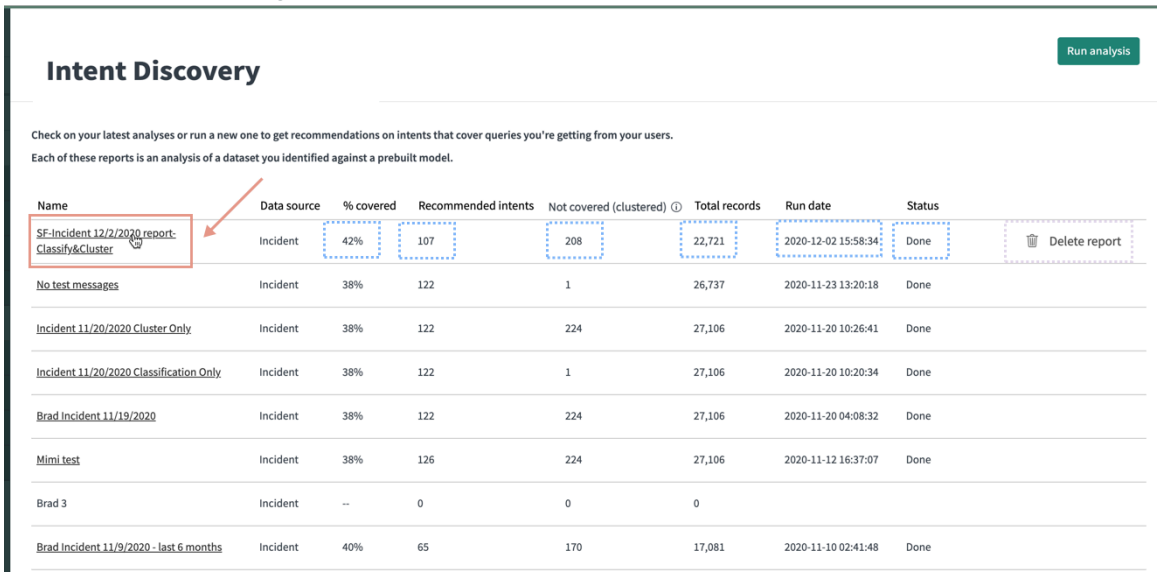


When the analysis is complete, the column values on the screen appear, with the **Status** column value set to **Done**, as shown in the image below.

Note:
If you want to delete the report and start over, point to the right of the Status column to invoke the **Delete report** icon.

3. Select the **Name** of your report.

A completed run analysis



Result: The screen refreshes, showing the analyzed incident records and the remaining incident records that were not classified.

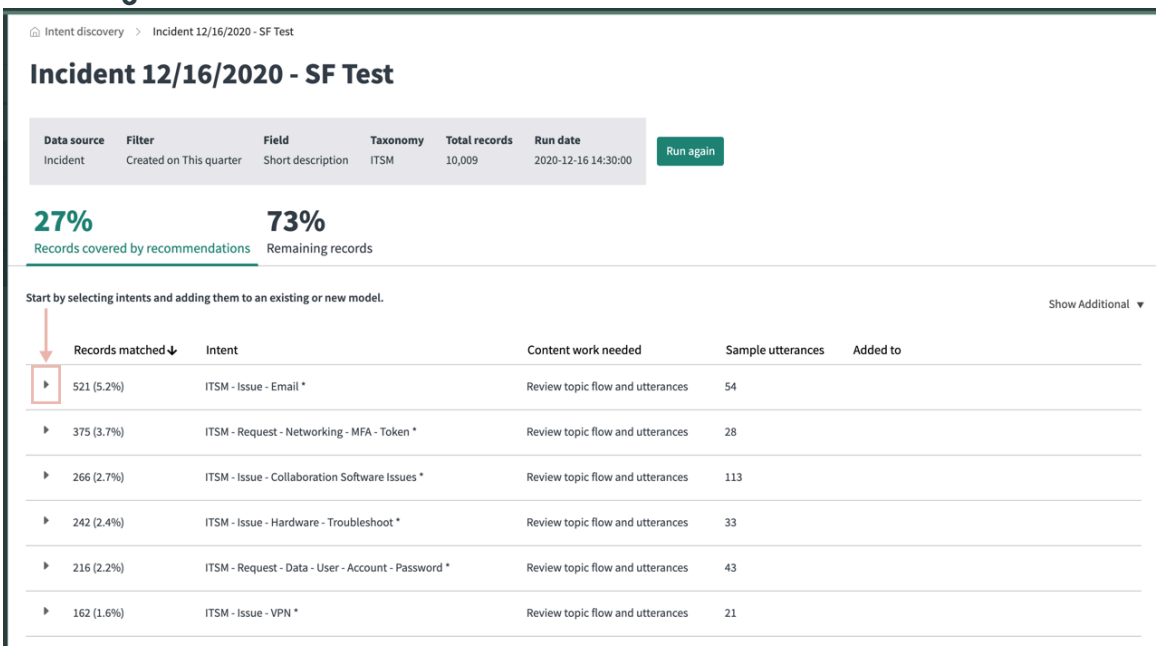
Importing recommended intents to new or existing custom models

Before importing intents to an NLU model, ensure that you are in the same application scope as the model. For more information, see .

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

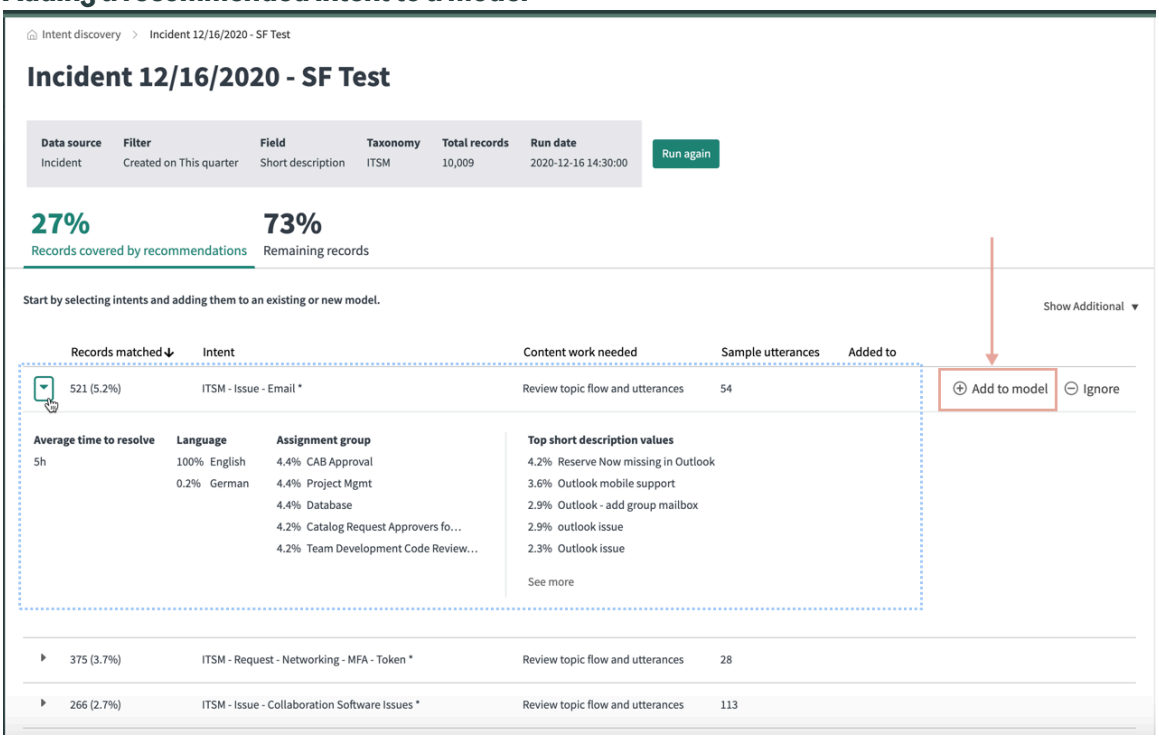
Reviewing a recommended intent



Result: The details of the recommended intent appear so you can review them, as shown in the image below.

2. Select **Add to Model**.

Adding a recommended intent to a model



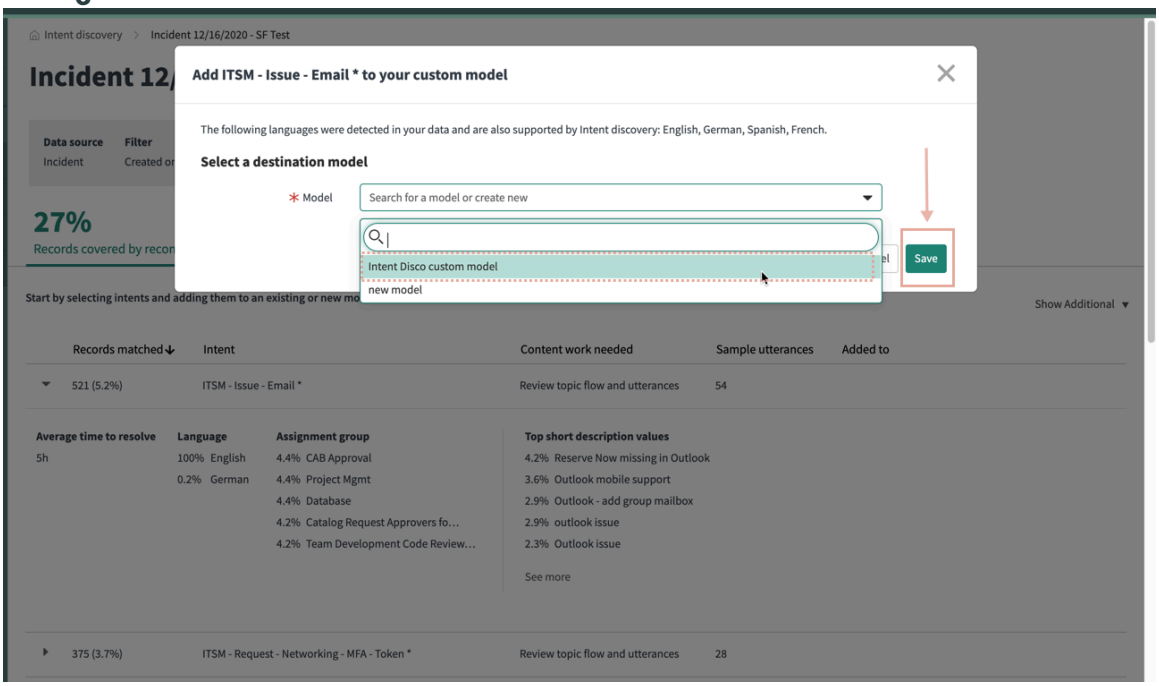
3. On the Select a destination model screen that appears, choose a model you want to add the recommended intent to. If you can't find an appropriate model, create a new one, return to the report, and add the new model.

Note:
The model you choose must have the same application scope as your current scope.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

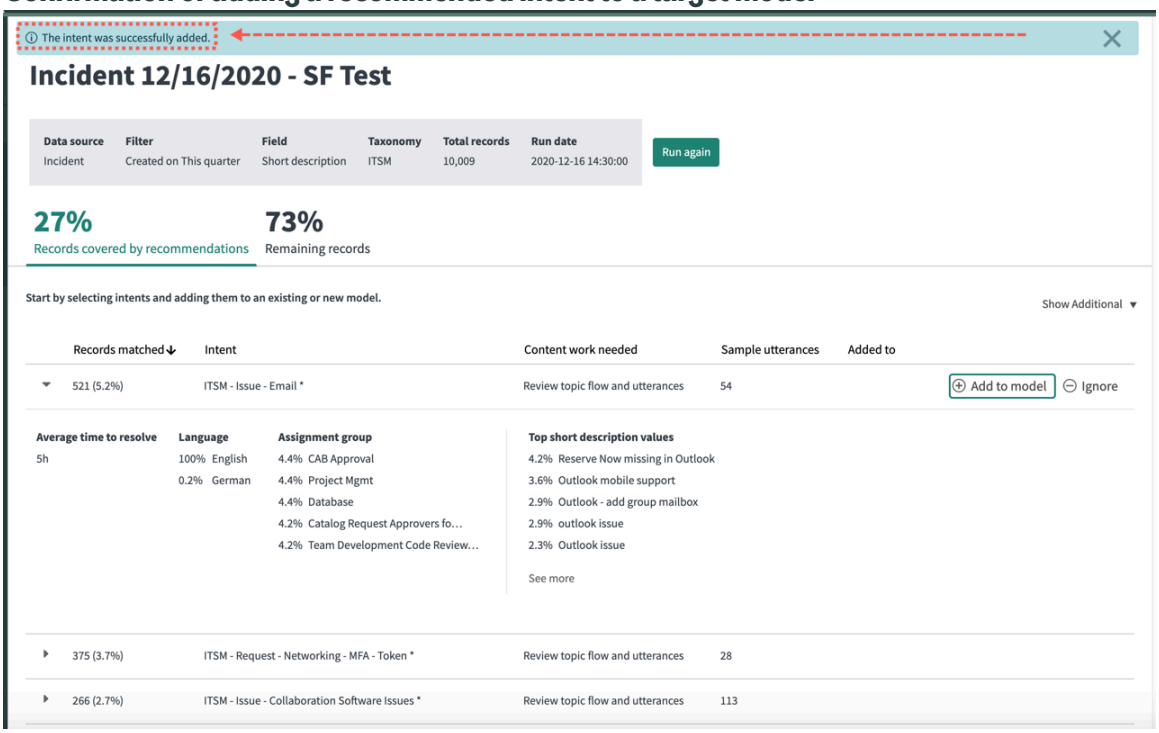
4. Select **Save**.

Saving a recommended intent to a model



Result: A banner appears on the screen, confirming the intent is added to the target model.

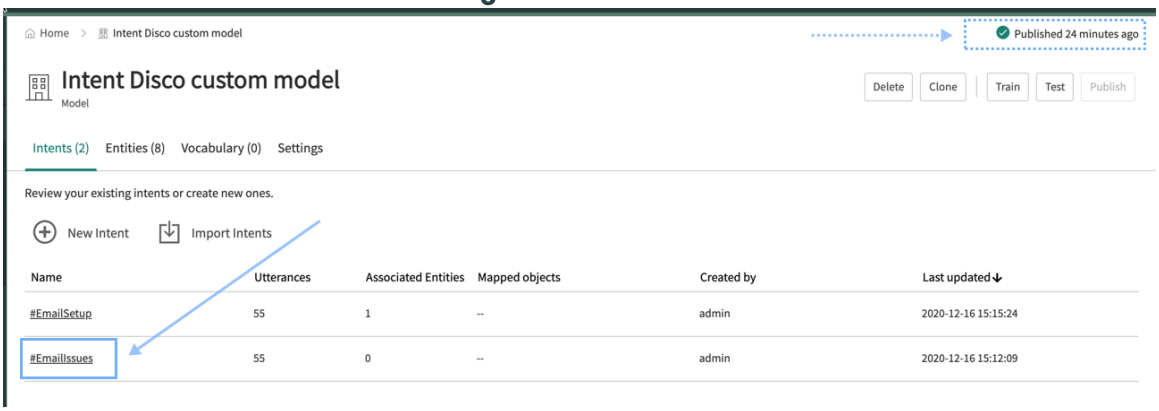
Confirmation of adding a recommended intent to a target model



The recommended intent also appears on the Model screen of the target model, as shown in the image below.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

View a recommended intent in the target model



Adding clustered utterances to an intent and its model

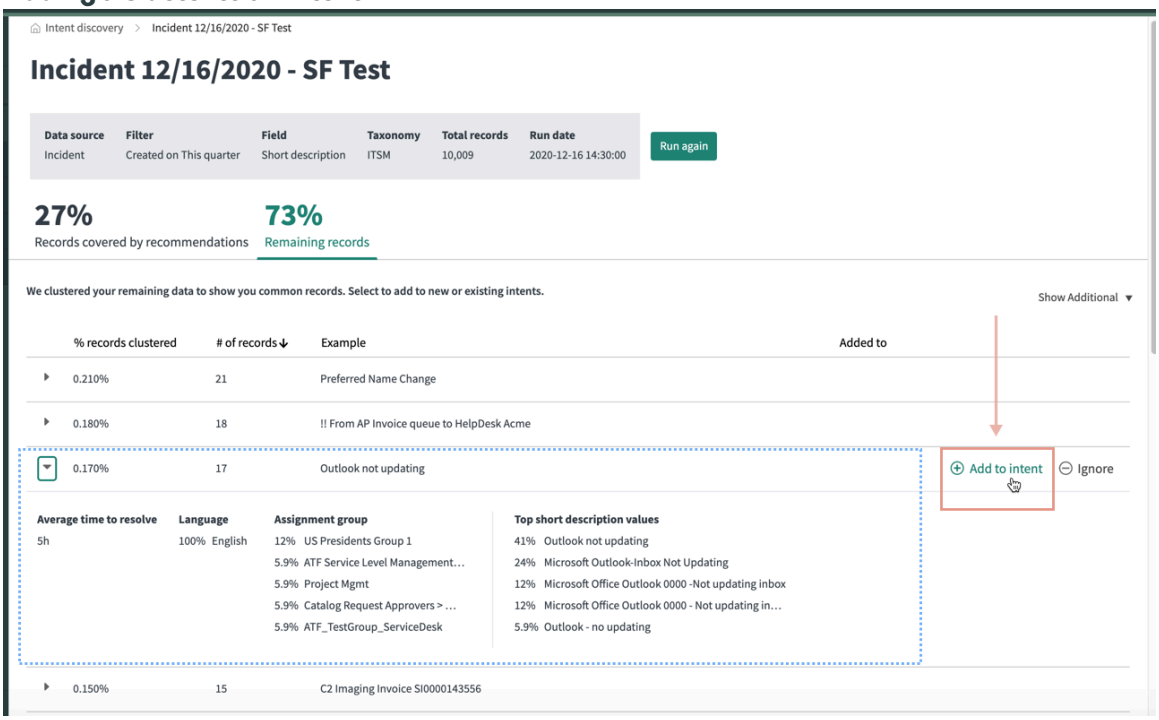
1. On the Remaining records section of the intent discovery records screen, select and open a cluster of utterance and short description data that you want to add to an intent and its associated model.

As you continue to build out new intents from these clusters, you can click the **Ignore** icon to remove any unwanted intents from the report.

There's also a **Show Additional** filter you can use to show or hide the added intents, and the ignored intents as well.

2. Select **Add to intent**.

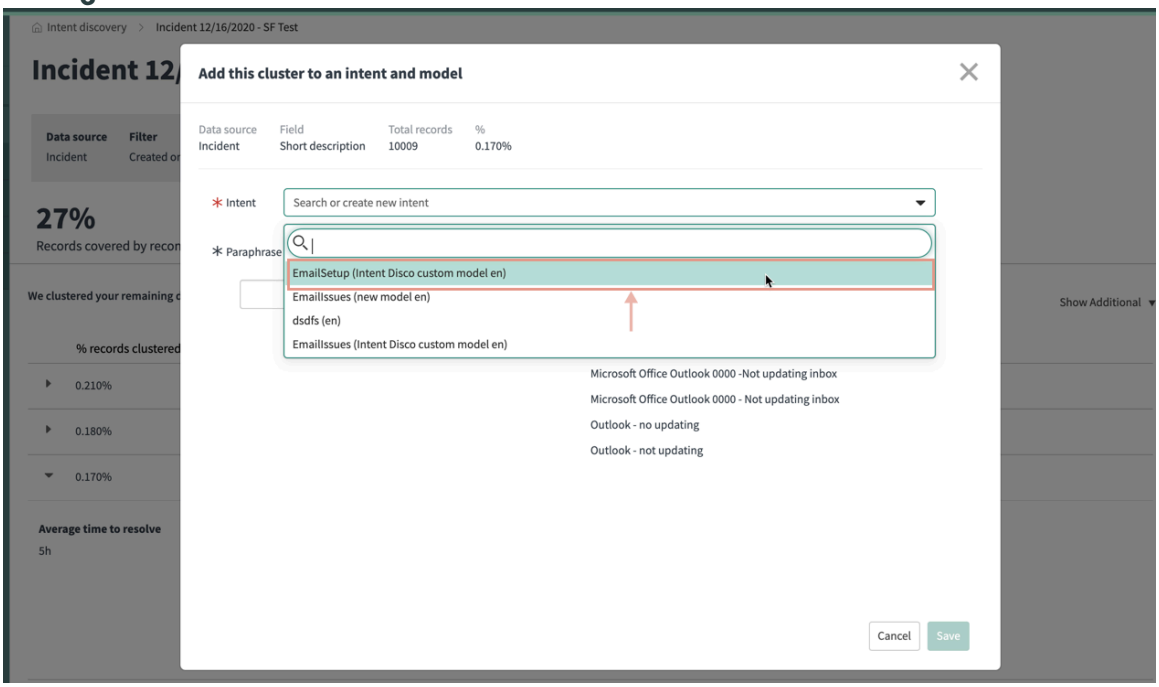
Adding a cluster to an intent



3. In the Add this cluster to an intent and model screen, select an intent and model pair you want to associate to this cluster.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

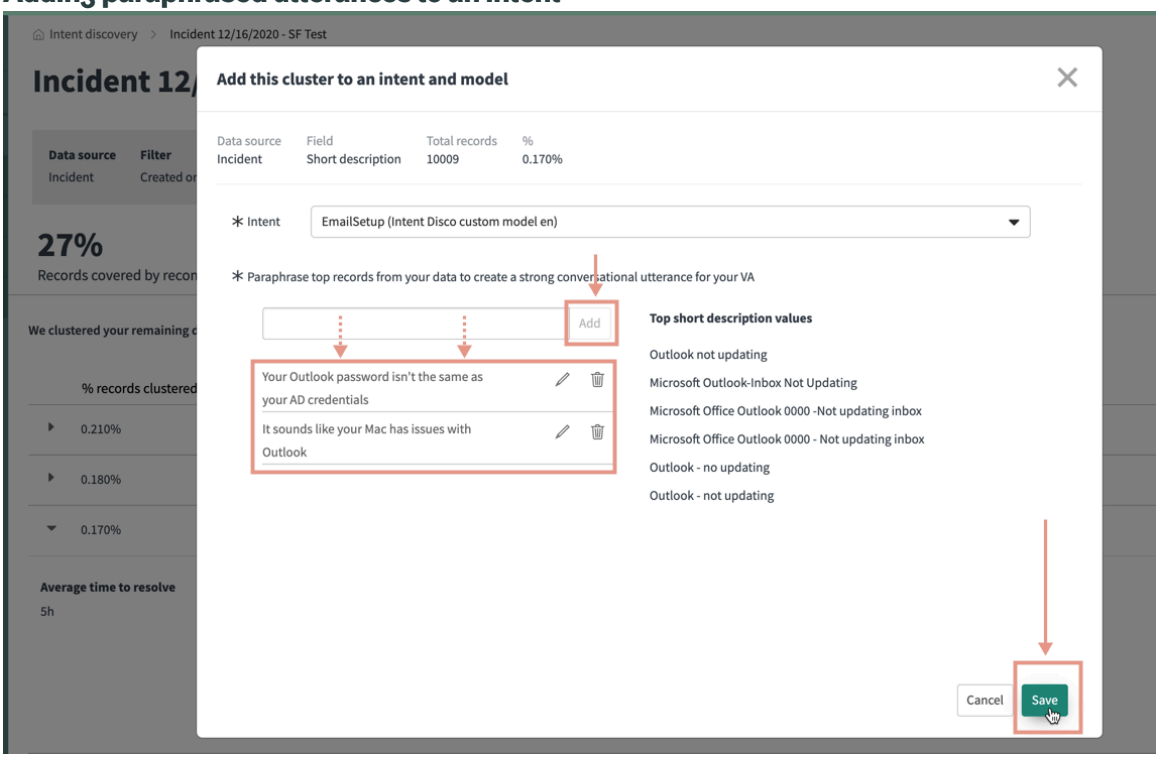
Adding a cluster to an intent and model



4. Enter a few utterance examples into the open text field. Select **Add** each time you complete your entry to save it in the system. Use the pencil icon or the trash can icon respectively to edit or delete your entry.

5. Select **Save**.

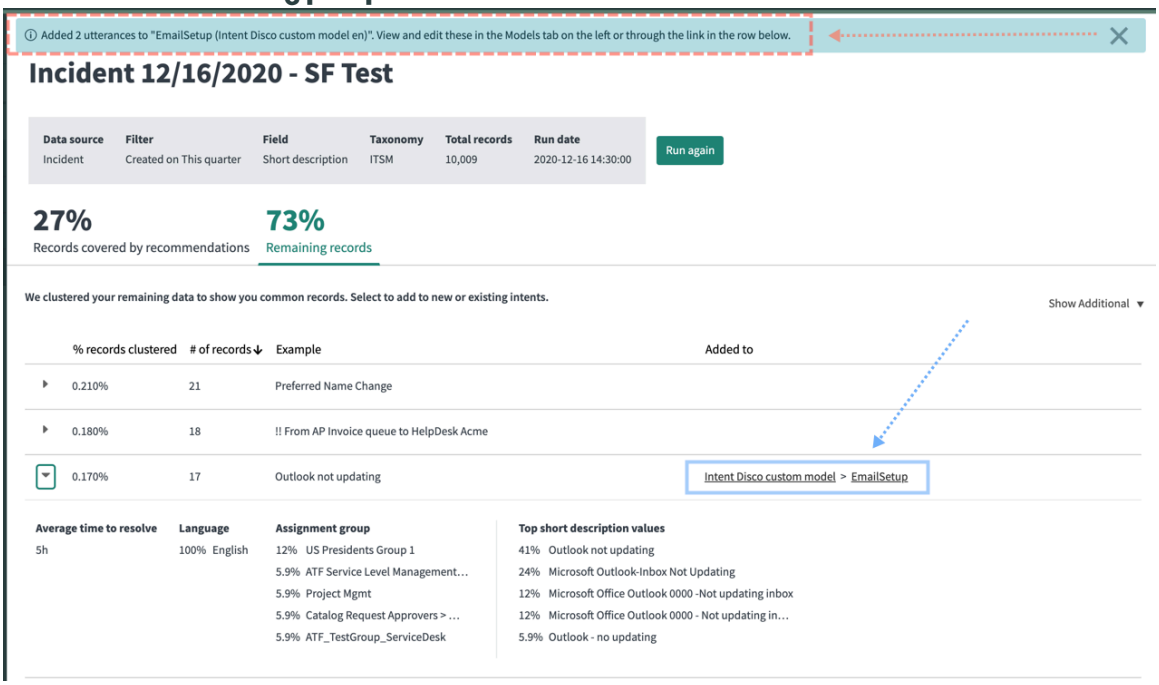
Adding paraphrased utterances to an intent



Result: The records screen appears, showing a banner confirming you added two new utterances to the target intent and its associated model. The model and intent pair appears in the **Added To** column, as shown in the image below.

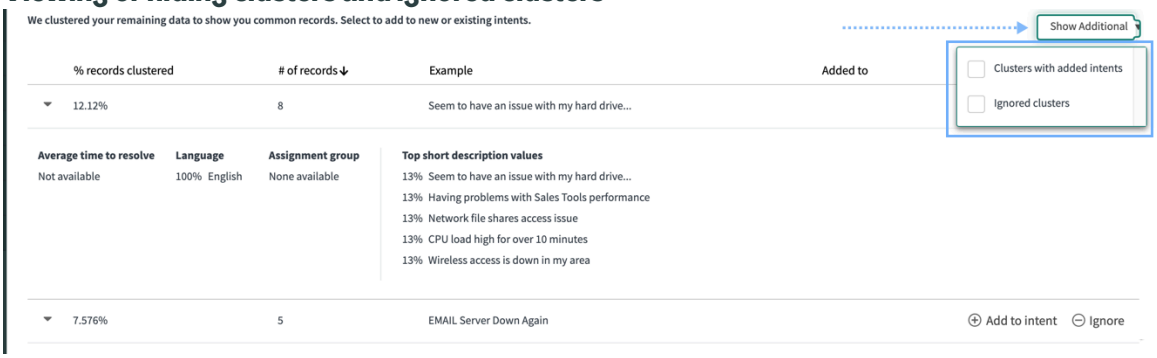
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Confirmation of adding paraphrased utterances to an intent



Use the **Show Additional** filter if you want to show or hide the clusters that have added intents, and the clusters that are ignored.

Viewing or hiding clusters and ignored clusters

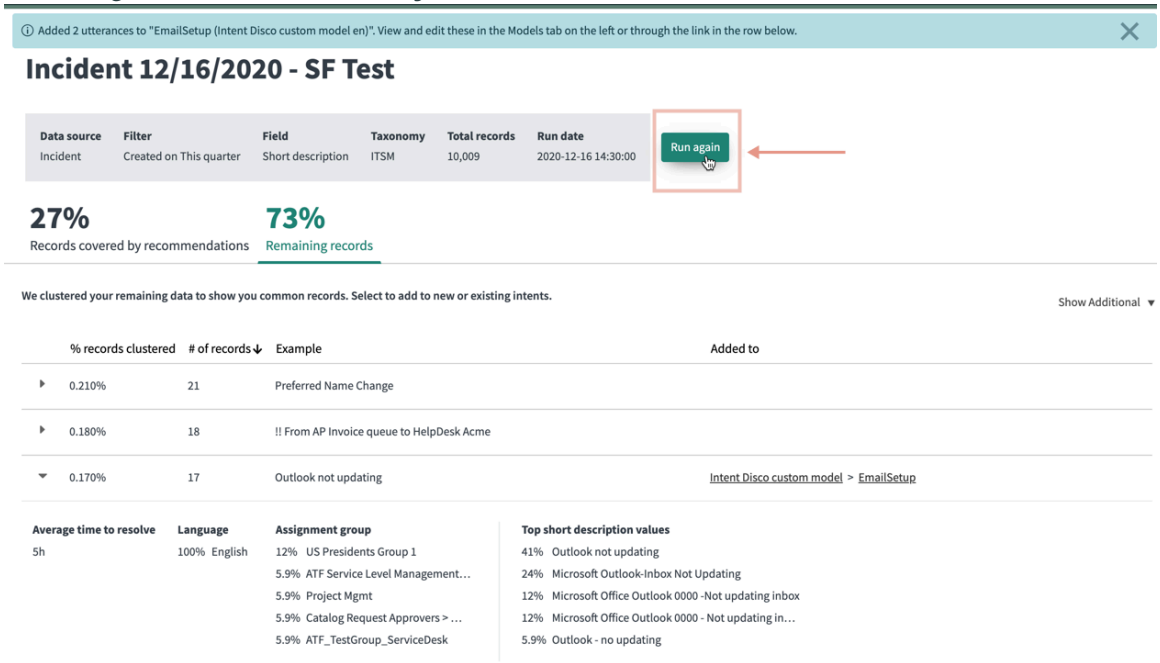


Running another analysis on your Intent Discovery report

1. Select **Run Again**.

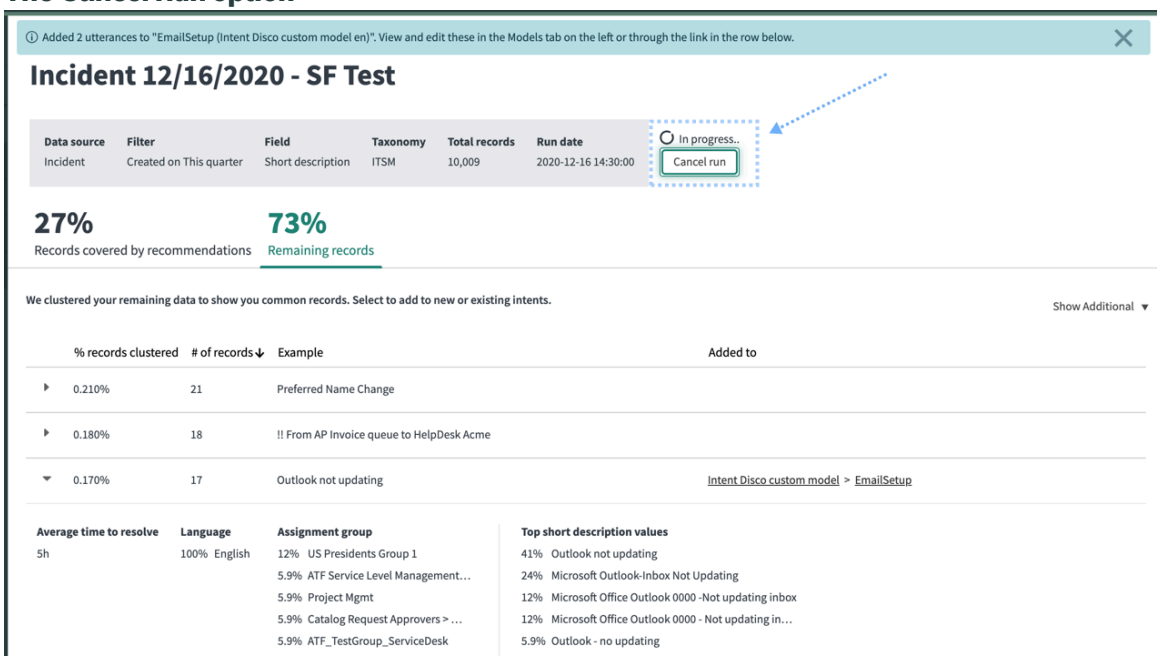
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Selecting the version of the analysis to run



Result: The new run begins. When it's in progress, the option to cancel the run appears, as shown in the image below.

The Cancel Run option

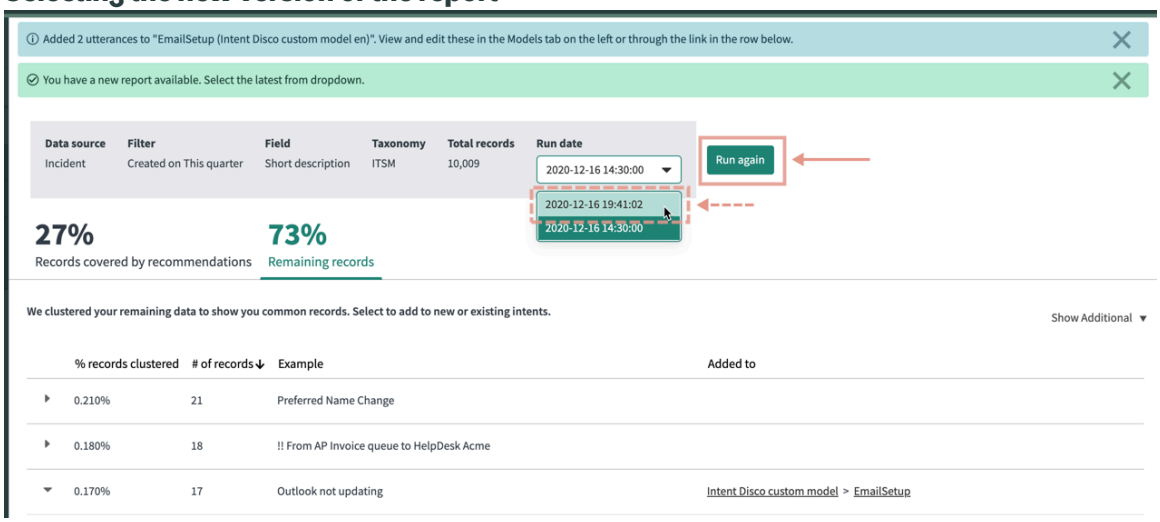


When the run is complete, a new banner appears that states you have a new version of the report.

2. Select the new version, then select **Run Again**.

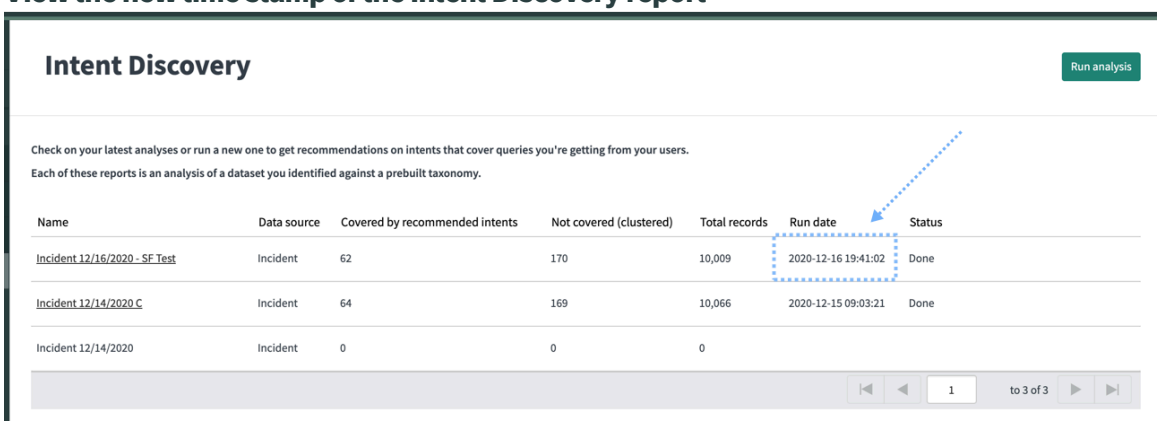
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Selecting the new version of the report



Result: The time stamp you selected for the most recent run appears in the **Run date** column of the Intent Discovery screen.

View the new time stamp of the Intent Discovery report



Install Intent Discovery

You can install the Intent Discovery application (sn_nlu_discovery) if you have the admin role.

Before you begin

- Ensure that the application and all of its associated ServiceNow Store applications have valid ServiceNow entitlements. For more information, see Get entitlement for a ServiceNow product or application.
- Review the Intent Discovery application listing in the ServiceNow Store for information on dependencies, licensing or subscription requirements, and release compatibility.
- Intent Discovery requires the following plugins. Ensure that these plugins are activated before you install Intent Discovery.

Required ServiceNow plugins

Predictive Intelligence (com.glide.platform_ml)

Enables various Predictive Intelligence and Machine Learning capabilities for training models. See Install Predictive Intelligence.

NLU Workbench - Core (com.glide.nlu)

Adds NLU Model capabilities. See Activate the NLU Workbench.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role required: admin

About this task

Tables are installed with Intent Discovery:

For more information, see Components installed with Intent Discovery.

Procedure

1. Navigate to **All** > **System Applications** > **All Available Applications** > **All**.
2. Find the Intent Discovery application (sn_nlu_discovery) using the filter criteria and search bar.

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

   Visit the ServiceNow Store 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
[sn_nlu_discovery_cluster]	
Discovery Report Definition [sn_nlu_discovery_report_definition]	Includes the necessary information for generating reports.

Natural Language Query

Natural Language Query (NLQ) enables you to query the data in your instance by entering plain language requests into the user interface.

Get started

Choose one of these tiles to get started.

Explore



Learn about NLQ concepts and features.

Use



Use NLQ to query your data with requests in natural language.

Configure



Set up NLQ for your environment.

References



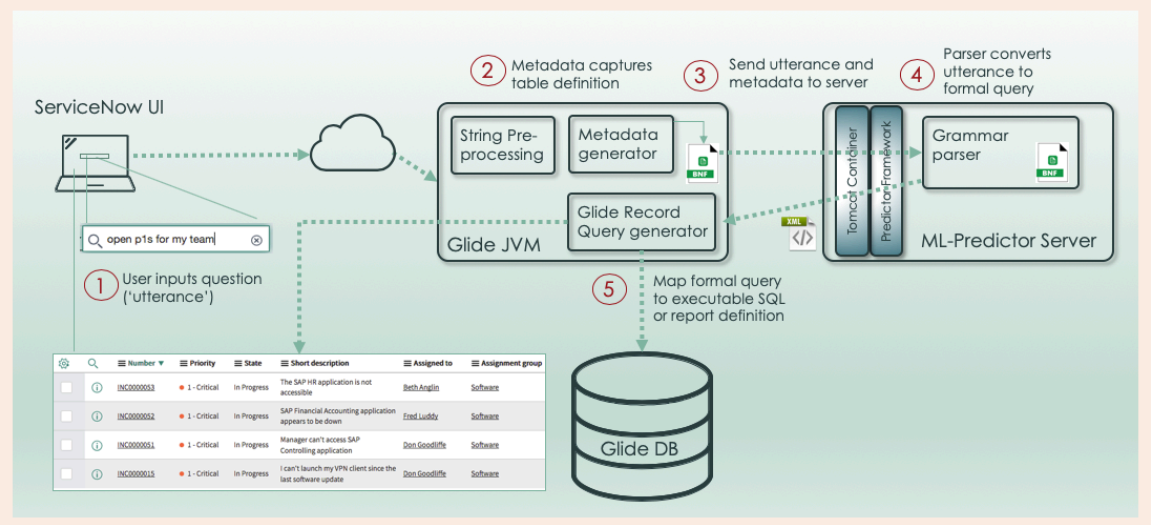
Get details about properties and roles installed with NLQ.

Exploring Natural Language Query

NLQ is a ServiceNow AI Platform feature that is active by default. Use NLQ to query the data in your instance by entering plain language requests into the user interface.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

NLQ data flow for user input and processing



ServiceNow® NLQ translates natural language user input into glide record queries. The queries are rendered into an executable structured format, such as a JavaScript Object Notation (JSON) file or a visual definition. The output, in whichever format, is the response to the user's request.

Note:
When a user enters a request directly into the user interface (UI), the text of the request is called an utterance in tables such as NLQ logs.

NLQ is a ServiceNow AI Platform feature that is active by default and supports the following data operations:

- Driving table suggestion
- Filtering
- Grouping and aggregations
- Sorting
- Data visualization (single score, list, bar chart, pie chart, time series)
- Business calendar
- Single number
- Multi-table

For more information, see Using Natural Language Query.

NLQ doesn't support domain separation. It also doesn't support on-premise instances.

Language support

NLQ supports American English by default. For all applications and features except CMDB, NLQ also supports queries in Spanish, French, French Canadian, German, and Japanese.

Note that you must first activate the languages on your instance for NLQ to parse queries in those languages. For more information, see Activate a language.

Using NLQ in other applications and features

Other ServiceNow® applications and features can consume NLQ functionality. See the following resources for more information.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Features that can consume NLQ

Application or feature	Information
Analytics Q & A	Create a report with Analytics Q&A
Configuration Management Database (CMDB)	Query your CMDB data without needing to know table relationships or data structures. • Querying the CMDB • Intelligent Search for CMDB English is the only supported language for CMDB.
NLQ with AI Search in global search	NLQ Genius Results in AI Search (supports English only)
List V2	Filter through any platform list by entering in a plain-language request.

Using Natural Language Query

With Natural Language Query (NLQ), you can query data in your tables by entering requests in natural, everyday language.

NLQ overview

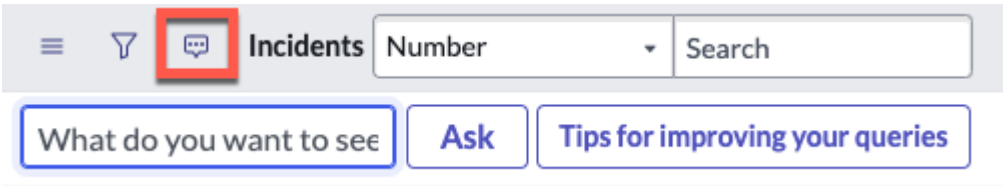
NLQ turns your plain-language requests into structured queries of your data. You don't need to know how to use the condition builder, because NLQ constructs and displays the conditions for you.

If you have a role such as itil that can view and interact with tables, you can use NLQ by selecting the natural language filter icon.

NLQ works on any list on the platform. It returns results only from the table or list you query on.

Find and use the natural language filter

Selecting the natural language filter icon brings up the NLQ interface.

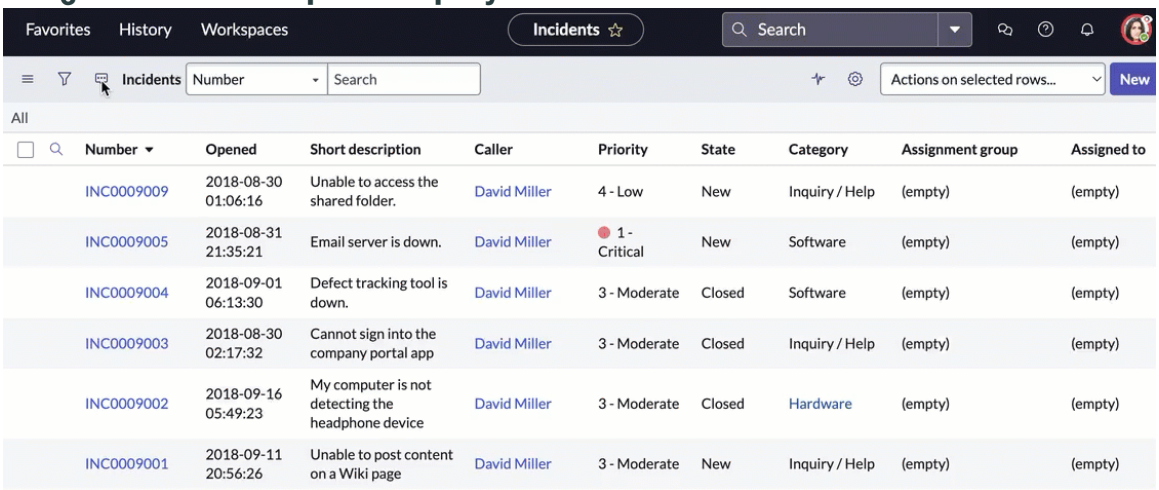


Enter your request in the What do you want to see field, then select **Ask**. NLQ parses your request, then displays the query in the condition builder. The results of the query are displayed in the list.

With the default configuration of NLQ, you can continue to refine your query by entering another request into the What do you want to see field.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Using NLQ to build on a previous query



Number	Opened	Short description	Caller	Priority	State	Category	Assignment group	Assigned to
INC0009009	2018-08-30 01:06:16	Unable to access the shared folder.	David Miller	4 - Low	New	Inquiry / Help	(empty)	(empty)
INC0009005	2018-08-31 21:35:21	Email server is down.	David Miller	1 - Critical	New	Software	(empty)	(empty)
INC0009004	2018-09-01 06:13:30	Defect tracking tool is down.	David Miller	3 - Moderate	Closed	Software	(empty)	(empty)
INC0009003	2018-08-30 02:17:32	Cannot sign into the company portal app	David Miller	3 - Moderate	Closed	Inquiry / Help	(empty)	(empty)
INC0009002	2018-09-16 05:49:23	My computer is not detecting the headphone device	David Miller	3 - Moderate	Closed	Hardware	(empty)	(empty)
INC0009001	2018-09-11 20:56:26	Unable to post content on a Wiki page	David Miller	3 - Moderate	New	Inquiry / Help	(empty)	(empty)

This example image and procedure illustrates how to build a query using NLQ:

1. On the Incident table in the natural language filter, enter `show me active hardware tickets` and select the **Ask** button.
2. The condition builder displays **All** > **Active = true** > **Category = Hardware** as the query. The filtered results are displayed in the list.
3. To narrow down the list of results further, enter `without assignment group`. Notice that as you type, NLQ displays possible matches for columns and fields. Select assignment group from the list of suggestions, and then **Ask**.
4. In the condition builder, NLQ adds **>Assignment group is empty** to the query. The list refreshes to display only the matching rows.
5. To reset and start a new query, delete everything in the condition builder so that only **All** remains.

Useful information

Keep the following information in mind when using NLQ.

- Your requests can contain periods and apostrophes, but not wildcard characters such as asterisks or regex.
- To group by a field or column, that column must be visible in the list view. Use the personalize list icon ( ) to hide or display columns.
- For information about querying CMDB tables, see Querying the CMDB.

Tips for improving your queries

When you select the **Tips for improving your queries** button, a modal window appears in the user interface. This window offers the following information about terms you can use in your NLQ questions and requests:

- Sorting or grouping: grouped by; sorted by; A-Z; z-a
- Dates: today; yesterday; last; this; next day(s); week(s); quarter(s); year(s)
- Filtering: starts with; ends with; more than; less than; empty; not empty; and; or
- Other information: my; my team; created by; unassigned

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Query	Example
Sorting or grouping	incidents grouped by category
Dates	created last month
Filtering	short description starts with computer
Other information	unassigned tickets
Single number	INC0777

Configuring NLQ

Enhance your users' query experience by supplementing NLQ with words and terms used in your environment. Review your users' actual requests in the NLQ logs to find possible synonyms and shortcuts to add.

Natural Language Query functionality is included in the base system. Admins can optionally expand NLQ's range of understanding by adding terms that commonly occur in users' requests. Review the following topics to learn more.

For information about system properties related to NLQ, see Natural Language Query References.

Create an NLQ synonym

Add synonyms to improve the ability of NLQ to recognize the various ways your users request data. With synonyms, you can map commonly used words or terms to table columns.

Before you begin

Role required: admin, nlq_admin, or pa_analyst

About this task

NLQ synonyms enable you to map common words in your users' requests to the columns in your instance tables. When these words are detected in natural-language queries, NLQ replaces them with actual column and table names, then submits the formal query.

Several synonyms are provided in the base system, and you can add more for your use cases and business requirements. Review NLQ logs of actual user utterances to find possible synonyms to add. For more information, see View NLQ logs.

You can create a new synonym, or modify an existing synonym, as follows.

- If any synonym already exists for your target table and column, you must add your new synonym to the current record. Separate multiple synonyms with a comma.
- Synonyms can point to reference fields, using dot-walking. For more information, see Dot-walking examples.
- Synonyms are case-insensitive in queries.
- Synonyms can contain apostrophes and periods, but not commas.
- Synonym records are associated to one table. You can use the same synonym term for more than one table, but you must create a record for each table.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Ensure that you are in the application scope you want for your synonym, then navigate to **All** > **NLQ** > **Synonyms**.
2. Select **New**.
   If you are updating an existing record, select its **Column name** in the list.
3. On the form, fill in the fields.

Field	Description
Synonym Type	Type of synonym. ○ **Table or Column**: Synonym for a value on a table or a specific column in that table. ○ **Record**: Synonym for a value on the CMDB tables [cmdb_rel_type_table]. For more information, see Querying the CMDB. **Note:** Supports cmdb_rel_type,cmdb_group, cmdb_ci_service_technical, cmdb_ci_service_discovered, cmdb_ci_query_based_service tables.
Application	[Read only] Application scope that can use the synonym in a query. The default is **Global**. When creating a synonym, ensure that you are in the scope you want for the synonym.
Table	Source table. You can use the same term for more than one table, but you must create a synonym record for each table. The synonym's mapping is unique to the table.
Column name	Specific column on the source table.
Language	Language of the synonym. Should match the language of the source value.
Synonyms	Words or short phrases that the system should map to column names or tables, when converting the utterance to a formal query. Separate multiple synonyms with a comma.
Extend to All Child Tables	**Table or Column** type only. Select this option to make any child tables of the source table inherit the synonyms.

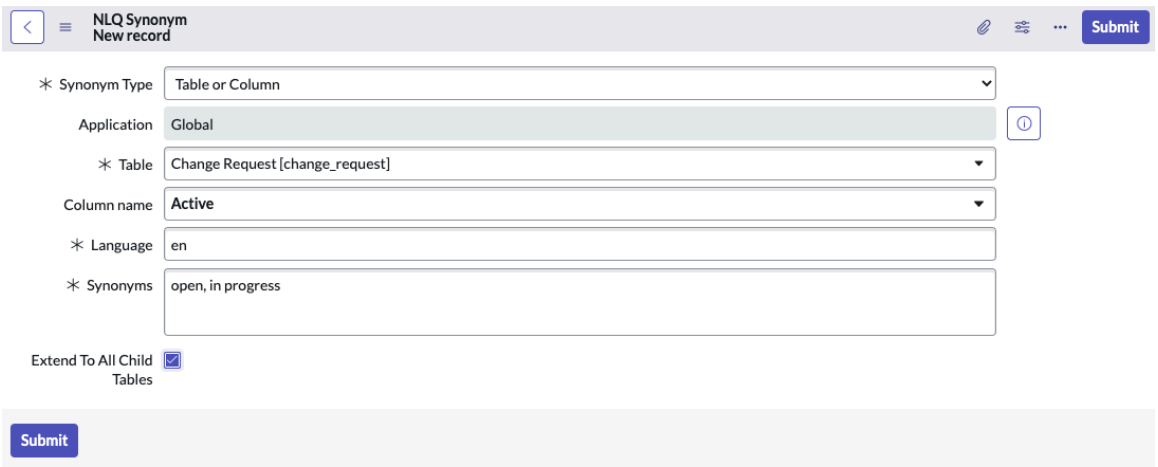
4. Select **Submit** if new, or **Update** if you are modifying an existing record.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
The new synonym is available to your users as soon as they refresh the browser window of the list.

Example: NLQ Synonym for the Active field

The following image shows an example of an NLQ synonym record for the **Active** field on a change request. The synonyms open and in progress are replaced by the term active when the system submits a query.



With this synonym record, a user can type `show open change requests` or `change requests in progress` and the system displays active change requests.

Create an NLQ shortcut

Create a semantic shortcut to help improve the ability of NLQ to recognize the various ways your users request data. Semantic shortcuts operate similarly to NLQ synonyms by mapping common words to columns, but for a selected table when certain conditions are met.

Before you begin
Role required: admin, nlq_admin, or pa_analyst

About this task
Like NLQ synonyms, semantic shortcuts enable you to map common words in your users' requests to the columns in your instance tables. When these words are detected in natural-language queries, NLQ replaces them with actual column and table names, then submits the formal query.

Semantic shortcuts provide a condition builder so that you can filter which records are covered by the terms you provide.

Some semantic shortcuts are provided in the base system, and you can add more for your use cases and business requirements. Review NLQ logs of actual user utterances to find possible terms to add. For more information, see View NLQ logs.

You can create a new shortcut, or modify an existing shortcut, as follows.

- If any shortcut already exists for your target table and filter conditions, you must add your new shortcut to the current record. Separate multiple shortcuts with a comma.
- Shortcuts can point to reference fields, using dot-walking. For more information, see Dot-walking examples.
- Synonyms are case-insensitive in queries.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Synonyms can contain apostrophes and periods, but not commas.
- Synonym and shortcut records are associated to one table. You can associate the same synonym term to more than one table, but you must create a record for each table.

Procedure

1. Ensure that you are in the application scope that you want for your shortcut, then navigate to **All** > **NLQ** > **Semantic Shortcuts**.
2. Select **New**.
   If you are updating an existing record, select its row in the list.
3. On the form, fill in the fields.

Field	Description
Table	Source table. You can use the same term for more than one table, but you must create a record for each table. The synonym's mapping is unique to the table.
Conditions	Conditions on the source table that must be met for the synonym to work.
Synonyms	Words or phrases that the system should map to column names or tables, when converting the utterance to a formal query. Separate multiple synonyms with a comma.
Application	[Read only] Application scope that can use the synonym in a query. Default is **Global**. Ensure that you are in the target application scope for your shortcut before creating it.

4. Select **Submit** if new, or **Update** if you are modifying an existing record.

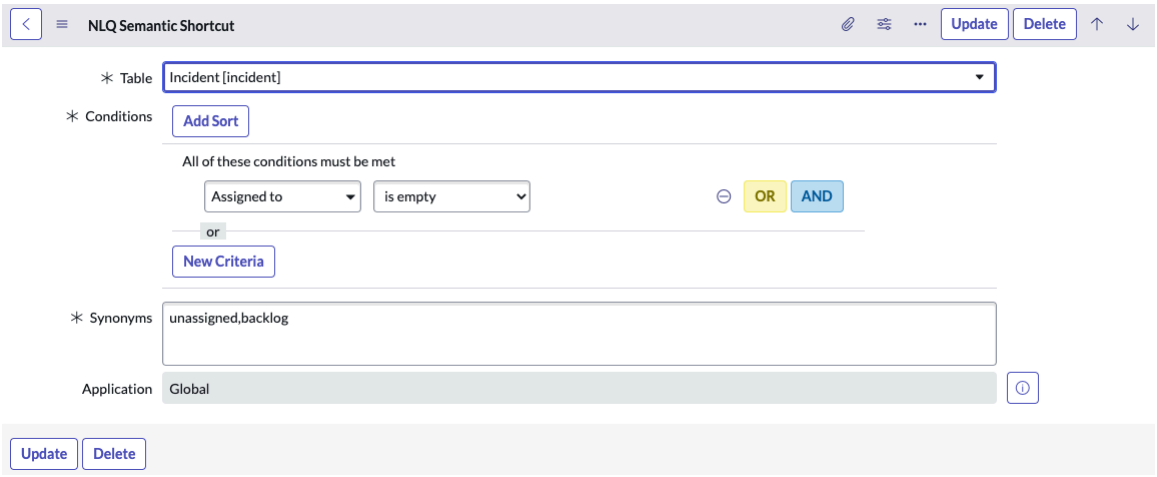
Result

The new shortcut is available to your users as soon as they refresh the browser window of the list.

Example: Semantic shortcut for Incident backlog

The following image shows an example of a semantic shortcut used on incident records. Incidents contain the Assigned to field. Using the condition builder, you can provide synonyms for when the field is empty. When NLQ detects unassigned and backlog in your users' input, it replaces them with the SQL clause `assigned_toISEMPTY`.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



A user can enter `show me incident backlog` or `unassigned incidents` and the system displays incidents with an empty Assigned to field.

View NLQ logs

Review NLQ logs to see how the system has handled your users' plain-language requests. Use log records from attempted requests to expand NLQ synonyms or shortcuts.

Before you begin

Role required: nlq_admin, pa_analyst, or admin

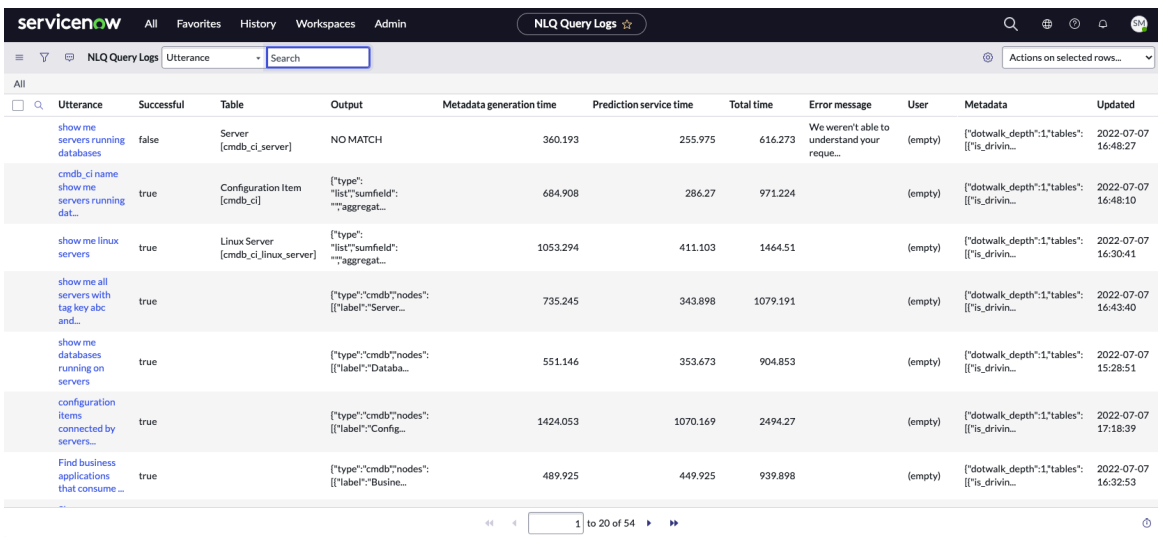
About this task

Every natural language query is logged in the *NLQ* **Query Logs** table [nlq_query_log]. Each log entry provides details such as the table that was queried, whether the query succeeded, and how the results were generated. Other available fields include the following.

- Output Source: how the results were generated. The value **BNF** indicates a rules-based method. The value **GAI** indicates the Now LLM Service (fallback) method.
- Source: the location from which the query was initiated. The value **AC** indicates Analytics Center. The value **CMDB_WS** indicates CMDB Workspace.
- Utterance: the original natural-language request which triggered NLQ.

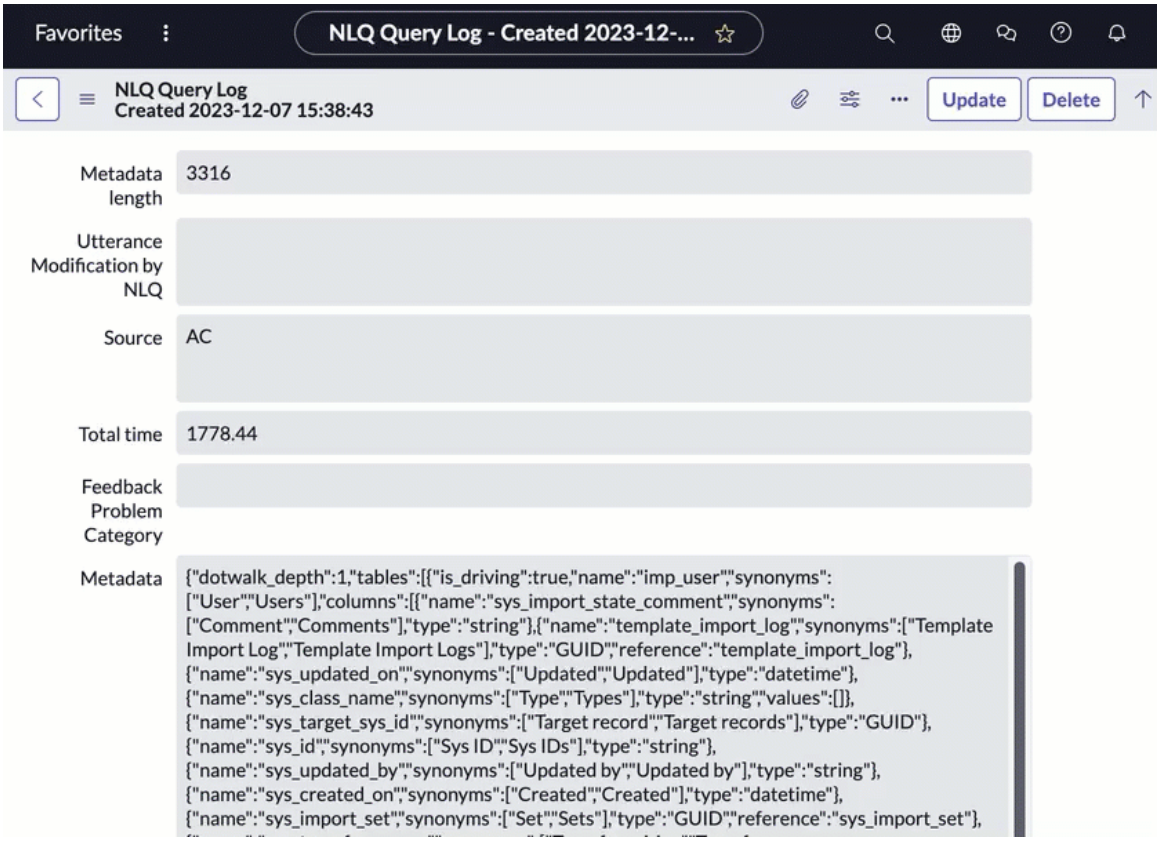
Procedure

1. Navigate to **All** > **NLQ** > **Logs**.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

2. Select the value in the **Utterance** column to open the full log entry.



What to do next

Based on your users' attempted queries, consider adding more synonyms or shortcuts.

View NLQ Table Guesser logs

Use the Table Guesser logs to review the CMDB tables that were picked by NLQ in response to plain-language queries.

Before you begin

This module is read-only.

Role required: nlq_admin.

About this task

When a plain-language request for CMDB data is parsed by NLQ, the system tries to determine the specific tables the user intended. The system's guesses, including confidence levels, are recorded in the NLQ Table Guesser logs module [nlq_table_guesser_log].

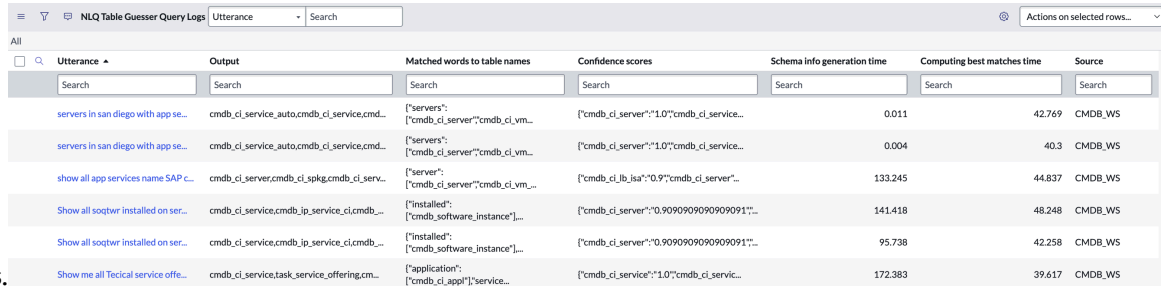
Review these logs to troubleshoot which CMDB tables were inferred by NLQ.

Note:
This list is read-only for NLQ admins. There are no buttons or actionable functions in the UI.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **NLQ** > **Table Guesser Logs**.
2. Select the personalize list icon ( ) to display the columns **Matched words to table names** and **Confidence scores**.



   The **Utterance** column shows the user's natural language query.
3. Use the information in the **Matched words to table names** column for troubleshooting if needed.
   In this column's value field, the first word was found in the utterance. Next is a list of the CMDB tables that were matched to that word.

What to do next

For more information on CMDB queries, see Querying the CMDB.

Natural Language Query References

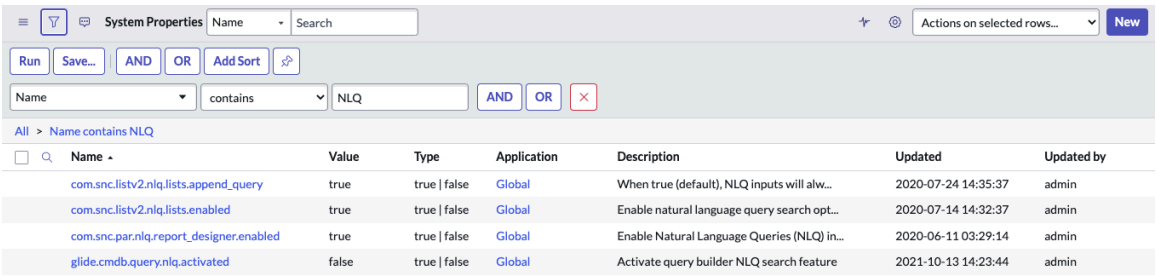
The following components are installed with Natural Language Query.

NLQ properties

The Natural Language Query (NLQ) properties control how and where NLQ operates.

Admins can edit properties of NLQ by navigating to **All** > **System Properties** > **All Properties**. Filter for the NLQ properties.

> **Note:**
> Editing these system properties requires the admin role. The nlq_admin role does not have permission to edit records in this table.



Property	Description
com.snc.listv2.nlq.lists.append_query	• True: NLQ inputs add onto existing queries via an "and" operator • False: New NLQ input replaces any existing queries

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Property	Description
	Example: You run two queries. • Query 1: Incidents with critical priority • Query 2: assigned to John Smith If the property is set to true, the results show incidents with critical priority that are assigned to John Smith. If the property is set to false, the results show only items assigned to John Smith.
com.snc.listv2.nlq.lists.enabled	• True: Enables NLQ search option for List v2 • False: Removes NLQ search option for List v2
com.snc.nlq.gai_enabled	• True: The Now LLM Service fallback is available • False: The Now LLM Service fallback is not available Initially, queries are interpreted using a rules-based method. If that method fails, queries are passed to the Now LLM Service as a fallback. Queries that fail both of these methods are marked as unsuccessful in the NLQ log.
com.snc.par.nlq.report_designer.enabled	• True: Enables NLQ in Report Designer • False: Removes NLQ in Report Designer
glide.cmdb.query.nlq.activated	• True: NLQ search feature is active in CMDB Query Builder • False: NLQ search feature is inactive in CMDB Query Builder
glide.service_portal.ais_nlq_enabled	• True: Enables NLQ in global search • False: NLQ is not available in global search

Natural Language Query roles

Natural Language Query (NLQ) is installed with these roles.

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

NLQ Admin [nlq_admin]

The administrator for Natural Language Query.

Module Access

Has full access to the following modules:

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- NLQ Cmdb Implicit Relationships. For more information see Intelligent Search for CMDB.
- NLQ Query Logs
- NLQ Semantic Shortcuts.
- NLQ Synonyms.

Has read-only access to the following module:

NLQ Table Guesser Query Logs.

Contains Roles

List of roles contained within the role.

None.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

To avoid granting an admin role, use the pa_analyst role for NLQ Cmdb Implicit Relationships, NLQ Semantic Shortcuts, and NLQ Synonyms.

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.
- Estimate time for case resolution.
- Identify major outages based on incoming incidents.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

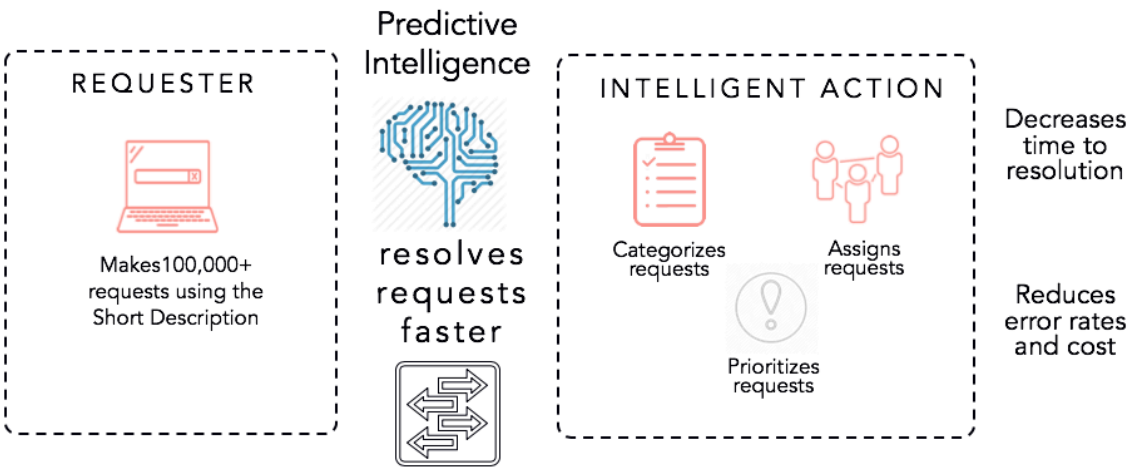
Explore Learn about Predictive Intelligence and machine learning	**Install** Install Predictive Intelligence and its associated apps	**Configure** Configure Predictive Intelligence on the platform
Train Create and train machine learning solutions.	**Use** Use Predictive Intelligence for machine learning.	**Reference** Get details about Predictive Intelligence components such as properties and language support.

Explore Predictive Intelligence

ServiceNow® Predictive Intelligence is a platform function that provides a layer of artificial intelligence that empowers features and capabilities across ServiceNow® applications to provide better work experiences.

Overview of Predictive Intelligence

Predictive Intelligence is a powerful set of tools 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 applications.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The following 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- The 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 Washington DC 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 Washington DC 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Predictive Intelligence 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 reusing 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

Important:
Support for creating new regression solutions was removed in the Washington DC release. You will still be able to train and edit existing solutions, but you won't be able to create new ones.

Regression is a machine-learning framework that you can train with historic data to predict numeric outputs, such as a temperature or a stock price. Regression frameworks can be used directly to measure success. For example, you can use regression to estimate the time it takes to resolve an incident or a case.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

When you 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 the 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. Because 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. For more information, see Create and train a classification solution.
5. On the non-production instance, train the solution definition records.
6. Test the solution predictions by either creating test records or importing more records from production.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

7. 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're 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 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Most raw data sets contain dirty and unusable data. Reviewing your input sets before training is essential to keeping accurate predictive models.

It is recommended to use approximately 80% of your input data to train your model and about 20% of the data to evaluate whether the model is accurate. You can compare the model's predicted results against the real values for the 20% of remaining data.

Solution training

Issue	Resolution or suggested action
The solution training remains in Waiting for Training status for too long, as the scheduler job is using an incorrect Glide callback instance URL.	Ensure the *glide.servlet.uri* property in the Glide instance is set to the correct instance URL. This issue can occur when: • An instance is cloned from production, yet it still refers to the production URL for the *glide.servlet.uri* property. • The Glide instance is provisioned and runs the training for the first time.
New categories have been added and aren't yet having an impact on training.	This is expected behavior, as the new categories may not yet have sufficient data until the solution is retrained.
The solution training fails.	When the training fails, click the **Show Training Progress** related link on the solution screen to determine where the potential problem resides.
The solution training fails due to user authentication.	Navigate to **System Security> Users** and ensure the sharedservice.worker user is set to Active.
The model training returns saying the model cannot be created. The training fails and shows the “Error while training solution” message. The training progress window shows this message: “Solution training failed as either the data used isn't sufficient or the input field isn't predictive of the output field."	This issue can occur when the data quantity or the distribution of field values isn't sufficient for a model to build successfully. Follow these steps to troubleshoot: **1.** Ensure the distribution of the output field isn't skewed. **2.** Retrain the model by changing the date filters to use a larger amount of data. **3.** If the input fields are not fully populated, add a filter to remove null records.
The solution has data in multiple languages but the coverage and precision results are poor.	Use the following options to help improve your metrics. Option 1: Update the processing language of the solution to the most prominent non-English language.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Issue	Resolution or suggested action
	Note: English is applied by default for all datasets. Option 2: If there's sufficient data for each language/region: 1. Add a filter criteria for a specific language/region where the primary language can be identified (Dutch, English, French, German, Japanese, or Spanish). 2. Generate a solution for each language/region and apply the proper processing language to each solution.

Solution prediction

Issue	Resolution or suggested action
The prediction fails and returns a Java exception where the cause is unknown.	1. Search for the exception in the Predictive Intelligence Glide logs. 2. Submit an Incident record for Predictive Intelligence including all relevant details, such as the exception, the impacted instance, the solution name, and the input string.
There is no prediction applied to the incident/case record but the prediction returns a value when tested in the Rest API Explorer.	This can occur when the confidence of the prediction is less than the threshold required to make a prediction. After your solution is trained, use the following steps to confirm if your solution settings need adjusting. 1. Navigate to **System Web Services > REST > REST API Explorer** to find the confidence level for the prediction. See Test a classification solution prediction. 2. On your ML Solution Definition record, check the threshold set for your outcome class that was returned in the prediction by clicking on the name of the class. The **Class** page appears. 3. Check the **Estimated Precision** and **Estimated Coverage** values. If the corresponding threshold is more than the prediction confidence of the outcome, this is the root cause for why you did not see any prediction. 4. Adjust the class precision and coverage values to increase coverage or precision. See Tune a trained classification solution.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Instance 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.

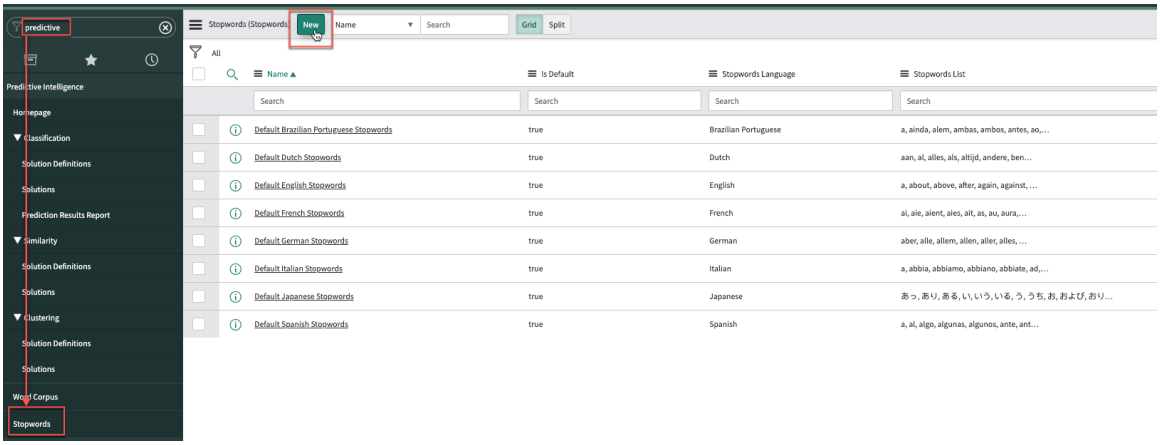
The custom list you provide works alongside those that the system already uses by default. For example, if incident records are used in a classification solution, and a company name is used in those records, consider adding that name to your list, as it's unlikely to provide any relevant information for the solution you're building.

In this example scenario, you create a custom stopwords list for the Brazilian Portugese language.

Procedure

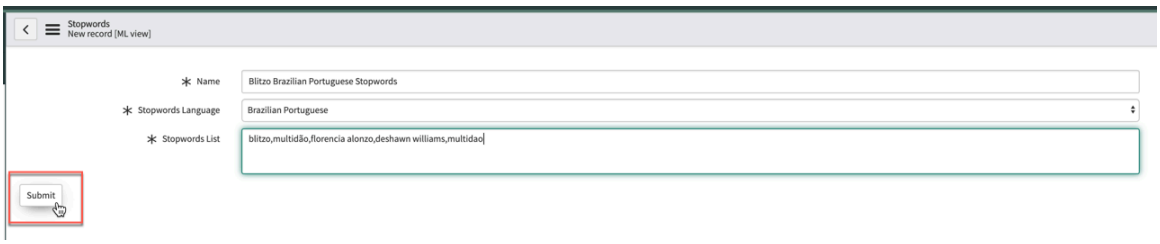
1. Navigate to **All** > **Predictive Intelligence** > **Stopwords**.
2. On the Stopwords list, click **New**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

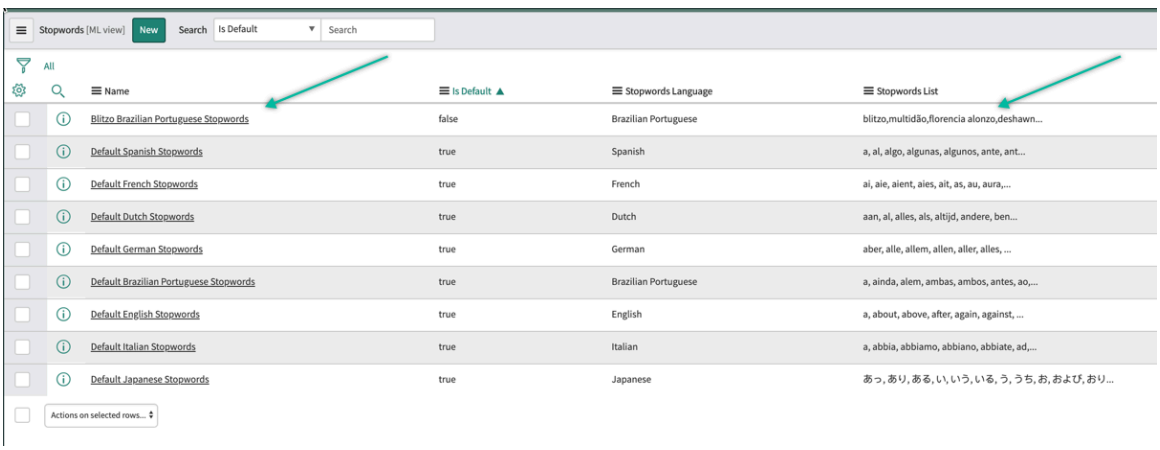


3. In the Stopwords form, configure these fields.

Field	Value
Name	Enter a unique name for the list, such as the name of your company and the processing language. For example, `Blitzo Brazilian Portuguese Stopwords`.
Stopwords Language	Select **Brazilian Portuguese**
Stopwords List	Manually enter the stopwords using a comma-separated format. For more examples of stopwords, see the image in Step 2 of this procedure.

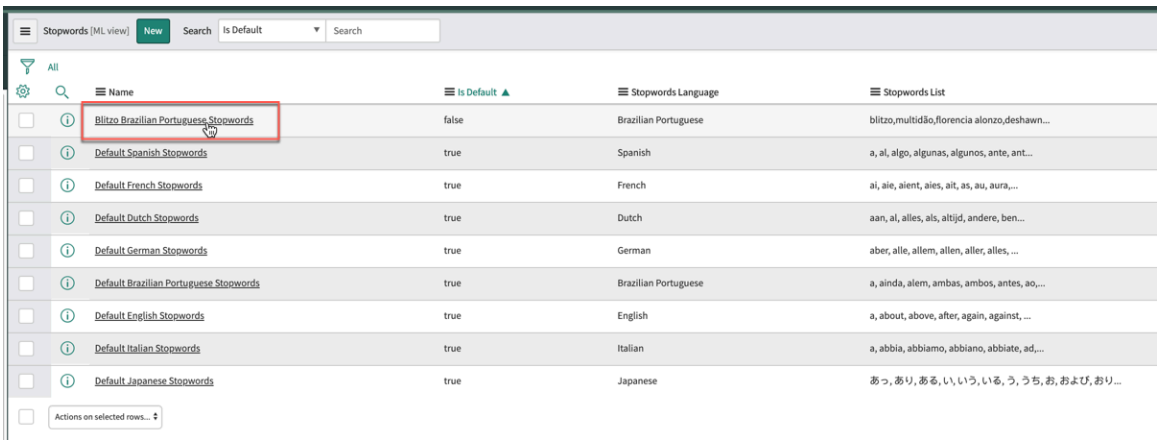


4. Click **Submit**.
Your custom stopwords list appears in the Stopwords list view.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

5. **Optional:** If you need to update your stopwords list, just click its Name, add or remove words from the list, and click **Update**.



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 Washington DC release, models in the classification, clustering, and similarity frameworks use Workflow solutions. These are pre-trained, so a word corpus isn't needed for your new solutions.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Word Corpus**.
2. In the Word Corpus form, click **New**.
3. Configure these fields according to the following guidance.

Field	Description
Name	A unique title that references the contents of your corpus. For example, in this use case you could enter a name such as `Active Incidents and Published KBs`, as the name indicates the tables that your corpus will mine to help create your solution.
Active	Select this check box if you're creating more than one word corpus at a time and you plan to configure their detail components later. Otherwise, leave it blank because you can select it in a later step.

4. Select **Submit**.
5. In the Word Corpus list view, locate your new word corpus and click its **Name** value to open the record.
6. In the Word Corpus Contents section, Click **New**.
7. In the Word Corpus Content form, configure these fields per the following guidance to define a content component for your word corpus.

Field	Description
Name	Enter a title that references the data you want to add to your corpus, such as `Incidents closed in last 6 months`.
Table	Select the table that contains the data you want to include in your word corpus. For this use case, select **Incident [incident]**. **Note:** The number of records per table for Word Corpus creation used in Similarity and Clustering solutions is limited to 300,000.
Filter	Select the following filter condition values: **[Closed] [is not empty] and [Created in last 6 Months]**.
Field List	For this use case, select **Short description**, **Description**, and **Resolution notes**.
Domain	The system automatically displays the user group for your corpus. For example, in this use case it shows the global user group. You can select other user groups as well.

8. Select **Submit**.
9. In the Word Corpus Details section, select **New**.

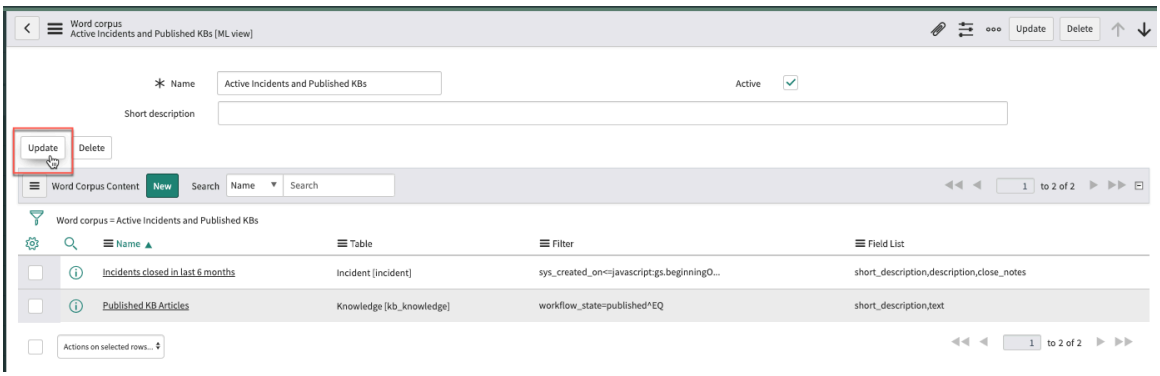
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

10. Configure these fields according to the following guidance to define a second content component for your word corpus.

Field	Description
Name	Enter a title that references the data you want to compare to your first content component, such as `Published KB Articles`.
Table	Select the table that contains the data you want to compare to your first content component. For this use case, select **Knowledge [kb_knowledge]**. **Note:** The number of records per table for word corpus creation used in Similarity and Clustering solutions is limited to 300,000 records per table.
Filter	Select the following Filter Condition values: **[Workflow] [is] [Published]**.
Field List	Select **Short description** and **Article body**.

11. Select **Submit**.

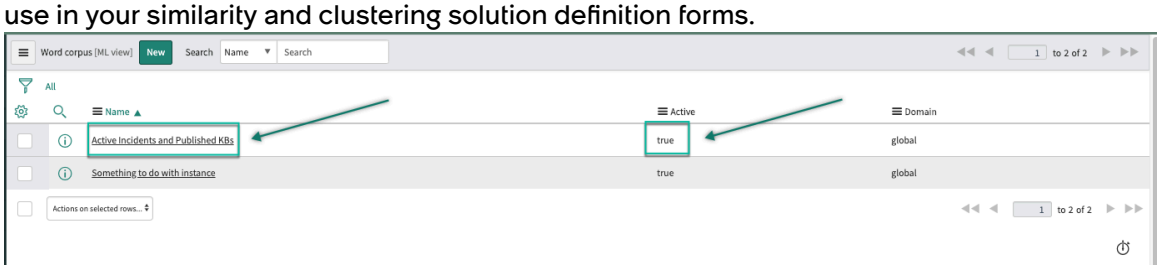
Your two word corpus content components appear on the word corpus form.



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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Related topics

- 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 solution version

The system activates the most recent version of the solution when it completes training a solution, and only allows one solution version to be active at a time. However, you can activate any previously trained solution version you want Predictive Intelligence to use to make predictions.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- 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 only manually create a new solution version 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 ML Solutions list view, click the Reference Lookup icon for the trained solution that you want to activate.
   
3. Click **Open Record**.
4. In the solution record, click **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**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
Your 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, similarity, and clustering capabilities.

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 about several of these parameters, see the Dive deeper with Clustering Advanced Parameters article on ServiceNow Community.

Classification, similarity, and clustering 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. These settings offer functionality that is targeted to specific use cases.

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. To enable one or more, 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

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

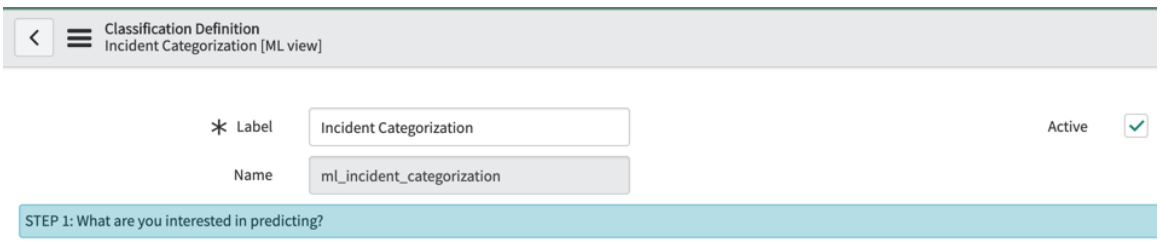
- Create and save a classification solution definition or use an existing one.
- Role required: admin or ml_admin

About this task

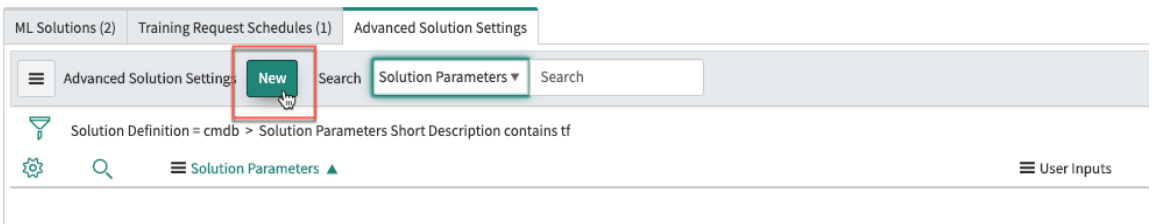
The class recall solution parameter enables you to steer a solution's training to bias a specific class. For example, classifying an incoming email as a Phish or not can be an important use case in a security-related machine learning solution. In this situation, it's very important to identify every 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**.

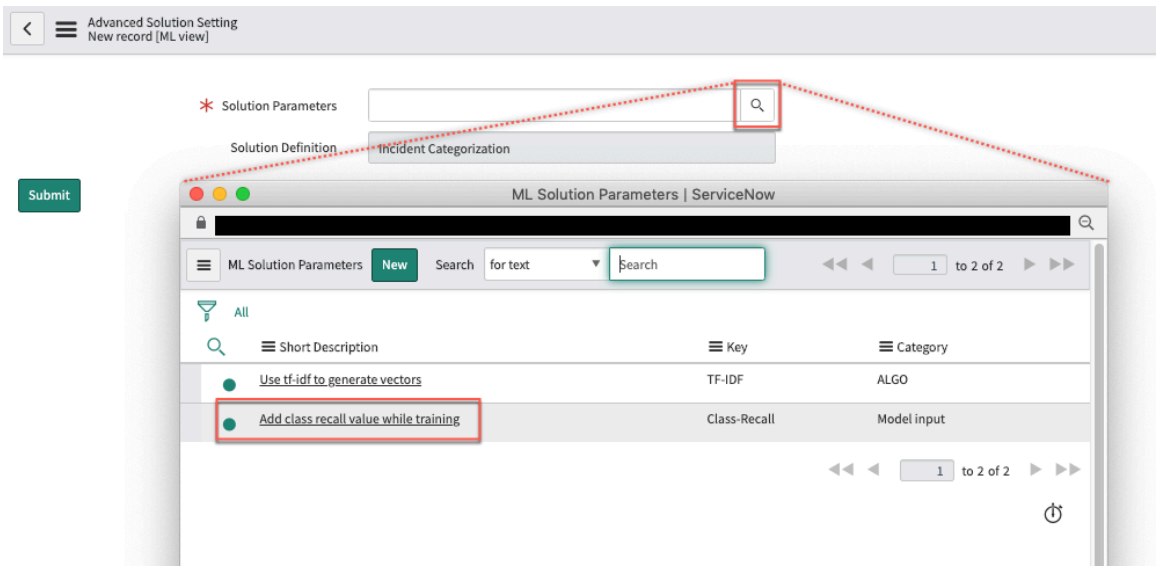


4. Create a parameter record.

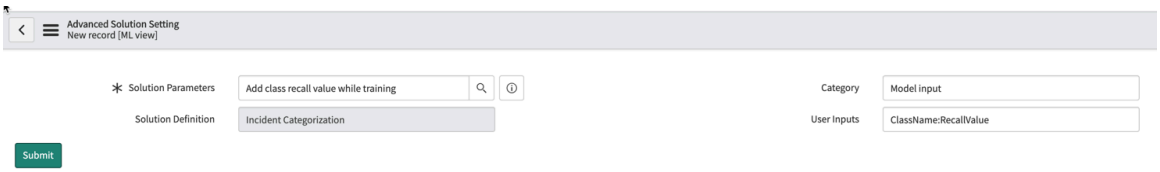
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the **Solution Parameters** field, click the search icon.

b. In the ML Solution Parameters screen, select **Add class recall value while training**.



5. Click **Submit**.
The Advanced Solution Setting record appears.

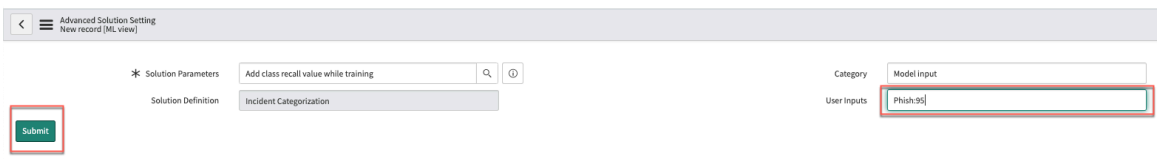


6. Configure the **User inputs** field.

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**.

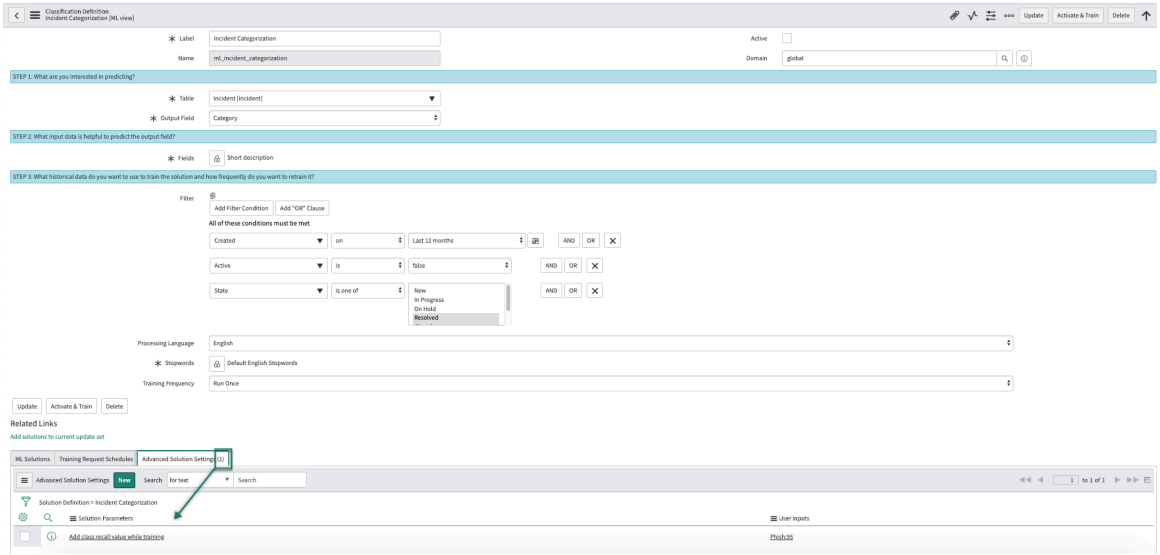
Here you're specifying Phish as the target class, and 95 is the recall percentage you're requesting the system to deliver during solution training.



7. Click **Submit**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result: 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

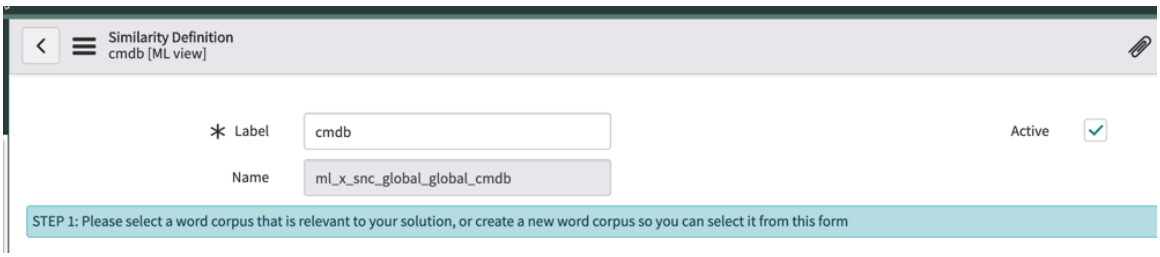
Predictive Intelligence uses paragraph vector word embedding by default in its classification and similarity frameworks, which is highly effective for processing data comprised of primarily human-readable content. However, TF-IDF might return better prediction results for records that have machine-generated 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

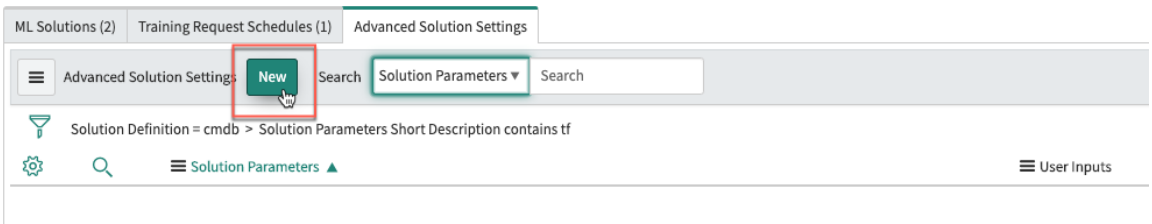
Note:
The steps for configuring TF-IDF are the same for all model frameworks, but TF-IDF support for clustering solution definitions is applicable only if you have a Professional subscription.

Procedure

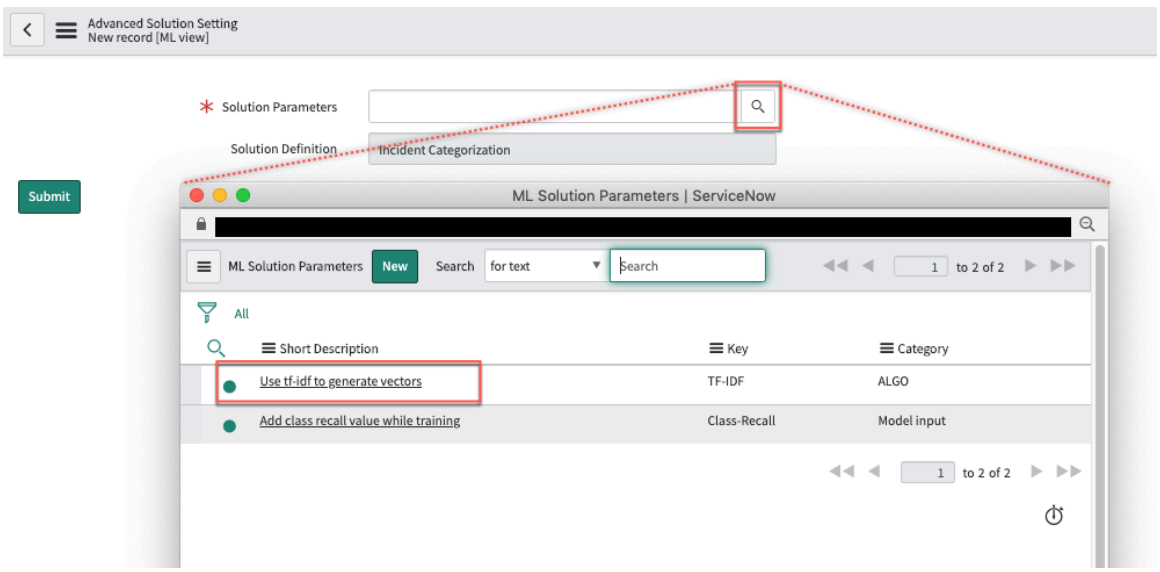
1. Navigate to a **Solution Definition**, such as **All** > **Predictive Intelligence** > **Similarity** > **Solution Definitions**.
2. Open a solution definition form.
   In this example scenario, you use a CMDB similarity definition form.



3. On the Advanced Solution Settings tab in the Related Links section of the form, click **New**.



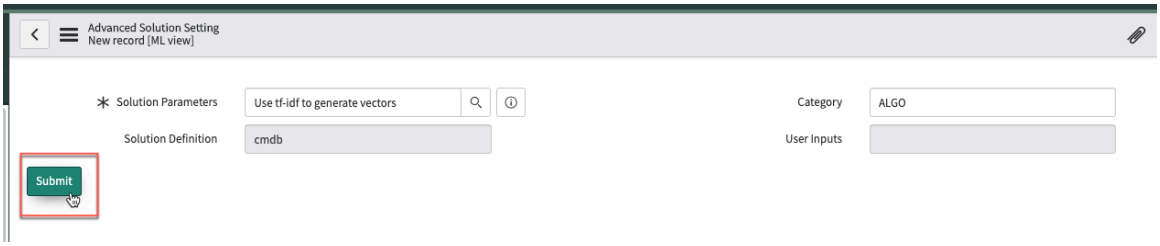
4. Create a parameter record.
   a. In the **Solution Parameters** field, click the search icon.
   b. In the ML Solution Parameters screen, select **Use tf-idf to generate vectors**.



5. Click **Submit**.

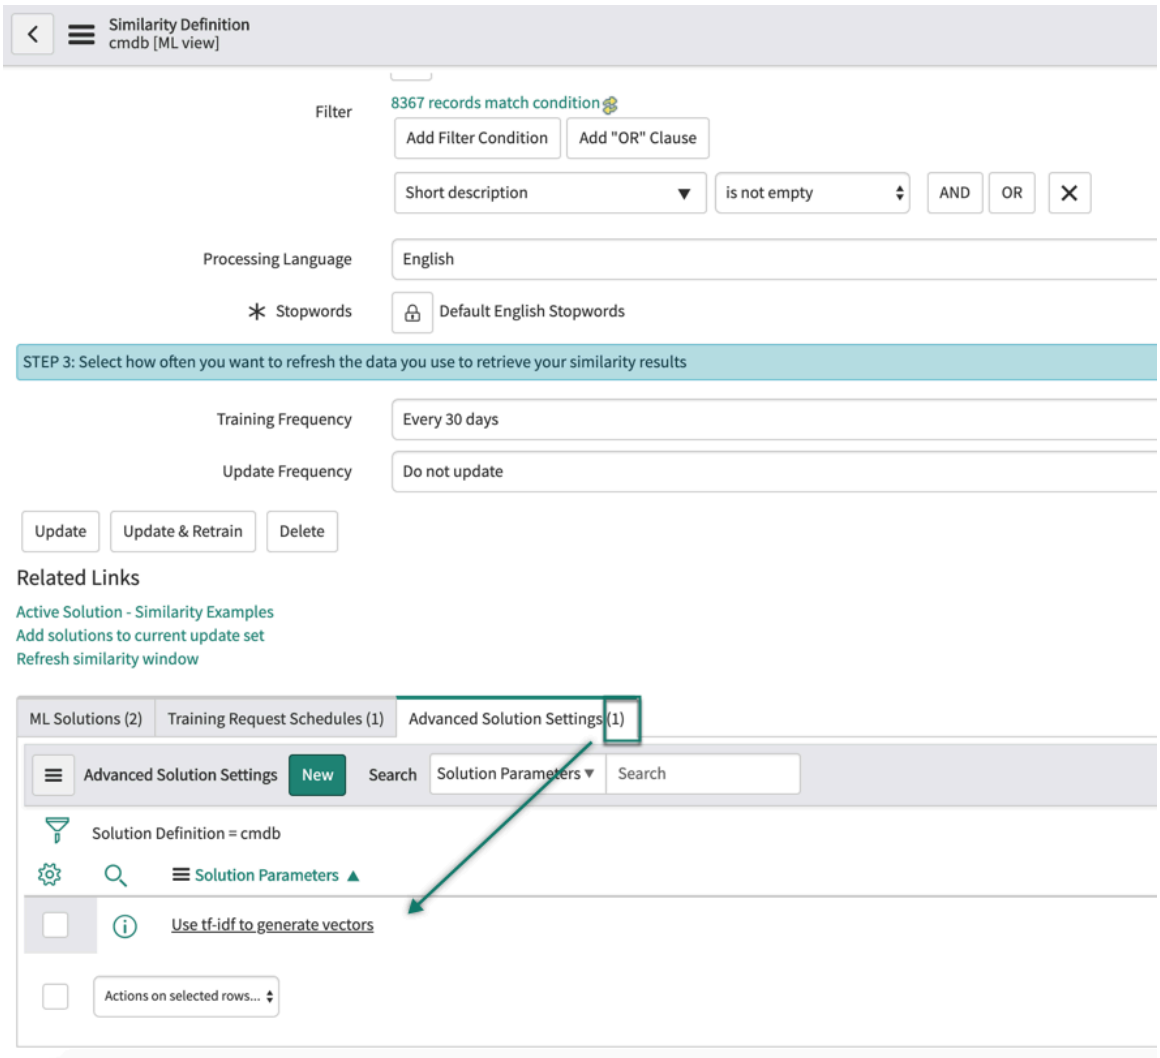
   The Advanced Solution Setting record screen refreshes.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



6. Click **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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

Note:
Configuring advanced settings on your ML solutions is optional. If you choose to configure any of these settings, make sure you're well informed regarding the technology you're enabling in the solution, and that 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

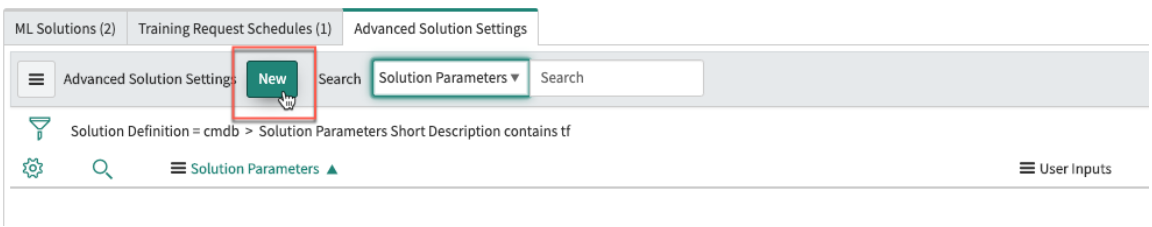
About this task

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.

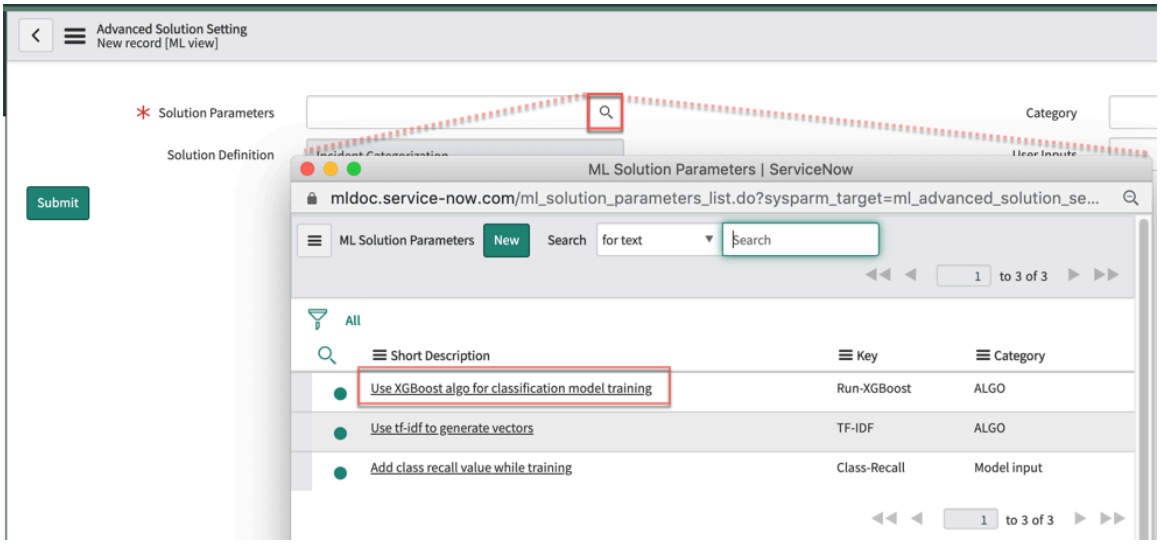
Procedure

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**.

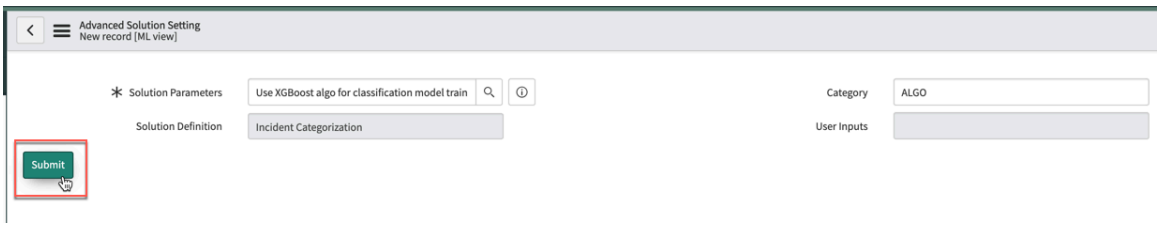


4. Create a parameter record.
   a. In the **Solution Parameters** field, click the search icon.
   b. In the ML Solution Parameters screen, select **Use XGBoost algo for classification model training**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

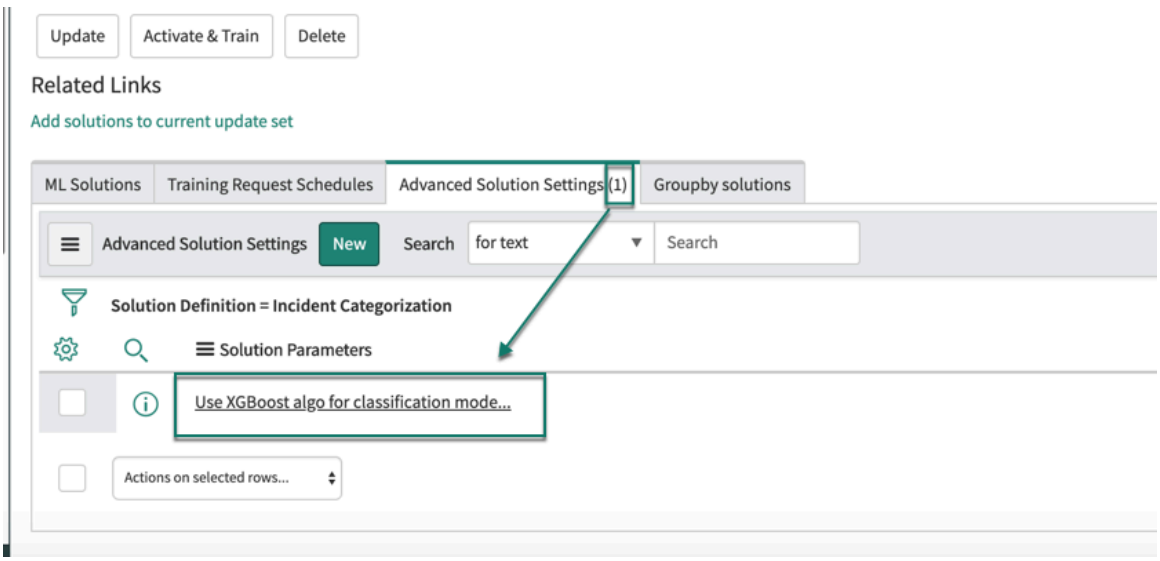


5. Click **Submit**.
   The Advanced Solution Setting record screen refreshes.



6. Click **Submit**.

   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**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

8. In 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

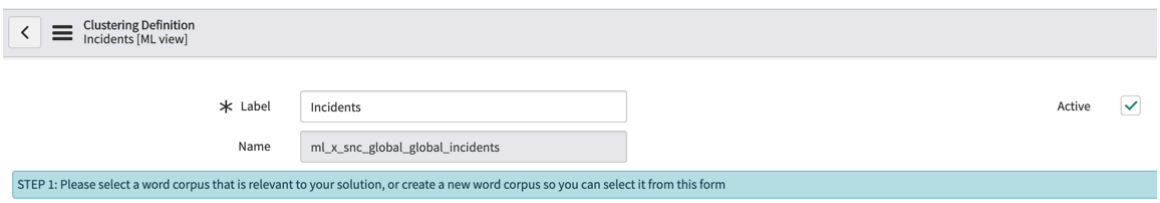
About this task

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.

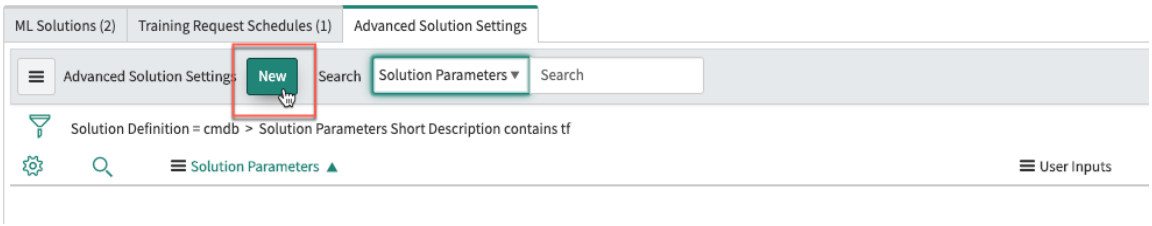
Procedure

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.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

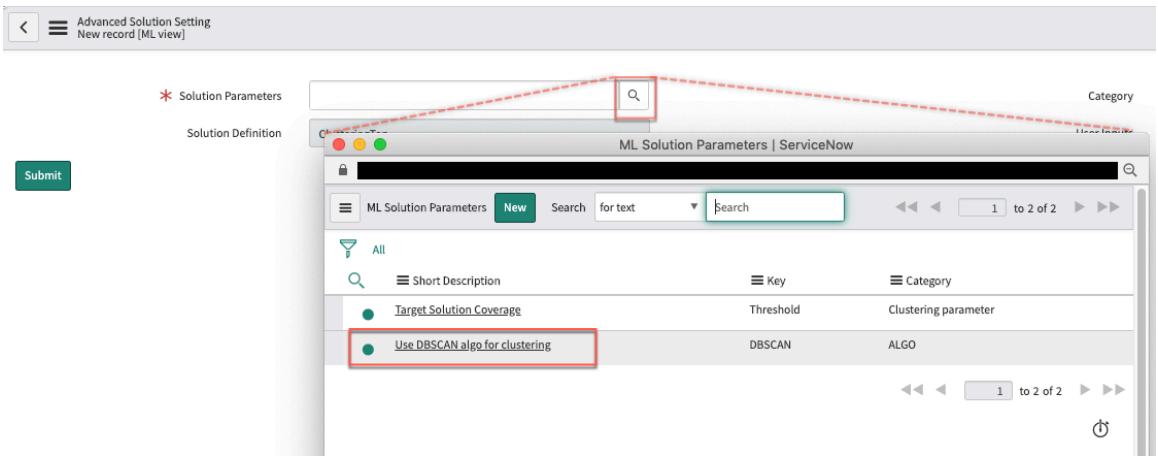
3. On the Advanced Solution Settings tab in the Related Links section of the form, select **Solution Parameters** from the picker, then click **New**.



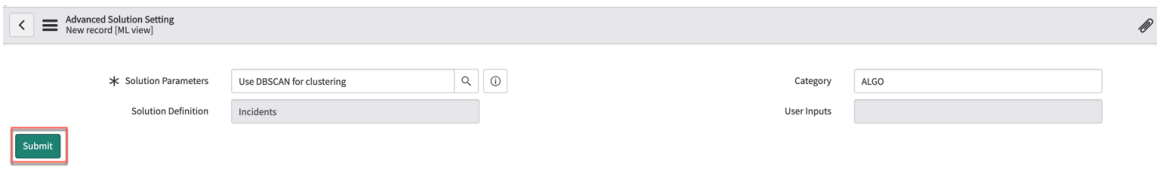
4. Create a parameter record.

   a. In the **Solution Parameters** field, select the search icon.

   b. In the ML Solution Parameters screen, select **Use DBSCAN algo for clustering**.



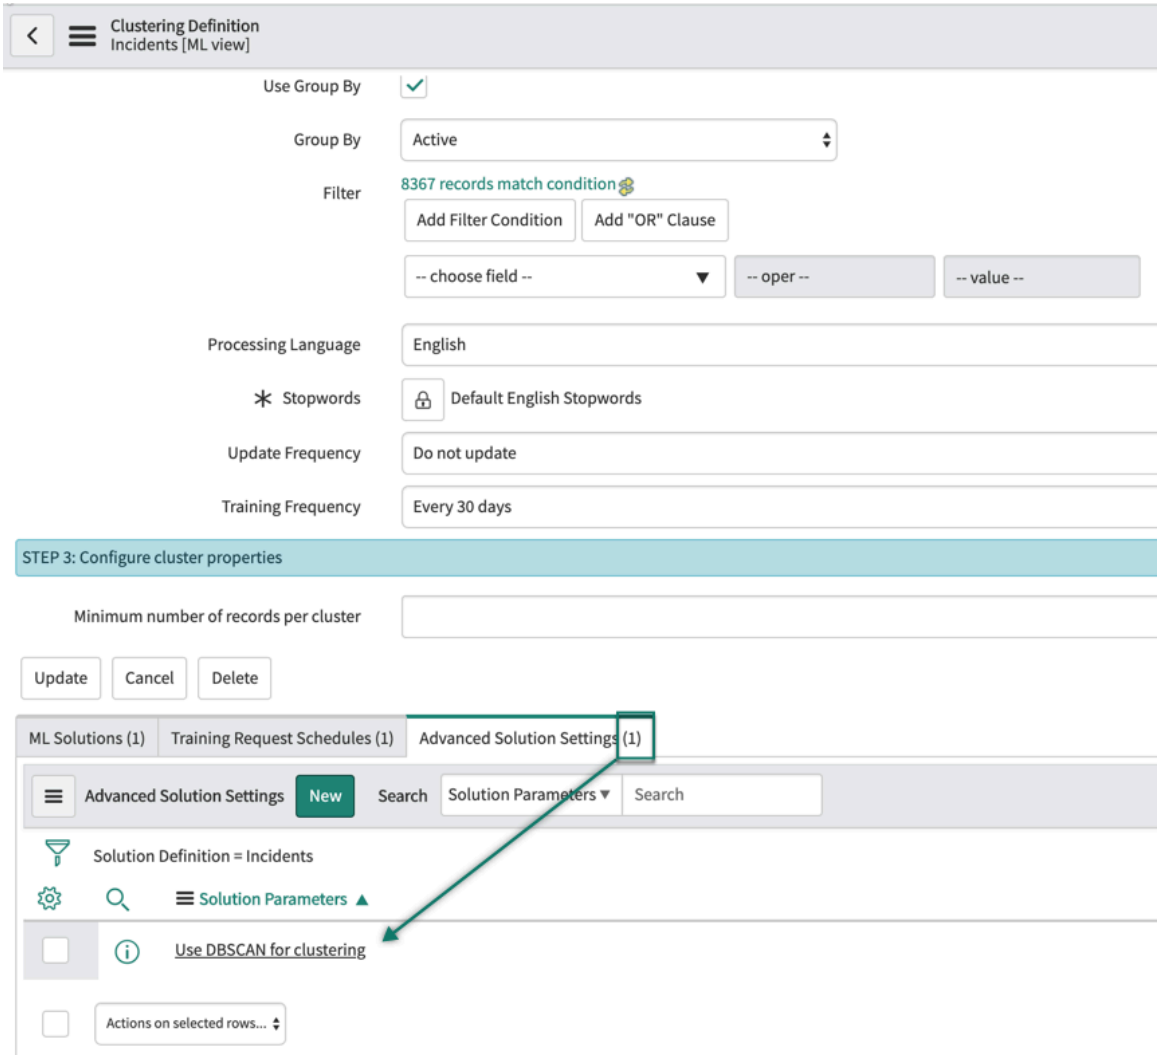
5. Select **Submit**.
   The Advanced Solution Setting record appears. The field **User Inputs** is grayed out because it does not apply to this algorithm.



6. Select **Submit**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result: DBSCAN is configured for your clustering solution. Its solution parameter appears on the Advanced Solution Settings tab of your clustering solution definition form.



Related topics

Create and train a clustering solution

Configure HDBSCAN for a clustering solution

Consider applying the Hierarchical Density Based Spatial Clustering of Applications with Noise (HDBSCAN) algorithm to your clustering solution. HDBSCAN is available as an alternative to the default clustering algorithm, k-means.

Before you begin

Note:
Configuring advanced settings on your ML solutions is optional. If you choose to configure any of these settings, make sure you're well informed regarding the technology you're enabling in the solution, and that your use case benefits from what the technology offers. For more information, see the Dive deeper with Clustering Advanced Parameters article on ServiceNow Community.

- Create a clustering solution definition or use an existing one.
- Role required: admin or ml_admin

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

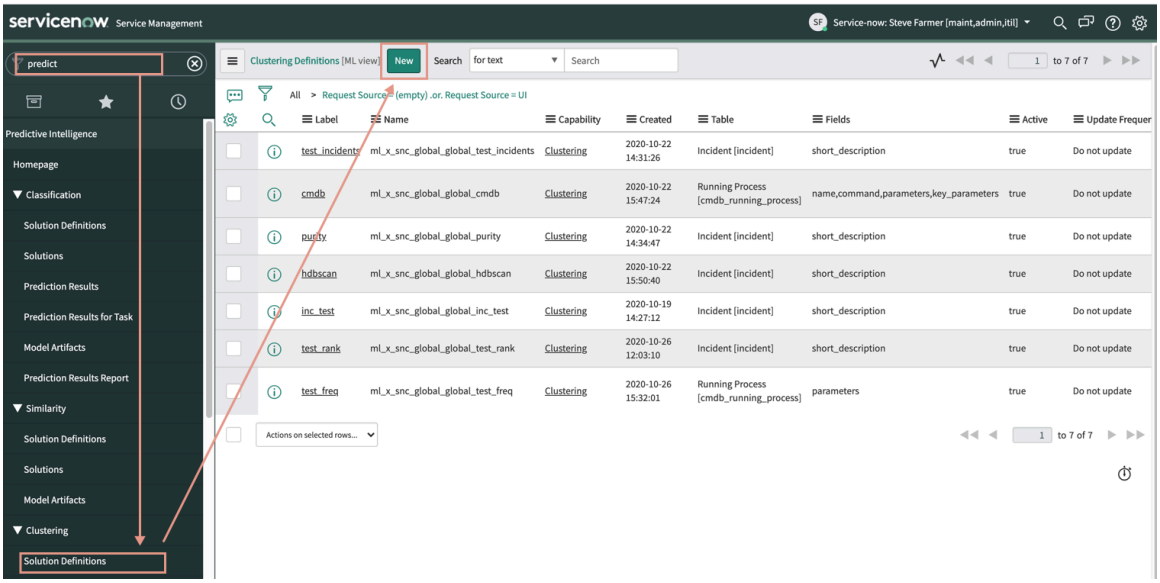
You can apply the HDBSCAN algorithm to help the system identify data samples that aren't assigned to any cluster. For example, you can apply HDBSCAN to support Topic Discovery.

Predictive Intelligence implements the k-means algorithm by default in its clustering framework. HDBSCAN is similar to the DBSCAN clustering algorithm except that it works with minimum-sized clusters and can help deliver more stable and persistent clusters. For a summary of how HDBSCAN works, see this article. For a comparison between DBSCAN and HDBSCAN, see this article and this article.

Note:
Clustering solutions trained with HDBSCAN do not support cluster updates. Updates on these solutions fail and the solutions are not logged in the ml_cluster_detail_table. Use DBSCAN or k-means training methods if you want to enable cluster updates.

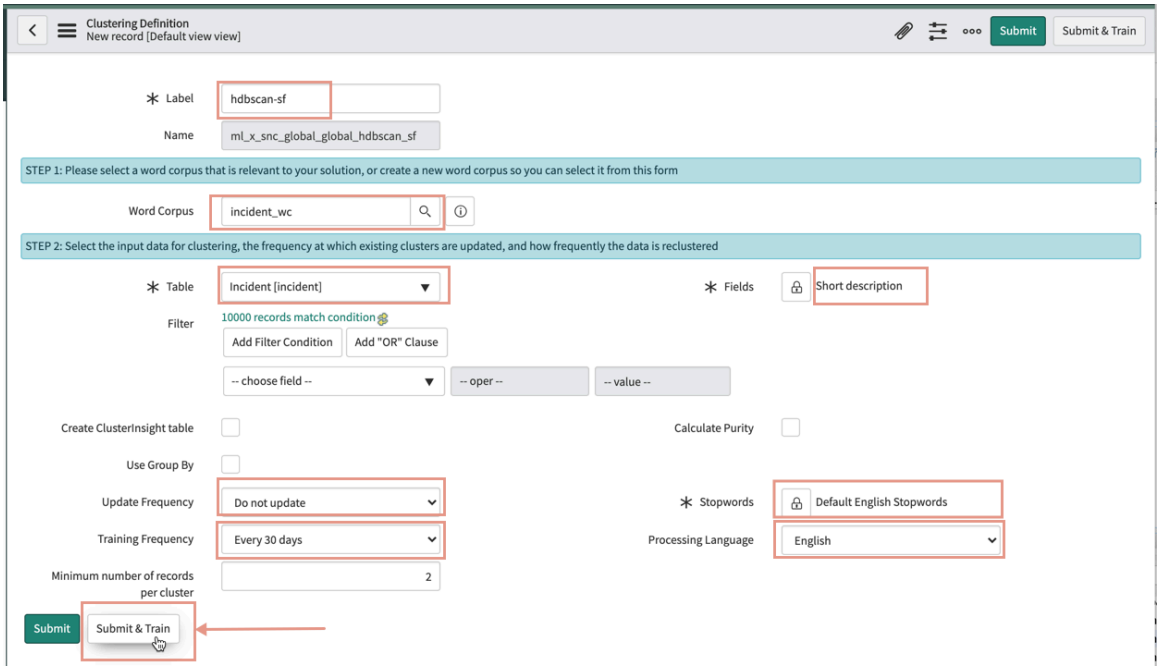
Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solution Definitions**.

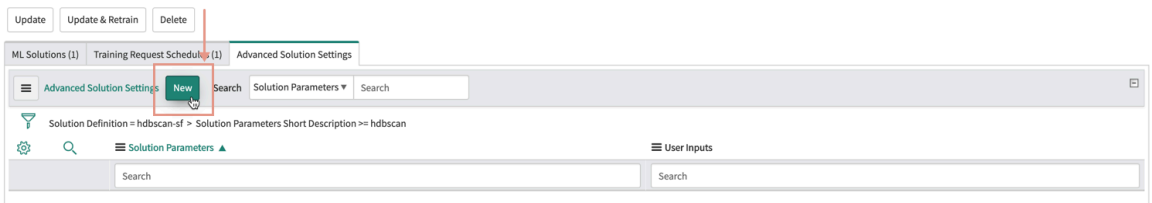


2. Select **New**.
3. Create a new clustering solution definition form or use an existing one.
   In this example scenario, you create the hdbscan-sf clustering definition form as in the image below. Configure the fields as follows:
   - **Label**: `hdbscan-sf`
   - **Word Corpus**: `incident_wc`, or any other word corpus that has incident record data (from the Washington DC release, a word corpus is not needed, so this field does not appear).
   - **Table**: Incident [incident]
   - **Fields**: Short description
   - **Update Frequency**: Do not update
   - **Stopwords**: Default English Stopwords
   - **Training Frequency**: Every 30 days
   - **Processing Language**: English

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. Select **Submit & Train**.
5. On the Advanced Solution Settings tab in the Related Links section of the trained form, select **Solution Parameters** from the picker, then select **New**.

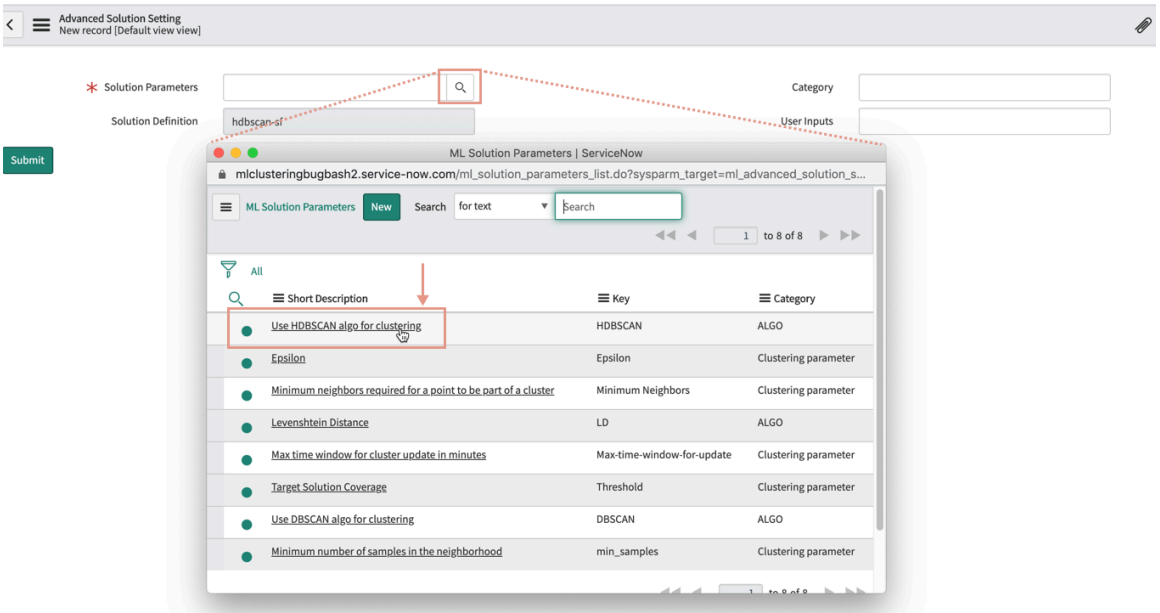


6. Create a parameter record.

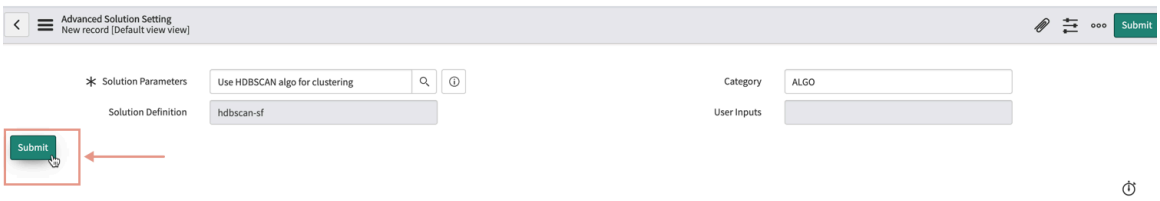
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. In the **Solution Parameters** field, click the search icon.

b. In the ML Solution Parameters screen, select **Use HDBSCAN algo for clustering**.



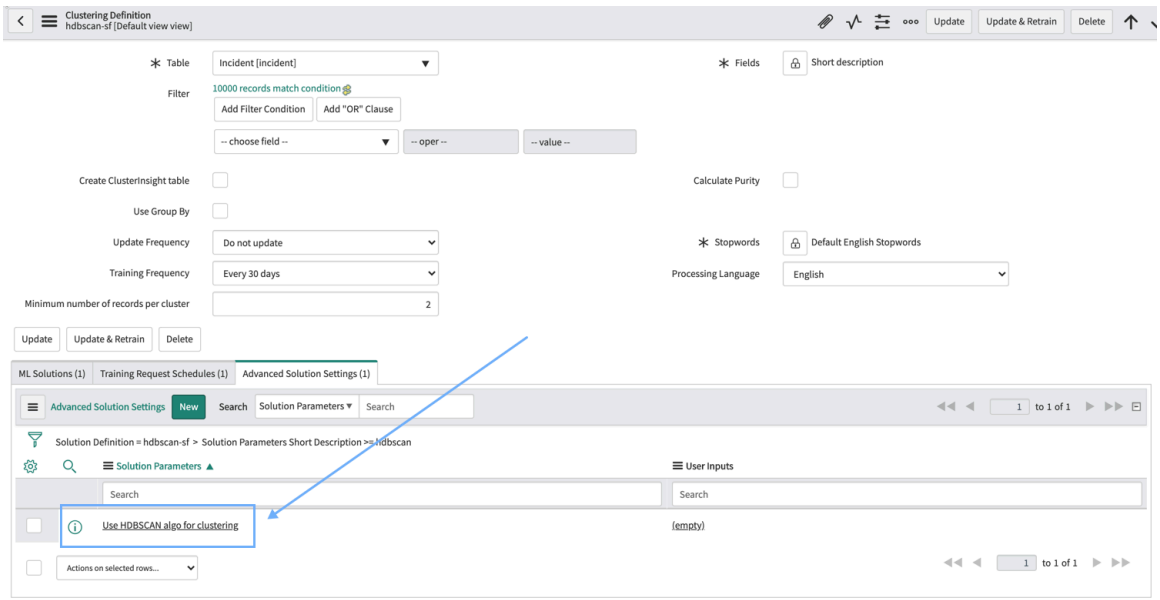
7. Select **Submit**.
The Advanced Solution Setting record appears with the HDBSCAN algorithm applied to the record. The field **User Inputs** is grayed out because it does not apply to this algorithm.



8. Select **Submit**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result: HDBSCAN is configured for your clustering solution. Its solution parameter appears on the Advanced Solution Settings tab of your clustering solution definition form.



Related topics

Create and train a clustering solution

Configure Connect Component algorithm and Levenshtein Distance method for a clustering solution

Apply Configure Connect Component and Levenshtein Distance method encoding to optimize the training for your clustering solutions.

Before you begin

Role required: admin or ml_admin

> **Note:**
> Configuring advanced settings on your ML solutions is optional. If you choose to configure any of these settings, make sure you're well informed regarding the technology you're enabling in the solution, and that your use case benefits from what the technology offers. For more information, see the Dive deeper with Clustering Advanced Parameters article on ServiceNow Community.

- Create and train a clustering solution definition or use an existing one.
- Role required: admin or ml_admin

About this task

When training clustering solutions, you have the following three options.

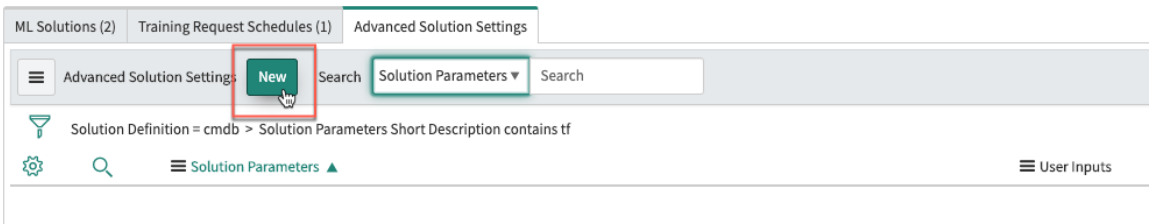
- Use the default k-means algorithm.
- Use the optional DBSCAN solution parameter with the Euclidean distance method as a metric.
- Use the optional DBSCAN, Minimum Neighbors, and Levenshtein Distance solution parameters. Connect Component is enabled by DBSCAN and Minimum Neighbors, and supports both Paragraph Vector-based text and Levenshtein Distance-based text. If you train your solution using the Levenshtein Distance method, you don't need to use a word corpus in your clustering solution.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

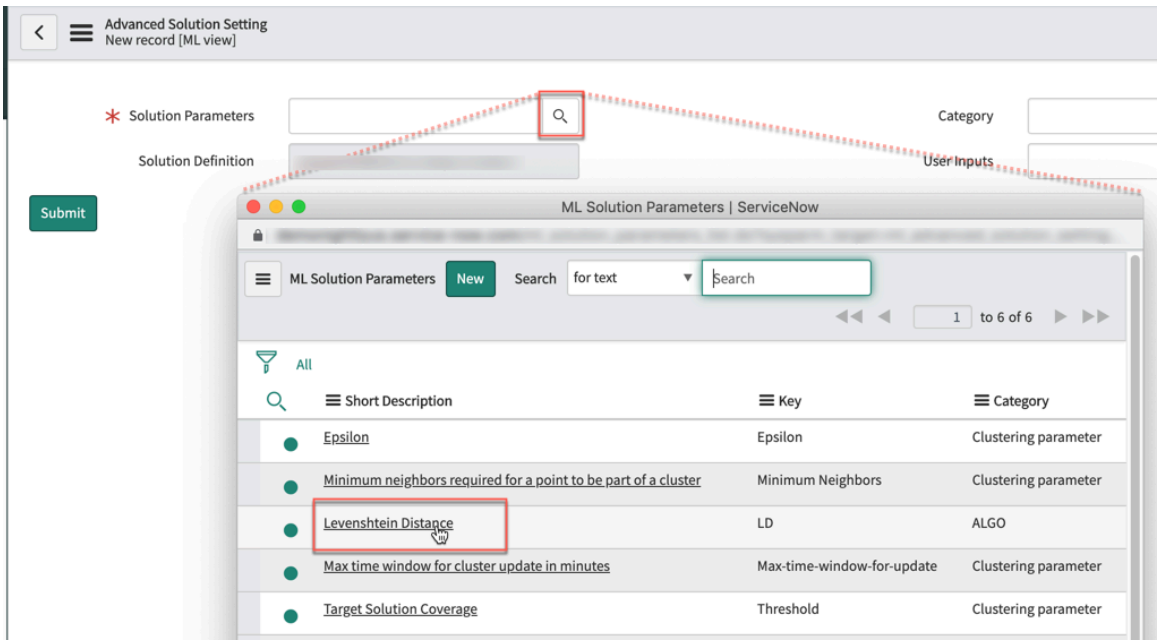
In this example scenario, you train your solution definition by using the third option referenced above.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solution Definitions**.
2. Open a trained clustering solution definition form.
3. On the Advanced Solution Settings tab in the Related Links section of the form, select **New**.

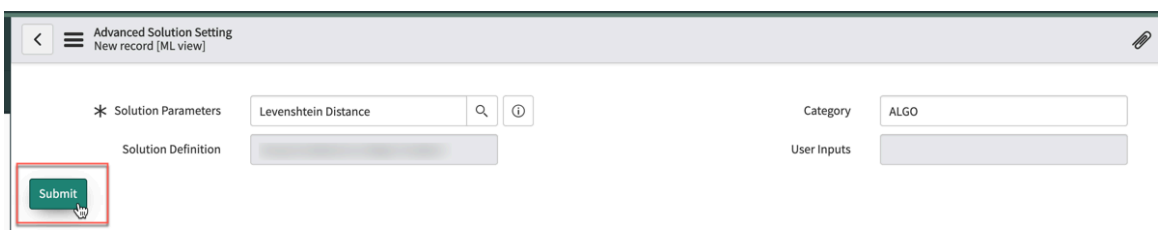


4. Create a parameter record.
   a. In the **Solution Parameters** field, select the search icon.
   b. In the ML Solution Parameters screen, select **Levenshtein Distance**.



5. Select **Submit**.

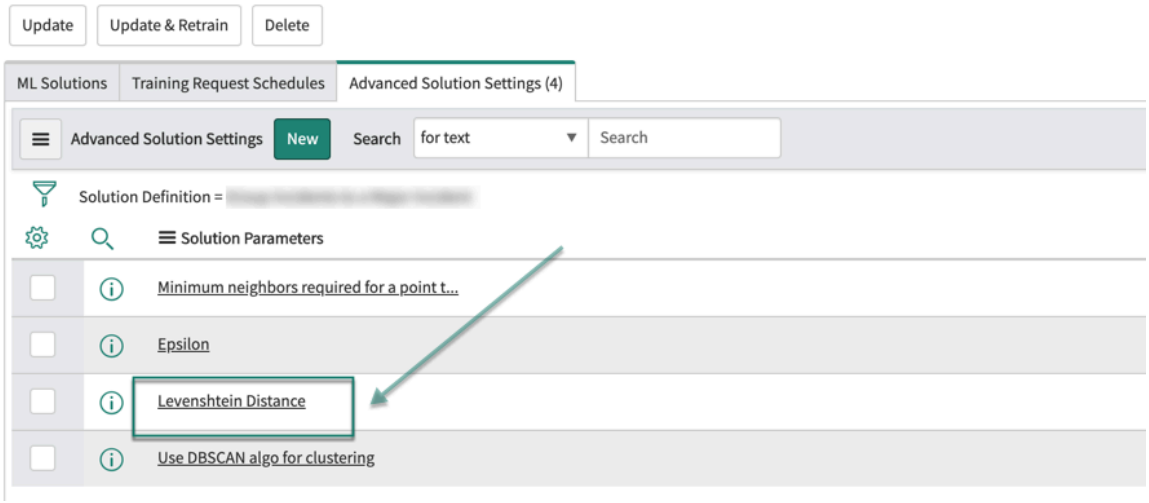
   The Advanced Solution Setting record screen refreshes.



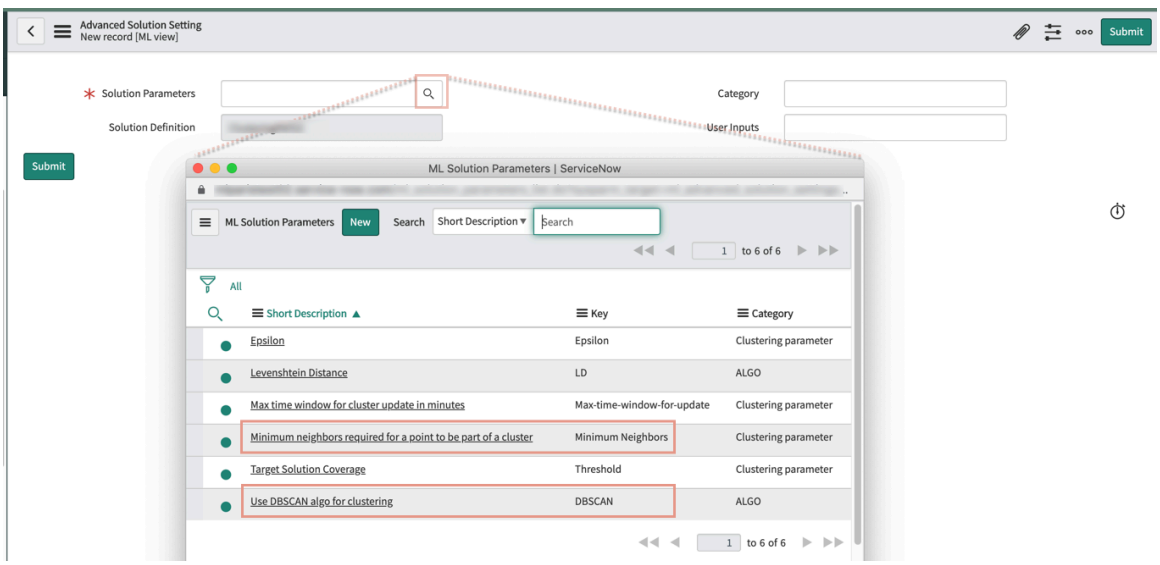
6. Select **Submit**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

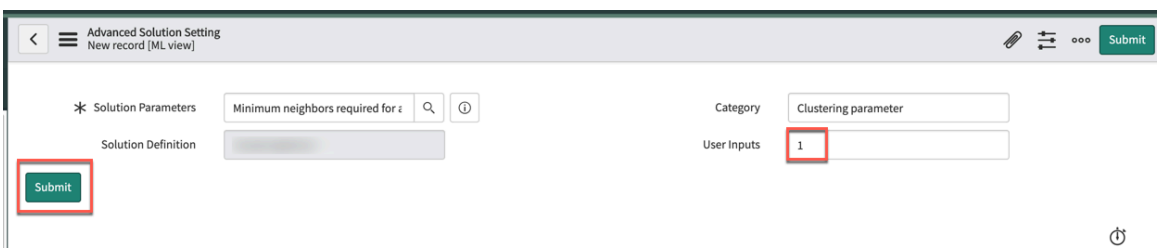
Result: Levenshtein Distance is configured for your clustering solution. Its solution parameter appears on the Advanced Solution Settings tab of your clustering definition form.



7. Repeat steps 1-6 from the previous Levenshtein Distance example, except this time you're creating the **Minimum Neighbors** and **DBSCAN** solution parameters, which together enable the Connect Component feature.



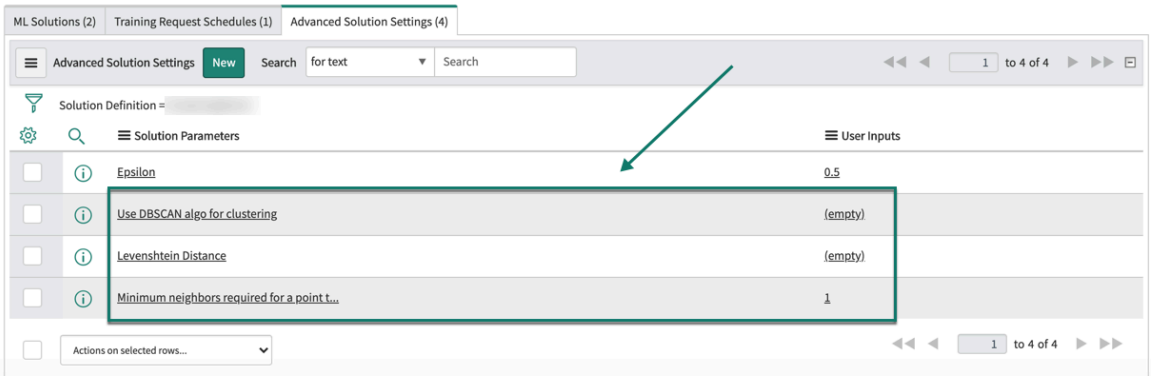
When you select, configure, and submit the **Minimum Neighbors** solution parameter, be sure to set the **User Inputs** field with a value of 1. Only some parameters have a User Inputs field.



Result:

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Connect Component is configured for your clustering solution. Its two solution parameters appear on the Advanced Solution Settings tab of your clustering definition form, alongside the Levenshtein Distance parameter you configured in steps 1-6 of this procedure.



Related topics

Create and train a clustering solution

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:

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Important:
With the Washington DC release, support for creating new regression solutions was removed. You can still edit and train existing regression solutions, but you won't be able to create new ones.

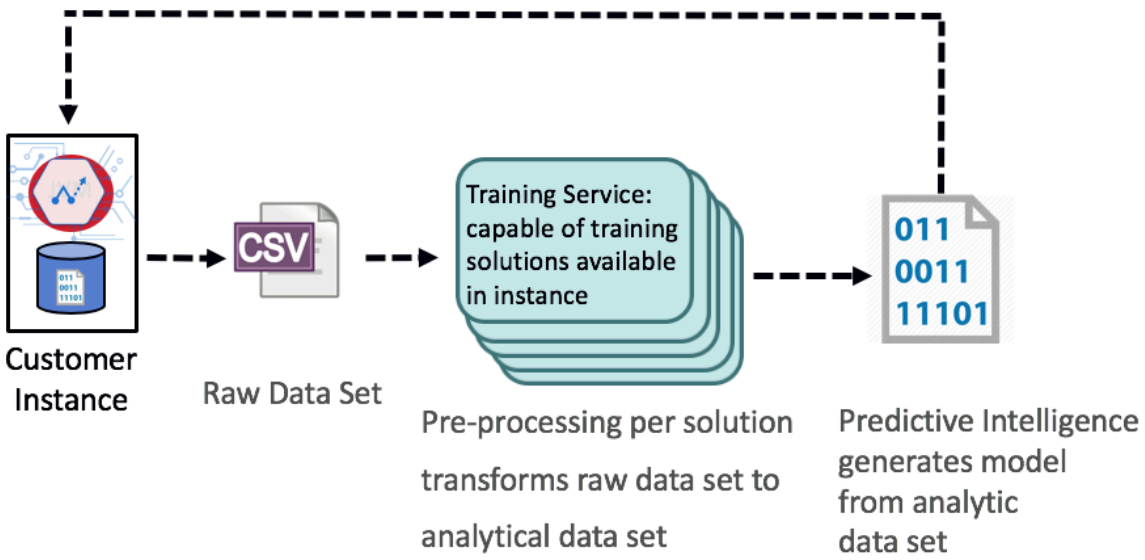
Uses historic data to predict numeric outputs, such as estimating the time it takes to resolve an incident or case. See Create and train a regression solution.

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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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 Washington DC release, models in the classification, clustering, and similarity frameworks use Workflow solutions. These are pre-trained, so a word corpus isn't needed for your new solutions. When your existing solutions with a word corpus are re-trained after upgrading, they become Workflow solutions, and the Word Corpus field is removed from the form.

About this task

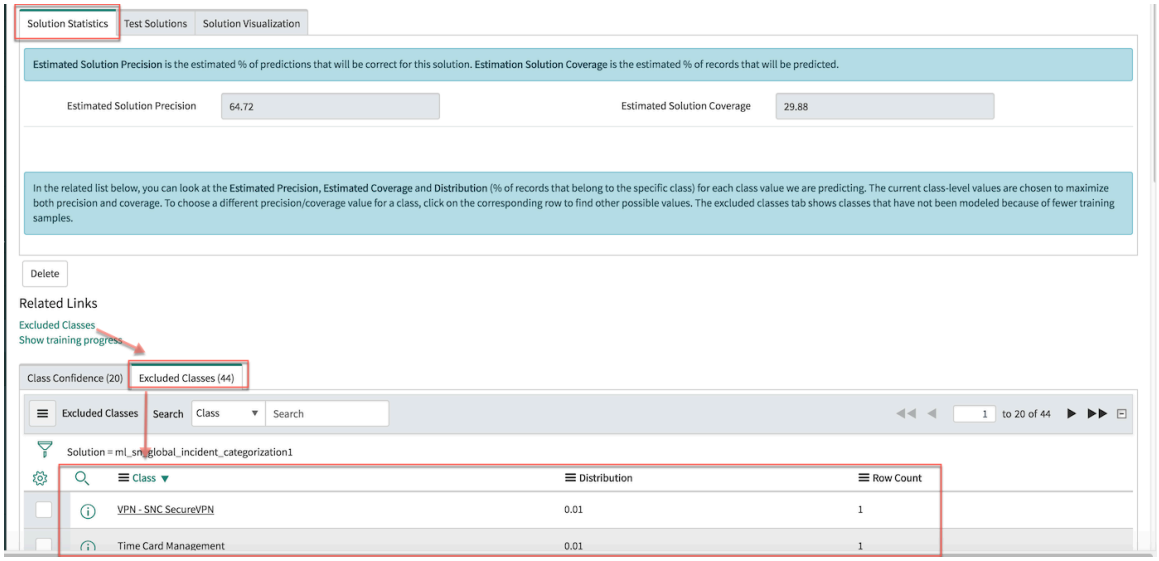
A predictive model is only as good as the data that you use to train it. To select appropriate records for training, examine the table's database dictionary as well as the current quality of the record values that you want to use.

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.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



You 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 Create a word corpus. **Note:** Word Corpus is not a required field for customers implementing Predictive Intelligence for the first time starting in Utah. A pre-trained model is used instead. The Word Corpus field is removed for pre-trained models.
Table	Select the table containing the target records that you want the system to predict.
Output Field	Select the field whose value you want the predictive model to set. In general, a good output field has these characteristics. ○ It's a choice field or a string field with a finite set of possible values. ○ It has some causal connection to the input fields.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	For example, on the default **Incident Categorization** solution definition, the output field is set to **Category**.
Fields	Select the input fields that you want the solution to use to generate a prediction. Input fields are fields within a record that may contain the classification information your prediction solution needs to succeed. For example, if you're predicting the correct class for triaging an incident record, the prediction should gather records containing text that references the class. Most records have contextual text in their **Short description** field, so it's a great input field to use in general. You could also use **Resolution notes** as an input field, as it too might reference the incident class in the detailed notes for the incident. In general, good input fields have these characteristics. ○ The fields are available to users when creating records. ○ The field data type can be string, reference, choice, or HTML. The more information that a field provides, the more often a solution can make a prediction, and the more often predictions are accurate. ○ The field has a default value and should not be blank. All default solution definitions use the **Short description** field.
Filter	Click **Add Filter Condition** to apply conditions to the records you're training. For example, the **Incident Categorization** solution definition uses a filter with these conditions: **[Created][on][Last 12 months] AND [Active][is][false] AND [State][is one of][Resolved | Closed]** To train a solution, the filter must return at least one record. If your filter returns no records, update it until it returns records for training. **Note:** The recommended number of records for training a good solution is from 30,000 through 300,000. If you submit more than 300,000 records, the most recent 300,000 records are used to train the solution. Use only authentic records from the database. In general, a good filter has these characteristics. ○ The training records are inactive and their states indicate work completed within your standard process, such as resolved or closed. ○ The target fields contain only correct values. Filter out records with unreliable target field values. For example, if you're predicting the assignment group/ category and your historic incident data contains assignment groups/ categories that are no longer used, add a filter to remove such records from the training. ○ The training records contain multiple examples of each target field value you want the solution to predict. ○ The training records include common variations of the input fields.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	Use relative date filters such as last 3 months or last 12 months. Don't use hard-coded dates because these filters aren't updated when the solutions are retrained, unless you update them manually.
Processing Language	Select the dominant language of the dataset that you're training on the solution definition. If the dataset language is Italian, choose `Italian`. Also, English processing is applied to all datasets by default. For example, if you select Italian, the system processes the data in both English and Italian. **Note:** The term processing indicates some of the language-specific steps used as part of training a solution. For example, tokenizing words, removing stop words, and stemming.
Stopwords	When you select your processing language, the system automatically adds a Stopwords list for that language. For example, if your processing language is Italian, the **Default Italian Stopwords** list appears. The **Default English Stopwords** list is also included. If you create a custom stopwords list, you can select it from the Stopwords field to add to your solution.
Training Frequency	Select how often the system regenerates the solution. The available options range from **Run once** up to **Every 180 days**. **Note:** The minimum number of records required for classification solution training is set at 10,000. By default, the system runs training once. This provides you time to review and update the solution definition until it provides acceptable coverage and precision values. When your solution definition is fairly stable, consider scheduled trainings, as data can age over time, thus degrading the accuracy of your prediction model. **Note:** The ML scheduler limits the number of trainings an instance can commit to 50 new ML training requests per instance within a 24-hour window. This limit excludes scheduled retraining requests, clustering updates, and similarity updates, even if the new training requests exceed 50 within a 24-hour window.

4. Click the appropriate context menu option or button for your solution definition.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

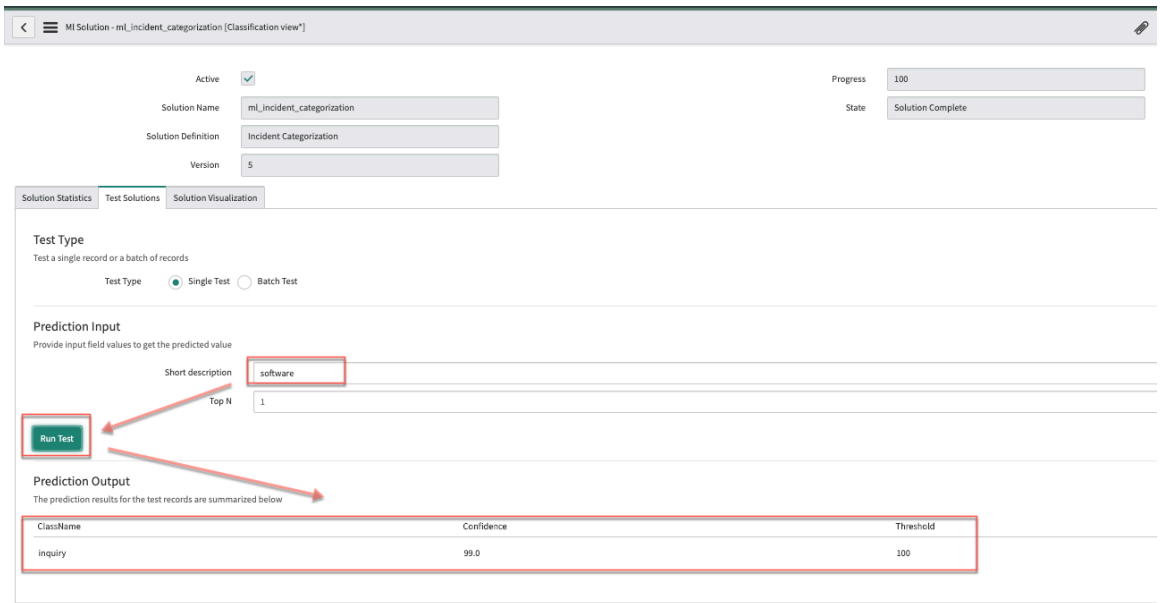
- A bubble chart populates the Solution Visualization tab of your solution form, showing the estimated precision and coverage for each of the classes covered by the solution. The size of the bubble indicates the percent of records (distribution) that belong to the class. When you point to a bubble, you 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

- 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

This 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.

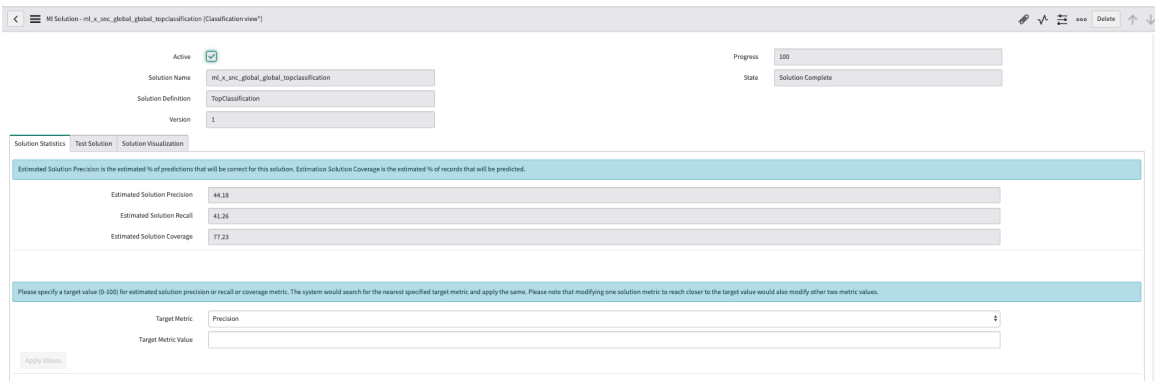
When you apply metric values at the solution level, the system automatically sets the appropriate values at the class level.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Here are the 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.

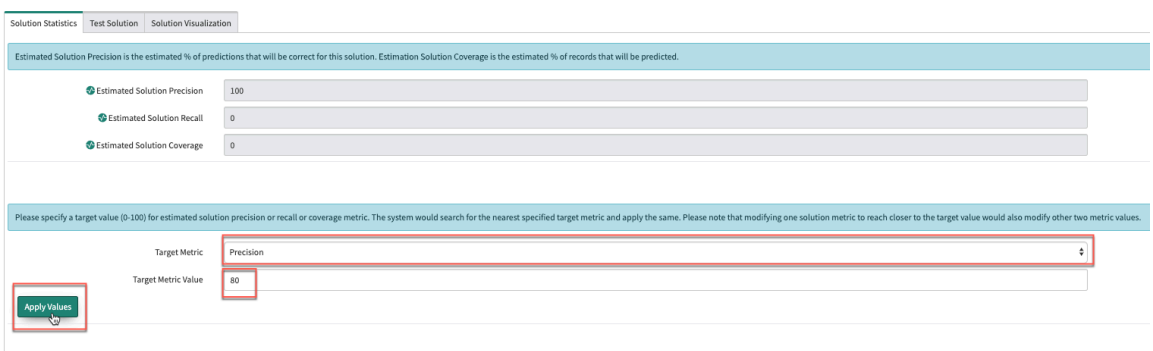
Here's a sample landing page for a recently trained classification solution. As you can see, the precision metric is 44.18, recall is 41.26, and coverage is 77.23.



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.

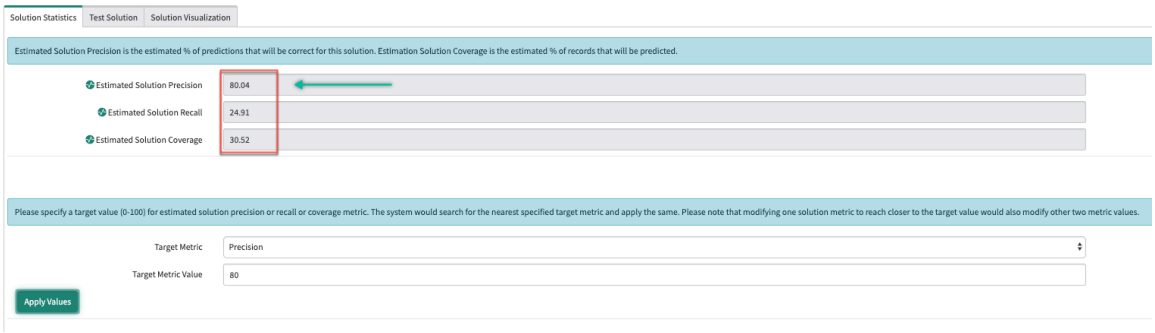
Precision configuration example

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**.



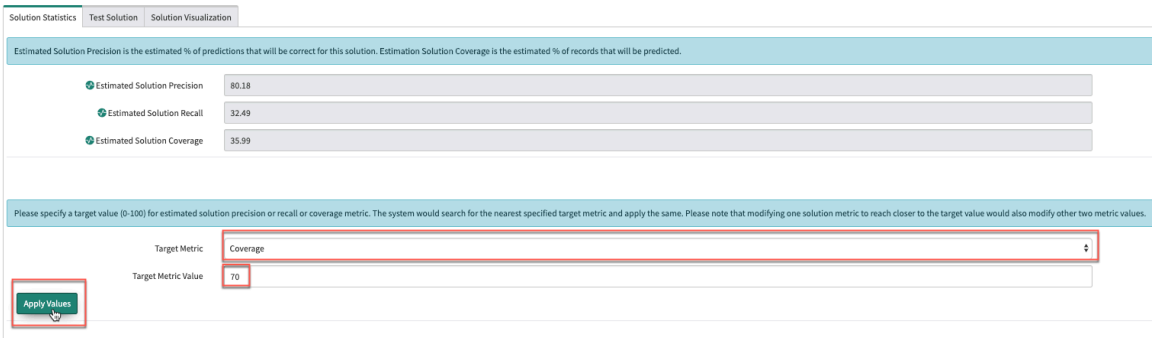
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Here are the metric values the system applied to the solution. In this scenario, the precision value of 80.04 slightly exceeded your request for 80%, so you're likely satisfied with that value.

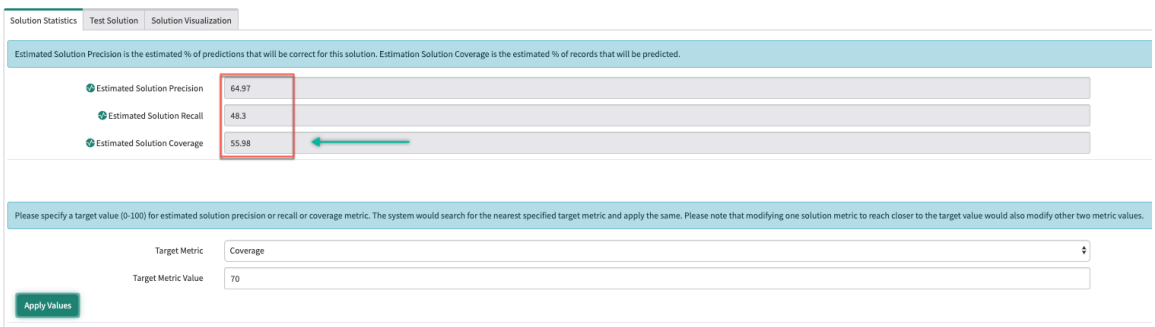


Coverage configuration example

In another example scenario where you're replacing a manual triage process for routing incident records, your minimum goal is to predict at least 70% of incoming incidents in the first quarter of the year. So you set the coverage metric value to 70 and click **Apply Values**.



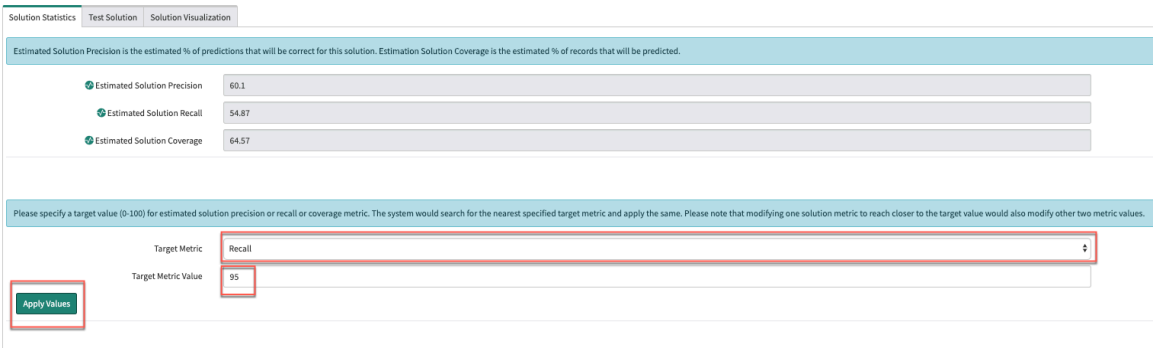
The metric values the system applied to the solution are shown in the following image. The coverage metric value increased from 35.99 to 55.98. However, the precision metric decreased from 80.18 to 64.97. This could be because you set the coverage metric to a relatively high value of 70, or perhaps because of how the data you're training is distributed.



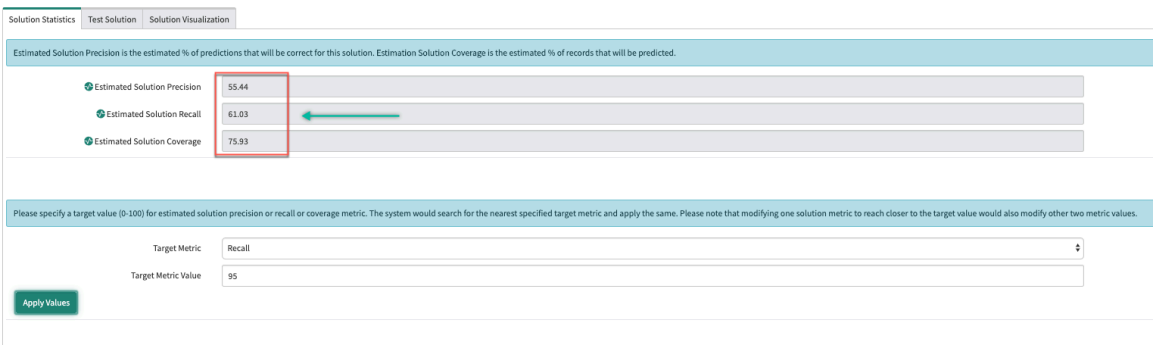
Recall configuration example

In another scenario, classifying if an incoming email is a Phish or not can be an important use case in a security-related machine learning solution. In this situation, it's very important to identify every Phish, and it may be okay to report a non-Phish as a Phish occasionally. However, no real Phish should be classified as a non-Phish. In such situations, the recall metric must have a high value, which might lead to lower percentages for precision and coverage. So here you can set the recall metric to 95 and click **Apply Values**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



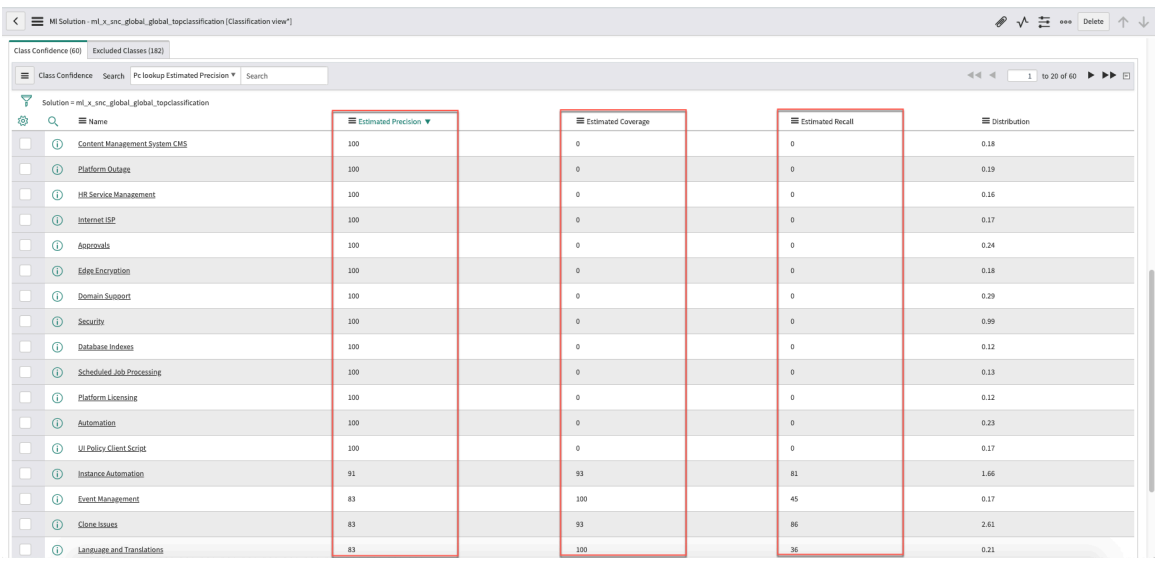
Here are the metric values the system applied to the solution. The recall metric value increased from 54.87 to 61.03. However, the precision metric decreased from 60.1 to 55.44. This is likely because you set the recall metric to the high value of 95.



Class-level results for the solution metric values you apply to your solution

The following image shows an example of the class-level results the system applied to a solution's precision, coverage, and recall statistics for 37 classes. You can keep modifying the metric values until you're fully satisfied with the results.

By Sorting (z to a) on the Estimated Precision column you can see which classes have the highest precision for the solution.



Using Group By for classification

Use APIs to simultaneously submit multiple classification solutions for training based on the Group By field.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can use 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();
```

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

```
// 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

Important:
In the Washington DC release, models in the classification, clustering, and similarity frameworks use Workflow solutions. These are pre-trained, so a word corpus isn't needed for your new solutions. When your existing solutions with a word corpus are re-trained after upgrading, they become Workflow solutions, and the Word Corpus field is removed from the form.

About this task

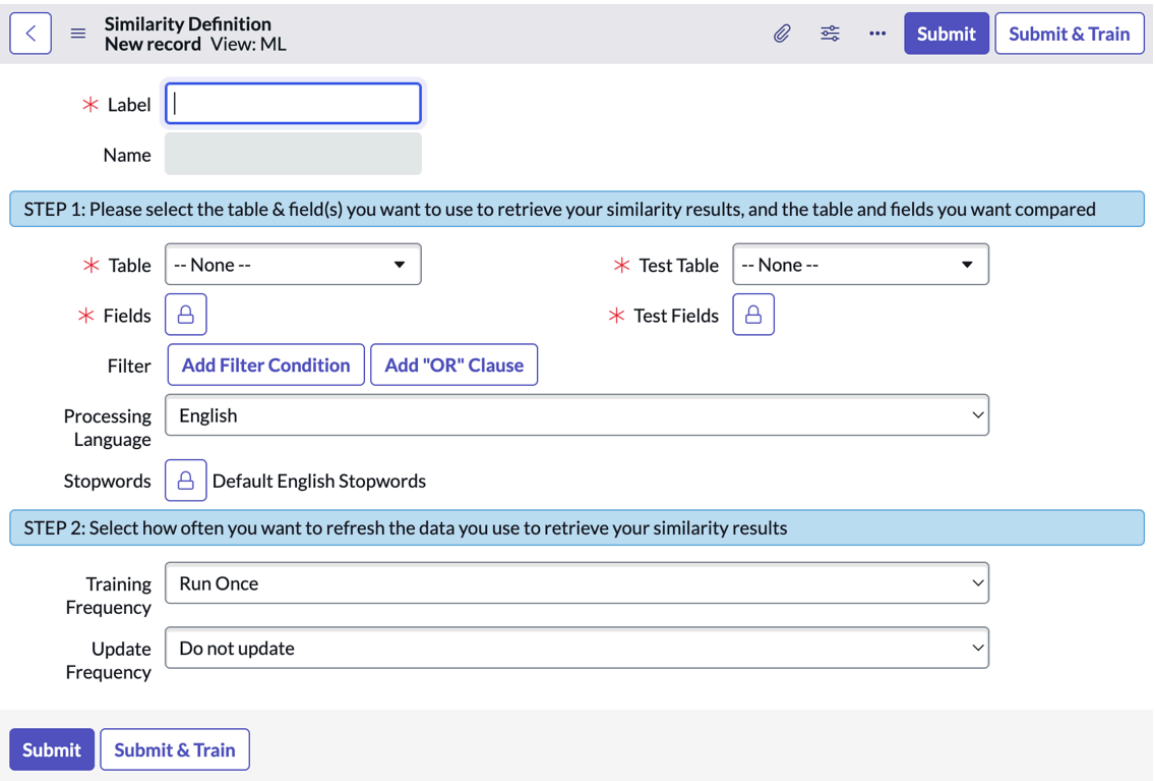
After comparing your existing records based on similarity, the system recommends examples that you can review and reuse in your solution.

When applied in your forms and flows, similarity solutions are domain-aware, so records from other domains on the instance are not displayed to users. For more information, see KB article Similarity prediction behavior in domain separated environment on Now Support.

For information on using encrypted training data, see Data Encryption in Predictive Intelligence.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In this example procedure, you're working on Incident records and you want to locate relevant Knowledge Base articles that could provide resolutions to those incidents.



Procedure

1. Ensure that you are in the application scope that you want for your solution definition, then navigate to **All** > **Predictive Intelligence** > **Similarity** > **Solution Definitions**.
2. In the Similarity Definitions list, select **New**.
3. On the Similarity Definition form, fill in the fields.

Field	Value
Label	Enter a unique name for your similarity solution. For example, in this use case you could enter `Match Knowledge Articles to Incidents`.
Name	As you enter a Label value, this field automatically populates with a system-assigned, read-only name based on your Label value.
Word Corpus	If you have a legacy similarity solution, you can select a relevant word corpus from the **Word Corpus** field in the definition form. **Note:** Starting from the Washington DC release, a word corpus is not required because a pre-trained model is used instead. The **Word Corpus** field is not visible in the definition form for pre-trained models.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	For more information, see Create a word corpus.
Table	In the Table field, select the table that contains records you want to use as an information source. In this use case you select the **Knowledge [kb_knowledge]** table, because its KB Article records might provide information relevant to the Incidents that you're trying to resolve. After you assign a Table, the number of records matching your filter conditions is displayed as a link. Select this link to view the list of records.
Test Table	In the Test Table field, select the table containing the records that you want to target. In this use case, you select the **Incident [incident]** table, as it contains the Incident records that you're trying to resolve. **Note:** You can select the same table for Table and Test Table. For example: using filter conditions, you could collect information from recent Incidents to help with target Incidents.
Fields	For the Table that you selected, enter fields that are likely to contain words and phrases relevant to the Incidents you're trying to resolve. In this example, you choose **Short description** and **Article body**. Including the article body increases your chances of capturing informative details regarding the subject. **Note:** **Journal Type** is not a supported data type.
Test Fields	For the Test Table that you selected, enter fields that contain text that you want to compare to other similar records. In this example, you choose the **Short description** of the Incident records you're trying to resolve.
Filter	Select **Add Filter Condition** to apply conditions to the Fields records you're using as an information source. For example, in this use case you could set a **workflow_state=published** condition to retrieve published KB articles only. **Note:** Script includes can't be referenced from the Filter. Use database views as an alternative.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
Processing Language	Select the dominant language of the dataset you're training on. Also, English processing is applied to all datasets by default. For example, if you select Italian, the system processes the data in both Italian and English. **Note:** The term processing indicates some of the language-specific steps used as part of training a solution, such as tokenizing words, removing stop words, and stemming.
Stopwords	When you select your processing language, the system automatically adds a Stopwords list for that language. For example, if your processing language is Italian, the **Default Italian Stopwords** list appears. The **Default English Stopwords** list is also included. To use a custom stopwords list, select the lock icon( ) and then search in the **Select target record** field.
Training Frequency	Select a retraining frequency. The available options range from **Run Once** up to **Every 180 days**.
Update Frequency	Select how often you want to refresh the data you use to retrieve your similarity results. For example, for open incident records, you could select an update frequency of **Every 15 minutes**, as new incidents typically occur frequently throughout the day. This frequency may increase the likelihood that newly opened records are included in the refresh. However, for KB Knowledge article records, which are typically not created often, you could choose a less frequent update frequency such as **Every 1 day**. **Note:** The ML scheduler limits the number of trainings an instance can commit to 50 new ML training requests per instance within a 24 hour window. This excludes scheduled re-training requests. In addition, clustering and similarity updates are also excluded from this limit, even if the new training requests exceed 50 within a 24 hour window.

4. Select the appropriate button for the solution definition.

5. If you submitted the solution for training, select **OK** on the **Training Activation** window to confirm.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result

- The system 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

Note:
In the Washington DC release, models in the classification, clustering, and similarity frameworks use Workflow solutions. These are pre-trained, so a word corpus isn't needed for your new solutions. When your existing solutions with a word corpus are re-trained after upgrading, they become Workflow solutions, and the Word Corpus field is removed from the form.

About this task

In this example procedure, you're creating a solution to identify a major incident by grouping similar incidents that have occurred recently.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For information on using encrypted training data, see Data Encryption in Predictive Intelligence.

Procedure

1. Ensure that you are in the application scope that you want for your solution definition, then navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solution Definitions**.
2. On the Clustering Definitions list, select **New**.
3. On the Clustering Definition form, configure the fields according to the following guidance.

Field	Value
Label	Enter a unique name for your clustering solution. For example, in this use case you could enter `Group Incidents to a Major Incident.`
Name	As you enter your solution Label, this field automatically populates with a system-assigned name based on your Label value.
Word Corpus	If you have a legacy clustering solution, you can select a relevant word corpus from the **Word Corpus** field in the definition form. **Note:** With the Washington DC release, a word corpus is not required, because a pre-trained model is used instead. The **Word Corpus** field is not visible in the definition form for pre-trained models. For more information, see Create a word corpus.
Table	Select the table that contains record types that you want to group into one or more clusters. For example, in this use case you select the **Incident [incident]** table as it contains incident records you want to group together for a major incident analysis. When you assign a table value, a link appears in the form that shows the number of records that match your current conditions.
Fields	Select one or more input fields types that help the system identify the records you want to include in your cluster. In this use case, use **Short description**. **Note:** When selecting a reference type field, you must dot-walk to the field's property name. For example, instead of `short_description`, enter `short_description.name`.
Use Group By	Select this check box only if you want to group input records by a field before creating clusters.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	Note: Selecting this check box activates the **Group By** list. If you don't select the check box, all the table records are grouped into clusters.
Group By	Selecting a value from this list is optional. If you do so, the system groups records into one or more clusters based on your selection.
Purity Fields	Choose fields from your table that can help the system identify the class that is most frequent in the cluster. In this example scenario, select **Category** and **Assignment group.Name**.
Filter	Add filter conditions to apply to the input field records that you want to include in your clusters. The number of records for clustering is limited to 300,000. **Note:** Script includes can't be referenced from the Filter. Use database views as an alternative.
Processing Language	Select the dominant language of the dataset you're training on the solution definition. If the dataset language is Italian, choose **Italian**. Also, English processing is applied to all datasets by default. For example, if you select Italian, the system processes the data in both English and Italian. **Note:** The term processing indicates some of the language-specific steps used as part of training a solution. For example, tokenizing words, removing stop words, and stemming.
Stopwords	When you select your processing language, the system automatically adds a Stopwords list in that language. For example, if your processing language is Italian, the **Default Italian Stopwords** list appears. The **Default English Stopwords** list is also included. If you create a custom stopwords list, you can select it from the Stopwords field to add to your solution.
Update Frequency	Select how often you want the system to update your clusters with new and updated records. **Note:** The system pulls records based on the Group By filter conditions that you set on your clustering solution, if any. For example, if you select **Every 15 minutes**, the system identifies which records have arrived within that time

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	frame. The system tries to assign them to the existing clusters, or creates a new cluster if possible. In this example, 20 new records arrive. If 16 of those records make it into an existing cluster and 4 don't, the system forms a new cluster for the four unassigned records. You can also choose not to update your clusters at all.
Training Frequency	Select how often you want the system to discard all previous cluster results and recreate clusters from the beginning. Your options range from daily, every third day, every seven days, or monthly. You can also choose to train your cluster once. **Note:** The ML scheduler limits the number of trainings an instance can commit to 50 new ML training requests per instance within a 24-hour window. The limit excludes scheduled retraining requests. In addition, clustering and similarity updates are also excluded from this limit, even if the new training requests exceed 50 within a 24-hour window.
Minimum number of records per cluster	Enter the minimum number of records you want a cluster to contain. The value you enter must be 2 or higher.

4. Select the appropriate context menu option or button for your solution definition.

5. If you submitted the solution for training, click **OK** on the Training Activation window to confirm.

Result

The system trains the solution and notifies you in real time when the training completes.

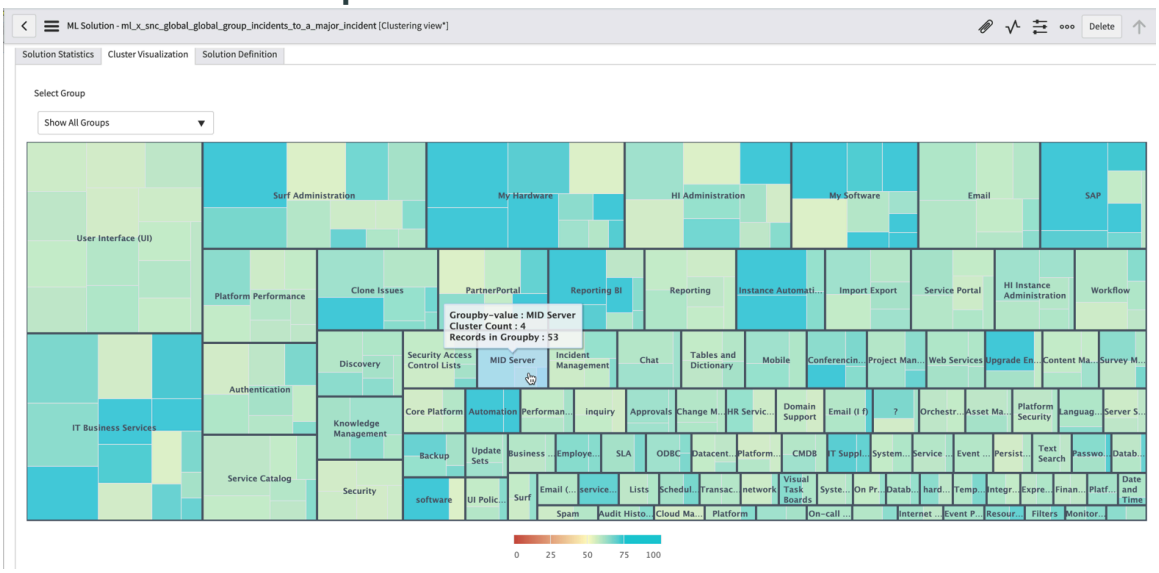
A treemap plot appears on the Cluster Visualization tab of your Clustering Solution Definition form. The plot shows the clusters the system formed for your solution in descending order from the top-left corner to the bottom-right corner. The treemap node labels are the Cluster Concept, which is created 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cluster visualization example



To open a cluster, you can click it, or select it from the Show All Groups filter.

Inside the cluster grouping, you can filter the results further by using the two slide bars for cluster size and cluster quality, respectively. You can also navigate backward by clicking the Back button, which only appears when a clustering hierarchy is present. When you point to a cluster at this level, the Purity field percentile values appear along with the Cluster Concept, Quality, and Size values.

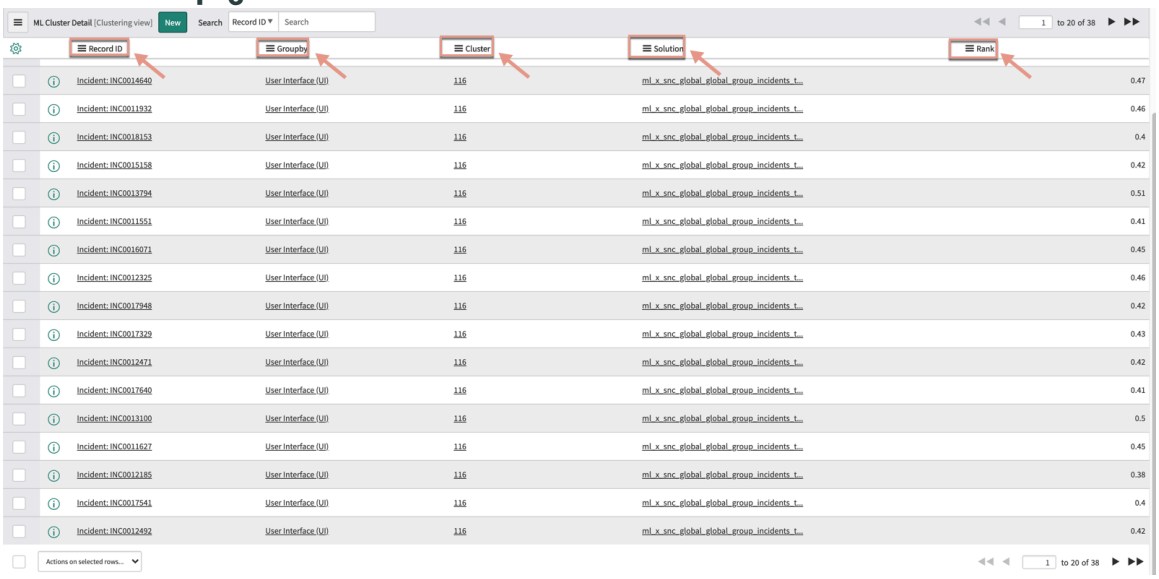
Cluster group example



When you click a cluster node, its ML cluster details appear in a list view format.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Cluster details page

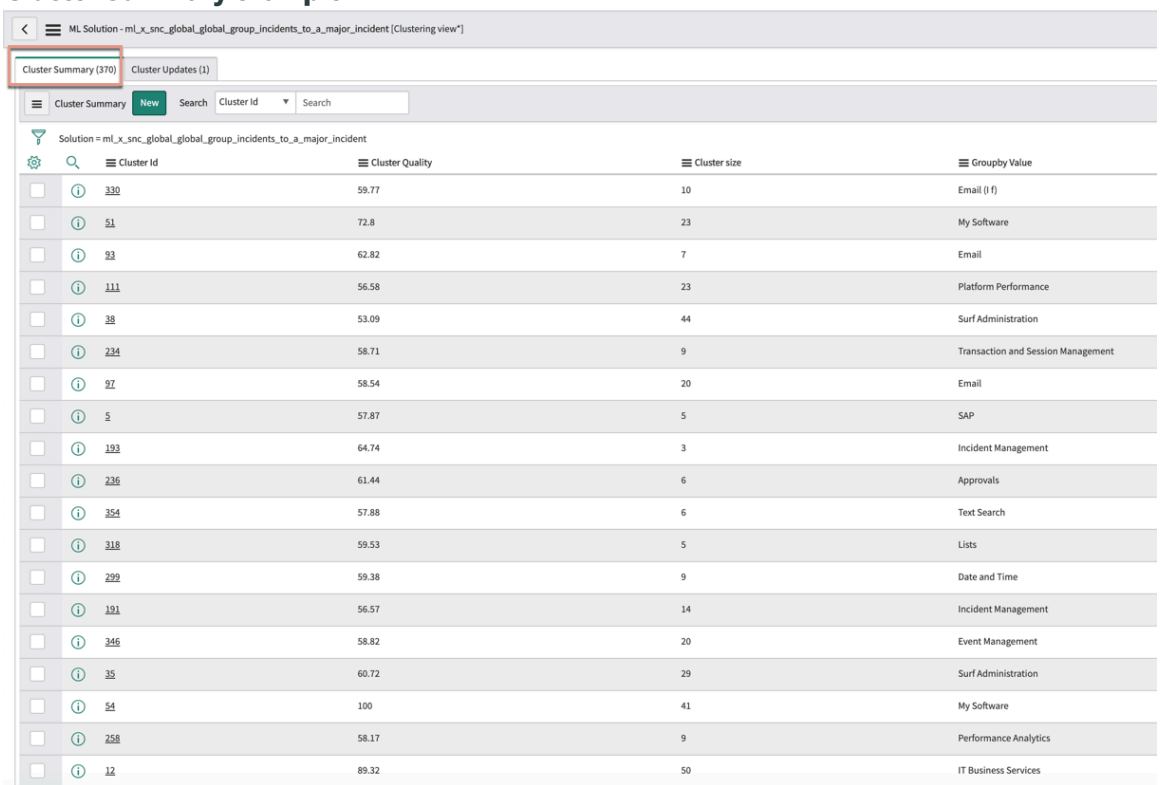


What to do next

- Review the solution output on the Solution Statistics tab of your solution. If you aren't satisfied with your clustering solution results, reconfigure the values you've set to your solution and retrain it until the results are to your satisfaction.

- Review the Cluster Summary tab for a list view of the cluster IDs, quality size, and Groupby values.

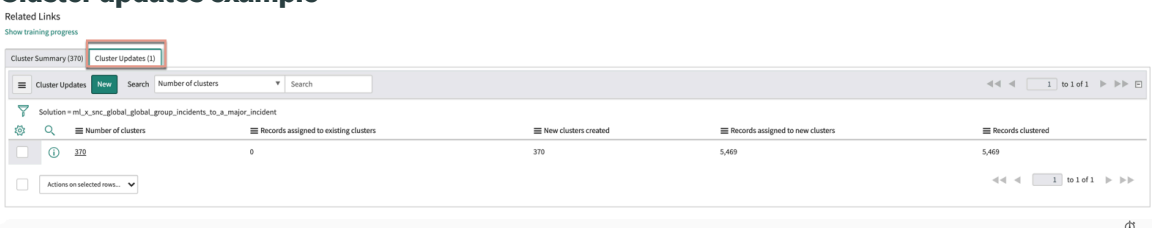
Cluster Summary example



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- On the Cluster Updates tab, review the summary of changes to clusters for each cluster update interval you configured in the solution definition.

Cluster updates example



Assign a name to a cluster

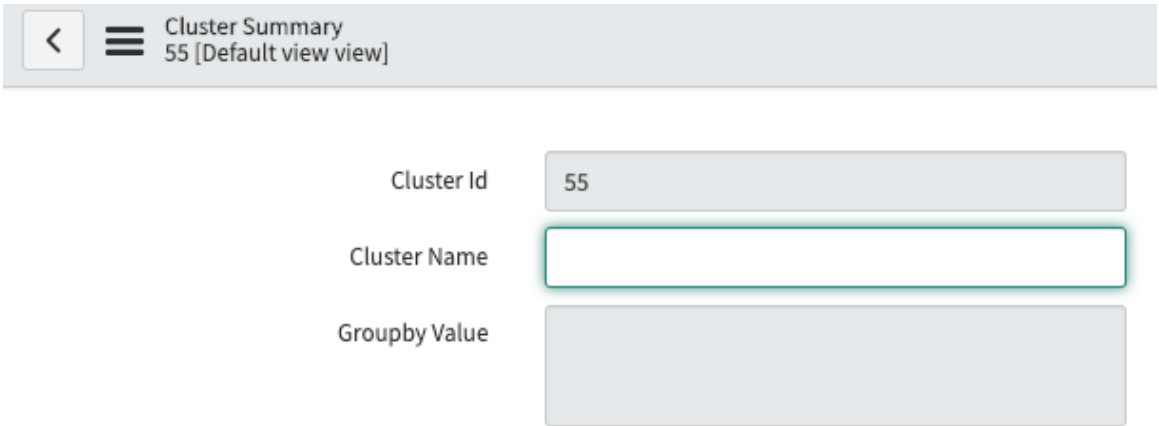
Name your clusters to help identify and organize them.

Before you begin

- Create a clustering solution definition or use an existing one.
- Role required: admin or ml_admin

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solutions**.
2. Select a solution.
3. Under the **Cluster Summary** related list, select a cluster ID.
4. Enter a name in the **Cluster Name** field.



5. Select **Update**.

Generate a representative sample of a cluster

View the top 25 most representative records of a cluster.

Before you begin

- Create a clustering solution definition or use an existing one.
- Role required: admin or ml_admin

About this task

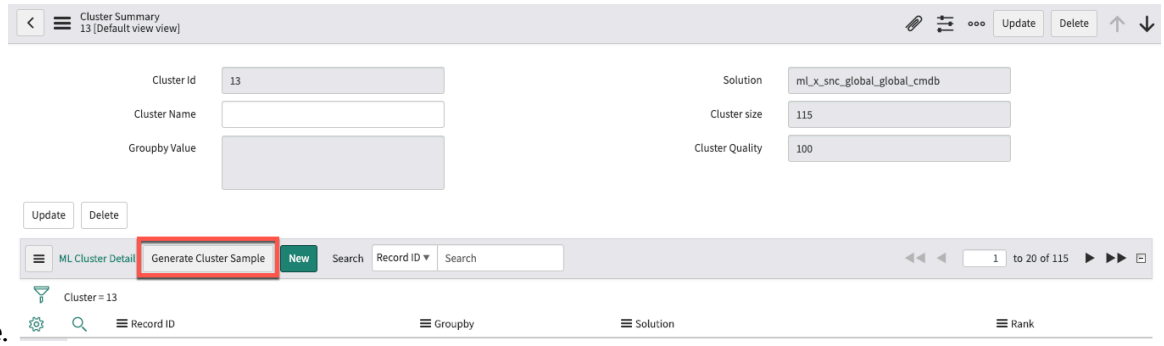
Generating a sample of a cluster filters the records for that cluster. Generating a sample works on clusters with over 25 records.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solutions**.
2. Select a clustering solution.
3. Under the **Cluster Summary** related list, select a cluster ID.
4. Select **Generate Cluster**

Sample.



Result

The **ML Cluster Detail** list shows the top 25 records in the cluster. The cluster sample also applies to the cluster visualization.

Apply purity on a clustering solution

Apply purity to learn details about the composition of each cluster. For example, see what percentage of incidents in a cluster have the same value for Assignment Group. You can specify which fields to focus on, or you can let auto-purity display fields by default.

Before you begin

- Create a clustering definition solution or use an existing one.
- Role required: admin or ml_admin

About this task

With purity, you can choose field distribution metrics for your clusters to display. The table assigned to your solution determines which fields are available. For more information about purity, see the following articles on ServiceNow Community:

- https://www.servicenow.com/community/intelligence-ml-articles/using-purity-fields-to-better-understand-your-clusters/ta-p/2302441
- https://www.servicenow.com/community/intelligence-ml-articles/predictive-intelligence-using-the-cluster-insight-table-to/ta-p/2301006

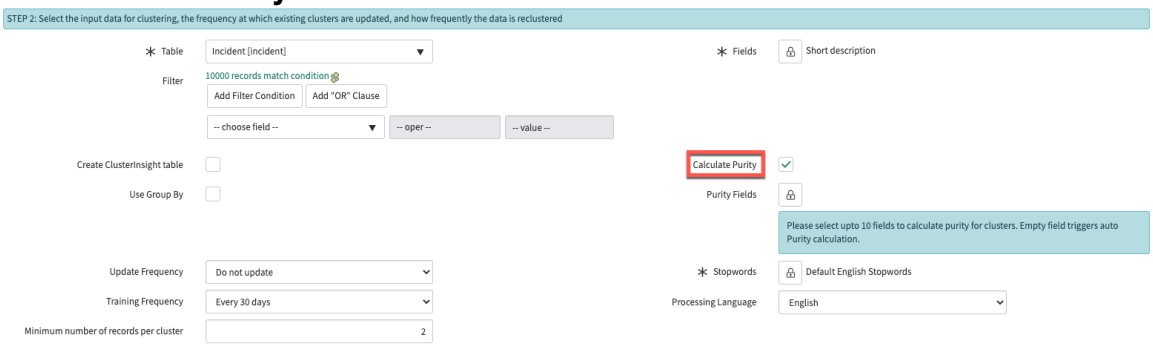
If you do not select any fields, auto-purity automatically determines which insights to show based on distribution significance. You can change the default auto-purity selections by editing the ML Auto Purity Configuration [ml_autopurity_config] table.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Clustering** > **Solution Definitions**.
2. Select a solution definition or create a new one.

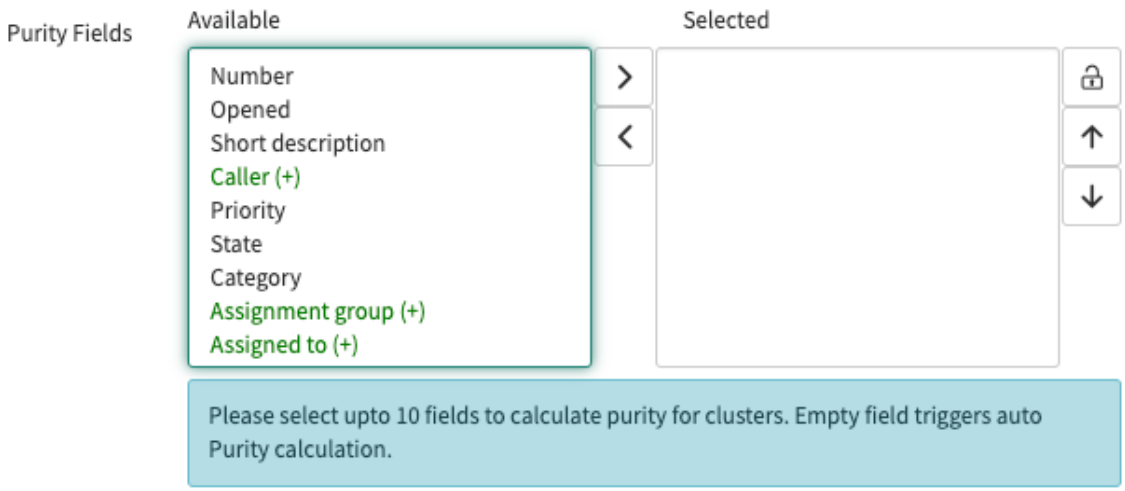
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. Select the **Calculate Purity** check box.



4. **Optional:** Select the **Purity Fields** lock icon and choose the fields that you want displayed in cluster results.

> **Note:**
> If you don't select any fields, auto-purity displays the most significant fields based on distribution.



5. Select **Update** or **Update & Retrain**.

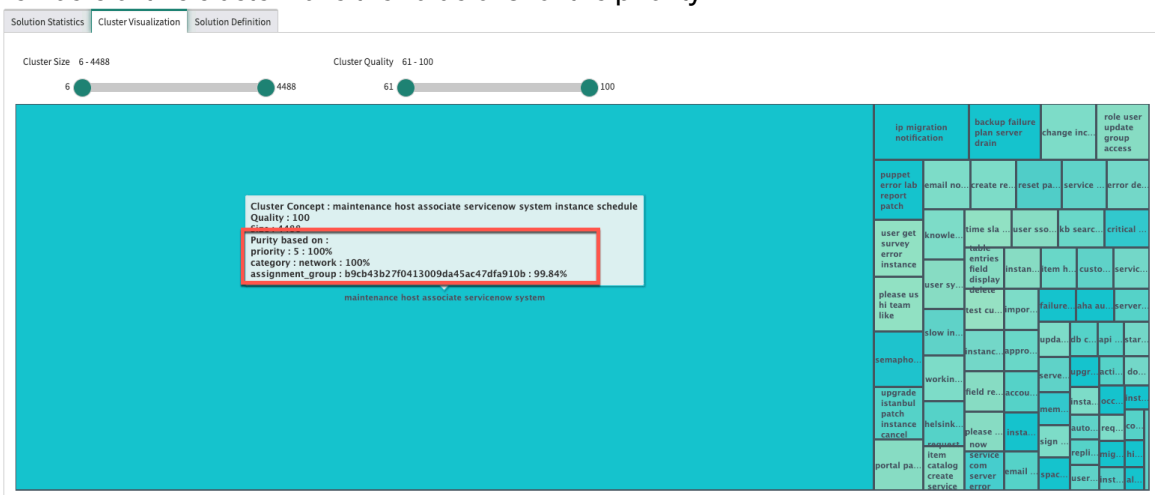
Result

View your insights by opening the **Cluster Visualization** tab of your retrained solution, then pointing to a cluster. Under **Purity based on**, the format displays as:

field name : value : percent.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, the row priority : 5 : 100% means that 100% of the members of this cluster have the value of 5 for the priority field.



Analyze a cluster with a data source

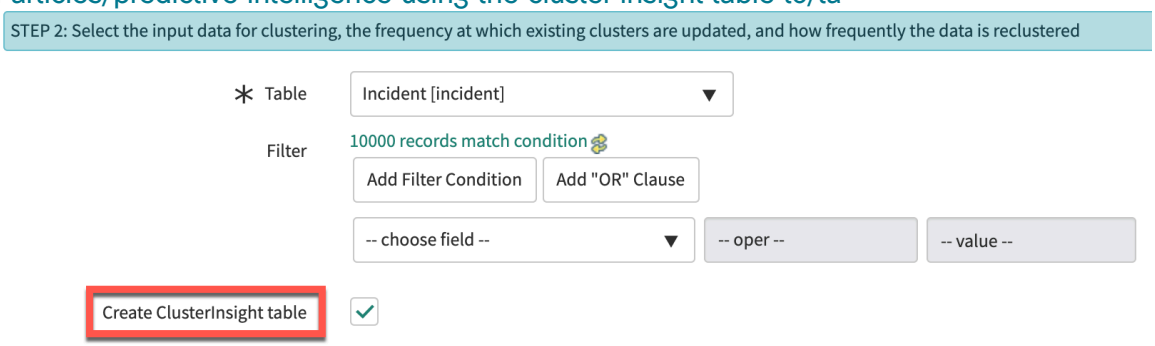
Analyze a cluster by any of the dimensions available on the source table.

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/intelligence-ml-articles/predictive-intelligence-using-the-cluster-insight-table-to/ta-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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create and train a regression solution

Train your solution by using historical data to predict numeric outputs, such as temperature or 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 Washington DC release. You can edit and train any existing solutions, but you can't create new ones. This information is provided for legacy context.

Role required: ml_admin or admin

About this task

Regression solutions enable you to predict a point estimate and prediction interval. The resulting model delivers the following statistics:

- Mean Absolute Error (MAE), which measures the mean deviation of a predicted value from the actual value. This metric is useful as it's easy to understand as its scale is the same as that of its target. However, MAE is unbounded, making it difficult to compare across models.
- Symmetric Mean Absolute Percentage Error (SMAPE) is a percentage value of the deviation from the predicted to the actual. SMAPE is a bounded version of MAE except that it has a value range between 0 and 100. The lower the SMAPE value, the better the model accuracy.
- Range Accuracy is the percentage of actual values between a predicted range. In other words, it's the range between the upper and lower bounds of the prediction. For example, if four out of five 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, select **New**.
3. On the Regression Definition form, configure these fields according to the following guidance.

Field	Value
Label	Enter a unique name for your regression solution. For example, in this use case you could enter `Regression Test for DB Restore`.
Name	As you enter your solution Label, this field automatically populates with a system-assigned name based on your label value.
Word Corpus	Select an existing word corpus that's relevant to your solution. For example, in this use case you select a word

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	corpus that has a title such as **Incidents in the last 3 months**. If you don't have a relevant word corpus, follow the steps to create a word corpus first. When the word corpus is complete, you can select it from the Word Corpus field in your Regression Definition form. However, the word corpus selection is optional. If your input data has text columns and you don't choose a word corpus, your regression solution trains a new word corpus model by using the text columns in your input data. The resulting word corpus can be reused in any other regression solution or other ML solution type. **Note:** A pre-trained model is used instead of the Word Corpus for users who activated Predictive Intelligence starting in Utah. The Word Corpus field is removed for pre-trained models.
Table	Select the database table on which you're applying regression. The table should contain historical records the system can use to predict the length of time for its database restore.
Output Field	Select the field whose value you want the predictive model to set. In general, a good output field is a numeric, integer, or floating point field. In this example scenario, you use the **Duration** field to measure a length of time. The output field should generate a numeric value.
Fields	Select one or more field types that help the system identify the records you want to train using regression. In this example scenario, you use **Short description**, **Source datacenter**, **Target datacenter**, and **Database size**. (short_description, Sourcedc, Targetdc, and Dbsize.) Input field types can be string, nominal, or numeric.
Filter	(Optional) Add filter conditions to the output field records that you want to train using regression. **Note:** ○ The minimum number of records for regression training is 10,000 records. ○ The maximum number of records for regression training is limited to 300,000.
Processing Language	Select the primary language of the dataset that you're training on the solution definition. If the dataset language

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Value
	is Italian, choose **Italian**. Also, English processing is applied to all datasets by default. For example, if you select Italian, the system processes the data in both English and Italian. **Note:** The term processing indicates some of the language-specific steps used as part of training a solution. These steps include tokenizing words, removing stop words, and stemming.
Stopwords	When you select your processing language, the system automatically adds a Stopwords list for that language. For example, if your processing language is Italian, the **Default Italian Stopwords** list is displayed. The **Default English Stopwords** list is also included. If you create a custom stopwords list, you can select it from the Stopwords field to add to your solution. In this scenario, you use the **Default English Stopwords** list.
Training Frequency	Select how often the system regenerates the solution based on records matching the **Filter**. The available options range from **Run once** up to **Every 180 days**. In this scenario, you select **Every 30 days**. 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. **Note:** ○ The minimum number of records required for regression solution training is set at 10,000. ○ The ML scheduler limits the number of trainings an instance can commit to 50 new ML training requests per instance within a 24-hour window. This limit excludes scheduled requests for retraining. In addition, clustering and similarity updates are also excluded from this limit, even if the new training requests exceed 50 within a 24-hour window.

4. Click the appropriate context menu option or button for your solution definition.

5. If you submitted the solution for training, click **OK** on the **Training Activation** window to confirm.

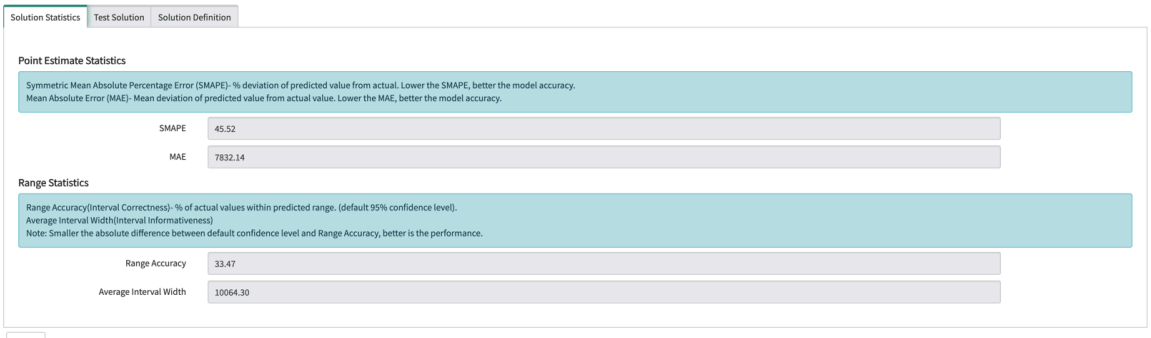
   The system schedules the solution for training with the nearest training service. The system sends you a notification when the training completes, including any errors that may have occurred in the training. Other users can subscribe to the Predictive Intelligence Notifications category. When training completes, the system uploads the solution as an Attachment record.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

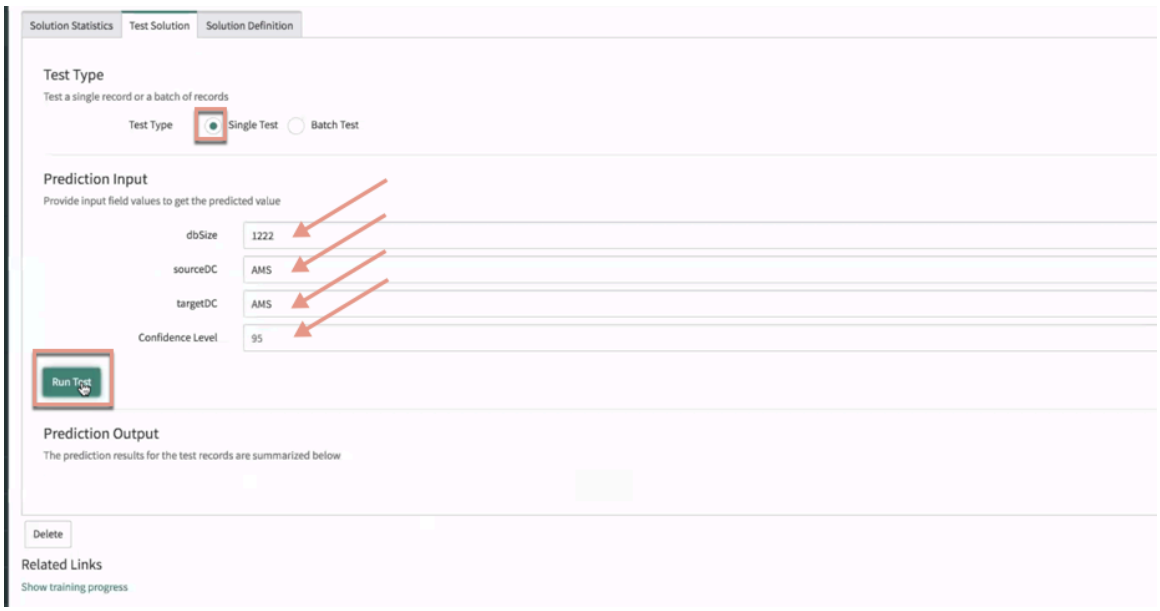
What to do next

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.

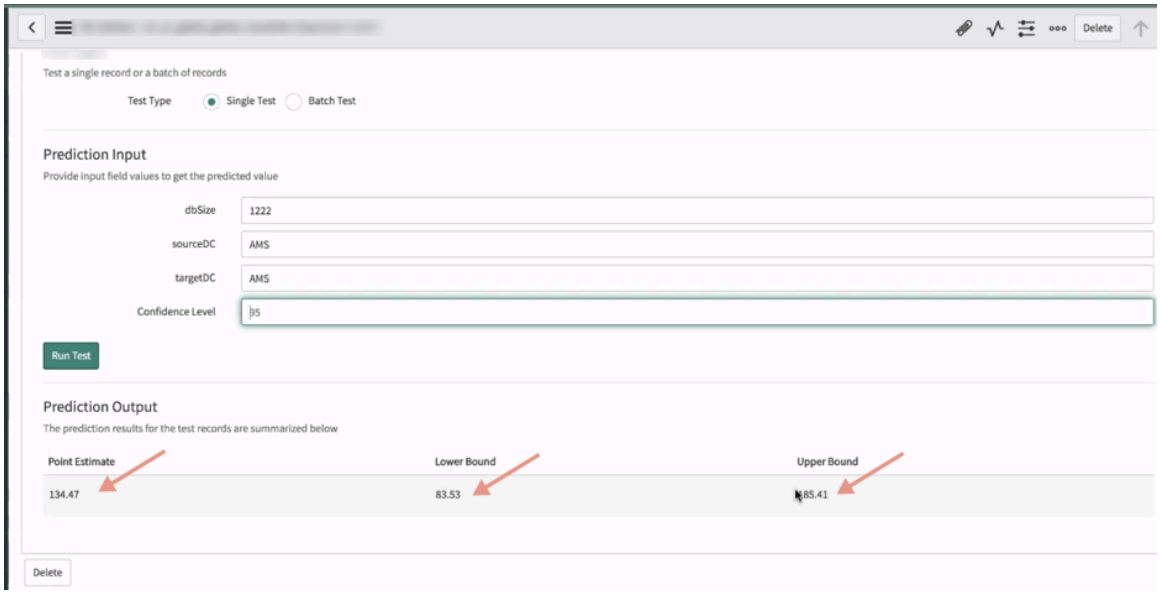


On the Test Solutions tab of your solution, you can test the prediction output for the records you used as input to the prediction by entering values for the input fields, such as the **Source datacenter**, **Target datacenter**, and **Database size**. You can also use the default prediction confidence 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.



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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



View 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.

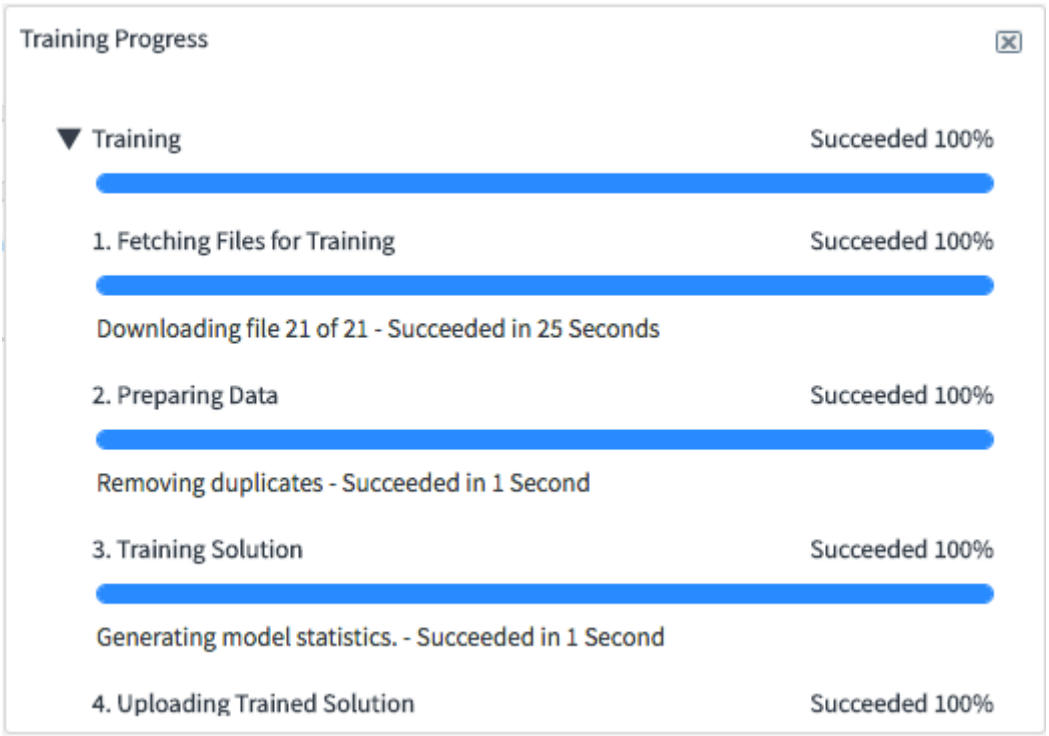
Procedure

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The system shows 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

Use the Solution Statistics dashboard to determine if a classification solution has sufficient precision and coverage for each class. Identify classes that require configuration or retraining with a new solution definition filter.

Before you begin

- To ensure optimal dashboard display, enable responsive dashboards or change the default dashboard layout.
- Role required: admin, ml_admin, or ml_report_user

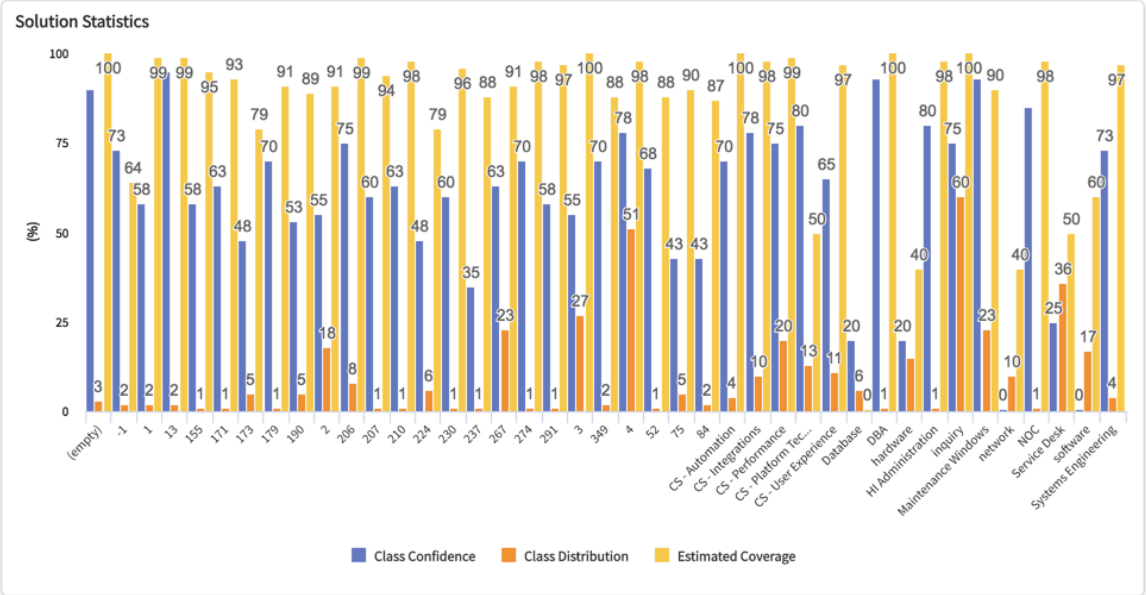
About this task

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. Note that the number of classes predicted may be less than 50, and can skip a class if there is not enough historical data to build a solution that can predict the class confidently.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



5. 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, it will already be active and ready to use if it's the latest version. 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **Predictive Intelligence** > **Similarity** > **Solutions**.
2. In the ML Solutions list, locate your solution 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

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For more information, see 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- 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 do not 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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- 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

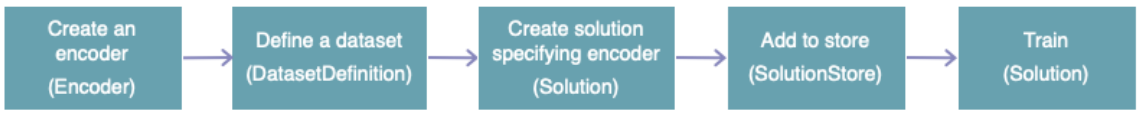


Note:
The encoder definitions support multiple dataset definitions, but have the same training flow.

To train a solution with an encoder, create the encoder first, then include the encoder in the solution configuration.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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()*

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. 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()*

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- *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"}
```

getLatestVersion()	getStatus()
ClassificationSolution - getLatestVersion()	*ClassificationSolutionVersion - getStatus()*

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

getLatestVersion()	getStatus()
ClusteringSolution - getLatestVersion()	*ClusteringSolutionVersion - getStatus()*
Encoder - getLatestVersion()	*EncoderVersion - getStatus()*
PredictabilityEstimate - getLatestVersion()	*PredictabilityEstimateVersion - getStatus()*
RegressionSolution - getLatestVersion()	*RegressionSolutionVersion - getStatus()*
SimilaritySolution - getLatestVersion()	*SimilaritySolutionVersion - getStatus()*

Get predictions using a solution version

After the solution has been trained, get the trained version and run a prediction on it. Get the solution you created from the store. Next, choose the trained version and predict the trained version.



Note:
Predictions cannot be made on encoders and predictability estimates.

1. Get the solution from the classification solution store using the *get()* method.

```
// Get the solution created in the first example from
 the classification solution store
var mlSolution =
 sn_ml.ClassificationSolutionStore.get('ml_x_snc_globa
l_global_my_solution_definition');
```

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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

ClassificationSolution - 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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Monitoring 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, you can 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, ml_admin, or admin

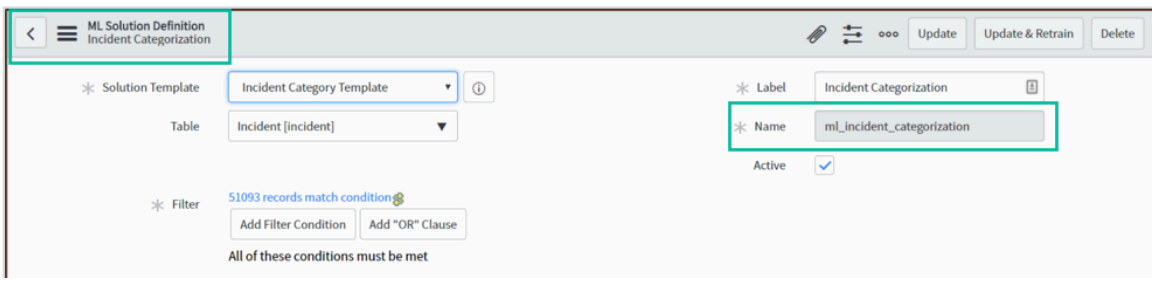
About this task

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.

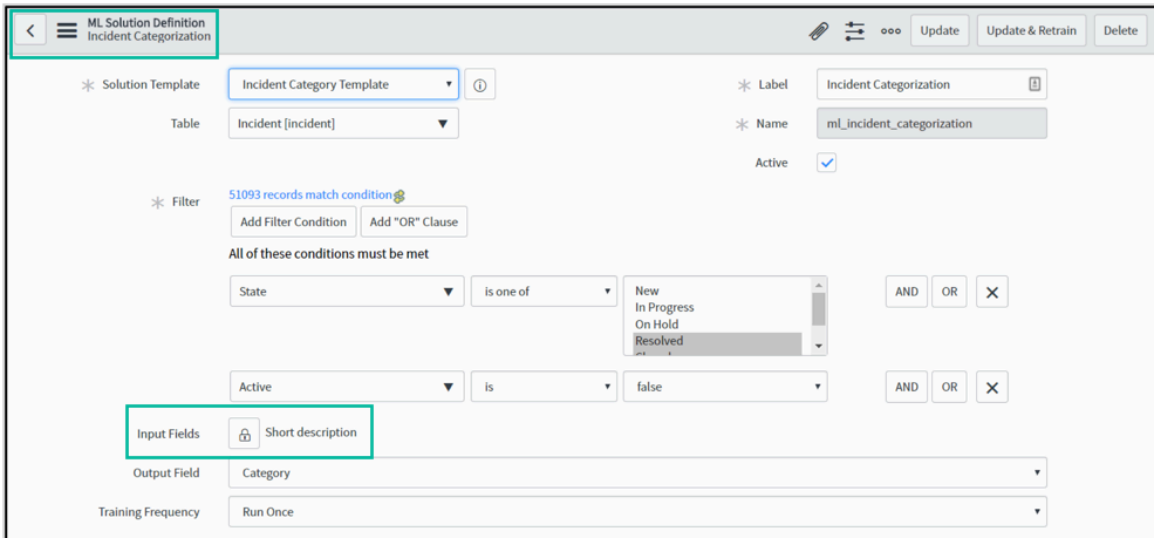
Procedure

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, you can use the **short _description** field, as the prediction model has been trained to use this field to learn its category definition.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. Navigate to **System Web Services** > **REST** > **REST API Explorer**.
5. Set these choice fields as follows.
   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 input is the Name value you captured in Step 1 of this procedure.

7. Select **Add query parameter**.
   The Predictive Intelligence form refreshes to show the **Query parameters** section.
8. Type `short_description` in the first field.

   Note:
   This field 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, type `Unable to connect`.
10. Select the **Send** button.
    The REST API Explorer sends your request to the Predictive Intelligence API.
    The system predicts the output value in the Response Body section of the API output. You can use other short descriptions to test what the solution is predicting.
11. **Optional:** Send a different request to the Predictive Intelligence API so that you can test the prediction model again.

    a. Return to the **Query parameters** section of the Predictive Intelligence form.

    b. Type a short description that references a different kind of incident in the second field.
    For example, type `Unable to connect to MSSQL`.

    c. Select the **Send** button.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The 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 an 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 doesn't 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.

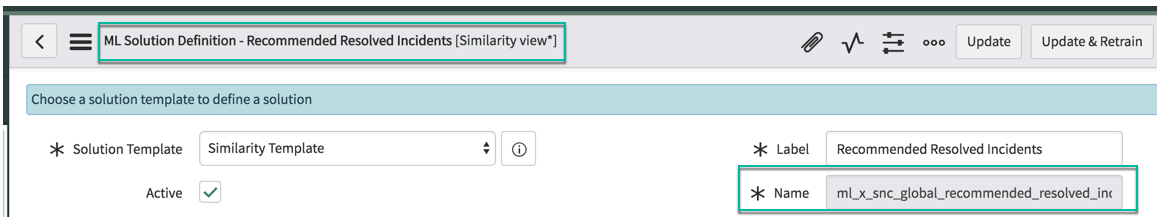
This scenario illustrates a similarity solution prediction for a hypothetical ML solution that you have previously created and trained. You can also use the REST API Explorer to test a classification solution prediction.

Procedure

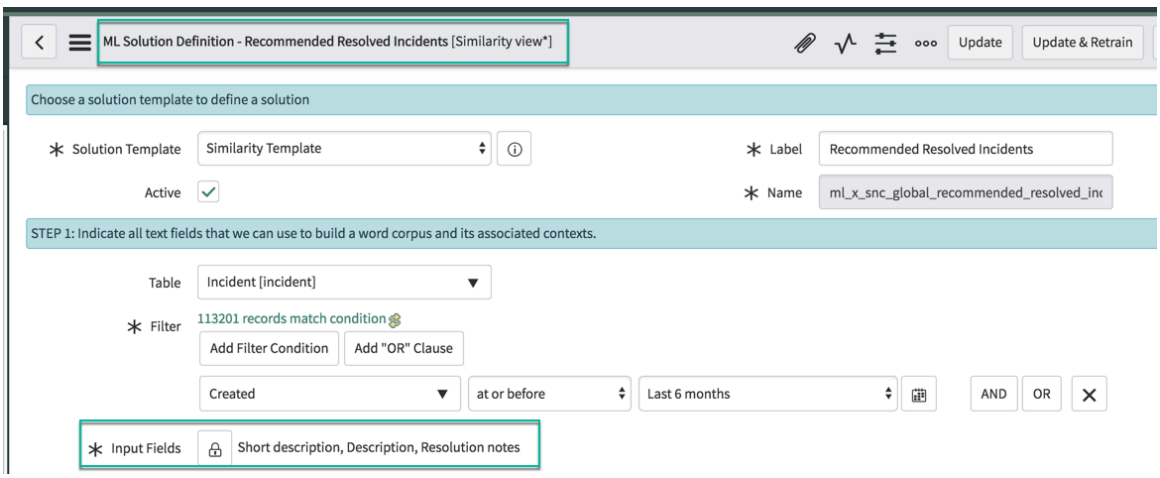
1. Navigate to **All** > **Predictive Intelligence** > **Similarity** > **Solution Definitions**.
2. Locate the ML solution definition whose prediction you want to test, and copy its **Name** value to your clipboard or a Notepad file.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

In this 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.
   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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can use 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.



14. In your original browser tab, navigate to **Servicedesk** > **Incidents**.

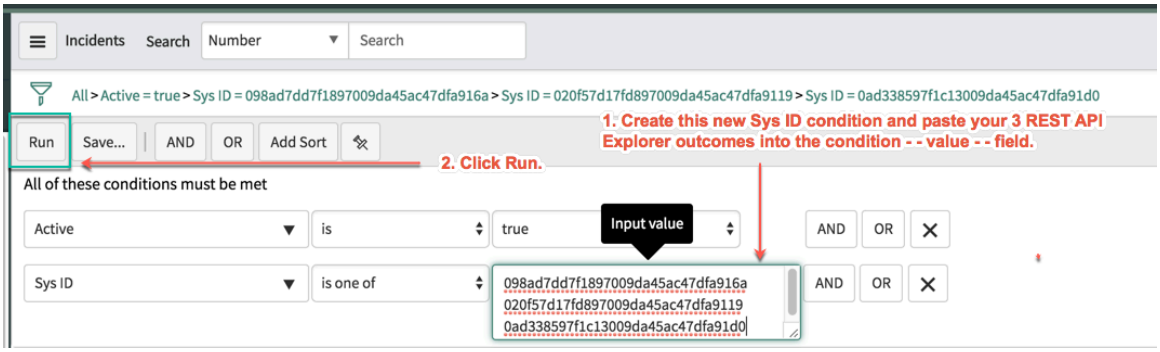
15. As shown in the image below, set filter conditions for the three REST API outcomes to the Incidents table list view.

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.

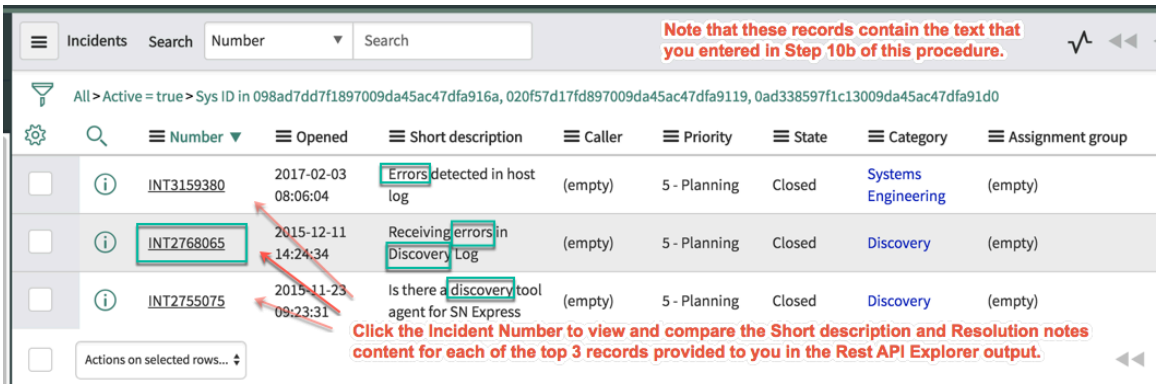
c. Click **Run**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

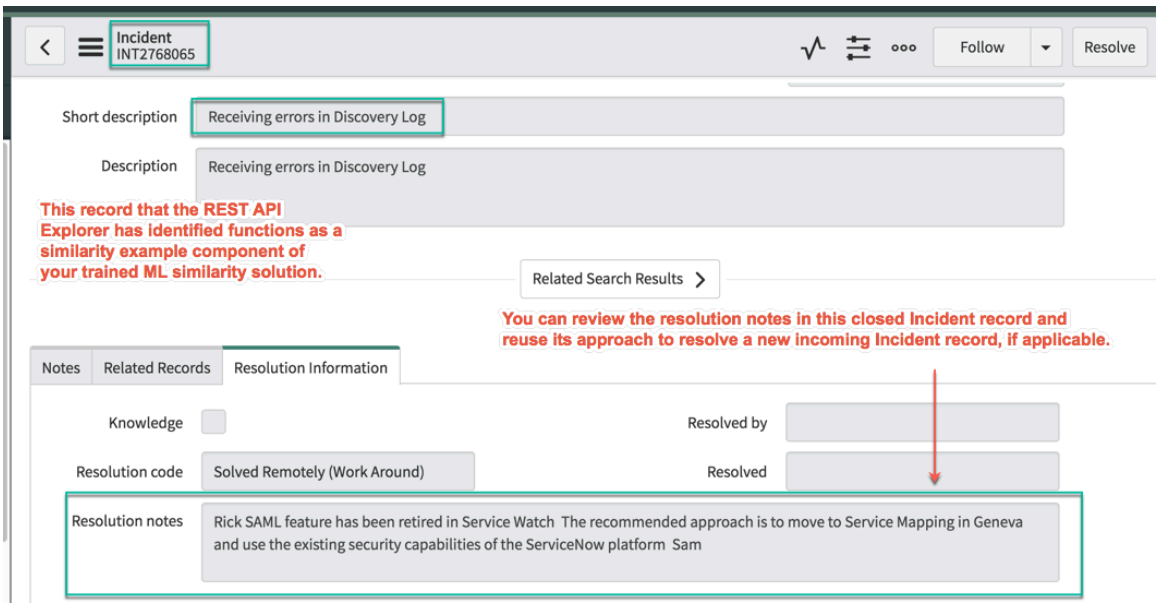


16. Per the image below, compare the returned list of incidents with the input for the prediction output in the REST API Explorer.

    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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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.

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,

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

the training can take hours or sometimes days to complete. The ML Training Jobs view helps you stay on top of all in-progress and completed ML solution training jobs in your instance.

To access this view, you use the admin or ml_admin role and the following navigation path: **Predictive Intelligence** > **Training Jobs**.

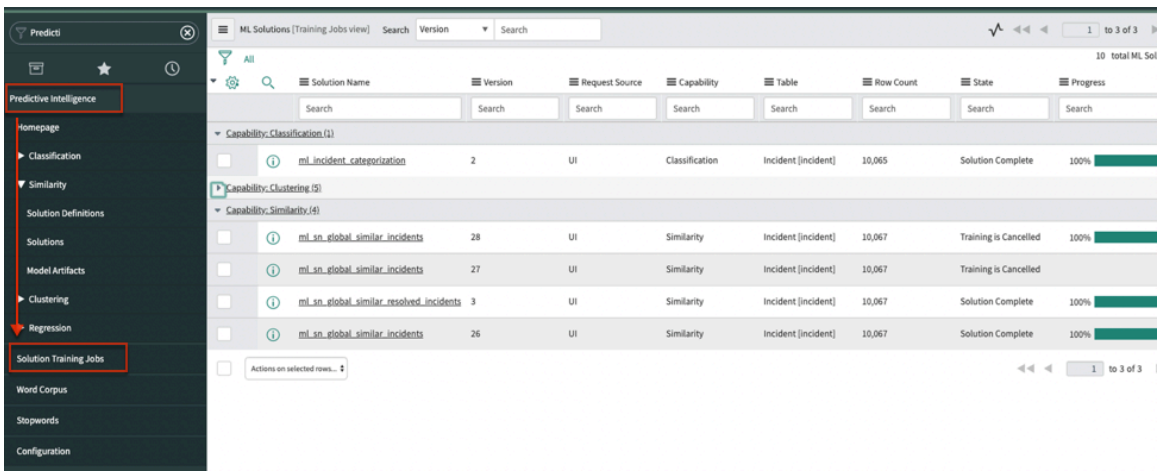
Note:
The ML scheduler limits the number of trainings an instance can commit to 50 new ML training requests per instance within a 24 hour window. This excludes scheduled re-training requests. In addition, clustering and similarity updates are also excluded from this limit, even if the new training requests exceed 50 within a 24 hour window.

ML Training Jobs view summary

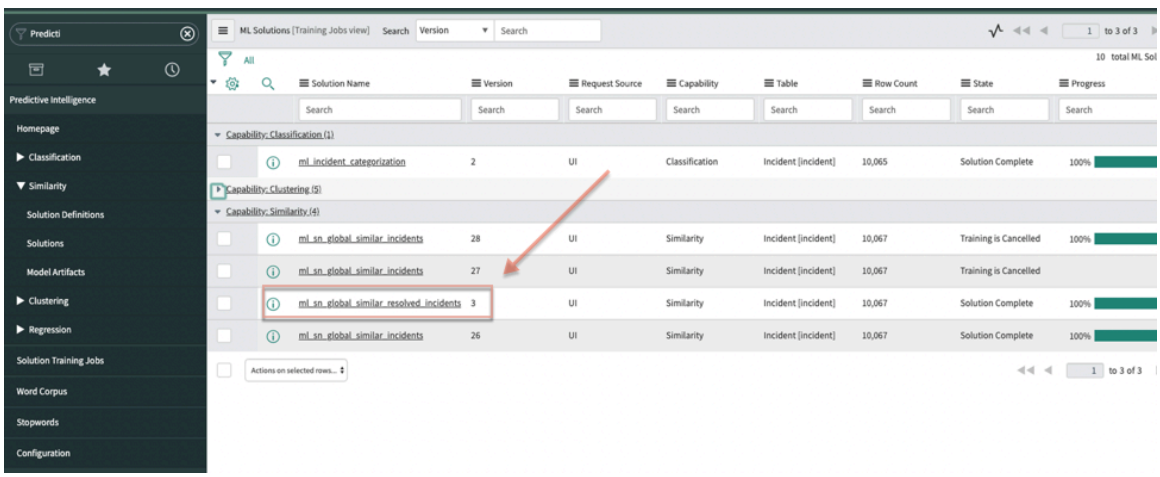
The view shows all ML training jobs grouped by the four Predictive Intelligence capability frameworks: classification, similarity, clustering, and regression.

Each record displays values such as solution name, version, training state, and training completion percentage.

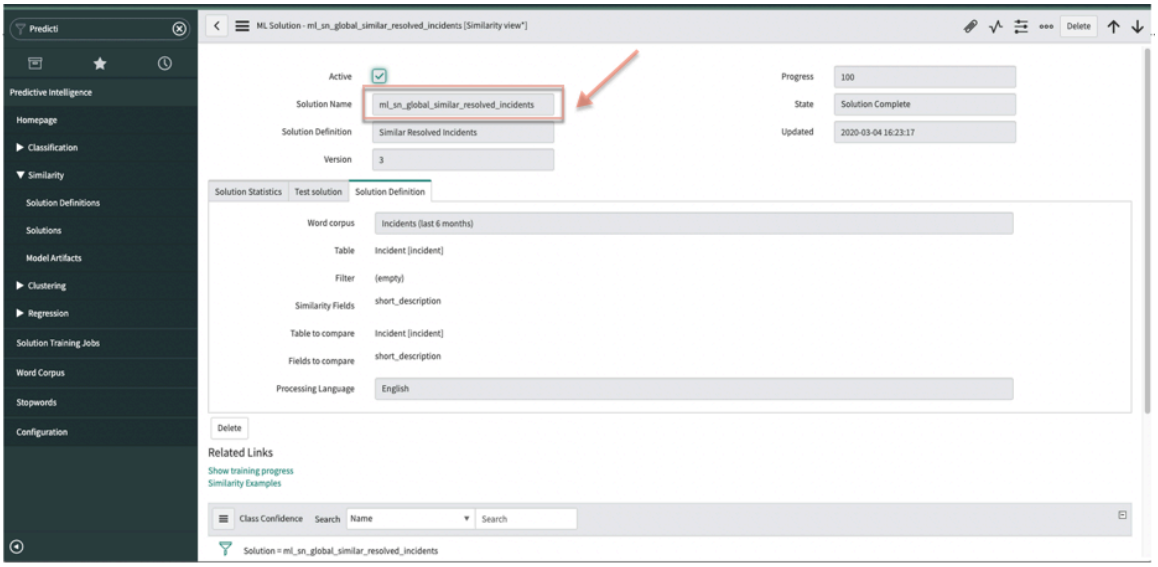
If you don't have any training jobs for a particular capability, the list doesn't display a group for that capability. For example, in this scenario, there is no group for regression because you don't have any regression solutions that you've submitted for training yet.



Select the Solution Name to see the details of the ML solution, as demonstrated in the images below.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Domain separation and 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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Tenant domains can manage their own application data
- Application properties are domain-aware when needed

Related topics

Domain separation for service providers

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

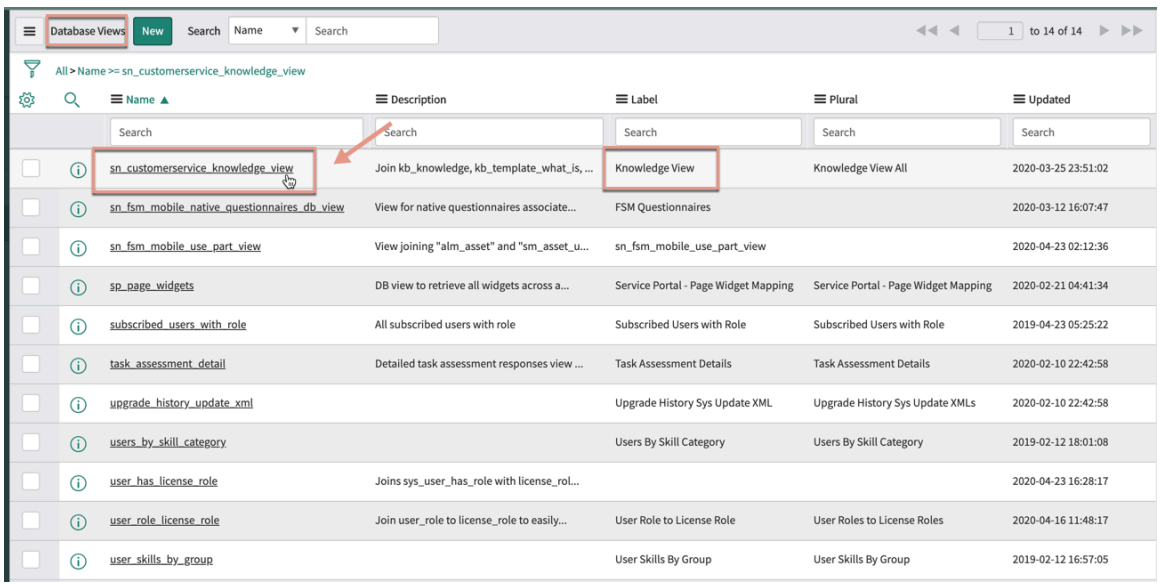
Database views help expand the amount of fields your solution can use for training. By using more than one table in your ML solution definition form, you can access more input data to help enrich the solution outcome.

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**.

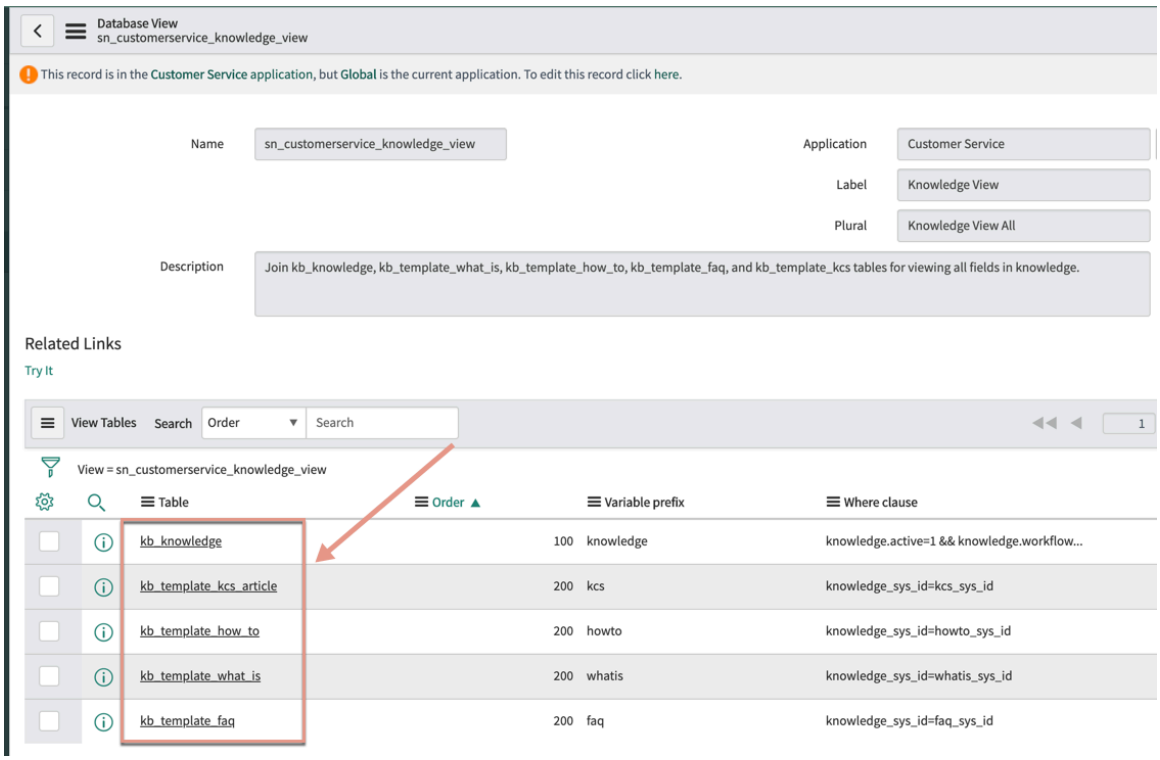


When you click the record **Name**, its content appears, as shown in the image below. Within the database view content, the five Knowledge Base tables you've joined to the view are listed. Most of these tables contain different Knowledge article template types, such as an FAQ or a How To article.

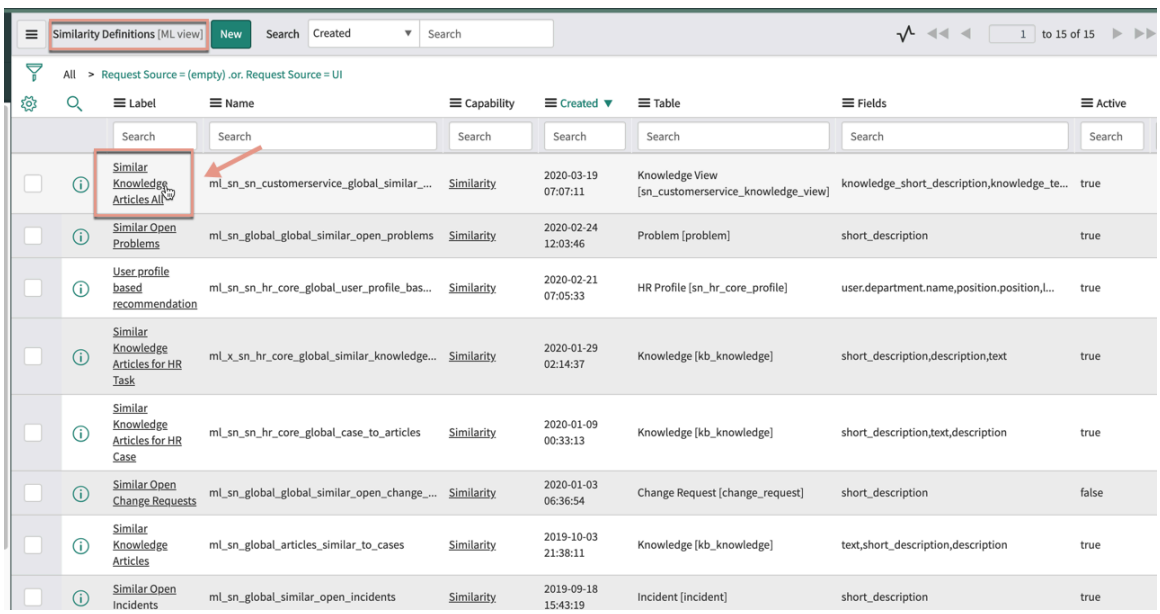
When a database view is used as input to a similarity solution, each of the tables that constitute the view must have at least the required number of records set in the configuration of your ServiceNow instance. The default minimum number required is 10,000. For example, the **Knowledge View** database view has five tables and each table must have 10,000 records. If

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a 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.



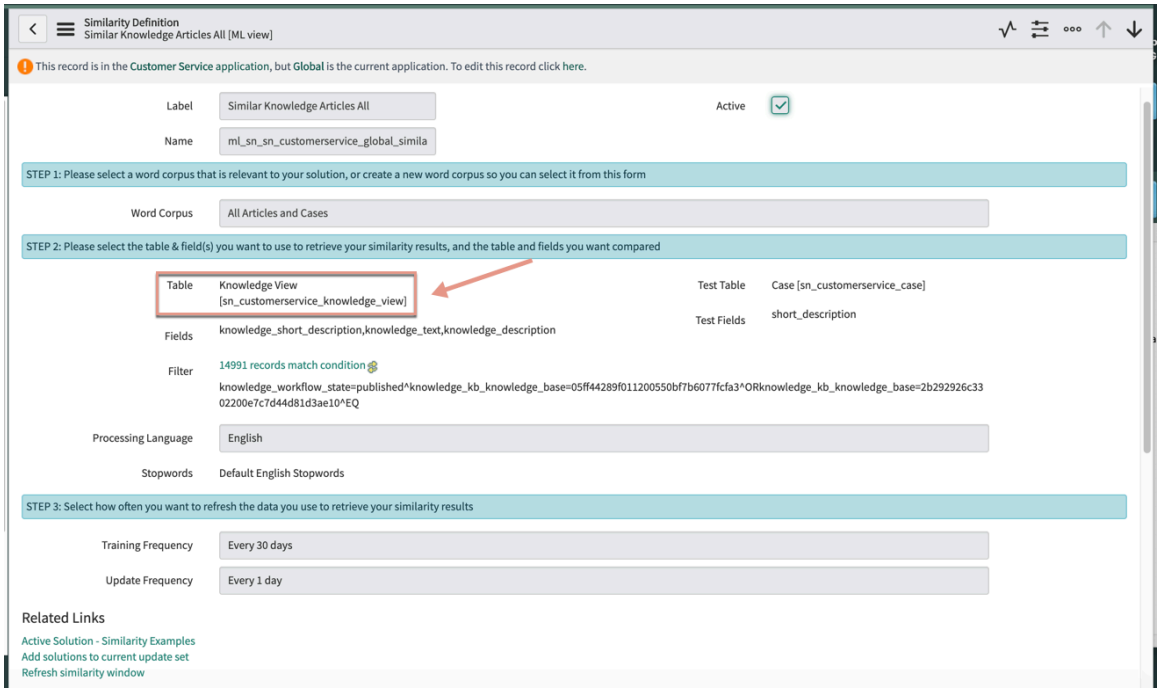
In the image below you can see the similarity solution definition record you've already created, which you plan to associate to your database view. When you click the **Label** for your similarity record, its Similarity Definition form appears.



Similarity Definition forms compare your existing table records based on their similarity by using a table in the **Table** field and another table in the **Test Table** field.

To use a database view in your similarity solution, instead of selecting a table in the **Table** field, you select the database view you created, which in this example scenario is the **Knowledge View** database view. This configuration increases the number of records your solution uses in training because the system compares and processes five tables of data instead of one.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



ServiceNow® 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.

Predictive Intelligence for Application Portfolio Management	Machine Learning solutions for Knowledge Management	Machine Learning solutions for IT Service Management
Machine Learning solutions for HR Service Delivery	Machine Learning solutions for HR Service Delivery Agent Workspace	Issue assignment using the Governance, Risk, and Compliance Predictive Intelligence plugin
Machine Learning solutions for Vulnerability Response	Machine Learning solutions for Customer Service Management	Machine Learning solutions for Search administration
Machine Learning solutions for Field Service Management	Machine Learning solutions for Event Management	Machine Learning solutions for Flow Designer

Use Predictive Intelligence in Workflow Studio with ML actions

Use Predictive Intelligence actions in Workflow Studio to create flows that incorporate your model predictions.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin

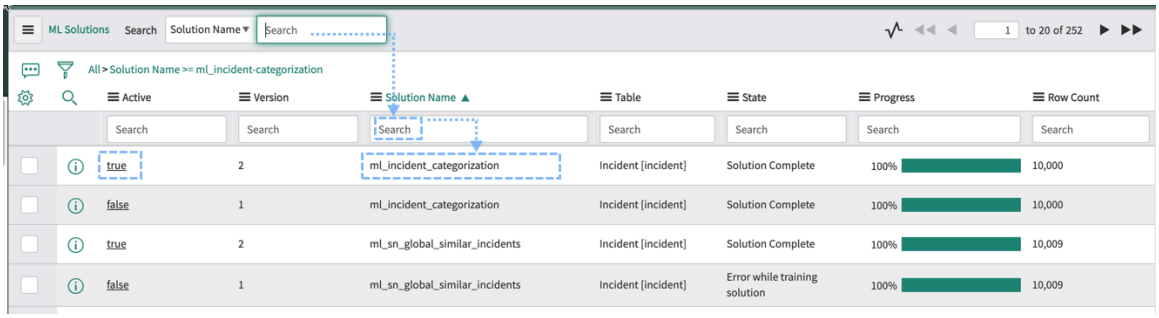
- Make sure the following plugins are activated in your instance: Predictive Intelligence (com.glide.platform_ml) and Predictive Intelligence for Flow Designer (com.snc.ml_flowdesigner).
- Create or use an existing trained Predictive Intelligence solution.
- Roles required: admin or ml_admin, and flow_designer 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.

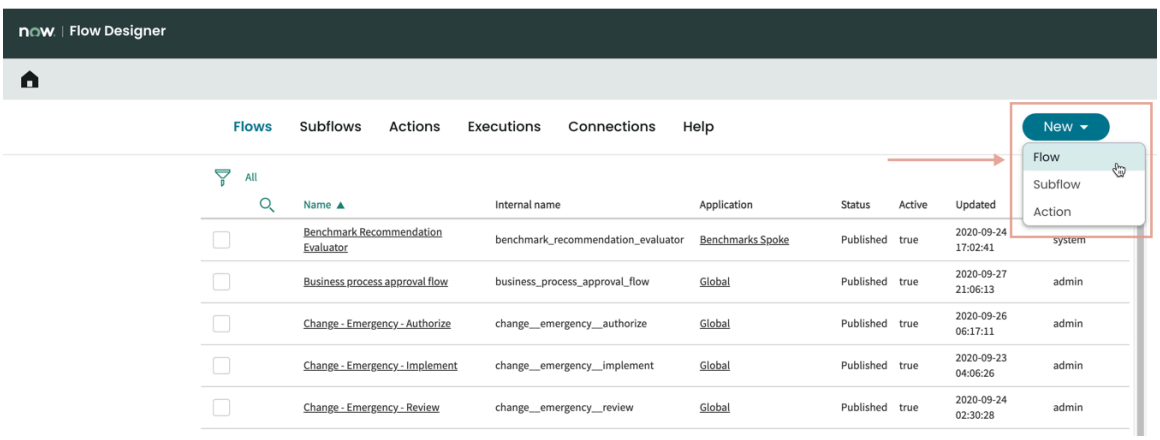
In this example procedure, you create a flow that deploys the ml_incident_categorization solution in a Flow Designer action. You can find this solution by searching on the ML Solutions table, as shown in the image below. Confirm that the solution you use has been trained and its **Active** value is set to **true**.



For more information on how to use Flow Designer in Workflow Studio, see . For information about the Actions included in Predictive Intelligence for Flow Designer, 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**.

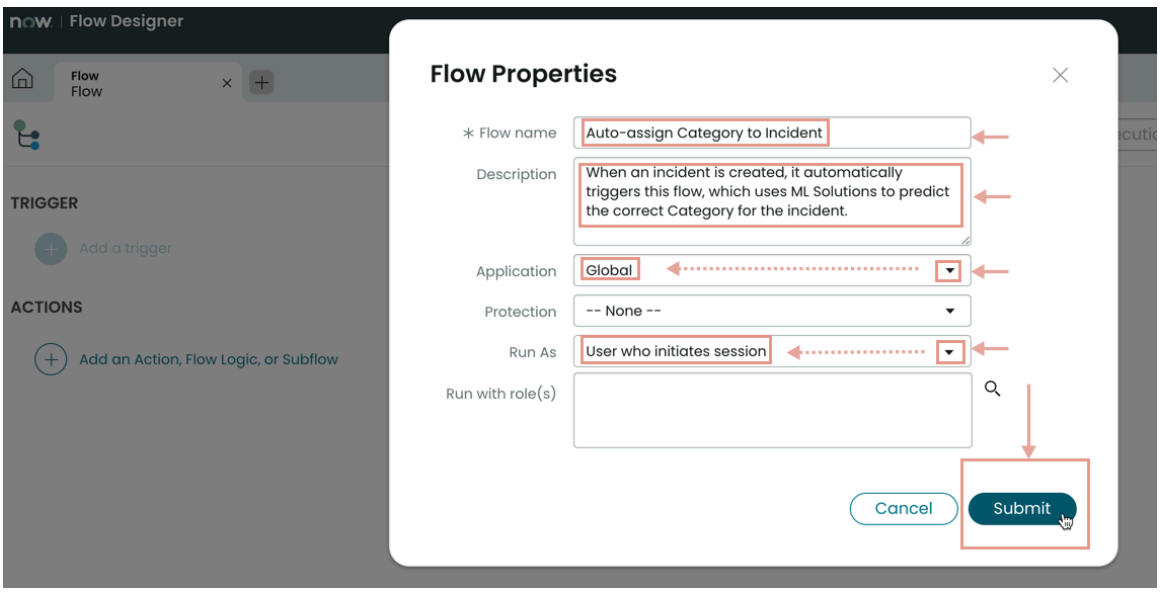


© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

3. On the **Let's get the details for your flow** screen, configure the following fields.

Field	Description
Flow Name	Provide a name for the flow. In this scenario, you enter `Auto-assign Category to Incident`.
Description	Enter a brief summary description of what the flow delivers. For example, you enter the following: `When an incident is created, it automatically triggers this flow, which uses ML Solutions to predict the correct Category for the incident.`
Application	Select **Global**.
Protection	Select **--None--** or **Read-only**. In this scenario, you select **--None--**.
Run As	Select **User who initiates session**.
Run with role(s)	Leave blank.
Flow priority default	Medium (default).

4. Select **Submit**.

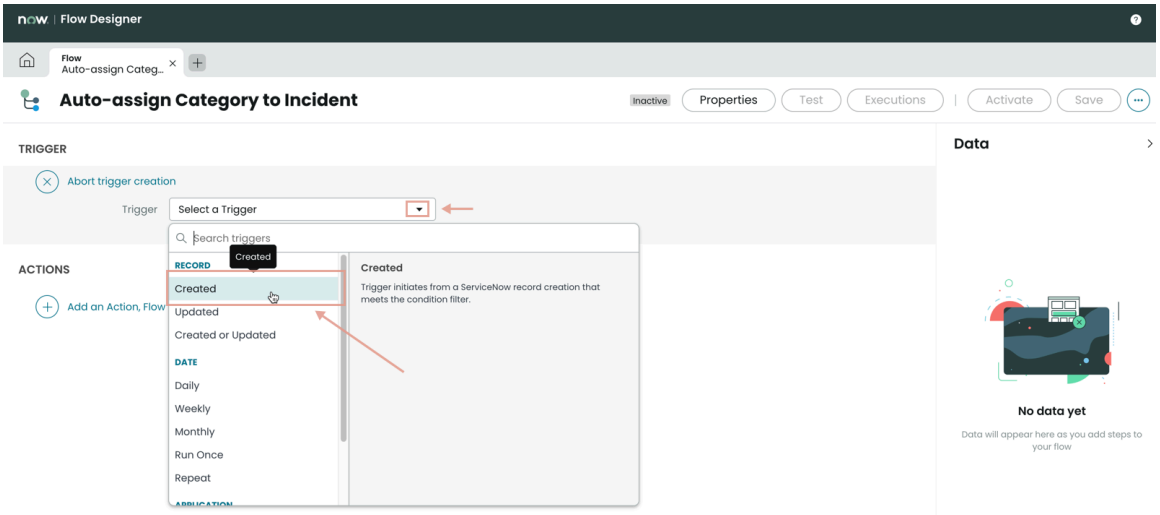


The Flow screen appears, showing the Auto-assign Category to Incident name you assigned to the flow. If a Getting started screen appears, select **Skip tour**.

5. In the TRIGGER section of the screen, configure the following fields to create a trigger for the flow.

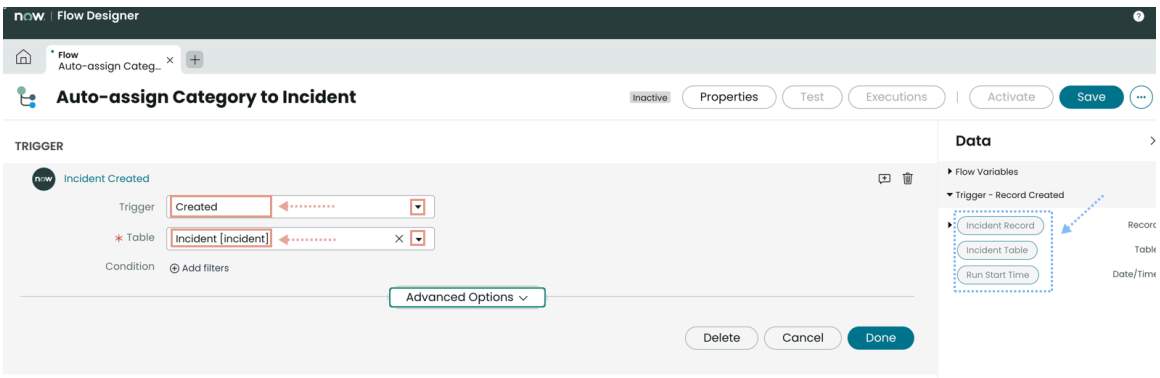
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

a. Add a trigger: Select **Created**.



b. Table: Select **Incident [incident]**.

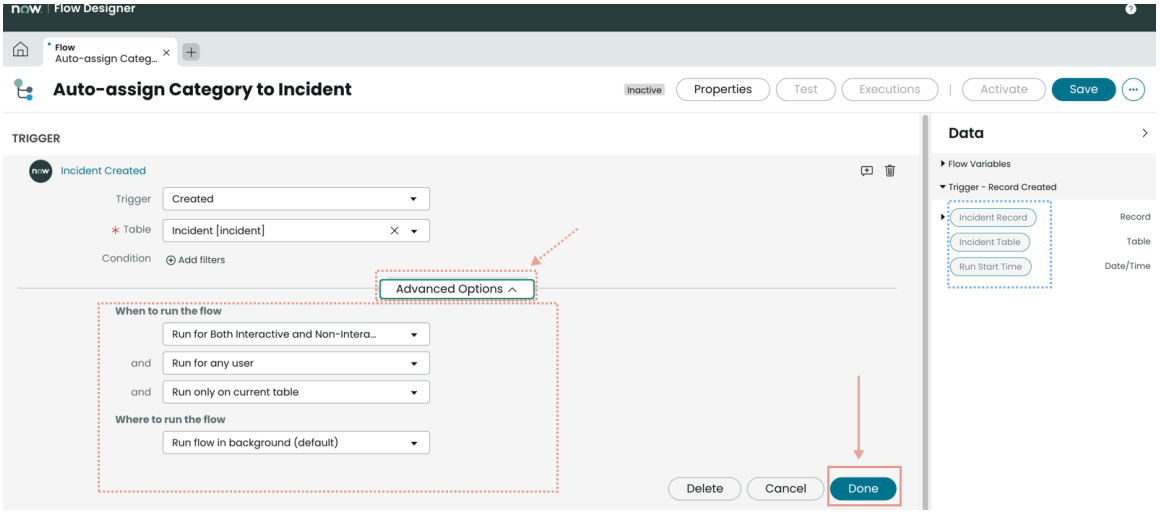
Note:
When you configure both the **Trigger** and **Table** fields, record data pills appear in the **Data** section of the screen so you can use them in your flow.



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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

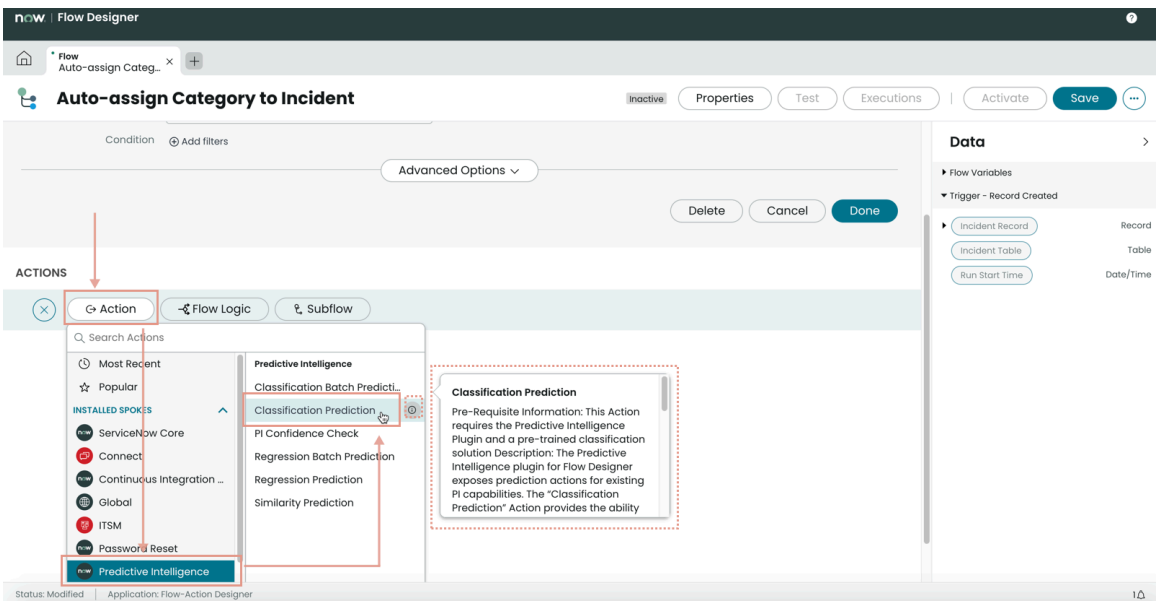


e. **Optional:** To close the panel, select **Advanced Options**.

f. Select **Done**.

6. In the ACTIONS section of the screen, configure the following fields to create a Classification Prediction action.

a. **Add an Action, Flow Logic, or Sub-Flow**: Select **Action** > **Predictive Intelligence** > **Classification Prediction**.



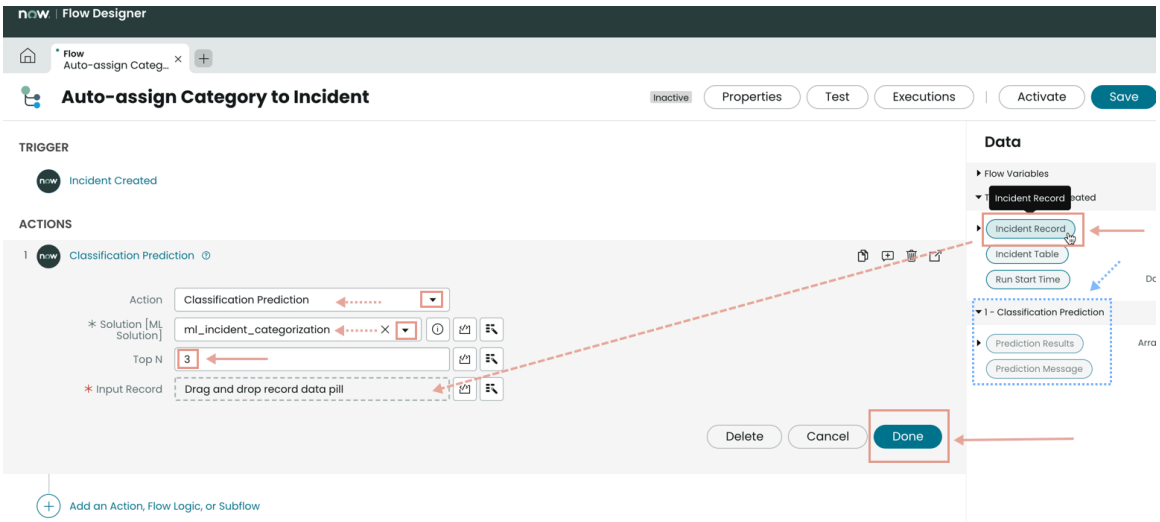
Select the tool tip to see a description of a Classification Prediction.

b. **Solution Name [ML Solution]**: Select **ml_incident_categorization**.

c. **Top N**: Enter 3.
When you enter a number, such as 3, the system uses the top 3 ML predictions that have the highest prediction confidence score. If you don't enter anything, the system sets the default value to 1.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

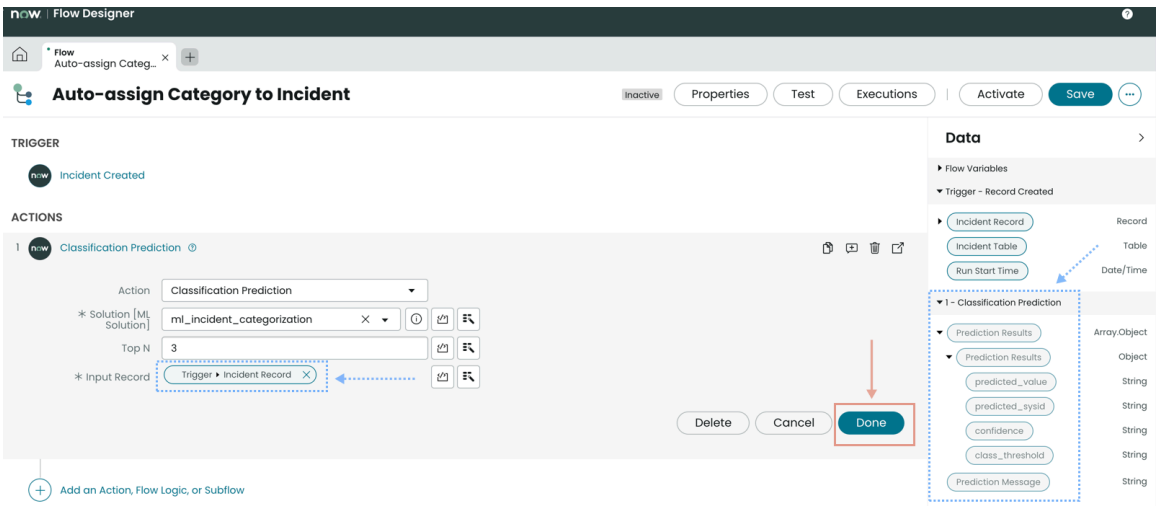
d. Input Record: Drag and drop your **Trigger → Incident Record** data pill into the Input Record field.



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.

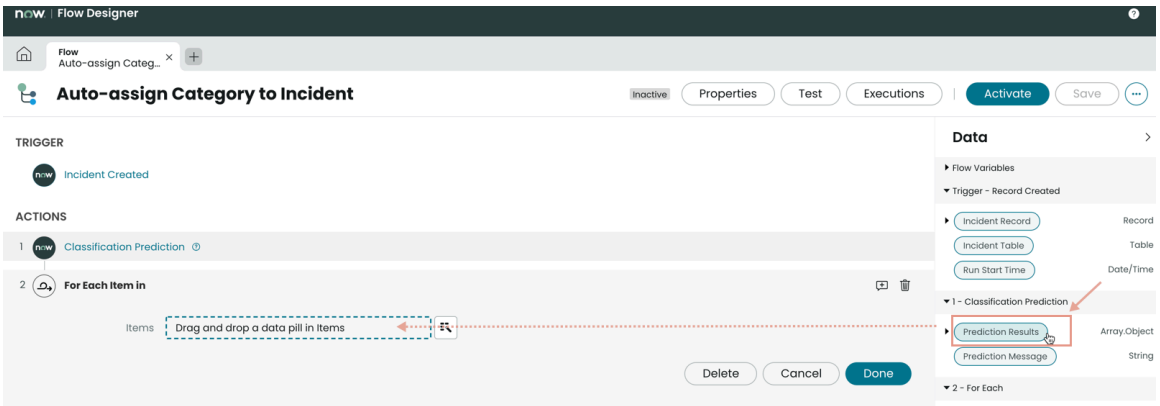


7. In the ACTIONS section of the screen, use the following steps to create actions and flow logic for the incident's Prediction Results.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

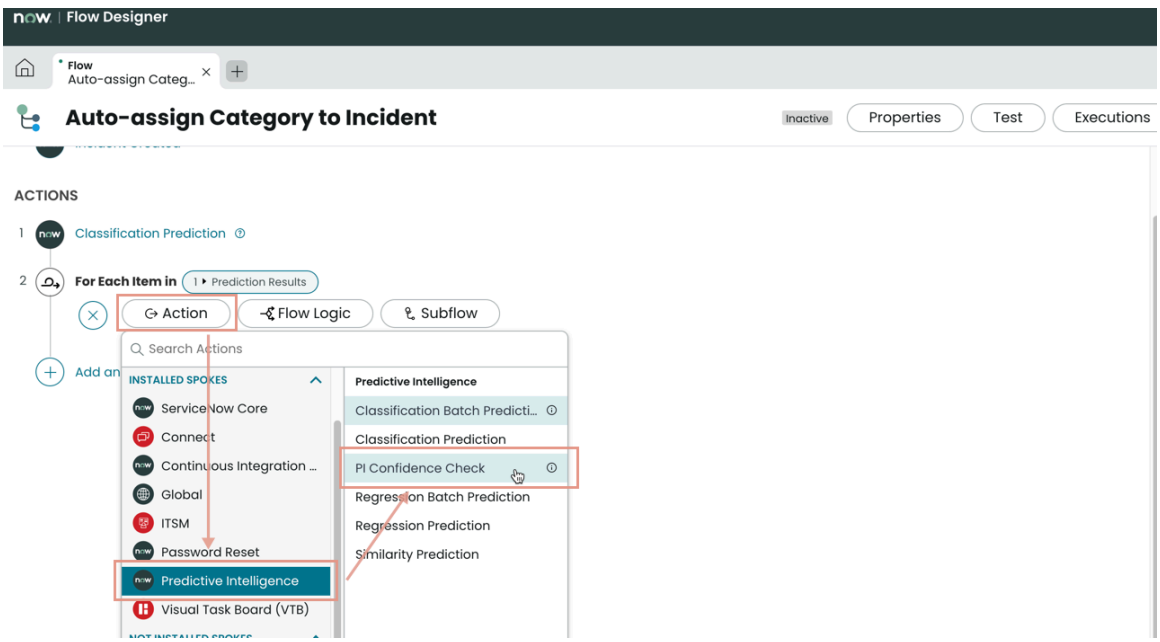
a. **For each item in list of items**: Drag and drop the **Prediction Results** data pill into the **Items** field.



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. Click **Done**.
Result: The Prediction Results action is started in the flow and its data pills appear in the Data section of the screen.

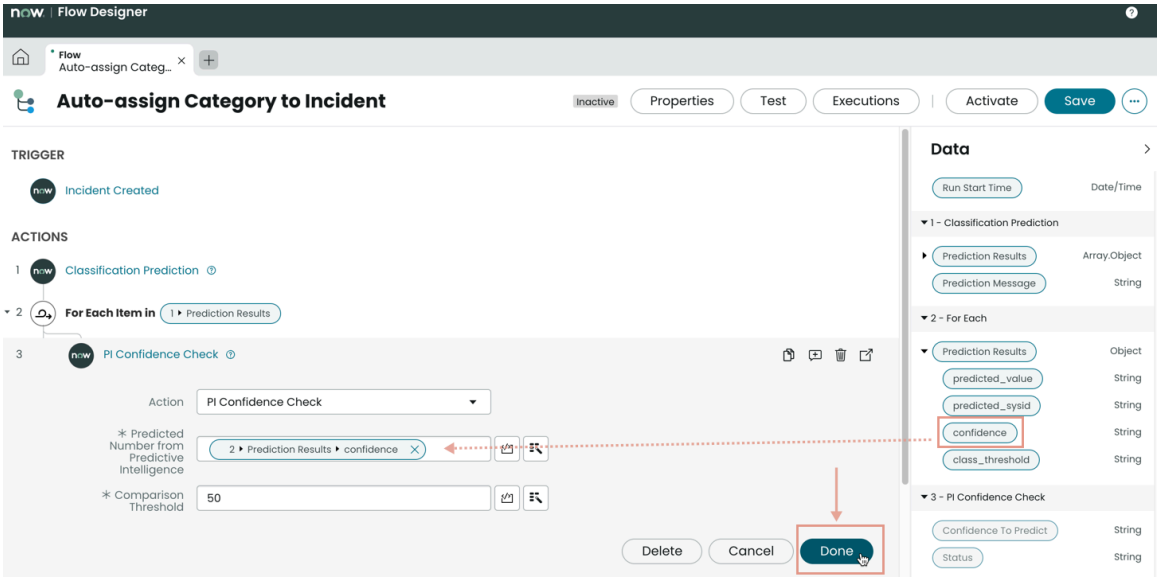
8. In the ACTIONS section of the screen, select **Action** > **Predictive Intelligence** > **PI Confidence Check**.
The PI Confidence Check is a tool you can use to compare values in a flow. In this use case, it compares prediction result values, and the output from the check is either True or False.



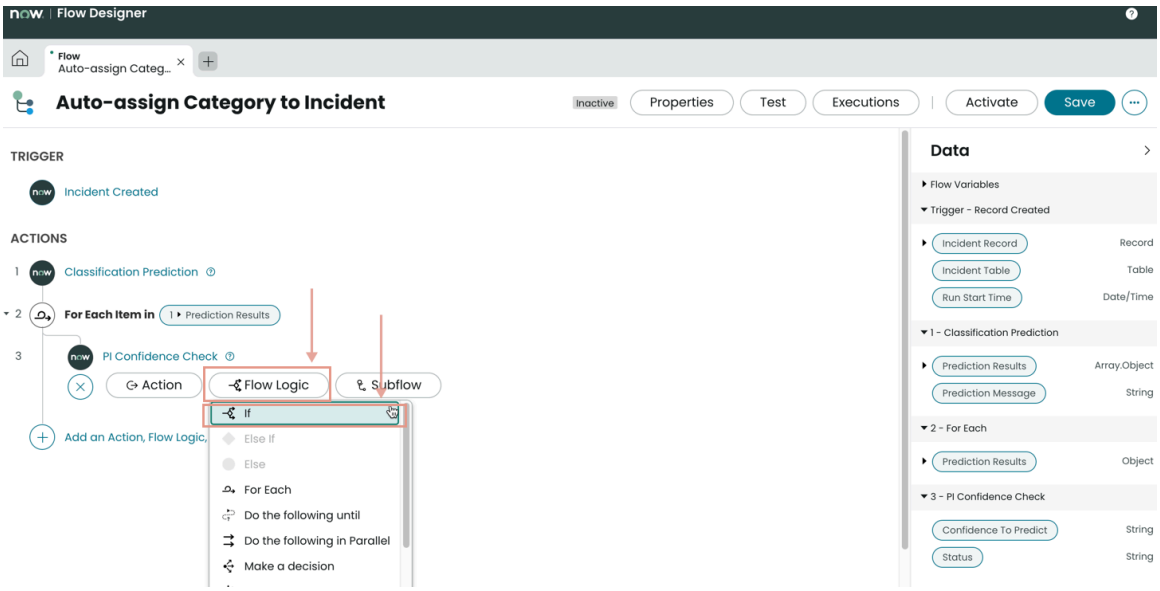
9. Drag and drop the **confidence** data pill into the **Predicted Number from Predictive Intelligence** field.
10. Enter 50 in the **Comparison Threshold** field.
In this example scenario, you enter the number 50, which tells the system to use predictions that have a confidence score above 50%.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

11. Select **Done**.



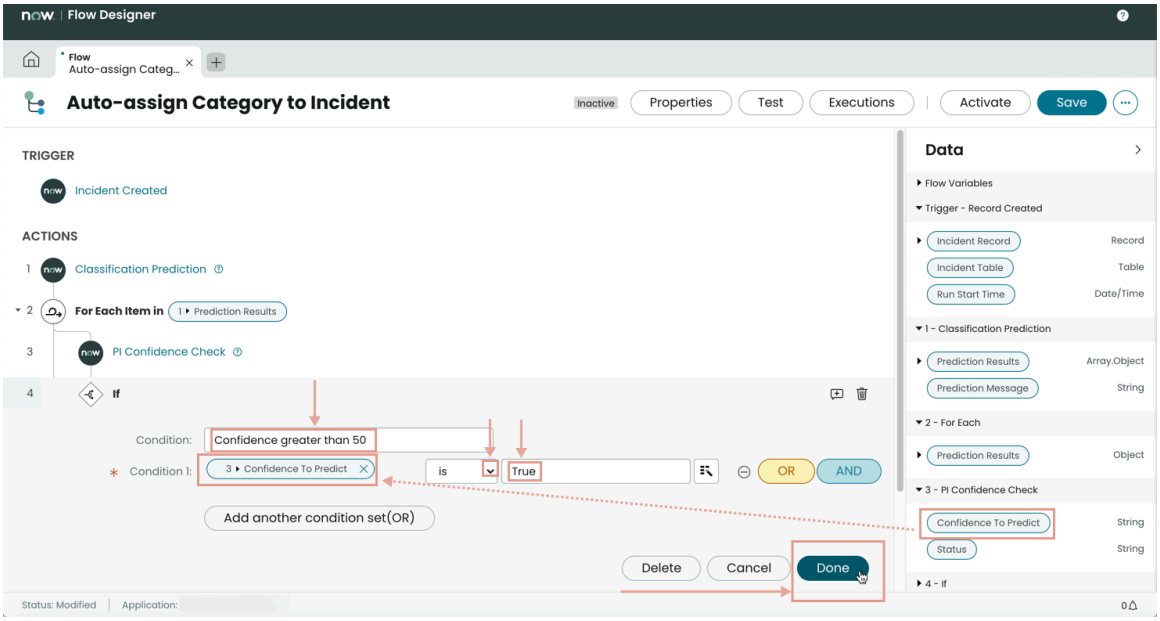
12. Select **Flow Logic** > **If** to add a condition to the flow.



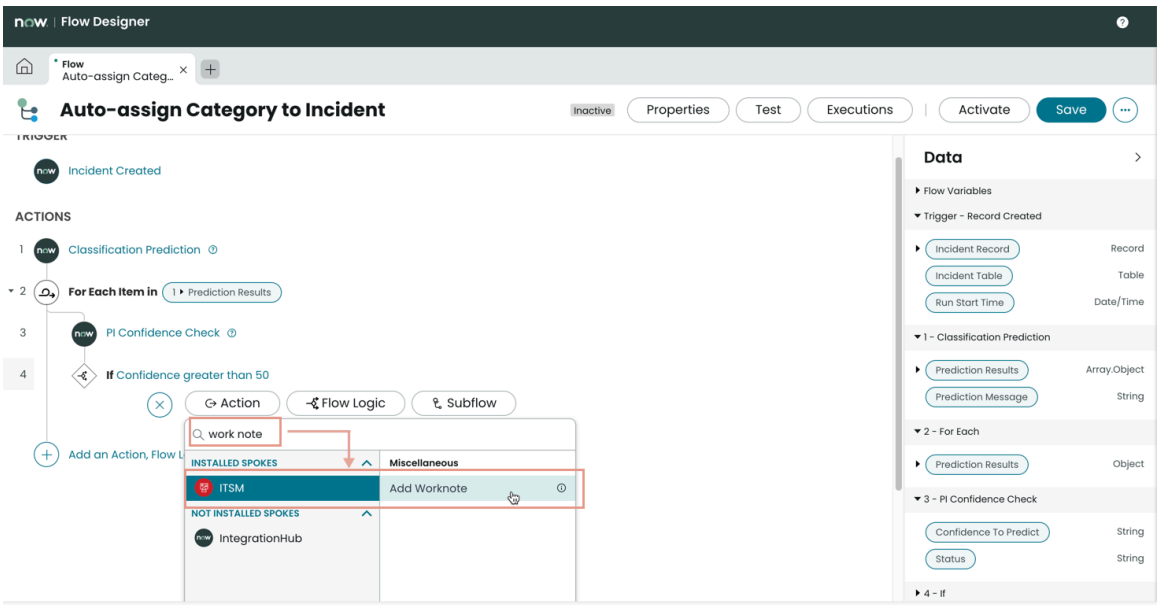
13. Configure the following fields to define the first part of the condition flow logic.

- **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**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

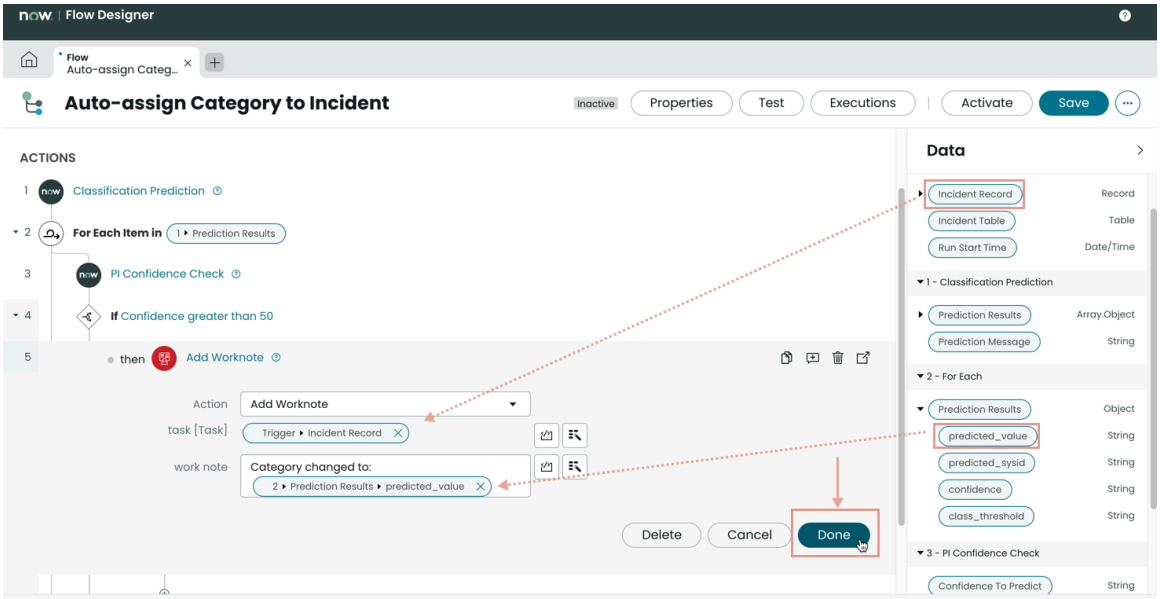


14. Select **Action** and enter `worknote` into the search field.
15. Select **ITSM** > **Add Worknote** to add a work note as the second part (the conclusion) of the condition.

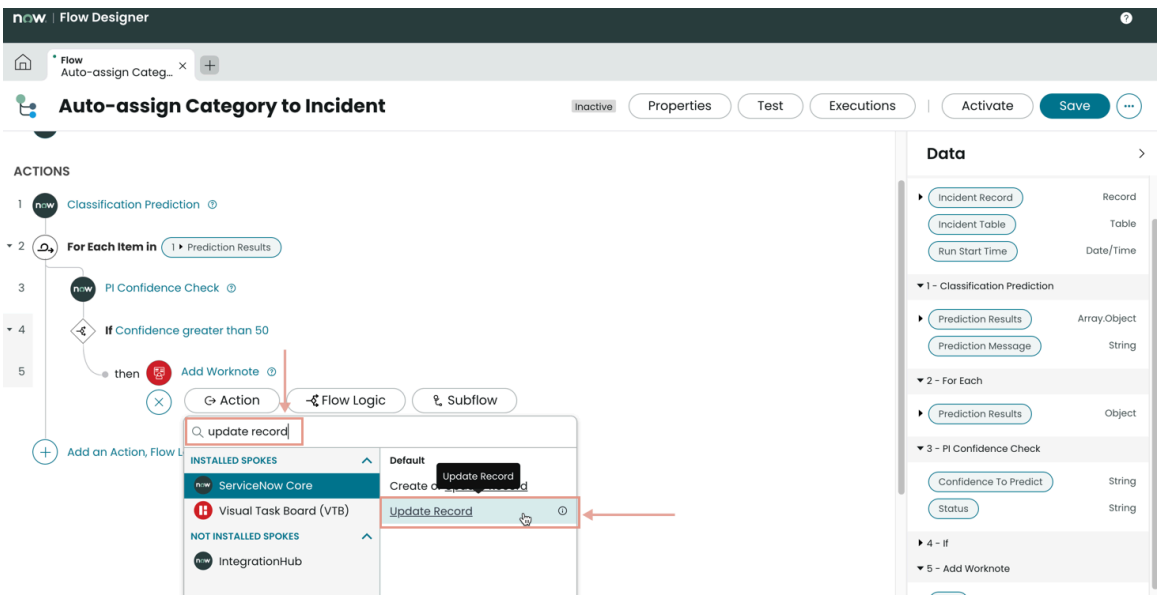


16. Configure the following fields to define the second and final part of the condition flow logic.
    - **Action**: As a result of Step 14 above, **Add Worknote** appears automatically in this field.
    - **task [task]**: Drag and drop the **Incident Record** data pill into the field.
    - **work note**: Drag and drop the **predicted_value** data pill into the field. This step completes the condition flow logic conclusion.
    - Select **Done**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



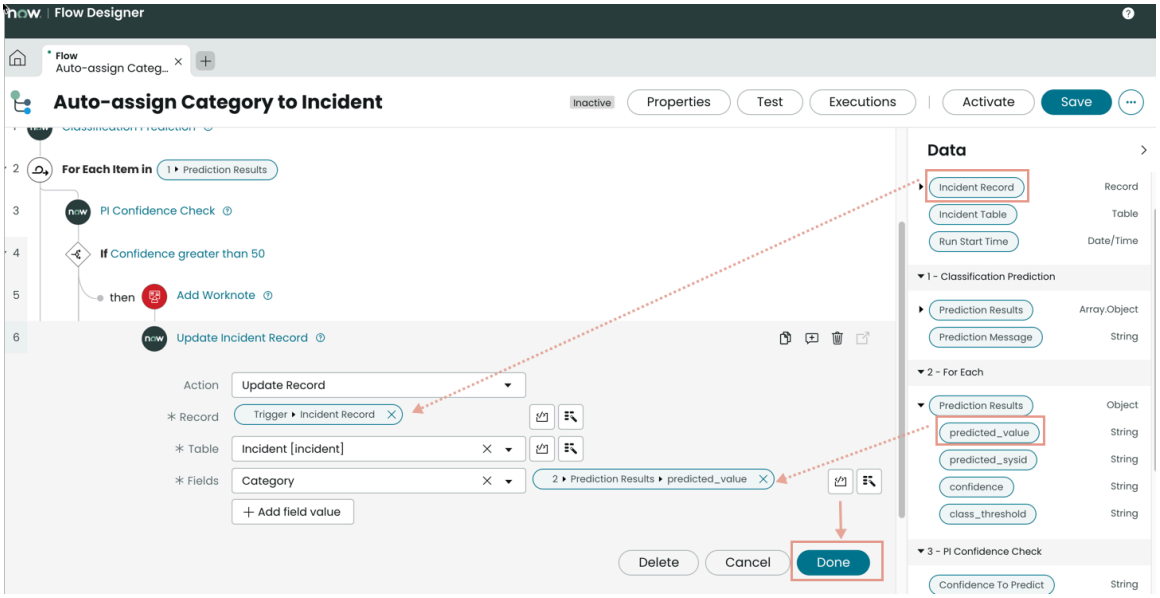
17. Select **Action** and enter `update record` into the search field.
18. Select **Update Record**.



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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

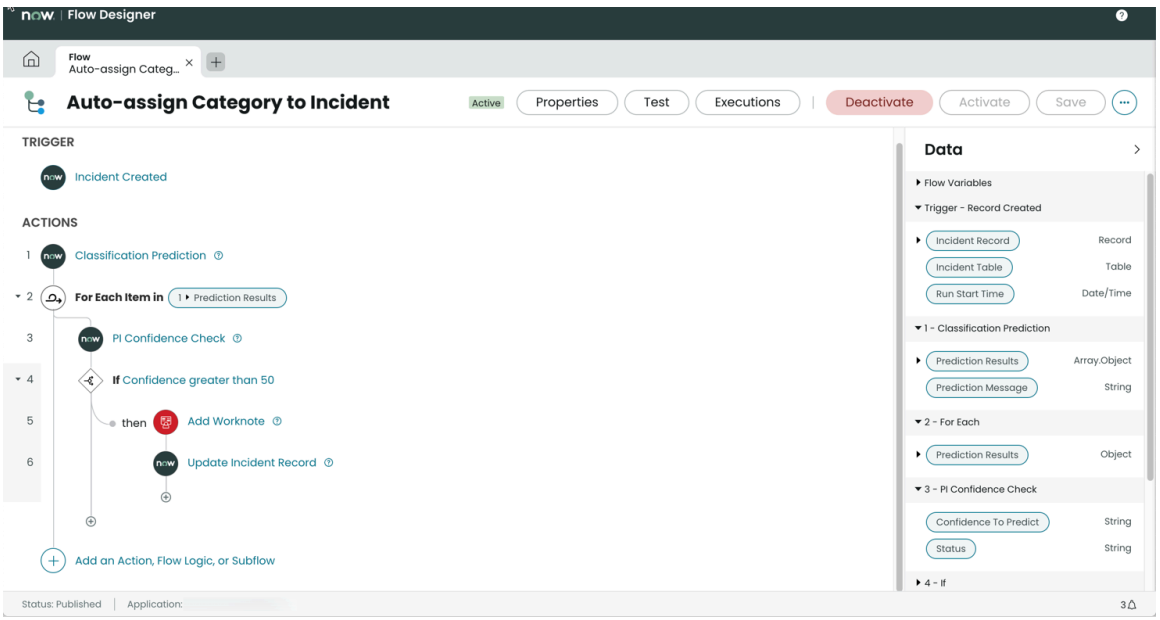
- Click **Done**.



20. Select **Save**.
21. Select **Activate**.

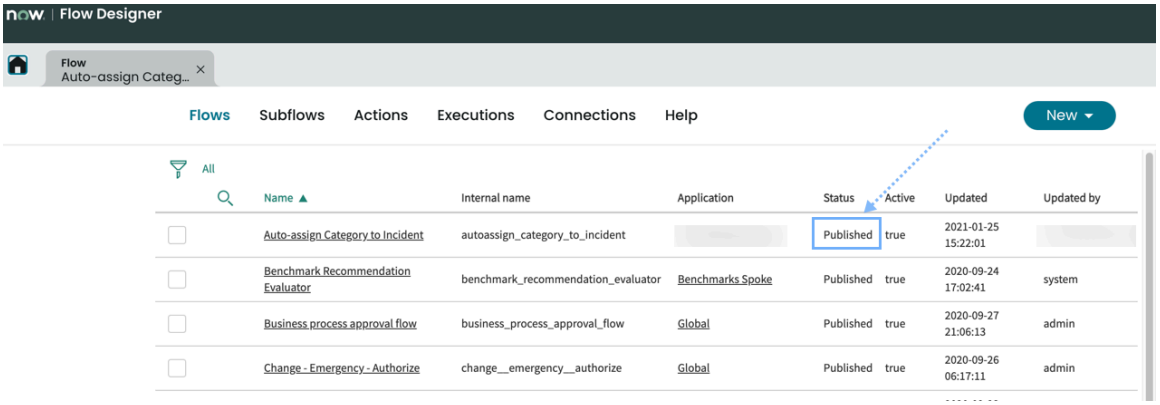
Result

- Your Auto-assign Category to Incident flow is activated and complete.

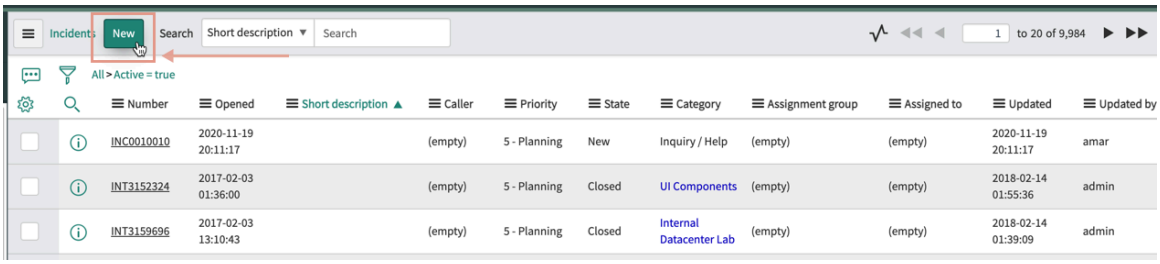


© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

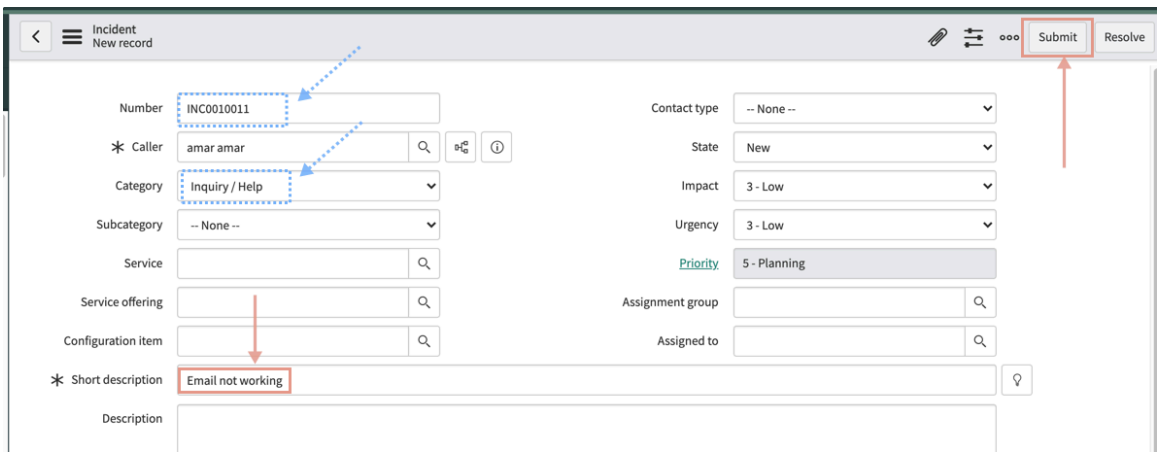
- It also appears as published in the Flows column on the Flow Designer home screen.



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.



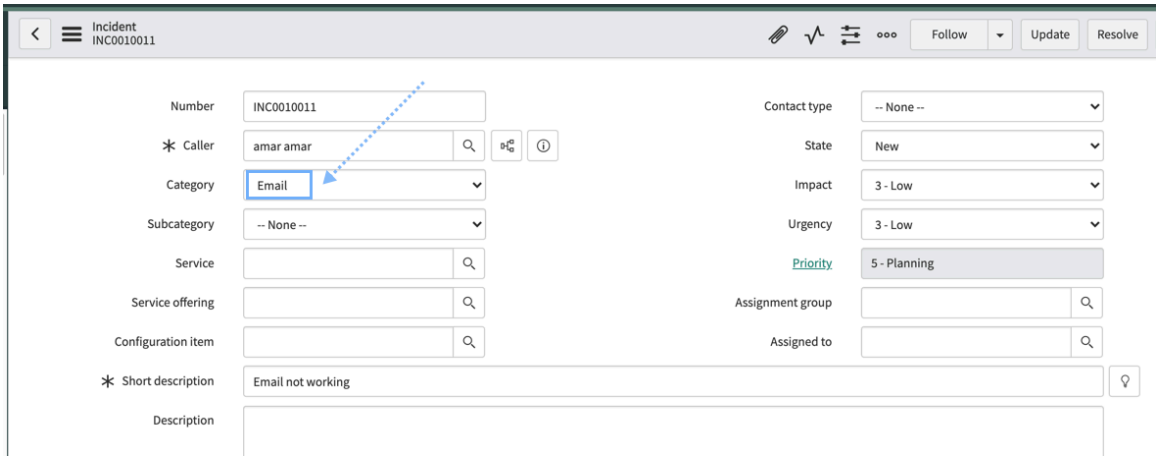
24. In the record you created, note in the following image that the Category value is set to **Inquiry / Help**.
25. In the **Short description** field, enter `Email not working`.
26. Select **Submit**.

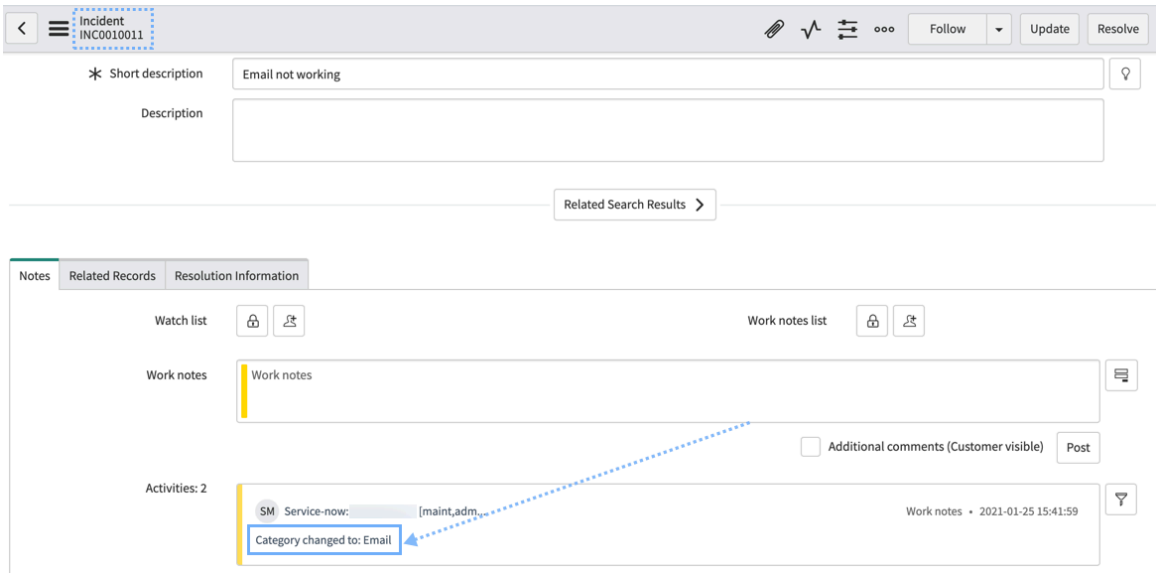


Result

The system updates the incident record to show that its Category value has changed from **Inquiry / Help** to **Email**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- 

- 

Related topics

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.

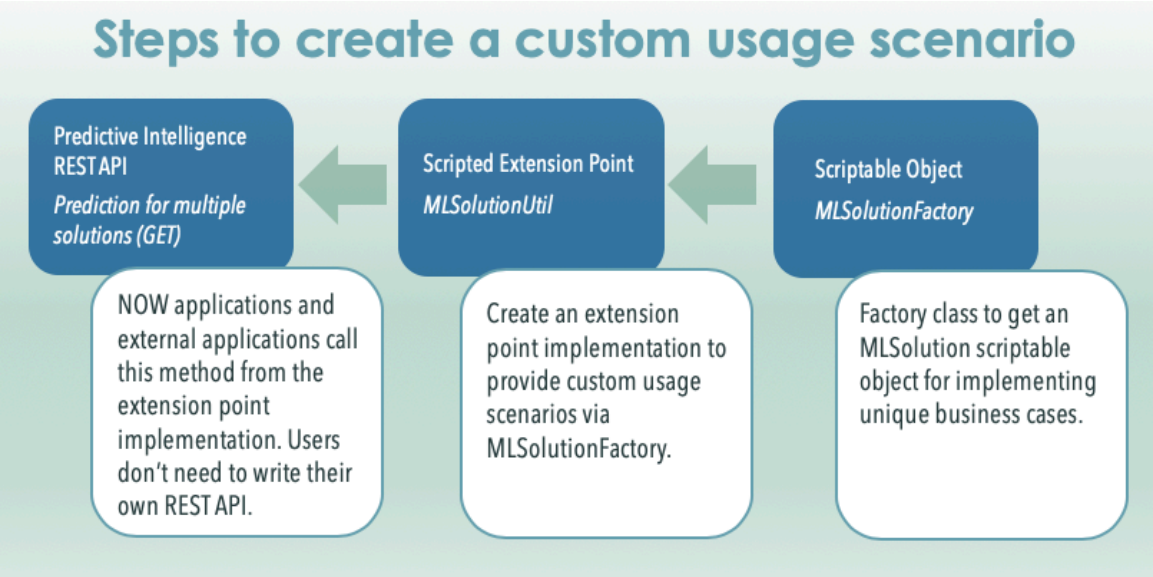
You can use scriptable objects in a scripted extension point to modify the Predictive Intelligence REST API method to address unique business use cases.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

1. 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Enter `sys_properties.list` in the application navigator to access the System Properties list.
2. Ensure the *glide.platform_ml.clone_artifacts* system property is set to **True**.
3. If you want to preserve only the ML solution records and not the numerous other records in the sys_attachments table, exclude the sys_attachments table from your clone run.
4. Request a system clone.
   The system preserves your ML solution records during the system clone.

Predictive Intelligence references

Reference pages for components of Predictive Intelligence.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Using 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	• True: Overrides the **readonly** state during training, enabling you to edit ml_solution_definition. • False: ml_solution_definition remains in the **readonly** state.
Minimum number of records for Regression	glide.platform_ml.api.min_regression_records	Sets the minimum number of records required for Regression solution training. • Type: Integer • Default value: 10000 **Note:** 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.
Maximum number of records for Regression	glide.platform_ml.api.max_regression_records	Sets the maximum number of records that can be used in Regression solution training. • Type: Integer • Default value: 300000

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Predictive Intelligence properties (continued)

Property	Property Name	Description
		Note: 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.
The time (in ms) an in-training solution will wait without updates before timing out	glide.platform_ml.training_timeout	Sets the time-out time in milliseconds for an in-training solution. • Type: Integer • Default value: 21600000 (ms)
Maximum model size	glide.platform_ml.api.model_size	Sets the maximum number of records you can use to train a model. • Type: Integer • Default value: 524288000
Maximum number of records used in training	glide.platform_ml.api.csv_max_line	Sets the maximum number of records that can be used in Classification solution training. • Type: Integer • Default value: 300000
Minimum number of records used in training	glide.platform_ml.api.csv_min_line	Sets the minimum number of records required for Classification solution training. • Type: Integer • Default value: 10000
Maximum number of days worth of records on request can retrieve	glide.platform_ml.api.csv_split_days	Sets the maximum number of days a single request can retrieve from your records.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Predictive Intelligence properties (continued)

Property	Property Name	Description
		• Type: Integer • Default value: 30
Maximum number of records per table for word corpus	glide.platform_ml.api.max_wordcorpus_records	Sets the maximum number of records per table for Word Corpus creation for Similarity and Clustering solutions. • Type: Integer • Default value: 300000 **Note:** 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.
Maximum number of records for similarity window (to lookup results)	glide.platform_ml.api.max_similarity_window_records	Sets the maximum number of records that the Similarity Window can retrieve to look up results. • Type: Integer • Default value: 100000
Maximum number of records for Clustering	glide.platform_ml.api.max_clustering_records	Sets the maximum number of records you can include in a cluster. • Type: Integer • Default value: 100000

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Shared 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

None.

Special considerations

None.

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 artificial intelligence (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

Explore	Configure	Integrate
		
Learn about DocIntel concepts and features.	Set up DocIntel to process different types of documents.	Integrate your DocIntel application with other applications.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Use	Monitor	References
		
Use the DocIntel application to process documents by detecting and analyzing text in documents.	Track DocIntel usage and effectiveness.	Get details about the properties installed with DocIntel.

Exploring Document Intelligence

Document Intelligence helps you to quickly and accurately classify and extract information from documents using artificial intelligence (AI).

Overview

Many organizations today use simple optical character recognition (OCR) solutions to extract data from documents. This requires significant manual configuration and often requires manual adjustment as the documents evolve. Document Intelligence extends beyond simple OCR, using AI to identify, understand, and extract text and data from documents. This enables you to accurately extract information to automate document processing, even when the documents have varied text, formatting, and templates.

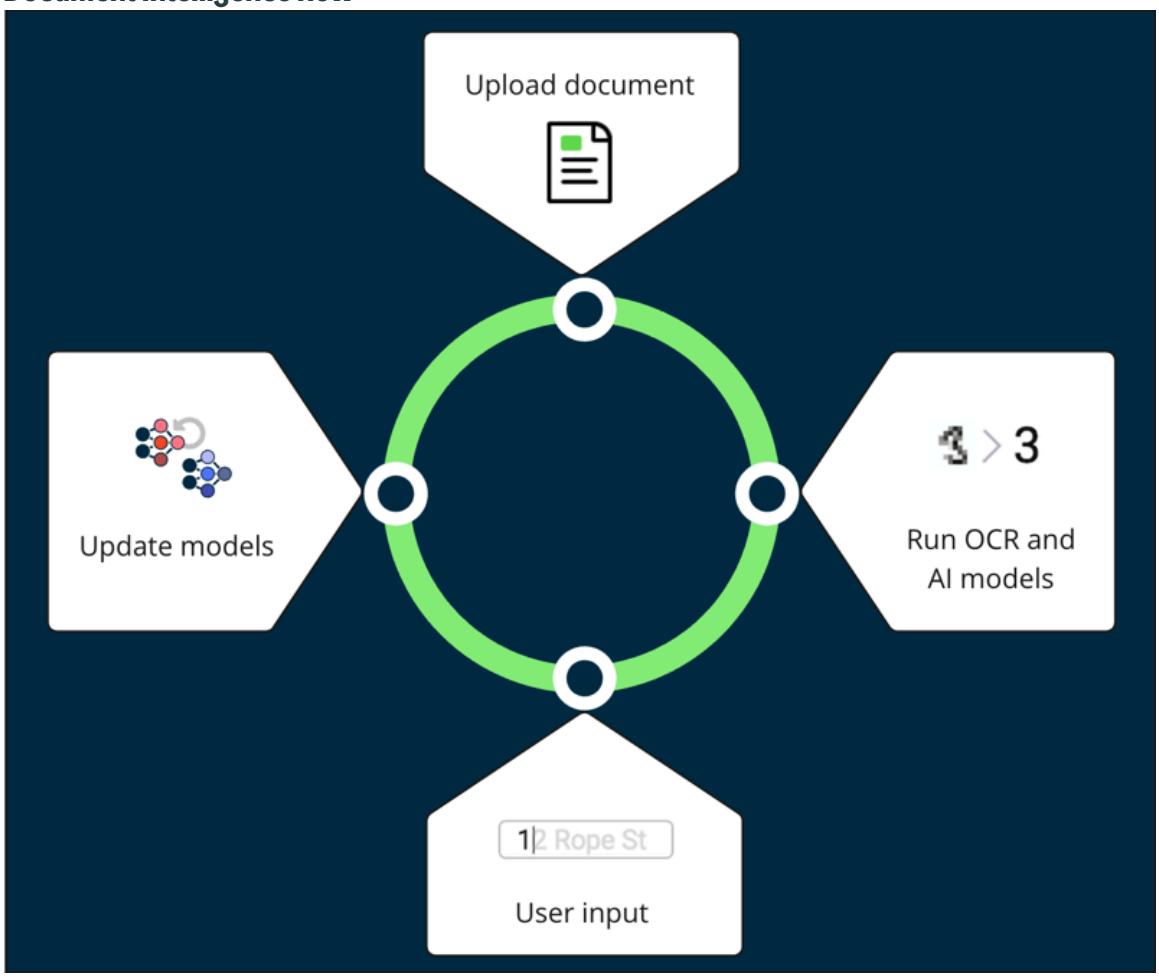
Document Intelligence workflow

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence flow



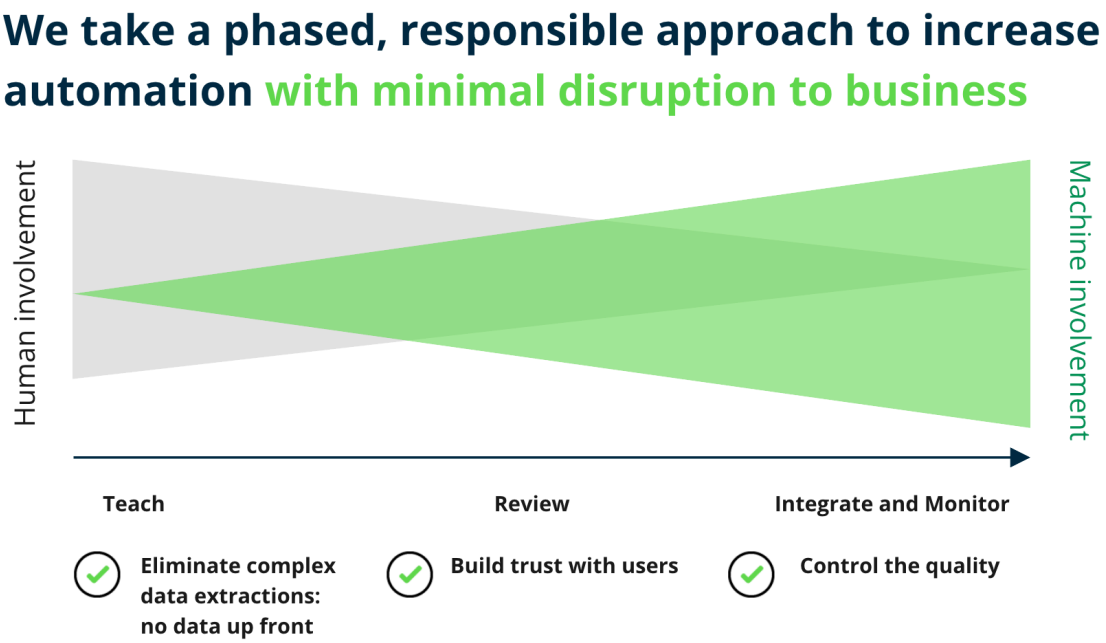
In this workflow:

1. A document is uploaded for processing in a document task.
2. DocIntel extracts the data from the document using OCR and AI models.
3. The user provides input to validate or correct the DocIntel recommendations.
4. The models are updated and trained to provide more accurate results.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence benefits

Benefits of Document Intelligence



Document Intelligence provides the following benefits.

Benefit	Feature	User
Start fast with a no-code set-up that enables data extraction from many document types including PDF and scanned paper documents.	Set up document extraction use cases	DocIntel Admin [sn_docintel.admin] DocIntel Manager [sn_docintel.manager]
Enable categorization for any type of document you define.	Set up document classification use cases	DocIntel Admin [sn_docintel.admin] DocIntel Manager [sn_docintel.manager]
Automate intelligently with responsible, feedback-driven AI for continual learning.	Configure data extraction modes	DocIntel Admin [sn_docintel.admin] DocIntel Manager [sn_docintel.manager]
Seamlessly integrate document processing steps into workflows.	Integrating Document Intelligence with other applications	DocIntel Admin [sn_docintel.admin] DocIntel Manager [sn_docintel.manager]

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Benefit	Feature	User
Accelerate extraction of structured and semi-structured documents such as forms, invoices, IDs, and more.	Extract fields using the Document Intelligence workspace	DocIntel Creation Agent [sn_docintel.creation_agent] DocIntel Extraction Agent [sn_docintel.extraction_agent]
Accelerate classification of single and multi-page documents.	Classify documents using the Document Intelligence workspace	DocIntel Creation Agent [sn_docintel.creation_agent] DocIntel Extraction Agent [sn_docintel.extraction_agent]

Configuring Document Intelligence

Activate Document Intelligence on your instance and get started with basic configuration.

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.
   - Set up document classification use cases.

     Create classification use cases, add fields (classes), and train your use cases to begin classifying documents.
   - Manage document extraction use cases.

     Manage your Document Intelligence use cases to support your document processing requirements.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Set up Document Intelligence

Review the following information before you start setting up Document Intelligence.

Checklist

Setup task	Description
Verify that the ServiceNow core applications or plugins that are required to support Document Intelligence are installed and activated.	Verify that the following applications or plugins are installed and activated from the ServiceNow Store. When you activate the first plugin, its dependent plugins are activated automatically. If not installed, install and activate one application at a time in the following order to ensure a smooth implementation. • Predictive Intelligence (com.glide.platform_ml) • Platform Document Intelligence (com.glide.platform_ml_di) • Document Intelligence UIB Component (sn_docintel_iframe) For more information, see Install Document Intelligence.
(Optional) Install the Document Intelligence Admin application for an improved user experience for process owners to set up, configure, and monitor their document extraction solutions.	Install the Document Intelligence Admin application from the ServiceNow® Store. Review the listing for information on dependencies, licensing or subscription requirements, and release compatibility. Key features of the application include: • Easily configure your document extraction solutions by defining what you want to extract from your documents. • Monitor the performance of your document extraction solutions and adjust the level of automation.
Verify that you've assigned the required ServiceNow AI Platform roles.	The following roles are used across the Document Intelligence features: • The administrator (admin) installs the application from the ServiceNow Store and assigns the roles. • sn_docintel.admin • sn_docintel.creation_agent • sn_docintel.extraction_agent

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Checklist (continued)

Setup task	Description
	• sn_docintel.manager • sn_docintel.viewer
Domain separation	Review the domain separation topic information if you intend to separate data, processes, and administrative tasks. See Domain separation and Document Intelligence.
Verify that file extensions required to support Document Intelligence are included in any glide.attachment.extensions system property customizations.	The glide.attachment.extensions system property is empty by default. When customizing it, ensure that the following file extensions are included: • jpeg • jpg • png • pdf For more information, see Restrict attachment file extensions.
Copy any use cases needed to support your Document Intelligence implementation.	Follow the steps in Duplicate a use case to copy a use case along with its fields, field groups, integrations, flows, and all the related machine learning (ML) models.
Import any use cases needed to support your Document Intelligence implementation.	Follow the steps in Import a use case to import a use case along with its fields, field groups, integrations, flows, and all the related machine learning (ML) models.

Install Document Intelligence

You can install Document Intelligence (sn_docintel) and Document Intelligence Admin (sn_docintel_admin) if you have the admin role. The sn_docintel_admin application installs related ServiceNow® Store dependencies if they aren't already installed.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Review the 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 Install dialog box that is displayed, any dependencies that are installed along 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
Choose a time to schedule the upgrade	Avoid performing the upgrade during the run time of the nightly task definition (use case) upgrade job. Upgrading the application at that

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence 3.0 or later pre-upgrade checklist (continued)

Pre-upgrade task	Description
	time would prevent use cases from upgrading until the run time of the next day.
Update any custom code that points directly to DocIntel database tables	After upgrading to Document Intelligence 3.0 or later from version 2.4 or earlier, DocIntel will use new flow actions. DocIntel will also use ServiceNow AI Platform database tables (sys_di_) in place of the scoped application tables (di_). All DocIntel data will be migrated automatically to these platform tables. **Note:** Document extraction use cases may not be available for use until the migration is complete. If you have custom code in your instance that points directly to the scoped application tables (di_), it should be updated, preferably to the platform component using DocIntel APIs. If you have custom code using DocIntel 2.4 or earlier flow actions, it should be updated to use 3.0 or later flow actions. For the list of tables, see Components installed with Document Intelligence.
Add cross-scope records for integrations	For any integrations with Document Intelligence, add cross-scope records for the new DocIntel database tables. **Note:** Ensure that the existing cross-scope records for the old tables are not removed. See the Cross-scope records section for the list of records to be added. For more detail on cross-scope records, see Cross-scope privilege record.

During the upgrade

Document Intelligence 3.0 or later upgrade checklist

Upgrade task	Description
Avoid using the Document Intelligence workspace to extract fields	Wait until a document task is migrated to the sys_di_task table before completing it

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence 3.0 or later upgrade checklist (continued)

Upgrade task	Description
	using the Document Intelligence workspace. Completing a task during the upgrade may lead to data loss.

After the upgrade

Document Intelligence 3.0 or later post-upgrade checklist

Post-upgrade task	Description
Test DocIntel integrations and custom workflows	Test your use case integrations to ensure they function with the new database schema.
Check for data loss	Some cases where you can lose data include: • Any action resulting in the attachment deletion in an old di_task record would result in a missing attachment for the new sys_di_task record. • Deletion of any new tasks created during the data migration (that is, during or right after the 3.0 or later upgrade) would result in permanent deletion as those records might not have entries in di_task tables. • Deletion of an old "process_task" flow (or any action resulting in that flow deletion) would result in an empty flow reference in the new integration setup records.
Post-upgrade migration of legacy use cases imported through update sets	At any time after the completion of the Document Intelligence 3.0 or later upgrade and data migration, you can rerun the migration batch script that ran during the upgrade. You may want to run this script if you must migrate any imported use cases requiring migration to the platform tables (sys_di_). 1. Navigate to **All** > **System Definition** > **Scheduled jobs**. 2. Open the *`DocIntel migrate remaining data`* scheduled job. 3. Select the **Active** check box. 4. Select **Execute Now**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence 3.0 or later post-upgrade checklist (continued)

Post-upgrade task	Description
	Warning: Never use the *DocIntel migrate remaining data* batch job to rerun migration after deleting all records from sys_di tables. You should only use it to migrate missing data from di_ tables to sys_di tables.

Cross-scope records

Cross-scope records for Document Intelligence 3.0 or later

Source scope	Target scope	Target name	Operation	Target type	Status
Scope of integrating BU	global	sys_di_task	Read	Table	Allowed
Scope of integrating BU	global	sys_di_key	Read	Table	Allowed
Scope of integrating BU	global	sys_di_key_group	Read	Table	Allowed
Scope of integrating BU	global	sys_di_ocr_input	Read	Table	Allowed
Scope of integrating BU	global	sys_di_pdf_input	Read	Table	Allowed
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 **Auto-sync** to enable the exact match option. This lets the user select an icon to

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Setting	Description
	view only the recommendations that match exactly what they type.
Candidate score	Select **Auto-sync** to make confidence scores visible next to each recommendation. For more information, see Confidence scores.
Side panels	Select the location of the navigation panel and document fields panel.
Default width of the extraction panel	Enter the default width of the document fields (extraction) panel in pixels.
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.

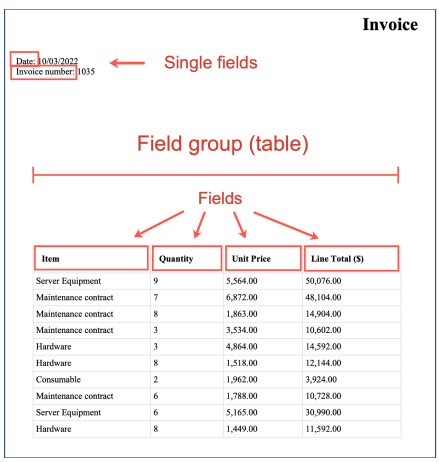
Set up document extraction use cases

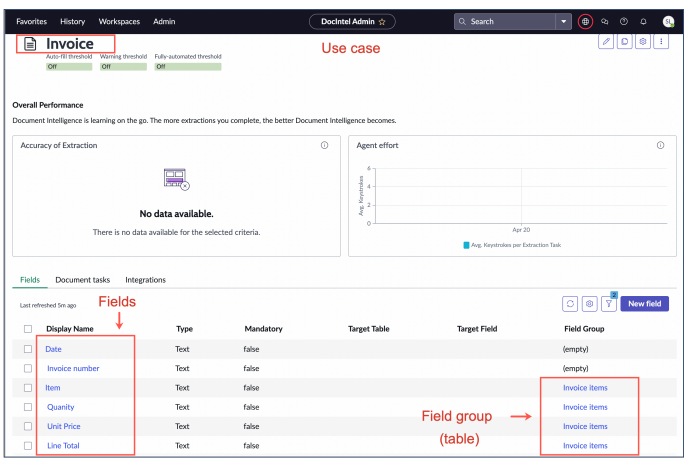
In Document Intelligence, a use case is a template used to define the structure of a type of document you want to process. It's made up of the use case record and its related fields, field groups, integrations, flows, and all the related machine learning (ML) models.

Overview

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.





© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

After you've defined a document 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 and target table for the use case.

2. Create a field.

   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. Select **Save**.

What to do next
After creating a use case, finish setting it up by adding fields and field groups and configuring the data extraction modes.

Create a field

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: sn_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.

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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

The 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	DocIntel provides recommendations for the fields in the Document Intelligence workspace. Choose the recommendation or manually enter the value. All fields must be reviewed. Recommendations are ordered based on how confident the AI is in the prediction. As DocIntel continues processing your documents, recommendations improve over time.
Auto-fill	DocIntel auto-fills the fields in the Document Intelligence workspace. All fields must be reviewed. Auto-fill works only if the AI has enough confidence to make the prediction. You can change the confidence threshold by updating the **Auto-fill threshold** field in the use case.
Fully automated (Straight-through processing)	DocIntel automatically extracts the data for all fields and processes the document task if the confidence scores for all required fields are above the defined confidence threshold. Fields don't need to be reviewed. DocIntel becomes more confident over time, as it processes more and more documents. Choose Fully automated mode for frequently processed documents or if you're confident in the system.

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Use Cases**.
2. Select the use case that you would like to configure.
3. Select the settings icon ().
4. Select the extraction mode for the use case.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If the DocIntel AI model needs further training to make recommendations with higher confidence scores, leave the default recommendation mode in place without selecting any other mode.
- If DocIntel provides recommendations with confidence scores above the specified threshold, select the **Auto-fill mode** option.
- If DocIntel auto-fills the required fields with very high confidence scores, select the **Fully Automated mode** option.

5. **Optional:** Adjust the auto-fill threshold and warning threshold for Auto-fill mode.

   - (Optional) **Auto-fill threshold**: DocIntel only auto-fills the fields if the confidence score of the top recommendation is at or above the percentage you define.

     Fields with a confidence score lower than the threshold are left empty in the Document Intelligence workspace, and the recommendation mode is available to extract these fields.

   - (Optional) **Warning threshold**: DocIntel shows a warning for empty fields and auto-filled fields with a confidence score at or below the percentage that you define.

6. **Optional:** Adjust the confidence threshold for required fields in Fully automated mode.

   - (Optional) DocIntel automatically completes and submits the document task if all required fields have a confidence score at or above the percentage you define in the **Fully-automated threshold** field.

   - (Optional) If any required field is below the threshold, the document task isn't submitted automatically and requires agent review. Fields that aren't required may be left empty or unreviewed.

   - (Optional) If there are no fields defined as required for the document task, then DocIntel automatically completes and submits the document task.

7. Select **Save**.

Manage document extraction use cases

Manage your Document Intelligence use cases to efficiently support your document processing 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 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

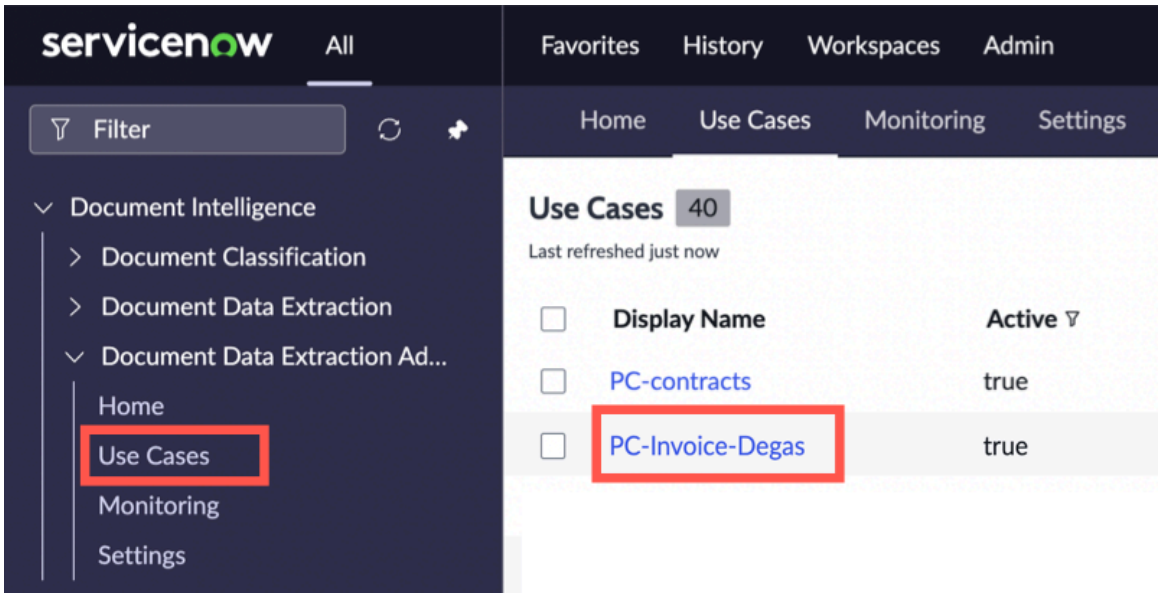
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

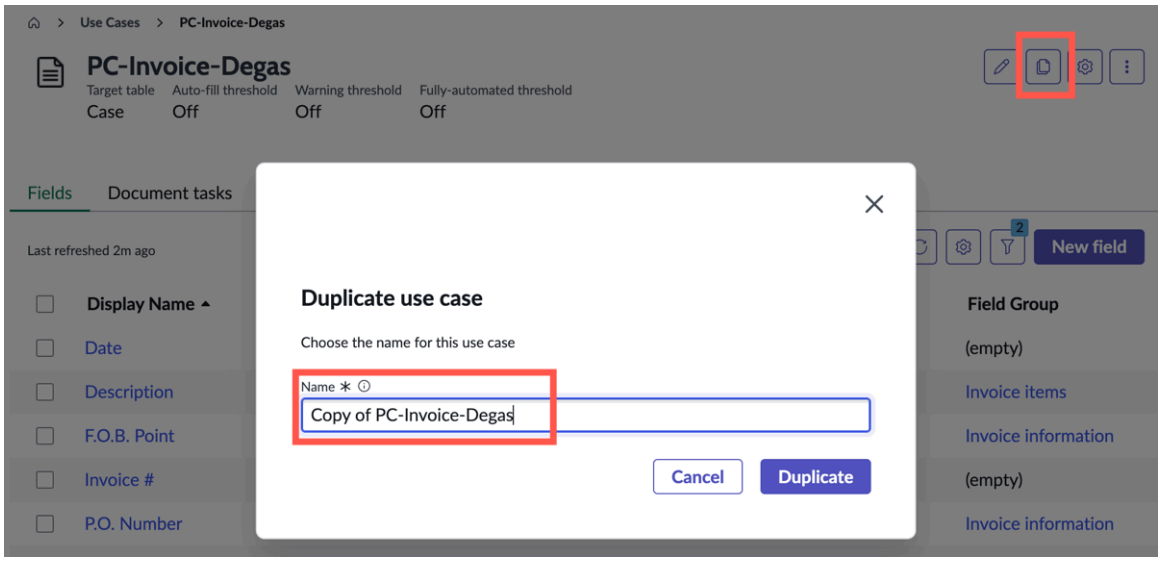
Follow these steps to create a copy of a use case along with its fields, field groups, integrations, flows, and all related machine learning (ML) models.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Export a use case

Export a data 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 data 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 use case

Import a data extraction use case for use in your ServiceNow instance.

Before you begin

- Update sets for a data extraction use case are downloaded according to the steps provided in Export a 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 data 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Procedure

1. Navigate to **All** > **System 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 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

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

For example, if you 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. Enter a name for the use case.
4. Select **Submit**.
5. Navigate to **All** > **Document Intelligence** > **Document Classification** > **Use Cases**.
6. Select the use case that you created.
7. On the form, edit the fields as needed.

Use case form

Field	Description
Display Name	The name for the use case as it appears in the Document Intelligence workspace.
Autofill	Option to enable Auto-fill mode. Otherwise, the default recommendation mode remains enabled.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	DocIntel auto-fills the classes (fields) for each page of the document task in the Document Intelligence workspace. All fields must be reviewed. Auto-fill works only if the AI has enough confidence to make the prediction. You can change the confidence threshold by updating the **Autofill Threshold** field in the use case.
Autofill Threshold	DocIntel only auto-fills the classes (fields) if the confidence score of the top recommendation is at or above the percentage you define. Fields with a confidence score lower than the threshold are left empty in the Document Intelligence workspace, and the recommendation mode is available to extract these fields. This field is available only if the Auto-fill mode is enabled.
Warning Threshold	DocIntel shows a warning for empty fields and auto-filled fields with a confidence score at or below the percentage that you define. This field is available only if the Auto-fill mode is enabled.
Enable Straight Through Processing	Option to enable the Straight Through Processing mode. DocIntel automatically classifies all document pages and processes the document task if the confidence scores for all pages are above the **Straight Through Processing Threshold** defined in the use case. The classifications do not need to be reviewed. DocIntel becomes more confident over time, as it processes more and more documents, and those documents are used to Train the use case. Choose the Straight Through Processing mode (Fully automated mode) for frequently processed documents or if you're confident in the system.
Fully Automated Threshold	Confidence score threshold for document classifications, which enables a document

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
	task to be fully automated in the Fully Automated (Straight Through Processing) mode.
Target Table	The table that stores the document processing results for this use case.

8. Select **Update**.

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.

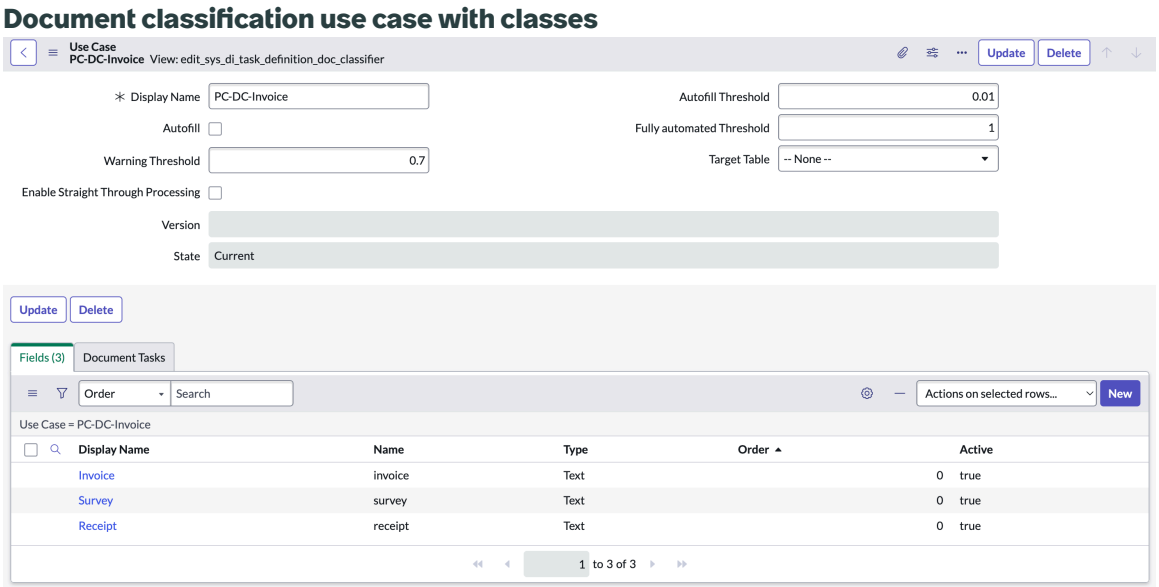
Define a new field form

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



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.

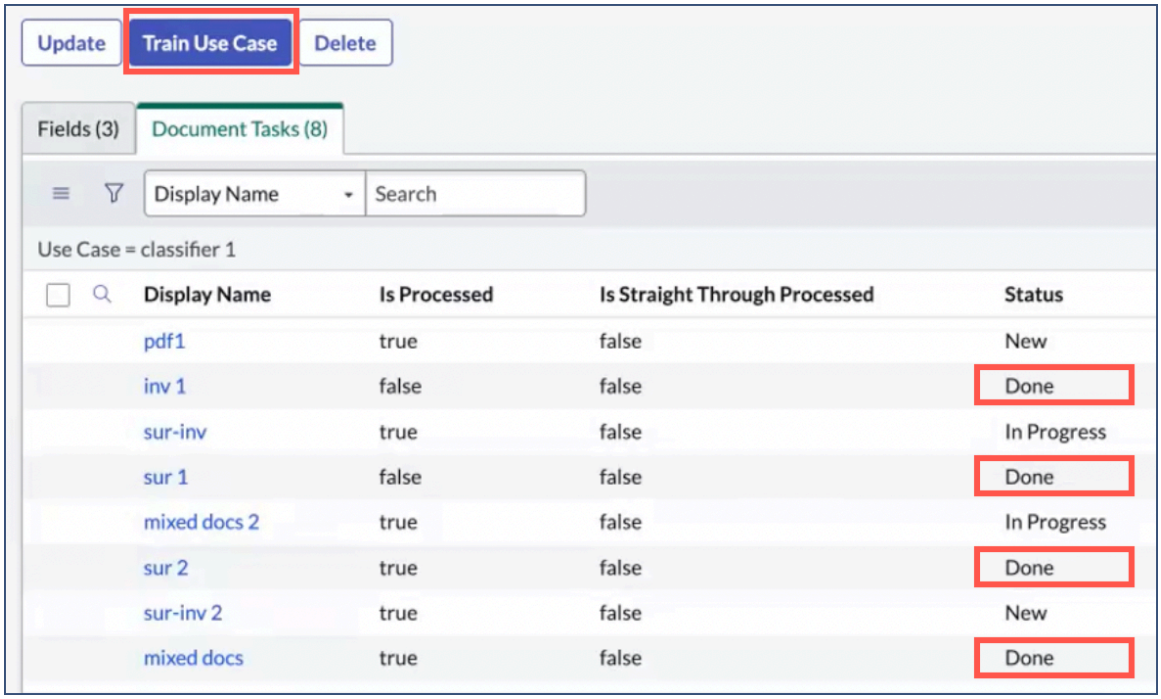
Note:
To reduce server load and minimize performance issues, the default limit for training a use case is once every 30 days.

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Classification** > **Use Cases**.
2. Select a use case in the list.
3. Select **Train Use Case**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Is Reviewed	Indicates whether this field value has been reviewed by a user.
Candidate ID (Recommendation ID)	Internal system ID for the selected recommendation.
Index	For fields that are part of a field group: the order of the field value in reference to other field values for the same field. For regular fields, the index is always 0.
Exact Match of Candidate (Recommendation)	Indicates whether the field value exactly matches the AI's top recommendation.
Field	Field from the use case that the value belongs to.
Candidate (Recommendation) Rank	Rank that the AI assigned to the selected recommendation.
Is Flagged	Whether a user has flagged this field value in the Document Intelligence workspace.
Document task	Document task with the attached document from which the data value was extracted.
Availability	Indicates whether this field value was available or missing in the document.
Keystrokes	Number of keystrokes that were needed to extract this field value.
Target Record	The record where the field value is used.
Metadata	The metadata associated with the field value.
Target Table	The table that stores the field values.
Domain	Domain linked to this field value. See Domain separation and Document Intelligence.

4. Select **Update**.

Integrating Document Intelligence with other applications

Extend the capabilities of Document Intelligence to other ServiceNow applications. Other applications can take advantage of document extraction using Document Intelligence.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Integrate with 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.
- Role required: sn_docintel.admin or sn_docintel.manager

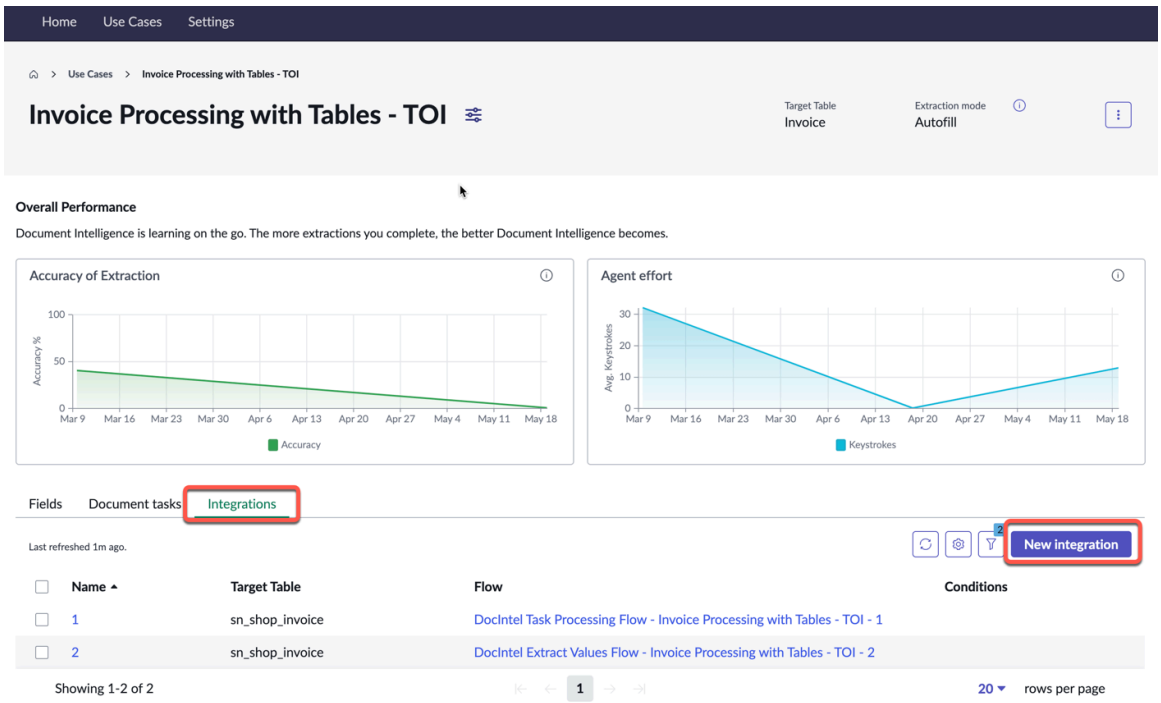
About this task

Define integration points for your Document Intelligence solution. Two integration points are provided: 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.

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Use Cases**.
2. Select the use case for which you want to set up integration points.
3. 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**.



4. On the form, fill in the fields.

Field	Description
Name your Integration	Name for the integration.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Field	Description
Use case	[Read only] Use case to use for the integration task.
Where do you want to take the documents from and store the extracted data?	[Read only] Table to receive data from or send data to. **Note:** The target table is taken from the use case.
What type of integrations you want to set?	**Process Task** or **Extract Values**. The **Process Task** type creates an integration point to automatically create and process DocIntel document tasks based on specific triggers happening in the target table. The **Extract Values** type creates an integration point to automatically propagate the extracted values to the target table when extraction has been completed in DocIntel (that is, when the document task status changes to Done).
Conditions	Filters used to select certain fields to use as specific triggers for the integration. Process Task only.
Create Flow	Select this option to create a Flow Designer flow for this integration. This option should be selected, unless you're planning to write your own custom script to set up the integration.

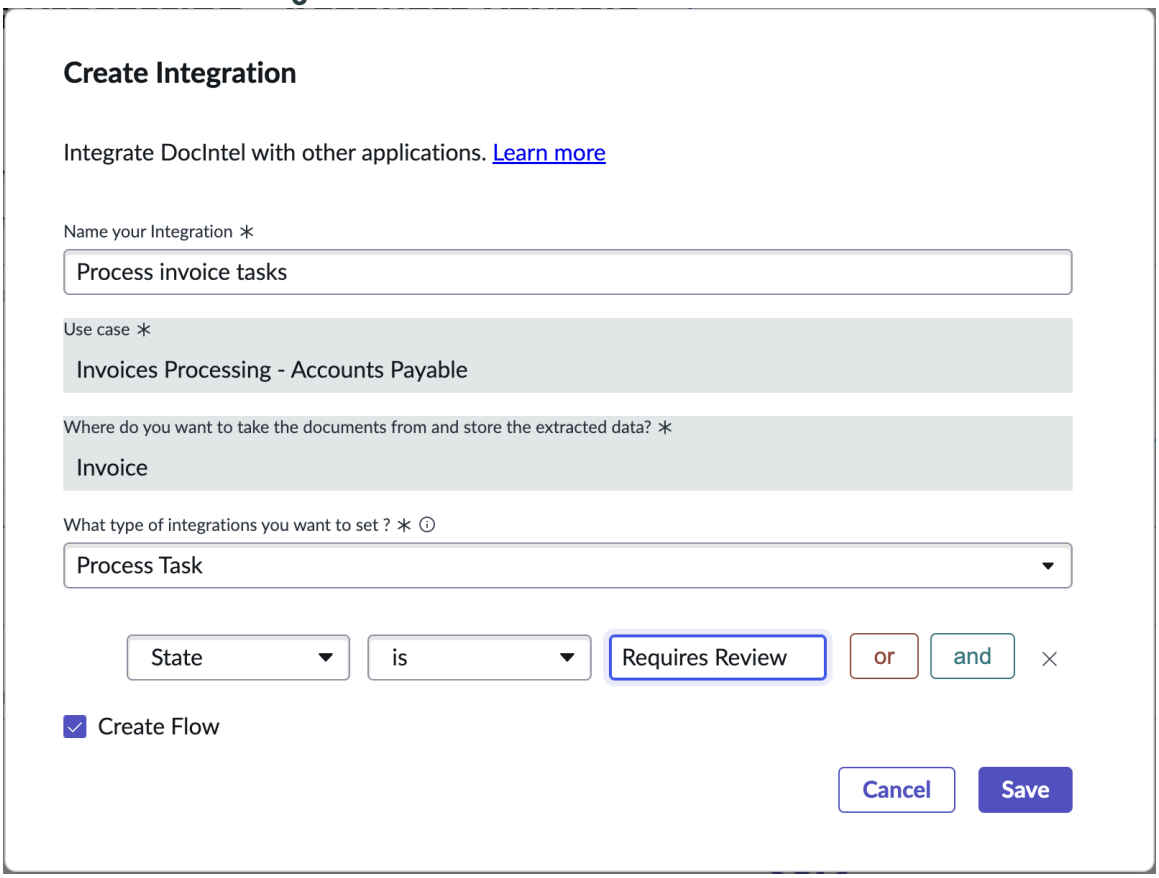
5. Select **Save**.

Example: Integration

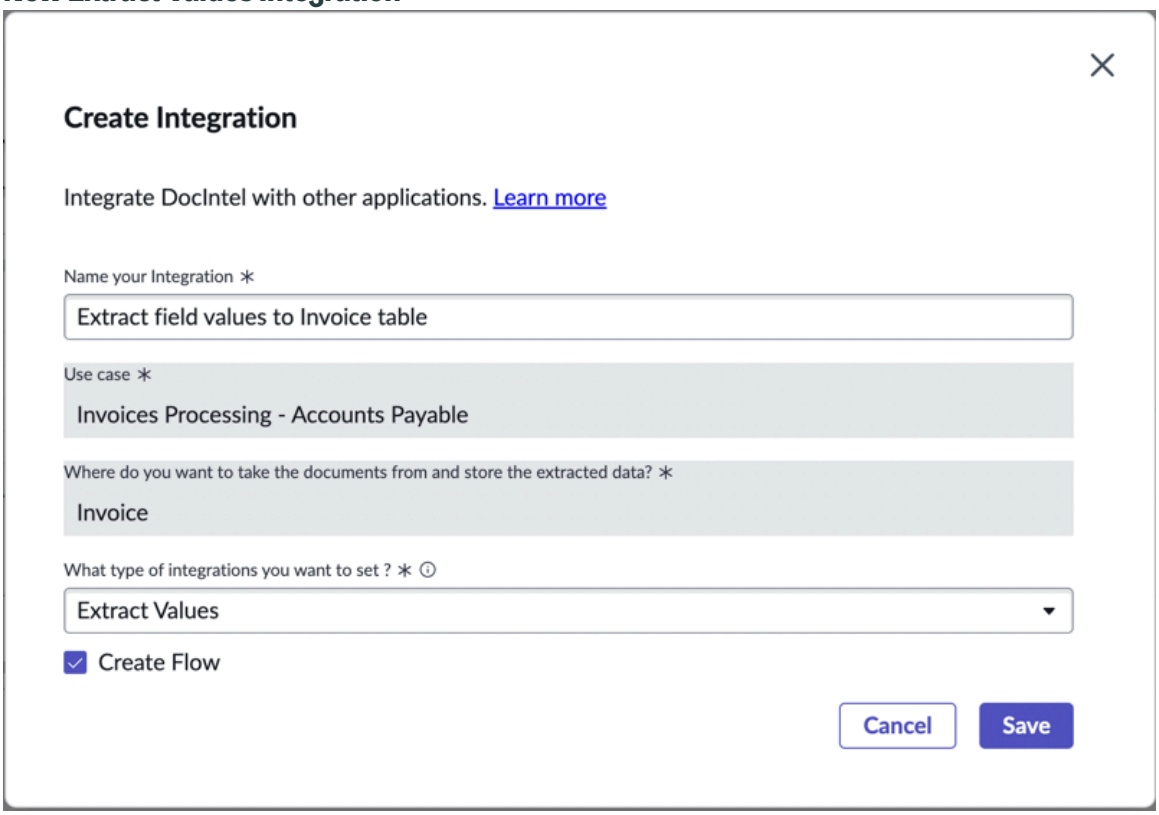
The following images show two example integrations. The first image is a Process Task integration that triggers when a record needs review. The second image is an Extract Values integration that can automatically send extracted fields to the invoice table.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

New Process Task integration



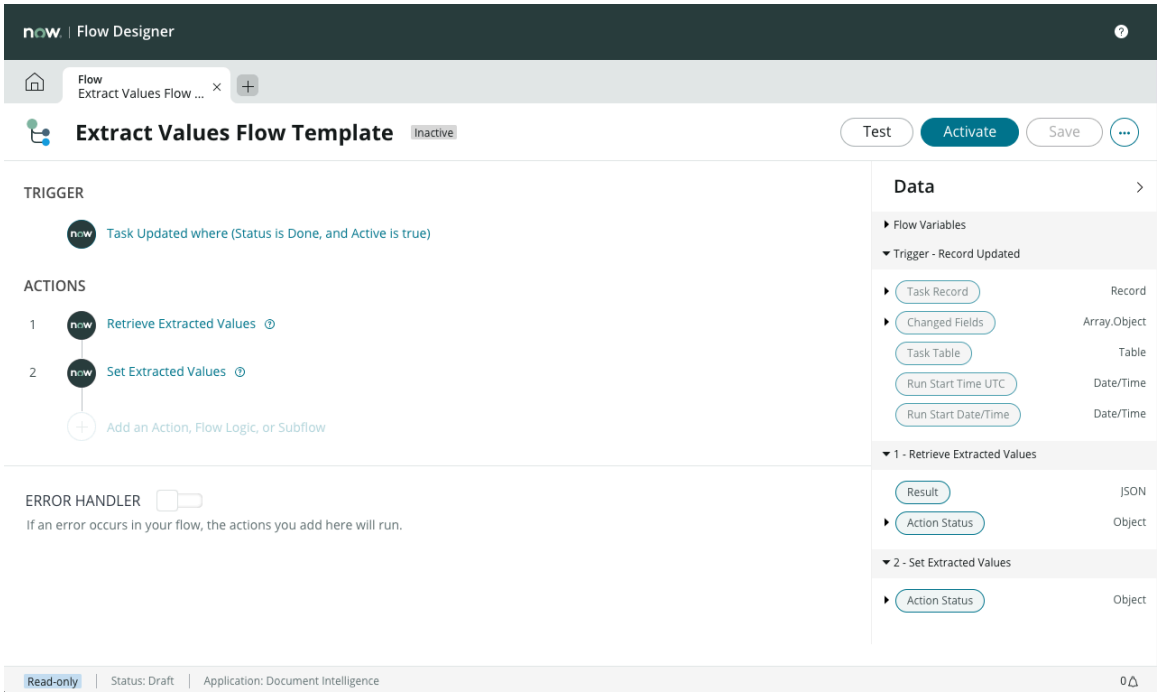
New Extract Values Integration



What to do next

If you selected to create a flow, finish the activation in Flow Designer.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



For more information, see Building flows 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- If the 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).

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Accounts Payable Operations uses the capabilities of Document Intelligence to extract information quickly and accurately from invoice documents that are received as email attachments and then create invoice records in the Accounts Payable Operations application.

For more information, see Accounts Payable Operations integration with Document Intelligence.

How Document Intelligence works with Accounts Payable Operations

DocIntel for Accounts Payable Operations extracts the required information from invoices received as email attachments, creates invoice records in Accounts Payable Management, and adds the extracted information into the invoice records.

For more information, see How Accounts Payable Operations integration with Document Intelligence works.

Enable Document Intelligence for Accounts Payable Operations

Ensure all the necessary applications and plugins are installed and activated to enable Document Intelligence in the Accounts Payable Operations application plugin.

See Components installed with Accounts Payable Operations integration with Document Intelligence.

Integrate with Automation Center

Use Automation Center to discover opportunities to automate document processing with Document Intelligence.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create a document task

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.

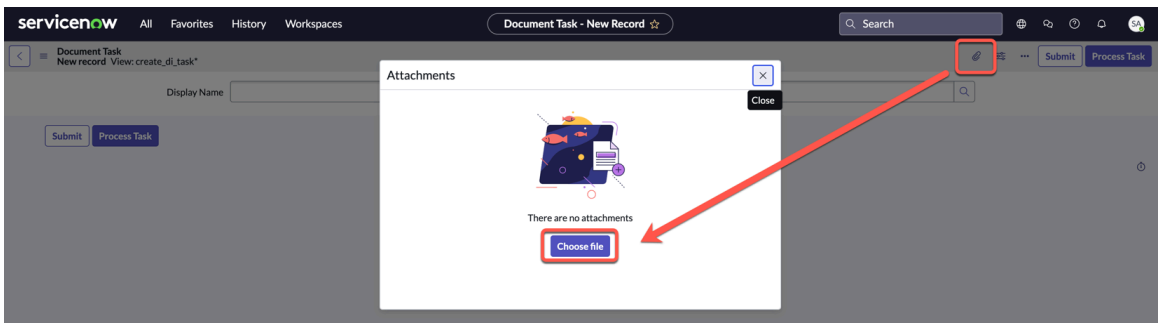
Procedure

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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

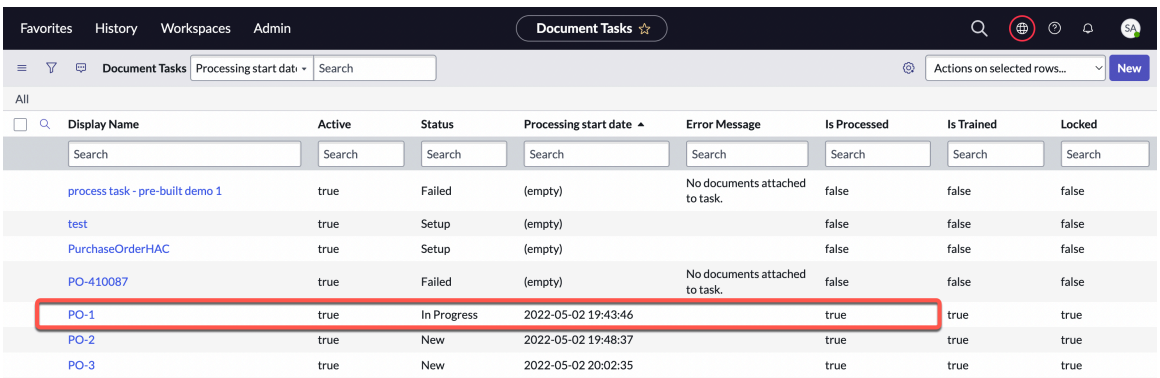
Procedure

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.

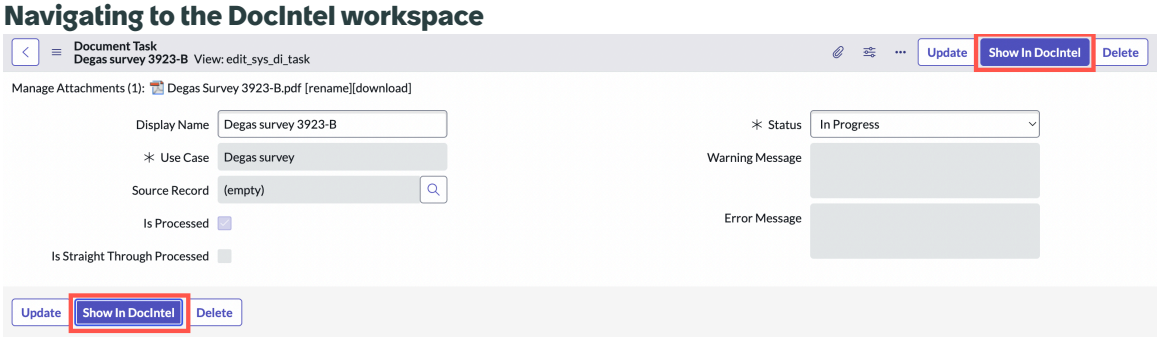
   If the status appears as Setup, go to the document task record and select **Process Task**.

Document tasks list



3. Select the processed document task.
4. Select **Show in DocIntel**.
   The following image shows how to navigate to the Document Intelligence workspace for data extraction.

Navigating to the DocIntel workspace



   The Document Intelligence workspace opens in a new tab.
5. Use the Document Intelligence workspace to review the document fields and select the appropriate recommendations.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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.

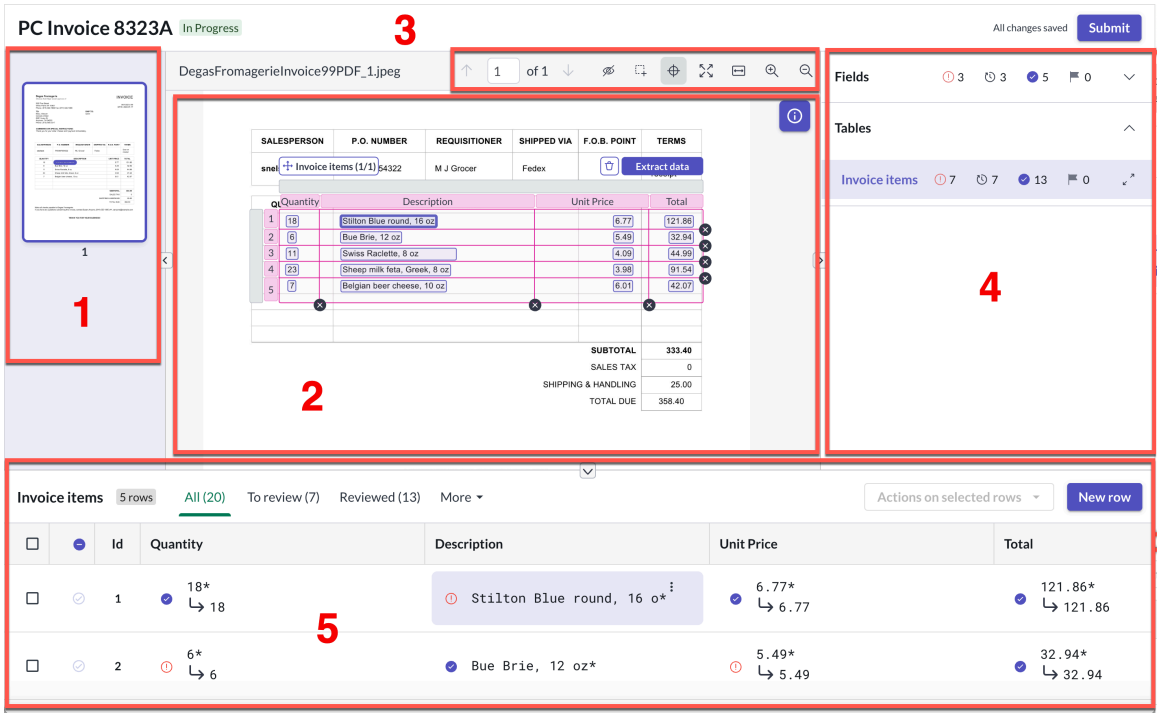
Layout of the Document Intelligence workspace

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

Note:
In this view, the document fields panel is expanded so that the fields are visible.

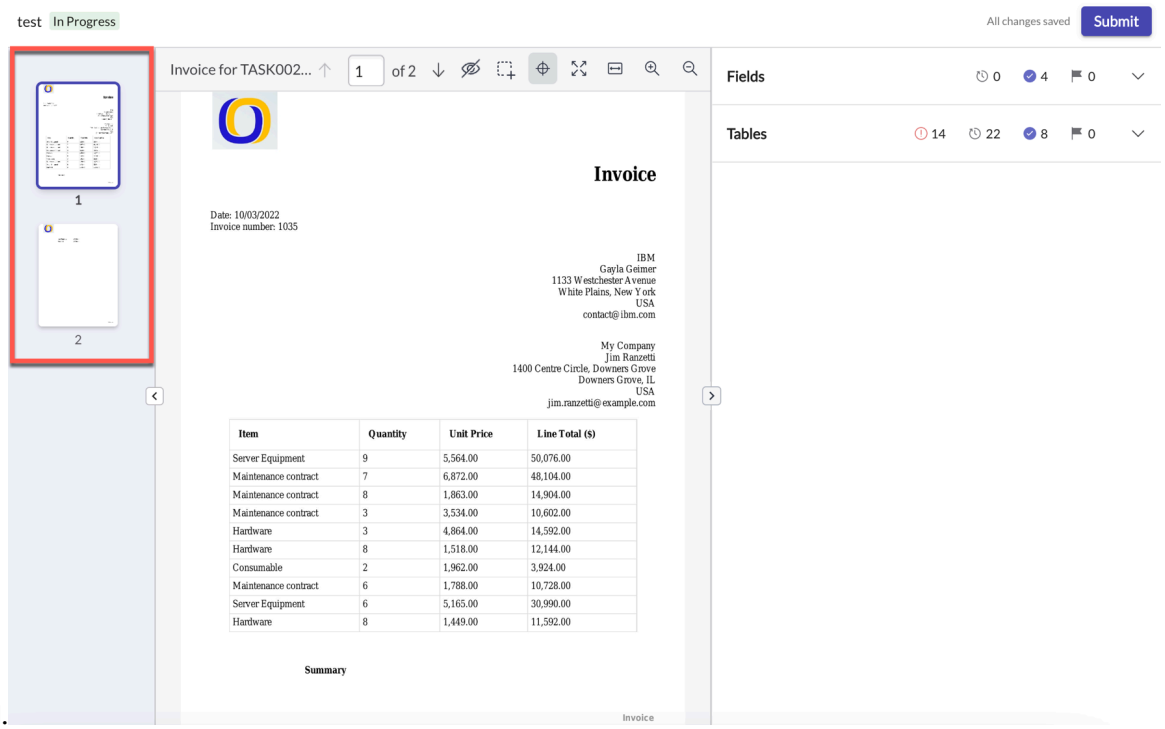
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Thumbnail panel

In the thumbnail panel of the workspace, you can select one page from a multiple-page document. The selected page is displayed in the document image panel. Selecting a page doesn't affect what is displayed in the document fields panel.

The following image shows a document with two pages in the thumbnail panel.

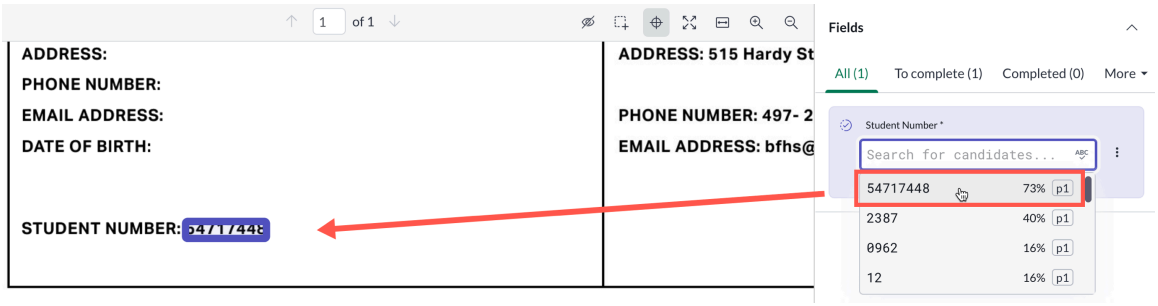


Document image panel

The document image panel displays the page selected in the thumbnail panel.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

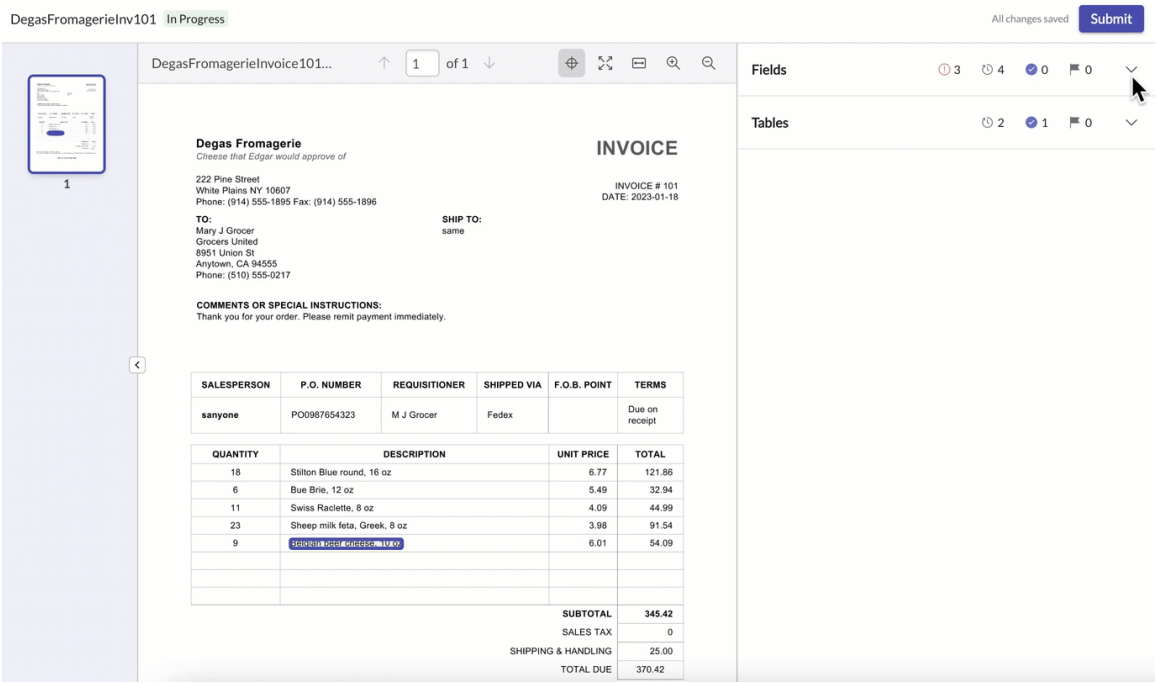
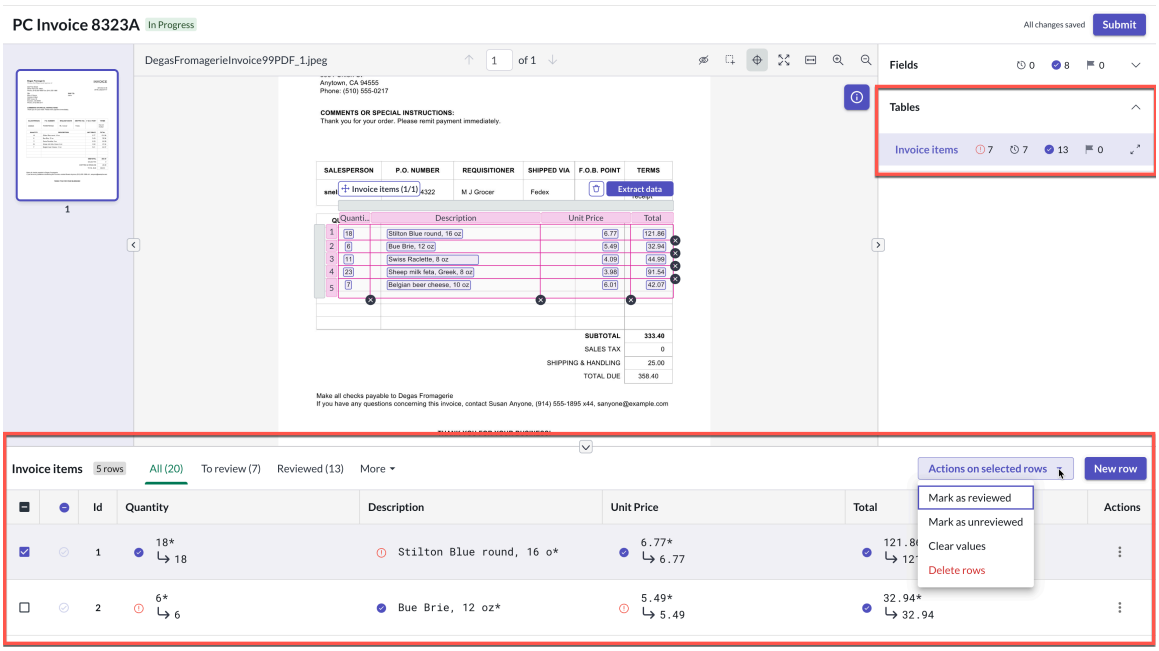


Table panel

The table panel enables you to open table rows for review, including viewing the AI's recommendations. You can also flag fields or mark fields as missing in the document. Other table controls enable you to insert rows and resize columns as needed.

Note:
An asterisk indicates a required field.

The following illustration shows the different features of the table panel.



Extract single fields

Extract text and number fields from your document in the document fields panel.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: sn_docintel.extraction_agent

About this task
Use the following steps to extract single fields from a document.

If the fields are already auto-filled with values, you can review them to ensure they are correct or adjust the fields as needed.

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.

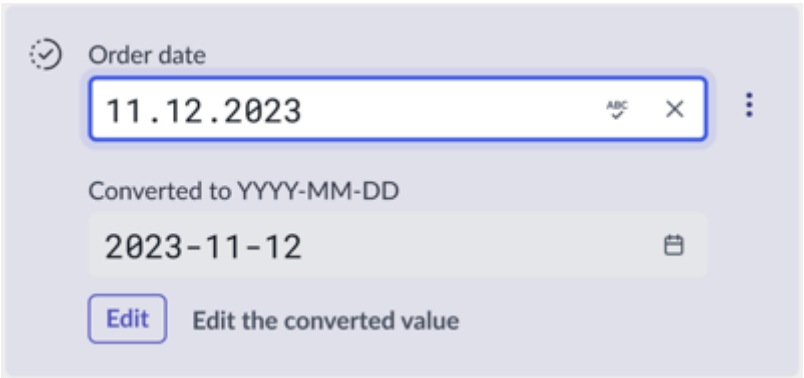
   The recommendation with the highest confidence score displays at the top of the list.

   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.

   The field will be marked as complete when you move to the next field.
5. **Optional:** If needed, correct any field value conversions.

   (Optional) Some field types convert the extracted value into a standard format. See Data normalization.

   a. Select **Edit** in the converted value field.

   b. Enter the updated conversion value.

   c. Select **Save**.

Extract check box fields

Extract check box fields from your document in the document fields panel.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Before you begin
Role required: sn_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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

About this task

Use the following steps to 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.

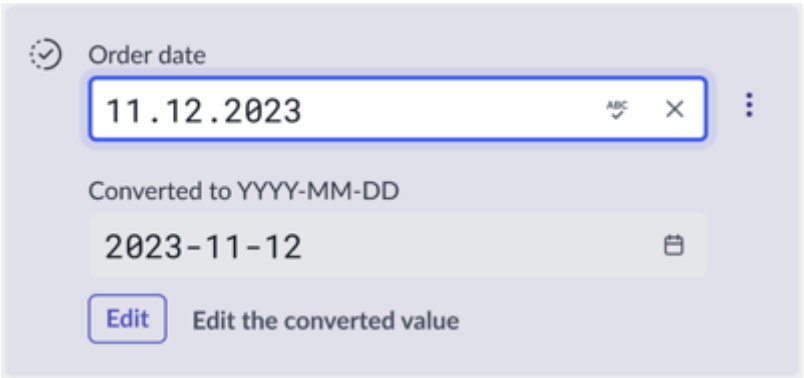
Procedure

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.
4. **Optional:** If needed, correct any field value conversions.

   (Optional) Some field types convert the extracted value into a standard format. See Data normalization.

   a. Select **Edit** in the converted value field.
   b. Enter the updated conversion value.
   c. Select **Save**.
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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- To add 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.

6. **Optional:** [store-future: BEGIN review] If needed, show or hide columns in the table panel.

   a. Select the settings icon ().

   b. Select the columns that you want to display.
   Deselect the columns that you want to hide.

   c. Select **Apply changes**.

   [End]

7. 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 tables on the document image.

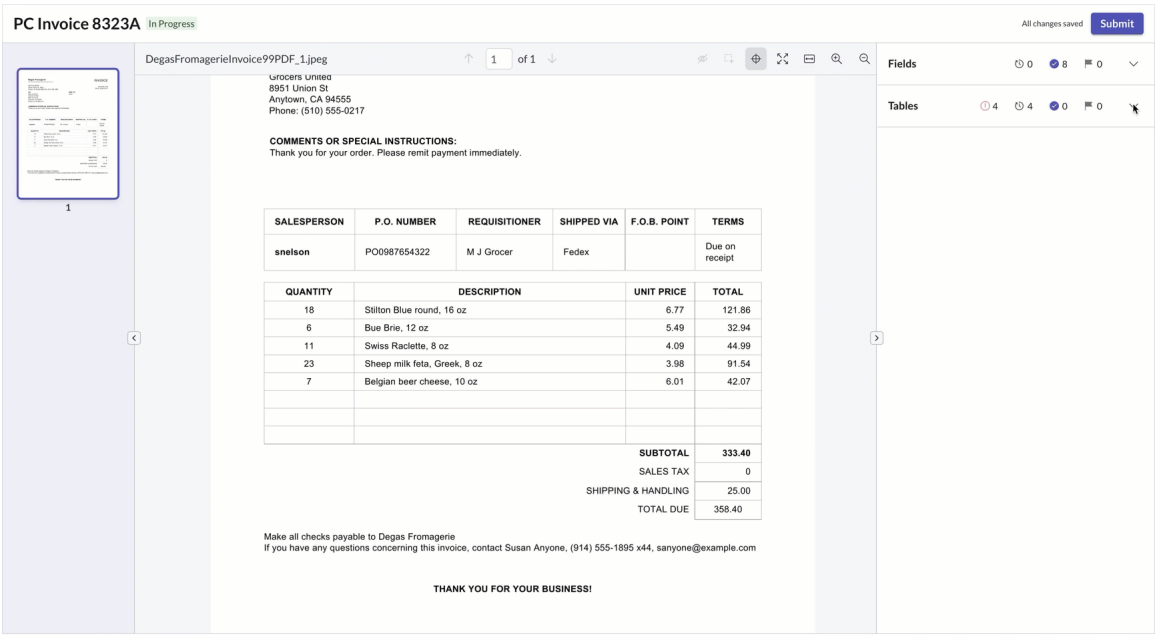
Overview of the draw tool

When viewing a processed document in the Document Intelligence workspace, you can use the draw tool to review and extract 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. You can review the extracted fields right from the grid 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Extract 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.

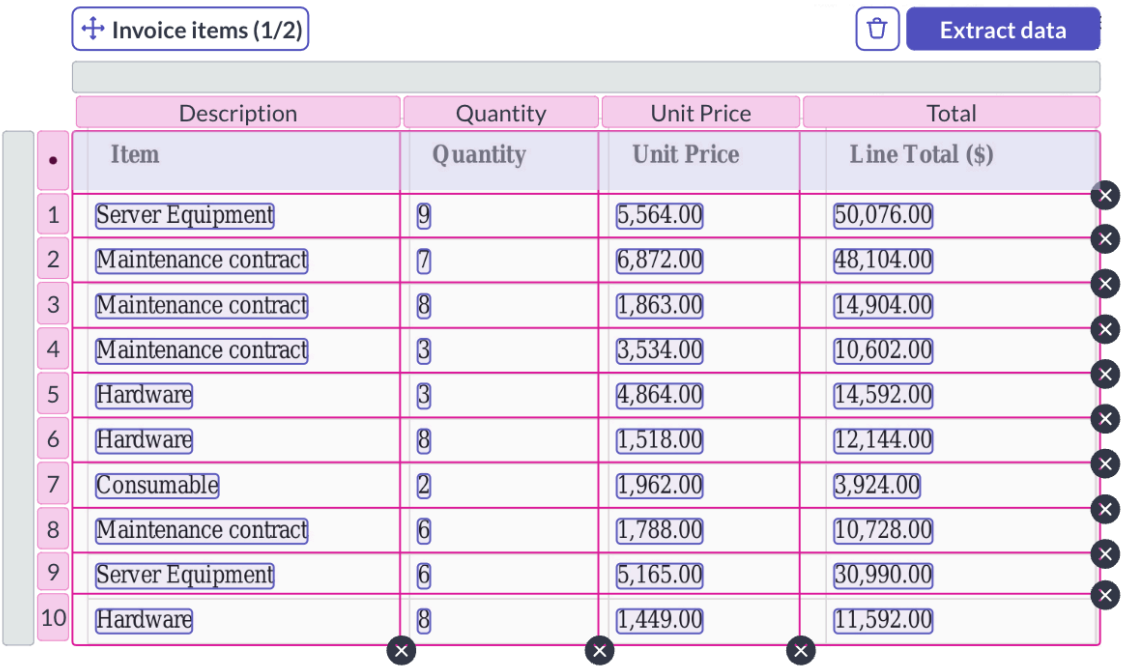
   Tip:

   Include the header row of the table in the selection area.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

An 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.



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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- Select **Extract data**.

  This button displays if no data has been extracted yet.

- 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

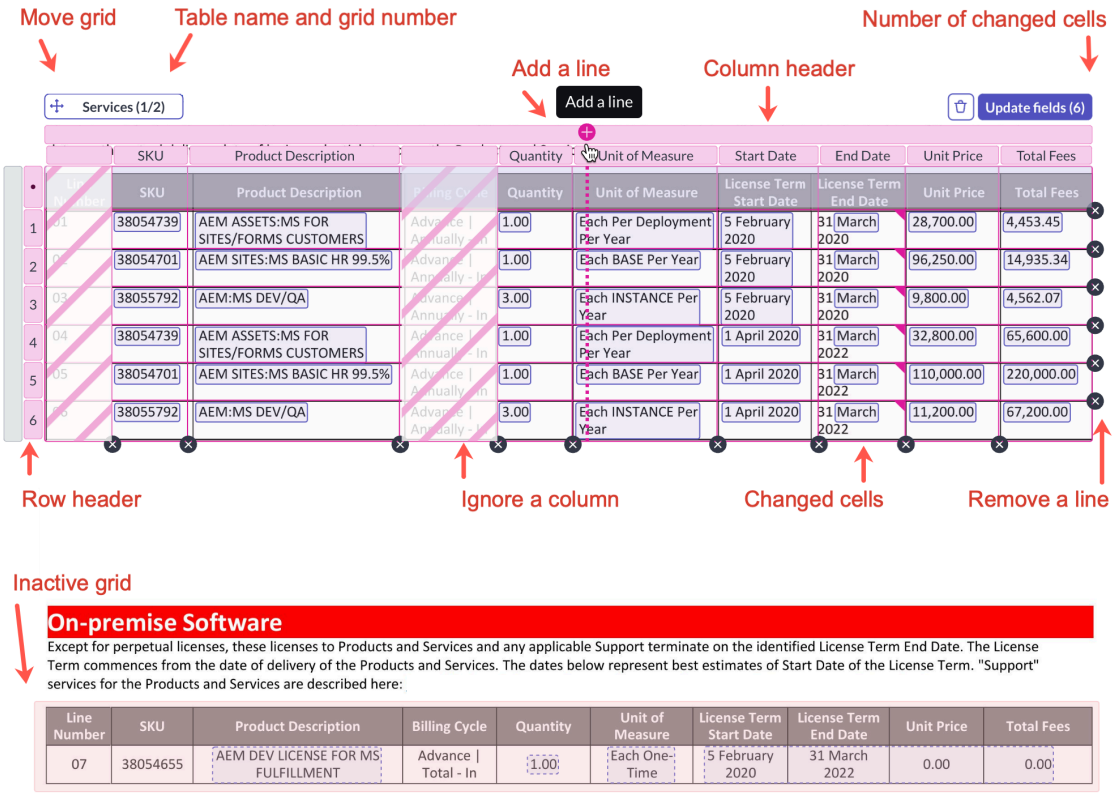
Use the following steps to make changes to a grid. You can then extract the updates to the related table fields.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



2. Adjust the grid as needed.

   Adjustments are auto-saved.

   Any adjustments that produce a change to previous extractions are indicated in the cell corner.

   The text in the cell appears surrounded by a solid-line box if it's already extracted. It appears in a dotted-line box if it isn't extracted yet.

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

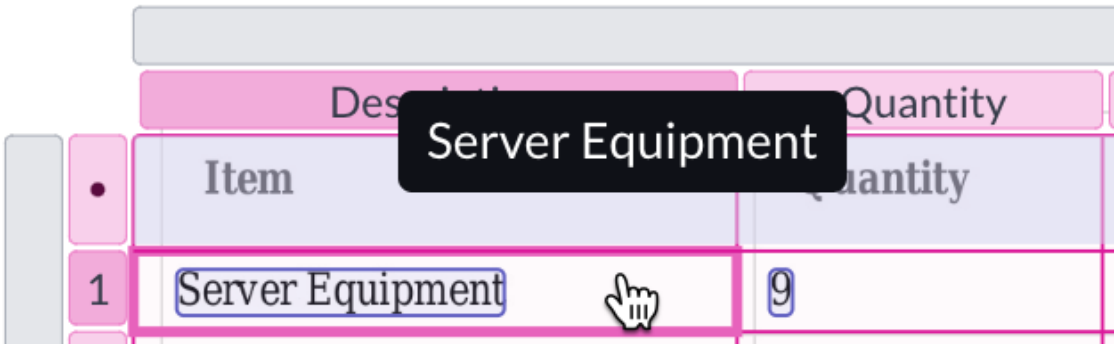
- These steps apply to tables that have a draw tool grid on the document. For more information, see Extract a table using the draw tool.
- Role required: sn_docintel.extraction_agent

Procedure

1. On the document image, point to a cell you want to review in the grid.

   If there's an extracted value, the text in the cell appears surrounded by a box.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



2. Verify the value that appears in the field's tooltip.
   - If the value is accurate, move on to review the next cell.
   - 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 on the document. For more information, see Extract a table using the draw tool.
- Role required: sn_docintel.extraction_agent

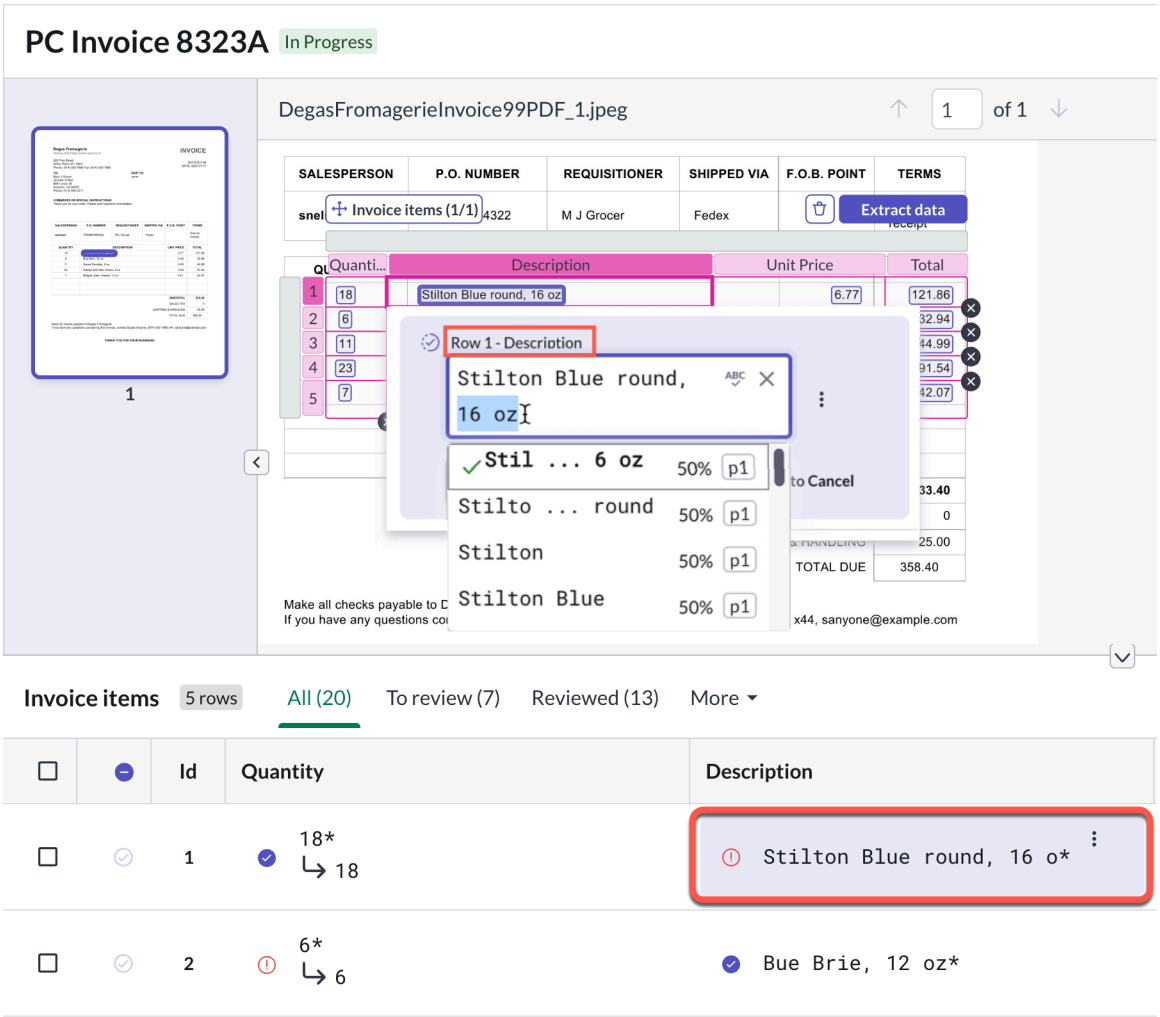
About this task

Use the following steps to 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 cell that you want to edit in the grid.
   A dialog box appears, showing the related table field.
2. Enter a value in the text box or select another recommendation from the list.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



3. Select **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.

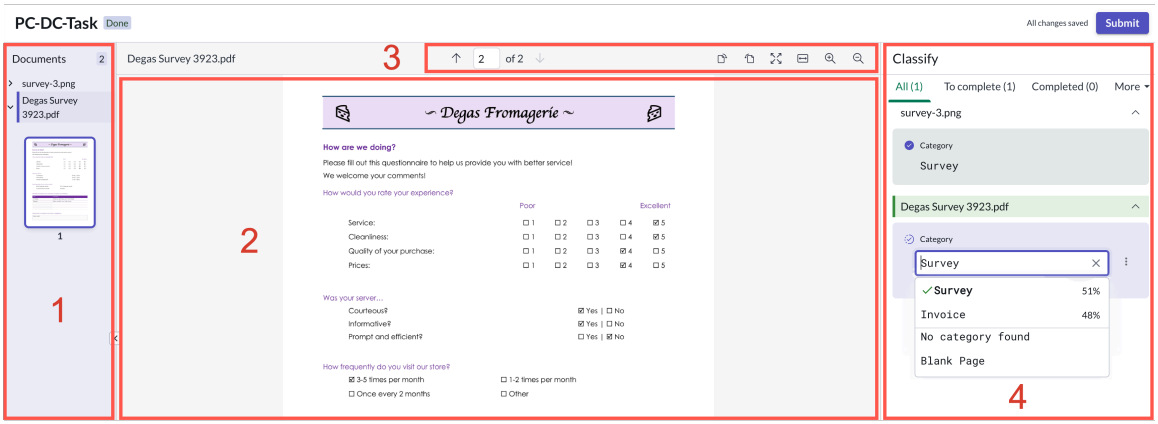
Layout of the Document Intelligence workspace

The following image shows the Document Intelligence workspace for a document task. The workspace includes the following areas:

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- 1 - 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

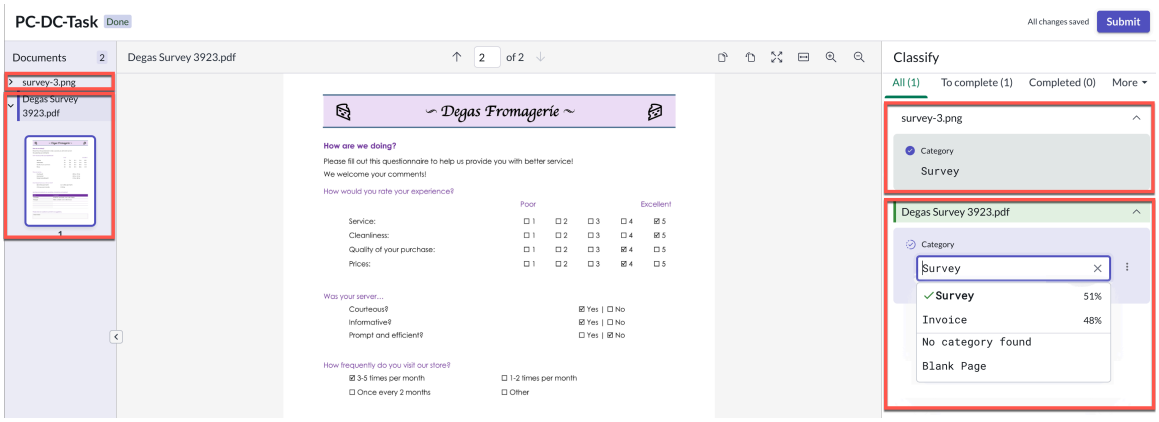
When you're reviewing a document for classification, you can use various controls to rotate, maximize the viewing area, zoom, or focus on the areas that you need.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Classify documents and document pages

Classify your documents in the Classify panel.

Before you begin

Role required: sn_docintel.extraction_agent

About this task

Use the following steps to apply categories to document attachments and each of their pages.

In this context, a document page is the individual page object of the digital document file, as displayed in the thumbnail panel of the workspace. A document file may contain multiple logical groupings of documents (like a PDF file with several invoices, each spanning several pages). Based on the contents, apply the categorization to the document file and the file pages.

If the categories are already auto-filled with values, you can review them to ensure they are correct or adjust them as needed.

Procedure

1. In the Classify panel, review any auto-filled categories for each document to ensure they're correct.
2. Select a category for a document by selecting the **Category** field.
3. Select an appropriate recommendation for the document using the recommendation list.
   - Begin typing in the **Category** field and select a recommendation from the list.

     The recommendation with the highest confidence score displays at the top of the list.
   - Select **Mixed categories** for a multi-page document that includes more than one document type.

     Otherwise, if you select a category for a multi-page document, that category is applied to all of the document's pages.
   - Select **Blank Page** for a blank page in a document.
   - Select **No category found** if you can't find a correct match.

     The document will be marked as complete when you move to the next **Category** field.
4. Select a recommendation for each page of a Mixed categories document.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The page will be marked as complete when you move to the next page.

5. **Optional:** If needed, flag the documents or pages for later attention by selecting **Flag for follow-up** in the options menu from the options icon ( ).

Monitoring Document Intelligence performance

Track document extraction performance in Document Intelligence to understand its usage and effectiveness.

Monitor performance

Monitor document extraction performance in Document Intelligence over time in 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.

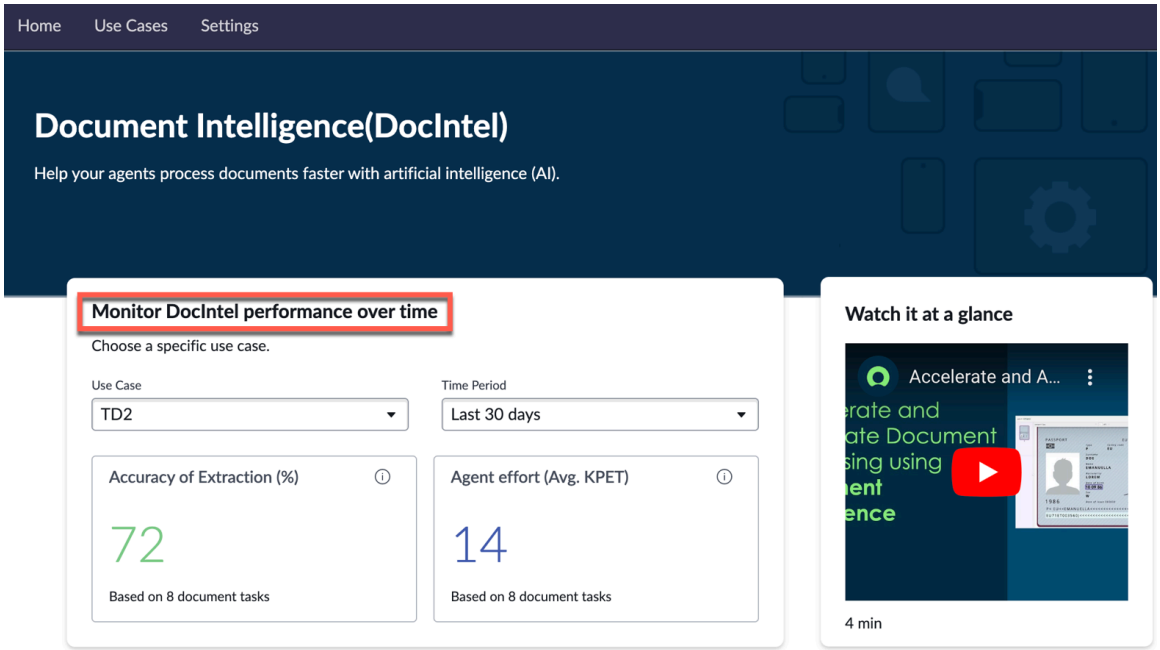
About this task

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction Administration** > **Home**.
2. Expand the **Use Case** list and select your use case.
3. Expand the **Time Period** list and select a date range.
4. Review the results displayed in the **Accuracy of Extraction** and the **Agent effort** widgets.

   You can also use the Document Intelligence dashboard to review other key metrics such as monthly transactions. For more information, see Using the Platform Document Intelligence Usage dashboard.

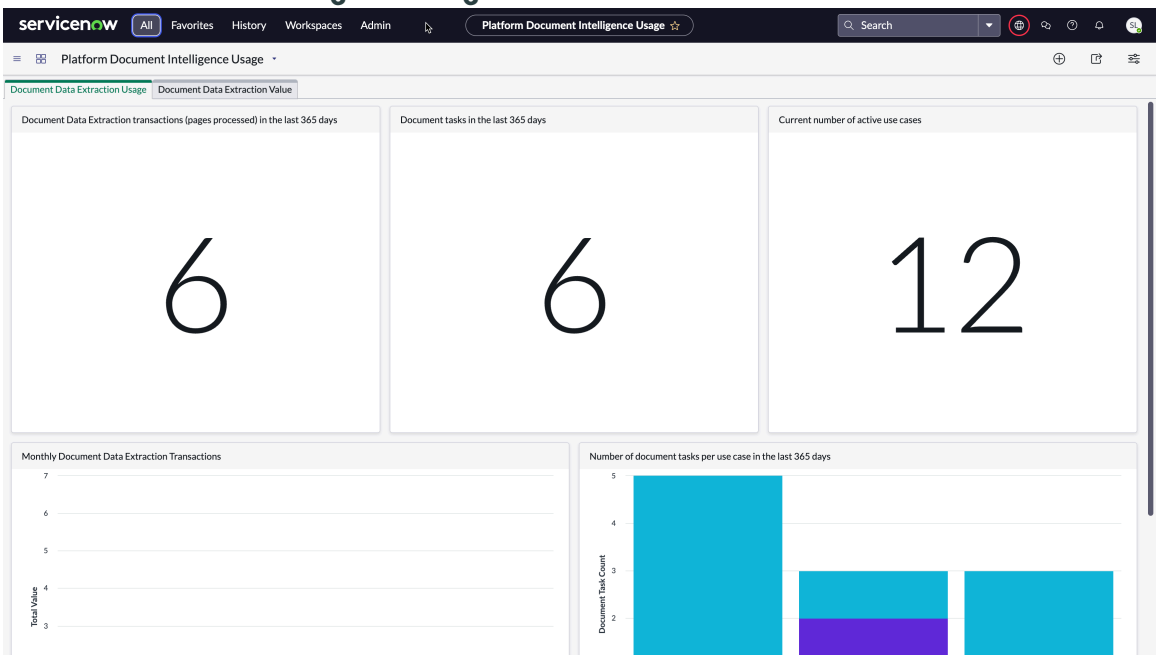
Using the Platform Document Intelligence Usage dashboard

The Platform Document Intelligence Usage dashboard provides a high-level overview of your Document Intelligence usage and value in the form of graphical charts. To view the data, drill down from each chart.

- The **Document Data Extraction Usage** tab provides key insights like pages processed and monthly transactions.
- The **Document Data Extraction Value** tab shows the number of field values extracted in the last 365 days, the accuracy per use case, and the number of keystrokes required per use case.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Platform Document Intelligence Usage dashboard



Analyze Document Intelligence usage and value

Understand the usage and effectiveness of your Document Intelligence application.

Before you begin

Role required: sn_docintel.admin, sn_docintel.manager

Procedure

1. Navigate to **All** > **Document Intelligence** > **Document Data Extraction** > **Dashboard**.
2. Select any of the following widgets to see additional details:

Category	Widget	Description
Document Data Extraction Usage	Document Data Extraction transactions (pages processed) in the last 365 days	The number of pages processed for data extraction using Document Intelligence in this instance in the last 365 days.
Document Data Extraction Usage	Document tasks in the last 365 days	The number of Document Intelligence document tasks created in this instance in the last 365 days.
Document Data Extraction Usage	Current number of active use cases	The number of document extraction use cases currently active in this instance.
Document Data Extraction Usage	Monthly Document Data Extraction transactions	The number of pages processed using Document Intelligence in this instance, per month.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Category	Widget	Description
Document Data Extraction Usage	Number of document tasks per use case in the last 365 days	The number of document tasks processed using Document Intelligence in this instance in the last 365 days, broken down by document extraction use case and task status.
Document Data Extraction Value	Number of field values extracted using Document Data Extraction in the last 365 days	The number of field values extracted using Document Intelligence in this instance in the last 365 days.
Document Data Extraction Value	Average accuracy per use case	The average extraction accuracy per document extraction use case, per day.
Document Data Extraction Value	Average number of keystrokes per use case	The average number of keystrokes needed per document extraction use case, per day.

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.	• platform_ml_di.admin • action_designer • flow_designer • sn_ace.ace_user • canvas_user • usage_admin
DocIntel Viewer [sn_docintel.viewer]	Has view-only access on Document Intelligence document tasks that they are authorized to view.	• snc_read_only • platform_ml_di.viewer
DocIntel Extraction Agent	Extracts data and text from documents using the	platform_ml_di.extraction_agent

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Role title [name]	Description	Contains roles
[sn_docintel.extraction_agent]	Document Intelligence workspace.	
DocIntel Creation Agent [sn_docintel.creation_agent]	Extracts information from documents using the Document Intelligence workspace. Also enables users to create Document Intelligence document tasks and submit them for processing.	platform_ml_di.creation_agent
DocIntel Manager [sn_docintel.manager]	Creates and edits use cases, fields, field groups, and document tasks. Views, measures, and analyzes the usage and effectiveness of Document Intelligence using the Platform Document Intelligence Usage dashboard. Grants access to submit document tasks and interact with the Document Intelligence workspace.	• platform_ml_di.manager • action_designer • flow_designer • sn_ace.ace_user • canvas_user • usage_admin

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
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 [sys_di_metrics_job_log]	[Internal table] Contains a log of the metrics daily aggregation jobs that happened in the instance.
DocIntel Raw Metrics [sys_di_metrics_raw]	[Internal table] Contains a list of raw metrics collected by Document Intelligence. The records are deleted when metrics aggregation happens.
DocIntel OCR Input [sys_di_ocr_input]	[Internal table] Contains image data to be sent to the OCR module for processing. The records are deleted when processing is complete.
PDF [sys_di_pdf]	[Internal table] Contains a list of PDF files stored across tasks.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table	Description
DocIntel 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.

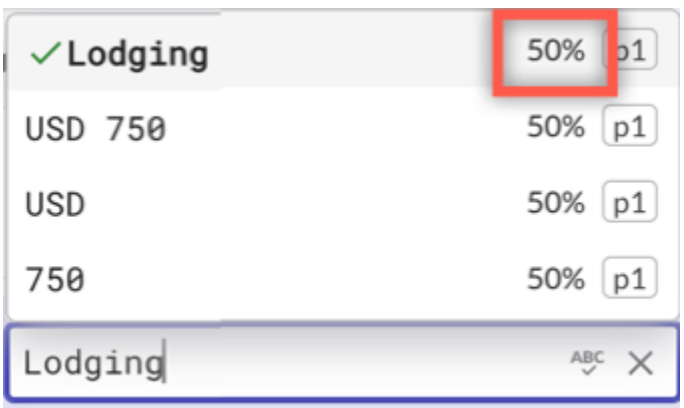
Confidence scores

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

- 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 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

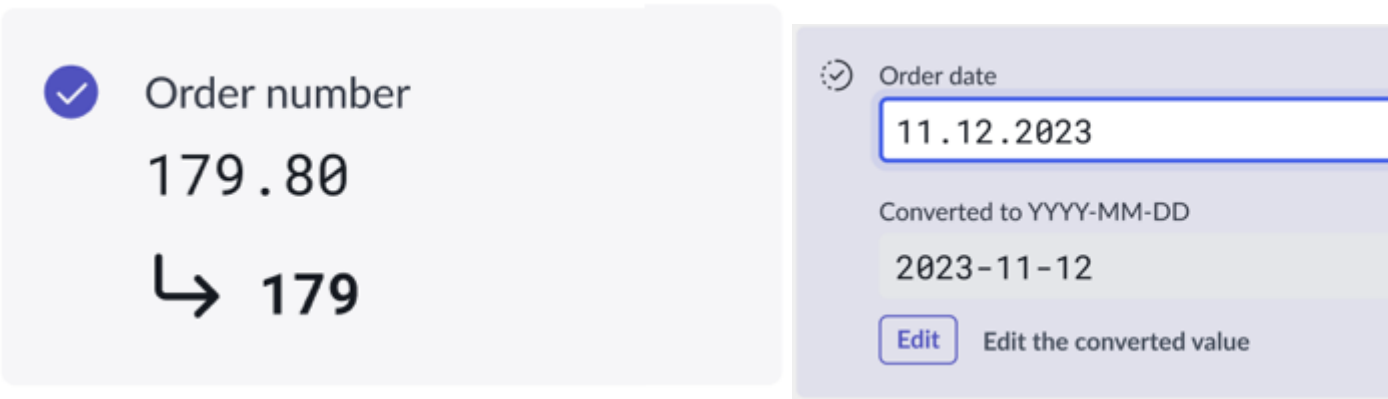
The following field types are converted to support data normalization:

Field type	Description
Date	Standard date format. For example, YYYY-MM-DD.
Integer	Whole number. For example, 12.
Decimal	Number with up to two decimal places. For example, 12.5 or 12.55.
Floating point number	Number with up to seven decimal places. For example, 12.0 to 12.0000000.

To set the field type, see Create a field.

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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence 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 name	The name for the check box list as it appears in the Document Intelligence workspace.
Target Table	The table that stores the document processing results for the check box list.
Parent Mapping to Field	Field on the target table you want to align this check box list with. **Note:** You must first select a target table.
Check box name	The name for the check box as it appears in the Document Intelligence workspace.
Select Target Field	Field on the target table that you want to align this field with. This field is used for integration with other applications. See Integrate with a custom application or workflow.
New check box	Option to add a check box to the list. Use the reorder icon (⠿) to reorder a check box in the list.
Create multiple check box lists	Option to keep the pop-up window displayed on the screen. Enable this option if you're adding more than one check box list to the use case.

Single field form

The Single field form enables you to define a single field for extraction.

The Single field form includes the following fields.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Single field form

Field	Description
Field Name	The name for the field as it appears in the Document Intelligence workspace.
Target Table	The table that stores the document processing results for this use case.
Type	The type of the field. For example, a text or date field. Some field types convert the extracted value into a standard format. See Data normalization.
Target Field	Field on the target table that you want to align this field with. **Note:** The use case must have a target table selected. This field is used for integration with other applications. See Integrate with a custom application or workflow.
Required	Option to make the field as required. Required fields can't be left empty or unreviewed.
Create more 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 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Single field group form (continued)

Field	Description
Parent Mapping to Field	Field on the target table that you want to align this field group with. **Note:** You must first select a target table.
Required	Option to make the field as required. Required fields can't be left empty or unreviewed.
Single field name	The name for the field as it appears in the Document Intelligence workspace.
Field type	The type of the field. For example, a text or date field. Some field types convert the extracted value into a standard format. See Data normalization.
Select Target Field	Field on the target table that you want to align this field with. This field is used for integration with other applications. See Integrate with a custom application or workflow.
New single field	Option to add a field to the group. Use the reorder icon (⠿) to reorder a field in the group.
Create multiple field groups	Option to keep the pop-up window displayed on the screen. Enable this option if you're adding more than one single field group to the use case.

Table form

The Table form enables you to define a table for extraction.

The Table form includes the following fields.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Table form

Field	Description
Table name	The name for the table as it appears in the Document Intelligence workspace.
Target Table	The table that stores the document processing results for these table fields.
Parent Mapping to Field	Field on the target table that you want to align this table with. **Note:** You must first select a target table.
Required	Option to make the field as required. Required table fields can't be left empty or unreviewed.
Column title	Name of the column header in the table.
Field type	The type of the field in the table. For example, a text or date field. Some field types convert the extracted value into a standard format. See Data normalization.
Select Target Field	Field on the target table that you want to align this field with. This field is used for integration with other applications. See Integrate with a custom application or workflow.
New column	Option to add a column to the table. Use the reorder icon (⠿) to reorder a column in the table.
Create multiple tables	Option to keep the pop-up window displayed on the screen. Enable this option If you're adding more than one table to the use case.

Document Intelligence properties

Document Intelligence (DocIntel) system properties control the behavior of the Document Intelligence application.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The DocIntel Manager (sn_docintel.manager) can modify the system properties for DocIntel.

Note:

Document Intelligence 3.0 includes new and updated system properties. For more detail about the upgrade, see Upgrade to Document Intelligence 3.0 or later from version 2.4 or earlier.

Document Intelligence 2.4- system properties

Property name	Description	Values
sn_docintel.default_field_sidebar_width	The default sidebar width for the document fields panel in the Document Intelligence workspace.	Allowed: Format of `[integer]px` Default: 416px
sn_docintel.default_image_fit	The default image fit for the document panel viewer in the Document Intelligence workspace.	Allowed: fit_to_page, fit_to_width Default: fit_to_page
sn_docintel.default_thumbnail_sidebar_width	The default thumbnail sidebar width for the navigation panel in the Document Intelligence workspace.	Allowed: Format of `[integer]px` 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	Allowed: Number

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence 2.4- system properties (continued)

Property name	Description	Values
	the document fields panel in the Document Intelligence workspace.	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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence 3.0+ system properties (continued)

Property name	Description	Values
glide.platform_ml_di.doc_data_extractor.show_candidate_score	Document extraction property Show confidence scores on the recommendations selection menu in the document fields panel in the Document Intelligence workspace.	Allowed: true, false Default: true
glide.platform_ml_di.doc_data_extractor.field_sidebar_layout_position	Document extraction property The layout position for the document fields panel in the Document Intelligence workspace, in relation to the document fields panel sidebar.	Allowed: right, left Default: right
glide.platform_ml_di.doc_data_extractor.default_thumbnail_sidebar_width	Document extraction property The default thumbnail sidebar width for the navigation panel in the Document Intelligence workspace.	Allowed: Format of [integer]px Default: 167px
glide.platform_ml_di.doc_data_extractor.default_image_fit	Document extraction property The default image fit for the document panel viewer in the Document Intelligence workspace.	Allowed: fit_to_page, fit_to_width Default: fit_to_page
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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Document Intelligence 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

Important:
Avoid granting an admin role when more specialized roles are available.

DocIntel Viewer [sn_docintel.viewer]

Has view-only access on Document Intelligence document tasks that they are authorized to view.

Contains Roles

List of roles contained within the role.

- snc_read_only
- platform_ml_di.viewer

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Important:
Avoid granting an admin role when more specialized roles are available.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

DocIntel Extraction Agent [sn_docintel.extraction_agent]

Extracts information from documents using the Document Intelligence workspace.

Contains Roles

List of roles contained within the role: platform_ml_di.extraction_agent.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Important:
Avoid granting an admin role when more specialized roles are available.

DocIntel Creation Agent [sn_docintel.creation_agent]

Extracts information from documents using the Document Intelligence workspace. Also enables users to create Document Intelligence document tasks and submit them for processing.

Contains Roles

List of roles contained within the role: platform_ml_di.creation_agent.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Important:
Avoid granting an admin role when more specialized roles are available.

DocIntel Manager [sn_docintel.manager]

Creates and edits use cases, fields, field groups, and document tasks. Views, measures, and analyzes the usage and effectiveness of Document Intelligence using the Platform Document Intelligence Usage dashboard. Grants access to submit document tasks and interact with the Document Intelligence workspace.

Contains Roles

List of roles contained within the role.

- platform_ml_di.manager
- action_designer
- flow_designer
- sn_ace.ace_user
- canvas_user
- usage_admin

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Important:
Avoid granting an admin role when more specialized roles are available.

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 **Process Task**. The documents for the task are being processed and the fields aren't yet populated.
In Progress	The documents for the task are processed. The fields are populated through auto-fill or by a user, who reviews and updates the predictions generated by the AI as needed.
Done	The task is completed after having been automatically processed in Fully Automated mode or after a user has reviewed the fields and selected **Submit**.
Failed	The task encountered an error during its processing and failed to process.

Domain separation and Document Intelligence

This domain separation overview relates to Document Intelligence. Domain separation enables you to separate data, processes, and administrative tasks into logical groupings called domains. You can then control several aspects of this separation, including which users can see and access data.

Support level

Support: Basic.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

How 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

Hard limits in Document Intelligence

The following is a list of hard limits in Document Intelligence, which can't be changed.

Hard limits in Document Intelligence

Limits	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	Document Intelligence supports documents in English, French, German, Spanish, Portuguese, Dutch, Italian, Czech, Danish, Finnish, Norwegian and Swedish.
Document rotation	Document Intelligence supports rotating in 90-degree increment.
Text alignment	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: • a-z, A-Z, 0–9 • á à â ä ã å ß ç é è ê ë î ï í ñ ó ò ô ö õ ú ù ü û œ • Á À Â Ä Ã Å ß Ç É È Ê Ë Î Ï Í Ñ Ó Ò Ô Ö Õ Ú Ù Ü Û Œ • , : ; . ' \ " ! ? ¿ ¡ + - * ( ) [ ] } { & % @ # / | ~ ^ < > ` ± = _ $ £ ¢ €
Sync/async operations	Async
Maximum number of document tasks processed per instance per day	The maximum number of document tasks processed per instance per day is 2000.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Hard limits in Document Intelligence (continued)

Limits	Description
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.

Terminology

Before getting started with Document Intelligence (DocIntel), it's important to understand some key concepts used in the application.

Document Intelligence terminology

Document task

A document processing activity. Includes the information that you want to extract from the document or documents.

Use case

A use case is a template used to define the structure of a specific 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. The use case also includes the mode for how the extraction should occur.

Field

A single piece of information to extract from a document. For example, the date on a document. (Fields are sometimes called keys or attributes.)

Field group

A group of fields that belong together. Field groups are used to extract information from lists and tables. For example, a group named item includes the description, quantity, and unit price fields from a purchase order.

Recommendation

A bit of text found on a document, with information about its location in the document, meaning the page of a document and the location on that page.

The recommendations that the AI provides to the user are sorted based on how likely the AI believes a given recommendation to be the correct value for the current field.

Field values

The final output of the Document Intelligence application. The output includes the values for the specified fields that were extracted for a given document task.

Document class

A field used to apply a category or label to a document and pages within a document. For example, for an identity document use case, the classes might be passport, driver's license, birth certificate, and so on.

New terminology

Document Intelligence v2.3 and later includes updated terminology for Document Intelligence features.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Terminology updates

New term	Old term
Document task	Task, extraction task
Field	Key, attribute
Field group	Key group
Field value	Extracted value
Integration	Integration setup
Recommendation	Candidate, suggestion
Fully automated	Straight through processing
Use case	Task definition

Task Intelligence

Task Intelligence uses machine learning to train solutions with your data and achieve important outcomes.

Overview

Task Intelligence enables you to set up machine learning models to automate task creation, triage, and investigation. You can then track and assess the performance of your solutions to optimize their potential to help you work more efficiently and lowering the mean time to resolve (MTTR).

Get started

Explore	Install	Configure
		
Explore Task Intelligence and machine learning models.	Install Task Intelligence and related applications.	Configure machine learning models.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Create models	Analyze models	References
		
Create and edit Task Intelligence models.	Monitor and analyze model performance.	References for components installed with Task Intelligence and language support.

Troubleshoot and get help

- Search the Now Community portal for AI & Intelligence
- Search the Known Error Portal for known error articles
- Contact Customer Service and Support

Exploring Task Intelligence

Learn more about Task Intelligence and how machine learning models can learn from your data to make predictions and achieve important outcomes.

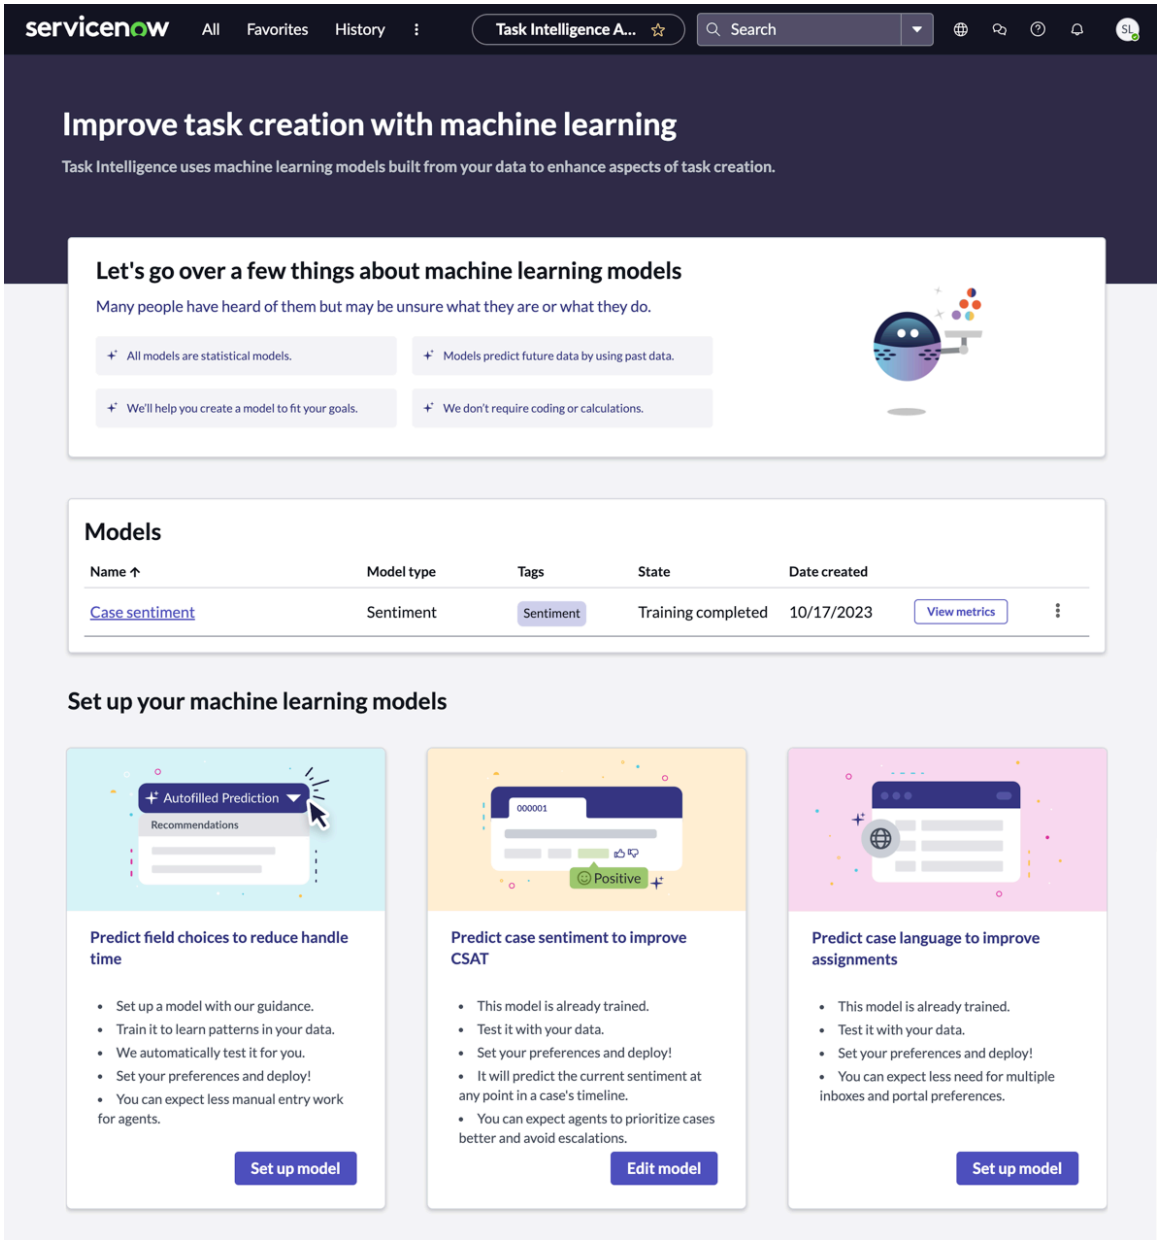
Task Intelligence overview

With Task Intelligence, you can easily set up machine learning solutions that interact with your data. Then, track how the solutions are impacting creation, deflection, triage, remediation, and optimization moments. With automated task creation, triage, and investigation, agents spend less time receiving and preparing information, so tasks can proceed faster. These models can help you work more efficiently, helping to lower the mean time to resolve (MTTR) of tasks.

With Task Intelligence, agents:

- Solve issues faster for better service and experiences
- Focus on meaningful, high-value work
- Reduce error rates and drive down costs
- Improve time to value with intuitive experiences

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



The Admin Console supports the following model types:

Model	Application	Description
Incident field prediction	Task Intelligence for ITSM	Create an incident field prediction model
Case field prediction	Task Intelligence for CSM	Predicts output fields on case forms, including categorization. Record categorization
Sentiment prediction	Task Intelligence for CSM	Predicts current and trending user sentiment.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Model	Application	Description
		Note: Model is pre-trained but can be tested and edited. Only one sentiment model can be configured.
Language detection	Task Intelligence for CSM	Detects language which language is being used. **Note:** Model is pre-trained but can be tested and edited. Only one language detection model can be configured.

Note: To learn more about setting up and testing Task Intelligence Admin Console in a production environment, see https://www.servicenow.com/community/ai-intelligence-articles/testing-field-predictions-in-task-intelligence/ta-p/2495989

Task Intelligence workflow

The Task Intelligence Admin Console uses machine learning models. The models are statistical and can predict future data by training with your past data.

With the Admin Console, create, configure, train, test, and deploy predictive models used for automation in other ServiceNow features and applications, such as Customer Service Management (CSM).

Training a machine learning model is when the model learns patterns in past data to make predictions for new data. Models are trained using a lot of data so that they can learn patterns and the large data set makes the learned patterns statistically significant. By answering questions about your information systems, business process, and service operations, the system actively learns from your responses.

Task Intelligence benefits

The Task Intelligence Admin Console provides admins with a no-code experience for deploying Task Intelligence solutions. The seamless experience helps you automate and optimize task creation, deflection, triaging, and resolution.

Features are implemented by two applications: Task Intelligence for Customer Service and Task Intelligence for ITSM.

Benefit	Feature	Users
Predict form field values to auto-populate or recommend as suggestions	Create a field prediction model	Admins, agents
Categorize emails and cases based on language and attachment content	Record categorization	Agents
Predict incident category and priority based on previous incidents to lower time to resolution	Create an incident prediction model	Agents
Analyze initial and ongoing sentiment during customer service cases	Sentiment Analysis	Agents

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Benefit	Feature	Users
Identify the language used to create customer service cases	Language detection	Agents
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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

You can 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 Install dialog box that is displayed, any dependencies that are installed along 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.

   Important:
   If you don't load the demo data during installation, it's unavailable to load later.

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

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Edit the 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.

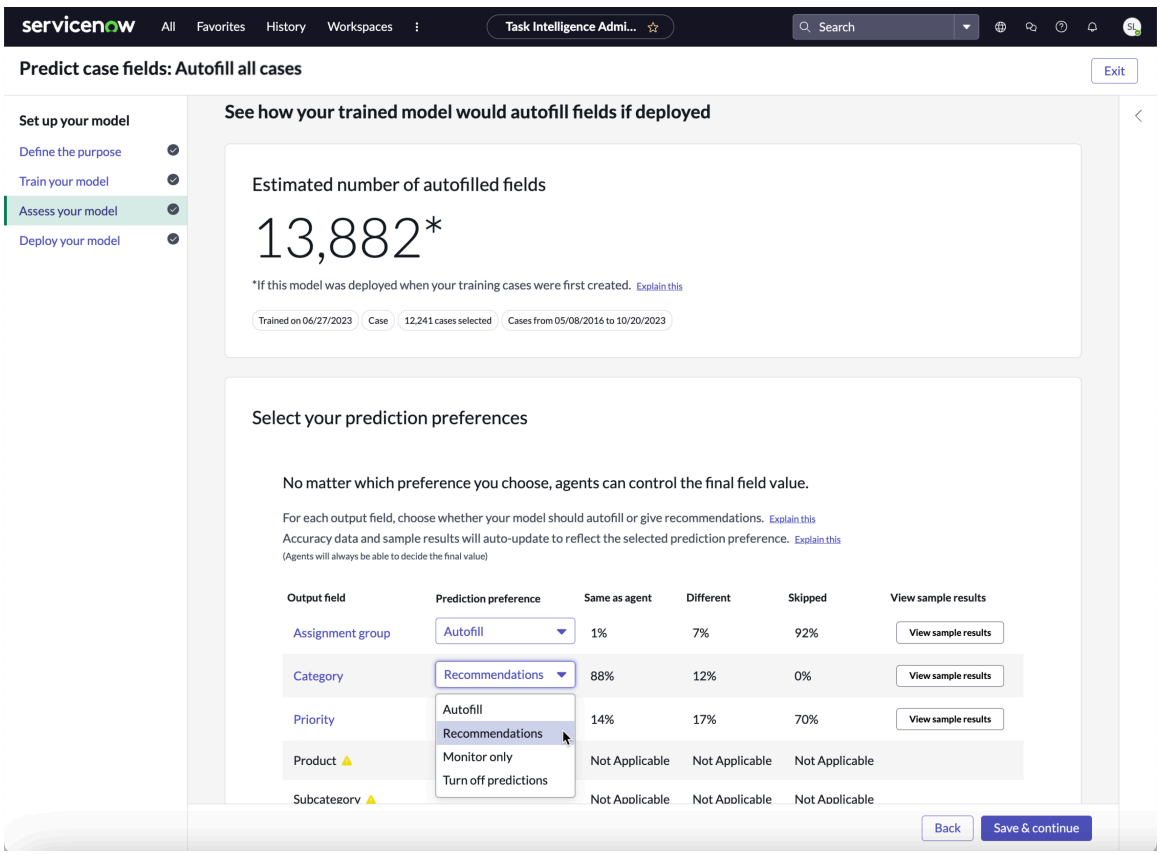
Assess a Task Intelligence model

Assessing a machine learning model's performance helps determine how to use and train the model to achieve your desired outcomes.

Overview

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.



© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

The 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.

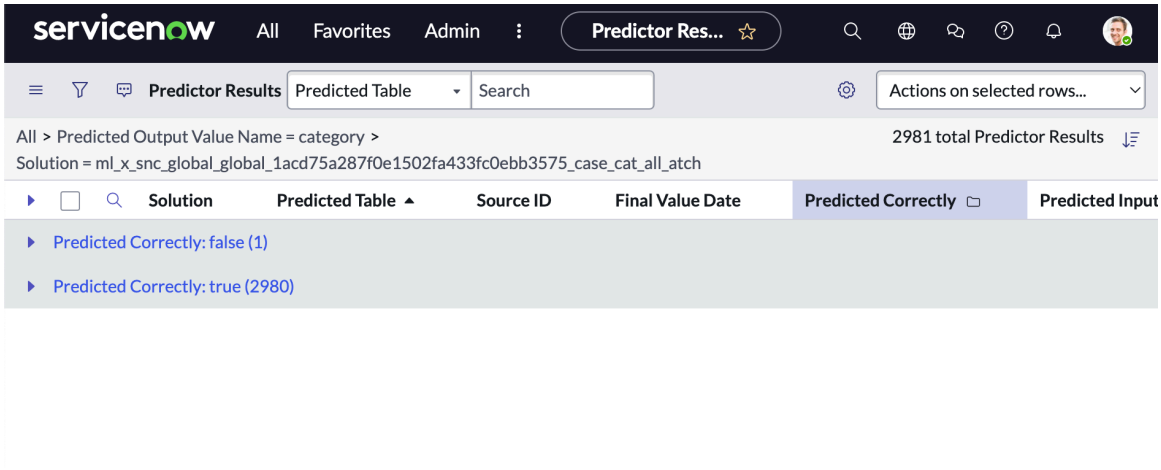
If you have not already created a model, see either Create a model to predict record fields or Create a model to predict incident fields to make one.

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.

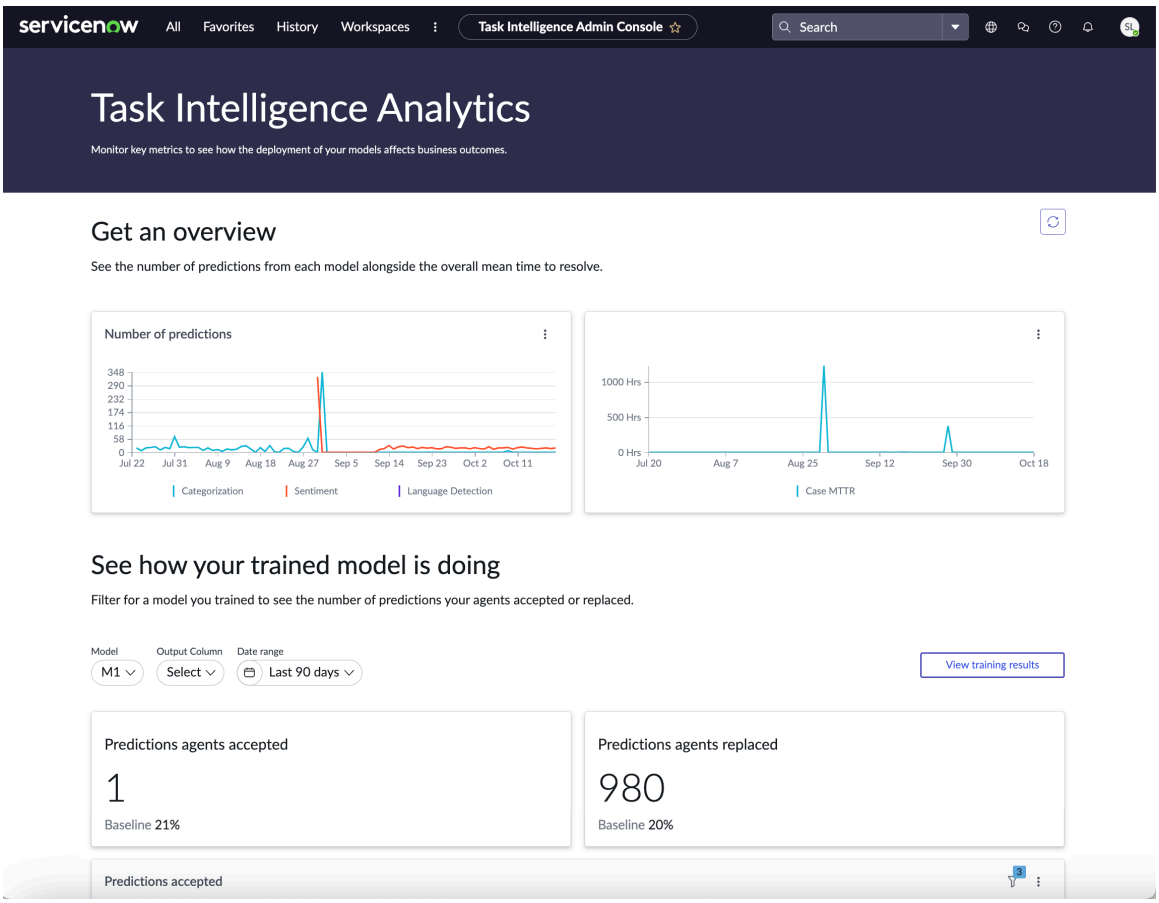
© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.



4. 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

c. Select 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

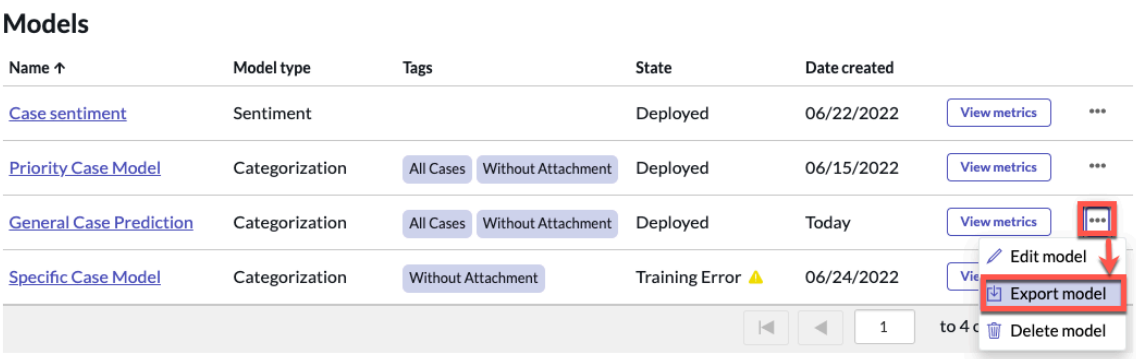
Export a Task Intelligence model to another instance.

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.



3. Click **Export model**.
   A new tab opens with authentication steps.
4. Click **Allow**.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Result
The 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.

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries. Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Task Intelligence User [tia_user]

Views and monitors Task Intelligence models. Requires additional accesses to see training.

Contains Roles

List of roles contained within the role: ml_report_user.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Important:
Avoid granting an admin role when more specialized roles are available.

Task Intelligence Analyst [tia_analyst]

Monitors Task Intelligence models through the Analytics dashboard.

Contains Roles

List of roles contained within the role: ml_admin.

Groups

List of groups this role is assigned to by default.

None.

Special considerations

Important:
Avoid granting an admin role when more specialized roles are available.

Languages supported by Task Intelligence

The Task Intelligence application provides support for predictions in different languages.

Each of the supported languages has an associated two-character code. For example, `English - en`.

The following table lists the languages supported for Task Intelligence along with the models they support.

Note:
Only languages included in the training dataset are supported.

Language name	Language code	Incident fields (ITSM)	Case fields (CSM)	Sentiment prediction (CSM)	Language detection (CSM)
Arabic	ar				X
Brazilian Portuguese	pb	X			

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Language name	Language code	Incident fields (ITSM)	Case fields (CSM)	Sentiment prediction (CSM)	Language detection (CSM)
Chinese (simplified)	zh	X			X
Czech	cs				X
Danish	da	X			
Dutch	nl	X			X
English	en	X	X	X	X
Estonian	et				X
Finnish	fi	X			X
French	fr	X	X		X
French Canadian	fq	X			
German	de	X	X		X
Hebrew	he				X
Hungarian	hu				X
Italian	it	X			X
Japanese	ja	X			X
Korean	ko	X			X
Norwegian	nb	X			X
Polish	pl	X			X
Portuguese	pt	X			X
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
Document Intelligence	Release notes	Document Intelligence release notes

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource links (continued)

AI features or products	Resource type	Resources
Document Intelligence	Getting Started	Getting started with Document Intelligence Document Intelligence quick start guide (ServiceNow Community)
Document Intelligence	Tips and examples	Document Intelligence examples
Document Intelligence	Troubleshooting	Troubleshooting and best practices for Document Intelligence
Document Intelligence	FAQs	Document Intelligence FAQs
Natural Language Understanding and Natural Language Query	Getting Started	Getting started with NLU NLU quick start guide (ServiceNow Community)
Natural Language Understanding and Natural Language Query	Best practices	Best practices for NLU
Natural Language Understanding and Natural Language Query	Tips and examples	NLU examples
Natural Language Understanding and Natural Language Query	Troubleshooting	Troubleshooting for NLU
Natural Language Understanding and Natural Language Query	FAQs	NLU FAQs
Natural Language Understanding and Natural Language Query	Release notes	Natural Language Query release notes
Now Assist	Release notes	Now Assist release notes Now Assist for Creator release notes Now Assist for Customer Service Management (CSM) release notes Generative AI Controller release notes

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource links (continued)

AI features or products	Resource type	Resources
Now Assist	Getting Started	Getting Started with Now Assist Getting started with Now Assist for Code Getting Started with Now Assist for Creator Getting Started with Now Assist for CSM Getting Started with Now Assist for ITSM Getting Started with Now Assist for AI Search Now Assist quick start guide (ServiceNow Community) Now Assist panel fundamentals
Now Assist	Best practices	Now Assist best practices (ServiceNow Community) Knowledge article best practices (ServiceNow Community)
Now Assist	Tips and examples	Now Assist conversational catalog request in practice (ServiceNow Community) Now Assist examples
Now Assist	FAQs	Generative AI Controller FAQ Now Assist FAQ Now Assist in AI Search FAQ
Predictive Intelligence	Release notes	Predictive Intelligence release notes
Predictive Intelligence	Getting Started	Getting started with Predictive Intelligence Predictive Intelligence quick start guide (ServiceNow Community)
Predictive Intelligence	Best practices	Predictive Intelligence best practices

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.

Resource links (continued)

AI features or products	Resource type	Resources
Predictive Intelligence	Tips and examples	Predictive Intelligence examples
Predictive Intelligence	FAQs	Predictive Intelligence FAQs
Task Intelligence	Getting Started	Getting started with Task Intelligence Task Intelligence quick start guide (ServiceNow Community)
Task Intelligence	Best practices	Troubleshooting and best practices for Task Intelligence
Task Intelligence	FAQs	Task Intelligence FAQs

© 2025 ServiceNow, Inc. All rights reserved.
ServiceNow, the ServiceNow logo, Now, and other ServiceNow marks are trademarks and/or registered trademarks of ServiceNow, Inc., in the United States and/or other countries.
Other company names, product names, and logos may be trademarks of the respective companies with which they are associated.